The Formation of Director Networks and Their Effect on Governance: Evidence From Corporate Relocations*

Sam Gifford†

October 20, 2024

Abstract

Research on networks of the board of directors of corporations has generally found positive effects for firm performance, but causal evidence has been sparse. Using data on corporate relocations, I find that following a relocation, incumbent firms in the relocation area are more likely to have future movements of directors between not only the relocating firm, but also their primary and secondary networks. This suggests the formation and usage of networks, which influences board composition. While there are changes to board composition following relocation, this network formation does not influence outcomes related to governance and management quality, or to overall firm performance. This null effect could be a consequence of either 'rubber stamp' boards that do not have power or could be masked by heterogeneity of network effects across firms.

Keywords: Corporate Governance; Board Behavior; Network Formation

JEL classification codes: G34, D85

*I would like to thank Ben Feigenberg and Ben Ost for their helpful comments throughout the process of writing this paper.
†University of Illinois at Chicago. sgiffor2@uic.edu

1 Introduction

Research on the governance practices of corporations has seen increasing attention in recent years due to high-profile fraud cases and the general trends toward environmental, social, and governance (ESG) investing. Despite this, causal evidence on boards of directors and their efficacy remains relatively rare in part due to the endogenous circumstances that lead to selection of board members. Given that company boards generally select their own members subject to approval by shareholders, studying the circumstances that lead to this selection process provides a unique opportunity to tackle the selection problem that plagues research in this field. I study how directors of companies form networks by exploiting plausibly exogenous variation in distance between firms, then use these networks to estimate the effect of director networks on governance quality and firm performance, providing insights into the process of how governance impacts firms.

Studying director networks is difficult because social interactions themselves are not observed. Instead, networks may be indirectly observed by looking at repeated patterns of movement of directors across companies or inferred based on directors who simultaneously serve on multiple boards, which almost all existing research uses[1]. While repeated interactions between two firms may be suggestive of the existence of a network, the competing explanation that the two firms are simply similar in ways that cannot be observed in data cannot easily be ruled out and is rarely addressed by existing research[2]. To isolate causal evidence of network effects I employ a difference-in-differences design using relocation of corporate headquarters. I compare the probability of a director movement from incumbent firms towards firms likely to be in the relocator's network before and after the relocation event, with the difference being attributed to network effects. I then measure outcomes related to firm performance and governance quality before and after relocation to see the net effect resulting from this change in governance.

The intuition behind the relocation design is that firms tend to sort into geographic locations based on preferences, leading to similar firms in the same area. These similar firms may therefore have similar directors who are likely to move to and from these companies, creating the illusion of a network. When a company relocates, they are similarly selecting into an area due to changes in either their own preferences or because of geographic trends. The relocating firm's network, however, does not choose to relocate. Once the relocating firm creates new network connections in their new headquarters location, if network effects are present then there should be an increased flow of directors to and from the incumbent firms in the

[1] See for instance Akbas et al (2016), Bakke et al (2024), Li et al (2023), Liu et al (2022), Mol (2001), and Renneboog & Zhao (2011)

[2] Khwaja & Mian (2005) uses a difference-in-difference design using political connections in Pakistan while Bakke et al (2024) uses a difference-in-difference design using deaths of directors in Canada. These are the only causal papers on the topic I am aware of.

relocation area and the network of the relocating firm. The relocation breaks the established homophily patterns and allows for an exogenous connection between these previously disjoint networks.

My key assumption for this design is the parallel trends assumption, which requires that incumbent firms that are located near relocating firms would have their outcomes trend similarly to other firms in the absence of this relocation. I also require a similar assumption for the presumed network of the relocating firm, requiring them to not have trends in their characteristics over time in a manner that would increase their probability of a director movement toward firms in the relocation area. To assess these concerns my analysis proceeds in the following manner. I first study the variation in director movement that exists to provide evidence that firms with repeated movements are likely operating through networks. In this process I find that while movements are strongly related to company characteristics, the actual firm matching process is surprisingly idiosyncratic, which rules out some of the most obvious violations of the parallel trends assumptions. I then assess how distance affects director movements to provide support that firms in close proximity are more likely to form networks and that variation through relocation is similarly operating through this channel. I then assess geographic variables to rule out concerns that observed changes in director movements are being driven by changing geographic preferences, and finally estimate my model and assess the trends of outcomes prior to relocation.

I find several results related to how directors move across companies. I find that firms have a preference towards electing directors who are within commuting distance to their firm. In particular, I find that following relocation, the relocating firm is more likely to elect directors near their new headquarters location at the expense of firms in the old locale. More importantly, I find that incumbent firms in the receiving area are much more likely to acquire directors from the network of the relocating firm regardless of whether these network firms are local or not. Taken together, I find that director networks play an important role in determining the composition of boards of directors. Despite these compositional changes, I find little evidence that governance quality or firm performance is affected by these changes, though I cannot rule out small to moderate effects.

My analysis contributes to the literature on director networks, on environmental, social, and governance (ESG) investing, as well as the literature on networks in labor markets. This analysis is of direct importance to the finance literature as ESG investing has gained a large market premium in recent years (Gillan et al, 2021) and networks provide an understudied way to assess the governance component of this trend. This also relates to issues of public economics and contract theory regarding shareholder welfare, executive compensation, and managing the principal-agent problem. My main contribution is the addition of causal evidence to the effects of director networks.

The remainder of this paper is organized as follows: section 2 describes the setting and background in more detail, section 3 describes the data, section 4 contains a theoretical framework, section 5 contains the empirical methodology, section 6 contains first stage results relating to network changes following relocation, section 7 contains reduced form results of relocation's effect on firms, and section 8 concludes.

2 Institutional Details

2.1 The Board of Directors

The board of governors performs several functions for a corporation, including determining the CEO, setting executive salaries, and voting on key strategic decisions of the company. Modern boards are often split out into separate committees that perform specific roles, such as audit or compensation committees (Becht et al 2003) . In theory, the board of governors can have a large ability to influence the company's direction through these channels. In practice board effectiveness can be limited in real-world scenarios as shareholders may have limited voting options (Cai et al 2013) and executives may have undue influence over the board (liu et al, 2022). One relatively common scenario is when the CEO also serves as chairman of the board or has a majority stake in the company. In these situations, the CEO has substantial influence over the selection of board members, creating conflicts of interest. Such boards are sometimes termed 'rubber stamp' boards as they effectively have no power and simply agree with whatever the CEO decides (kemp, 2006).

Even if the board is independent, there are similar reasons why boards may fail to govern a company effectively. Board members do not run the company's day-to-day operations and rely on top executives to provide them with relevant information. Since executives have more information, this gives rise to a classic principal-agent concern where the executive can manage the flow of information to benefit themselves at the expense of shareholders (Jensen & Meckling, 2019). If boards are generally ineffective, then changes to boards may have little impact on overall firm metrics. Changes in director appointments may also contribute to either increases or decreases in this efficacy based on whether members are appointed based on their governance ability or their likelihood to collude with existing directors or executives to enrich themselves at the expense of shareholders.

Unlike executives, directors of a company tend to be fluid and move between companies relatively frequently or serve on multiple boards simultaneously, with roughly 10% of directors changing appointments each year. These movements of directors create potential linkages between firms whereby information and potential director recommendations can flow. Descriptive evidence suggests that linked

firms are several orders of magnitude more likely to have future movements of directors between them, even accounting for observable similarities, which is explored in more detail in the data section. It is generally believed that networks play a role in corporate governance nominations as these networks can enhance the strategic position of a company (Mol, 2001).

Governors are appointed based on the bylaws of the corporation charter, but generally follow one of a few methods. One of the board's committees will include a nominating committee that determines who will be on the ballot for governor appointments, whom shareholders of the corporation will then vote on. Different corporations have different rules for elections, but the two basic forms are plurality voting and majority voting. In plurality voting schemas, the individuals with the most votes will be elected, creating the potential for governors to win elections with a very small share of the vote. Under a majority voting schema, an individual is only elected if they receive a majority of the votes. This creates a potential for contested seats, for instance, if 1 board position was up for election but no member received a majority vote. However, in practice this is extremely rare with Cai et al (2009) finding that this only occurred in 4 out of 2488 cases in their setting. How these are handled in practice can vary, but one common method is for the person with the most votes to temporarily fill this position then offer their resignation. This resignation is then typically declined by the board, and the director continues in their position despite not having met the majority voting requirements (Cai et al 2009).

Given the potential for abuse of voting procedures outlined above, and given that the board itself determines nominees, it is worth questioning whether shareholders have any power to contest abuses that may arise. In practice, shareholders have several methods of remedy, with one particularly effective method being the threat of selling their shares, reducing the price of the stock. Large shareholders may have some influence over boards of directors in this capacity. In the case of the contested board situation above, shareholders can also file shareholder lawsuits, and board and executive behavior has generally been shown to act in a manner that avoids the potential for litigation (Cai et al, 2009). Finally, even very small shareholders can go through proxy filings where they can force a vote on certain issues, including writing in of director appointments though evidence on the efficacy is mixed with Karpoff et al (1996) finding little effect and Clifford (2008) finding that activist investors improve the operating performance of firms.

2.2 Firm Relocations

Due to the high fixed costs of relocating, firms generally remain in the initial corporate headquarters for long periods of time; in the sample of S&P 500 firms, more than 95% never change headquarters in the period from 2014 to 2022. Among large corporations, research on corporate relocations finds that, on

average, only half a percent to two percent relocate in any given year (Fatemi et al, 2024) which is consistent with data in my setting. An executive survey by West Monroe indicates that relocations mostly occur due to considerations over cost issues such as taxes and local labor markets. Given the large fixed costs and long lifetimes of corporations, relocations should only be expected to occur when these tax or labor conditions are expected to be relatively permanent.

Given the considerations above, it is expected that relocating firms may differ substantially from incumbent firms. In particular their strategic vision may systematically differ, which may also reflect a difference in governance quality, and a relocation event may also indicate a differential trend in governance quality over time rather than merely a level difference.

3 Data and Descriptive Statistics

3.1 Data Sources

I use data provided by publicly traded companies in the US through their SEC 10-K filings provided by SEC EDGAR and services that aggregate from this data, including Compustat. The universe of companies which this includes is generally large companies, with a median revenue slightly over 1 billion USD.

In order to identify director networks, I first make use of listings of all members of boards of governors from publicly traded companies on a year-by-year basis using data from Compustat Execucomp. From this data I am able to identify individual directors who move between companies over time and use this to create a potential network of firms. In general, movements of directors are common, and since companies often have boards that consist of between 5 and 7 governors, about a third of all firms will have a director move to or from another company in a given year. In my sample, roughly 500 directors move across 3000 companies per year. Director positions tend to only require a few hundred hours of work per year (McGregor, 2016), so it is not uncommon for these movements to be a director adding an additional appointment rather than actually changing companies.

