

Interlocking Directorates and Innovation

Mirko De Maria*

October 18, 2026

[Click here for the latest draft.](#)

Abstract

Connections between competing companies usually affect innovation, but their net effects and mechanisms remain unclear. This paper studies these effects by analyzing a prevalent type of corporate connection: shared board members. To account for persistent differences in board member ability, I introduce a new empirical strategy that uses estimated board member fixed effects from the AKM model as control variables in a second stage. I find that interlocked firms innovate more and cite their interlocked partners' patents more often. However, these new patents do not have higher market-based values and are increasingly likely to be continuation patents. Cross citations also shift increasingly toward older and more technologically distant partner patents.

*Department of Economics and Public Policy, Imperial College London; e-mail: m.de-maria@imperial.ac.uk. I am grateful to Jonas Lieber and Tommaso Valletti for their continued support and guidance, and to Simon Quinn for his detailed feedback on an earlier version of the draft. For helpful comments that improved the paper, I would also like to thank Pierre Azoulay, Nick Bloom, Claudia Custodio, Jacob Dorn, Gaia Dossi, Florian Ederer, Mireia Giné, Jillian Grennan, Laura Grigolon, Cristina Gualdani, Peter John Lambert, Josh Martin, Keaton Miller, Andrea Prat, Carol Propper, Jens Prüfer, Patrick Rey, and Joel Stiebale, as well as seminar, workshop, and conference participants at Imperial College London, the 2025 CMA-Durham Workshop, the 2025 CEPR Paris Symposium, the 2025 European Winter Meeting of the Econometric Society, the 2026 North American Summer Meeting of the Econometric Society, the 2026 Annual Conference of the Royal Economic Society, and the 2026 EARIE Conference. All mistakes are exclusively my own.

1 Introduction

Firms' incentives to develop new technologies depend not only on competition with other firms, but also on the connections that exist between them. Because innovation requires large investments in R&D and produces knowledge that can spill over across firms, connections between competitors could either dampen innovation incentives (by reducing rivalry) or enhance them (by sharing information and reducing duplicative R&D). These competing forces make the relationship between firm connectedness and innovation theoretically ambiguous. A prevalent form of connection among competing companies happens at the corporate level through shared board members. Such board interlocks can help reduce uncertainty and facilitate information exchange (Stein, 2008), but may also support coordination among rivals (Battaglion and Cerasi, 2020, Herrera-Caicedo, Jeffers and Prager, 2026). However, despite rich theoretical background, whether these organizational connections promote innovation or instead alter firms' strategic incentives remains an open empirical question.

This paper addresses the gap in the literature. The way it does it is by introducing a new way to account for non-random director assignment across firms. As an intermediate step, I estimate an AKM model (Abowd, Kramarz and Margolis, 1999) for my sample of board members and firms, and then use the resulting board member fixed effects as controls for latent director ability in a second stage. The AKM model exploits the connectedness of the director-firm matching network to separate persistent board member heterogeneity from persistent firm heterogeneity. This allows me to mute a fundamental—and typically unobserved—endogenous channel. The empirical design pairs these AKM controls with a comprehensive set of control variables, which are designed to control for all the known channels that make the relationship between interlocking directorates and innovation endogenous. The remaining variation is therefore intended to be as close as possible to selection on observables. As key robustness tests, I complement the baseline regressions with both (i) sensitivity analyses that quantify how strong residual selection on unobservables would have to be to overturn the coefficients of interest, and (ii) a set of event studies around the formation of a new interlock. I also analyze the external validity of the estimated board member fixed effects.

I find that interlocking directorates are associated with higher innovative output, mea-

sured as increases in both raw and citation-weighted counts of patents. At the same time, I find no evidence that these interlocks are associated with higher R&D spending or average market-based values. These estimates are robust to a wide variety of alternative specifications, including different estimators, alternative definitions of market competition, and other vectors of control variables. The sensitivity analyses also indicate that unobservables would need to be much stronger than the rich set of included observables to overturn the headline estimates, indicating that omitted variable bias is an unlikely driver of these estimates.¹ For instance, the test from [Cinelli and Hazlett \(2020\)](#) indicates that one would need an unobservable covariate that is approximately 9.2 times stronger than firm size or 46 times stronger than R&D spending to nullify my estimate on the number of patents in the OLS regression.

To understand the mechanisms behind why interlocked firms patent more, I run a series of additional regressions and event studies to distinguish between strategic patenting behavior as opposed to technological learning. I first show that, when two firms interlock with each other for the first time, they start cross-citing each other more in their patents. However, although interlocked firms patent more after treatment, these new patents are mostly continuation patents, which are commonly associated with strategic patenting and more litigation ([Berger, Blind and Thumm, 2012](#), [Hegde, Mowery and Graham, 2009](#), [Righi, 2026](#), [Righi and Simcoe, 2023](#), [Righi, Cannito and Vladasel, 2023](#)). Also, although they cite each other more often after treatment, these citations do not systematically expand into new partner technology classes. Nor do the data show a clear increase in citations to technologically proximate partner patents. Instead, the stronger post-interlock responses are concentrated in citations to partner patents that are older and not technologically close to the citing patent. This would also explain why, in the baseline regressions, there is no statistically significant effect on the market-based values of these new patents: if they were truly innovative, we would expect to see a value-based response.

