Collusion through Common Leadership*

Alejandro Herrera-Caicedo
University of Wisconsin-Madison

Jessica Jeffers
HEC Paris
and CEPR

Elena Prager
Simon Business School
University of Rochester
and NBER

October 10, 2024

Abstract

This paper examines whether common leadership, defined as two firms sharing executives or board directors, contributes to collusion. Using an explicit measure of collusion from unsealed court evidence, we find that the probability of collusion between two firms increases by 12 percentage points after the onset of common leadership. These results are not driven by closeness of product or labor market competition, and are often driven by common leadership that involves an executive rather than an independent board director. Our findings are consistent with the increasing attention toward common leadership under Clayton Act Section 8. They also highlight the gap between stronger antitrust enforcement tools for collusion in product markets, and weaker tools in input markets.

*Jeffers gratefully acknowledges funding from the Rothschild & Co Data and Impact Investment Chair. Prager gratefully acknowledges funding from the Washington Center for Equitable Growth. Views expressed in this paper are those of the authors and do not necessarily represent the views of these funders. All errors are our own. We thank seminar and conference audiences at Dartmouth, DOJ, EIEF, ERIM, Imperial College London, Harvard Law School, IWH, Mannheim, Northwestern, Rochester, Texas A&M, UBC, and UCLA. We also thank Peter Haslag, Eric Posner, Andrew Sweeting, and Alminas Zaldokas for helpful feedback and discussions. Dev Burman, Josephine Freis, Chengdi Hu, and Yijin Wang provided excellent research assistance.

One third of public companies share a high-level leader with another firm, disproportionately firms in the same industry. In principle, antitrust law prohibits competing firms from sharing leaders because of concerns that leadership overlap facilitates collusion. This prohibition is implemented by Clayton Act Section 8, which bars two firms with substantial competitive overlap from sharing any members of their boards of directors and, since 1990, any officers of the corporation. Yet until very recently, enforcement of this prohibition was dormant. In 2022, the Department of Justice (DOJ) signaled a change in policy, stating that it had historically underused Section 8 and was "ramping up efforts to identify violations" (DOJ 2022).

Despite the renewed scrutiny, there is little evidence on whether common leadership increases collusion in practice.[1] On one hand, common leaders may provide firms with the means to communicate, align incentives, or enforce collusive agreements. On the other hand, firms have other means for colluding, and collusion has been documented in many settings where common leaders were apparently absent. Since Section 8 is a *per se* statute, the government can legally enforce the prohibition even if there is no evidence that common leadership is anticompetitive. If it is not anticompetitive, however, then enforcing a strict prohibition could unduly harm firms by restricting the pool of qualified leadership candidates or dampening productivity-enhancing information flows.

This paper therefore studies whether common leadership contributes to collusion. Using an explicit measure of collusion from unsealed court evidence, we find that entry into collusive agreements is indeed more likely among pairs of firms that share common leaders. Our empirical context is the largest known case of modern US labor market collusion. Starting in the early 2000s, several dozen Silicon Valley firms entered into agreements not to recruit one another's employees. A series of lawsuits ultimately revealed detailed evidence about these illegal "no-poaching" agreements. The unsealed information includes email trails and human resources policy documents that explicitly list partner firms with agreements not to recruit.

We take advantage of the unusual richness of the unsealed court evidence to construct a direct measure of collusion at the firm pair level. This direct measure means we do not need to indirectly infer collusion from observed prices, quantities, or auction behavior. Moreover, by constructing a sample consisting of all pairwise permutations of known colluders, we minimize the rate of false negatives. Discussions with enforcers indicate that if two firms were each identified as colluding with

[1] One exception is Gopalan et al. (2024), whose appendix documents a correlation between cartel convictions and degrees of separation of boards.

others, then any collusion between the two firms themselves was also likely to be uncovered. Our sample construction therefore reduces the rate of false negatives (true collusive agreements that are missed by the econometrician) compared to more typical contexts. We also corroborate our main findings using just the firm pairs that colluded, a sample definition that eliminates false negatives by construction. Our sample of colluders includes 45 firms with an average of 2.7 agreements per firm.

To measure common leadership, we use data on professional histories from BoardEx, which contains detailed information on senior personnel at publicly traded companies and significant private entities. We focus on individuals in leadership roles, such as board directors and senior executives. For each firm pair in a given year, we assess whether any individual held leadership positions at both firms simultaneously. For instance, in 2005, Arthur Levinson served as CEO of Genentech and board director for Apple and Google. We consider Genentech, Apple, and Google connected through common leadership at that time. Genentech, Apple, and Google all entered into no-poaching agreements with each other.

In the sample of known colluders from the Silicon Valley case, 62 percent of firms share common leaders with at least one other firm in the sample during our main sample period of 2000–2009. Pairs of firms that never share leaders have a collusion rate of 5 percent; pairs of firms that share leaders have a collusion rate of 35 percent. To contextualize the incidence of common leadership, we repeat our calculations for all publicly-held US companies. During the same period, 38 percent of public firms share a leader with at least one other public firm. Fourteen percent share a leader with at least one other firm in the same state of headquarters, and 8 percent with at least one other firm in the same 3-digit NAICS industry.

Using difference-in-differences regressions estimated on the Silicon Valley no-poaching case sample, we find that the probability of a collusive agreement between a pair of firms rises by 12 percentage points following the onset of common leadership. In event studies, we find that the effect is largest within one to three years after the onset of common leadership. Our preferred specification includes fixed effects at the firm pair level, so that the coefficient on common leaders is identified from changes within a pair of firms over time. This specification minimizes bias from collusion and common leaders both being caused by an omitted variable at the firm pair level, such as close competition between the two firms.

Nevertheless, our estimates will not have a causal interpretation if firms appoint common leaders for the purpose of facilitating previously planned collusion. While our data do not allow us to rule out such reverse causality, we present several arguments against it. If the collusive intent predated the appointment of a common leader, we would expect some run-up of collusive agreements, or at least faster implementation of collusion. In addition, subpoenaed documents explicitly invoke common leaders as reasons for collusion (Figure 1b). Some of these colluding pairs have no meaningful competitive overlap in the labor or product market, making it less likely that common leaders were installed for the purpose of implementing planned collusion. Lastly, court documents suggest that the incentives for collusion also existed among other firm pairs, but entry into collusion was more successful among pairs with common leaders.

Common leadership is a stronger predictor for collusion in our sample than other measures of firm overlap. Firms that compete more closely or share the same owners may have greater incentives to collude. We therefore check whether collusion is more likely among firm pairs with greater product or labor market overlap, or greater common ownership. We estimate a precise zero for the relationship between common ownership and collusion. We propose a novel measure of labor market overlap based on vector similarity of other past and future employers of a firm's employees. This measure predicts collusion, but the implied relationship is half the magnitude of the relationship between common leadership and collusion.[2] Moreover, the coefficient on common leadership is virtually unchanged when we control for labor market overlap. We also test the predictive power of product market overlap using standard definitions of industry (NAICS codes) and product similarity (Hoberg and Phillips 2016). We find no incremental predictive power of product market overlap. This suggests that Clayton Act Section 8, which only prohibits common leadership among competitors in the product market, may not be suited to enforcing against collusion in input markets.

Economy-wide effects may be smaller than the ones we estimate within our sample. Our estimates are specific to a particular set of firms and time in which the conditions may have been especially conducive to collusion, as the explosive growth of the tech industry in Silicon Valley in the mid- to late 2000s created fierce competition for the best talent. We define our sample to

[2] A one-standard deviation increase in labor market overlap implies a one percentage point increase in the probability of collusion, while a one-standard deviation increase in common leadership implies a two percentage point increase in the probability of collusion.

comprise firms involved in at least one no-poaching agreement, resulting in a sample selected on revealed willingness to collude. Consequently, we interpret our results as directionally supporting closer scrutiny of common leadership.

