Title
Essays in Labor Economics and Networks

Author
Nadler, Carl Marc

Publication Date
2016

Peer reviewed|Thesis/dissertation

Essays in Labor Economics and Networks

by

Carl Marc Nadler

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Card, Chair
Associate Professor Bryan Graham
Associate Professor Patrick Kline
Associate Professor Jeremy Magruder

Spring 2016

Essays in Labor Economics and Networks

Abstract

Essays in Labor Economics and Networks

by

Carl Marc Nadler

Doctor of Philosophy in Economics

University of California, Berkeley

Professor David Card, Chair

Network-based connections are pervasive in hiring and mobility patterns. While the theoretical impacts of network connections on the inequality of labor market outcomes are well known (Calvo-Armengol and Jackson, 2004), the empirical magnitude of actual network effects is less certain. A key issue is the difficulty of disentangling the causal effect of network connections from differences in the characteristics of workers in better and worse networks.

In the first chapter of my dissertation, I study this question using data on freelance workers in Hollywood, who, like many workers in the "freelance economy," are hired for short-term jobs through an informal process that relies in part on previous connections. In such a labor market, the fortunes of an individual worker are closely linked to the careers of key agents who make the hiring decisions for jobs. In Hollywood, workers who know position supervisors who manage more jobs will have more job opportunities. To measure the size of the network effects, I follow cohorts of freelance grips and lighting technicians who first work on a major movie production between 1988 and 2002. I develop two alternative models for the probability that workers are hired on subsequent productions—based on random effects and fixed effects specifications—that incorporate network effects, experience effects, and unobserved heterogeneity. Both models yield large estimates of the effect of experience-based connections on hiring outcomes. I then use the random effects model to develop simulation-based estimates of the fraction of overall inequality in job outcomes for workers in a given cohort that is attributable to inequality in the career success of the key supervisors they met during their initial year of work. I find that about half of the wide dispersion of career outcomes in Hollywood is generated by differences in the career trajectories of these initial key supervisors.

In the second chapter, I study the asymptotic properties of the fixed effects estimator I employ in the first chapter. The estimator is robust to unobserved heterogeneity across workers and movies. It is based on subgraphs of worker-movie dyads that I call pairs. Inference is non-standard, because pairs within a sample are only independent when they do not share any workers or movies in common. The underlying criterion is a two-sample

U-process. I show that the U-statistic derived from the estimator's first order condition is asymptotically equivalent to a certain projection which involves summation over all the worker-movie dyads in the sample. I use this result to derive a consistent estimator of the variance of the estimator.

To Mom, Dad, and Elyssa

Acknowledgments

There are many people who supported me as I worked on this project that I wish to acknowledge. I will begin with David Card, who chaired my committee. Over the course of the three years I worked on this project, I went to David Card many times for advice and encouragement. His belief in this project was a frequent source of strength. I am also extremely grateful to Bryan Graham and Pat Kline. Bryan taught me how to think about the complicated endogeneity issues that arise in studying networks. His approach for modeling network formation strongly influenced the hiring model I estimate. Pat helped me breathe new life into this project at a time during my fourth year when I was ready to abandon it. When I came to him with an idea to frame the project around inequality, he suggested that I look at the relationship between a worker's initial network of connections and their long term outcomes. This relationship eventually became the focus of my first chapter.

I began this project during my third year for the qualifying oral exam. My oral committee members met with me many times during this formative period and provided important guidance. In addition to Bryan, these members included Enrico Moretti, Jesse Rothstein, and Sameer Srivastava. I would also like to thank several other faculty who provided key advice at various stages of the project: John Abowd, Sylvia Allegretto, Hilary Hoynes, Matthew Jackson, Jeremy Magruder, Michael Reich, Ian Schmutte, Chris Walters, and Andrea Weber.

I first learned about key grips and gaffers from Bernard Hunt. We were a year or two out of New York University (NYU), and at the time Bernard worked as a lighting technician in the city. Over the past few years, I have spoken with Bernard over the phone several times to learn more about how grips and lighting technicians build their careers. Without Bernard, I doubt I would have written this dissertation. My friend Madeline Austin-Kulat, who was a member of an IATSE local in New York, also spoke to me several times about what life is like working on movie productions. I was also fortunate to speak informally with two members of the IATSE-International union, Vanessa Holtgrewe and Russell Nordstedt. In addition to helping me understand how grips and lighting technicians find their jobs, they provided important background on the role of the union itself.

I am extremely grateful to my friends, who provided constructive feedback at various stages. Eric Auerbach, in particular, was extremely helpful. I have learned a lot from Eric's deep understanding of random networks. I am very fortunate to have had as office mates David Silver and Moises Yi, who created a supportive and fun work environment. Kyle Carlson, Sydnee Caldwell, Andrew Foote, and Kevin Todd gave helpful comments on earlier versions of my first chapter. Sanaz Mobasseri introduced me to many articles in the sociology literature. Hedvig Horvath and Attila Lindner provided helpful advice when I was preparing for the job market and cared for me when I was injured during the summer before my final year. Austin Chase provided valuable research assistance. There are many others who provided incredible feedback and support, and I am sincerely grateful to have them all in my life.

I would also like to acknowledge the financial support I have received. I was very fortunate to obtain a National Science Foundation Graduate Research Fellowship, which funded my first three years of graduate school. I would also like to thank David Card for supporting me at various times via research and teaching assistantships and Thomas Marschak for employing me as a teaching assistant in his class. Sylvia Allegretto and Michael Reich supported my fifth year via a research assistantship at the Institute for Research on Labor and Employment (IRLE). I am extremely grateful to the IRLE, not only for what I learned while working there, but also for the community it provided. I also thank the economics department and the university for the Dean's Completion Fellowship and summer support.

My journey to become an economist began at NYU. I was originally a film student, with little interest in research. That changed when Alessandro Lizzeri employed me as a research assistant during the summer before my senior year. He later introduced me to Matthew Wiswall, who became my senior thesis adviser. Matthew sparked my initial interest in labor economics and worked with me for three years to publish a paper based on my thesis research. Both of them were extremely helpful in advising me on how to get into graduate school. I am also very grateful to Mary Burke and Jane Katz who hired me at the Federal Reserve Bank of Boston. Many of my interests in labor economics and network effects can be traced back to the research they worked on. I am also grateful to the other economists I worked for at the Fed, including Julian Jamison and Daniel Cooper.

Finally, I would like to thank my family, who have supported me in every endeavor. I dedicate this dissertation to them.

Chapter 1

Networked Inequality: Evidence from Freelancers

1.1 Introduction

Young workers increasingly rely on their social networks when looking for jobs. Between 2005 and 2011, the fraction of new entrants who report contacting friends or family during their job search rose from about 15 percent to 25 percent.[1] Today, over a quarter of internet users are members of the website LinkedIn, which enables workers to search for members of their network who are connected to job vacancies (Duggan et al., 2015). It is not surprising that networks are so frequently used during job search. At least half of all jobs are found through personal contacts (Topa, 2011), and a growing body of evidence confirms that referred applicants are much more likely to be hired than those without referrals.[2]

In this chapter, I present new evidence on the effect of networks formed early in workers' careers on their long-term labor market outcomes. Models of networked labor markets, such as Calvo-Armengol and Jackson (2004), suggest that differences in initial network quality may be an important factor in explaining disparities in employment outcomes between groups.[3] Yet, true network effects are difficult to distinguish from worker attributes that may be correlated with their network quality. Social networks often form in environments, such as at school or in the workplace, that select workers based on their abilities. This non-random assignment of workers into networks biases estimates based on common measures of network quality such as the network's employment rate.

I study a unique labor market that is well-suited for assessing the role of networks: the "below-the-line" freelancers who work on Hollywood movies.[4] New crews are assembled for

[1]These statistics are based on tabulations performed by the Bureau of Labor Statistics. See their annual average data on characteristics of the unemployed, Table 34: "Unemployed jobseekers by sex, reason for unemployment, and active job search methods used." Historical data are available going back to 1995: http://www.bls.gov/cps/tables.htm. Last accessed July 2, 2015.

[2]See, for example, Brown et al. (2016), Burks et al. (2015), and Fernandez et al. (2000).

[3]Models in which social networks affect the distribution of labor market outcomes include Krauth (2004), Ioannides and Soetevent (2006), Montgomery (1994), and Zenou (2013). Granovetter (1973, 1974) provides a number of insights on how networks affect labor markets that continue to inform this literature.

[4]In film production the above-the-line workers include the cast, producers, writers, and directors. Every-

each film, and *key supervisors* typically decide which workers are hired. As a result, workers who know supervisors who manage more jobs will have more job opportunities. I assemble a dataset of screen credits on over 3700 major movie productions and follow cohorts of freelance grips and lighting technicians who first take a job on a major movie production between 1988 and 2002. l develop two alternative models for the probability that workers are hired on subsequent productions—based on random effects and fixed effects specifications—that incorporate network effects, state dependence, and unobserved heterogeneity. I then use the random effects model to develop simulation-based estimates of the fraction of overall inequality in job outcomes for workers in a given cohort that is attributable to inequality in the career success of the key supervisors they met during their initial year of work.

I find large returns to experience-based connections to key supervisors on hiring outcomes. In Hollywood, careers among grips and lighting technicians are highly skewed. After a first job on a major movie production, many never appear again, and a typical worker has only a handful of subsequent credits. A small group of technicians, however, go on to work on 1–2 movies per year.[5] Results from my simulations indicate that about half of this dispersion is generated by differences in the career trajectories of the key supervisors who manage workers during their initial year.

This research contributes to a large empirical literature assessing the impact of social networks on labor market outcomes.[6] My approach is similar to one recent strand that exploits matched worker-firm panel data and tests whether higher quality co-worker networks improve job search outcomes following a mass layoff (e.g., Cingano and Rosolia, 2012; Glitz, 2013; and Saygin et al., 2014).[7] These papers find that better networks help workers find employment, but evidence on wage effects is mixed. An important issue confronting this literature is disentangling the causal effect of network connections from differences in the characteristics of workers in better and worse networks.

I directly address this issue using methods from an emerging literature that studies economic models of how networks form.[8] The econometric challenge of identifying network effects given potentially unobserved differences in workers' abilities is similar to the issue of distinguishing state dependence from unobserved heterogeneity (Heckman, 1981). My

one else is below-the-line.

[5]Previous research studying the career outcomes in Hollywood using screen credit data have found similar patterns. See, for example, Faulkner and Anderson (1987).

[6]A closely related literature studies why firms use social networks when hiring workers. This research is motivated by theories of employee referrals that argue that network-based hiring reduces uncertainty in the worker-firm match (e.g., Simon and Warner, 1992) or helps firms find higher ability workers (e.g., Montgomery, 1991). Recent contributions include Beaman and Magruder (2012), Brown et al. (2016), Burks et al. (2015), Dustmann et al. (Forthcoming), Heath (2013), Hensvik and Skans (Forthcoming), and Pallais and Sands (Forthcoming).

[7]Past research has also considered network measures based on labor market outcomes of nearby residents (e.g, Bayer et al., 2009; Hellerstein et al., 2015; and Schmutte, 2015), workers of the same nationality or ethnic group (e.g., Åslund et al., 2014 and Beaman, 2012), and family (e.g., Kramarz and Skans, 2014 and Magruder, 2010). See Ioannides and Loury (2004) and Topa (2011) for reviews, focusing on the earlier economics literature. Granovetter (1995) and Marsden and Gorman (2001) review the related sociology literature.

[8]See Chandrasekhar (2015) and Graham (2015b) for recent reviews. This research is related to a statistics and computer science literature that characterizes the distribution of random networks (e.g., Goldenberg et al., 2009).

main approach for estimating network effects builds on recent work by Goldsmith-Pinkham and Imbens (2013). They develop a random effects framework and model the formation of the endogenous network. I extend this approach to a multi-period, labor market setting. I corroborate the main conclusions from this analysis using a fixed effects approach (Charbonneau, 2014; Graham, 2014, 2015a).

I begin by documenting career inequality among workers in my setting and relating it to differences in the quality of workers' initial network of key supervisors. I define cohorts by year of initial employment and follow each cohort for the subsequent 10 years. I measure a worker's initial network quality by the number of jobs that a worker's initial supervisors can pass on to them during the subsequent decade. I call this measure a worker's *initial connections*. A simple least squares fit suggests that almost 1 in 10 initial connections led to work.[9]

Next, I present descriptive evidence suggesting that work experience with key supervisors increases the probability that a worker is hired on a given movie. I include in a worker's *key network* all key supervisors who managed them in an earlier year, and I measure direct and indirect connections to supervisors on a movie. Since networks depend on earlier positive employment outcomes, simple comparisons of workers with and without connections overstate the network effect. I investigate this issue two different ways and conclude that this bias is modest relative to the large impact of connections in my setting. First, I compare workers who have been in the sample for the same number of years and who have worked on the same number of previous movies. Second, I compare workers who have been in the sample for the same number of years and who are directly connected to the same number of movies that are hiring that year. Both approaches lead to similar results. Though the bias strengthens over time as those with more successful employment outcomes sort into higher quality networks, the overall magnitude of the adjusted impacts suggest powerful network effects at all stages of workers' careers.

Motivated by these facts, I develop a model for the probability a movie hires workers based on network connections as well as other, potentially unobserved, factors. I first take a random effects approach to addressing this unobserved heterogeneity, specifying hiring models with up to 3 types of worker effects. Similar to the correlated random effects models employed in the panel literature (Chamberlain, 1984), the probability of a worker's type depends on their job outcomes during their initial year and their initial connections.[10] The results strongly confirm the importance of unobserved heterogeneity: About 60 percent of workers in my sample are extremely unlikely to be hired even with a direct connection. The "highest" type of worker, representing about 14 percent of the sample, is about 3 log odds points more likely to be hired than an observably similar "lowest" type.

[9] This finding is similar to earlier research on Hollywood's labor force by Jones (1996) and Jones and Walsh (1997) who find that network position during the late 1970's predict career outcomes for various occupations over the next 10 years. In their analysis, network position is measured using a k-core decomposition. More recently, Lutter (2015) and Rossman et al. (2010) show that different measures of network structure predict actors' Academy Award nominations and survival. I contribute to this literature by using relationships with supervisors to assess the effects of networks on careers.

[10] In my sample, I observe each worker for 11 years. I take as given the hiring outcomes during their initial year, and model their outcomes over the next 10 years. Wooldridge (2005) proposes modeling the random effects conditional on initial conditions. A worker's initial connections are a function of these initial conditions and the hiring outcomes of the key supervisors over the next 10 years, which I assume are exogenous.

Results from the random effects models confirm there are large returns to network connections that increase with the number of times a worker has been managed by a key supervisor. Nevertheless, a concern with the random effects specifications is that it assumes worker unobservables that may bias my estimates do not change over time and that I observe all relevant attributes of a movie. I assess the robustness of these estimates to more general forms of heterogeneity by developing an alternative fixed effects estimator that allows for time-varying unobserved worker effects as well as unobserved movie effects. I propose an estimator based on pairs of worker-movie observations that contain exactly two workers and two movies. I show that, conditioning on pair outcomes in which each worker is hired on one movie and each movie hires one worker, the marginal likelihood does not depend on the unobserved worker and movie effects. The marginal likelihood admits a familiar difference-in-difference interpretation, in which the network effects are identified off of the relative strength of the workers' connections on the two movies. Comparing estimates using the fixed effects estimator to what I found using the random effects approach, I conclude that my model's estimates are robust to unobserved heterogeneity among workers and movies.

In the final part of the chapter, I estimate the fraction of overall inequality in job outcomes for workers in a given cohort that is attributable to inequality in the career success of the key supervisors they met during their initial year. I use the random effects model to simulate a counterfactual that removes variation in the initial supervisors' career trajectories. I find the counterfactual reduces the variance of workers' career outcomes about 50 percent and reduces the mean by 18 percent. A decomposition of the effects by type reveals that these reductions are driven by effects on the upper tail of the distributions of high and medium type workers. Intuitively, more active key supervisors provide opportunities for the higher ability workers to develop stronger relationships with them, which increases the chance the workers will be hired on the jobs the supervisor manages. Over time, the connected workers accumulate a network that can provide them jobs independent of the supervisors they were initially attached to. The counterfactual disrupts this process by reducing the number of highly active initial supervisors.

The remainder of the chapter is structured as follows. Section 1.2 provides background on the two groups of workers that I study. Section 1.3 explains how I construct my sample, and Section 1.4 examines the relationship between initial network quality and workers' career outcomes. Section 1.5 provides descriptive evidence of the importance of connections. Section 1.6 describes the model. Section 1.7 reports the empirical results. Section 1.8 examines the robustness of these results to unobserved worker and movie heterogeneity. Section 1.9 uses the estimates to perform the counterfactual analysis. Section 1.10 concludes.

1.2 Grips and lighting technicians

Movie production in Hollywood has been project-based since the fall of the Studio System in the 1950s (Christopherson and Storper, 1989; Scott, 2005). Crews of freelance technicians and actors are brought together on a temporary basis for the purpose of shooting a single film. This paper focuses on two specific groups of workers, *grips* and *lighting technicians*, who take responsibility for the arrangement and movement of different kinds of equipment.[11]

Lighting technicians (also called "electrics" or "sparks") set up and operate the lighting equipment. When movies are shot on location, lighting technicians also find the sources of electricity for operating the lights or set up independent generators.

Grips (sometimes called "hammers") are responsible for the non-lighting equipment, including flags, camera rigs, dollies and the dollies' tracks. An important position on the movie set is the dolly grip, who moves the dolly on which the camera rests during tracking shots.

Although general demographic portraits of these workers are unavailable, informal conversations with members of the union suggest a large range in age and educational backgrounds. Some grips and lighting technicians enter the business young with the help of family connections, while others start after studying filmmaking at degree or certificate programs. Some enter with the specific intent of making a career as a grip or lighting technician, while others begin while pursuing other, higher profile filmmaking jobs, such as directing or cinematography.

1.2.1 Key supervisors

Figure 1.1 shows the standard organization of the grip and lighting technician crews.[12] Position-specific key supervisors are responsible for the management of these crews. I often call these supervisors a worker's *keys*. These supervisors—the key grip for the grips and the chief lighting technician for the lighting technicians[13]—are hired before production begins and work with the director of photography and unit production manager to assess the personnel and equipment needs for the production. The key grips and chief lighting technicians are supported in their supervision of the rest of their crews with the help of their "best boys."[14]

Ethnographic and narrative accounts suggest that keys prefer to hire those who they have worked with before.[15] Table 1.1 provides some direct evidence. It shows the average number of grips and lighting technicians by their relationship to their respective key supervisors on major California movies released between 1988 and 2012. I detail how I construct this sample of movies in Section 1.3.

On average, there are about 9 grips and lighting technicians on movies released during this period. About a third of this average is composed of new workers, who have never

[11] See Taub (1994) and Crouch (2002) for background on the grip and lighting technician crews based on interviews.

[12] See, for example, Bechky (2006).

[13] On some movie sets, the key grip is referred to by their official title on the union contract, the first company grip. The chief lighting technician is often called the gaffer.

[14] As the title best boy suggests, most grips and lighting technicians are men. For example, Lauzen (2009) finds that about 1 percent of key grips and chief lighting technicians on the top 250 grossing movies of 2008 were women.

[15] See, for example, Bechky's (2006) ethnographic description of life on movie sets. Relationships with chief lighting technicians and key grips are emphasized as important sources of work in Crouch's (2002) compilation of interviews on building careers in Hollywood. Taub (1994) describes how one chief lighting technician directly hires his crew before movie production occurs. Jones's (1996) qualitative research on careers in the film industry emphasizes personal contacts generally, along with reputation and skills as important factors in building a career in project-based industries.

appeared on a major California movie in a prior year in the position. Over half of the remaining crew is composed of workers who have either been supervised by a key in a prior year ("directly connected workers"), or have been supervised by another key who has worked alongside one of these keys ("indirectly connected workers"). Grips and lighting technician crews have similar relationships with their respective key supervisors. Table 1.1 suggests that networks play an important role in assembling a crew, although we might find large shares of connected workers for other reasons as well.

1.2.2 Institutional setting

Grips and lighting technicians, like most production workers in Hollywood, are represented by separate locals of the International Alliance of Theatrical and Stage Employees (IATSE) union. It is rare for workers to switch between these jobs once they become members of their respective local. The union contract guarantees a high hourly wage on movies produced by the major studios and their subsidiaries[16] plus health benefits for those who work a minimum amount. Independent low budget productions typically pay lower rates.[17]

Membership into the unions is difficult, requiring completing a minimum number of hours on union productions, which usually only hire members. Membership is granted for applicants during a union meeting, during which current members are called on to speak on behalf of the applicant. According to public records, since 2000, membership in the grip and lighting technician locals have been fairly steady at around 2600 and 2100 active members, respectively. Each year there are between 40 and 200 non-members paying dues, and between 10 and 50 workers are granted membership.[18]

When the Studio System ended in the 1950s, the IATSE locals originally played a direct role in the assignment of production workers onto movies (Christopherson and Storper, 1989; Amman, 1996). The unions operated rosters that ranked members according to their experience working in the industry. Similar to a hiring hall, available members were placed onto movie sets based on their position on the roster. This system faced criticism as the unions expanded their membership over the next couple of decades, taking in workers from television and independent movie production, and was dismantled (Amman, 1996). The

[16]For example, the 2012–2015 contracted minimum rate for grips and lighting technicians for movies produced by the major studios is 36 dollars an hour.

[17]During the 1980s, independent production companies began to produce a larger share of Hollywood movies. The major studios would then buy these movies for distribution. While some of these films would still employ union members, the terms of employment were not governed by the union contract unless there was a successful organizing drive on the movie set Amman (1996). IATSE unionized more of these independent production companies during the 1990s, in part, by agreeing to lower wage scales for low budget film and television productions (Cox, 1994; Johnson and Weiner, 1996; Madigan, 1998).

[18]These membership numbers are based on tabulations from the Labor Organization Annual Reports (2000–2013), reported by the US Department of Labor. Public records are available going back to 2000 at `http://www.dol.gov/olms/regs/compliance/rrlo/lmrda.htm` (last accessed: July, 8, 2015). The grip and lighting technician locals have slightly different reporting practices. The numbers I report on non-members paying dues come from the grip union, IATSE Local 80, while the numbers I report on new members come from the lighting technician union, IATSE Local 728. Since movies usually have a similar number of grips and lighting technicians on set, I believe Local 80's larger membership is due to their representation of several other occupations in Hollywood, including crafts service, marine, First Aid employees, and warehouse workers.

roster system does not operate during the period of film production that I study (Amman, 1996; Scott, 2005).

1.3 Data

1.3.1 Sample construction

I measure workers' networks and career outcomes using screen credit data from the Internet Movie Database (IMDb).[19] IMDb was created in 1994 and is considered the most complete, reliable aggregator of information on movies and the people who create them.[20] The website collects and verifies data from studios, filmmakers, on-screen credits, press kits, and fans, among other sources. Each title in the database is accompanied with cast and crew credits as well as release dates, genre, filming dates, shooting locations, and other trivia, when available. Today, many workers in the film industry use their IMDb filmography—the webpage IMDb creates for each person in their database that lists all their credits—as a resumé.