I define two firms as being connected if they have a common director between the two in the past 5 years, regardless of the direction of the movement or whether the director simultaneously served on both boards, known as board interlock. I do not differentiate between the direction of this movement as is typical in the literature. While I use the prior 5 years when defining these connections, in practice companies that are linked based on this definition have a high propensity towards having a future movement of directors between them, so changing the definition from 5 years to an indefinite time period is unlikely to have any

major impact on any of the proceeding analysis due to the serial correlation of these movements. A minimum of 5 years is required, however, since on average an individual firm will only have 1 movement every 3 years. By using 5 years or more this, results in a connected graph, ensuring that every firm has an associated network. Since Execucomp data is populated started in 2009, my analysis sample is from 2014-2022 to allow for the initial graph to be populated.

To obtain data on corporate relocations, I downloaded the annual 10-K statements for all firms in my sample using the SEC's EDGAR tool. Each 10-K statement contains the address of the headquarters for the executives of the company, and I use regular expression based text parsing to extract this zip code. Starting in the year 2020, these documents have XML tags that allow the zip code to be directly parsed. For these years I was able to confirm that the regular expression based parsing algorithm matched the tagged data in over 95% of the sample. Headquarter relocations were then constructed by looking at changes in zip codes. The distance between zip codes is measured using the Department of Housing and Urban Development's zip code centroids, and a firm is classified as relocating in a given year if their address moves at least 50 miles. This threshold is mostly enforced to prevent very small movements that result from P.O. box changes or construction of new buildings near existing headquarters and also to match the treatment definition provided later. Of the relocations that were excluded based on this definition, the vast majority moved by less than 5 miles.

General data on firm performance, such as revenue and return on equity, is obtained from Compustat. Compensation data for directors and executives is obtained from Compustat Execucomp. Environmental, social, and governance (ESG) data comes from Thomson-Reuters Refinitiv ESG dataset, which contains both aggregated ESG component scores as well as individual sub-components. For instance, the number of shareholder lawsuits and controversies over executive compensation are among several dozen components that are used to compute the overall governance score.

3.2 Director Movements

Given the nature of director appointments, it is also worth considering what leads to directors moving across companies or serving on multiple boards at once. Due to anticompetition and noncompete agreements, it is typically not possible for directors to move between companies within the same industry, so almost all movements occur across broad industries. Director characteristics are generally not easily observed within data, but the characteristics of the firms they serve on are readily observed. I use a lasso regression to descriptively characterize which pairs of companies are most likely to have movements between directors based on their firm characteristics. This is presented in table 1. Observable

characteristics play a smaller role than might be expected in this, with lasso and xgboost models giving area under the ROC curve (AUC) values of less than 0.53. For reference, random guessing would produce an AUC of 0.5, a perfect classifier would produce 1.0, and a common threshold for a "satisfactory" classifier is 0.6. Despite this, several patterns emerge from this data.

First, whether two firms have a director movement between them is strongly predicted by past director movements. The bivariate model given in table 1 that only uses past director movement gives an AUC of 0.62, indicating that this is satisfactory in predicting director movements. This movement probability is roughly symmetric, which supports the graph connection definition that I use and that is commonly used in the literature. This is suggestive of network formation whereby information about suitable candidates flows between the two companies, but a competing explanation is simply that these companies are similar in unobserved dimensions that affect both past and future flows of directors. I provide evidence that this is operating through networks in section 6.

There are also additions to boards of directors who have not previously served on another board. In some of these cases the director may have come from a firm in the company's network but from a non-director position within that company. Unfortunately, this is not immediately observable in the data without additional text mining and these members are simply excluded from the network calculations and should be viewed as a form of measurement error.

Given this definition of a connection between two firms, I arrive at a graph. On average, each firm is directly connected to 10 other firms, which I refer to as a firm's direct or primary network. I can then define a firm's secondary network as all firms that are directly connected to the firm's primary network (excluding the primary network), which is, on average, 25 members. The secondary network has some useful features that are worth highlighting. First, when companies nominate board members it is likely that characteristics of the firm where they came from are directly considered. For this reason, unobservable characteristics of a firm and their primary network could potentially be similar if homophily-based matching is occurring, which may potentially raise concerns over endogeneity depending on the context. Unless firms actively seek out board members with the aim of improving their secondary network, however, secondary network characteristics are less likely to have these types of endogeneity concerns. Second, because the secondary network is much larger than the primary network it will give greater power in analyzing director networks. Given their distance to the originating firm, secondary networks are unlikely to have much or any influence over company decisions, but they may still have direct network effects by being referred to for new governor appointments through 'friend of a friend' arrangements.

Tertiary connections and beyond are also calculated, but the number of connections for each firm grows

rapidly, and the likelihood that these operate through the intended channels decreases as this distance increases. Each firm has an average of over 100 tertiary connections, and all firms form a connected graph with a diameter of 14. To simplify matters I only consider the primary and secondary connections.

The three most predictive characteristics of director movement using lasso regression are eigenvector centrality, distance between headquarters, and executive compensation. Eigenvector centrality, as calculated using the entire graph of directors over the preceding 5 years, effectively characterizes how connected a company is, with high values being more central to the network. While it is not entirely clear what contributes to eigenvector centrality, this result is an indicator that large, important companies that have high director churn or large boards are more likely to have a director movement between other companies, which is partially a mechanical effect. Executive compensation likely also acts as a proxy for company size, though this may also get at the demand side of the equation in terms of the desirability of working for the company. Distance is an interesting predictor as companies that are close may have more opportunities to network, though this is only a suggestive exercise as competing explanations cannot be immediately ruled out. The biggest competing explanation being that firms sort into geographic locations based on unobservable characteristics that also make them more likely to have similar director preferences. Due to the potential for idiosyncratic networking opportunities, I use variation in distance for my research design and devote most of my research efforts towards demonstrating that this is operating through network effects and ruling out competing omitted variables concerns.

3.3 Relocation

Using distance as a source of variation for networks directly has a number of issues due to potentially sorting based on geographic regions described above. To use variation in distance in a causal setting, I need to isolate changes in these distances over time, which only occurs due to new firms entering the market or due to existing firms relocating headquarters. Entry of new firms has a number of issues including the inability to observe any firm characteristics prior to the firm going public, which may be related to networks and geographic preferences.

Of course relocation also presents an opportunity to select into geographic regions based on firm characteristics. This creates an identification concern but is mitigated in my setting. The network of the relocating firm chooses not to relocate, and therefore is unlikely to have significant changes in their characteristics. This means that the relocating firm should have a network that is more closely related to their origin location compared to their destination location. This means that any changes to director movement between incumbent firms in the relocating area and the network of the relocating firm should be

occurring due to network effects rather than through sorting on unobservable characteristics. Changes in network characteristics are also testable since trends prior to relocation can be observed, ruling out concerns that these firms are also changing but not enough to trigger a relocation event. Relocators themselves, however, are likely contaminated for the reasons described above and are removed from my sample, which is described in more detail in section 4.

In total, I find that approximately 0.5-2% of publicly listed corporations relocate every year, which is consistent with existing literature. This means that over 95% of firms in the sample never relocate. Patterns of movement over time are shown in figure 6. As would be expected from a sample of large corporations, the majority of relocations occur within densely populated metropolitan areas. There are not any strong patterns of movement over time, except that densely populated areas tend to have consistently larger magnitudes, whether positive or negative. Of note is that within an individual state, all treatment tends to be populated within a single metropolitan area with little variation based on distance within that area.

One potential concern with using corporate relocation data is that firms preferentially sort into more valuable geographic locations so that any changes to company behavior is driven by local geographic trends. To assess these types of concerns, I gather state-level data from FRED concerning unemployment, population, and price levels to assess differential trends in treated areas, which are provided as a robustness check in the results section.

4 Theoretical Framework

There are a number of ways in which director networks may end up impacting the eventual hiring of directors to a board of directors. Since the board of directors itself appoints nominees for governors, networks may play an important role in determining who is actually considered for these positions. Unlikely regular jobs where there are public postings of job applications, the nominating committee needs to be aware of any qualified candidates for the role, which makes the potential for contacts of these directors to play a role. While directors can also be placed on a ballot through proxy statements, direct nomination by the nominating committee is by far the most commonly used method (Cai et al, 2009). While these nominees also have to win the vote, in practice it is rare for nominees to lose the vote given the voting schemas used by directors (Cai et al 2013). While institutional investors have been more active in voting for directors and appointing directors directly through proxy statements in recent years, it is still the case that networks may play a large role in determining whether individual directors end up serving on a board.

While networks may play a role in director appointments, ex-ante it is ambiguous how an exogenous increase in potential network size would affect the match quality of a firm's governors. First, while the pool of potential governors increases with an expanded network, the incumbent directors of a firm also have the potential to be poached by more companies who may provide a better outside option. Additionally, if directors end up having multiple simultaneous appointments through this channel, it may decrease performance since that director has less time to spend on the original firm (Brown et al, 2019). Second, the quality of a match can occur either through compatibility with other directors or with compatibility with firm objectives. In the latter case, it would be expected that overall governance quality may improve, but in the former it may only be expected to improve compensation packages for directors and executives. Liu et al (2022) and Akbas et al (2016) both find negative effects of board connectedness while most of the literature tends to find positive effects, as in (Bakke et al 2024; Li et al 2023; Mol, 2001). In the broader literature on job networks, Beaman and Magruder (2012) demonstrate that the incentive structures surrounding job referrals play an important role in whether networks lead to qualified candidates, and in the case of corporate boards, it is unclear whether directors are adequately incentivized to nominate qualified candidates.

This suggests that heterogeneity is expected on two major dimensions. First, firms that are better able to compensate their directors or otherwise have attractive qualities for directors are likely to find better match quality compared to low-quality firms that may bleed talent. Second, companies with better governance quality may find they are better able to match along management quality dimensions while those with poorer governance may instead nominate directors who support increasing compensation and decreasing oversight at the expense of shareholders. This suggests that institutional ownership, board independence, and similar metrics may be important variables for future heterogeneity analysis, though these are not explored at present. To the extent that this heterogeneity is present it may attenuate estimates as the positive and negative effects of expanded networks partially cancel each other out.

In my research design, I isolate changes to director networks that occur due to corporate relocation. Conceptually, there are multiple mechanisms through which distance could operate to expand networks. While closer corporate headquarters may facilitate social exchanges that lead to network formation, it's also plausible that these exchanges occur along other dimensions, such as attendance at conferences. Given that I only observe firm location, and not director location, directors themselves could live anywhere in the city, introducing measurement error. Any direct effect of proximity of corporate headquarters may be less impactful given that the directors may not spend the majority of their time near the corporate headquarters, but more importantly, would be difficult to isolate. For these reasons, I use a binary

treatment definition based on a 50 mile cutoff distance, which allows me to be inclusive of all of these potential mechanisms for how directors may form networks while being agnostic as to which these are operating through. While the 50 mile cutoff is arbitrary, this is effectively operating through metropolitan area and almost all of this variation occurs within 10 miles. This means that results should be insensitive to either changing this cutoff or changing the treatment definition to a metropolitan-based one

When defining treatment, I consider any incumbent firms within the relocating area to be treated, regardless of whether they directly had a director move to or from a relocating firm or its primary or secondary network. While I only observe networks indirectly through their associated director movements, it is plausible that directors across companies interact much more frequently. In this case there is still an exchange of information even if it does not lead to a change in board composition, so this form of effect would not be captured. Furthermore, companies may potentially compete for these directors, so if one company elects a director through their newly acquired network, this may lead to spillovers to other local firms, particularly if they have preferences for local directors. Finally, restricting to only firms who end up electing a new director may introduce sample selection as these firms may systematically be changing aspects of their governance or corporate strategy.