This interpretation carries some policy implications. If interlocks were mainly promoting knowledge diffusion, their effect on patenting could be viewed as at least partly innovation-enhancing. By contrast, the evidence here provides evidence that most of the

¹The use of OLS regressions with a saturated vector of controls and sensitivity analyses in the absence of clear exogenous shocks is not new in empirical research in Economics. Arguably, the most famous example is [Altonji, Elder and Taber \(2005\)](#), in which the authors estimate the impact of Catholic high school attendance on college outcomes.

post-interlock increase in patent activity reflects strategic portfolio management. No result presented in this paper suggests that board interlocks promote the production of higher-valued or more innovative patents. That matters because strategic patenting may raise legal congestion or increase the complexity of overlapping intellectual property claims (?). Therefore, the costs of interlocks may extend beyond product market coordination or information sharing in the usual antitrust sense. This may imply that regulators evaluating interlocks would want to pay close attention not only to product market effects but also to how board connections may alter innovation strategy, which in turn may affect market entry.

This paper contributes to three strands of literature. The first one is the literature on interlocking directorates. In the recent empirical branch of this literature, [Poberejsky \(2024\)](#) uses exogenous firm mergers to find that horizontal interlocks increase cumulative returns on assets. He also documents that, when a horizontal interlock breaks, patenting decreases.² [Barone, Schivardi and Sette \(2025\)](#) exploit the effect of a ban on interlocking directorates in the Italian banking industry to find that, as a result of the law, previously interlocked banks experienced a decrease in their interest rates compared to banks that were never interlocked. [Herrera-Caicedo, Jeffers and Prager \(2026\)](#) provide direct evidence that firms become substantially more likely to collude after they begin sharing executives or board directors. Whereas they study explicit coordination in labor markets, I examine whether the same organizational link reshapes firms' innovation behavior and the strategic use of intellectual property. [Cabezon and Hoberg \(2026\)](#) show that firms embedded in denser competitor board networks experience innovation herding, lower product differentiation, lower patent values, and poorer performance. [Guiso, Schivardi and Zaccaria \(2026\)](#) study the spillover effects of board gender quotas in Italy and find that firms connected to regulated companies through interlocking directorates increase their female board representation despite not being subject to the quotas themselves. [Gualdani \(2021\)](#) develops a game of network formation with an application to interlocking directorates and finds that a firm's payoff from having an executive sitting on the board of a competitor increases with the number of other competitors that have an

²Apart from the empirical strategies being radically different, while [Poberejsky \(2024\)](#) looks at the effect of *losing* an interlocking directorate on financial returns using a difference-in-differences strategy, mine focuses on what happens when you *form* a new interlocking directorate. It is unlikely that the two phenomena have symmetric effects.

executive appointed to that firm's board. Earlier theoretical and qualitative work in sociology focused on the determinants of interlocking directorates and tried to uncover factors that are correlated with them (Dooley, 1969, Mizruchi, 1996). These include studies on industry concentration and horizontal interlocks (Burt, 1983, Pennings, 1980), cooptation (Allen, 1974, Bunting, 1976, Palmer, Friedland and Singh, 1986, Pfeffer, 1972, Pfeffer and Salancik, 1978, Selznick, 1949, Stearns and Mizruchi, 1993, Thompson and McEwen, 1958), career advancement (Stokman, Knoop and Wasseur, 1988, Zajac, 1988), and social ties (Mills, 1956).

This paper contributes to this literature in several ways. First, it provides new firm- and pair-level evidence that interlocks reshape innovative activity not only on the intensive margin of patenting, but also in the direction of technological interaction between connected firms. Existing work has argued that interlocks may transmit information, facilitate coordination, and affect strategic behavior (Burt, 1983, Mizruchi, 1996, Pfeffer and Salancik, 1978). Nevertheless, we lack empirical evidence on how these channels affect firm innovation. I address this by analyzing patenting behavior, citations, and the timing of first interlock formation and showing that board connections are associated with a reorganization of firms' innovation activity. Second, the paper opens the black box of the mechanisms more directly than the existing literature by distinguishing between genuine knowledge transmission and defensive use of patents. At the firm level, I show that the increase in patenting is concentrated in continuation patents, which are commonly linked with strategic patenting motives. Using pair-level event studies, I then show that post-interlock citation responses are disproportionately concentrated in older and non-technologically proximate partner patents. Third, this paper offers a new organizational perspective to the literature on innovation by showing that the consequences of interlocking are not restricted to product market coordination but extend to how firms manage patenting portfolios and position themselves in the technological space. In that sense, the paper bridges the literatures on corporate governance, innovation, and intellectual property (Cohen, Nelson and Walsh, 2000, Hall and Ziedonis, 2001, ?) by documenting that interlocks can operate through a strategic channel that has largely been overlooked in prior work.