In July 2013, I searched the website for all theatrically released feature films produced in the United States that were released between January 1980 and December 2012. The search provided a list of IMDb web addresses for over 26 thousand titles. I recorded the non-cast credits found on the "Full Cast and Crew" webpage for each of these titles. The screen credits on these webpages have a unique identifier for each person, which I use to create a worker-movie dataset for all workers on these movies. I merge into this dataset movie characteristics (e.g., filming locations, theatrical release dates, ticket sales, production budgets) from IMDb's Plain Text Files, available for download on their FTP site,[21] and AC Nielsen.[22] My focus is on the careers of the grip and lighting technician crews, and I drop screen credits for work outside of these positions. I include in my sample screen credits on only major, theatrically released movies, which I measure using the information on the movie's budget, genre, and crew size that is available.[23] Figure 1.2 shows that, starting in the late 1980s, the resulting number of major movies corresponds roughly to the Motion Picture Association of America (MPAA)'s count of feature films distributed by the major studios and their subsidiaries. I make this restriction primarily because the kinds of productions listed on IMDb change over time. Today, many independent filmmakers use IMDb for self-promotion, and included among the titles I originally found through my search are movies that were shown only at film festivals. These independently exhibited productions make up most of the theatrically released productions on IMDb during the past decade, but make up a much

[19] `http://www.imdb.com`.

[20] Recent papers studying the movie industry using IMDb data include Button (2015), Lutter (2015), and Rossman et al. (2010).

[21] Links to the public FTP servers that hold the Plain Text Data Files are currently available on IMDb's Alternative Interfaces webpage: `http://www.imdb.com/interfaces` (last accessed January 14, 2013).

[22] The AC Nielsen data are from the replication files to Moretti (2011).

[23] In order to qualify as a major motion picture, the movie production must be a feature film shown in theaters with non-missing ticket sale information. I also drop productions with budgets smaller than 1 million (2012) dollars and the following genres: documentaries, game shows, short films, and pornography. I drop films that employ cinematographers and camera crew that do not appear on any other productions. I also drop films that do not have any grip or lighting technician workers.

smaller share during the 1980s and early 1990s, before IMDb was created. Restricting the sample to major films provides a consistent measure of career outcomes and networks over time.

A second reason for focusing on major, theatrically released movies is that these tend to be the best paying jobs for the workers I study. The typical Hollywood movie is covered by the IATSE union contract and provides several weeks of well-paid work. For example, at the 2012–2015 contracted minimum rate of 36 dollars an hour, taking into account overtime pay, 12 hour days, and 6 day workweeks, a grip or lighting technician earns over 3 thousand dollars a week. A filmography with a number of major movie credits signals a successful career in Hollywood.

I focus on movies filmed in California, which produces most of the theatrically released films during the period I study.[24] My impression, based on informal conversation with workers in this industry, is that while movies filmed outside of California might employ California residents for leading cast and department head positions, movies typically hire local crews, especially for the non-supervisory positions. Geographic distance between the worker's residence and the production's location could bias the estimated effect of network connections on hiring if not accounted for. Because shooting location information is missing for most of the movies in my sample, I define California productions to be those that list a California shooting location on IMDb or employ at least one non-supervisory grip, lighting technician, or camera worker who also worked on a movie with a California shooting location released the same year.

Table 1.2 shows descriptive statistics for the resulting *major California* movie productions released between 1980 and 2012 that I use in my analysis. Figure 1.2 shows that generally between 120 and 150 are released each year. The typical movie employs between 3 and 11 non-supervisory grips and lighting technicians. These crews on average are supervised by about 2 chief lighting technicians and 2 key grips. In total, I observe nearly 19 thousand workers working in grip and lighting technician positions during this period, including 4268 in a key supervisor role.

1.3.2 Grip and lighting technician careers

Workers enter my sample the first year they appear on one of these major California movie productions—what I call their *initial year*—as either a grip or lighting technician.[25] I observe production dates for only a small fraction of movies. Because most movies take about a year to be released after production ends, I use instead the release year as a rough proxy for when they worked on the film.

I create a subsample of grips and lighting technicians (the *career sample*) to study early career outcomes during the decade that follows their initial year. I exclude grips and lighting

[24]For example, in Appendix Figure A1, I show that California's share of US motion picture production employment between 1980 and 2012 is above 50 percent. Nearly all of this employment is in establishments located in Los Angeles County. For comparison, New York State generally accounts for only 15–20 percent of employment.

[25]Some workers appear in both grip and lighting technician crews. In these cases, I assign workers to the position they appear in first. A small number of workers appear in both grip and lighting technician crews during their initial year. I assign these workers to be grips.

technicians who appear before 1988 in order to create a more consistent sample of workers at an early stage of their career. Although I have screen credits for movies released between 1980 and 2012, as we will soon see, a few years can pass between work on the movies in my sample, during which grips and lighting technicians subside on smaller scale productions, such as commercials, video, or television. Workers who first appear on movies in the early 1980's may have actually appeared in a major release on IMDb in the 1970's. These workers are less likely to be at an early stage of their career than those that appear later in my sample. I also exclude workers who appear after 2002 so that each worker has 11 years of data, including their initial year. The career sample of grips and lighting technicians is balanced in terms of the number of years included in the sample, but career outcomes are based on the number of movies released, which varies somewhat across cohorts. Finally, I drop a small fraction from the career sample who appear in a key supervisor position during their initial year since these workers are already at an advanced stage of their career.[26]

Table 1.3 describes year-level job outcomes for the 6826 grips and lighting technicians in the career sample, tabulated by the number of years since the worker's initial year (i.e., year 0). Job outcomes are the total number of movies that year on which the worker is credited in their assigned position. *Key jobs* are credits in a key supervisory position (i.e., key grips for grips and chief lighting technicians for lighting technicians). *Regular jobs* are non-supervisory credits.[27] By construction, all workers in their initial year have at least 1 regular job and 0 key jobs.

I find that employment on these major California movies is low. At the beginning of their career, each year only about 20 percent of grips and lighting technicians work on 1 movie or more. This fraction falls over the decade. I also find that over the decade a small fraction of workers are promoted and begin working as keys.

Figure 1.3 examines the distribution of the total jobs worked over the decade following the worker's initial year (i.e., years 1–10).[28] I exclude key jobs from this count since I am interested in the effect of network connections to keys on career outcomes. I find that while about half of workers will not work another job after their initial year, a small fraction work on average roughly 1–2 a year.[29] The high skew of the distribution implies a large variance relative to its mean, nearly 7 to 1.[30] The top 6 percent of workers who work 10 or more jobs account for nearly 70 percent of this variance.[31] The main goal of this paper is to understand how networks formed by the end of the initial year contribute to this dispersion.

[26] I also drop workers from the career sample who work as a key supervisor in a year before their initial year on a movie outside of my sample of major California productions.

[27] I count best boy credits as regular jobs.

[28] I tabulate the complete distribution of total jobs worked during years 1–10 in Appendix Table A1.

[29] Previous research studying Hollywood's labor force using screen credit data has also found a large number of workers who acquire a single credit. See, for example, Faulkner and Anderson (1987).

[30] The sample mean of the total number of jobs worked is 2.3. The variance is 15.6.

[31] Because the total number of jobs is discrete, I can decompose the variance into the weighted average of the squared deviation of the count from the mean. Let Y_i denote the total number of jobs worked by worker i. Then $\mathbb{E}\left[(Y_i - \mathbb{E}(Y_i))^2\right] = \sum_y (y - \mathbb{E}(Y_i))^2 P(Y_i = y)$.

1.3.3 Key networks

A worker's *key network* is composed of all keys who have supervised the worker on a movie released in an earlier year. Figures 1.4 and 1.5 show how these networks evolve over time. At the beginning of year 1, most workers have been supervised by only 1 to 3 keys. This variation is generated by differences in both the number of keys per movie as well as the number of movies a worker works on during their initial year. At the beginning of year 10, most workers have only worked an additional movie since their initial year, and their networks are still small. However, the top performing workers build a large network of keys that can potentially provide them work. For instance, the top 1 percent of workers have worked with 40 or more keys by the beginning of year 10.

1.4 Initial network quality and career outcomes

I measure a worker's initial network quality by their *initial connections*: the number of movies released during years 1–10 that are managed by a key who supervised the worker during the worker's initial year. This quantity represents the number of jobs that a worker's initial supervisors can pass on to them.

The top panel of Figure 1.6 shows that a worker's number of initial connections is strongly correlated with the number of movies they work over the next decade. In this figure, I group initial connections into 10 decile bins and then plot the average number of regular jobs worked among the workers in that bin. I find that workers who start in the upper half of initial network quality distribution work on about 84 percent more movies than those in the lower half. A least squares fit suggests that almost 1 in 10 initial connections become jobs.[32]

In Figure 1.3, I find that most workers in my sample work on at most 1 more movie after their initial year, while a small group work on average 1–2 jobs a year. The bottom panel of Figure 1.6 examines whether initial connections are related to this pattern. The points labeled "None" plot the fraction of workers who work 0 jobs during their decade in the career sample. The points labeled "10 or more" plot the fraction who work 10 jobs or more. Although workers who start their careers in higher quality networks are more likely to work on at least 1 movie during the next decade, large fractions fail to find any work regardless of their initial connections. In contrast, workers who begin their careers in higher quality networks are much more likely to work on 10 or more jobs. For instance, workers in the eighth quantile are only 20 percent less likely than those in the third quantile to never find work again, but they are nearly 100 percent times more likely to work on 10 or more jobs.

A limitation of this analysis is that workers who start in higher quality networks may be more likely to find work than those in lower quality networks for reasons other their initial connections. For example, the number of key supervisors in a worker's initial network is based on the number of movies they work on during their initial year, which suggests they are more employable for potentially unobserved reasons. In an analysis not shown, I find similar relationships between career outcomes and initial connections even controlling for

[32]In a bivariate regression of total jobs on initial connections, initial connections have an effect of 0.080 (S.E. 0.006). The R^2 of this model is 0.05.

the number of movies they work on during their initial year. Nevertheless, the nonrandom assignment of workers into initial networks is a potentially serious issue. I examine this issue more closely in the next section.

1.5 Job outcomes of worker-movie dyads

I create a dataset of worker-movie dyads for assessing the importance of network relationships with keys on workers' job outcomes (the *dyad sample*). The dyad sample contains every pairwise combination of the workers in the career sample with the set of major California movies released during years 1–10. For example, if a grip first appears in 1988, I include in this sample all the dyads constructed between this grip and the movies released during the years 1989 through 1998.

I measure two kinds of relationships to the key supervisors on a movie. The first kind, a *direct connection*, is illustrated in Figure 1.7. A worker has a direct connection if a key on the movie is a member of the worker's network. Again, a worker's network includes all keys who have supervised the worker in an earlier year. I measure the strength of a direct connection based on the number of times a worker has been supervised by a key. For example, a grip has a direct connection in which they "worked together 3 times" with a key on a movie released in 1995 if there is a key grip on the movie that has supervised the grip on 3 movies released in 1994 or earlier. Connection strength categories are mutually exclusive. Table 1.4 shows descriptive statistics of network characteristics and job outcomes for the worker-movie dyad sample. Direct connections are rare, which is a result of the large numbers of workers in the career sample who never work on more than 1 or 2 movies. Under 2 percent of dyads have some direct connection.

I also measure *indirect connections*. A worker has an indirect connection if they do not have a direct connection, but there is a key in their network who has worked alongside a key on the movie. Figure 1.8 illustrates this type of connection. Indirect connections are much more common than direct connections. Table 1.4 shows that workers have an indirect connection on 7.5 percent of worker-movie dyads.

Table 1.5 provides strong evidence that these connections matter for finding work. Indirectly connected workers are over 3 times more likely to get hired in a regular position than a worker with no direct or indirect connection (i.e., *unconnected workers*).[33] Direct connections are even more powerful: Workers who have been supervised once with a key are more than 30 times more likely to get hired. The effect of direct connections increases rapidly with the number of times workers have been supervised. For instance, movies hire over 1 in 5 workers who have been supervised 3 or more times with a key on the movie.

An important issue with this interpretation is that key networks accumulate as workers gain more experience working on movies. As a result, the workers with the most connections

[33] I call a worker-movie dyad unconnnected if the worker has neither a direct nor an indirect connection on that movie. Many of these workers are more distantly connected to the movie. For example, a worker may have worked with a key who has worked with a key who has worked with a key on the movie. I am also excluding indirect connections based on experience working alongside a regular worker who has been supervised by a key on the movie. In future work, I will quantify the prevalence of these other types of connections and examine their relationship with job outcomes.

are also more likely to be hired for other, potentially unobserved, reasons. One way to assess this bias is to compare the probability of being hired on a given movie among workers with the same number of years in the sample and who have worked the same number of jobs. Among subgroups, some workers will have a direct connection to a key on a movie and some will not. Table 1.6 tabulates the job outcomes of connected workers 1, 5, and 10 years following their initial year. I measure previous jobs by the total number of regular jobs in past years. For each year, I group together job counts at the upper tail of the distribution so that I have a relatively large number of workers in each bin. I present the percent of dyads that are employed (1) without any connection (direct or indirect) on the movie, (2) with an indirect connection on the movie, and (3) with a direct connection on the movie. For indirect and direct connections, I also calculate the implied change in the log odds of employment relative to unconnected workers.[34] In the model in the next section, I assume hiring probabilities take a logit form. If the positive effect of connections is due to bias, then I should find a large drop in the change in the log odds of employment once I condition on the number of previous jobs.

Table 1.6 shows that connections have a large impact on the job outcomes of otherwise similar workers. Although workers who have worked on more jobs are more likely to be hired than those who have worked less, connected workers—even inexperienced ones—are much more likely to be hired than unconnected workers. For example, in year 5, a directly connected worker who has worked only 1 job is more than 3 times as likely to be hired than an unconnected worker who has worked 9 jobs.

Table 1.6 also provides evidence that simple comparisons of workers with and without connections overestimate the importance of connections. The change in the log odds of employment from an indirect or direct connection is smaller conditioning on the number of previous jobs than for the entire sample of workers in the same year of their career. This bias strengthens over the course of a worker's career as those with more successful employment outcomes over time sort into higher quality networks. Nevertheless, the over all magnitude of these changes in the log odds conditioning on the number of previous jobs suggests that connections have a powerful effect for workers at various stages of their career.

An alternative approach for assessing this bias is to compare the probability of being hired among workers who are directly connected to the same number of movies during the year the movie is produced. I call the number of movies a worker is directly connected to that year their *total connections*. Relationships with keys should affect employment only on the jobs they supervise, but nonrandom sorting of more employable workers into networks that are connected to many movies will make it appear as if connections to keys matter even on jobs they do not supervise. Table 1.7 shows the results of this exercise. Each row presents the percent of dyads in which the worker is hired conditioning on the number of years since they have entered the sample and their total connections. I find the return to having a key supervisor in a worker's network is specific to jobs that the key manages. For example, during year 1, a worker with 1 total connection that year, but no direct or indirect connection

[34] Formally, let $p(t, x, 0)$ denote the fraction of dyads with a worker in year t of their career, number of previous jobs x, and no direct or indirect connection that are hired on the movie. Let $p(t, x, s)$ denote the fraction with a connection of strength $s \in \{indirect, direct\}$ that are hired. Then, the change in the log odds is $\Delta(log\, odds) \equiv \log\left(\frac{p(t,x,s)}{1-p(t,x,s)}\right) - \log\left(\frac{p(t,x,0)}{1-p(t,x,0)}\right)$.

on the particular movie, is only about 0.03 percentage points more likely to be hired than an unconnected worker without any total connections. But, if the direct connection is to a key on the movie, then they are more than 8 percentage points more likely to be hired.

This analysis of the dyad sample confirms that relationships with keys on movies play an important role in the hiring process. Observably similar workers experience different employment outcomes depending on the strength of their connection to a key on the movie. This effect is specific to the movie the key is on, suggesting a causal relationship. We also find some evidence of modest bias from nonrandom sorting of workers into higher quality networks.

1.6 Model

1.6.1 Setup

I develop a model for workers' hiring outcomes during the decade after their initial year of work on a major California movie. The model describes how workers' networks of relationships with keys evolve over time as they and the keys accumulate experience working on movies. It also allows for hiring outcomes to depend on time invariant unobserved factors.

Each year, t, a set of movies, $\{m \in \mathbb{M}(t)\}$, hires crews. To simplify the exposition, assume that there are exactly M movies each year. Crews consist of key and regular jobs. First, movies hire key supervisors. Then, movies hire regular jobs based on the recommendations of the key supervisors. I assume that these two types of jobs are performed by two mutually exclusive groups: keys and (regular) workers. Ignoring the possibility that workers transition into supervisory key roles allows me to model key hiring outcomes as exogenous to the hiring of the regular workers. As we saw in Table 1.3, only a small fraction of workers in my career sample move between regular and key jobs. Let K denote the number of keys.

Let d_{km} be an indicator that equals 1 if key k works on movie m. Let D_{kt} be the $M \times 1$ vector of key k's stacked hiring outcomes on movies in year t. Let $\mathbf{D}_s^t$ denote the $K \times M(t-s+1)$ matrix of key hiring outcomes on all movies produced between years s and t:

$$\mathbf{D}_s^t \equiv \begin{bmatrix} D'_{1s} & \cdots & D'_{1t} \\ \vdots & \ddots & \vdots \\ D'_{Ks} & \cdots & D'_{Kt} \end{bmatrix}$$

Each movie m has a vector of R observable attributes, z_m (e.g., genre, budget, number of keys hired). Let $\mathbf{Z}_s^t$ denote the $R \times M(t-s+1)$ matrix of attributes on all movies produced between years s and t.

I observe hiring outcomes for a sample of regular workers, indexed $i = 1, ..., N$. A worker enters my career sample in year $\tau(i)$, hired on at least one movie that year. Let a_{im} be an indicator that equals 1 if worker i works on movie m. Let A_{it} be the $M \times 1$ vector of worker i's stacked hiring outcomes on movies in year t. Each worker has a time invariant unobserved effect, α_i, that affects their probability of being hired on all movies. I assume the effect is drawn from a multinomial distribution with Q points of support.

I model regular worker i's hiring outcomes between years $\tau(i)+1$ and $\tau(i)+10$ jointly with their unobserved effect, α_i. Given worker i's hiring outcomes during their initial year, $A_{i\tau(i)}$ and all movie attributes and key hiring outcomes, the joint probability of worker i's outcomes and unobserved effect is

$$P\left(A_{i\tau(i)+10},...,A_{i\tau(i)+1},\alpha_i|A_{i\tau(i)},\mathbf{D}_0^{\tau(i)+10},\mathbf{Z}_0^{\tau(i)+10}\right)$$
$$=\left[\prod_{t=\tau(i)+1}^{\tau(i)+10}P\left(A_{it}|A_{it-1},...,A_{i\tau(i)},\alpha_i,\mathbf{D}_0^{\tau(i)+10},\mathbf{Z}_0^{\tau(i)+10}\right)\right]P\left(\alpha_i|A_{i\tau(i)},\mathbf{D}_0^{\tau(i)+10},\mathbf{Z}_0^{\tau(i)+10}\right) \tag{1.1}$$

Notice that equation (1.1) includes key hiring outcomes starting in year 0, the first year I observe outcomes on major California movies.[35] Assuming hiring outcomes depend on key outcomes dating back to year 0 allows for indirect relationships between worker i and keys that are based on keys working together on movies that were produced before worker i arrives in the sample.

Key hiring outcomes and movie attributes in years $s>t$ do not affect regular worker i's hiring outcomes in year t once I condition on hiring outcomes and movie attributes in earlier years. In addition, I assume worker i's hiring outcomes are independent conditional on outcomes on movies produced in prior years, their unobserved effect, α_i, key outcomes, and movie attributes:

$$P\left(A_{it}|A_{it-1},...,A_{i\tau(i)},\alpha_i,\mathbf{D}_0^{\tau(i)+10},\mathbf{Z}_0^{\tau(i)+10}\right)=\prod_{m\in\mathbb{M}(t)}P\left(a_{im}|A_{it-1},...,A_{i\tau(i)},\alpha_i,\mathbf{D}_0^{t},\mathbf{Z}_0^{t}\right) \tag{1.2}$$

Equation (1.2) implies that workers accept all work regardless of the number they have already accepted that year—their schedules cannot become too "busy." Although in reality workers cannot work on two jobs at once, I find that this model fits the data well. One reason is that few workers find any work in my sample at all, so most workers are available to accept an additional offer.

I find worker i's contribution to the observed sample likelihood by plugging equation (1.2) into equation (1.1) and averaging over all Q types of worker effect, α_i:

$$\mathcal{L}_i\equiv\sum_{q=1}^{Q}\left[\prod_{t=\tau(i)+1}^{\tau(i)+10}\prod_{m\in\mathbb{M}(t)}P\left(a_{im}|A_{it-1},...,A_{i\tau(i)},\alpha_q,\mathbf{D}_0^{t},\mathbf{Z}_0^{t}\right)\right]P\left(\alpha_q|A_{i\tau(i)},\mathbf{D}_0^{\tau(i)+10},\mathbf{Z}_0^{\tau(i)+10}\right) \tag{1.3}$$

The sample log likelihood is then $\ln\mathcal{L}=\sum_{i=1}^{N}\ln\mathcal{L}_i$.

1.6.2 Hiring model

I posit a simple threshold model for $P\left(a_{im}|A_{it-1},...,A_{i\tau(i)},\alpha_i,\mathbf{D}_0^{t},\mathbf{Z}_0^{t}\right)$, the probability that movie m hires regular worker i. Hiring outcomes depend on: (1) the unobserved worker

[35] In my sample, year 0 corresponds to 1980.

effect, α_i, (2) observable worker characteristics, x_{it}, (3) the strength of worker i's connection to the keys on movie m, C_{im}, and (4) movie attributes, z_m.

Worker characteristics, x_{it}, include measures of worker i's experience on different types of movies and the number of movies worked on in recent years. Therefore, x_{it} is implicitly a function of worker i's hiring outcomes and movie attributes in prior years. Similarly, the strength of worker i's connection to the keys on movie m, C_{im}, is a function of worker i's hiring outcomes in past years, key hiring outcomes in past years, and movie m's choice of keys.

Finally, for each movie produced in year t, worker i draws an i.i.d. productivity shock, ϵ_{im}, from a logistic distribution. Movie m hires worker i if and only if

$$\alpha_i + x_{it}'\beta + f(C_{im};\theta) + g(x_{it}, z_m;\gamma) - \epsilon_{im} > 0 \tag{1.4}$$

Equation (1.4) implies that the likelihood that movie m hires worker i is

$$\begin{aligned} P\left(a_{im}|A_{it-1},...,A_{i\tau(i)},\alpha_i,\mathbf{D}_0^t,\mathbf{Z}_0^t\right) &= \Lambda\big(\alpha_i + x_{it}'\beta + f(C_{im};\theta) + g(x_{it}, z_m;\gamma)\big)^{a_{im}} \\ &\times \Big(1-\Lambda\big(\alpha_i + x_{it}'\beta + f(C_{im};\theta) + g(x_{it}, z_m;\gamma)\big)\Big)^{1-a_{im}} \end{aligned} \tag{1.5}$$

where $\Lambda(\cdot)$ is the CDF of the logistic distribution. The vector θ is the parameter of interest. It governs the causal effect of the network connections to keys on hiring outcomes. $f(\cdot)$ is a function of indicator variables for connections of different strength, as in Table 1.5. The parameter β is the effect of the other observable worker-level variables. $g(\cdot)$ is a function that measures the quality of the match between worker i and movie m based on their observable characteristics. For instance, we might expect workers who have worked on a number of action movies to be more qualified to work on action movies than workers whose experience is mainly on comedies. In addition, $g(\cdot)$ includes movie-level characteristics, such as budget or genre, that might affect the expected number of workers hired. The parameter γ measures the importance of these match and movie effects.

1.6.3 Type model

I address the network-version of the initial conditions problem (Heckman, 1981; Wooldridge, 2005) by specifying $P\left(\alpha_i|A_{i\tau(i)},\mathbf{D}_0^{\tau(i)+10},\mathbf{Z}_0^{\tau(i)+10}\right)$, the distribution of the worker effect conditional on job outcomes during their initial year and key outcomes during the period they are in my sample.