This paper provides the “first stage” for a recent body of evidence documenting apparent anticompetitive outcomes among companies with board overlap, including in financing (Barone et al. 2022; Eldar et al. 2023), product markets (Gopalan et al. 2024; Geng et al. 2021), innovation (Cabezon and Hoberg 2024), and, closest to our setting, labor markets (Begley et al. 2023). Our findings provide the first link in a potential causal chain from common leadership to collusion to observed anticompetitive outcomes. We also depart from prior work by including firm executives in our definition of common leadership. We find suggestive evidence that common leadership involving executives, rather than only independent board members, is at least equally predictive of collusion.

The rationale for prohibiting common leaders is that they may provide firms with the means to communicate, align incentives, or enforce collusive agreements. These mechanisms are qualitatively similar to common ownership—the same investors owning shares in two ostensibly competing firms—which has been found by some researchers to lead to coordination among the co-owned firms (Azar et al. 2018; Boller and Morton 2020; Anton et al. 2023) and has attracted substantial attention from regulators (FTC 2018; DOJ and FTC 2023; FTC and DOJ 2024). Although common leaders provide an even more direct link between firms than common owners, common leadership has received comparatively little regulatory attention.

Turning to our empirical context, two other recent papers examine the Silicon Valley no-poaching case, focusing on the eight firms targeted by the DOJ's 2010 case and the presumed end of collusion at that time. Gibson (2024) finds evidence of lower salaries and other benefits prior to the DOJ's lawsuit, and Ferrés et al. (2024) find that colluding firms posted jobs at higher rates, with lower benefits, and spent less on R&D while producing more innovation. This paper focuses on the origins of the collusion and expands the sample to the dozens of other firms implicated in collusive agreements across the multiple lawsuits that followed the DOJ's investigation.

The remainder of the paper proceeds as follows. Section 1 describes our empirical setting and data, and summarizes the prevalence of common leadership. Section 2 describes our sample construction and regression specification. Section 3 presents the results.

1 Data

This section describes the collusive agreements revealed by the DOJ's investigation and subsequent series of lawsuits in the Silicon Valley no-poaching case. We then describe our measure of common leadership and summarize the prevalence of common leadership both in the set of firms involved in the Silicon Valley case and in the set of all US public firms. For the set of firms involved in the case, we also summarize the prevalence of collusive agreements. Finally, we introduce a novel measure of the strength of labor market competition that enters into the regressions in Section 2.

1.1 Measure of Collusive Agreements

The Silicon Valley no-poaching case provides an unusual level of transparency into collusive agreements among a large group of firms. Modern collusion is typically difficult to measure due to its illegality and the secretive nature of such behavior. In typical research contexts, the presence of collusion must be inferred from a combination of data and models of conduct, such as when bidding rings are inferred from observed bids in a procurement auction. In our context, evidence emerging from multiple court cases provides a direct measure of which pairs of firms had no-poaching agreements. This obviates the need to impose assumptions on conduct or firm objective functions before inferring collusion.

We measure collusion using evidence unsealed from the court case the DOJ brought in 2010, as well as several civil lawsuits that followed. Evidence entered into court exhibits contains lists of firm pairs with agreements, along with bounds on the timing of the agreements. Figure 1 shows excerpts from two such court exhibits, which were typically drawn from subpoenaed emails or internal company documents. Figure 1a contains a list of Google's earliest five agreements, and Figure 1b lists Apple's active agreements a few months before the DOJ investigation. Notably, the Apple document explains that the agreements with firms in seemingly unrelated industries, such as Nike and J. Crew, involve common leadership. We use all available court evidence from both the DOJ suit and the multiple follow-on civil suits to construct a dataset of no-poaching agreements at the firm pair-year level. We provide summary statistics for the resulting data later in this section, after describing the institutional details of the Silicon Valley no-poaching case.

In the tech sector in the 2000s, talent was scarce (Helft 2007) and cold-calling other firms' current employees was a common recruiting tactic. Employees hired via cold-calling were perceived as higher-quality than those who actively applied for a job.[3] When a rival firm tried to poach an employee, the poacher and the incumbent employer could enter a bidding war, which could result in losing an employee or paying higher compensation. A mutual agreement not to cold-call the employees of a close competitor could therefore benefit both firms by reducing labor costs. The agreements we study all contained such a no-cold-calling provision, commonly called a "no-poaching" agreement. The agreements typically applied to all the employees of the rival.[4] Even so, they still allowed each firm to hire the other's employees if those employees independently applied for a job. Some agreements went further with a provision that the firms would refuse to engage in a bidding war even if an employee actively applied to the rival's job. However, the details of individual agreements are not always observable from the court documents. We provide summary statistics on these collusive agreements in Section 1.3.

The agreements in the case began in earnest with a wave of no-poaching arrangements in 2005. Figure A.1 shows an email exchange between Google co-founder Sergey Brin and his executive team in February 2005. Apple CEO Steve Jobs approached Brin with a request to stop recruiting Apple employees (Figure A.1a). Within two weeks, Apple and Google had an agreement not to cold-call one another's employees (Figure A.1b). After the Apple-Google agreement, the list of no-poaching agreements exploded. These agreements are illegal under antitrust law because they amount to market allocation; that is, firms agreeing not to compete aggressively for certain subsets of the market, in this case the market for labor. However, antitrust law had rarely been enforced in labor markets in the decades before 2005, and company leaders appear to have concluded that the legal risk from these agreements was tolerable.

1.2 Measure of Common Leadership

We build a measure of common leadership using data from BoardEx on firms' leadership structures and their leaders' careers. BoardEx is a data analytics company that systematically collects information about the affiliations and professional histories of firms' leaders and repackages the information for sale to business clients. The BoardEx data include the professional histories of

[3] For example, Google VP Jonathan Rosenberg testified, "We don't want the people who are applying, we want the best in the world [...] the people who apply, on average, aren't as good." This view is a recurring theme in testimony and other court evidence.

[4] Some of the court evidence suggests that the agreements were enforced more aggressively for employees in senior or technical roles, but enforcement was not limited to those roles. For example, Apple's recruiters questioned whether they could hire a chef who was applying for a food service job while working as a sous-chef at a Google cafeteria.

Figure 1: No-Poaching Lists from Court Exhibits

(a) Google's October 2005 List

From: Shona Brown [mailto:shona@google.com]
Sent: Tuesday, October 04, 2005 7:58 PM
To: EMG
Cc: Judy Gilbert; Stacy Sullivan; Arnnon Geshuri
Subject: Protocol for "Do Not Cold Call" and "Sensitive" Companies ---please comment to Arnnon ASAP if you have any changes

[OMID -- check with Arnnon before you go "live" with this] [JOAN - please confirm with Arnnon that you are comfortable with this final wording]

Special Protocol for "Do Not Cold Call" and "Sensitive" Companies

The following companies are part of the "Do Not Cold Call" list:

* Genentech
* Intel
* Apple
* Paypal
* Comcast

(b) Apple's July 2009 List

Hands Off (Do Not Call List) :

Microsoft – Mountain View (exchange group and Mac group)
Garmin
Palm
Adobe (Software partner)
Aspyr
AMD/ATI
Best Buy
CDW
Cingular
Comp USA (product re-seller)
Foxconn
Genentech (CEO sits on our board)
Google
Ingram Micro
Intel
Intuit (Common board members)
JCrew (Common board members)
Mac Zone
Nike (Common board members)
Nvidia
PC Connection
PC Mall
Pixar
Lucas
Quanta
Tech Data
Zones

Each figure is drawn from an exhibit presented as evidence in court. Exhibits such as these form the basis of our data on the pairwise existence and timing of collusive agreements.

nearly all individuals who ever served on a board or in a moderately senior employee role at any publicly traded firm, or at a large or otherwise significant private entity. As a result, the data also include information on the many additional entities with which these individuals were ever affiliated, covering more than 2 million entities. The professional histories record the name and identifier of the entity, the nature of the individual's affiliation with it (including job title), and the dates of the affiliation.