The second literature studies common ownership among rival firms and, more gener-

ally, the consequences of firms sharing the same resources.³ Work in this area analyzes how horizontal links created by overlapping shareholders reshape competition, managerial incentives, innovation, and merger incentives. [Azar \(2022\)](#) shows that firms with a high degree of common ownership are more likely to share directors, while [Azar, Schmalz and Tecu \(2018\)](#) argue that common ownership can soften product market competition and raise prices. [Antón et al. \(2023a\)](#) develop a mechanism through top management incentives and show that, when firms share important shareholders, managerial pay becomes less tied to aggressive relative performance. This in turn can dampen competitive pressure. Relatedly, [Antón et al. \(2023b\)](#) show that common ownership can affect merger incentives and acquisition activity, while [Ederer and Pellegrino \(2026\)](#) quantify the welfare costs of the interaction between ownership overlap and product market rivalry. [Eldar and Grennan \(2021\)](#) study common ownership in startups backed by venture capitalists and argue that it differs from passive institutional ownership of public firms, since venture capitalists are active owners and may create channels through which information and coordination pass across firms. This literature is relevant for my paper as it shows that, in general, horizontal links among rivals need not have a uniform effect on innovation. For instance, [Antón et al. \(2025\)](#) shows that the sign of this effect depends on whether the dominant force is the internalization of technological spillovers or the softening of product market rivalry. Whereas common ownership links rivals through shared shareholders and changes incentives from above, interlocking directorates link rivals inside the boardroom and may operate through information transmission as well as through competition.

I contribute to this literature by showing that horizontal board links shape innovation not through broad increases in measured R&D, but through changes in the direction and strategic use of firms' patenting behavior. In particular, I show that while newly interlocked competitors increase cross citation activity, this response is concentrated in the partner patents that are older and non-technologically close to the citing patent, a pattern more consistent with strategic patenting than with technological learning. This paper complements the common ownership literature by identifying a board-level channel through which shared resources among competing firms can affect innovation and the management of patent boundaries. In my series of robustness checks, I also test how

³A related literature about mergers and innovation similarly emphasizes that horizontal integration can alter innovation incentives through the interaction of reduced competitive pressure and innovation externalities ([Federico, Langus and Valletti, 2017, 2018](#)).

the estimated coefficients would change if I were to control for average levels of common ownership at the firm level, and I find that the results would remain unchanged. This suggests that shared board members may shape firms' innovation strategies in addition to the influence of common owners.

The third and final strand of literature studies AKM models. This class of models was originally introduced by [Abowd, Kramarz and Margolis \(1999\)](#). Since its introduction, the AKM model has provided researchers with a way to quantify how much of wage inequality is attributable to worker-specific factors, firm-specific factors, and the sorting between high-wage workers and high-wage firms. In the recent past, this framework has been employed in several empirical applications (see [Card et al. \(2018\)](#) and [Kline \(2024\)](#) for a comprehensive overview), as well as sparking a lively literature on the estimation of the model. For instance, [Andrews et al. \(2008\)](#) showed that the so-called limited mobility bias can depress the estimated correlation between worker and firm effects, and they propose a bias-correction method.⁴ [Andrews et al. \(2012\)](#) report clear evidence that this bias masks positive assortative matching. [Kline, Saggio and Sølvesten \(2020\)](#) develop a leave-one-out correction for estimating the variance of firm effects and the covariance. The current consensus, accounting for the latest bias-correction methods, is that there is positive assortative matching in most labor markets. In the empirical application of this paper, however, I am not concerned about the estimation of such sorting. Instead, I use the AKM model only to estimate the individual fixed effects of the board members. In a correctly specified AKM model, worker and firm fixed effects are unbiased regardless of positive or negative sorting. Therefore, the correction approaches that the literature has discussed to correctly estimate sorting do not apply to my case.

To my knowledge, this is the first paper to use the AKM model in an intermediate step to control for an endogenous and unobservable channel. In my empirical strategy, I first estimate the AKM model for the largest connected sample of board members and firms. Then, I collect the estimated director fixed effects –which are, by definition, a proxy variable for their time-invariant ability– and aggregate them at the firm-year level to con-

⁴Limited mobility bias can arise because worker and firm fixed effects in AKM models are identified from workers who move across firms. The connectedness of the worker-firm network is particularly weak when the mobility of workers across firms is sparse, which means that the sampling noise in estimated firm effects can attenuate the estimated covariance between worker and firm effects. In other words, if not too many workers move across firms in the data, it becomes difficult to distinguish whether high wages are due to the worker or the firm, which biases downward the estimates of assortative matching.