Of course, relocation may affect incumbent firms in ways other than through director networks. An increase in local firms also reflects an increase in the local supply of directors. To the extent that I observe director movement between incumbent firms and the relocating firm this could simply be due to local labor market forces. I therefore use the primary and secondary networks of relocating firms to confirm that changes to board structures of incumbent firms are occurring through networks resulting from the increased proximity of directors rather than through other methods. With that said, this only confirms that this is one channel through which this operates and does not rule out other channels, such as changes to local labor markets. As will be shown in section 6, the majority of these board changes are occurring through secondary networks rather than the incumbent firm itself, which shows that most of the effect is operating through this channel. Unfortunately, because I cannot rule out other channels entirely, I cannot credibly use relocations as an instrument for network changes, which would allow me to estimate the returns to director networks.

One other concern worth pointing out is that since corporate relocations can involve years of planning, there may be anticipatory effects when using an event study to measure impacts. While there is some evidence that the relocating firm changes its preference for local directors before actually relocating, any anticipatory effects of the secondary network and incumbent firms appear to be more muted, as shown in more detail in the results sections. Results are generally observed at the end of year, meaning that on

average the year of relocation will have 6 months where the firm had already relocated, plus some potential additional times for directors of incumbent firms to anticipate this change and network with the incoming directors. Thus, while it is plausible that networks may take time to form, in practice these appear to quickly manifest in the data. Firm outcomes following changes to directors, on the other hand, tend to show a delayed response.

5 Methodology

To find the effect of director networks on firm outcomes I use a difference-in-differences design using corporate relocations as the event. To support this design I need to first show that a first stage exists, i.e. that relocation actually changes director networks. I must then show that the parallel trends assumption holds, namely that the considered outcomes would have trended similarly for high-treatment status firms compared to low-treatment status firms in the absence of corporate relocations. In practice this means I must show that relocation is not driven by factors that systematically relate to either potential director movement probability or firm outcomes, with the biggest concern being that important firms systematically relocate to valuable locations.

Evidence for this design starts with descriptive patterns of director networks, showing that local proximity affects the probability of director movements, that this occurs entirely through commuting zones, and that connected firms have an order of magnitude higher likelihood of having a future movement. These suggest the use of a mover design using corporate relocations.

To use a corporate relocation design, I am broadly interested in two effects: the first stage of how relocations impact networks and subsequent board composition and the reduced form of how relocation changes firm outcomes, namely governance quality and bottom line and efficiency metrics. These are estimated using the following regression equations.

$$(1) connected_{ijt} = \beta_0 + \sum_{k=-4}^{4} \beta_k localfirms_{it+k} + \theta_i + \lambda_t + \varepsilon_{ijt}$$

$$(2) y_{it} = \beta_0 + \sum_{k=-4}^{4} \beta_k localfirms_{it+k} + \theta_i + \lambda_t + \varepsilon_{it}$$

Here, the treatment variable $localfirms_{it+k}$ is the number of corporations within 50 miles of firm i in year $t+k$. When interpreting coefficients in equation 1, $\hat{\beta}_k$ represents the increase in probability of firm i matching to firms in set j k years following relocation. The main set of firms considered is the secondary network of relocating firms[3]. Firm dyads, indexed by j, can be parsed in a few ways. This can include the probability of having any director movement with any firms in sample, having any movement with firms

[3]To avoid selection issues I keep the firms in network static as of the time of relocation

outside of their own primary network, having any movement with firms outside of their local geography, or having any movement with either relocators or their networks. When looking at relocator matching, this acts as a binary variable, e.g. following relocation, a firm is 2 percentage points more likely to match with their secondary network 1 year following relocation. For general movements, the coefficient represents the increase in probability of a director movement in a given year when the number of firms within 50 miles increases by 1. Given that the mean number of firms within 50 miles of a given firm is 40 with a standard deviation of 40, this number could be multiplied by a factor of 40 to get the effect of a 1 standard deviation increase in the firm density of a location on outcomes.

For firm-specific outcomes, this carries the same meaning, though it's worth noting that it may be expected that most changes to governance are likely to accrue only to firms that end up actually having a change in their board as a result of the relocations. In an instrumental variables framework, these could then be combined to yield an estimate of the returns to director networks. A relocating firm increases the probability of having a movement with the secondary network of that firm by approximately 1.2 percentage points, indicating a scaling factor of 83. Of course, this scaling only works to the extent that exclusion holds: if relocation affects firms through any channel other than direct matching with secondary networks of relocators, then this can significantly bias these estimates. In particular, I find that the number of changes to boards of directors does not substantially change following relocation, suggesting that a change to one board may impact who is elected on a separate board if firms are competing over the same local supply of directors. Similarly, changes to the quality of governance of one firm may influence other firms through competitive channels, or through information channels that may operate through networks. For this reason I will not focus on these estimates, but I will calculate them as a ballpark exercise.

Since relocation is likely strongly related to potential firm outcomes, relocators are removed from the sample such that only incumbent firms are considered. Because the sample is restricted to incumbent firms who never move, the firm fixed effect also acts as a narrow geographic fixed effect. This means that all variation in the number of local firms is occurring through corporation relocations. The time fixed effect then makes this a difference-in-differences estimator. The effect that is being identified is the differential trend, relative to the global time trend, for an individual firm's outcome following an increase in the total number of firms in their local proximity.

In order for this estimate to hold, it needs to be the case that, in the absence of any relocation, firms in areas that received relocations would have trended similarly to firms who did not receive such treatment. Ex-ante, there are a number of reasons why this might be expected to fail. Geographic areas may have characteristics that attract similar firms, leading to apparent increases in matching probability based on

locality. This by itself does not bias the difference-in-difference estimate, but to the extent that there are changes in these qualities over time and relocation matches firms to areas with similar preferences, this may bias the results. The specific concerns vary depending on whether the outcome of interest is the network formation or firm-level outcomes.

Identification issues regarding director movement tend to be easier to assess. The concern is that a geographic location changes characteristics over time and that this leads to relocation by firms who value these characteristics and then match with the remaining incumbent firms who have chosen not to relocate. In this case it would be observed that following relocation there is an increase in director movement caused entirely by unobserved confounding variables. This is unlikely to be significant for a few reasons. First, since more than 95% of firms have never relocated, a change to geographic characteristics means that firms that relocate to this area were likely attracted for different reasons than the incumbent firms that chose to stay. This still may pose an issue if existing characteristics were merely amplified[4], but these sorts of concerns would typically occur due to a trend, so looking at the trend in outcomes prior to treatment should assuage these concerns. Another important point is that the network of relocators are used for director movement outcomes, and since these firms did not choose to relocate themselves, these types of endogenous selection concerns are less compelling.

For firm-level outcomes, the major concern is that some geographic locations become more desirable over time, which benefits firms in these areas while simultaneously attracting more firms, creating an upward bias in the estimates. Outcome trends prior to relocation (pre-trends) can be used to assess any of these concerns that result from gradual shifts over time, e.g. resulting from local labor market conditions, particularly since it often takes several years for a company to actually relocate after initial planning. Fortunately trends for outcomes such as revenue are extremely stable in the pre-period, suggesting that such types of bias are not influencing the results. Rapid changes to a geographic environment, such as a sudden tax incentive, are more problematic since they may not appear in the pre-period if firms are able to relocate quickly in response. While these can't be ruled out in all cases, there is substantial variety in the timing and location of heavily treated locations in the data, as shown in figure 6, making these systematic concerns over timing less of an issue. Local unemployment and price levels are also considered as outcomes to rule out these concerns, and no differential trends are observed in either before or after relocation.

While difference-in-differences designs have raised concerns recently over issues of weighting of different treatment and control groups, this is a staggered, continuous treatment, dynamic model of differences-in-differences. As is shown in de Chaisemartin & D'Haultfoeuille (2023), the negative weighting

[4]for instance if low tax rates were made even lower, inducing initially marginal firms to relocate

issues in Goodman-Bacon (2021) are unlikely to occur when there are no time periods where most groups are treated, and no group that is treated most of the time, which are satisfied by my data. For this reason, I present the base difference-in-differences estimates.

6 Local Preferences and Mover Analysis

In order to use relocations as a source of variation I must first show that variation in distance is actually operating through the intended channel of local proximity rather than through other confounding variables. To do this, I first show that distance operates through commuting zones, then show that firms that relocate change their director movement preferences towards their new locality in favor of their old one. I then show that this increases the likelihood of matching through secondary networks and rule out concerns related to geographic trends over time that may occur if firms strategically relocate based on factors that are correlated with local firm characteristics.

If local proximity of companies make it more likely that directors interact, thereby increasing future likelihood of movement between these companies, it should be expected that most of this effect occurs within a distance that would make encounters more frequently, e.g. by being in the same commuting zone. As an illustrative example, if looking at a company located on the west coast, it would be unreasonable to expect that directors would be more likely to move to a company in Chicago compared to New York City solely through the distance channel despite being nearly a thousand miles closer as realistically the only way there would be an in-person interaction would be from flying out to a common location. For two companies located in San Francisco, however, it may be more reasonable for directors to naturally interact due to attendance at local conventions or chance meetups at local networking locations.

The regression $match_{ij} = \beta_0 + \beta_1 d_{ij} + \varepsilon_{it}$ which indicates the match probability between firms i and j who are located d miles apart from each other therefore tells us very little by itself. This can instead be broken out into distance rings to see where this effect is operating, i.e. $match_{ij} = \beta_0 + \sum_j \beta_j I(d_{ij} \leq \delta_j) + \varepsilon_{ij}$ where δ_j indicate binary distance cutoffs. Running the regression with cutoffs for 50, 100, 200, 500, 1000, and 2000 miles yields results whose magnitudes are by far the largest at at the d=50 level, suggesting that most of this effect is operating through reasonable commuting zones rather than larger distances that are likely driven by confounding variable concerns. These results are provided in table 2 which should not be interpreted causally, but merely as a way of describing the variation in director movements based on distance.

Distances less than 50 miles could also be used, and doing so yields more significant results compared to

using the 50 mile cutoff. With that said, there are theoretical reasons why it may be a bad idea to isolate the variation in distance below this threshold. While firms located within very close proximity to each other may interact more, the measurement that is of interest here is the proximity of directors, and their location of residence is unobserved. Furthermore, it is unlikely that directors are literally more likely to interact if their headquarters are very close, as they are presumably not meeting at each other's office building but rather at neutral places such as convention centers or country clubs. Combining these factors, it may truly be the case that headquarters being physically closer within the same commuting zone will increase interaction likelihood, but the effect is likely relatively small but, more importantly, difficult to isolate. Given this, I choose a binary 50 mile cutoff, which effectively acts as a metropolitan area indicator.