In the type model, I include information about the quality of worker i's network known at the end of their initial year to account for nonrandom sorting of more employable workers into higher quality networks. In Section 1.4, I introduced a measure of a worker's initial network quality called their initial connections. In this context, initial connections are the number of movies during years $\tau(i)+1$,...,$\tau(i)+10$ that hire a key who supervised worker i in year $\tau(i)$. In contrast to the connection variables included in the hiring model, initial connections are only a function of variables that I take as given: worker i's hiring outcomes in year $\tau(i)$, $A_{i\tau(i)}$, and key hiring outcomes, $\mathbf{D}_0^{\tau(i)+10}$. I also include in the type model a worker's *initial* *indirect connections*: The total number of movies during years $\tau(i)+1, ..., \tau(i)+10$ that do not have a direct connection to worker i but have a key who worked in the past with a key who supervised i during year $\tau(i)$. Let IC_i denote the vector containing the values of worker i's initial direct connections and their initial indirect connections.

Finally, I include in the type model a vector of indicator variables for the number of movies worker i works on during year $\tau(i)$, w_i. I expect workers who are hired on multiple movies in year $\tau(i)$ to have higher values of the worker effect, α_i, since they are able to acquire this work without observable network connections.

I assume the type model takes the form of a multinomial logit:

$$P\left(\alpha_i | A_{i\tau(i)}, \mathbf{D}_0^{\tau(i)+10}, \mathbf{Z}_0^{\tau(i)+10}\right) = \frac{\exp\left(\pi_q^0 + w_i'\pi_q^w + IC_i'\pi_q^c\right)}{1+\sum_{s<Q}\exp\left(\pi_s^0 + w_i'\pi_s^w + IC_i'\pi_q^c\right)} \tag{1.6}$$

Equation (1.6) allows for higher ability workers to sort into networks that will produce more connections during their career. Modeling this sorting explicitly mitigates the selection bias in estimates of the network effect, θ. Wooldridge (2005) argues for specifying the distribution of the unobserved effect conditional on initial conditions. Nevertheless, Equation (1.6) is admittedly ad hoc. As a robustness check, in Section 1.8, I estimate the model's network effects employing a fixed effects approach.

1.7 Results from random effects models

I estimate a series of models based on Equations (1.1)–(1.6) on the worker-movie dyad sample I introduced in Section 1.5. There are 15 cohorts of grips and lighting technicians in this sample, who first appear on movies released between 1988 and 2002. Each cohort is in the sample for 10 years. I exclude dyads from the sample log likelihood when the worker is hired in *any* key role, regardless of position.[36] For example, a dyad for a grip is excluded if the grip is hired on the movie as either a key grip or a chief lighting technician. Appendix Table A2 provides additional summary statistics for this sample. I estimate standard errors using the sandwich variance formula and cluster at the worker-level.[37]

[36]In addition, when computing variables such as the number of previous jobs or connections, I use only regular jobs. For example, if a grip has worked with a key grip on a movie twice, once as a regular grip and once as a key grip, I would consider this relationship a direct connection in which the grip and key have worked together only once. I explicitly model hiring outcomes in only regular jobs, and I would be treating this key experience as exogenous if I included it in measuring these variables. Key jobs make up a small fraction of the number of jobs worked in the career sample, so I do not expect this choice to qualitatively affect my results.

[37]I compute the standard errors using the following expression. Let $\hat{\varphi} \equiv \left(\hat{\beta}', \hat{\theta}', \hat{\gamma}', \hat{\pi}'\right)'$ denote the vector of estimated parameters. Let $\widehat{\mathbf{H}(\hat{\varphi})}$ denote the negative of the average Hessian over the sample of N workers. Let $\mathbf{s}_i(\hat{\varphi})$ denote worker i's estimated score vector, $\mathbf{s}_i(\hat{\varphi}) \equiv \left.\frac{\partial \mathcal{L}_i(\varphi)}{\partial \varphi}\right|_{\varphi=\hat{\varphi}}$. Then I estimate the variance of $\hat{\varphi}$ using $\hat{\mathbb{V}}(\hat{\varphi}) = \frac{1}{N}\widehat{\mathbf{H}(\hat{\varphi})}^{-1}\frac{1}{N}\left(\sum_{i=1}^N \mathbf{s}_i(\hat{\varphi})\,\mathbf{s}_i(\hat{\varphi})'\right)\widehat{\mathbf{H}(\hat{\varphi})}^{-1}$.

1.7.1 Parameter estimates

Tables 1.8–1.10 show results from fitting models that assume different numbers of types in the distribution of the unobserved worker effect, $P(\alpha_i)$. Model (1) assumes only 1 type of worker, which is the same as fitting a simple logistic regression on the dyad sample. Models (2) and (3) assume 2 and 3 types of workers, respectively. Other than the number of types, all models include the same sets of variables in the hiring and type models. In addition to the sets of worker characteristics, connection variables, match effects, and movie characteristics shown, all models include indicators for each release year and each year of a worker's career they are in my sample.

Table 1.8 shows estimates of the parameters in the hiring model described in Section 1.6.2. I find network connections have a precisely estimated and powerful effect on hiring outcomes. This effect increases in the strength of the connection. For example, consider a grip in year 1 who has worked only 1 job. Ignoring match effects,[38] the estimates in Model (1) imply that an indirect connection raises the probability of employment on a low budget movie with only 1 key grip from about 0.09 to 0.23 percentage points. A grip who has worked only 1 time with a key on the movie has about a 2 percentage point chance of employment. I find in Models (2) and (3) that adding more types of workers to the model does not significantly attenuate the impacts of these connection variables. If anything, adding more types marginally strengthens their impact.

In contrast, adding multiple types to the model significantly changes the effect of the number of previous jobs worked. I include in the model indicator variables for each number of jobs between 2 and 4 and bin 5–9 and 10 or more. By construction, everyone in the sample has worked at least 1 regular job. In Model (1), more jobs worked increases the probability of employment. The largest increase comes from the second job. One interpretation of this estimate is that workers are more likely to qualify for membership in the union after working a couple of movie jobs. Another interpretation is that most workers in these positions are only loosely attached to the job, lacking the ability or interest in having a movie career, and the estimated returns to experience are biased by the selection of more employable workers into movie jobs. Consistent with this latter interpretation, the impact of the second job drops to nearly 0 when we assume 3 types of workers in Model (3). In this model the probability of employment is decreasing in the number of jobs worked after the second movie. One explanation of this result is that keys are reluctant to hire unconnected experienced workers.

I include among the worker characteristics time-varying measures of the worker's attachment to the industry, measured by the number of previous jobs worked in (1) the prior year and (2) the year before the prior year. These variables capture a form of state dependence in my setting. For example, workers who do not work any jobs in the prior year may have switched to another sector, like television, or retired. In both cases they would be less likely to work on movies in the current year. I find positive effects of these variables that are robust to the number of types in the models. I also include an interaction of these two variables. I estimate a small negative coefficient on their interaction.

The intercepts in the hiring model in Model (3) for the three types are all large and neg-

[38] In these specifications, ignoring match effects for calculating the probability of getting hired on, for example, a low-budget movie is equivalent to assuming the worker's experience is on a non-low budget movie in a non-overlapping group of genres.

ative, reflecting the small probability of employment on any particular movie. Nevertheless, I find large differences in employment propensities across types. The "highest" type, Type 1, worker has a 3 higher log odds of employment than an otherwise identical "lowest" type, Type 3.

In Table 1.9, I use the distribution of worker characteristics during their initial year to estimate the proportions of each type. First, I calculate the workers' prior probabilities of each type, as implied by the estimated type model. I then take the average of the type probabilities across the sample of 6826 workers. The majority of workers are the lowest type, Type 3. A small fraction, about 14 percent, are the highest type, Type 3.

Table 1.9 also shows estimates of the parameters in the Type Model in Equation (1.6). The model is a multinomial logit, where I have normalized the parameters for the third type to be 0. Workers who work on multiple jobs their initial year tend to be higher type. Nevertheless, holding constant job outcomes, workers who initially sort into high quality networks, measured by the number of initial connections and initial indirect connections, are most likely to be Type 2, rather than Type 1.

Table 1.10 shows additional estimates for other covariates included in the hiring model. These variables are observable characteristics of the movie and worker-movie match. In addition to these variables and those shown in Tables 1.8, all models include a full set of year dummies for 1989 through 2012 and dummies for each year of the worker's career.

I include in each model in Table 1.10 movie-level characteristics based on their production budget and genre. I partition budget into 6 categories.[39] I find that bigger budget movies hire more workers. I include in the models the 12 most prominent genres in my sample. IMDb assigns movies genres using overlapping tags. Controlling for budget, comedies, dramas, and horror movies use smaller crews. I also include counts of the number of keys in the worker's position. Movies that hire more keys are also more likely to hire larger crews of regular workers.

I create measures of the quality of the worker-movie match by interacting these movie-level characteristics with worker-level experience on movies of the same budget or genre. Including these variables controls for the most obvious forms of worker-movie-level heterogeneity that might bias my estimates of the network effects. I find some evidence of significant match effects, but the magnitudes of these impacts are smaller than what I found for network connections. For example, an additional job working on a movie in the same budget category increases the hiring probability about 3 percent. I find larger match effects for genres that tend to use more specialized and sophisticated equipment, such as adventure, sci-fi, and fantasy movies.

1.7.2 Model fit

Table 1.8 shows large increases in the log likelihood when I add multiple types to the model. Figure 1.9 provides visual evidence for this improvement in fit. It plots a histogram of total number of regular jobs worked during years 1–10 of the worker's career. I compare the empirical distribution, the bars, to predictions based on an average of 100 simulations of

[39] I measure production budgets in 2012 dollars, adjusted using the CPI. Budget information is missing for about 20 percent of major California movies released between 1988 and 2012.

Models (1) and (3).

I perform 100 simulations of each model for each of the 15 cohorts who enter my career sample between 1988 and 2002. For a cohort who arrives in year c, I first draw their type from the prior distribution of the worker effects, using the observed hiring outcomes during the workers' initial year. At the end of each year, I update workers' experience measures and key networks. Then at the start of the next year, I calculate connections on movies taking as given all key job assignments on the movies released that year. I then assign workers regular jobs using productivity shocks drawn from a logistic distribution.

Figure 1.9 shows that Model (1), which assumes only 1 type of worker, understates the dispersion of job outcomes in my sample. It under predicts both the number of workers who do not work at all during the decade following their initial year as well as those who on average work 1 or more job a year. Adding more types to the model improves the fit. We see that Model (3) closely matches the empirical distribution through the 90th percentile, though it still tends to over predict the number of workers hired on 18 or more jobs.

The models are also able to capture the positive relationship between job outcomes and workers' initial connections. Figure 1.10 plots the observed average number of jobs worked among workers by initial connection decile, as in Figure 1.6. The lines labeled "1 type" and "3 types" connect the predicted number of jobs in that decile based on Models 1 and 3, respectively.

I find that adding more types to the model improves its ability to predict this relationship. Relative to Model (1), Model (3) more closely predicts the average number of jobs worked in 8 out of 10 decile bins. The improvement is particularly notable in the 10th decile, though Model 3 still over predicts the average number of jobs worked by about 18 percent.

1.7.3 Nonrandom sorting of workers into initial networks

The significant estimates of initial connections in Table 1.9 suggest that higher type workers are more likely than lower type workers to be supervised in their initial year of work by keys who will have more successful careers. Table 1.11 provides more direct evidence. To construct this table, I create a sample of *initial keys*: The set of keys who supervise workers during their initial year. I select each cohort's initial keys by position (i.e., grip or lighting technician), so in total the sample is composed of $(15 \times 2 =)\, 30$ cohort-position specific subsamples. Key supervisors can appear in multiple subsamples. I then divide this sample into 5 quintiles based on the number of movies managed by a key during the decade their corresponding cohort appears in the dyad sample. For example, if a key grip supervises a grip who initially works in 1990, I would determine which quintile the key grip is in by counting the number of movies they manage during years 1991–2000. The column labeled "Key jobs, years 1–10" shows the average number of movies managed by each key-cohort observation by quintile.

In a final step, I find the average characteristics of the workers initially supervised by the keys. The columns labeled "Type 1," "Type 2," and "Type 3" show the estimated proportion of each type of worker initially supervised by quintile. I find the proportion by computing each worker's posterior type probabilities[40] for the three types in Model (3) in Table 1.9 and taking the average across all workers connected to a key in that quintile group.

Table 1.11 shows that keys who have the strongest careers supervise fewer lowest type workers than keys who have less successful careers. Relative to workers initially supervised by a key in the lowest quintile, workers initially supervised by a key in the highest quintile are about 15 percent more likely to be the highest type, 30 percent more likely to be the medium type, and 16 percent less likely to be the lowest type.

1.8 Fixed effects approach

In Sections 1.6 and 1.7, I assume each worker's unobserved effect is the same each year and that I know its distribution conditional on what I observe about the worker during their initial year. In addition, I assume that I observe all characteristics of the movie that could affect the size of the regular worker crew. In this section, I relax these assumptions and develop a fixed effects estimator for the effect of network connections that places no restrictions on distribution of the unobserved worker and movie effects conditional on a worker's network connections. Comparing my estimates of the network effects using this alternative approach to what I find in Section 1.7, I conclude that my model's estimates of the network effects are robust to unobserved heterogeneity among workers and movies.

By assuming in Section 1.6 that unobserved worker effects do not change each year, I rule out some important behaviors that could bias my estimates of the network effects. For instance, Table 1.3 shows that the number of workers finding work on major California movies declines each year. Workers may be dropping out of the labor market or finding steady work in other sectors like television. If these changes in workers' circumstances are correlated with similar changes happening to the key supervisors in their networks, then I would find positive network effects even if connections do not affect hiring outcomes.

Unobserved movie-level factors could also bias the estimated network effects. Movies that are expected to hire larger crews might hire more experienced keys with larger networks. Because more workers are connected to the keys on these movies, estimates of network effects based on comparisons of hiring outcomes across movies released in the same year would be positively biased. In addition, these unobserved factors may depend on worker characteristics, such as the year they enter. Some movies might prefer to hire younger or older workers.

I incorporate these forms of unobserved worker and movie heterogeneity into my hiring model. Consider a worker, i, who enters my sample in year $\tau(i) = c$. Let W_{im} denote a vector of all the worker-movie dyad level covariates between worker i and movie m produced in year t. In the models estimated in Section 1.7, these covariates include indicators for different

[40] I find a worker i's posterior type probability applying Bayes' Rule:

$$P\left(\alpha_i = \alpha_q | A_{i\tau(i)+10}, ..., A_{i\tau(i)+1}, A_{i\tau(i)}, \mathbf{D}_0^{\tau(i)+10}, \mathbf{Z}_0^{\tau(i)+10}\right) = \frac{\mathcal{L}_{iq} P\left(\alpha_i = \alpha_q | A_{i\tau(i)}, \mathbf{D}_0^{\tau(i)+10}, \mathbf{Z}_0^{\tau(i)+10}\right)}{\mathcal{L}_i}$$

where $\mathcal{L}_{iq} \equiv P\left(A_{i\tau(i)+10}, ..., A_{i\tau(i)+1} | A_{i\tau(i)}, \alpha_i = \alpha_q, \mathbf{D}_0^{\tau(i)+10}, \mathbf{Z}_0^{\tau(i)+10}\right)$, the probability of observing their sequence of hiring outcomes conditional on their type.

strengths of a worker i's connection on movie m, as well as match characteristics, such as interactions between the number of action movies worked on in the past and whether movie m is an action movie. Let η_{it} denote the value of worker i's effect in year t. η_{it} captures both observed and unobserved factors that affect i's hiring outcome on any movie produced in year t. Let δ_{cm} denote a movie-level effect specific to all worker's who enter the sample in year c. δ_{cm} includes attributes about the movie I observe, like its budget, as well as ones I do not, like the characteristics of the keys it hires.

I now assume the hiring model for a movie m produced in year t is:

$$P\left(a_{im}|A_{it-1},...,A_{ic},\eta_{it},\delta_{cm},\mathbf{D}_0^t,\mathbf{Z}_0^t\right)=\Lambda\left(\eta_{it}+W_{im}'\phi+\delta_{cm}\right)^{a_{im}}\times \left(1-\Lambda\left(\eta_{it}+W_{im}'\phi+\delta_{cm}\right)\right)^{1-a_{im}} \tag{1.7}$$

where $\phi\equiv(\theta',\gamma')'$ are all the worker-movie level parameters and $\Lambda(\cdot)$ is the CDF of the logistic distribution. In contrast to Equation 1.5, this model allows for unobserved time-varying worker effects and unobserved cohort-specific movie effects.

In logit model panel settings with only one unobserved effect, consistent estimation of the parameters of the time-varying covariates is possible by conditioning on a sufficient statistic for the effect (Chamberlain, 1980). More recently, Graham (2014, 2015a) and Charbonneau (2014) extend this approach to settings with multiple unobserved effects. I build on their insight that identification is possible in these settings by focusing on subgraphs of dyads.

I propose an estimator of ϕ based on pairs of worker-movie dyads. Figure 1.11 illustrates a pair. A pair contains exactly two workers, i and j, who first take a job on a movie in the same year, c, and two movies, m and n that are produced in the same year, t.[41] Each pair contains 4 worker-movie dyads: im, jn, in and jm. There are 16 possible configurations of hiring outcomes of the two workers on these two movies, ranging from no movie hiring either worker to both movies hiring both workers. However, by conditioning on a particular subset of configurations occurring, the unobserved effects in Equation (1.7) do not affect the probability of the pair's outcome. Let

$$Y_{ijmn}=a_{im}a_{jn}(1-a_{in})(1-a_{jm})-(1-a_{im})(1-a_{jn})a_{in}a_{jm} \tag{1.8}$$

where a_{im}, a_{jn}, a_{in}, and a_{jm} are indicators for whether a worker is hired on the particular movie. As illustrated in Figure 1.12, $Y_{ijmn}=1$ when worker i is hired on movie m and worker j is hired on movie n, and $Y_{ijmn}=-1$ when worker i is hired on movie n and worker j is hired on movie m. For any other of the remaining 14 possible configurations of the two workers on the two movies, Y_{ijmn} takes a value of 0.

Notice that in both events, $Y_{ijmn}=1$ and $Y_{ijmn}=-1$, each movie hires exactly one worker and each worker is hired on exactly one movie. Given the logit form of the hiring model, the unobserved worker and movie effects then do not affect the relative probability of one of these two events occurring. In addition, the likelihood of these outcomes conditional

[41] I am ignoring worker i's and j's positions to simplify the exposition. In practice, I allow for position-specific unobserved movie effects by focusing only on pairs of worker-movie dyads in which the workers enter the sample in the same year *and* are in the same position.

on one of them occurring also takes a logit form:

$$l^{cond}_{ijmn} = \Lambda \left(\left((W'_{im} - W'_{in}) - \left(W'_{jm} - W'_{jn} \right) \right) \phi \right)^{\mathbb{I}(Y_{ijmn}=1)} \times \left(1 - \Lambda \left(\left((W'_{im} - W'_{in}) - \left(W'_{jm} - W'_{jn} \right) \right) \phi \right) \right)^{\mathbb{I}(Y_{ijmn}=-1)} \tag{1.9}$$

where $l^{cond}_{ijmn} \equiv P\left(Y_{ijmn} | A_{it-1}, A_{jt-1}, ..., A_{ic}, A_{jc}, \eta_{it}, \eta_{jt}, \delta_{cm}, \delta_{cn}, \mathbf{D}^t_0, \mathbf{Z}^t_0, Y_{ijmn} \in \{1, -1\}\right)$ is the marginal likelihood of the pair's hiring outcome conditional on either $Y_{ijmn} = 1$ or $Y_{ijmn} = -1$. Importantly, l^{cond}_{ijmn} does not depend on the worker effects or movie effects. Intuitively, when each worker is hired on a different movie, the particular arrangement is affected only by the relative strength of the connections and match quality of the workers on the two movies. The unobserved worker and movie effects do not matter.

Equation (1.9) motivates an estimator for ϕ that is the solution to a maximization problem in which the criterion sums over all sampled pairs' log marginal likelihoods:

$$\hat{\phi} = \arg\max_{\phi} \sum_{c=1988}^{2002} \sum_{t=c+1}^{c+10} \sum_{i:\tau(i)=c} \sum_{j<i} \sum_{m \in \mathbb{M}(t)} \sum_{n<m} \log l^{cond}_{ijmn} \tag{1.10}$$

In other words, I find $\hat{\phi}$ by fitting a binary logit model on the sample of all pairs available in my dyad sample. A pair contributes to the criterion only when $|Y_{ijmn}| = 1$. The dependent variable is an indicator that equals 1 when $Y_{ijmn} = 1$ and 0 otherwise. The covariates are a transformation of the dyad-level variables: $((W_{im} - W_{in}) - (W_{jm} - W_{jn}))$.

Notice that the criterion in Equation (1.10) is not the log likelihood for the pair sample. Pairs are only conditionally independent when they are constructed from mutually exclusive sets of worker-movie dyads. Rather, this criterion is a U-process. Honoré and Powell (1994) prove consistency and asymptotic normality of U-process minimizers under regularity conditions. Graham (2015a) extends these results to network settings. I construct asymptotic standard errors following Graham (2015a) in Chapter 2.

Table 1.12 compares estimates of the effect of network connections using this estimator to those I found using the random effects approach in Table 1.8. I construct a sample of all pairs based on the dyads in my dyad sample. As in Section 1.7, I exclude dyads in which the worker is hired as a key in either position. Each pair is composed of two movies released in the same year, t, and two workers who are both in the same position and enter in the sample in the same year, c. In order to satisfy the condition that the pair's hiring outcomes be one of the two configurations describes above, both workers must be hired on at least 1 movie in year t. As a result, the pair sample excludes workers who are not hired on any movie during the decade they appear in the dyad sample. The fourth column of Table 1.12 shows that I drop 3381 workers when employing the fixed effects estimator for this reason. In addition, each movie in the pair sample must hire at least 1 worker in the dyad sample, and I drop 574 movies. In total, there are over 400 thousand pairs constructed from the dyads in the dyad sample.

As a baseline, I include in Table 1.12 estimates of the network connections from a simple logit estimated on the dyad sample. Other than the connection variables, the only other covariates in the model are career-year effects and a control for the worker's position. The second and third columns display the results from Models (1) and (3) from Table 1.8, respectively. The final column, labeled "FE logit," presents estimates employing the fixed effects approach described above. In contrast to the previous models, this estimator controls for year-specific worker effects and cohort-specific movie effects.[42] In addition to the, appropriately transformed, connection variables, the model estimated in the fourth column includes all the match quality variables shown in Table 1.10.

Table 1.12 shows that the network effects presented earlier in Table 1.8 are robust to unobserved time-varying worker and movie heterogeneity. Comparing the estimates in the fourth column to those in the third, I find that most coefficients are only about 0.2 to 0.4 smaller in magnitude. By this measure, the effect of indirect connections suffers from the most bias relative to the effects I presented earlier. Nevertheless, the estimated effect controlling for these unobserved effects remains about 0.8. Comparing the results in the second, third, and fourth columns to those estimated in the first column, the baseline model, suggests that the observable covariates that I included in the models presented in Section 1.7 absorb most of the heterogeneity across workers and movies that can positively bias estimates of the network effects.

1.9 The role of initial network quality in career inequality

The model I presented in Section 1.6 enables me to predict worker job outcomes in counterfactual environments. In this section, I estimate the fraction of overall inequality in job outcomes for workers in a given cohort that is attributable to inequality in the career success of the keys who supervise them during their initial year of work by simulating a counterfactual that removes variation in the initial keys' career trajectories.