struct a measure of aggregate board ability for every company in each year. Using this firm-level aggregated AKM measure as a control variable allows me to mute the board member ability channel, which would otherwise complicate the analysis. In the corporate finance literature, [Bertrand and Schoar \(2003\)](#) exploit the movement of top executives across firms to estimate manager fixed effects and show that managers account for persistent differences in corporate performance. [Graham, Li and Qiu \(2012\)](#) use the AKM model to study how firm and executive-specific effects explain executive compensation, and find that executive fixed effects explain a large portion of the variation in pay. [Cavaco et al. \(2017\)](#) applies the AKM model to find that having more independent directors on the board correlates with lower firm performance. Importantly, [Bertrand and Schoar \(2003\)](#) and [Graham, Li and Qiu \(2012\)](#) focus on top executives, while [Cavaco et al. \(2017\)](#) focuses on French board members, using independence as the key margin of interest. My AKM model differs from theirs both in the sample, as I look at the universe of executives and independent directors in public U.S. firms, and in the purpose, as I use the AKM fixed effect estimates as control variables in a second stage and not as an outcome of interest per se.

The rest of this paper is organized as follows. Section 2 illustrates the institutional context, the data I use, and a series of descriptive statistics. Section 3 lays out my identification strategy. Then, Section 4 shows the baseline results. After that, Section 5 discusses and shows results related to the mechanisms that may explain the results I document. Section 6 discusses the robustness of these results to alternative specifications. Finally, Section 7 concludes.

2 Institutional Context and Data

2.1 Section 8 of the Clayton Act and Recent Ramp-Up in Enforcement

Because of collusion concerns, Section 8 of the Clayton Antitrust Act deemed interlocking directorates illegal in 1914, and they have remained so ever since. Not all forms of interlocking directorates are unlawful: only those between competing companies (i.e., in the

same line of industry) that exceed size-based thresholds.⁵ Nevertheless, antitrust enforcement against interlocking directorates has been notoriously lax since the establishment of this law. Between 1899 and 2022, the U.S. Department of Justice (DoJ) was involved in just eleven cases involving interlocking directorates (Poberejsky, 2024). Several reasons may explain this lenient antitrust stance, including for example ambiguity over the relevant notion of competition and the FTC’s broad discretion in exercising its enforcement powers (Nili, 2022).

That said, under the Biden administration in particular, enforcement of Section 8 of the Clayton Act ramped up. In 2021, the DoJ started renewing its scrutiny of interlocking directorates. In June of that year, they broke an interlock between Endeavor and Live Nation, two sports and entertainment companies that closely competed in several markets. Scrutiny of interlocking directorates only grew stronger after that, and in October 2022, seven more directors from five different companies also resigned in response to DOJ concerns about potentially illegal interlocks.⁶ In 2023, the FTC followed suit and announced an order that would prohibit Quantum from obtaining a seat on EQT’s Board of Directors following EQT’s acquisition of certain natural gas assets from private equity firm Quantum. Behind this decision was the FTC’s explicit worry that this board appointment would create a potentially illegal interlock.⁷

2.2 Data on Boards and Interlocking Directorates

The main source of information on boards in this paper is BoardEx, which assembles board rosters for U.S. public companies from disclosures filed with the Securities and Exchange Commission (SEC). For each appointment to a board, BoardEx records the names and the titles of the board members along with their demographic attributes, compensa-

⁵This revenue threshold is set at \$54,402,000 (for Section 8(a)(1) of the Clayton Act) as of the time of this writing. Interlocking directorates are still permissible for competing firms if and only if the sales of either corporation are less than \$5,440,200 (Section 8(a)(2)(A) of the Clayton Act), the competitive sales of either corporation constitute less than 2% of that corporation’s total sales, or the competitive sales of both corporations constitute less than 4% of each corporation’s total sales. These thresholds are reviewed each year—usually in January—by the Federal Trade Commission (FTC).

⁶Directors Resign from the Boards of Five Companies in Response to Justice Department Concerns about Potentially Illegal Interlocking Directorates ([justice.gov](https://www.justice.gov/opa/pr/directors-resign-from-the-boards-of-five-companies-in-response-to-justice-department-concerns-about-potentially-illegal-interlocking-directorates))

⁷FTC Approves Final Order to Prevent Interlocking Directorate Arrangement, Anticompetitive Information Exchange in EQT, Quantum Energy Deal ([ftc.gov](https://www.ftc.gov/press-release/ftc-approves-final-order-prevent-interlocking-directorate-arrangement-anticompetitive-information-exchange-eqt-quantum-energy-deal))

tion components, and employment histories. From these microdata, I reconstruct board composition over time and identify interlocking directorates by identifying the firms that share at least one board member. To then classify horizontal ties, I implement a rule-based Section 8 screen. I first restrict attention to firm pairs that are competitors at the chosen industry cut⁸ and that share at least one board member (either an independent director or an executive) in a given year. I then merge the FTC’s annual Section 8 thresholds with this information and use Compustat to determine whether each firm is covered by the statute in a given year and is not exempt under absolute de minimis sales cutoff. Finally, I classify an observed horizontal interlock as Section 8–violating when both firms are eligible and neither belongs to an excluded entity type, such as banks or trust entities. Appendix A.1 provides the full algorithm and implementation details.