We use the BoardEx data to construct measures of common leadership between pairs of firms in 2000–2019. For each firm pair in each year, we measure whether any of the high-level leaders at one firm simultaneously served as a high-level leader at the other, and if so, how many.

Definition of a leader. Following Clayton Act Section 8, we define a high-level leader as either a member of the board of directors or a senior executive. Board members are straightforward to identify in the data using a BoardEx-provided indicator. However, the data do not distinguish between senior executives and more junior employees such as regional directors. Perhaps for this reason, the literature on common leadership almost exclusively focuses on shared board members, also known as interlocking directorates (Azar 2022; Barone et al. 2022; Ghezzi and Picciau 2022;

Lemley et al. 2022; Begley et al. 2023). We use the term "common leadership" in part to distinguish from definitions that exclude senior executives.

In order to include senior executives in our definition, we classify employees' roles within each firm as senior or other. In the senior category are C-suite executives and, for firms with particularly small C-suites, roles at or above the level of a senior vice president. By allowing the mapping from job title to senior executive to depend on the structure of a firm's leadership team, we account for the fact that some firms use job titles like "vice president" to refer to numerous upper middle managers, whereas other firms reserve such titles for a handful of truly senior employees who are central to the firm's strategic decision-making.[5] Our classification results in a mean of 9.6 senior leaders per firm-year among all US public firms during our primary sample period.

Samples. We create two separate samples of firms. The first sample, which we call the no-poaching court case sample, consists of all firms implicated in at least one no-poaching agreement according to the unsealed documents from any of the Silicon Valley no-poaching court cases. This sample forms the basis of our regressions estimating the relationship between common leadership and collusion, described in Section 2. The second sample, which we call the public firms sample, consists of all publicly traded firms in the US. We use the public firms sample to illustrate the prevalence of common leadership across a broad set of firms.

Our primary sample period is 2000–2009. Although our common leadership data cover 2000–2019, the DOJ investigation was announced in 2009, and the evidence released from the court cases is backward-looking. It does not contain systematic information on agreements after 2009. In addition, even if firms entered into new collusive agreements after 2009, they may have become more careful not to leave paper trails. We begin the sample in 2000 because BoardEx substantially expanded its coverage of firms that year, and because the court cases did not involve investigations of collusion in prior decades.

Additional firm characteristics. We merge in firm characteristics from Compustat for publicly traded firms. Compustat provides information on public firms' financials, including revenues,

[5]For example, the economic consulting company Analysis Group has hundreds of employees and affiliates with the title "Vice President", and they are separated from the C-suite by at least one additional layer of hierarchy. By contrast, Apple has fewer than a dozen "Senior Vice Presidents," and they sit immediately below the CEO in the corporate hierarchy.

profitability, and market capitalization; and some competitive characteristics, including primary industry NAICS code. We also merge in a commonly used proxy for the closeness of product market competition between two firms developed by Hoberg and Phillips (2016). Hoberg-Phillips similarity measures the similarity of text descriptions of publicly traded firms' product portfolios from their 10K reports. We merge in common ownership profit weights from Backus et al. (2021) to measure overlap in firm ownership. For firms implicated in the no-poaching case, we also construct a novel measure of the closeness of labor market competition, described below in Section 1.4.

Aggregation to firm level. We aggregate entities from our various data sources to the level of the ultimate owner of the firm. For the public firms sample, we consider Compustat's Global Company Key to represent the appropriate level of aggregation. In the no-poaching court case sample, some firms are privately owned and therefore do not appear in Compustat. For this sample, we manually find all entities owned by the key owner of the firm, excluding ownership "in name only" via financial holding companies, and aggregate them into a single firm.

1.3 Prevalence of Common Leadership

We first provide summary statistics for common leadership alongside statistics on collusive agrements in the court case sample, then provide context for the prevalence of common leadership using the sample of all US public firms.

No-poaching court case sample. Table 1 summarizes key statistics for the court case sample. After aggregating related firms and merging with BoardEx, the sample consists of 45 firms. By construction, all of these firms eventually have an agreement with at least one other firm in the sample. Not by construction, 62.2 percent of them share a common leader with at least one other firm in the sample.

For each firm, we identify an average of 12.9 leaders, including 9.0 independent board directors and 3.9 executives. Note the distinction between two types of leaders: independent directors, i.e. board members who hold no employment relationship with the firm; and leaders who hold an executive position in the firm, with or without also holding a seat on the board. As we will discuss in Section 3.4, executives may have greater ability and incentives to collude than independent board directors. Summing both types of leaders, an average of 7.1 individuals per firm serve in leadership

Table 1: Summary Statistics for No-Poaching Court Case Sample

Panel A. Firm-level summary statistics	
Firms	45
Firm-years	408
Fraction of firms with...	
At least one common leader shared with another firm in the sample	0.622
At least one agreement with another firm in the sample (by construction)	1.000
Mean number of leaders per firm-year	12.912
Independent directors	9.044
Executives and interested directors	3.868
Mean number of common leaders per firm-year	7.051
Mean number of known collusive agreements per firm	2.667
Panel B. Firm pair-level summary statistics	
Firm pairs	980
Firm pair-years	8,123
Fraction of pairs ever having common leadership	0.050
Fraction ever having an agreement among pairs...	
Ever having common leadership	0.347
Never having common leadership	0.046

Sample is all possible pairs consisting of firms $i \neq j$ where each firm is implicated in at least one no-poaching agreement in the Silicon Valley case during 2000–2009. Agreements drawn from court evidence as described in Section 1.1. Common leadership calculated from BoardEx as described in Section 1.2. Firm pairs are undirected: i, j and j, i are counted as one pair.

roles at more than one firm in the sample. Firm pairs very rarely share more than one common leader at a time.

After accounting for periods of time overlap in BoardEx, there are 980 possible pairwise combinations of firms in the court case sample. Of these pairs, 5.0 percent are connected by common leadership at some point during 2000–2009. Collusive agreements occur at a rate of 34.7 percent among pairs that ever had common leaders during this period, compared to a rate of only 4.6 percent among pairs without common leaders. Each firm is involved in an average of 2.7 collusive agreements.

Public company sample. Table 2 summarizes key statistics for the public firm sample, separately for our primary 2000–2009 sample period and a more recent 2010–2019 period for comparison. There are 4,340 firms in the 2000–2009 decade and 5,120 firms in the 2010–2019 decade. In the first decade, 37.5 percent of firms share a common leader with at least one other firm in the public firm sample, compared to a common leadership rate of 62.2 percent among the court case sample in the same decade. (Note that if common leadership was randomly distributed, the larger set in the public firm sample would mechanically increase the chances of overlapping with at least one other firm.) In the most recent decade, 35.0 percent of firms share a common leader with at least one other firm.

Firm pairs in the same industry or location are disproportionately likely to share common leaders. In both decades, approximately 8 percent of firms (unconditionally) share leadership with a firm in the same 3-digit NAICS industry and 13 to 14 percent of firms share leadership with a firm headquartered in the same state. In 2000–2009, we identify an average of 9.6 leaders per firm-year, including 7.4 independent directors and 2.2 executives. Of these leaders, 3.7 on average serve in a leadership role at more than one firm in the sample.

There are 9.3 million possible pairwise combinations of firms in the first decade and 11.5 million combinations in the second decade. Of these pairs, 0.03 percent are connected by common leadership. The number is three times as high for firms in the same 3-digit NAICS industry, and four times as high for firms headquartered in the same state.