1.9.1 Setup

I first create a sample of initial key supervisors. To create this sample, I select for each cohort the set of keys who supervise the workers during their initial year. I select the initial keys by position (i.e., grip and lighting technician), so in total the sample is composed of $(15 \times 2 =)30$ cohort-position specific subsamples. All together, there are 4759 key-cohort observations.[43]

Figure 1.13 plots the distribution of the total number of movies managed by each initial key during the decade their cohort is in my sample (see the bars labeled "Observed"). This distribution is similar to what I observed in Figure 1.3 among regular workers, with a large number of initial keys not hired on any movies and a small minority hired on about 1 to 2 a year.

I simulate a counterfactual in which each initial key is replaced with one that instead works approximately the cohort-position-specific average number of jobs. For most cohorts of workers arriving between 1988 and 2002, the average initial key is hired on between 4 and

[42] In fact, because both workers in a pair belong to the same position as well as the same cohort, the fixed effects estimator controls for position-cohort-specific movie effects.

[43] I analyzed nonrandom sorting of workers into initial networks using the initial key sample in Section 1.7.3. Key supervisors can appear in multiple subsamples.

5 movies during the next decade. This procedure does not change the total number of jobs worked by initial keys, but the dispersion is nearly eliminated.

I perform the following steps for each cohort-position group of initial keys in implementing this counterfactual. Assume the group of initial keys are associated with a cohort of regular workers who initially take a job in year c:

1. I find all key jobs filled by an initial key on movies released during years $c+1, ..., c+10$.
2. I determine a minimum number of jobs (the *cap*) worked by each key under the counterfactual by dividing the total number of key jobs by the total number of initial keys in the group and rounding down to the nearest whole number.
3. I assign each job at random to an initial key who has not yet reached the cap. I repeat this step until each initial key in the group has reached the cap.[44]

The total number of jobs worked by a group of initial keys is not perfectly divisible by the size of the group, and there will be jobs remaining to be assigned after each initial key has reached the cap. Therefore, in a final step:

4. I randomly assign each remaining job to a different initial key.

Figure 1.13 also plots the distribution of the total number of movies managed by each initial key under the counterfactual. Most initial keys are now hired on either 4 or 5 jobs. Comparing the counterfactual distribution to that originally observed in the sample, the variance of total key jobs worked is reduced over 99 percent, from about 30 to 0.3.

1.9.2 Results

Table 1.13 shows that the counterfactual compresses the distribution of workers' initial connections. I compute descriptive statistics under the counterfactual by simulating the procedure above 100 times and taking averages across the replications. The dispersion in workers' initial connections is the result of two distinct sources of variation: (1) differences across workers in the number of keys who supervise them during their initial year (initial network size), and (2) differences in the number of movies managed by the initial keys. By construction, the counterfactual affects only the second source of variation. The rows in Table 1.13 labeled "Over all" shows that removing this source of variation cuts the total variation in initial connections roughly in half. Nevertheless, tabulating the descriptive statistics by initial network size confirms that the counterfactual essentially eliminates variation in initial connections within each bin. However, because the counterfactual does not change the total number of jobs worked by initial keys, it does not reduce the average number of initial connections.[45]

[44] Keys cannot occupy more than 1 slot on a movie.

[45] In fact, the average number of initial connections is slightly higher under the counterfactual. I believe this increase is due to a tendency for keys to work together over time. When two keys in a worker's initial network appear on the same movie, it counts as only 1 initial connection for a worker. When I randomly assign jobs to initial keys, I remove this redundancy.

Table 1.14 summarizes the effect of the counterfactual on worker outcomes by their type. I predict worker outcomes using the estimates from Model 3 in Tables 1.8, 1.9, and 1.10, following the steps I described in Section 1.7.2.[46] I compute the means and variances under the columns labeled "Observed" and "Counterfactual" by taking averages over 100 replications.[47]

I find the counterfactual reduces the variance of predicted number of jobs worked between years 1–10 about 50 percent, implying that half of the dispersion in career outcomes in my sample is generated by differences in the career trajectories among initial key supervisors. The counterfactual also reduces the mean by 18 percent. Examining the effects by type, I find that these reductions are driven by the effects on high and medium type workers, Types 1 and 2. For example, the model predicts that given the key assignments observed in the sample, Type 1 workers work on average 10 jobs during the decade after they enter my sample, but under the counterfactual, Type 1 workers work only on average 8 jobs. I find the effect on the variance of Type 1 workers jobs is even larger, falling over 60 percent.

Figure 1.14 compares the distribution of worker career outcomes under the counterfactual to the distribution that I predict using the observed key assignments in my sample. Reductions in the upper tail of the distribution generate most of the changes I observed in Table 1.14. Figures 1.15–1.17 plot the distribution by worker type and show a similar pattern. For example, the counterfactual reduces the 90th percentile for Type 1 workers from about 20 jobs to 13.

Table 1.14 also shows the counterfactual's effects on workers' total connections, the number of movies released between years 1–10 on which the worker has a direct connection. The negative effect on connections based on keys who have supervised the worker 3 or more times is large, nearly 70 percent. Intuitively, by capping the number of jobs worked by initial keys, I have reduced workers' ability to form strong relationships, which were an important source of jobs during the latter part of their career.

1.10 Conclusion

Network-based connections are pervasive in hiring and mobility patterns. While the theoretical impacts of network connections on the inequality of labor market outcomes are well known (e.g., Calvo-Armengol and Jackson, 2004), the empirical magnitude of actual network effects is less certain. A key issue is the difficulty of disentangling the causal effect of network connections from differences in the characteristics of workers in better and worse networks. I study this question using data on freelance workers in Hollywood, who are hired for short term jobs through an informal process that relies in part on previous connections. I assemble a dataset of screen credits for over 3700 major California productions and follow cohorts of freelance grips and lighting technicians who first take a job on a major movie production between 1988 and 2002. I develop two alternative models for the probability that workers are

[46]In each replication, I draw a worker's type from the prior distribution of the worker effects, using the observed hiring outcomes during the workers' initial year in my original sample. Drawing the worker's type instead from the posterior distribution of worker effects yields similar results.

[47]Each replication of the counterfactual uses a unique simulation of the key job assignments.

hired on subsequent productions—based on random effects and fixed effects specifications—that incorporate network effects, experience effects, and unobserved heterogeneity.

I find that in freelance labor markets like Hollywood the fortunes of an individual worker are closely linked to the careers of key agents who make the hiring decisions for jobs. In my setting, workers who know position supervisors who manage more jobs will have more job opportunities. Career outcomes are highly skewed among the workers in my sample. After they first appear credited on a movie, most accrue at most only 1 more credit during the next decade, while a small group works on average 1–2 movies a year. My results suggest that about half of this wide dispersion of career outcomes is generated by differences in the career trajectories of key supervisors who manage workers during their initial year.

My model-based approach is facilitated by focusing on a particular labor market in which most employment is temporary and unions negotiate a high pay standard. These features allow me to abstract away from issues such as job-to-job transitions or wage setting that would need to be modeled in a more general setting.[48] Nevertheless, my findings inform our understanding of careers outside of Hollywood as well. The recent emergence of the "on-demand economy" has renewed interest in the potential benefits and risks of all kinds of alternative employment arrangements (e.g., Hall and Krueger, 2015 and Weil, 2014). Some sectors that are prominent users of these more flexible forms of production, such as construction, rely heavily on networks.

[48]See Abowd et al. (2015) for such a model. They assess bias in estimates of worker and firm effects using linked employer-employee data due to endogenous mobility. In an earlier version of the paper (Abowd and Schmutte, 2012) they propose using co-worker networks as an instrument for the worker-firm assignment.

1.11 Tables

Table 1.1: Characteristics of grips and lighting technicians on California produced major movies

	Grips	Lighting Technicians
New workers	2.98	2.97
Experienced workers	6.26	6.25
No direct or indirect connection	2.50	2.40
Indirect connection	0.84	1.03
Direct connection	2.92	2.82
Worked together once	1.22	1.34
Worked together twice	0.54	0.55
Worked together 3+ times	1.16	0.92
Over all	9.25	9.22
Total number of movies	3432	3432

Notes: This table presents the average number of regular (i.e., non-supervisory) grips and lighting technicians on major California movies by their relationship to their key supervisors. Movies were released between 1988 and 2012. "New workers" are workers who first appear on a major California movie in the same year as the movie's release. "Experienced workers" include all other workers. A worker's network is composed of all keys who have supervised the worker in an earlier year. A worker has a direct connection on a movie if a key on the movie is a member of the worker's network. A worker has an indirect connection if they do not have a direct connection, but there is a key in their network who has worked with a key on the movie. For example, I find that, on average, 1.16 grips on a movie have been supervised at least 3 times by one of the movie's key grips on movies released in prior years.

Table 1.2: Descriptive statistics — California produced major movies, 1980–2012

	Mean	Std. Dev.	p25	p50	p75
Production budget (millions, $2012)	43.56	41.97	15.28	30.67	58.78
Domestic box office (millions, $2012)	49.77	74.45	4.49	24.82	64.44
Number of workers credited on film	300.25	263.87	151	226	348
Number of grips	8.72	7.18	4	7	11
Number of key grips	1.95	1.20	1	2	2
Number of lighting technicians	8.71	7.93	4	6	11
Number of chief lighting technicians	2.24	1.40	1	2	3
Total number of movies	3780				

Notes: This table presents descriptive statistics of the California produced major motion pictures released between 1980 and 2012 that are used in my analysis of Hollywood grips and lighting technicians. See Section 1.3.1 for background on the sources I use to construct this sample. Production budget information is missing for 753 movies. "Number of workers credited on film" counts all non-cast screen credits, excluding acknowledgements.

Table 1.3: Grip and lighting technician job outcomes (year-level)

Year	Regular jobs		Key jobs	
	Mean	1 or more (%)	Mean	1 or more (%)
(Initial year) 0	1.121	100.0	0.000	0.0
1	0.267	19.4	0.011	1.0
2	0.281	20.3	0.017	1.5
3	0.260	17.8	0.022	2.0
4	0.253	17.1	0.025	2.1
5	0.230	16.0	0.028	2.5
6	0.220	14.7	0.033	2.8
7	0.218	14.5	0.033	2.7
8	0.201	13.8	0.034	2.8
9	0.179	11.9	0.037	3.0
10	0.170	11.4	0.038	2.9

Notes: The career sample includes 6826 grip and lighting technician workers. The column labeled "Year" shows the number of years since the worker entered the sample. Year-level job outcomes are the total number of movies on which the worker is credited in their assigned position, based the position I observe them in during their initial year. The column labeled "1 or more" shows the percent of workers who work at least 1 job of the specified type. Key jobs are credits in a key supervisory position (i.e., key grips for grips and chief lighting technicians for lighting technicians). Regular jobs are non-supervisory credits. By construction, all workers in the career sample do not work any key jobs during their initial year. See Section 1.3.2 for background on how I construct the career sample.

Table 1.4: Characteristics of worker-movie dyads

	(%)
No indirect or direct connection	90.873
Indirect connection	7.548
Direct connection	
Worked together once	1.374
Worked together twice	0.124
Worked together 3 times	0.037
Worked together 4 times	0.019
Worked together 5+ times	0.024
Regular job	0.158
Key job	0.019
Number of worker-movie dyads	9814840

Notes: The dyad sample includes all pairwise combinations of the 6826 workers in the career sample with the movies released during years 1–10. For example, if a worker initially takes a job in 1988 as a grip, I include in this sample all the dyads constructed between this grip and the movies released during the years 1989 through 1998. A worker's network is composed of all keys who supervised the worker in an earlier year. A worker has a direct connection on a movie if a key on the movie is a member of the worker's network. A worker has an indirect connection if they do not have a direct connection, but there is a key in their network who has worked with a key on the movie. The row labeled "worked together once" shows the percent of dyads in which the worker has been supervised by a key on the movie only once. I define the variables "worked together twice," "worked together 3 times," etc., similarly.

Table 1.5: Job outcomes of worker-movie dyads

	(%)
No indirect or direct connection	0.076
Indirect connection	0.256
Direct connection	
Worked together once	2.457
Worked together twice	11.612
Worked together 3 times	21.101
Worked together 4 times	23.303
Worked together 5+ times	39.176
Over all	0.158

Notes: This table presents the percent of worker-movie dyads in which a worker is hired in a regular job conditioning on the type of connection on the movie. There are 9814840 worker-movie dyads in the dyad sample. See Table 1.4 for how I construct this sample and definitions of direct and indirect connections.

Table 1.6: Job outcomes of worker-movie dyads, by year and previous jobs

	Connection						
	None	Indirect		Direct			
Previous jobs	(%)	(%)	Δ(*log odds*)	(%)	Δ(*log odds*)	Workers	Dyads
Year 1							
1	0.11	0.42	*1.33*	5.93	*4.04*	6135	820041
2	0.32	0.45	*0.33*	10.24	*3.56*	578	76650
3+	0.43	0.92	*0.77*	11.76	*3.44*	113	15318
All	0.13	0.45	*1.22*	6.98	*4.03*	6826	912009
Year 5							
1	0.03	0.09	*1.12*	0.95	*3.47*	3833	552126
2	0.08	0.31	*1.38*	3.05	*3.69*	1221	175232
3	0.15	0.35	*0.87*	5.20	*3.63*	616	88882
4	0.15	0.28	*0.63*	4.71	*3.51*	421	60377
5+	0.23	0.52	*0.83*	6.91	*3.47*	735	106023
All	0.07	0.32	*1.47*	4.48	*4.16*	6826	982640
Year 10							
1	0.01	0.02	*0.91*	0.14	*3.01*	3299	489681
2	0.03	0.07	*0.94*	0.72	*3.23*	1034	153512
3	0.05	0.13	*0.99*	1.75	*3.64*	613	90546
4	0.07	0.10	*0.38*	2.59	*3.64*	420	62750
5	0.10	0.19	*0.65*	1.92	*3.00*	295	43459
6	0.07	0.30	*1.39*	2.99	*3.73*	271	40410
7	0.07	0.31	*1.52*	2.88	*3.78*	160	23586
8	0.09	0.24	*0.99*	4.42	*3.94*	132	19896
9	0.12	0.19	*0.42*	4.05	*3.52*	118	17593
10+	0.14	0.32	*0.80*	5.42	*3.68*	484	71664
All	0.03	0.17	*1.63*	3.22	*4.60*	6826	1013097
Years 1–10	0.08	0.26	*1.22*	4.43	*4.11*	6826	9814840

Notes: This table presents the percent of worker-movie dyads in which a worker is hired in a regular job conditioning on the type of connection on the movie. There are 9814840 worker-movie dyads in the dyad sample. See Table 1.4 for how I construct this sample and definitions of direct and indirect connections. Statistics are tabulated by the number of previous jobs and the number of years since the worker entered the career sample. I measure previous jobs using the total number of regular jobs worked in past years. The column labeled "None" shows the percent of dyads with no direct or indirect connection that result in a regular job for the worker. The columns labeled "Δ(*log odds*)" show the change in the log odds of a regular job outcome for workers with a connection of that type, relative to having no connection.

Table 1.7: Job outcomes of worker-movie dyads, by year and total connections

	Connection						
	None	Indirect		Direct			
Total connections	(%)	(%)	Δ(*log odds*)	(%)	Δ(*log odds*)	Workers	Dyads
Year 1							
0	0.10	0.65	*1.87*	—	—	3071	408868
1	0.13	0.30	*0.81*	8.82	*4.28*	1757	233472
2	0.17	0.49	*1.07*	6.66	*3.74*	976	130324
3	0.16	0.36	*0.80*	7.06	*3.86*	486	65242
4	0.17	0.36	*0.73*	5.29	*3.47*	293	40378
5+	0.32	0.56	*0.55*	6.48	*3.07*	243	33725
All	0.13	0.45	*1.22*	6.98	*4.03*	6826	912009
Year 5							
0	0.04	0.22	*1.64*	—	—	2548	364394
1	0.05	0.18	*1.18*	2.83	*3.98*	1274	182519
2	0.08	0.30	*1.35*	3.54	*3.84*	834	119741
3	0.09	0.17	*0.60*	4.76	*3.99*	623	89900
4	0.09	0.24	*1.00*	3.76	*3.81*	419	60746
5+	0.16	0.43	*0.98*	4.97	*3.47*	1128	165340
All	0.07	0.32	*1.47*	4.48	*4.16*	6826	982640
Year 10							
0	0.01	0.04	*1.20*	—	—	2694	397467
1	0.02	0.07	*1.20*	1.37	*4.17*	952	142478
2	0.03	0.08	*1.03*	1.30	*3.86*	770	114980
3	0.04	0.16	*1.40*	1.66	*3.77*	521	76959
4	0.05	0.09	*0.62*	1.36	*3.39*	350	51692
5	0.06	0.15	*0.97*	2.47	*3.81*	283	42186
6	0.08	0.15	*0.65*	3.69	*3.87*	185	27538
7	0.08	0.18	*0.86*	2.43	*3.48*	223	33346
8	0.13	0.29	*0.82*	3.04	*3.19*	115	16950
9	0.10	0.14	*0.34*	2.74	*3.34*	134	19918
10+	0.13	0.32	*0.91*	4.58	*3.61*	599	89583
All	0.03	0.17	*1.63*	3.22	*4.60*	6826	1013097
Years 1–10	0.08	0.26	*1.22*	4.43	*4.11*	6826	9814840

Notes: This table presents the percent of worker-movie dyads in which a worker is hired in a regular job conditioning on the type of connection on the movie. There are 9814840 worker-movie dyads in the dyad sample. Total connections are the total number of movies that a worker is directly connected to that year. See Table 1.6 for more information.

Table 1.8: Estimates from random effects models: Hiring model parameters

	(1)		(2)		(3)	
Type 1 Intercept	-7.020	(0.126)	-5.874	(0.127)	-5.483	(0.133)
Type 2 Intercept	—	—	-7.924	(0.163)	-6.363	(0.154)
Type 3 Intercept	—	—	—	—	-8.522	(0.215)
Network effects						
Indirect connection	0.952	(0.031)	0.996	(0.032)	1.002	(0.032)
Direct connection						
Worked together once	3.055	(0.032)	3.091	(0.033)	3.098	(0.034)
Worked together twice	4.547	(0.048)	4.612	(0.051)	4.635	(0.051)
Worked together 3 times	5.231	(0.064)	5.359	(0.069)	5.383	(0.069)
Worked together 4 times	5.391	(0.085)	5.515	(0.092)	5.581	(0.091)
Worked together 5 times	6.213	(0.079)	6.319	(0.083)	6.421	(0.089)
Worker characteristics						
State dependence						
Jobs in year $t-1$	0.238	(0.018)	0.201	(0.023)	0.183	(0.020)
Jobs in year $t-2$	0.150	(0.021)	0.122	(0.021)	0.104	(0.019)
Jobs in year $t-1 \times t-2$	-0.060	(0.011)	-0.043	(0.009)	-0.040	(0.008)
Previous jobs						
2 jobs	0.807	(0.036)	0.250	(0.079)	0.042	(0.080)
3 jobs	1.008	(0.044)	0.031	(0.113)	-0.201	(0.085)
4 jobs	1.061	(0.054)	-0.192	(0.127)	-0.367	(0.085)
5–9 jobs	1.076	(0.063)	-0.381	(0.116)	-0.598	(0.094)
10+ jobs	0.822	(0.096)	-0.726	(0.134)	-1.101	(0.124)
Grip	0.017	(0.023)	-0.038	(0.028)	0.001	(0.032)
Log likelihood	-88817.0		-88061.5		-87903.0	
Number of types	1		2		3	
Number of parameters	83		89		95	
Number of workers	6826		6826		6826	
Number of worker-movie dyads	9812873		9812873		9812873	

Notes: This table presents estimates from the hiring model presented in Section 1.6, in which I treat the unobserved worker effects as randomly drawn from a distribution that depends on the worker's characteristics during their initial year. I estimate the models on the dyad sample described in the notes to Table 1.4. The dependent variable in the hiring model is an indicator that equals 1 if the worker is credited on the movie in a regular job. See the notes to Table 1.4 for the definitions of the connection variables. The variables "jobs in year $t-1$" and "jobs in year $t-2$" are the number of jobs worked in the previous year and two years ago, respectively. The variable "jobs in year $t-1 \times t-2$" is the interaction of these two variables. The variable "2 jobs" is an indicator for whether the worker has worked on exactly 2 jobs in the past. I define the variables "3 jobs," "4 jobs," etc., similarly. "Grip" is an indicator that equals 1 if the worker is a grip and 0 if the worker is a lighting technician. All models also include worker-movie match characteristics, movie characteristics, career-year effects, and year effects. Standard errors clustered at the worker-level are in parentheses.

Table 1.9: Estimates from random effects models, *continued*: Type model parameters

	(1)		(2)		(3)	
Type 1						
2 jobs during initial year	—	—	1.269	(0.130)	1.819	(0.202)
3+ jobs during initial year	—	—	2.112	(0.285)	3.151	(0.427)
Initial indirect connections, years 1–10	—	—	-0.002	(0.001)	-0.002	(0.001)
Initial connections, years 1–10	—	—	0.015	(0.005)	0.011	(0.008)
Intercept	—	—	-0.993	(0.054)	-1.606	(0.150)
Type 2						
2 jobs during initial year	—	—	—	—	1.463	(0.175)
3+ jobs during initial year	—	—	—	—	1.825	(0.484)
Initial indirect connections, years 1–10	—	—	—	—	-0.001	(0.001)
Initial connections, years 1–10	—	—	—	—	0.022	(0.007)
Intercept	—	—	—	—	-1.076	(0.106)
ESTIMATED TYPE PROPORTIONS						
Type 1	1.00		0.31		0.14	
Type 2	—		0.69		0.28	
Type 3	—		—		0.58	

Notes: This table presents estimates from the type model presented in Section 1.6, in which I treat the unobserved worker effects as randomly drawn from a distribution that depends on the worker's characteristics during their initial year. See Table 1.8 for more information. "Initial connections, years 1–10" are the total number of movies during years 1–10 that are managed by a key who supervised the worker during their initial year. "Initial indirect connections, years 1–10" are the total number of movies during years 1–10 that do not have a direct connection but are managed by a key who worked in the past with a key who supervised the worker during their initial year. I estimate the type proportions by first calculating the workers' prior probabilities of each type, as implied by the type model presented in Section 1.6.3. I then take the average of the type probabilities across the sample of 6826 workers. Standard errors clustered at the worker-level are in parentheses.