Figure 1 summarizes the prevalence of Section 8–violating horizontal interlocks in the sample. Panel 1a shows that illegal interlocks are far from rare: in recent years, roughly 30% of public firms in the U.S. engaged in the practice with a competitor, while the average number of interlocks per firm was 0.67 in 2024. Panel 1b shows the cross section from 2024 which highlights that, conditional on interlocking, most firms only had one horizontal interlock. Only a handful collects three or more at the same time. In Appendix A.5.4, I relax the NAICS-2 assumption and focus on all interlocks, including those between firms with different NAICS-2 codes. Although the regression results do not change under this more relaxed definition, I do find that the estimated prevalence of interlocks would rise substantially.⁹ Finally, in Appendix A.2, I report three empirical facts about interlocking directorates that, to the best of my knowledge, are novel.


2.3 Measures of Innovation

In my analysis, I distinguish between firm-level innovation inputs, outputs, and values. Inputs reflect the resources that firms devote to the innovation process, outputs capture the technological products of that process, and value measures the market-based economic importance of the innovations that are ultimately produced. First, to measure in-

⁸I use NAICS-2 codes as a baseline but later relax this restriction to test for robustness.

⁹See Figure A.5.1 for the counts of industry-unrestricted interlocking directorates.

Figure 1: Firm-Level Statistics About Horizontal Interlocking Directorates, 2000-2024



Notes: Panel (a) shows the average number of horizontal interlocking directorates (blue line, left axis) across all firms in the sample by year, and the percentage of firms that have at least one interlocking directorate by year (red line, right axis). Panel (b) shows the cross sectional distribution of the number of horizontal interlocking directorates by firm in 2024, along with the percentage of firms for each category.

novation inputs, I use the R&D expenditure data from Compustat in U.S. Dollars¹⁰ To measure innovation outputs, I use two patent-based outcomes. The first is the raw count of new patents issued to a firm in a certain year. The second one is a citation-weighted count, defined as

$$CW \text{ Patents}_{jt} = \sum_{p \in P_{jt}} \left(1 + \frac{C_p}{\bar{C}_p} \right),$$

where P_{jt} is the set of patents granted to firm j in year t , C_p is the number of forward citations received by patent p , and $\bar{C}_p$ is the average number of citations received by patents granted in the same year as patent p . Both the firm-level patent counts and the number of citations come from the dataset built by Kogan et al. (2017), which has now been extended to the end of 2024, and builds on the patent data from Hall, Jaffe and Trajtenberg (2001).

¹⁰In Compustat, R&D spending is missing for many firms. Therefore, I follow the standard in the literature and replace missing R&D values with zeros and then use a dummy variable equal to one if the reported R&D is zero as a control variable (Antón et al., 2025).

Table 1: Mean Firm-Level Summary Statistics for All Firms and All Available Years

	(1) All Firms, 2000–2024
<i>Panel A: Firm-Level Statistics</i>	
Assets	14,173.82
Revenue	6,302.94
COGS	4,159.07
Employees (k)	17.16
Patenting Dummy	0.38
<i>Panel B: Outcome Variables</i>	
R&D	169.43
Patents	26.24
Citation-Weighted Patents	50.93
Average Patent Value	8.39
N (firm-years)	10,034
N (firms)	1,949

Notes: Columns report means across firm-years. Assets, Revenue, COGS, and R&D are in U.S. dollars (millions). Employees are in thousands. R&D Intensity is the ratio of R&D to assets. At Least One Patent Indicator is a dummy variable equal to 1 if a firm patented at least once during the year.

Finally, to measure the value of innovation, I again follow [Kogan et al. \(2017\)](#) and use stock market reactions around patent grants as a proxy for the private economic value of innovation ([Hall, Jaffe and Trajtenberg, 2005](#)). In particular, for each patent p , [Kogan et al. \(2017\)](#) estimate the firm’s abnormal stock-market reaction, denoted by ξ_p , over the three-day event window around the grant date. I then aggregate these patent-level values to the firm-year level by summing across all patents granted to firm j in year t :

$$\text{Total Value}_{jt} = \sum_{p \in P_{jt}} \xi_p,$$

where P_{jt} again refers to the number of patents granted to firm j in year t . This variable captures the total private economic value that the stock market attributes to the firm’s

newly granted patents in that year. I also construct an average-value measure,

$$\text{Average Value}_{jt} = \frac{\text{Total Value}_{jt}}{P_{jt}},$$

which captures the mean private economic value of the patents granted to firm j in year t . Table 1 summarizes the variables introduced above.