I find that following relocating, the relocating firm appears more likely to have director movements between firms in the relocating area at the expense of existing networks in the origin location, though the effect is noisy due to the small sample of relocators, with 527 relocators across 13 years, which is shown in figure 1. There is also some evidence that this change in preferences occurs prior to relocation, which supports the decision to remove relocators from the sample. These movement patterns by themself does not necessarily imply any network effects as this could be operating through local labor markets. Using the secondary network of these relocating firms dramatically increases my power.

Following relocation, incumbent firms in the relocating area are significantly more likely to have a director movement between the relocating firm's secondary network, increasing from around 0.4 percentage points per year to 1.6 percent, as shown in figure 2. While there is some evidence of an upward trend in the years preceding the movement, this jump is rapid following relocation and remains elevated afterward. Some of this trend in the pre-period may be anticipatory in nature as the moves would be announced in advance of the actual movement date. This suggests that firms are actually forming networks after increasing proximity to other firms rather than simply operating through any changes in local labor supply for directors.

7 Governance and performance outcomes

The results in section 6 provide evidence that networks are influencing the composition of boards of companies though it is unclear whether such changes would have an effect on the actual company's management and, if so, in which direction. When looking at firm outcomes, I consistently find statistically insignificant effects and have power to find moderate effect sizes. These are shown in table 4 and figures 3, 4, and 5 and are split between firm performance metrics, network effects, and governance quality metrics.

Results are consistently null, with 7 out of 72 point estimates between significant at the 0.05 level and none significant at the 0.01 level or below.

One of the largest point estimate increases comes from revenue, which, using the average of the values in the pre- vs post-periods gives an effect size of $2.2 million per additional firm within 50 miles. Translating this to a 1 standard deviation increase in firm density yields an estimate of around $89 million compared to a mean revenue within sample of $2.9 billion, or 3 percent. While this is not trivial, the lack of significance rules out all but small to moderate effect sizes. Using the suggestive instrumental variables scaling factor of 83 to get at the value of an additional firm in the network yields an estimate of $185 million, or 6.4% of revenue. Given that the average firm has 10 firms in their primary network, I cannot rule out meaningful effect sizes but given that exclusion is unlikely to hold in my setting this is likely an inflated number, even ignoring the lack of significance.

Outcomes related to governance quality similarly have point estimates that may be of marginal economic significance but lack statistical significance. If the results on executive compensation are taken at their face value and the instrumental variables scaling factor is used, this would suggest each additional network contact created through a relocation increases executive compensation by around $2 million, or 0.14 standard deviations. Similarly, the management score defined by Thompson Refinitiv increases by .054 per network connection, or 0.2 standard deviations. Due to limited power, I do not explore heterogeneity, but this effect is likely to vary based on existing levels of governance quality. Regardless, this does show some weak suggestive evidence that boards may be effective at improving shareholder value in some circumstances.

Outcomes related to overall movement of directors are small in magnitude and statistically insignificant. Given that firms are much more likely to have director movements to the secondary network of relocators, this suggests that a substitution effect is present since overall movement does not increase. The point estimate for any director movement is actually negative, and scaling by a 1 standard deviation increase in local firm density gives a decrease in movement propensity by 3.5 percentage points, which again is not statistically significant.

Table 3 gives estimates of the effect of relocation on state-level unemployment and population levels, which is only measured in the contemporaneous period due to power concerns. Ideally these results should not be significant as this may indicate an omitted variables concern whereby changes in firm performance outcomes are driven by changes in the local economic environment. I find statistically significant effects, but the actual magnitude of these changes is minuscule. Interpreting these results causally, 1 firm relocating decreases local unemployment by 0.14 basis points (0.0014 percentage points) and increases

population by 655 people. Of note is that the mean number of employees for relocating firms is 17281 compared to the estimated increase in population of 655. While the size of corporate headquarters is not observed, these are plausible causal estimates of the effect of corporate relocation and therefore should not pose an identification concern.

In total, I find that there is strong evidence that firm relocations change the composition of incumbent firms' directors through network effects, but there is only weak evidence that this results in any meaningful changes to governance quality within these firms. Given that ineffective rubber stamp boards are likely to have a null effect and therefore attenuate estimates towards zero, there may be room for heterogeneity analysis based on this type of classification of firms and their boards, suggesting that future analysis in this area may be fruitful.

8 Conclusions

Descriptive evidence suggests that networks may play a large role in how directors of companies move across firms. Using data on relocations, I find strong causal evidence that firms in close geographic proximity are more likely to interact and form networks and that these change the compositions of these boards of directors. I do not find strong evidence that these changing board compositions affect firm performance or governance quality, though I am unable to rule out small to moderate effects. Point estimates suggest that there may be moderate effects on governance quality and returns to social capital through networks. This contributes to the existing director network and governance literature by providing causal evidence for network formation and their impact on governance quality and firm performance. This also points to future research where heterogeneity along governance quality dimensions can be explored to estimate causal returns to director networks.

Tables

Table 1: Predictors of Director Movements

variable	(1)	(2)	(3)	(4)	(5)	(6)
(Intercept)	-6.264	-7.030	-7.322	-7.363	.	.
Same state	0.000	0.299	0.307	0.309	.	.
Neither missing employment	0.000	0.219	0.660	0.721	.	.
Minimum CEO Salary	0.000	0.138	0.158	0.163	.	.
Minimum DebtRatio	0.000	0.122	0.220	0.233	.	.
Neither missing CEO Age	0.000	0.121	0.160	0.232	.	.
Neither missing Total Compensation	0.000	0.114	0.128	0.210	.	.
Minimum Total Compensation	0.000	0.114	0.140	0.147	.	.
Neither missing CEO Gender	0.000	0.111	0.065	0.142	.	.
Neither missing CEO Salary	0.000	0.111	0.065	0.142	.	.
AUC	0.500	0.518	0.520	0.527	0.5186	0.6161

Each row represents a variable used to predict director movements. In total 238 variables were included, and the top 10 most predicted are displayed. Each variable represents either the minimum or the difference of a variable between a pair of firms. The variables that made it here are derived from headquarter state, number of employees, CEO salary and total compensation, CEO gender, and the debt/equity ratio of the firm. Each of the first 4 columns shows the standardized coefficients for lasso regression models with λ parameter equal to $0.001, 0.0001, 0.00001$, and 0, respectively. Each variable is standardized to be in the $[0, 1]$ interval to make direct comparisons across variables for variable importance. Model 5 represents an xgboost model whose parameters are not directly interpretable, and model 6 is a bivariate regression using whether the pair of firms had a prior director connection. The AUC metric represents the area under the curve of the receiver operating characteristic curve, which is a measure of classification strength. A value of 0.5 is as good as random assignment, while a value of 1.0 represents a perfect classifier. A value of 0.6 is considered to be the lower threshold for a "satisfactory" classifier.

Table 2: Association of Distance and Director Movement

Variable	Coefficient	Std. Error	tvalue
Intercept	0.0008893	0.0000153	58.1
Distance	-0.0000002	0.0000000	-23.9

Variable	Coefficient	Std. Error	tvalue
Intercept	0.0004	0.00002	19.8
D=50	0.0013	0.00009	15.2
D=100	0.0002	0.00009	2.2
D=200	0.0002	0.00005	3.0
D=500	0.0001	0.00003	4.8
D=1000	0.0002	0.00002	8.5
D=2000	0.0000	0.00003	-0.6

Table 2 displays regression results from regressing probability of pairwise director movement between firms and their corresponding distance. The units of the top table are in miles, and indicate that each extra mile between a pair of firms corresponds to a 2 in a thousand percentage point decrease in director movements each year. The bottom table decomposes this into binary distance thresholds and shows that most of this effect is operating within 50 miles. These are cumulative thresholds, so the average increase in connection probability for two firms is the sum of the coefficients from D=50 to D=2000. Virtually every value is highly statistically significant, but there is no causal interpretation here.

Table 3: State Variables and Relocation

Variable	Estimate	Std. Error	t value	Percent Change
Unemployment	-0.0014	0.00015	-9.28	-0.023
Population	0.6550	0.07945	8.24	0.004

This table represents the coefficients of the main specification for state level unemployment and population, but only for the contemporaneous year of relocation. If identification assumptions hold, these are the causal effects of relocation on population and unemployment. Unemployment is measured in percentage points (with a value of 1 representing 1%) and population is measured in thousands. The percent change indicates the percentage change from baseline following relocation in percent, i.e. while unemployment decreases it does so by 2.3 basis points.

Table 4a: Event Study Estimates: Network

T	EigenCentrality	NetworkSize	DirectorMovement
-4	-0.000121	0.018	0.003148
	(0.000271)	(0.012)	(0.002051)
-3	0.000285	-0.018	0.001627
	(0.00025)	(0.011)	(0.002052)
-2	-0.000005	0.002	-0.000349
	(0.000101)	(0.004)	(0.000735)
-1	0.000002	-0.005	0.000035
	(0.000098)	(0.004)	(0.000755)
0	0.00006	0.018	0.000694
	(0.00023)	(0.01)	(0.002103)
1	0.000126	0.003	0.000151
	(0.0002)	(0.008)	(0.001591)
2	0.000026	0.005	-0.002216
	(0.000205)	(0.008)	(0.001597)
3	-0.000577	0.02	-0.002343
	(0.000307)	(0.013)	(0.002358)
4	-0.000083	0.004	0.004865*
	(0.000205)	(0.009)	(0.001786)
Mean Value	0.683	10.28	0.152
n	38445	38445	38445

Significance: * .05 ** .01 *** .001

Each coefficient represents a causal estimate t years after corporate relocation with $T = 0$ being the year of relocation for each additional net firm within 50 miles. Standard errors are indicated below in parenthesis and are calculated using heteroskedastic autocorrelation robust errors. The mean value of the sample is presented at the bottom for comparison. Eigenvector centrality is a measure of network connectedness varying from 0 to 1. Network size is the number of firms within a firm's direct network, i.e. the number of firms with a common past director movement. Director movement indicates the probability that a director to or from any firm in the sample, measured in percentage points (a value of 0.01 indicates 1 percentage point).

Table 4: Event Study Estimates: Performance

T	Revenue	Debt/Equity	ROE
-4	0.483	0.007	0.00028
	(2.507)	(0.009)	(0.00274)
-3	-3.759	0.006	-0.00206
	(2.511)	(0.009)	(0.00272)
-2	-0.705	-0.002	-0.00039
	(0.899)	(0.003)	(0.001)
-1	-1.96*	0	0.00048
	(0.925)	(0.003)	(0.00104)
0	-3.198	0.021*	-0.00307
	(2.568)	(0.01)	(0.00288)
1	-1.501	-0.007	-0.00012
	(1.95)	(0.007)	(0.00217)
2	0.246	-0.01	-0.00036
	(1.956)	(0.007)	(0.00212)
3	7.017*	-0.021	0.00289
	(2.889)	(0.011)	(0.00323)
4	1.169	0.016	-0.00207
	(2.182)	(0.008)	(0.00237)
Mean Value	2915	0.784	0.165
n	38445	38445	38445

Significance: * .05 ** .01 *** .001

Each coefficient represents a causal estimate t years after corporate relocation with $T = 0$ being the year of relocation for each additional net firm within 50 miles. Standard errors are indicated below in parenthesis and are calculated using heteroskedastic autocorrelation robust errors. The mean value of the sample is presented at the bottom for comparison. Revenue is total firm revenue measured in millions of USD. Debt/Equity is the ratio total debt to equity, winzorized at the 5th and 95th percentiles to avoid ratio issues. ROE is return on equity, measured in percentage points (a value of 0.01 indicates 1 percentage point).