Table 1.10: Estimates from random effects models, *continued*: more hiring model parameters

	(1)		(2)		(3)	
Observable match characteristics						
Budget-type jobs	0.033	(0.008)	0.027	(0.008)	0.028	(0.008)
Thriller jobs	0.010	(0.009)	0.006	(0.009)	0.002	(0.009)
Action jobs	0.004	(0.010)	0.001	(0.010)	0.002	(0.010)
Adventure jobs	0.078	(0.017)	0.074	(0.016)	0.075	(0.019)
Comedy jobs	0.024	(0.007)	0.018	(0.007)	0.015	(0.007)
Crime jobs	0.019	(0.013)	0.019	(0.013)	0.023	(0.013)
Drama jobs	0.027	(0.007)	0.022	(0.007)	0.017	(0.006)
Sci-fi jobs	0.074	(0.023)	0.075	(0.024)	0.077	(0.023)
Horror jobs	0.129	(0.049)	0.113	(0.050)	0.099	(0.049)
Fantasy jobs	0.079	(0.022)	0.069	(0.022)	0.071	(0.023)
Romance jobs	-0.017	(0.012)	-0.022	(0.012)	-0.022	(0.012)
Mystery jobs	-0.001	(0.024)	-0.007	(0.024)	-0.007	(0.024)
Family jobs	0.047	(0.031)	0.036	(0.031)	0.028	(0.031)
Movie characteristics						
2 keys	0.127	(0.024)	0.130	(0.024)	0.130	(0.024)
3 keys	0.136	(0.028)	0.139	(0.029)	0.142	(0.029)
4 keys	0.285	(0.034)	0.284	(0.034)	0.287	(0.034)
5 keys	0.211	(0.044)	0.216	(0.044)	0.220	(0.044)
6+ keys	0.285	(0.044)	0.293	(0.044)	0.300	(0.045)
Missing budget	-0.010	(0.038)	-0.010	(0.038)	-0.011	(0.038)
\$12m–\$24m	0.100	(0.040)	0.112	(0.040)	0.112	(0.041)
\$24m–\$40m	0.061	(0.042)	0.070	(0.042)	0.069	(0.042)
\$40m–\$70m	0.117	(0.043)	0.125	(0.043)	0.126	(0.043)
≥\$70m	0.338	(0.042)	0.349	(0.042)	0.348	(0.042)
Thriller	-0.051	(0.029)	-0.047	(0.029)	-0.039	(0.029)
Action	0.005	(0.031)	0.007	(0.031)	0.007	(0.031)
Adventure	-0.014	(0.031)	-0.011	(0.031)	-0.017	(0.032)
Comedy	-0.074	(0.028)	-0.063	(0.028)	-0.054	(0.028)
Crime	-0.099	(0.027)	-0.094	(0.027)	-0.099	(0.027)
Drama	-0.114	(0.026)	-0.103	(0.026)	-0.092	(0.025)
Sci-fi	0.003	(0.033)	0.003	(0.033)	0.003	(0.033)
Horror	-0.115	(0.041)	-0.113	(0.042)	-0.108	(0.042)
Fantasy	0.031	(0.032)	0.038	(0.033)	0.037	(0.033)
Romance	-0.047	(0.025)	-0.042	(0.025)	-0.043	(0.025)
Mystery	0.094	(0.032)	0.097	(0.032)	0.098	(0.032)
Family	-0.007	(0.037)	-0.005	(0.037)	0.000	(0.037)

Notes: This table presents additional estimates from the hiring model presented in Section 1.6.2. See Table 1.8 for more information. Match characteristics measure the number of regular jobs worked in the past on movies of the same type. Standard errors clustered at the worker-level are in parentheses.

Table 1.12: Comparison of estimates of network effects under different models

	Logit (1)	Logit (2)	RE Logit (3)	FE Logit (4)
Indirect connection	1.433	0.952	1.002	0.781
	(0.031)	(0.031)	(0.032)	(0.040)
Direct connection				
Worked together once	3.618	3.055	3.098	2.758
	(0.025)	(0.032)	(0.034)	(0.043)
Worked together twice	5.424	4.547	4.635	4.232
	(0.043)	(0.048)	(0.051)	(0.088)
Worked together 3 times	6.196	5.231	5.383	4.964
	(0.059)	(0.064)	(0.069)	(0.159)
Worked together 4 times	6.376	5.391	5.581	5.219
	(0.082)	(0.085)	(0.091)	(0.213)
Worked together 5+ times	7.249	6.213	6.421	6.560
	(0.075)	(0.079)	(0.089)	(0.219)
Number of types	1	1	3	—
Worker and match controls	No	Yes	Yes	Yes
Worker-year effects	No	No	No	Yes
Movie-cohort effects	No	No	No	Yes
Log likelihood	-91673.3	-88817.0	-87903.0	-130213.8
Workers	6826	6826	6826	3445
Movies	3319	3319	3319	2745
Unit of observation	Dyad	Dyad	Dyad	Pair
Observations	9812873	9812873	9812873	426045

Notes: This table presents estimates of the effect of network connections from models that make different assumptions about how unobserved worker and movie effects affect workers' hiring outcomes. The dependent variable in the hiring model is an indicator that equals 1 if the worker is credited on the movie in a regular job. The first column from the left estimates a hiring model using a simple logit model without any additional controls other than for position and career-year effects. The second column corresponds to Model (1) in Table 1.8 and adds to the logit model observable worker and movie characteristics as well as measures of the quality of the worker-movie match and other controls. The third column ("RE Logit") corresponds to Model (3) in Table 1.8 and allows for 3 unobserved worker types. See Section 1.6 for how I specify this model. The model labeled "FE Logit" estimates the network connections allowing for unobserved year-specific worker effects and movie effects. I estimate this model using a fixed effects approach I develop in Section 1.8. See the notes to Table 1.4 for the definitions of the connection variables. In the first three columns from the left the row labeled "Observations" shows the number of dyads included in the model. In the fourth column, "Observations" shows the number of pairs included in the model. Standard errors are in parentheses.

Table 1.11: Estimated type proportions by quality of initial key supervisors

			Characteristics of connected workers			
Quintile	Keys	Key jobs, years 1–10	Regular jobs, years 1–10	Type 1	Type 2	Type 3
1	1402	0.00	2.01	0.13	0.26	0.61
2	618	1.00	2.21	0.13	0.29	0.58
3	1003	2.80	2.46	0.15	0.30	0.54
4	854	6.78	2.79	0.16	0.31	0.53
5	882	14.30	3.24	0.15	0.34	0.51
Over all	4759	4.59	2.28	0.14	0.28	0.58

Notes: This table presents characteristics of workers by the quality of the keys who supervise them during their initial year of work (initial keys). I measure quality by the number of movies managed by the key during the subsequent decade. I describe how I construct this sample in Section 1.7.3. The column labeled "Keys" shows the number of initial keys in each quintile, and the column labeled "Key jobs, years 1–10" shows the average number of movies managed by each initial key during their corresponding cohort's 10 year period. The column labeled "Regular jobs, years 1–10" shows the average number of jobs worked by the workers who are initially supervised by a key in the quintile group. The columns labeled "Type 1," "Type 2," and "Type 3" show the estimated proportion of workers of that type who are initially supervised by a key in the quintile group.

Table 1.13: Effects of counterfactual on initial connections, by initial network size

Initial network size	Workers	Mean	Var	p10	p25	p50	p75	p90
			Observed key assignments					
1	2253	4.6	27.9	0.0	0.0	3.0	7.0	14.0
2	2090	8.9	59.8	0.5	3.0	7.0	13.0	20.0
3	1101	12.6	105.6	1.0	4.0	10.0	19.0	28.0
4	632	16.4	118.6	3.0	9.0	13.0	23.0	31.0
5	272	21.9	154.4	7.0	12.0	22.0	30.0	39.0
6	179	25.0	250.3	5.0	13.0	25.0	32.0	49.0
7	105	25.5	273.7	9.0	12.0	22.0	35.0	48.0
8+	194	29.6	630.0	4.0	5.0	27.5	50.0	58.0
Over all	6826	10.6	130.8	0.0	2.0	7.0	15.0	25.0
			Counterfactual key assignments					
1	2253	4.6	0.3	4.0	4.0	5.0	5.0	5.0
2	2090	9.2	0.8	8.0	9.0	9.0	10.0	10.0
3	1101	13.8	1.2	12.0	13.0	14.0	15.0	15.0
4	632	17.9	2.5	16.0	17.0	18.0	19.0	20.0
5	272	22.6	2.7	20.8	21.4	22.5	24.0	25.0
6	179	27.0	3.2	25.0	25.8	26.9	28.0	29.4
7	105	32.2	5.5	29.3	30.7	32.4	34.2	34.7
8+	194	40.2	56.0	34.2	35.1	38.3	41.7	49.3
Over all	6826	11.5	65.6	4.0	5.0	9.6	14.3	21.6

Notes: This table compares summary statistics of workers' initial connections under the counterfactual distribution of key hiring outcomes ("Counterfactual key assignments") to that under the key hiring outcomes I observe in my original sample ("Observed key assignments"). Initial network size is the total number of keys who supervise a worker during their initial year. See Section 1.9 for details on the counterfactual. Statistics based on counterfactual key assignments are averages over 100 simulations.

Table 1.14: Effects of counterfactual on worker outcomes, by worker type

	Key assignments					
	Observed		Counterfactual		*Relative effect (%)*	
	Mean	Var	Mean	Var	Mean	Var
Type 1						
Total regular jobs, years 1–10	10.1	63.9	7.8	22.9	-23.2	-64.1
Total connections, years 1–10						
Worked together at least once	64.3	2227.1	49.3	971.1	-23.4	-56.4
Worked together at least 3 times	5.6	83.2	2.0	16.4	-65.0	-80.2
Type 2						
Total regular jobs, years 1–10	3.1	6.6	2.7	3.8	-12.7	-41.9
Total connections, years 1–10						
Worked together at least once	32.9	804.6	27.0	363.5	-18.1	-54.8
Worked together at least 3 times	1.4	16.1	0.3	1.9	-76.8	-88.0
Type 3						
Total regular jobs, years 1–10	0.3	0.3	0.3	0.3	0.6	-2.4
Total connections, years 1–10						
Worked together at least once	10.9	132.7	12.0	73.1	10.1	-44.9
Worked together at least 3 times	0.0	0.3	0.0	0.0	-78.0	-90.7
Over all						
Total regular jobs, years 1–10	2.4	22.1	2.0	11.0	-18.1	-50.3
Total connections, years 1–10						
Worked together at least once	24.5	963.7	21.4	449.0	-12.7	-53.4
Worked together at least 3 times	1.2	19.8	0.4	3.3	-69.0	-83.4

Notes: This table compares predicted worker outcomes under the counterfactual distribution of key hiring outcomes ("Counterfactual") to predicted outcomes under the key hiring outcomes I observe in my original sample ("Observed"). Statistics are averages over 100 replications. I predict worker outcomes using estimates from Model (3), presented in Tables 1.8, 1.9, and 1.10. Total connections are the number of movies during the period indicated on which the worker has a connection of the indicated strength. The columns labeled "Relative effect" show the relative difference of the statistic between the the counterfactual and original distributions of key hiring outcomes. For example, in the top row labeled "Total regular jobs, years 1–10" I find the relative effect of the counterfactual on the variance is $\frac{22.9-63.9}{63.9} \times 100 = -64.1$. See the notes to Table 1.4 for the definitions of the connection variables. See Section 1.9 for details on the counterfactual.

1.12 Figures

Figure 1.1: Organization chart for a typical movie, grip and lighting technician crews

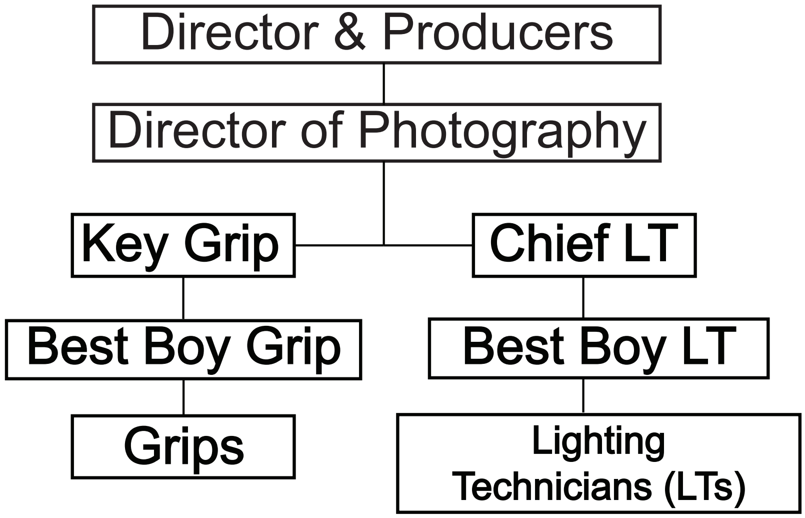

Notes See Section 1.2 for background on the grip and lighting technician crews.

Figure 1.2: Counts of major motion pictures, by year

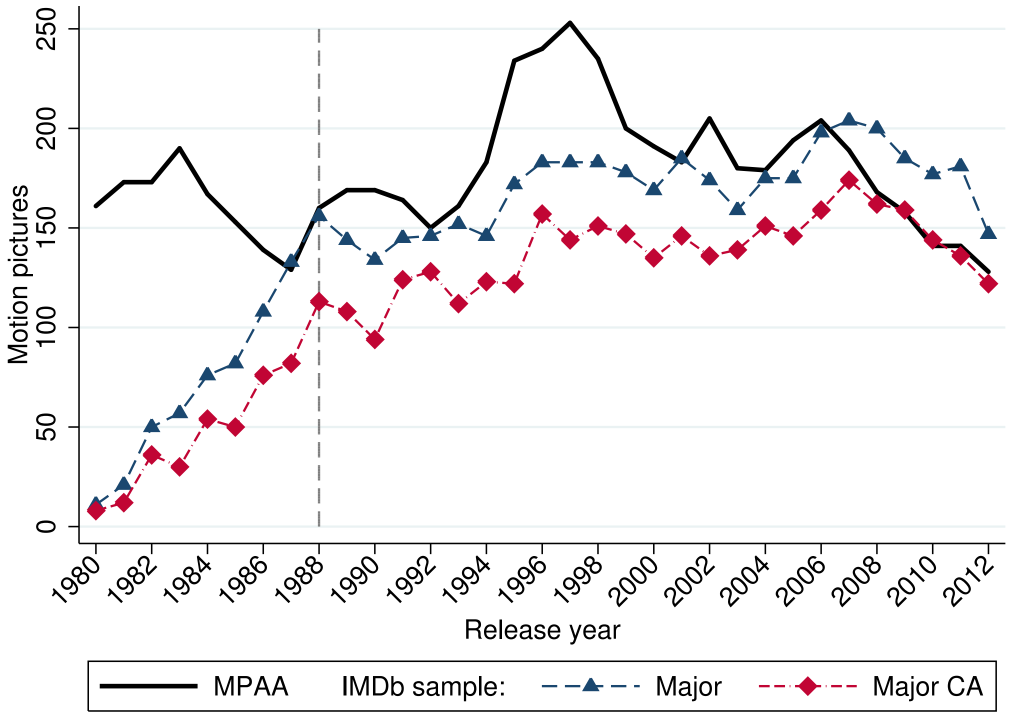

Notes: This figure shows the total number of movies released each year. The line labeled "MPAA" shows counts according to the Motion Picture Association of America, which counts the total number of movies released by the major studios and their subsidiaries. Counts in 2005 and later are from the association's Theatrical Market Statistics report for 2014 (MPAA, 2015). Earlier values are from Vogel (2014), Table 3.4. The remaining two lines show counts of movies in my screen credit data downloaded from the Internet Movie Database. "Major" includes all movies produced in the United States. "Major CA" includes only movies produced in California. See Section 1.3.1 for more on how I construct these samples. The grips and lighting technicians in the career sample first appear between 1988 and 2002.

Figure 1.3: Histogram of total jobs worked, years 1–10

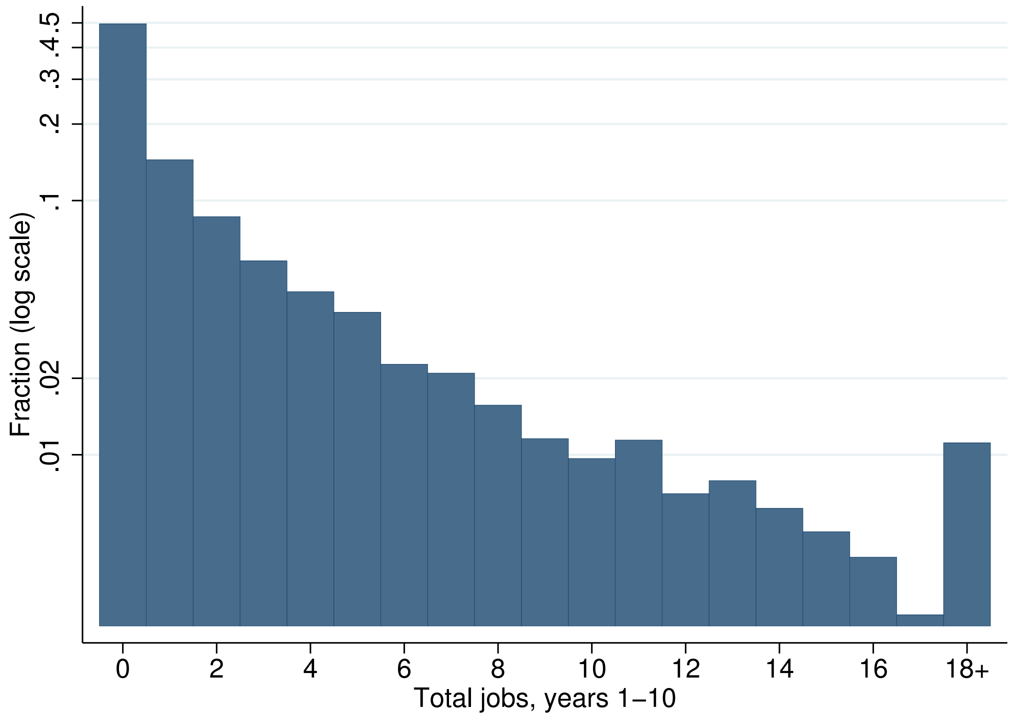

Notes: This figure is a histogram of the total number of regular (i.e., non-supervisory) jobs worked among the grips and lighting technicians in the career sample during years 1–10. The point on the horizontal axis labeled "18+" includes workers who work on 18 or more movies during this period. I plot fractions on a log scale. See Section 1.3.2 for construction of the career sample.

Figure 1.4: Histogram of network size, year 1

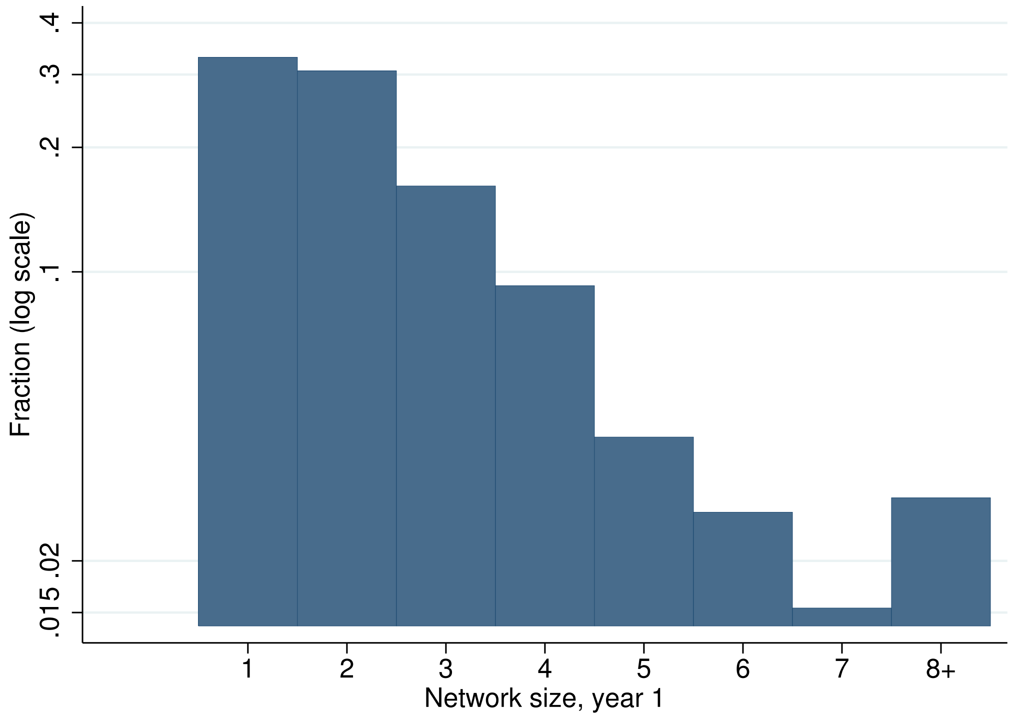

Notes: In this figure, a worker's network is composed of all keys who supervise the worker on a movie by the beginning of year 1 in the career sample. I plot fractions on a log scale. See Section 1.3.2 for construction of the career sample.

Figure 1.5: Histogram of network size, year 10

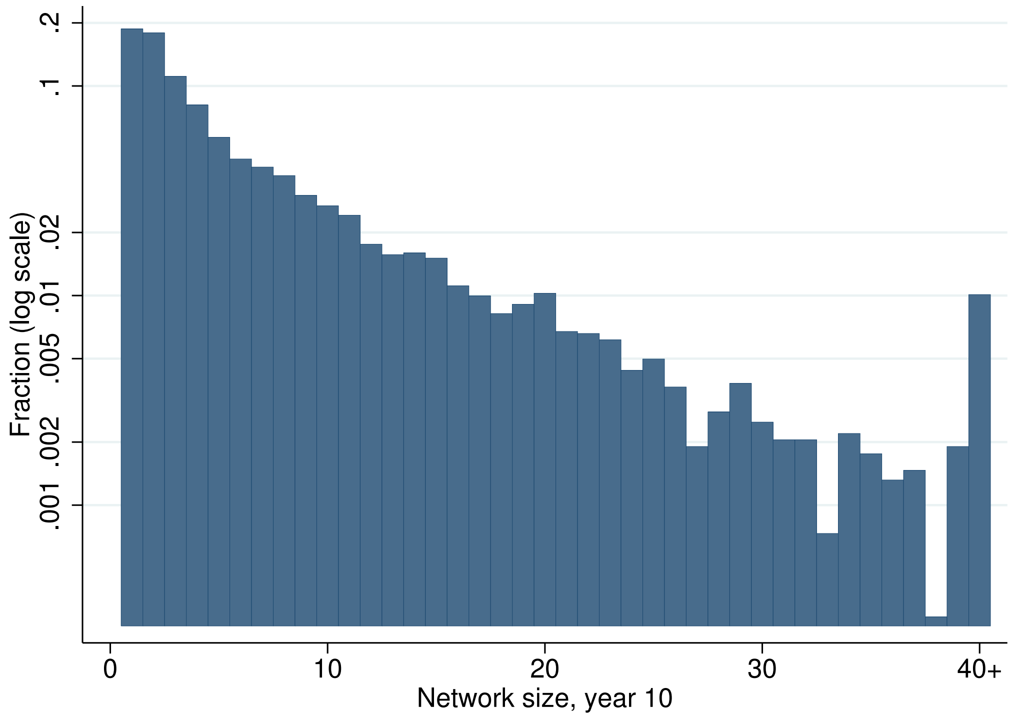

Notes: In this figure, a worker's network is composed of all keys who supervise the worker on a movie by the beginning of year 10 in the career sample. I plot fractions on a log scale. See Section 1.3.2 for construction of the career sample.