3 Empirical Strategy

Let j index firms and t years. Use Y_{jt} to denote an outcome of interest (e.g., patent counts, citation-weighted patents, log R&D, or patent-value measures). Let Interlock_{jt} denote the interlock treatment, measured as an indicator for whether firm j holds at least one relevant interlocking directorate in year t . The baseline estimating equation is

$$Y_{jt} = \beta \text{Interlock}_{jt} + X_{jt}\phi + A_{jt}\pi + \psi_j + \lambda_{s(j)t} + \eta_{r(j)t} + \varepsilon_{jt}, \quad (1)$$

where ψ_j are company fixed effects, $\lambda_{s(j)t}$ are industry×year fixed effects (with $s(j)$ denoting the firm’s industry), and $\eta_{r(j)t}$ are state×year fixed effects (with $r(j)$ denoting the firm’s headquarters U.S. state). The vector X_{jt} collects a rich set of time-varying firm and governance controls, chosen using double LASSO. These controls come from Compustat and BoardEx. The coefficient of interest, β , captures the association between interlocks and the outcome. In an alternative specification,

$$Y_{jt} = \beta \text{Interlock}_{jt} + X_{jt}\phi + A_{jt}\pi + \psi_j + \tau_t + \varepsilon_{jt}, \quad (2)$$

one may use company fixed effects ψ_j and year fixed effects, τ_t . Each table will report estimates from both Equation 1 and 2.

These equations also include A_{jt} , which denotes board ability, which means that the neglect of A_{jt} from the estimated regression would lead to omitted variable bias. Ability is something we usually cannot explicitly control for, as it is unobserved in most empirical settings. This is why, in this paper, I introduce a novel identification strategy that hinges on the AKM model to partial out the effect of A_{jt} on Y_{jt} . In particular, I first estimate

board member ability in an intermediate step using the AKM model. This is a framework that originates from [Abowd, Kramarz and Margolis \(1999\)](#), which they used to decompose wage variation between worker and firm effects. In its canonical form, the AKM is a two-way fixed effect model:

$$W_{ijt} = \gamma_i + \psi_j + X_{jt} \theta + Z_{it} \xi + \epsilon_{ijt}, \quad (3)$$

where W_{ijt} is the wage earned by worker i in firm j at time t . This wage is then regressed on both worker fixed effects and firm fixed effects, along with a vector of control variables (X_{jt}, Z_{it}) collecting firm- and director-level observables. In the presence of the standard assumptions of this model, namely, mean independence of the error term and no match effects, the worker and firm fixed effects are unbiased. In my application on board members, I use the total compensation package (which includes bonuses and options on top of the baseline salary) as my measure of W_{ijt} . I collect this compensation data from ExecuComp.

There is a large and rich literature that deals with bias in the measurement of sorting and variance decomposition in this model ([Andrews et al., 2008, 2012](#), [Bonhomme et al., 2023](#), [Card et al., 2018](#), [Card, Heining and Kline, 2013](#), [Eeckhout and Kircher, 2011](#), [Kline, Saggio and Sølvsten, 2020](#)). However, in my empirical framework, I do not compute plug-in variance decompositions. My exclusive purpose is to use the AKM model as illustrated in Equation 3 to recover $\hat{\gamma}_i$ as a proxy for director ability. After estimating the model, I collect the $\hat{\gamma}_i$ and use them to construct board-level ability aggregates at the firm-year level and obtain A_{jt} . These firm-level AKM aggregates enter downstream regressions as controls for the ability channel that simultaneously drives interlock formation and outcomes. I restrict these further analyses to firm-years in which the full board is observed. The individual fixed effects are known to always be unbiased and consistent in the presence of complete information about the workers' wages ([Kline, 2024](#)).

Even after controlling for A_{jt} and all other channels, one cannot rule out remaining selection on unobservables using OLS. We can, however, determine how large unobservables would have to be to nullify the estimates on β from Equations 1 and 2. Hence why, in the second part of my identification strategy, I complement the AKM and baseline regressions with a range of modern sensitivity analyses ([Altonji, Elder and Taber, 2005](#), [Cinelli and Hazlett, 2020](#), [Diegert, Masten and Poirier, 2025](#), [Masten and Poirier, 2026](#), [Oster, 2019](#)). Every regression that I estimate in this paper is complemented by several tests: in

the main part of the text, I report the sensitivity analyses from [Cinelli and Hazlett \(2020\)](#). In Appendix [A.4](#), I report the complementary Oster- β and Oster- δ estimates from [Oster \(2019\)](#), as well as the sign-change breakdown point from [Masten and Poirier \(2026\)](#). The same appendix section contains an in-depth mathematical illustration of each of these tests.

4 Results

4.1 AKM Model

The director fixed effects from the AKM model reflect the time-invariant ability of each board member. They are, effectively, the residuals of their compensation after partialling out time trends, company fixed effects, individual and firm characteristics, and other relevant covariates that may affect how much they get compensated. If better board members enjoy better careers and sort into more successful firms, as indeed several papers show ([Baltrunaite and Karmaziene, 2024](#), [Ferris, Jagannathan and Pritchard, 2003](#), [Masulis and Mobbs, 2014](#), [Yermack, 2004](#)), then simple regressions should show that directors with higher estimated AKM fixed effects should, for example, work for larger, productive firms and have more seats throughout their career.