Table 4: Event Study Estimates: Governance

T	GenderDiversity	Management	CEOCompensation
-4	-0.00049	-0.00043	-15.553
	(0.00112)	(0.00123)	(20.562)
-3	-0.00214*	-0.00185	-46.276*
	(0.00104)	(0.00114)	(19.319)
-2	0.00061	0.00016	4.358
	(0.00042)	(0.00046)	(7.737)
-1	0.0002	0.00048	-15.092*
	(0.00043)	(0.00047)	(7.571)
0	-0.00055	-0.00044	-17.512
	(0.00099)	(0.00109)	(18.122)
1	0.00071	0.00122	-0.835
	(0.00077)	(0.00084)	(14.327)
2	0.00111	0.00071	23.176
	(0.00079)	(0.00086)	(15.013)
3	-0.00021	0.00073	36.903
	(0.00124)	(0.00136)	(22.74)
4	-0.00055	-0.00099	-3.407
	(0.00099)	(0.00109)	(15.668)
Mean Value	0.553	0.593	6975
n	38445	38445	38445

Significance: * .05 ** .01 *** .001

Each coefficient represents a causal estimate t years after corporate relocation with $T = 0$ being the year of relocation for each additional net firm within 50 miles. Standard errors are indicated below in parenthesis and are calculated using heteroskedastic autocorrelation robust errors. The mean value of the sample is presented at the bottom for comparison. Gender diversity gives the percentage of males serving on the board of directors. Management is the management score provided by Thomson Refinitiv ESG and is an aggregation of metrics concerning shareholder welfare ranging from 0 to 1. CEO compensation is the total compensation of the CEO including bonuses, measured in thousands of USD

Figures

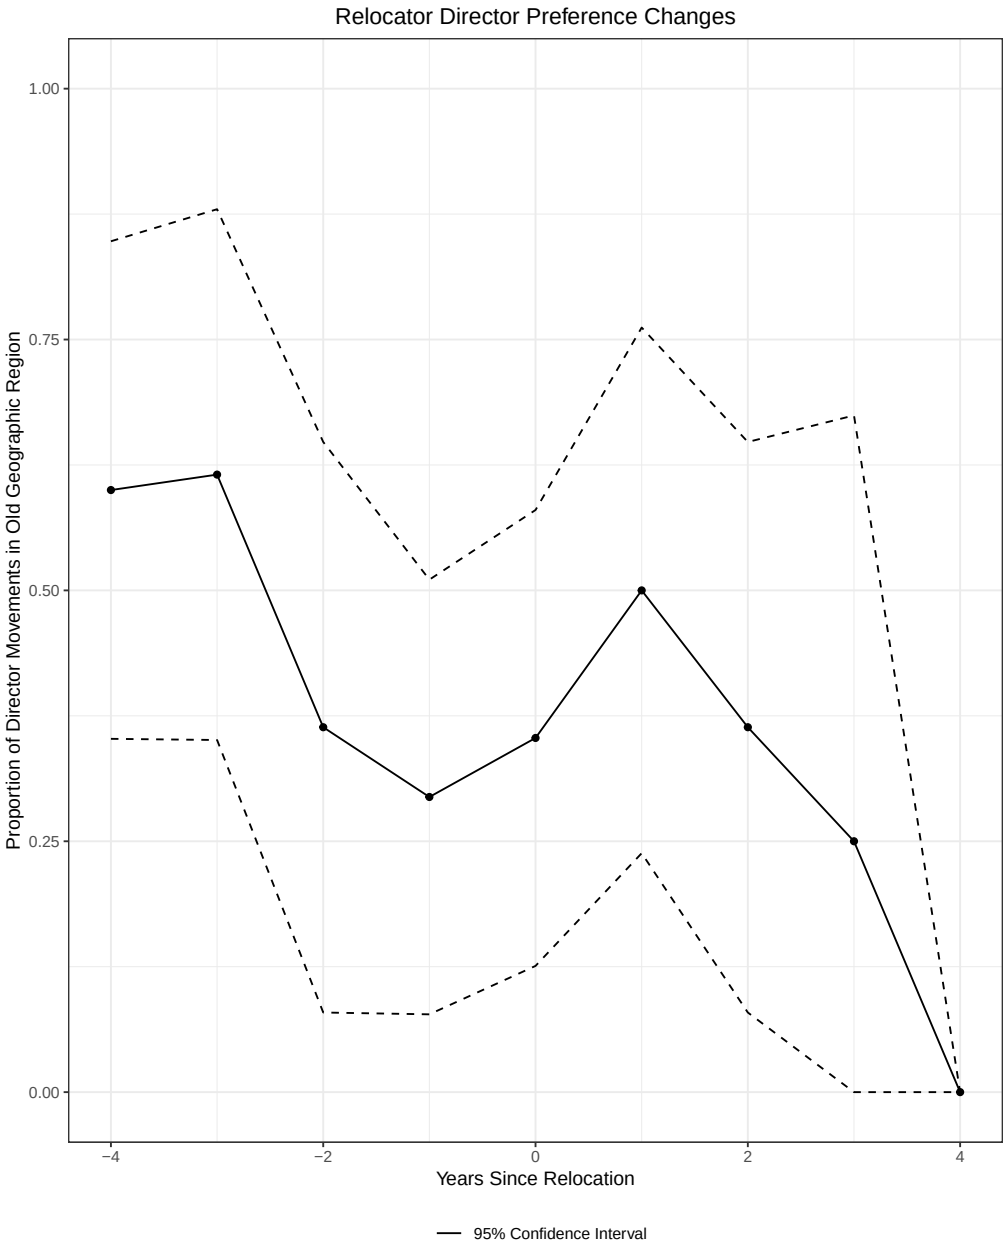

Figure 1: Proportion of director movements occuring in old headquarters location for relocators before and after relocation. Sample size is 674 relocators across 13 years

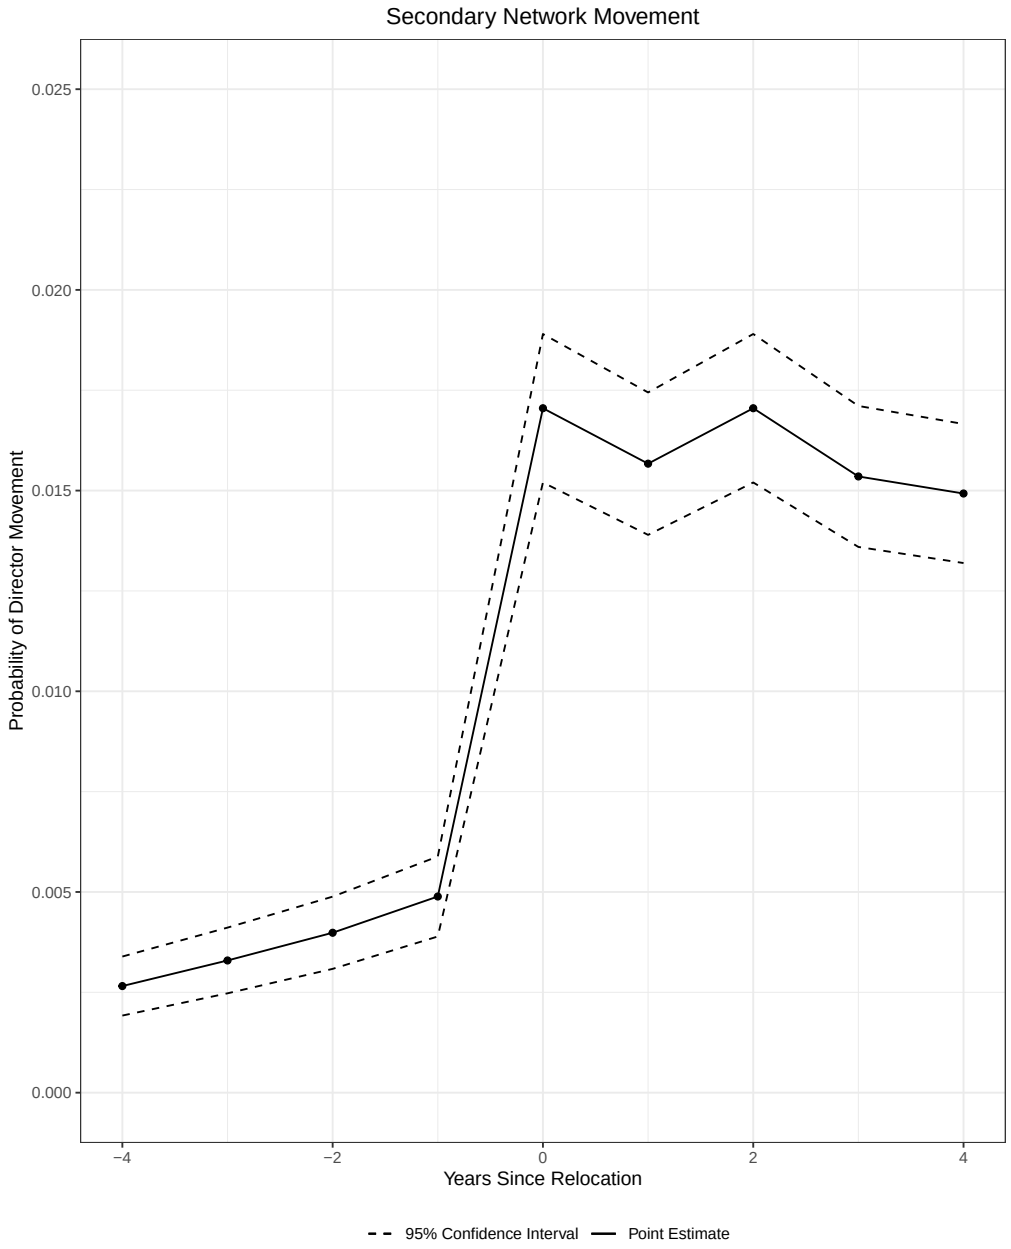

Figure 2: Probability of incumbent firm having a director move to or from the secondary network of relocations before and after relocation