Table 2: Summary Statistics for Public Firms Sample

	(1) 2000–2009	(2) 2010–2019
Panel A. Firm-level summary statistics		
Number of firm-years	36,830	34,263
Number of unique firms	4,340	5,120
Fraction of firm-years with		
At least one leader shared with another firm	0.375	0.350
At least one leader shared with firm in the same 3-digit NAICS	0.084	0.080
At least one leader shared with firm headquartered in the same state	0.144	0.132
Mean number of total leaders	9.583	7.844
Independent directors	7.404	6.035
Executives and interested directors	2.179	1.809
Mean number of common leaders	3.722	2.705
Panel B. Firm pair-level summary statistics		
Number of firm pair-years	68,053,457	58,720,711
Number of unique firm pairs	9,279,958	11,524,911
Fraction of firm pair-years with common leaders among		
All firm pairs	0.0003	0.0003
Firm pairs in the same 3-digit NAICS	0.0009	0.0009
Firm pairs headquartered in the same state	0.0013	0.0012

Sample is all publicly traded firms in the US. Common leadership calculated from BoardEx as described in Section 1.2. Firm pairs are undirected: i, j and j, i are counted as one pair.

1.4 Measure of Labor Market Competition

We construct a measure of how closely two firms compete for labor using LinkedIn data. As with product market collusion, the incentive to collude in the labor market is greater when two firms compete more closely. If firms with a preexisting incentive to collude are also more likely to appoint common leaders, then the coefficient on common leadership will be biased. We therefore control for how closely a firm pair competes in the labor market.

We propose a novel measure of the closeness of labor market competition: the similarity between the worker pools from which two firms draw their workers. Our goal is to measure whether firms compete for the same set of workers, and therefore have an incentive to allocate the market for workers by colluding. Our measure checks for high overlap between the sets of *other* firms their workers have worked at in the past or will work at in the future. For example, if many of both Google's and Microsoft's workers subsequently leave for Facebook, then Google and Microsoft will be coded as similar. Formally, we take the cosine distance between firm i's and firm j's worker histories:

$$\text{overlap}_{ij} = 1 - \frac{\sum_{k \notin \{i,j\}} s_{i \to k} s_{j \to k} + s_{i \leftarrow k} s_{j \leftarrow k}}{\|\vec{s}_i\| \, \|\vec{s}_j\|}$$

With some abuse of notation, $s_{i \to k}$ is the share of i's workers subsequently working at k, and $s_{i \leftarrow k}$ is the share of i's workers who formerly worked at k.

We compute this measure using data on worker histories in 2000–2010 from LinkedIn. Appendix B.1 describes the data and the labor market overlap measure in more detail. As we will show below, this measure better explains collusive agreements than standard product market-focused measures of competitive overlap such as NAICS industry codes or Hoberg-Phillips product description text similarity. For example, our measure implies that Apple's two closest competitors for labor are Intuit and Genentech. Despite Genentech being in an entirely different industry (biotech), Apple had agreements with both firms.

In addition to its relevance for explaining labor market collusion, our measure also has the advantage of being computable using data and variation typically available to researchers, including for non-publicly traded companies. While a cross-wage elasticity of labor supply may be more theoretically appealing, we do not have access to the deidentified firm wage data and natural experiment that would be required to estimate a model of labor supply.

2 Empirical Strategy

We use the court case sample to answer our primary research question, whether common leadership predicts collusion. Section 1.3 documents a strong relationship between common leadership and collusion in the raw data. In this section, we use difference-in-differences and event study regressions to study whether the timing of agreements is consistent with common leaders causing agreement. To the extent possible, we control for other causes of entry into collusive agreements, such as closeness of competition.

2.1 Regression Sample

Our regression sample is designed to minimize the risk of measurement error in the outcome variable: the presence of a collusive agreement between a pair of firms.[6] The regression sample consists of all possible pairwise combinations of firms in the court case sample. That is, each pair consists of two firms that were each implicated in at least one no-poaching agreement during the DOJ's investigation or follow-on lawsuits. The rationale is that if each firm in the pair was found to have any collusive agreements, then the existence of a collusive agreement between the two firms was also likely to be detected. Discussions with enforcers suggest that such detection was indeed likely, especially considering the multiple civil lawsuits that came after the DOJ's investigation and revealed additional evidence.

Of course, it is possible that some agreements between firms in the court cases did not come to light despite the investigations, leading to false negatives. In most contexts, however, the risk of false negatives is greater still: the researcher may have to assume that observations lacking information on collusion—such as pairs consisting of firms not investigated in these court cases—have no collusion. To further reduce concerns about measurement error, we also show estimates subsetting to firm pairs that are known to have colluded. By construction, this subsample cannot have false negatives in the agreement variable.

An implication of this sample construction is that we estimate the relationship between common leadership and collusion in a selected market using selected firms. The competition for talent in Silicon Valley in the 2000s was particularly acute. The agreements were entered into during a period of low antitrust enforcement in labor markets. The firms in the sample, which are all known

[6] The key regressor of interest, common leadership, is measured with minimal error as described in Section 1.2.

colluders, may be less risk-averse than the average firm. We therefore interpret our estimates as directionally correct but potentially too large in magnitude for the economy as a whole. This limitation of external validity is a cost of constructing our sample for the greatest possible internal validity.

2.2 Regression Specification

We estimate difference-in-differences and event study regressions at the level of a firm pair-year. We regress an indicator for the presence of a collusive agreement on time since the onset of common leadership (if any). For a pair of firms i and j in year t, we regress

$$\text{Agreement}_{ijt} = \beta\,\text{CommonLeader}_{ijt} + \gamma\,X_{ij} + \zeta_{it} + \eta_{jt} + \varepsilon_{ijt} \tag{1}$$

where $\text{CommonLeader}_{ijt}$ is an indicator that is equal to one if i and j have experienced the onset of common leadership.[7]

The level of observation is a directed firm pair in a given year. *Directed* means that we include separate observations for pair i, j and pair j, i. This allows us to include firm-by-year fixed effects for both firms, ζ_{it} and η_{jt}, without arbitrarily assigning which firm is i and which is j. The ability to include firm-by-year fixed effects, which would ordinarily absorb treatment status, is due to the pairwise structure of our data. Firm-by-year fixed effects allow us to control for changes within a firm that may drive common leadership and collusion across the pairs involving this firm, such as entry into new product markets or the arrival of an aggressive CEO such as Steve Jobs.

Standard errors are three-way clustered at the levels of firm i, firm j, and the *undirected* pair. The clustering at the level of undirected firm pairs (treating i, j and j, i as a single cluster level) avoids double-counting the two observations from a single undirected pair when performing inference.

In all but our most parsimonious specifications, we include measures of the closeness of product market and labor market competition at the firm pair level in X_{ij}. If close competition drives firm pairs to collude and to appoint common leaders, then omitting it from the regression will bias the main coefficient of interest on common leadership. In our preferred specification, observables in

[7] There are very few cases of firm pairs sharing more than one common leader at a time, limiting our ability to estimate the marginal effect of additional common leaders. There are also few cases of common leadership links dissolving during our relatively short sample period.

X_{ij} are replaced with firm pair fixed effects.

Staggered treatment timing. In the presence of staggered treatment timing and heterogeneous or dynamic treatment effects, the two-way fixed effects estimator for Equation 1 is biased. We therefore use the Sun and Abraham (2021) estimator to aggregate cohort-specific estimates into a single average treatment on the treated (ATT) coefficient estimate. We also show event study analogs to Equation 1. The Sun and Abraham (2021) estimator is particularly useful in our setting because, unlike some alternative estimators, it allows for treatment to affect outcomes with a lag. It is also less demanding of the data than de Chaisemartin and D'Haultfoeuille (2020), which is useful in our small, short sample.[8] We discuss other threats to identification, such as reverse causality, in Section 3.2.

ATT vs. ATE. Our sample yields an estimate of the average treatment effect on the treated rather than an average treatment effect (ATE), as is typical in difference-in-differences designs without a randomized treatment. The ATT measures the effect of common leadership on collusion among firm pairs that select into common leadership rather than for a randomly selected firm pair. For enforcement purposes, the ATT is the policy parameter of interest. Enforcers need to understand the efficacy of enforcing against precisely those firms that choose to appoint common leaders.