Table [2](#) tests precisely this hypothesis. For easier interpretation, in this set of regressions, I standardize the director fixed effects first. This way, each coefficient reflects a change in one standard deviation of the director’s ability on the outcome. In short, the estimated director ability is a strong predictor of career success. An increase in one standard deviation of the director’s ability is associated with, on average, 2.3% more simultaneous board seats. Better directors are also much more likely to work for larger firms, companies with bigger revenue, and firms that are in the S&P 500 and/or the Fortune 500. Finally, they enjoy slightly longer careers, too.^{[11](#)}

Throughout the years, several methods have been developed to test whether the assumptions of the canonical AKM model hold. I discuss these tests in Appendix [A.3.1](#). There I show, for example, that the additively separable and exogenous mobility assump-

¹¹In Appendix [A.3.2](#), I report correlations between the firm-level AKM controls and innovation measures.

Table 2: Director Career Outcomes and the AKM Director Fixed Effect

	(1) Log(Seats)	(2) Log(Firm Size)	(3) Log(Revenue)	(4) Log(1+S&P 500)	(5) Log(1+Fortune 500)	(6) Log(Career Length)
AKM Director FE (SD)	0.023*** (0.004)	0.230*** (0.015)	0.201*** (0.014)	0.057*** (0.003)	0.044*** (0.003)	0.016*** (0.002)
N	16,389	16,318	16,318	16,389	16,389	16,389
R-squared	0.222	0.052	0.049	0.046	0.070	0.399

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Robust standard errors are reported in parentheses. The table reports OLS estimates from six separate cross-sectional regressions. The unit of observation is a director. The sample is US directors for whom a director fixed effect is estimated from the AKM model, standardised to mean zero and unit variance over the estimation sample. Coefficients reflect the change in the outcome associated with a one standard deviation increase in the estimated director fixed effect. All specifications control for gender, decade of birth indicators, indicators for first-active decade, and years active. Years active is omitted from the career-length regression, where it is mechanically collinear with the outcome. *Seats* is the number of simultaneously held US board seats averaged over active years. *Firm Size* and *Revenue* are Compustat total assets and revenue, averaged over director-firm-year service years. *S&P 500* and *Fortune 500* are counts of appointments at firms holding that status in the appointment year. *Career Length* is the span in years from a director's first to last recorded role.

tions are satisfied in my data.

4.2 Baseline Regressions

Table 3 reports baseline firm-year estimates of Equation 1 and 2 for horizontal interlocks. The treatment variable, $Interlock_{jt}$, is an indicator equal to one if a firm j has at least one horizontal interlocking directorate in year t , and zero otherwise. All specifications condition on a rich set of firm controls chosen by double LASSO to control for observed channels on top of the above-mentioned vector of fixed effects.¹² Further, all these specifications always include the aggregated firm-level AKM estimates discussed in the previous subsection to partial out board member ability, and covariates that control for technological and product market spillovers.¹³

In this set of results, I find that the point estimates for R&D spending are close to zero and non-statistically significant. However, when we turn to innovation outputs, we can see that horizontal interlocks are positively associated with patenting, both in raw counts

¹²In Appendix A.8, I report an additional outcome variable which is outside of the scope of this paper but may nevertheless be of interest to the reader, which is firm-level markups. These markup measures are constructed using the De Loecker and Warzynski (2012) method. In short, I do not find statistically significant results for this outcome.

¹³In Appendix A.9, I report heterogeneity results by industry.

and when weighted by citations. The estimated coefficients range from 0.116 to 0.134 log points for the raw patent count and from 0.121 to 0.143 log points for the citation-weighted equivalent, implying roughly a 12–13% increase in raw patent output and a 12–14% increase in citation-weighted output, depending on the specification. Relative to the R&D spending results, the patent count estimates are both larger and more precisely estimated. Finally, in the last two columns of the same table, I look at the average value of those patents. The point estimates for those are very close to zero and non-statistically significant at conventional levels. Therefore, while interlocked firms patent more, the value of those patents does not increase on average.

Table 3: Effects of Horizontal Interlocks on Innovation Inputs, Outputs and Values

	(1) Log (R&D)	(2) Log (R&D)	(3) Log (Patents)	(4) Log (Patents)	(5) Log (CW Patents)	(6) Log (CW Patents)	(7) Log (Avg. Value)	(8) Log (Avg. Value)
Interlock	0.015 (0.027)	0.005 (0.024)	0.116*** (0.042)	0.134*** (0.043)	0.121** (0.049)	0.143*** (0.051)	0.005 (0.051)	0.007 (0.045)
N	4,283	4,528	4,575	4,810	4,575	4,810	4,246	4,461
R-squared	0.990	0.986	0.933	0.918	0.927	0.909	0.840	0.787
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one Section 8–violating horizontal interlock in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

4.3 Sensitivity Analyses

Even after controlling for director ability and all observable channels, one may still be worried about remaining selection on unobservables using OLS as an empirical framework. Nevertheless, we can use sensitivity analysis tests (Altonji, Elder and Taber, 2005, Cinelli and Hazlett, 2020, Diegert, Masten and Poirier, 2025, Masten and Poirier, 2026, Oster, 2019) to quantify how strong the remaining selection on unobservables would have to be to nullify our coefficients of interest. In this Section, I report the sensitivity analyses for the innovation outputs, since these are the outcomes that are statistically significant.