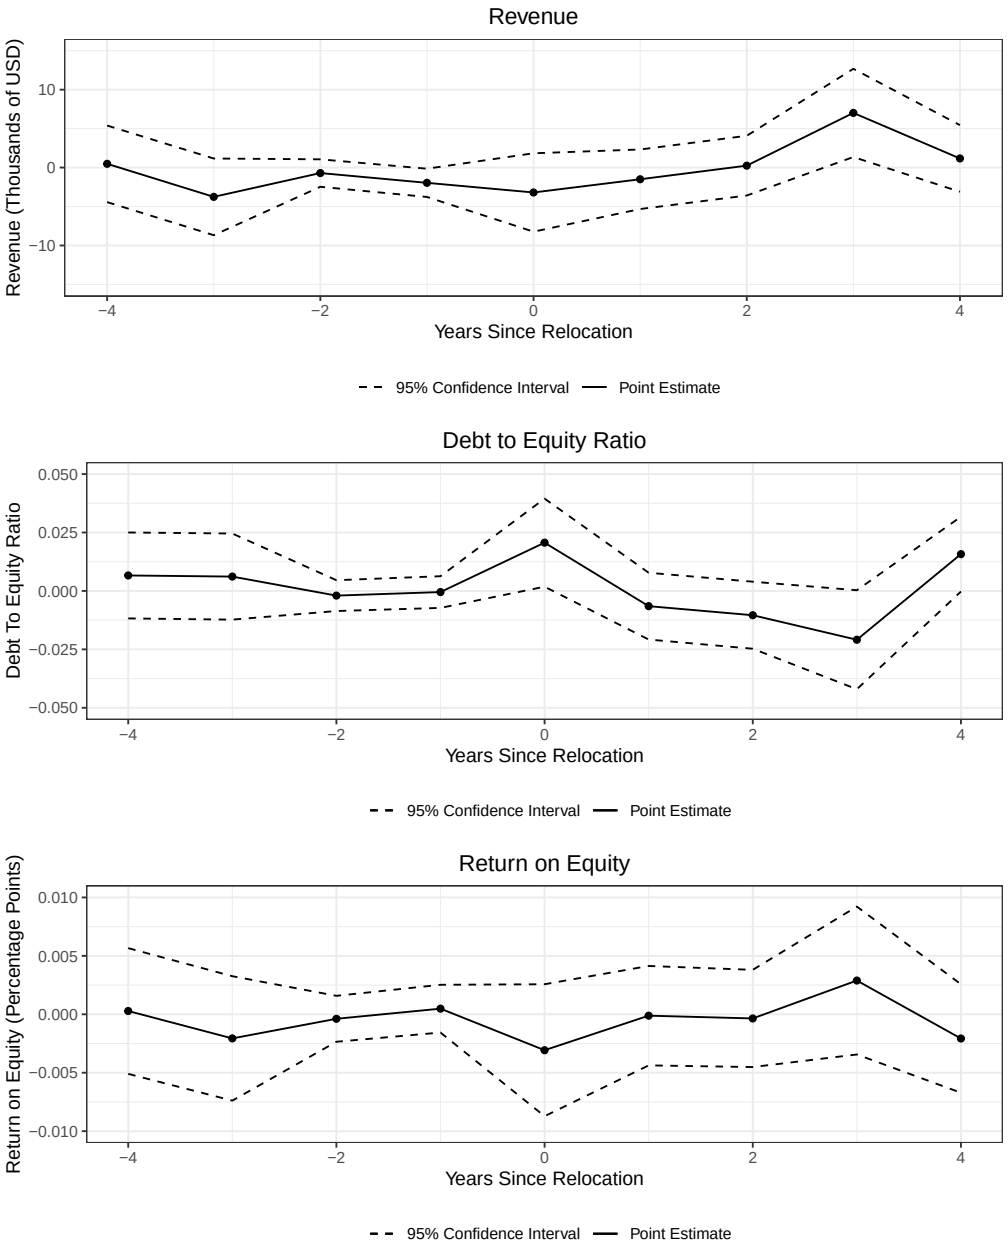

Figure 3: Firm Performance Metrics

Each point represents a coefficient from the main specification, namely the causal effect of an additional firm within 50 miles of the incumbent firm's headquarter locations t years after relocation. The year of relocation is defined at t=0. 95% confidence intervals are calculated using heteroskedastic and autocorrelation robust errors. Variable definitions are also included in table 4.

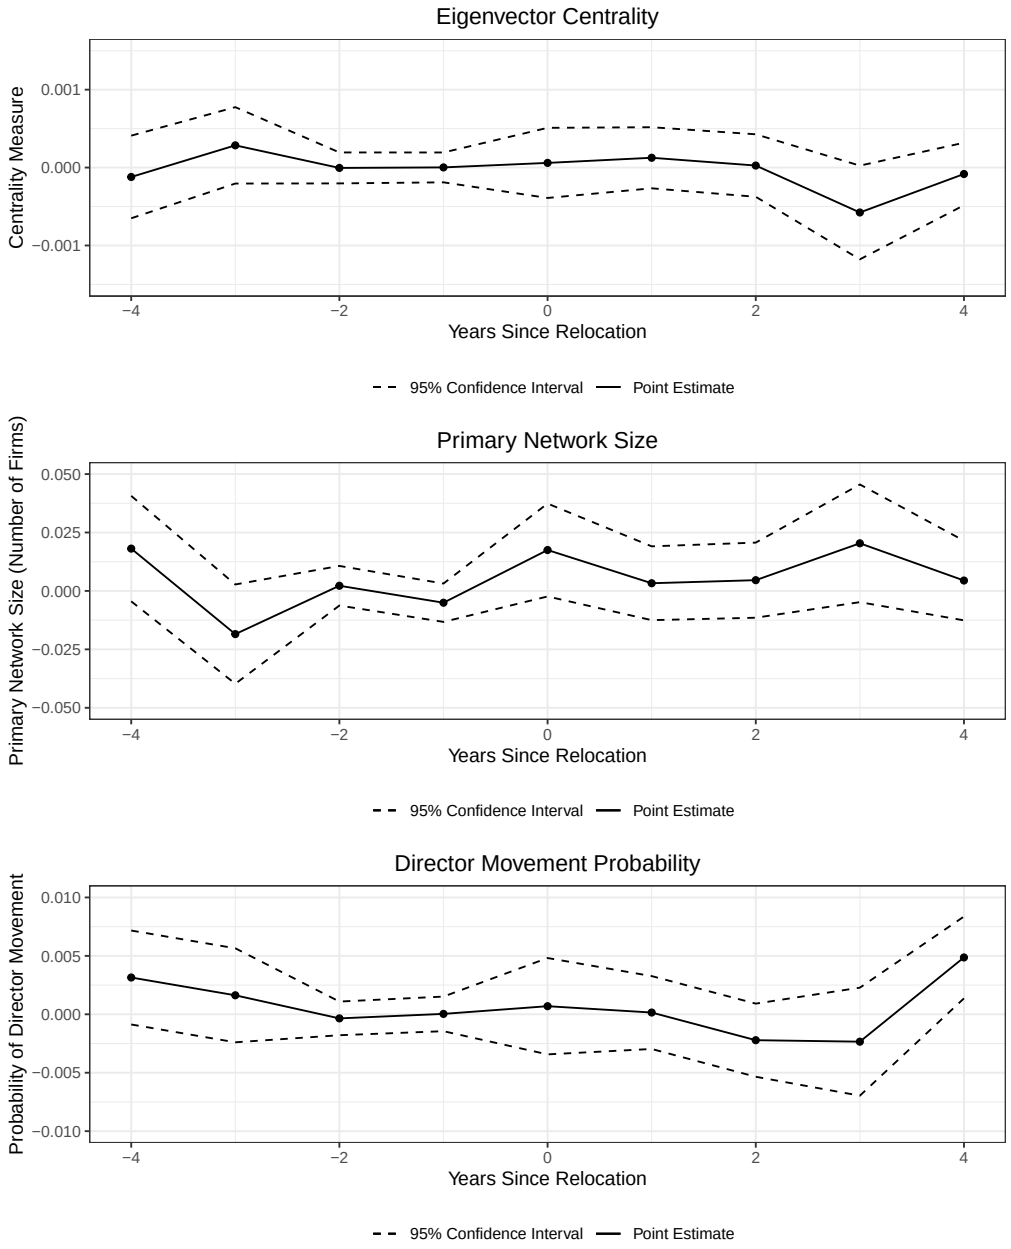

Figure 4: Firm Network Metrics

Each point represents a coefficient from the main specification, namely the causal effect of an additional firm within 50 miles of the incumbent firm's headquarter locations t years after relocation. The year of relocation is defined at t=0. 95% confidence intervals are calculated using heteroskedastic and autocorrelation robust errors. Variable definitions are also included in table 4.