3 Results

This section first presents the results of our main regressions studying collusive agreements in our no-poaching court case sample. We describe the main difference-in-differences results and the corresponding event studies in Section 3.1, and discuss threats to interpretation in Section 3.2. Section 3.3 addresses concerns about measurement error. Section 3.4 presents exploratory analyses of the importance of a common leader being an executive rather than merely a board member.

[8]In principle, a sensible alternative is Callaway and Sant'Anna (2021), which more explicitly allows for the inclusion of covariates than Sun and Abraham (2021). However, their estimator is not suited to settings where the level of the unit fixed effects does not correspond to the level of observation. Our preferred specification uses firm pair fixed effects rather than other observables, making it explicitly suitable for the Sun and Abraham (2021) estimator.

3.1 Main Results

Table 3 reports estimates of Equation 1. The first row reports the coefficient of interest on common leadership. In specifications with firm-by-year fixed effects, we find that the probability of a collusive agreement rises by 11 to 12 percentage points after the onset of common leadership, from a sample mean of 1.6 percent.

The magnitude of the coefficient estimate is stable as we add controls, up to and including fixed effects at the firm pair level. The coefficient is slightly larger in Column 1, which uses separate firm fixed effects and year fixed effects, and therefore does not control for time-varying factors at the firm level. Starting in Column 2, we add firm-by-year fixed effects and then successively more detailed controls. Columns 3 and 4 check whether collusion and common leadership are both explained by closeness of competition or other overlap. Column 3 uses three measures of overlap between a pair of firms. The first measures their common ownership using the weight one firm's objective function places on the other's profits (κ values), as implied by their common ownership structure and calculated by Backus et al. (2021). The second measures how closely they compete in the labor market according to our worker flows measure from Section 1.4. The third measures whether they are in the same industry, as categorized by BoardEx. Column 4 replaces the industry categorization from BoardEx with more conventional measures of product market overlap: same NAICS industry code and Hoberg-Phillips similarity.[9]

If collusion and the appointment of common leaders were both driven by closeness of competition or common ownership, then including these measures should reduce the magnitude of the coefficient on common leadership. Instead, the coefficient estimate is almost unchanged from Column 2 to Columns 3–4. Anecdotally, some of the agreements in the data are between pairs of firms that have negligible competitive overlap, but do share common leaders (e.g., Apple and Nike; see Figure 1). Closeness of labor market competition is, however, predictive of collusive no-poaching agreements. Common ownership and closeness of product market competition are not.

Column 5 presents our preferred specification, with a coefficient on common leadership of 0.119. It replaces observables at the firm-pair level with firm-pair fixed effects. If our measures of overlap in Columns 3–4 do not fully capture factors at the firm pair level that drive both common leadership and collusion, then the firm-pair fixed effects should mitigate those concerns. The coefficient

[9] These measures are only available for the subset of firm pairs consisting of two publicly traded firms.

on common leadership is estimated solely using variation over time within a pair of firms. The coefficient may still be biased if the factors driving collusion are changing simultaneously with the onset of common leadership. However, the change would need to happen quickly and would need to be after the onset of common leadership in order for this alternative explanation to be consistent with our event studies in Figure 2.

Table 3: Difference-in-Differences Regressions for Pr(Agreement): Main Sample

	(1)	(2)	(3)	(4)	(5)
Common leadership	0.135**	0.111**	0.108**	0.109**	0.119**
	(0.050)	(0.048)	(0.047)	(0.047)	(0.051)
Common ownership weight			0.004	0.004	
			(0.004)	(0.004)	
Missing common ownership			-0.008	-0.011	
			(0.006)	(0.010)	
Labor market overlap			0.067**	0.068**	
			(0.027)	(0.026)	
Missing lab. mkt. overlap			0.028**	0.029**	
			(0.013)	(0.012)	
Same BoardEx industry			0.003		
			(0.007)		
Same 4-digit NAICS				-0.033*	
				(0.017)	
Missing NAICS				0.009	
				(0.010)	
Hoberg-Phillips similarity				0.499	
				(0.513)	
Missing HP similarity				2.39×10^{-5}	
				(0.024)	
Separate firm FEs, year FEs	Yes				
Firm $\times$ year FEs		Yes	Yes	Yes	Yes
Firm pair FEs					Yes
Observations	16,245	16,245	16,245	16,245	16,245
R^2	0.18907	0.31205	0.31578	0.31709	0.57944
Within R^2	0.08457	0.07148	0.07652	0.07828	0.07624

Standard errors in parentheses, multi-way clustered on firm i, firm j, and firm pair $(i, j$ or $j, i)$. ATT summarized using Sun and Abraham (2021) to account for staggered treatment timing. Baseline mean of dependent variable in main sample = 0.016. Significance codes: *: 0.1; **: 0.05; ***: 0.01.

Figure 2 presents the event study version of our preferred specification from Table 3. The pre-trends leading up to the onset of common leadership are reasonably flat. Differential pre-trends would raise concerns about the required identifying assumption for the regression: that collusive agreements among the firm pairs with and without common leadership would have evolved similarly if not for the onset of common leadership. The probability of a collusive agreement increases slightly

Figure 2: Event Study Estimates: Agreements as a Function of Common Leadership

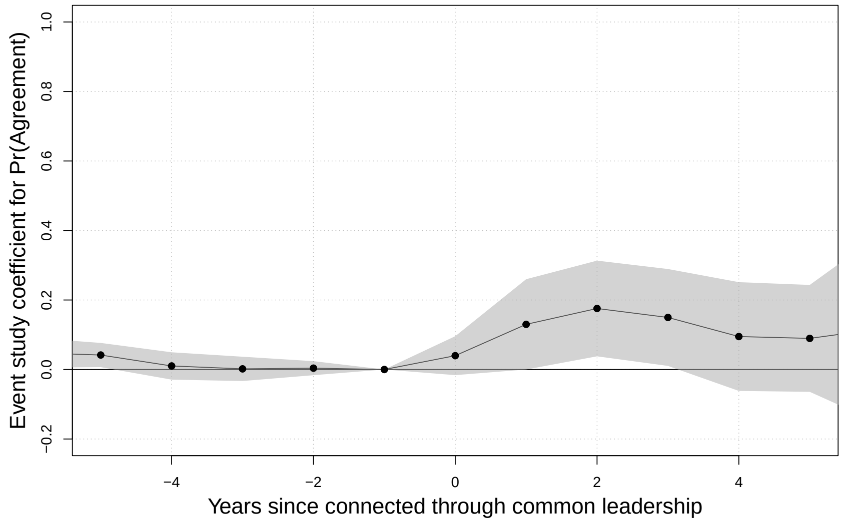

Figure plots event study coefficients from Equation 1 using the Sun and Abraham (2021) estimator to account for staggered treatment timing. Regression includes firm-by-year fixed effects and firm pair fixed effects. The sample is defined as described in Section 2.1.

in the year of onset (event time zero), but the peak is one to three years after event time. Figure A.2 plots the coefficients for more leads and lags around the event. It shows a large but imprecisely estimated increase in the point estimates in years seven and eight. We caution against taking those point estimates seriously because there are few treated firm pairs contributing data there due to the short sample period.

3.2 Interpretation

We interpret the results in Section 3.1 as evidence that common leadership raises the probability of collusion. It is also possible that firms appoint common leaders in order to collude or as a result of collusion. Both situations would affect the interpretation of our results and warrant discussion. In the first case, if common leaders are a mechanism for implementing collusion, then stopping common leadership should still reduce collusion by removing the tool that facilitates it. Firms would be unable to collude or be forced to resort to mechanisms that are, by revealed preference, more costly or less effective than common leadership. In the second case, if common leadership is itself a consequence of collusion, we would be concerned about reverse causality. Our setting lacks exogenous variation to rule out such reverse causality.