Figure 1.6: Career outcomes by initial network quality

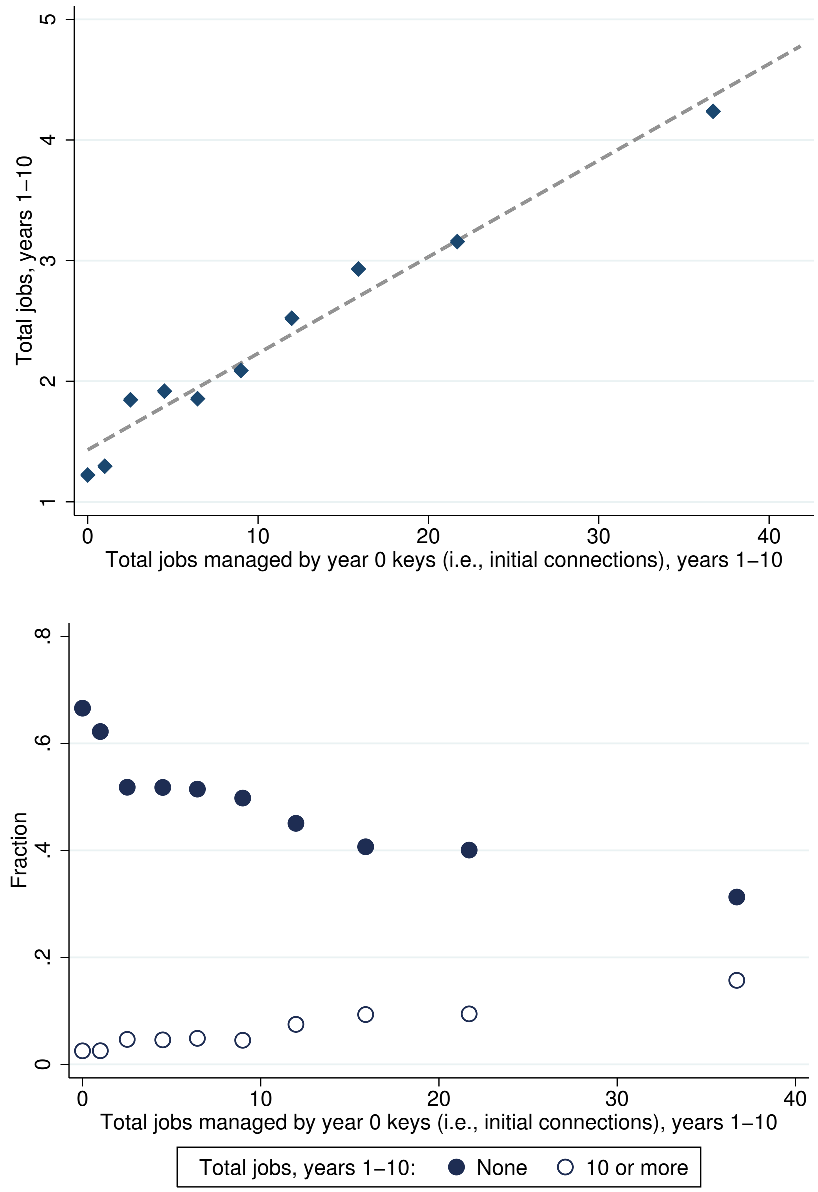

Notes: These figures present the relationship between the total number of regular jobs worked over years 1–10 and initial connections—the number of movies during years 1–10 that are managed by a key who supervised the worker during the worker's initial year (i.e., year 0). Initial connections represent the number of jobs that a worker's initial supervisors can pass on to them. I bin initial connections into 10 deciles. In the top panel, each point is the average number of jobs worked among the workers in that bin. The dashed line plots a least squares fit of total jobs on initial connections. In the bottom panel, the points labeled "None" plot the fraction of workers in the bin who work 0 jobs during the decade. The points labeled "10 or more" plot the fraction of workers who work 10 or more jobs. See Section 1.3.2 for background on how I construct the career sample.

Figure 1.7: A direct connection

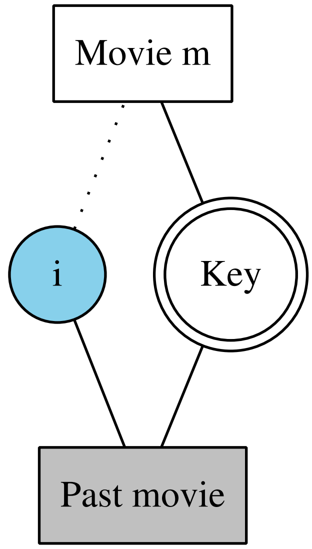

Notes: A worker i has a direct connection on a movie m if there is a key on movie m who has supervised worker i on a movie released in a prior year.

Figure 1.8: An indirect connection

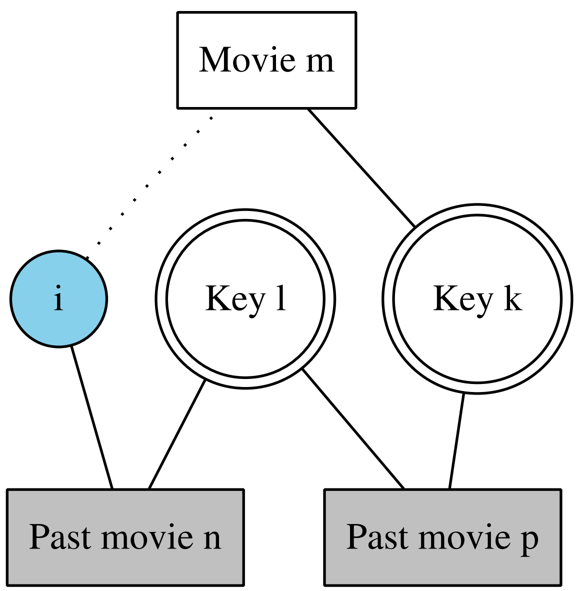

Notes: A worker i has an indirect connection on a movie m if they do not have a direct connection, but there is a key l who has (1) supervised worker i on a movie n released in a prior year, and (2) worked with a key k on movie m on another movie p released in a prior year.

Figure 1.9: Histogram of total jobs worked: real vs. predicted

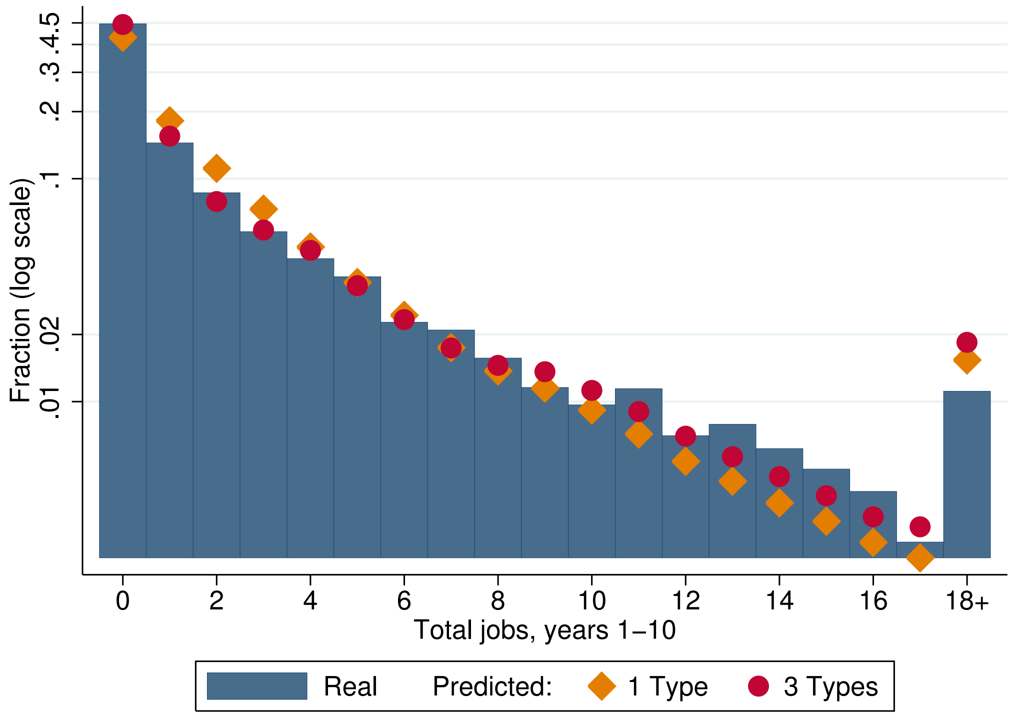

Notes: This figure is a histogram of the total number of regular jobs worked during years 1–10. The point on the horizontal axis labeled "18+" includes workers who work on 18 or more movies during this period. The bars labeled "real" show the empirical distribution. The points labeled "1 Type" and "3 Types" plot predicted fractions from the estimated logit models presented in Tables 1.8, 1.9, and 1.10. Predicted fractions are averages over 100 simulations. See Section 1.7.2 for background on how I simulate the models.

Figure 1.10: Initial connections and job outcomes: real vs. predicted

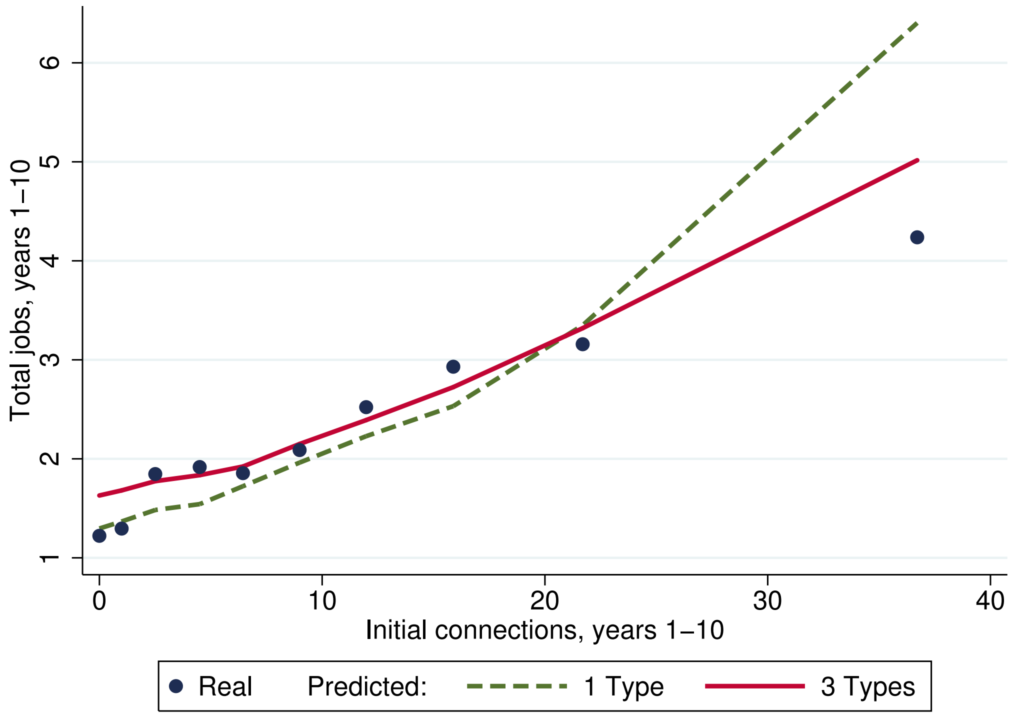

Notes: This figure presents the relationship between the total number of regular jobs worked during years 1–10 and initial connections—the number of movies during years 1–10 that are managed by a key who supervised the worker during the worker's initial year (i.e., year 0). Initial connections represent the number of jobs that a worker's initial supervisors can pass on to them. I bin initial connections into 10 deciles. Each point is the average number of jobs worked over 10 years among the workers in that bin. The lines labeled "1 Type" and "3 Types" connect the predicted average number of jobs in that bin from the estimated logit models presented in Tables 1.8, 1.9, and 1.10. Predictions are averages over 100 simulations. See Section 1.7.2 for background on how I simulate the models.

Figure 1.11: A pair of worker-movie dyads

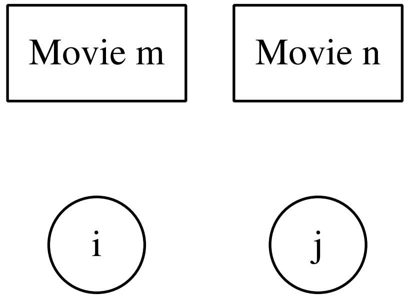

Notes: A pair contains exactly two workers, i and j, and two movies, m and n.

Figure 1.12: Pairs of worker-movie dyads for which Y_{ijmn} is non-zero

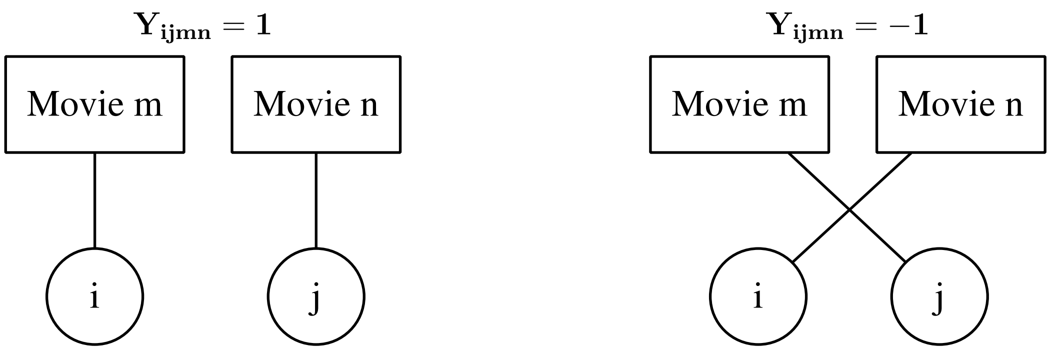

Notes: See Section 1.8 for definition of Y_{ijmn}.

Figure 1.13: Effects of counterfactual on total jobs managed by initial keys, years 1–10

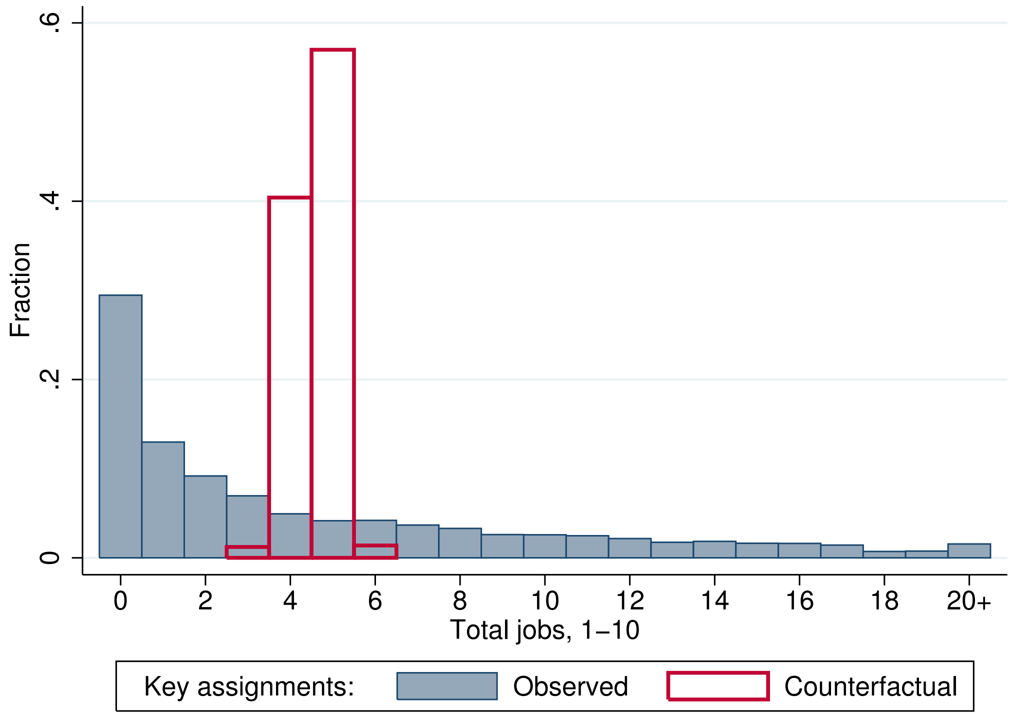

Notes: This figure compares the number of jobs worked by initial key supervisors in my original sample ("Observed") to those under my counterfactual ("Counterfactual"). I describe how I construct the sample of initial keys in Section 1.9.There are 4759 initial keys in this sample. I create this figure by counting the number of movies managed by each initial key in their corresponding cohort's decade in my original sample and under my counterfactual.

Figure 1.14: Effects of counterfactual on distribution of total jobs worked, years 1–10
All workers

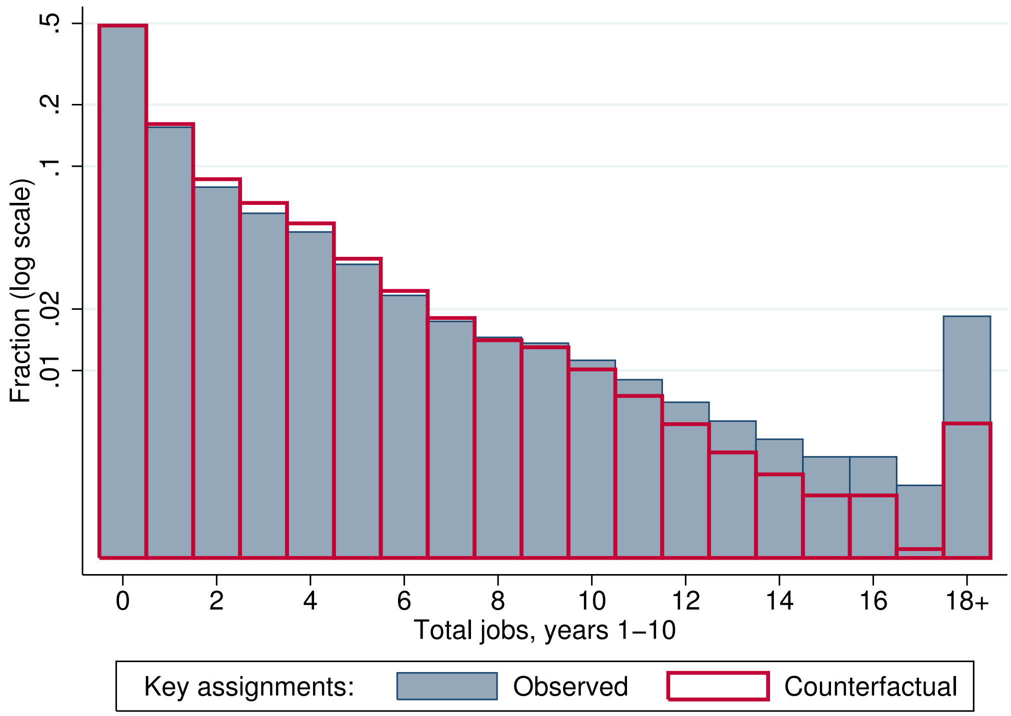

Notes: This figure compares predicted worker outcomes under the counterfactual distribution of key hiring outcomes ("Counterfactual") to predicted outcomes under the key hiring outcomes I observe in my original sample ("Observed"). Each bar is an average over 100 replications. I predict worker outcomes using estimates from Model (3), presented in Tables 1.8, 1.9, and 1.10. See Section 1.9 for details on the counterfactual. I plot fractions on a log scale.

Figure 1.15: Effects of counterfactual on distribution of total jobs worked, years 1–10
Type 1 workers

<u>Notes</u>: This figure compares predicted Type 1 worker outcomes under the counterfactual distribution of key hiring outcomes ("Counterfactual") to predicted outcomes under the key hiring outcomes I observe in my original sample ("Observed"). Each bar is an average over 100 replications. I predict worker outcomes using estimates from Model (3), presented in Tables 1.8, 1.9, and 1.10. I draw a worker's type in each replication from the prior distribution of the worker effects, using the observed hiring outcomes during the worker's initial year in my original sample. See Section 1.9 for details on the counterfactual. I plot fractions on a log scale.

Figure 1.16: Effects of counterfactual on distribution of total jobs worked, years 1–10
Type 2 workers

Notes: This figure compares predicted Type 2 worker outcomes under the counterfactual distribution of key hiring outcomes ("Counterfactual") to predicted outcomes under the key hiring outcomes I observe in my original sample ("Observed"). Each bar is an average over 100 replications. I predict worker outcomes using estimates from Model (3), presented in Tables 1.8, 1.9, and 1.10. I draw a worker's type in each replication from the prior distribution of the worker effects, using the observed hiring outcomes during the worker's initial year in my original sample. See Section 1.9 for details on the counterfactual. I plot fractions on a log scale.

Figure 1.17: Effects of counterfactual on distribution of total jobs worked, years 1–10
Type 3 workers

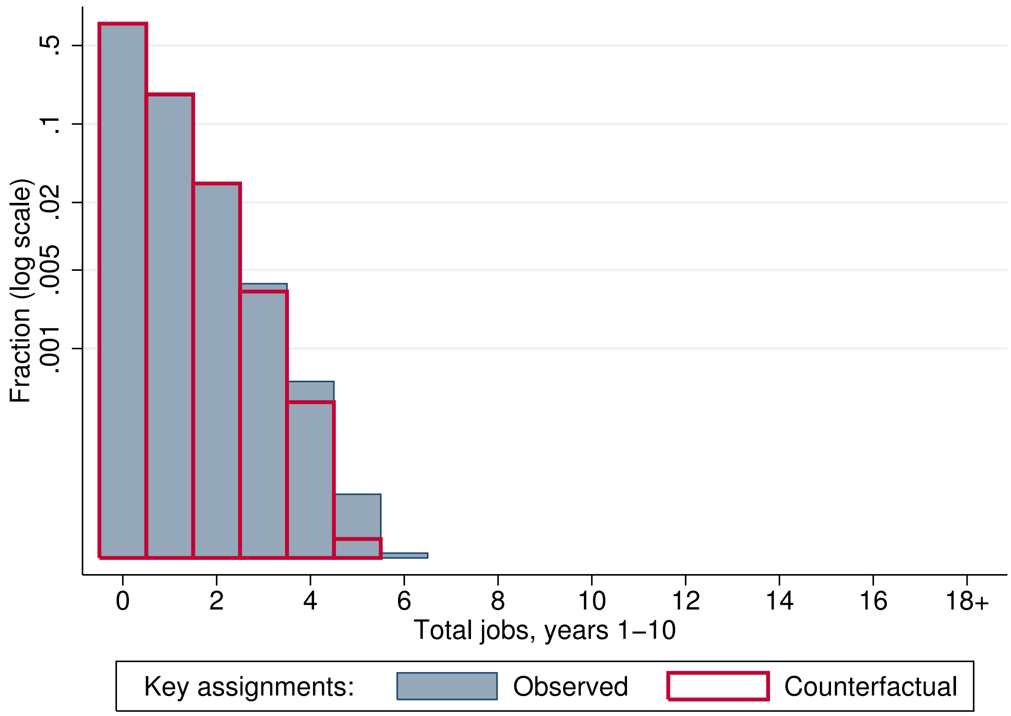

Notes: This figure compares predicted Type 3 worker outcomes under the counterfactual distribution of key hiring outcomes ("Counterfactual") to predicted outcomes under the key hiring outcomes I observe in my original sample ("Observed"). Each bar is an average over 100 replications. I predict worker outcomes using estimates from Model (3), presented in Tables 1.8, 1.9, and 1.10. I draw a worker's type in each replication from the prior distribution of the worker effects, using the observed hiring outcomes during the worker's initial year in my original sample. See Section 1.9 for details on the counterfactual. I plot fractions on a log scale.

Chapter 2

Asymptotic variance of the two-way fixed effects estimator

2.1 Introduction

In Chapter 1, I introduced a two-way fixed effects estimator for the effect of workers' network connections on their employment outcomes in a model that allows for unobserved heterogeneity across workers and movies. The estimator is based on subgraphs of worker-movie dyads that I call pairs. Each pair contains exactly two workers and two movies. I showed that when conditioned properly, the marginal likelihood of a pair's (joint) outcome does not depend on the unobserved worker and movie effects. This result motivated an estimator in which the criterion sums over all pairs of workers and movies available in the sample.

In this chapter, I study the asymptotic properties of this estimator. Inference is non-standard, because pairs within a sample are only independent when they do not share any workers or movies in common. The underlying criterion is a two-sample U-process. Although the properties U-process minimizers have been studied (e.g., Honoré and Powell, 1994), prior results do not apply.

The estimator's first order condition is asymptotically equivalent to a degenerate U-statistic. I adapt a proof presented by Graham (2015a) and show that the U-statistic is asymptotically equivalent to a certain projection which involves summation over all the worker-movie dyads in the sample. I use this result to derive a consistent estimator of the variance of the two-way fixed effects estimator.