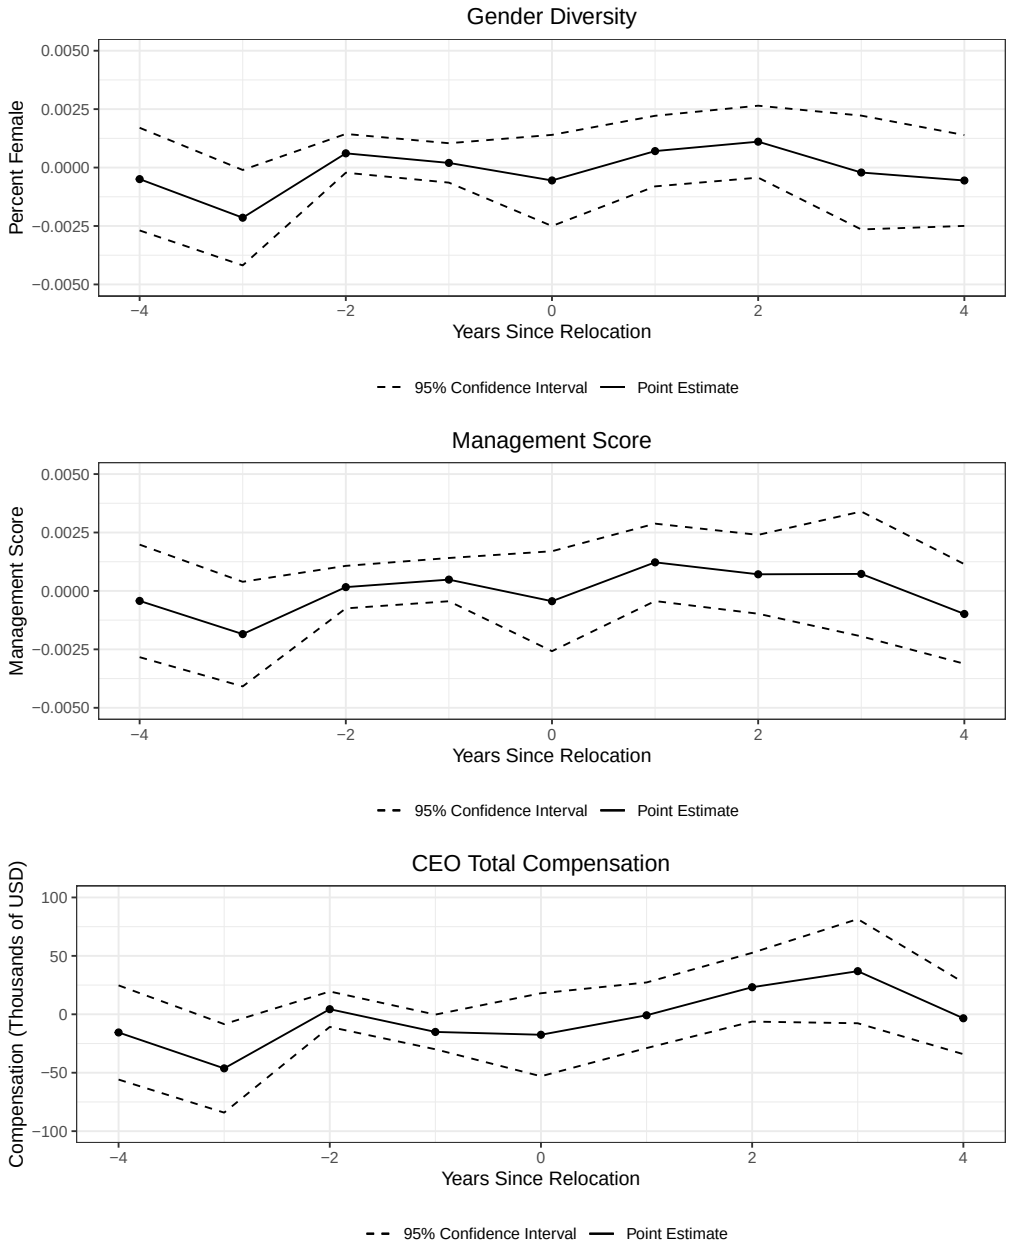

Figure 5: Firm Governance Metrics

Each point represents a coefficient from the main specification, namely the causal effect of an additional firm within 50 miles of the incumbent firm's headquarter locations t years after relocation. The year of relocation is defined at t=0. 95% confidence intervals are calculated using heteroskedastic and autocorrelation robust errors. Variable definitions are also included in table 4.

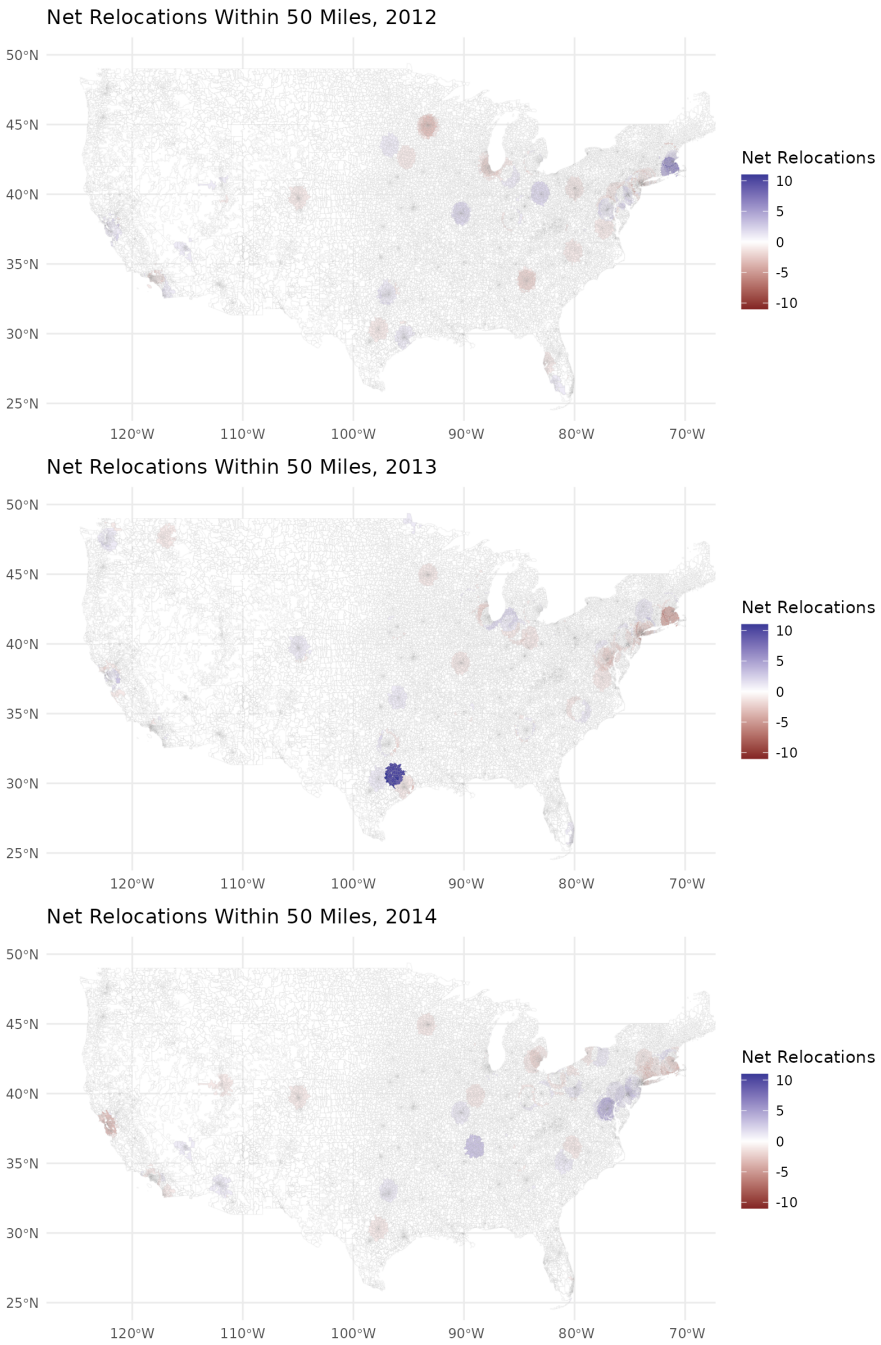

Figure 6a: Net Relocations 2012-2014

Number of firms relocating within 50 miles of the area minus number of firms leaving the area.

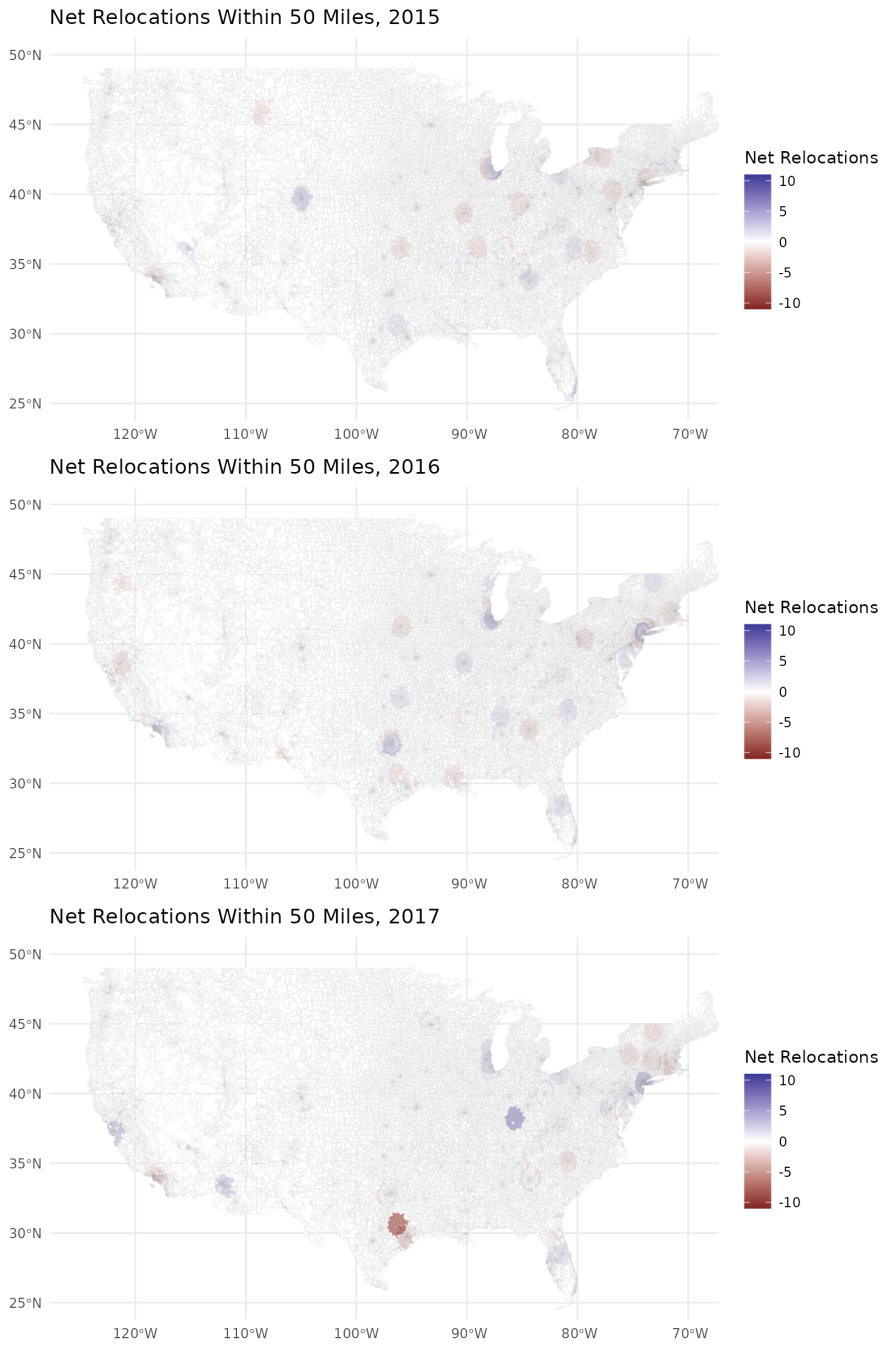

Figure 6b: Net Relocations 2015-2017

Number of firms relocating within 50 miles of the area minus number of firms leaving the area.