Nonetheless, several factors support interpreting the direction of causality as pointing from common leadership to collusion. First, if collusion was in process prior to the appointment of a common leader, we would expect to observe some preliminary steps towards collusive agreements before common leadership, or at least a rapid initiation of collusion following the onset of common leadership. Our event study estimates do not support this expectation. Second, court documents indicate that while incentives for collusion were present among various firm pairs, collusion was much more likely to be realized among those connected by common leaders. For example, Google attempted unsuccessfully to initiate collusion with Facebook, with which it did not share common leaders at the time.

Third, some firm pairs that engaged in collusion after the introduction of common leadership did not have significant competitive overlap in the labor or product markets, such as Apple and J. Crew. The negligible competition between these colluders makes it less plausible that the appointment of common leaders was a strategy for implementing intended collusion, because filling leadership positions is costly. Appointing one leader for implementing collusion carries the opportunity cost of appointing a different leader with potentially more valuable knowledge of the firm's sector. This cost is justified if the benefits from collusion are nontrivial, but when colluders are not competitors, we believe it is more likely that the collusion is a "collusion of convenience" resulting from the existing common leadership. This interpretation is consistent with the parenthetical notes explaining Apple's agreements in subpoenaed documents (Figure 1b): the document explicitly invokes common leaders to explain agreements with firms with little apparent competitive overlap.

Taking a step back from the relationship between common leadership and collusion, a welfare analysis of collusion itself is beyond the scope of this paper. Others have documented both costs (worker wages: Gibson (2024)) and benefits (innovation: Ferrés et al. (2024)) of this collusive episode, and we do not take a stance on the net effect.

However, if we take at face value the illegality of collusion, our results establish common leadership as a relevant and economically meaningful statistic for detecting illegal actions. We therefore interpret our results as supporting the increasing attention regulators are paying to common leadership under Section 8. Moreover, common leadership may be useful for detecting collusion in contexts not covered by Section 8, such as when firms compete in input markets rather than product markets.

3.3 Results with Alternate Sample Definitions

The unsealed evidence from the court documents varies in the degree of detail revealed about each agreement. For many firm pairs, the start date of the agreement is bounded within a range rather than pegged to a specific year. In the main results, we assign the start date for such pairs to the maximum of the bounds in order to avoid counting an agreement that does not yet exist. However, this results in some measurement error. For example, an agreement that actually began prior to the onset of common leadership could be miscoded as beginning after common leadership, inflating our coefficient estimate. Table 4 Column 2 therefore repeats the regression for just the subset of firm pairs where precise start dates are known. Pairs consisting of firms for which we only observe ranges are dropped. The estimated coefficient is slightly larger in magnitude than its full-sample analog (displayed in Column 1). This is the opposite of what would be expected if the potentially mismeasured agreements started prior to the onset of common leadership. Figure A.3 shows the corresponding event study estimates. There is a slight differential pre-trend that flattens out starting four years prior to event time.

As discussed in Section 2.1, our main regression sample is designed to minimize the risk of false negatives in the agreement variable. Despite the DOJ's investigation and additional evidence uncovered in follow-on litigation, there is a risk that we fail to detect some true agreements. If any such false negatives are disproportionately among firm pair-years that lack common leadership, then our main estimate will be biased upward. Table 4 Column 3 therefore repeats the regression using only the subsample of firm pairs that are known to reach eventual agreement. Within this subsample, false negatives are fully eliminated by construction. The resulting estimate in Column 3 is nearly triple the magnitude of our preferred specification. The larger magnitude is in part explained by this sample's omission of firm pairs that have common leadership but no observed agreement, pairs which mitigate the coefficient estimate in our preferred sample. Nevertheless, the estimate in Column 3 provides reassurance that our preferred specification finding is not driven by false negatives. Figure A.4 shows the corresponding event study estimates. The individual event time coefficients are noisy due to the small sample size.

Table 4: Difference-in-Differences Regressions for Pr(Agreement): Additional Samples

	Main Sample (1)	Certain Dates (2)	Agreeing Pairs (3)
Common leadership	0.119**	0.141**	0.430**
	(0.051)	(0.065)	(0.192)
Firm pair FEs	Yes	Yes	Yes
Separate firm FEs, year FEs			
Firm × year FEs	Yes	Yes	Yes
Observations	16,245	3,099	1,103
R^2	0.57944	0.71693	0.93155
Within R^2	0.07624	0.10333	0.19737

Standard errors in parentheses, multi-way clustered on firm i, firm j, and firm pair $(i, j$ or $j, i)$. ATT summarized using Sun and Abraham (2021) to account for staggered treatment timing. Baseline mean of dependent variable in main sample = 0.016. Significance codes: *: 0.1; **: 0.05; ***: 0.01.

3.4 Results on Source of Common Leadership

As discussed in Section 1.2, two firms may share a common leader through a board member or an executive. Executives participate more directly in the day-to-day operations of a firm than board members and typically have compensation tied to the performance of the firm.[10] It is therefore possible that a common leader who has an executive role in at least one of the firms in a pair would be better positioned or incentivized to collude.

Table A.1 tests this hypothesis and finds suggestive evidence in favor. Column 1 repeats our preferred specification with the pooled measure of common leadership. Columns 2–5 report separate estimates for common leadership through a shared independent board director (a board member who is not an executive at either firm in a pair), versus through a common leader who is an executive at one or both firms. The unusual structure of the table is due to the fact that, to our knowledge, no consistent estimator has yet been found for estimating treatment effects of two separate treatments in settings like ours.[11] We therefore use the Sun and Abraham (2021) estimator

[10] We distinguish between an independent director, i.e. a board member who holds no employment relationship with the firm; and an interested director, who is both a board member and an executive within the same firm. For example, CEOs often have a seat on the board. We pool interested directors and executives together, contrasting them as a group against independent directors.

[11] De Chaisemartin and D'Haultfoeuille (2023) propose a method for estimating the effects of multiple treatments in a staggered-treatment setting, but their method requires that the treatments always occur in the same order. In our setting, arrivals of board-only common leaders occur both before and after arrivals of executive common leaders. To our knowledge, no method has yet been developed for estimating difference-in-differences with multiple, non-ordered treatments and lagged treatment effects.

in separate regressions defining each of the two treatments as the primary treatment, in which we include the other treatment as a binary control variable.

In Table A.1, the coefficients on executive-involved common leadership are more often statistically significant than the coefficients on board-only common leadership. Among firm pairs that eventually reach an agreement, coefficient magnitudes for executive-involved common leadership are larger than for the board-only common leadership (Columns 4–5). In the full sample, however, the magnitudes are comparable (Columns 2–3). When board-only common leadership is the primary treatment, neither estimate is statistically significant, but when executive-involved common leadership is the primary treatment, both are. This discrepancy underscores the need for new estimators adapted to multiple non-sequential treatments.

Due to the econometric limitations just discussed, we view these results as merely suggestive. Taken at face value, they would imply that enforcement of Clayton Act Section 8 may be more efficient if it targeted common leadership involving executives in addition to shared board members. They also highlight the importance of expanding studies of common leadership beyond the shared board member definition used in the interlocking directorates literature.

4 Conclusion

In the largest modern case of US labor market collusion, collusion occurred disproportionately after firms began sharing common leaders. Collusive agreements typically began one to three years after the onset of common leadership, and the probability of collusion increased an average of 12 percentage points. This is a large effect, eight times the sample mean of 1.6 percent. Combined with the ease of detecting common leadership, this may make investigation of firm pairs with common leaders a useful tool for antitrust enforcers.

Our results are estimated in a context with a particularly dynamic labor market, but they dovetail with recent work documenting apparent anticompetitive outcomes for other types of firms with common board members. We provide evidence for one possible mechanism driving such outcomes, namely explicit, illegal collusion. Our results support stronger scrutiny of instances of common leadership, including where common leaders hold an executive position at one of the firms.

References

Anton, Miguel, Florian Ederer, Mireia Gine, and Martin Schmalz (2023) "Common Ownership, Competition, and Top Management Incentives," *Journal of Political Economy*, Vol. 131, No. 5, pp. 1294–1355.