All notation is as defined in Chapter 1, unless noted otherwise. A zero subscript on a parameter denotes its population value. Let $\mathbf{A}_{ijc}^{t-1}$ denote worker i's and j's stacked hiring outcomes on movies produced between years c and $t-1$. To ease the exposition, assume there are exactly N workers in each cohort $c = 1, ..., J$. Furthermore, assume there are exactly M movies released each year $t = 1, ..., T$.

Recall that given a pair of two workers in cohort c, i and j, and two movies produced in year t, m and n, we have:

$$P(Y_{ijmn}=1|\mathbf{A}_{ijc}^{t-1},\eta_{it},\eta_{jt},\delta_m,\delta_n,\mathbf{D}_0^t,Y_{ijmn}\in\{-1,1\})=\frac{\exp(\tilde{C}'_{ijmn}\theta_0)}{1+\exp(\tilde{C}'_{ijmn}\theta_0)} \tag{2.1}$$
$$\tag{2.2}$$

where $\tilde{C}_{ijmn}\equiv(C_{im}-C_{in})-(C_{jm}-C_{jn})$. In Section 1.8 of Chapter 1, I proposed an estimator for θ_0:

$$\hat{\theta}=\text{argmax}\frac{1}{JT}\binom{N}{2}^{-1}\binom{M}{2}^{-1}\sum_{c=1}^{J}\sum_{(i,j)\in\mathbb{C}_{2,\mathbb{N}(c)}}\sum_{t=c+1}^{T}\sum_{(m,n)\in\mathbb{C}_{2,\mathbb{M}(t)}}\log l_{ijmn}^{cond}(\theta) \tag{2.3}$$

where $\log l_{ijmn}^{cond}(\theta)$ denotes each pair's contribution to the criterion above evaluated at a potential value of θ, $\log l_{ijmn}^{cond}(\theta)\equiv|Y_{ijmn}|\log\left[\frac{\exp\left(\mathbb{I}(Y_{ijmn}=1)\left(\tilde{C}'_{ijmn}\theta\right)\right)}{1+\exp\left(\tilde{C}'_{ijmn}\theta\right)}\right]$.

Observe that the first order condition of the maximization problem in Equation (2.3) is:

$$\frac{1}{JT}\binom{N}{2}^{-1}\binom{M}{2}^{-1}\sum_{c=1}^{J}\sum_{(i,j)\in\mathbb{C}_{2,\mathbb{N}(c)}}\sum_{t=c+1}^{T}\sum_{(m,n)\in\mathbb{C}_{2,\mathbb{M}(t)}}\frac{\partial\log l_{ijmn}^{cond}(\hat{\theta})}{\partial\theta}\equiv 0 \tag{2.4}$$

where $\frac{\partial\log l_{ijmn}^{cond}(\theta)}{\partial\theta}$ is the score for a given pair (i,j,m,n) evaluated at θ. Let $h_{ijmn}\equiv\frac{\partial\log l_{ijmn}^{cond}(\theta_0)}{\partial\theta}$ denote the pair's score evaluated at θ_0.

To begin, I take a mean value expansion of a score evaluated at $\hat{\theta}$ around its population value, θ_0. After some manipulation, we have:

$$\sqrt{NM}\left(\hat{\theta}-\theta_0\right)=\hat{\Gamma}\left(\bar{\theta}\right)^{-1}\frac{1}{JT}\binom{N}{2}^{-1}\binom{M}{2}^{-1}\sum_{c=1}^{J}\sum_{(i,j)\in\mathbb{C}_{2,\mathbb{N}(c)}}\sum_{t=c+1}^{T}\sum_{(m,n)\in\mathbb{C}_{2,\mathbb{M}(t)}}h_{ijmn} \tag{2.5}$$

where $\bar{\theta}$ is between $\hat{\theta}$ and θ_0, and $\hat{\Gamma}(\theta)$ is a sample average over the pairs' Hessians:

$$\hat{\Gamma}(\theta)\equiv\frac{1}{JT}\binom{N}{2}^{-1}\binom{M}{2}^{-1}\sum_{c=1}^{J}\sum_{(i,j)\in\mathbb{C}_{2,\mathbb{N}(c)}}\sum_{t=c+1}^{T}\sum_{(m,n)\in\mathbb{C}_{2,\mathbb{M}(t)}}\frac{\partial^2\log l_{ijmn}^{cond}(\theta)}{\partial\theta\partial\theta'} \tag{2.6}$$

Assuming the usual conditions, $\hat{\Gamma}\left(\hat{\theta}\right)$ is a consistent estimator of the expectation of the population Hessian, $E\left[\frac{\partial^2\log l_{ijmn}^{cond}(\theta_0)}{\partial\theta\partial\theta'}\right]$. Equation (2.5) shows that the asymptotic properties of $\hat{\theta}$ will depend on the asymptotic distribution of the sample average over all the pairs' scores. Nevertheless, the problem is non-standard, because this average is a fourth order two-sample U-statistic.

Let U_{NM} denote the U-statistic of interest: the sample average of the pairs' scores evaluated at θ_0:

$$U_{NM} \equiv \frac{1}{JT}\binom{N}{2}^{-1}\binom{M}{2}^{-1}\sum_{c=1}^{J}\sum_{(i,j)\in\mathbb{C}_{2,\mathbb{N}(c)}}\sum_{t=c+1}^{T}\sum_{(m,n)\in\mathbb{C}_{2,\mathbb{M}(t)}} h_{ijmn} \tag{2.7}$$

Then, the variance of U_{NM} is

$$\begin{aligned}\operatorname{var}(U_{NM}) = \left(\frac{1}{JT}\right)^2\binom{N}{2}^{-2}\binom{M}{2}^{-2}\sum_{c=1}^{J}\sum_{(i,j)\in\mathbb{C}_{2,\mathbb{N}(c)}}\sum_{t=c+1}^{T}\times \\ \sum_{(m,n)\in\mathbb{C}_{2,\mathbb{M}(t)}}\sum_{d=1}^{J}\sum_{(k,l)\in\mathbb{C}_{2,\mathbb{N}(d)}}\sum_{s=d+1}^{T}\sum_{(p,q)\in\mathbb{C}_{2,\mathbb{M}(s)}}\operatorname{cov}(h_{ijmn}, h_{klpq})\end{aligned} \tag{2.8}$$

I will show that U_{NM} is asymptotically equivalent to a projection which involves summation over all the worker-movie dyads in the sample. I then use this result to construct a consistent estimator of the variance of the two-way fixed effects estimator.

The rest of the chapter proceeds as follows. In Section 2.2, I examine the properties of the covariance terms in Equation (2.8), $\operatorname{cov}(h_{ijmn}, h_{klpq})$. In Section 2.3, I use these properties to find the asymptotic variance of $\sqrt{NM}U_{NM}$. In Section 2.4, I introduce a projection of the U-statistic. In Section 2.5, I show asymptotic equivalency of $\sqrt{NM}U_{NM}$ and this projection. In Section 2.6, I derive a consistent estimator of the variance of $\hat{\theta}$.

2.2 Properties of cov (h_{ijmn}, h_{klpq})

The covariance of two scores will depend on the number of workers and movies they have in common. For example, when two scores share no workers or movies in common, their covariance is zero, because workers and movies are assumed to be independently sampled. In fact, the covariance of two scores is non-zero only if they share *both* a worker *and* a movie in common. In other words, the covariance is non-zero only if the two pairs share at least one worker-movie dyad. This property is important, as I will show that it implies that the leading term of the variance of U_{NM} is inversely proportional to the number of worker-movie dyads in the sample.

I show this property in three steps. The argument requires some nuance, because of the longitudinal structure of my sample of pairs. First, I show that the covariance of two scores that belong to the same year is zero whenever they do not share any worker-movie dyads in common. Second, I show that this result implies that the expectation of a pair's score conditional on only one of its worker's attributes is zero. Third, I show that the covariance of two scores that belong to different years is zero whenever they do not share any worker-movie dyads in common. To reduce notational burden, I use bold symbols to denote stacked matrices of multiple workers' or movies' attributes.

Observe that each score, h_{ijmn}, is implicitly a function of the year-specific attributes of two workers, i and j, and the attributes of two movies, m and n. Let $W_{it} = \left(A_{it}, \mathbf{A}_{i0}^{t-1}, \eta_{it}\right)$

denote worker i's attributes that determine h_{ijmn}: their current and past hiring outcomes and unobserved year-specific effect. Let $Z_m = (D_m, \delta_m)$ denote movie m's attributes: its choice of keys and unobserved movie effect. Then, $h_{ijmn} = h\left(W_{it}, W_{jt}, Z_m, Z_n; \mathbf{D}_0^{t-1}\right)$. Moreover, it can be shown that this function $h(\cdot)$ is symmetric regarding its worker and movie components. This symmetry derives from the fact that $h(\cdot)$ is the partial derivative of the log of the marginal likelihood of a pair, $\log l_{ijmn}$, and the value of $\log l_{ijmn}$ does not depend on the order in which the workers or movies are specified.

Step 1: Covariance of pairs that belong to the same year

First, consider two pairs, (i, j, m, n) and (i, k, p, q), that share only one worker in common, i, and all movies are produced in the same year, t. Then:

$$
\begin{aligned}
&\mathrm{cov}(h_{ijmn}, h_{ikpq}) = \\
&=E\left[h_{ijmn}h_{ikpq}\right] \\
&=E\left(E\left[h_{ijmn}h_{ikpq}|\mathbf{A}_{ijkc}^{t-1}, \boldsymbol{\eta_{ijkt}}, \boldsymbol{\delta_{mnpq}}, \mathbf{D}_0^t, Y_{ijmn} \in \{-1,1\}, Y_{ikpq} \in \{-1,1\}\right]\right) \\
&=E\Bigg(E\left[h_{ijmn}|\mathbf{A}_{ijc}^{t-1}, \boldsymbol{\eta}_{ijt}, \boldsymbol{\delta}_{mn}, \mathbf{D}_0^t, Y_{ijmn} \in \{-1,1\}\right] \times \\
&\quad E\left[h_{ikpq}|\mathbf{A}_{ikc}^{t-1}, \boldsymbol{\eta}_{ikt}, \boldsymbol{\delta}_{pq}, \mathbf{D}_0^t, Y_{ikpq} \in \{-1,1\}\right]\Bigg) \\
&=0
\end{aligned}
$$

The first line follows, because the expectation of each score is zero evaluated at the true value of θ by the population first order condition: $E\left[h_{ijmn}\right] = 0$. The second line follows by the law of iterated expectations. The third line follows, because worker-movie dyads' outcomes are conditionally independent. The fourth line follows, because:

$$
\begin{aligned}
&E\left[h_{ijmn}|\mathbf{A}_{ijc}^{t-1}, \boldsymbol{\eta}_{ijt}, \boldsymbol{\delta}_{mn}, \mathbf{D}_0^t, Y_{ijmn} \in \{-1,1\}\right] \\
&=E\left[|Y_{ijmn}|\left(\mathbb{I}\left(Y_{ijmn}=1\right) - \frac{\exp\left(\tilde{C}_{ijmn}'\theta_0\right)}{1+\exp\left(\tilde{C}_{ijmn}'\theta_0\right)}\right) \times \right. \\
&\quad \left. \tilde{C}_{ijmn}|\mathbf{A}_{ijc}^{t-1}, \boldsymbol{\eta}_{ijt}, \boldsymbol{\delta}_{mn}, \mathbf{D}_0^t, Y_{ijmn} \in \{-1,1\}\right] \\
&=\left[E\left[\mathbb{I}\left(Y_{ijmn}=1\right)|\mathbf{A}_{ijc}^{t-1}, \boldsymbol{\eta}_{ijt}, \boldsymbol{\delta}_{mn}, \mathbf{D}_0^t, Y_{ijmn} \in \{-1,1\}\right] - \frac{\exp\left(\tilde{C}_{ijmn}'\theta_0\right)}{1+\exp\left(\tilde{C}_{ijmn}'\theta_0\right)}\right]\tilde{C}_{ijmn} \\
&=0
\end{aligned}
$$

Analogous arguments can be used to show that the covariance of any two pairs in which all movies are produced in the same year have zero covariance if they do not share any worker-movie dyads in common.

Step 2: Expectation of the score conditional on one worker's attributes

Again, consider two pairs, (i,j,m,n) and (i,k,p,q), that share only one worker in common, i, and all movies are produced in the same year, t. Since $\text{cov}(h_{ijmn}, h_{ikpq}) = 0$, it then follows that the variance of h_{ijmn} conditional on only one of its worker's attributes is also zero:

$$\begin{aligned}0 = \text{cov}(h_{ijmn}, h_{ikpq}) =& E\left[h_{ijmn}h_{ikpq}\right] \\ =& E[E\left[h_{ijmn}h_{ikpq}|W_{it}=w\right]] \\ =& E\left[E\left[h_{ijmn}|W_{it}=w\right]E\left[h_{ikpq}|W_{it}=w\right]\right] \\ =& E\left[E\left[h_{ijmn}|W_{it}=w\right]^2\right] \\ =& \text{var}(E\left[h_{ijmn}|W_{it}=w\right])\end{aligned}$$

If $\text{var}(E\left[h_{ijmn}|W_{it}=w\right]) = 0$, then

$$\begin{aligned}E\left[h_{ijmn}|W_{it}=w\right] =& E\left[E\left[h_{ijmn}|W_{it}=w\right]\right] \\ =& E\left[h_{ijmn}\right] \\ =& 0\end{aligned}$$

Step 3: Covariance of pairs that belong to different years

Consider two pairs that share only one worker in common and the movies are drawn from different years: $(m,n) \in \mathbb{C}_{2,\mathbb{M}(t)}$ and $(p,q) \in \mathbb{C}_{2,\mathbb{M}(s)}$, where $s \neq t$. Recall that W_{it} is a vector of worker i's current and past hiring outcomes and unobserved year-specific effect: $W_{it} = \left(A_{it}, \mathbf{A}_{i0}^{t-1}, \eta_{it}\right)$.

Without loss of generality, suppose $s < t$. Then, conditional on $W_{it} = w_t$, $W_{is} = w_s$ does not affect the distribution of h_{ijmn}, because (1) all prior hiring outcomes are already included in $W_{it} = w_{it}$, and (2) I have assumed that worker ability in previous years does not affect the probability of i's hiring outcome in year t conditional on η_{it}. Then

$$\begin{aligned}\text{cov}(h_{ijmn}, h_{ikpq}) =& E\left[h_{ijmn}h_{ikpq}\right] \\ =& E[E\left[h_{ijmn}h_{ikpq}|W_{it}=w_t, W_{is}=w_s\right]] \\ =& E\left[E\left[h_{ijmn}|W_{it}=w_t, W_{is}=w_s\right]E\left[h_{ikpq}|W_{it}=w_t, W_{is}=w_s\right]\right] \\ =& E\left[0 \times E\left[h_{ikpq}|W_{it}=w_t, W_{is}=w_s\right]\right] \\ =& 0\end{aligned}$$

where the fourth line follows, because $E\left[h_{ijmn}|W_{it}=w_t, W_{is}=w_s\right] = E\left[h_{ijmn}|W_{it}=w_t\right]$, and I showed in Step 2 that $E\left[h_{ijmn}|W_{it}=w_t\right] = 0$.

Analogous arguments can be used to show that pairs that share both workers in common also have zero covariance when the movies are produced in different years.

2.3 Asymptotic variance of $\sqrt{NM}U_{NM}$

Equation (2.8) expresses the variance of U_{NM} as the weighted average of covariance terms. In the previous section, I showed that these covariance terms are zero whenever the pairs do not share any worker-movie dyads in common. To find the asymptotic variance of $\sqrt{NM}U_{NM}$ I will first show that the form of the covariance of two pairs depends only on the number of workers and movies that they share in common. To start, I introduce notation for the expectation of a pair's score conditional on a particular number of its worker and movie attributes:

$$\begin{aligned}
\bar{h}_{11}(w,z) \equiv& E\left[h_{ijmn}(W_{it},W_{jt},Z_m,Z_n)\,|W_{it}=w,Z_m=z\right]\\
\bar{h}_{21}(w_1,w_2,z) \equiv& E\left[h_{ijmn}(W_{it},W_{jt},Z_m,Z_n)\,|W_{it}=w_1,W_{jt}=w_2,Z_m=z\right]\\
\bar{h}_{12}(w,z_1,z_2) \equiv& E\left[h_{ijmn}(W_{it},W_{jt},Z_m,Z_n)\,|W_{it}=w,Z_m=z_1,Z_n=z_2\right]\\
\bar{h}_{22}(w_1,w_2,z_1,z_2) \equiv& E\left[h_{ijmn}(W_{it},W_{jt},Z_m,Z_n)\,|W_{it}=w_1,W_{jt}=w_2,Z_m=z_1,Z_n=z_2\right]
\end{aligned}$$

Let $\Omega_{wz} \equiv \text{var}\left(\bar{h}_{wz}(\cdot)\right) = E\left[\bar{h}_{wz}(\cdot)^2\right]$ denote the variance of the respective conditional expectation. For example, $\Omega_{12} = \text{var}\left(\bar{h}_{12}(W,Z_1,Z_2)\right)$.

An example suffices to show that for any two pairs with non-zero covariance, I can express their covariance as one of these four variances: $\Omega_{11},\Omega_{12},\Omega_{21}$, or Ω_{22}. Consider two pairs that share one worker and one movie in common: h_{ijmn} and h_{ikmp}. Then:

$$\begin{aligned}
\text{cov}(h_{ijmn},h_{ikmp}) =& E\left[h_{ijmn}h_{ikmp}\right]\\
=& E\left[E\left[h_{ijmn}h_{ikmp}|W_{it}=w,Z_m=z\right]\right]\\
=& E\left[E\left[h_{ijmn}|W_{it}=w,Z_m=z\right]E\left[h_{ikmp}|W_{it}=w,Z_m=z\right]\right]\\
=& E\left[\bar{h}_{11}(W,Z)^2\right]\\
=& \Omega_{11}
\end{aligned}$$

where the third line follows, because once I condition on the worker and movie attributes in common, the remaining terms are independent by assumption.

Therefore, the expression for $\text{var}(U_{NM})$ in Equation (2.8) reduces to a weighted sum of the four variances above. Table 2.1 reports the count of each variance term in the pair sample. Plugging in these counts into Equation (2.8) and simplifying, I find that:

$$\begin{aligned}
\text{var}(U_{NM}) =& \frac{16}{JT}\frac{N-2}{N(N-1)}\frac{M-2}{M(M-1)}\Omega_{11}+\\
& \frac{8}{JT}\frac{N-2}{N(N-1)}\frac{1}{M(M-1)}\Omega_{12}+\frac{8}{JT}\frac{1}{N(N-1)}\frac{M-2}{M(M-1)}\Omega_{21}+\\
& \frac{4}{JT}\frac{1}{N(N-1)}\frac{1}{M(M-1)}\Omega_{22}
\end{aligned} \tag{2.9}$$

Therefore, $\text{var}\left(\sqrt{NM}U_{NM}\right) = \frac{16}{JT}\Omega_{11}$ as $N\to\infty, M\to\infty$.[1]

[1] In my sample, the number of workers per cohort and movies per year are similar magnitudes. Both are about 160–170 each year.

2.4 Projection of U_{NM}

Let U^*_{NM} denote the following statistic that averages over the sample of worker-movie dyads:

$$U^*_{NM} = \frac{4}{JTNM} \sum_{c=1}^{J} \sum_{i \in \mathbb{N}(c)} \sum_{t=c+1}^{c+T} \sum_{m \in \mathbb{M}(t)} \bar{h}_{11}\left(W_{it}, Z_m\right) \tag{2.10}$$

In Section 2.2, I showed that the covariance of two pairs belonging to two different years is always zero even when they share a worker in common. A similar argument shows that $\text{cov}\left(\bar{h}_{11}\left(W_{it}, Z_m\right), \bar{h}_{11}\left(W_{is}, Z_p\right)\right) = 0$ for $t \neq s$. Therefore,

$$\begin{aligned}
\text{var}\left(\sqrt{NM} U^*_{NM}\right) =& \frac{16}{(JT)^2 NM} \sum_{c=1}^{J} \sum_{i \in \mathbb{N}(c)} \sum_{t=c+1}^{c+T} \sum_{m \in \mathbb{M}(t)} \text{var}\left(\bar{h}_{11}\left(W_{it}, Z_m\right)\right) \\
=& \frac{16}{JT} \Omega_{11}
\end{aligned}$$

Thus, $\text{var}\left(\sqrt{NM} U^*_{NM}\right) = \text{var}\left(\sqrt{NM} U_{NM}\right)$ as $N \to \infty, M \to \infty$.

2.5 Asymptotic equivalence of $\sqrt{NM} U_{NM}$ and $\sqrt{NM} U^*_{NM}$

To show that $\sqrt{NM} U_{NM}$ and $\sqrt{NM} U^*_{NM}$ are asymptotically equivalent, I prove that $\sqrt{NM} U_{NM}$ converges in mean square to $\sqrt{NM} U^*_{NM}$ as $N \to \infty$ and $M \to \infty$. Observe that:

$$NME\left[\left(U_{NM} - U^*_{NM}\right)^2\right] = \underbrace{NME\left(U^2_{NM}\right)}_{\to \frac{16}{JT}\Omega_{11} \text{ as } N \to \infty, M \to \infty} - 2NME\left(U_{NM} U^*_{NM}\right) + \underbrace{NME\left(U^{*2}_{NM}\right)}_{= \frac{16}{JT}\Omega_{11}}$$

It then suffices to show that $E\left(U_{NM} U^*_{NM}\right) = \frac{16}{JTNM} \Omega_{11}$. To see this, first note that:

$$\begin{aligned}
E(U_{NM} U^*_{NM}) =& \frac{4}{JTNM} \sum_{c=1}^{J} \sum_{i \in \mathbb{N}(c)} \sum_{t=c+1}^{c+T} \sum_{m \in \mathbb{M}(t)} \times \\
& \frac{1}{JT} \binom{N}{2}^{-1} \binom{M}{2}^{-1} \sum_{d=1}^{J} \sum_{(k,l) \in \mathbb{C}_{2,\mathbb{N}(d)}} \sum_{s=d+1}^{d+10} \sum_{(p,q) \in \mathbb{C}_{2,\mathbb{M}(s)}} \text{cov}\left(\bar{h}_{11}\left(W_{it}, Z_m\right), h_{klpq}\right)
\end{aligned} \tag{2.11}$$

Observe that $\text{cov}\left(\bar{h}_{11}\left(W_{it}, Z_m\right), h_{klpq}\right) > 0$ if and only if $\bar{h}(\cdot)$ and h_{klpq} share a worker-movie dyad in common. Without loss of generality, suppose that $k = i$ and $p = m$. Then the covariance is equal to the variance of a expectation of a score conditioning on one worker and one movie attribute, Ω_{11}:

$$\begin{aligned}
\operatorname{cov}\left(\bar{h}_{11}\left(W_{it}, Z_m\right), h_{klpq}\right) =&E\left[\bar{h}_{11}\left(W_{it}, Z_m\right) h_{klpq}\right] \\
=&E\left[E\left[\bar{h}_{11}\left(W_{it}, Z_m\right) h_{klpq} | W_{it}=W_{kt}=w, Z_m=Z_p=z\right]\right] \\
=&E\left[\bar{h}_{11}\left(W_{it}, Z_m\right) E\left[h_{klpq} | W_{it}=W_{kt}=w, Z_m=Z_p=z\right]\right] \\
=&E\left[\bar{h}_{11}\left(W_{it}, Z_m\right)^2\right] \\
=&\Omega_{11}
\end{aligned}$$

There are $JN(N-1)TM(M-1)$ terms of $\operatorname{cov}\left(\bar{h}_{11}\left(W_{it}, Z_m\right), h_{klpq}\right)$ that share a worker-movie dyad in common in Equation (2.11). Therefore,

$$\begin{aligned}
E(U_{NM}U_{NM}^{*}) =&\frac{4}{JTNM} \times \frac{1}{JT}\binom{N}{2}^{-1}\binom{M}{2}^{-1} \times JN(N-1)TM(M-1)\Omega_{11} \\
=&\frac{16}{JTNM}\Omega_{11}
\end{aligned}$$

Hence, $NME\left[(U-U^{*})^2\right] \rightarrow 0$ as $N \rightarrow \infty$ and $M \rightarrow \infty$.