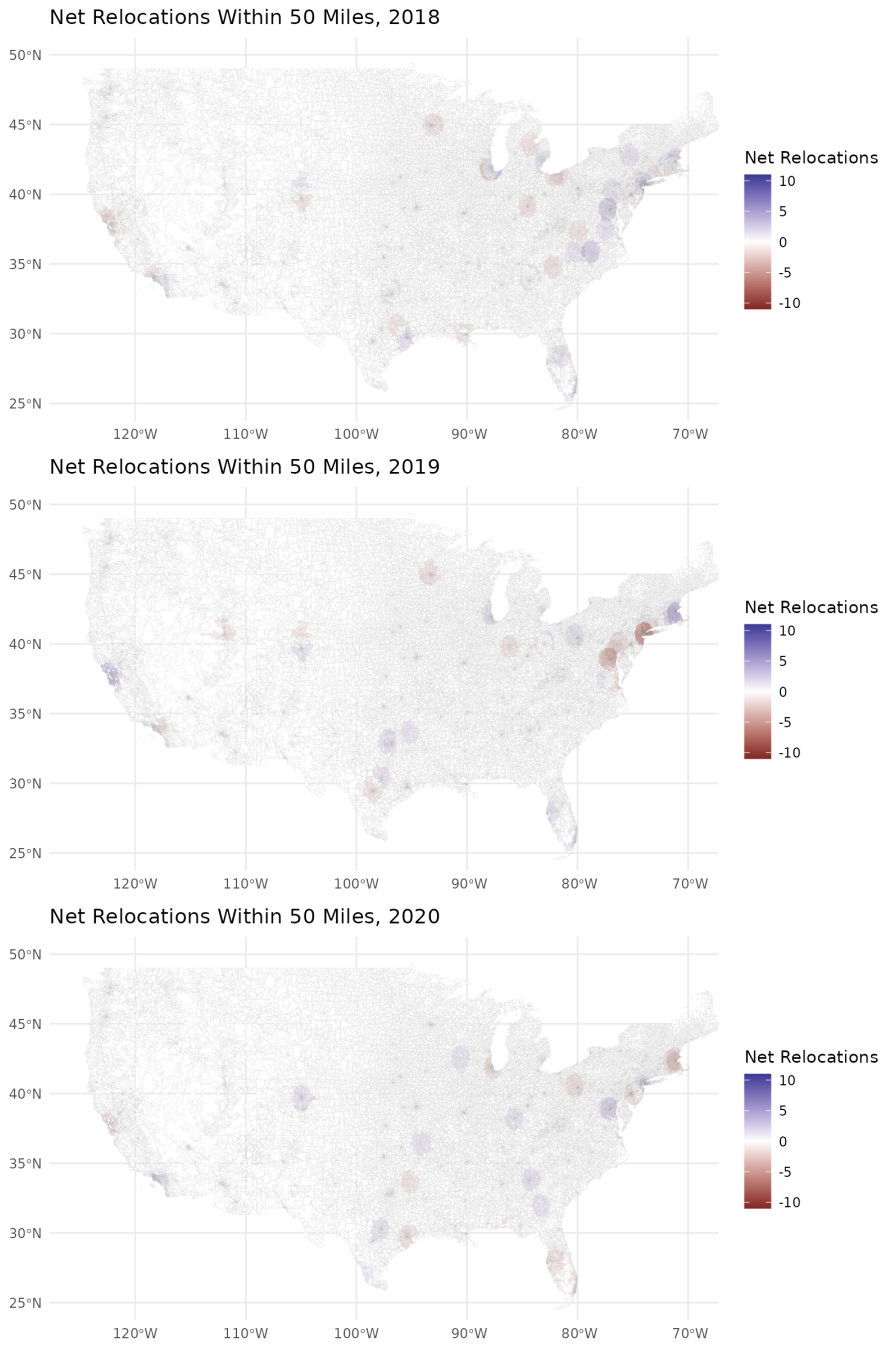

Figure 6c: Net Relocations 2018-2020

Number of firms relocating within 50 miles of the area minus number of firms leaving the area.

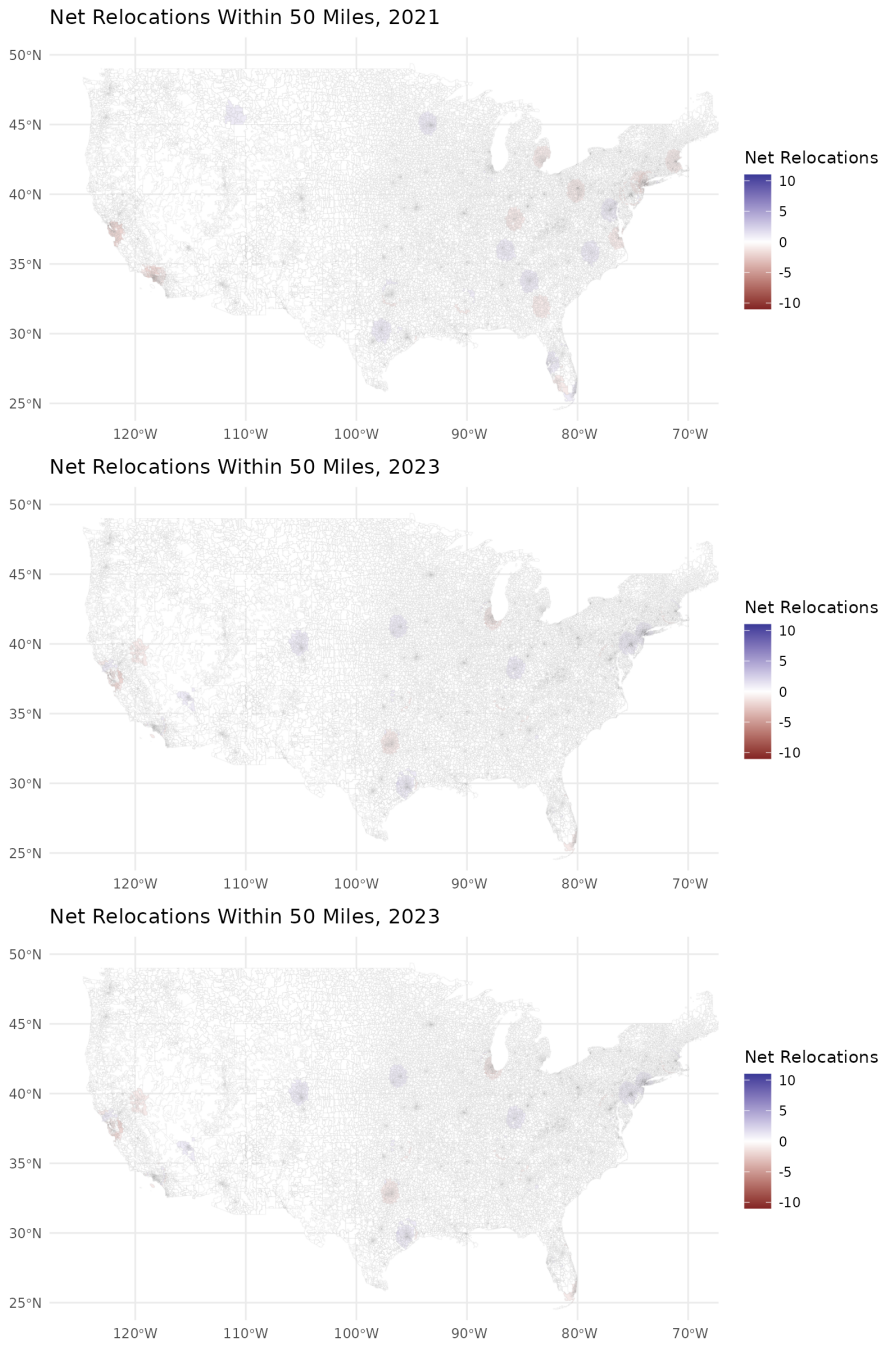

Figure 6d: Net Relocations 2021-2023

Number of firms relocating within 50 miles of the area minus number of firms leaving the area.

References

Akbas, F., Meschke, F., & Wintoki, M. B. (2016). Director networks and informed traders. *Journal of Accounting & Economics*, 62(1), 1–23.

Bakke, T.-E., Black, J. R., Mahmudi, H., & Linn, S. C. (2024). Director networks and firm value. *Journal of Corporate Finance (Amsterdam, Netherlands)*, 85, 102545-.

Beaman, L., & Magruder, J. (2012). Who Gets the Job Referral? Evidence from a Social Networks Experiment. *The American Economic Review*, 102(7), 3574–3593.

Becht, M., Bolton, P., & Röell, A. (2003). Corporate governance and control. In Handbook of the Economics of Finance (Vol. 1, pp. 1-109). Elsevier.

Brown, A. B., Dai, J., & Zur, E. (2019). Too busy or well-connected? Evidence from a shock to multiple directorships. *The Accounting Review*, 94(2), 83-104.

Cai, J., Garner, J. L., & Walkling, R. A. (2009). Electing directors. *The Journal of Finance*, 64(5), 2389-2421.

Cai, J., Garner, J. L., & Walkling, R. A. (2013). A paper tiger? An empirical analysis of majority voting. *Journal of Corporate Finance (Amsterdam, Netherlands)*, 21, 119–135.

Clifford, C. P. (2008). Value creation or destruction? Hedge funds as shareholder activists. *Journal of Corporate Finance (Amsterdam, Netherlands)*, 14(4), 323–336.

Chaisemartin, C. & D'Haultfœuille, X. (2023). Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey, *The Econometrics Journal*, Volume 26, Issue 3.

Fatemi, A., Fooladi, I., Sy, O., & Zaman, A. (2024). Corporate headquarters relocation and CSR performance. *International Review of Economics & Finance*, 89, 622–637.

Gee, L. K., Jones, J., & Burke, M. (2017). Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network. *Journal of Labor Economics*, 35(2), 485–518.

Gillan, S. L., Koch, A., & Starks, L. T. (2021). Firms and social responsibility: A review of ESG and CSR research in corporate finance. *Journal of Corporate Finance*, 66, 101889.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225, 254–277.

Jensen, M. C., & Meckling, W. H. (2019). Theory of the firm: Managerial behavior, agency costs and ownership structure. In Corporate governance (pp. 77-132). Gower.

Karpoff, J. M., Malatesta, P. H., & Walkling, R. A. (1996). Corporate governance and shareholder initiatives: Empirical evidence. *Journal of financial economics*, 42(3), 365-395.

Kemp, S. (2006). In the driver's seat or rubber stamp?: The role of the board in providing strategic guidance in Australian boardrooms. *Management Decision*, 44(1), 56–73.

Khwaja, A. I., & Mian, A. (2005). Do lenders favor politically connected firms? Rent provision in an emerging financial market. *The quarterly journal of economics*, 120(4), 1371-1411.

Khwaja, Asim Ijaz and Mian, Atif R. and Qamar, Abid, Bank Credit and Business Networks (February 17, 2011). HKS Working Paper No. RWP11-017

Khwaja, Asim, Atif Mian, and Abid Qamar. 2005. "Identifying Business Networks In Emerging Economies".

Li, W., Zhang, J. Z., & Ding, R. (2023). Impact of Directors' Network on Corporate Social Responsibility Disclosure: Evidence from China. *Journal of Business Ethics*, 183(2), 551–583.

Liu, Y., Ma, S., & Tang, X. (2022). Independent Director Networks and Executive Perquisite Consumption–"Collusion" or "Coordination" in Governance? *Emerging Markets Finance & Trade*, 58(13), 3824–3839.

McGregor, J. (2016, July 20). The part-time job that can pay more than $250,000 a year. The Washington Post. https://www.washingtonpost.com/news/on-leadership/wp/2016/07/20/pay-for-corporate-directors-creeps-higher-to-263500-a-year/

Mol, M. J. (2001). Creating wealth through working with others: Interorganizational relationships. *Academy of Management Perspectives*, 15(1), 150-152.

Renneboog, L., & Zhao, Y. (2011). Us knows us in the UK: On director networks and CEO compensation. *Journal of Corporate Finance (Amsterdam, Netherlands)*, 17(4), 1132–1157.

"West Monroe's Quarterly Executive Poll – Q1 2021". https://www.westmonroe.com/perspectives/report/quarterly-executive-poll-q1-2021.