Azar, Jose (2022) "Common Shareholders and Interlocking Directors: The Relation Between Two Corporate Networks," *Journal of Competition Law & Economics*, Vol. 18, No. 1, pp. 75–98.

Azar, Jose, Martin C. Schmalz, and Isabel Tecu (2018) "Anticompetitive Effects of Common Ownership," *The Journal of Finance*, Vol. 73, No. 4, pp. 1513–1565.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson (2021) "Common Ownership in America: 1980-2017," *American Economic Journal: Microeconomics*, Vol. 13, No. 3, pp. 273–308.

Barone, Guglielmo, Fabiano Schivardi, and Enrico Sette (2022) "Interlocking Directorates and Competition in Banking," *SSRN Electronic Journal*.

Begley, Taylor A., Peter H. Haslag, and Daniel Weagley (2023) "Directing the Labor Market: The Impact of Shared Board Members on Employee Flows," *SSRN Electronic Journal*.

Boller, Lysle and Fiona Scott Morton (2020) "Testing the Theory of Common Stock Ownership," *Working paper*.

Cabezon, Felipe and Gerard Hoberg (2024) "Leaky Director Networks and Innovation Herding,"Technical report.

Callaway, Brantly and Pedro H. C. Sant'Anna (2021) "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, Vol. 225, No. 2, pp. 200–230.

de Chaisemartin, Clement and Xavier D'Haultfoeuille (2020) "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, Vol. 110, No. 9, pp. 2964–2996.

De Chaisemartin, Clement and Xavier D'Haultfoeuille (2023) "Two-way fixed effects and differences-in-differences estimators with several treatments," *Journal of Econometrics*, Vol. 236, No. 2, p. 105480.

DOJ (2022) "Assistant Attorney General Jonathan Kanter Delivers Opening Remarks at 2022 Spring Enforcers Summit," April.

DOJ and FTC (2023) "2023 Merger Guidelines,"Technical report, U.S. Department of Justice and the Federal Trade Commission.

Eldar, Ofer, Yaron Nili, and James Pinnington (2023) "Common Ownership Directors,"Technical report, Working paper.

Ferrés, Daniel, Gaurav Kankanhalli, and Pradeep Muthukrishnan (2024) "Anti-Poaching Agreements, Corporate Hiring, and Innovation: Evidence from the Technology Industry,"Technical report.

FTC (2018) "FTC Hearing #8: Common Ownership," November.

FTC and DOJ (2024) "FTC, DOJ Submit Joint Comment to FERC Warning of Common Ownership Competition Risks in the Public Utilities Industry | Federal Trade Commission," April.

Ge, Chunmian, Ke-Wei Huang, and Ivan P. L. Png (2016) "Engineer/scientist careers: Patents, online profiles, and misclassification bias," *Strategic Management Journal*, Vol. 37, No. 1, pp. 232–253.

Geng, Heng, Harald Hau, Roni Michaely, and Binh Nguyen (2021) "Does Board Overlap Promote Coordination Between Firms?" *SSRN Electronic Journal*.

Ghezzi, Federico and Chiara Picciau (2022) "Evaluating the Effectiveness of the Italian Interlocking Ban: An Empirical Analysis of the Personal Ties Among The Largest Banking and Insurance Groups in Italy," *Journal of Competition Law & Economics*, Vol. 18, No. 1, pp. 29–74.

Gibson, Matthew (2024) "Employer market power in Silicon Valley,"Technical report.

Gopalan, Radhakrishnan, Renping Li, and Alminas Zaldokas (2024) "Board Connections, Firm Profitability, and Product Market Actions," *European Corporate Governance Institute–Finance Working Paper*, No. 996.

Helft, Miguel (2007) “In Fierce Competition, Google Finds Novel Ways to Feed Hiring Machine,” *The New York Times*.

Hoberg, Gerard and Gordon Phillips (2016) “Text-Based Network Industries and Endogenous Product Differentiation,” *Journal of Political Economy*, Vol. 124, No. 5, pp. 1423–1465, Publisher: The University of Chicago Press.

Lemley, Mark A., Anoop Manjunath, Nathan Kahrobai, and Ishan Kumar (2022) “Analysis of Over 2,200 Life Science Companies Reveals a Network of Potentially Illegal Interlocked Boards,” *SSRN Electronic Journal*.

Sun, Liyang and Sarah Abraham (2021) “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, Vol. 225, No. 2, pp. 175–199.

A Additional Figures and Tables

Table A.1: Difference-in-Differences Regressions for Pr(Agreement): By Type of Common Leader

	All Pairs			Agreeing Pairs	
	Any Leader	Brd. Only	Exec	Brd. Only	Exec
	(1)	(2)	(3)	(4)	(5)
Main coefficient: Column heading variable	0.119^{**}	0.105	0.123^{**}	0.018	0.816^{***}
	(0.051)	(0.065)	(0.050)	(0.150)	(0.211)
Control coef: Exec-involved common leader		0.094		0.559^{**}	
		(0.070)		(0.269)	
Control coef: Board-only common leader			0.133^{*}		0.306
			(0.076)		(0.191)
Firm pair FEs	Yes	Yes	Yes	Yes	Yes
Separate firm FEs, year FEs					
Firm $\times$ year FEs	Yes	Yes	Yes	Yes	Yes
Observations	16,245	16,244	16,245	1,103	1,103
R^2	0.57944	0.56714	0.58204	0.93744	0.93949
Within R^2	0.07624	0.04920	0.08204	0.26879	0.29274

First row reports ATT coefficient on the type of common leader in the column name. Control coefficient rows report coefficient on other type of common leader, included as a binary regresor. Standard errors in parentheses, multi-way clustered on firm i, firm j, and firm pair $(i, j$ or $j, i)$. ATT summarized using Sun and Abraham (2021) to account for staggered treatment timing. Baseline mean of dependent variable in main sample = 0.016. Significance codes: *: 0.1; **: 0.05; ***: 0.01.

Figure A.1: Court Exhibits About Establishment of Apple-Google Agreement

(a) Apple CEO Steve Jobs Approaches Google Co-Founder Sergey Brin

From: Sergey Brin <sergey@google.com> on behalf of Sergey Brin
Sent: Thursday, February 17, 2005 8:20 PM
To: emg@google.com; joan@google.com; Bill Campbell
Cc: arnnon@google.com
Subject: Re: FW: [Fwd: RE: irate call from steve jobs]

So I got another irate call from jobs today.
I don't think we should let that determine our hiring strategy but thought I would let you know.
Basically, he said "if you hire a single one of these people that means war".
I said I could not promise any outcome but I would discuss it with the executive team again.
I asked if he expected us to withdraw offers and he said yes.

In reviewing the data below again, I do think this could be treated as not just an employee referral since he referred essentially a whole team. So a compromise would be to continue with the offer we have made (to Redacted?) but not to make offers to any of the others unless they get permission from Apple.

In any case, lets not make any new offers or contact new people at Apple until we have had a chance to discuss.

--Sergey

(b) Apple and Google Enter Into Agreement

Subject: Google
From: "Danielle Lambert" <lambert@apple.com>
Received(Date): Sat, 26 Feb 2005 05:28:46 +0000
To: <usrecruitingall@group.apple.com>

All,

Please add Google to your "hands-off" list. We recently agreed not to recruit from one another so if you hear of any recruiting they are doing against us, please be sure to let me know.

Please also be sure to honor our side of the deal.

Thanks,
Danielle

Each figure is drawn from an exhibit presented as evidence in court. Apple and Google entered into an agreement in 2005. At the time, they shared two common leaders: Genentech CEO Arthur Levinson sat on both boards, and Intuit CEO Bill Campbell was chair of Apple's board and a Senior Advisor to Google. According to additional court exhibits, Bill Campbell facilitated the communication between Apple and Google that established the agreement.