2.6 Estimating the asymptotic variance of $\hat{\theta}$

Recall the influence function representation of $\hat{\theta}$:

$$\sqrt{NM}\left(\hat{\theta}-\theta_0\right) =\hat{\Gamma}\left(\bar{\theta}\right)^{-1} U_{NM} \tag{2.12}$$

Therefore, as $N \rightarrow \infty$ and $M \rightarrow \infty$:

$$\operatorname{var}\left(\sqrt{NM}\left(\hat{\theta}-\theta_0\right)\right) =\frac{16}{JT}E\left[\frac{\partial^2 \log l_{ijmn}^{cond}(\theta_0)}{\partial\theta\partial\theta'}\right]^{-1}\Omega_{11}E\left[\frac{\partial^2 \log l_{ijmn}^{cond}(\theta_0)}{\partial\theta\partial\theta'}\right]^{-1} \tag{2.13}$$

To estimate this asymptotic variance, I replace the population expectation of the Hessian with $\hat{\Gamma}\left(\hat{\theta}\right)$, defined in Section 2.1. To construct an estimate of Ω_{11}, I first compute for each dyad for a given worker i in cohort c and movie m in year t:

$$\hat{\bar{h}}_{11}\left(W_{it}, Z_m\right) \equiv\frac{1}{N-1}\frac{1}{M-1}\psi_{im} \tag{2.14}$$

where $\psi_{im} \equiv \sum_{j\in\mathbb{N}(c), j\neq i}\sum_{n\in\mathbb{M}(t), n\neq m}\frac{\partial \log l_{ijmn}^{cond}(\hat{\theta})}{\partial\theta}$. I then calculate:

$$\hat{\Omega}_{11} \equiv \frac{1}{JTNM} \sum_{c=1}^{J} \sum_{i \in \mathbb{N}(c)} \sum_{t=c+1}^{c+T} \sum_{m \in \mathbb{M}(t)} \hat{\bar{h}}_{11}\left(W_{it}, Z_m\right) \hat{\bar{h}}_{11}\left(W_{it}, Z_m\right)' \tag{2.15}$$

Replacing the terms in Equation (2.13) with their sample analogs and simplifying, I arrive at the following expression for an estimator of the asymptotic variance of $\hat{\theta}$:

$$\begin{aligned} \widehat{\text{var}\left(\hat{\theta}\right)} \equiv & \left(\sum_{c=1}^{J} \sum_{(i,j) \in \mathbb{C}_{2,\mathbb{N}(c)}} \sum_{t=c+1}^{T} \sum_{(m,n) \in \mathbb{C}_{2,\mathbb{M}(t)}} \frac{\partial^2 \log l_{ijmn}^{cond}(\hat{\theta})}{\partial\theta\partial\theta'}\right)^{-1} \times \\ & \left(\sum_{c=1}^{J} \sum_{i \in \mathbb{N}(c)} \sum_{t=c+1}^{c+T} \sum_{m \in \mathbb{M}(t)} \psi_{im}\psi_{im}^{-1}\right) \times \\ & \left(\sum_{c=1}^{J} \sum_{(i,j) \in \mathbb{C}_{2,\mathbb{N}(c)}} \sum_{t=c+1}^{T} \sum_{(m,n) \in \mathbb{C}_{2,\mathbb{M}(t)}} \frac{\partial^2 \log l_{ijmn}^{cond}(\hat{\theta})}{\partial\theta\partial\theta'}\right)^{-1} \end{aligned} \tag{2.16}$$

Inspection of Equation (2.16) reveals that the estimator admits a familiar sandwich variance formula. In this case, the "meat" of the sandwich is a summation over all the worker-movie dyads in the pair sample. The sample estimate of the expectation of the population Hessian is computed in standard statistical packages when performing a logistic regression. To compute ψ_{im}, I first reshape the pair sample to the worker-movie dyad-level. I then sum the scores by dyad.

2.7 Tables

Table 2.1: Number of covariance terms by number of workers and movies in common

Workers in common	Movies in common	$\text{cov}\left(h_{ijmn}, h_{klpq}\right)$	Number in sample
1	1	Ω_{11}	$J\binom{N}{2}\binom{2}{1}\binom{N-2}{1}T\binom{M}{2}\binom{2}{1}\binom{M-2}{1}$
1	2	Ω_{12}	$J\binom{N}{2}\binom{2}{1}\binom{N-2}{1}T\binom{M}{2}$
2	1	Ω_{21}	$J\binom{N}{2}T\binom{M}{2}\binom{2}{1}\binom{M-2}{1}$
2	2	Ω_{22}	$J\binom{N}{2}T\binom{M}{2}$

Notes: This table presents the number of covariance terms in the sample of pairs according to the number of workers and movies the two pairs share in common. See Section 2.3 for more information.

Bibliography

Abowd, John M. and Ian M. Schmutte, "Endogenous Mobility," *Unpublished Manuscript*, 2012.

_ **, Kevin L. McKinney, and Ian M. Schmutte**, "Modeling Endogenous Mobility in Wage Determination," *Unpublished Manuscript*, 2015.

Amman, John, "The Transformation of Industrial Relations in the Motion Picture and Television Industries: Craft and Production," *Under the Stars: Essays on Labor Relations in Arts and Entertainment*, 1996, pp. 113–155.

Åslund, Olof, Lena Hensvik, and Oskar Nordström Skans, "Seeking Similarity: How Immigrants and Natives Manage in the Labor Market," *Journal of Labor Economics*, 2014, *32* (3), 405–441.

Bayer, Patrick, Giorgio Topa, and Stephen Ross, "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market," *Journal of Political Economy*, 2009, *116*, 1150–1196.

Beaman, Lori A., "Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the US," *The Review of Economic Studies*, 2012, *79* (1), 128–161.

Beaman, Lori and Jeremy Magruder, "Who Gets the Job Referral? Evidence from a Social Networks Experiment," *The American Economics Review*, 2012, *102* (7), 3574–3593.

Bechky, Beth A., "Gaffers, Gofers, and Grips: Role-based Coordination in Temporary Organizations," *Organization Science*, 2006, *17* (1), 3–21.

BLS, "Contingent and Alternative Employment Arrangements February 2005," Technical Report, US Department of Labor July 2005.

Botsman, Rachel and Andrew Keen, "Can the Sharing Economy Provide Good Jobs?," *Wall Street Journal*, May 2015.

Brown, Meta, Elizabeth Setren, and Giorgio Topa, "Do Informal Referrals Lead to Better Matches? Evidence from a Firm's Employee Referral System," *Journal of Labor Economics*, 2016, *34* (1).

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The Value of Hiring Through Employee Referrals," *The Quarterly Journal of Economics*, 2015, p. qjv010.

Button, Patrick, "Do Tax Incentives Affect Business Location? Evidence from Motion Picture Production Incentives," *Unpublished Manuscript*, 2015.

Calvo-Armengol, Antoni and Matthew O. Jackson, "The Effects of Social Networks on Employment and Inequality," *The American Economic Review*, 2004, *94* (3), 426–454.

Chamberlain, Gary, "Analysis of Covariance with Qualitative Data," *The Review of Economic Studies*, 1980, *47*, 225–238.

_ , "Panel Data," in Zvi Griliches and Michael D. Intriligator, eds., *Handbook of Econometrics*, Vol. II, Elsevier Science Publishers, 1984, chapter 22, pp. 1247–1318.

Chandrasekhar, Arun G., "Econometrics of Network Formation," in Yann Bramoullé, Andrea Galeotti, and Brian W. and Rogers, eds., *Oxford Handbook on the Economics of Networks*, Oxford University Press, 2015.

Charbonneau, Karyne B., "Multiple Fixed Effects in Binary Response Panel Data Models," *Bank of Canada Working Paper 2014-17*, 2014.

Christopherson, Susan and Michael Storper, "The Effects of Flexible Specialization on Industrial Politics and the Labor Market: The Motion Picture Industry," *Industrial & Labor Relations Review*, 1989, *42* (3), 331–347.

Cingano, Federico and Alfonso Rosolia, "People I Know: Job Search and Social Networks," *Journal of Labor Economics*, 2012, *30* (2), 291–332.

Costantino, Nicola and Roberto Pietroforte, "Subcontracting Practices in USA Homebuilding: An Empirical Verification of Eccles's Findings 20 Years Later," *European Journal of Purchasing & Supply Management*, 2002, *8* (1), 15–23.

Cox, Dan, "IA Hits 'Hidden' Pix," *Variety*, May 18 1994, p. 1.

Crouch, Tanja L., *100 Careers in Film and Television*, Barron's Educational Series, 2002.

Duggan, Maeve, Nicole B. Ellison, Cliff Lampe, Amanda Lenhart, and Mary Madden, "Social Media Update 2014," Technical Report, Pew Research Center January 2015.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Bruecker, "Referral-based Job Search Networks," *Review of Economic Studies*, Forthcoming.

Eccles, Robert G., "The Quasifirm in the Construction Industry," *Journal of Economic Behavior & Organization*, 1981, *2* (4), 335–357.

Faulkner, Robert R. and Andy B. Anderson, "Short-term Projects and Emergent Careers: Evidence from Hollywood," *American Journal of Sociology*, 1987, *92* (4), 879–909.

Fernandez, Roberto M., Emilio J. Castilla, and Paul Moore, "Social capital at work: Networks and Employment at a Phone Center," *American Journal of Sociology*, 2000, pp. 1288–1356.

Glitz, Albrecht, "Coworker Networks in the Labour Market," *IZA Discussion Paper*, 2013.

Goldenberg, Anna, Alice X. Zheng, Stephen E. Fienberg, and Edoardo M. Airoldi, "A Survey of Statistical Network Models," *Foundations and Trends in Machine Learning*, 2009, *2* (2), 129–233.

Goldsmith-Pinkham, Paul and Guido W. Imbens, "Social Networks and the Identification of Peer Effects," *Journal of Business & Economic Statistics*, 2013, *31* (3), 253–264.

Graham, Bryan S., "An Empirical Model of Network Formation: Detecting Homophily When Agents are Heterogenous," *NBER Working Paper*, 2014, (w20341).

_ , "An Econometric Model of Link Formation with Degree Heterogeneity," 2015.

_ , "Methods of Identification in Social Networks," *Annual Review of Economics*, 2015, *7.*

Granovetter, Mark S., "The Strength of Weak Ties," *American Journal of Sociology*, 1973, pp. 1360–1380.

_ , *Getting a Job: A Study of Contacts and Careers*, University of Chicago Press, 1974.

_ , "Afterward: Reconsiderations and a New Agenda," in "Getting a job: A Study of Contacts and Careers," 2 ed., University of Chicago Press, 1995.

Hall, Jonathan V. and Alan B. Krueger, "An Analysis of the Labor Market for Ubers Driver-Partners in the United States," 2015.

Heath, Rachel, "Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garmet Factories," *Mimeo*, 2013.

Heckman, James J., "The Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process.," in Charles F. Manski and Daniel L. McFadden, eds., *Structural Analysis of Discrete Data and Econometric Applications*, Mit Press, 1981, chapter 4.

Hellerstein, Judith K., Mark J. Kutzbach, and David Neumark, "Labor Market Networks and Recovery from Mass Layoffs Before, During, and After the Great Recession," NBER Working Paper 21262, National Bureau of Economic Research 2015.

Hensvik, Lena and Oskar Nordström Skans, "Social Networks, Employee Selection and Labor Market Outcomes: Toward an Empirical Analysis," *Journal of Labor Economics*, Forthcoming.

Honoré, Bo E. and James L. Powell, "Pairwise Difference Estimators of Censored and Truncated Regression Models," *Journal of Econometrics*, 1994, *64* (1), 241–278.

Ioannides, Yannis M. and Adriaan R. Soetevent, "Wages and Employment in a Random Social Network with Arbitrary Degree Distribution," *The American Economic Review*, 2006, *96* (2), 270–2774.

_ **and Linda Datcher Loury**, "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 2004, *42* (4), 1056–1093.

Johnson, Ted and Rex Weiner, "Indies, IATSE Reach for Detente," *Variety*, February 19–February 25 1996, p. 9.

Jones, Candace, "Careers in Project Networks: The Case of the Film Industry," in Michael B. Arthur and Denise M. Rousseau, eds., *The Boundaryless Career: A New Employment Principle for a New Organizational Era*, Oxford University Press, 1996, pp. 58–75.

_ **and Kate Walsh**, "Boundaryless Careers in the US Film Industry: Understanding Labor Market Dynamics of Network Organizations," *Industrielle Beziehungen/The German Journal of Industrial Relations*, 1997, *4*, 58–73.

Kramarz, Francis and Oskar Nordström Skans, "When Strong Ties are Strong: Networks and Youth Labour Market Entry," *The Review of Economic Studies*, 2014, *81* (3), 1164–1200.

Krauth, Brian V., "A Dynamic Model of Job Networking and Social Influences on Employment," *Journal of Economic Dynamics and Control*, 2004, *28* (6), 1185–1204.

Lauzen, Martha M., "The Celluloid Ceiling II: Production Design, Production Management, Sound Design, Key Grips, and Gaffers," Technical Report, Center for the Study of Women in Televsion & Film 2009.

Lutter, Mark, "Do Women Suffer from Network Closure? The Moderating Effect of Social Capital on Gender Inequality in a Project-Based Labor Market, 1929 to 2010," *American Sociological Review*, 2015, *80* (2), 329–358.

Madigan, Nick, "IATSE Confab: A Harmonious Union," *Variety*, July 27–August 2 1998, p. 12.

Magruder, Jeremy R., "Intergenerational Networks, Unemployment, and Persistent Inequality in South Africa," *American Economic Journal: Applied Economics*, 2010, *2* (1), 62–85.

Marsden, Peter V. and Elizabeth H. Gorman, "Social Networks, Job Changes, and Recruitment," in "Sourcebook of Labor Markets," Springer, 2001, pp. 467–502.

Montgomery, James D., "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis," *The American Economic Review*, 1991, *81* (5), 1408–1418.

Montgomery, James D, "Weak ties, employment, and inequality: An equilibrium analysis," *American Journal of Sociology*, 1994, pp. 1212–1236.

Moretti, Enrico, "Social Learning and Peer Effects in Consumption: Evidence from Movie Sales," *The Review of Economic Studies*, 2011, *78* (1), 356–393.

MPAA, "Theatrical Market Statistics 2014," Technical Report, Motion Picture Association of America 2015.

Pallais, Amanda and Emily Glassberg Sands, "Why the Referential Treatment? Evidence from Field Experiments on Referrals," *Journal of Political Economy*, Forthcoming.

Rossman, Gabriel, Nicole Esparza, and Phillip Bonacich, "I'd Like to Thank the Academy, Team Spillovers, and Network Centrality," *American Sociological Review*, 2010, *75* (1), 31–51.

Saygin, Perihan Ozge, Andrea Weber, and Michele A. Weynandt, "Coworkers, Networks, and Job Search Outcomes," *Unpublished Manuscript*, 2014.

Schmutte, Ian M., "Job Referral Networks and the Determination of Earnings in Local Labor Markets," *Journal of Labor Economics*, 2015, *33* (1), 1–32.

Scott, Allen J., *On Hollywood: The Place, the Industry*, Princeton University Press, 2005.

Simon, Curtis J. and John T. Warner, "Matchmaker, Matchmaker: The Effect of Old Boy Networks on Job Match Quality, Earnings, and Tenure," *Journal of Labor Economics*, 1992, pp. 306–330.

Singer, Natasha, "In the Sharing Economy, Workers Find Both Freedom and Uncertainty," *New York Times*, August 2014.

Taub, Eric, *Gaffers, Grips and Best Boys: From Producer-Director to Gaffer and Computer Special Effects Creator, a Behind-the-Scenes Look at Who Does What in the Making of a Motion Picture*, Macmillan, 1994.

Topa, Giorgio, "Labor markets and Referrals," *Handbook of Social Economics, Amsterdam: Elsevier Science*, 2011, pp. 1193–1221.

Vogel, Harold L., *Entertainment Industry Economics: A Guide for Financial Analysis*, 9 ed., Cambridge University Press, 2014.

Weil, David, *The Fissured Workplace*, Harvard University Press, 2014.

Wooldridge, Jeffrey M., "Simple Solutions to the Initial Conditions Problem in Dynamic, Nonlinear Panel Data Models with Unobserved Heterogeneity," *Journal of Applied Econometrics*, 2005, *20* (1), 39–54.

Zenou, Yves, "Spatial Versus Social Mismatch," *Journal of Urban Economics*, 2013, *74*, 113–132.

Appendix A

Additional tables and figures

Table A1: Distribution of total jobs worked, years 1–10

Total jobs	Frequency	(%)
0	3381	49.53
1	988	14.47
2	590	8.64
3	395	5.79
4	299	4.38
5	248	3.63
6	155	2.27
7	143	2.09
8	107	1.57
9	79	1.16
10	66	0.97
11	78	1.14
12	48	0.70
13	54	0.79
14	42	0.62
15	34	0.50
16	27	0.40
17	16	0.23
18	14	0.21
19	11	0.16
20	8	0.12
21	13	0.19
22	5	0.07
23	6	0.09
24	3	0.04
25	8	0.12
26	3	0.04
27	1	0.01
30	2	0.03
33	1	0.01
59	1	0.01
Total	6826	100.00

Notes: This table presents the distribution of the total number of regular jobs worked among the grips and lighting technicians in the career sample during years 1–10. Jobs are screen credits in the worker's assigned position. The sample mean of the total number of jobs worked is 2.3. The variance is 15.6. See Section 1.3.2 for how I construct the career sample.

Table A2: Additional descriptive statistics of grips and lighting technicians

	Mean	Std. dev.	p25	p50	p75
Worker-level (6826 obs)					
Total regular jobs in years 1–10	2.279	3.953	0	1	3
Indirect connections in years 1–10	62.231	74.754	7	37	88
Direct connections in years 1–10	10.552	11.439	2	7	15
Regular jobs during initial year	1.121	0.391	1	1	1
Grip	0.518	0.500	0	1	1
Worker-year-level (68260 obs)					
Total regular jobs	0.228	0.620	0	0	0
Network size	4.846	5.723	2	3	6
Total indirect connections	10.854	14.153	0	5	16
Total direct connections					
Worked together at least once	1.976	2.918	0	1	3
Worked together at least twice	0.179	0.701	0	0	0
Worked together at least 3 times	0.053	0.356	0	0	0
Worked together at least 4 times	0.027	0.252	0	0	0
Worked together at least 5 times	0.035	0.299	0	0	0
Previous jobs (regular)	2.246	2.503	1	1	2
Previous jobs by type of movie					
Missing budget	0.308	0.597	0	0	1
Low budget	0.307	0.573	0	0	1
Medium budget	1.072	1.468	0	1	1
Big budget	0.559	1.219	0	0	1
Thriller	0.770	1.224	0	0	1
Action	0.555	1.031	0	0	1
Adventure	0.333	0.709	0	0	0.0
Comedy	0.960	1.347	0	1	1
Crime	0.489	0.881	0	0	1
Drama	1.128	1.467	0	1	1
Scifi	0.258	0.613	0	0	0
Horror	0.174	0.440	0	0	0
Fantasy	0.240	0.567	0	0	0
Romance	0.470	0.816	0	0	1
Mystery	0.280	0.609	0	0	0
Family	0.199	0.501	0	0	0

Notes: Total jobs are the total number of regular jobs worked over the period indicated. "Grip" is an indicator that equals 1 if the worker is a grip and 0 if the worker is a lighting technician. A worker's network is composed of all keys who supervise the worker on a movie released in an earlier year. See the notes to Table 1.4 for the definitions of the connection variables. I discretize production budgets into 4 categories: low ($\leq$ \$12m), medium (\$12m–\$70m), big ($\geq$ \$70m), and missing. Movies can be assigned multiple genres. See Section 1.3 for background on constructing the sample.

Figure A1: Motion picture production employment, by year

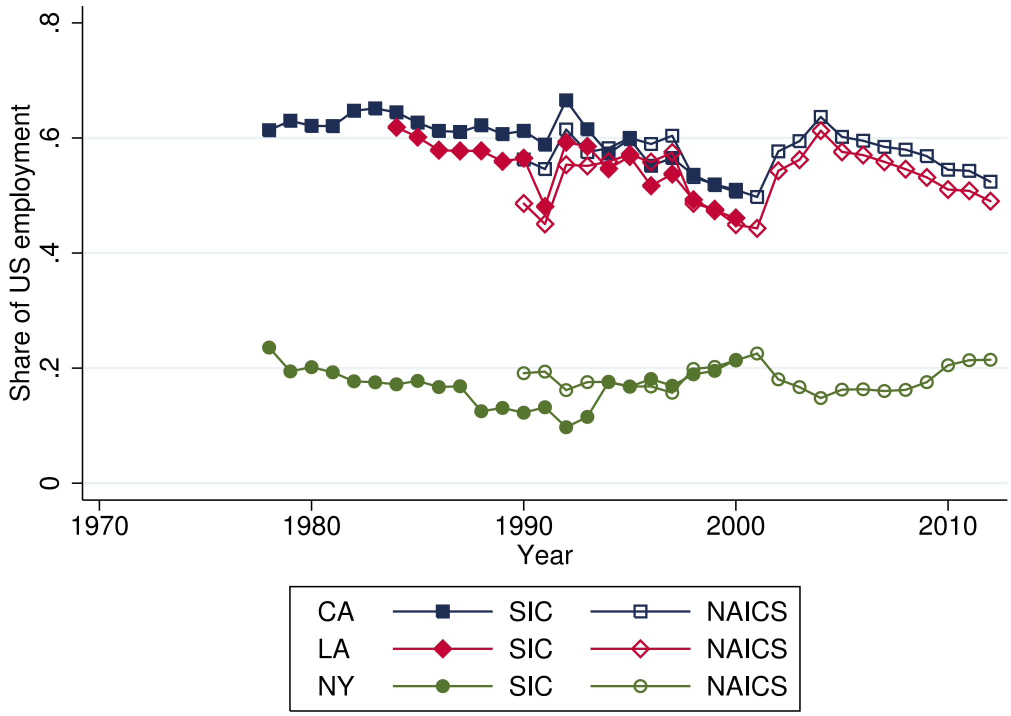

Notes: This figure shows the share of US motion picture production employment in California (CA), Los Angeles County (LA), and New York State (NY) over time. Counts of employment are from Bureau of Labor Statistics's Quarterly Census of Employment and Wages (QCEW). Annual averages of the QCEW by industry can be downloaded at `http://www.bls.gov/cew/datatoc.htm` (last accessed: May 16, 2015). The labels "SIC" and "NAICS" refer to SIC or NAICS-based definitions of the motion picture production industry, respectively. I measure employment using the 5 digit NAICS 51210, "Motion pictures and video production," and by subtracting SIC 7819, "Services allied to motion pictures," from the 3 digit SIC 781, "Motion Picture Production and Services."