Figure A.2: Event Study Estimates: Additional Leads and Lags

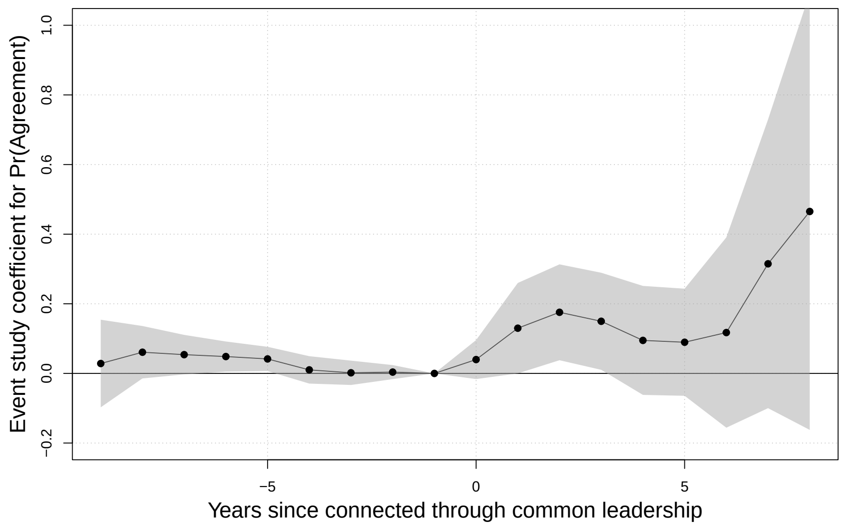

Figure plots additional leads and lags for event study coefficients from Figure 2. We caution that very few treated firm pairs are observed seven or more years after event time due to the short sample period.

Figure A.3: Event Study Estimates: Precise Agreement Dates

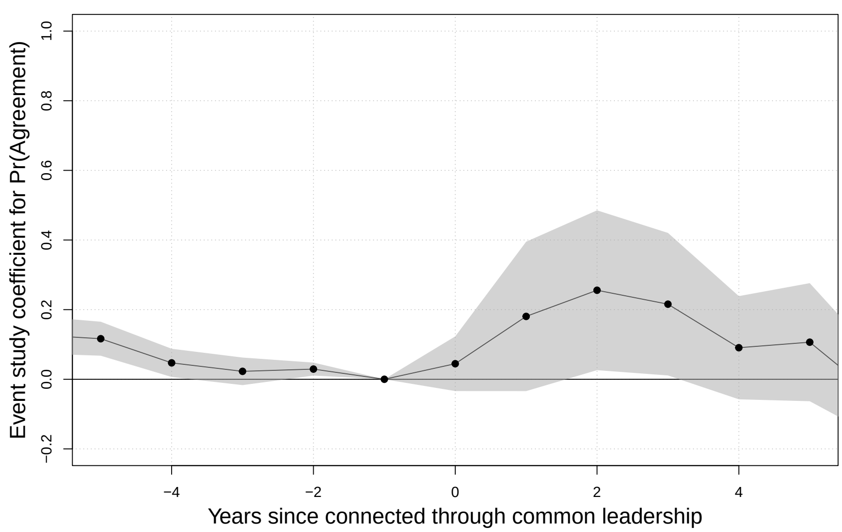

Figure plots event study coefficients from Equation 1 using the Sun and Abraham (2021) estimator to account for staggered treatment timing. Regression includes firm-by-year fixed effects and firm-pair fixed effects. The sample subsets to firms for which we observe a precise agreement start date.

Figure A.4: Event Study Estimates: Firm Pairs with Eventual Agreement

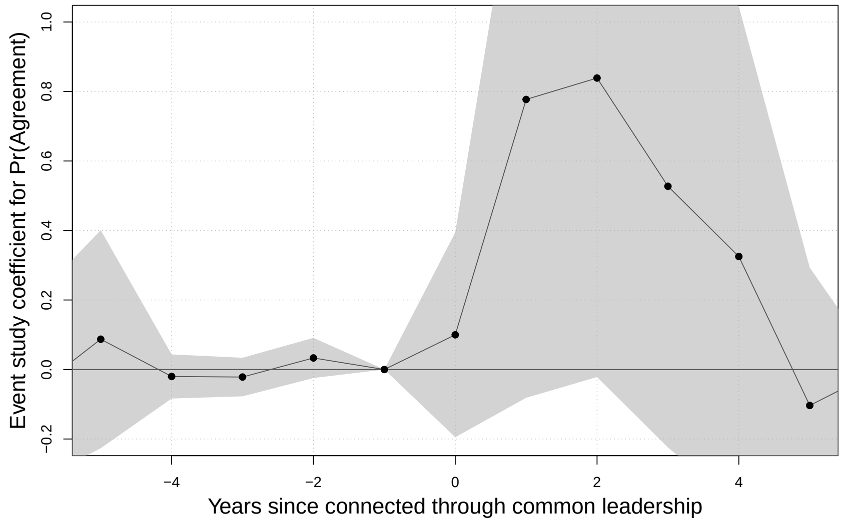

Figure plots event study coefficients from Equation 1 using the Sun and Abraham (2021) estimator to account for staggered treatment timing. Regression includes firm-by-year fixed effects and firm-pair fixed effects. The sample consists of pairs of firms that are observed to collude.

B Data Appendix

B.1 LinkedIn Data and Measure of Labor Market Overlap

Source of LinkedIn data. The measure of labor market overlap in Section 1.4 uses worker history data from LinkedIn. We obtain LinkedIn data from two sources. The first source is the Bright Initiative, a researcher-facing arm of the analytics firm BrightData. The BrightData sample was collected in 2022. Because it consists of publicly viewable LinkedIn profiles, which only show the most recent few entries in the user's job history, many worker histories are left-censored. We therefore supplement with a second source of LinkedIn data generously shared by Ge et al. (2016). The LinkedIn data in Ge et al. (2016) were collected in 2013 and improve our coverage of worker histories during our sample period.

Matching no-poaching court case sample firms to LinkedIn employers. In their raw form, LinkedIn data are partially structured. Key variables such as job title, employer name, and job dates are structured as separate variables. However, employer names are a free-text field and user inputs are not necessarily harmonized. For example, Apple workers may list their employer as Apple; Apple, Inc.; Apple Computer; Apple iPhone Group; and so on. We therefore clean employer names in a multi-step process in order to match the LinkedIn data to our regression data. First, we find the most unusual word in each employer name, relative to the corpus of all employer names in the LinkedIn data.[12] This produces a set of candidate matches of varying quality. For example, "Google" is an unusual word that rarely appears in entries not related to Google or Alphabet, so this step is highly informative for Google. For firms like Apple, which find many candidate matches related to fruit and apple cider, our next two steps are critical. In the second step, we assign a master employer name for each firm in our no-poaching case sample, and calculate string distances between the master name and candidate matches from the LinkedIn data. Finally, in the third step, we manually remove unrelated firms.

Worker sample for labor market overlap measure. Some of the firms with no-poaching agreements have substantial presence in the retail sector, such as Apple and J. Crew. The no-poaching agreements did not apply to store clerks. We therefore only use non-retail workers to calculate labor market overlap. To do this, we must first drop retail workers from the LinkedIn data. There are 24 million unique job titles in the raw LinkedIn data, so finding the retail workers manually is infeasible. We train a neural net to classify raw LinkedIn job titles into retail versus other job types. Specifically, we use a multilayer perceptron classifier with 1-grams and 2-grams (that is, one- and two-word strings) with five-fold cross-validation. The training dataset consists of 3,099 titles classified by ChatGPT 3.5 with extensive human supervision. We confirm high performance quality of the final classifier algorithm using manual checks of a random sample. Only workers classified as non-retail are kept for calculating the labor market overlap measure in Section 1.4. In practice, the list of implied closest competitors changes little regardless of whether retail workers are kept or dropped.

[12]We drop geographic place names to avoid matching together firms such as "First Kansas Bank" and "University of Kansas". The firms in our no-poaching court case sample all have sufficiently unusual names for our procedure to work well.