Finally, for brevity, in the main text I will explain these analyses mostly in words. For an in-depth review of the mathematics behind these tests, I refer the reader to Appendix A.4.

Table 4 shows the statistics from the state-of-the-art sensitivity analysis test from Cinelli and Hazlett (2020) for the innovation output coefficients. Column (1) reports the same point estimate from Table 3, while columns (2) and (3) report its associated homoskedastic standard errors and t-statistics. Column (4) is the share of the residual variance of the outcome, after the controls and fixed effects are partialled out, that interlocks explain. The small values are consistent with the saturated design, as the fixed effects absorb the majority of the variance. Column (5) is the robustness value, or the minimum strength a confounder would need to have with both the treatment and the outcome, in equal measure, to drive the corresponding estimate to zero. For example, the first row tells us that an unobservable would have to explain 6.25% of the residual variance of the interlock status and also 6.25% of the residual variance of log patents—after partialling out all the covariates and fixed effects—to drive the coefficient to zero.

Is this a small or large number? Cinelli and Hazlett (2020) give no universal threshold, and the number is meaningless without a benchmark. Their advice is to benchmark this estimate against strong observed confounders. This is where the comparison with columns (7), (8), (9), and (10) is useful. Column (7) reports the share of the residual variance of the outcome that would be explained by an unobserved confounder assumed to be exactly as strong as firm size, measured by total assets, and column (8) reports the share of the residual variance of the treatment that the same confounder would explain. Columns (9) and (10) report the corresponding estimates for a confounder as strong as R&D. These results tell us that driving the estimate to zero requires an unobservable covariate explaining at least 6.25% of the residual variance of both the outcome and the treatment status. But a confounder as influential as total assets would explain 2.79% of the residual variance in patenting and 0.17% of that in interlock status, which is roughly 44.64% or 2.72% of what is required on each margin. These imply a Cinelli–Hazlett bias factor equal to about 10.7% of that required to eliminate the estimated coefficient. In other words, one would require an unobservable some 9.2 times stronger than firm size to overturn the results for patent counts. Similarly, a confounder as influential as R&D spending would explain 2.17% of the residual variance in patenting and very little of that in interlock status. Because bias requires association with both, such a confounder accounts for roughly

1.6% of the estimated coefficient according to the [Cinelli and Hazlett \(2020\)](#) bias expression, and confounding some 46 times stronger than R&D spending would be needed to overturn the result. Confounding on that scale, operating within firm, industry-by-year and state-by-year variation, is not something the observed covariates give any reason to expect. Very similar observations apply to the other three rows. In [Appendix A.4](#), I also report the results using the sensitivity analysis tests from [Oster \(2019\)](#) and [Masten and Poirier \(2026\)](#), which confirm that these estimates are robust to omitted variable bias.

4.4 Event Studies

As a further robustness check, we can also inspect patent outcomes around the time when two firms interlock with each other for the first time. If interlocking directorates do indeed spark an increase in innovation, we should be able to see a larger increase in the number of total patents issued by the two interlocking firms compared to a control group formed by non-interlocking dyads. In particular, I estimate a dyad-level event study using the staggered adoption difference-in-differences estimator from [Callaway and Sant’Anna \(2021\)](#). Let G_{fj} denote the first year in which dyad (f, j) becomes interlocked. For each treatment cohort g and calendar year t , the estimator estimates a group-time average treatment effect, $ATT(g, t)$, comparing dyads first treated in g with never-treated dyads. I then aggregate these $ATT(g, t)$ ’s by event time $k = t - g$:

$$\theta_k = \sum_{g: g+k \leq T} w_{g,k} ATT(g, g+k), \quad (4)$$

with $k = -1$ normalized to zero. $w_{g,k} = \Pr(G = g \mid G + k \leq T)$ is the relative share of treatment cohort g among cohorts observed at event time k .


The results are shown in [Figure 2](#) for the patent counts, both raw and weighted by citations. In particular, [2a](#) shows the total number of patents issued by the dyads and [2b](#) shows the total number of citation-weighted patents, again at the dyad level. In both cases, patent counts increase in the medium-to-long run following the establishment of a new interlock between two firms. Specifically, there is a non-statistically significant effect in the two years after the interlocking event, but then the count of patents starts increasing from the third year and becomes statistically significant, which signals a slow,

Table 4: Cinelli-Hazlett Sensitivity Analyses on the Effects of Horizontal Interlocks on Innovation Inputs, Outputs and Values

Outcome	(1)	(2)	(3)	Treatment: Interlock		(5)	(6)	(7)	(8)	(9)	(10)
	Estimate	Std. err.	t-value	$R^2_{Y \sim D X}$ (%)	RV (%)	df	$R^2_{Y \sim Z D,X}$ (%)	$R^2_{D \sim Z X}$ (%)	$R^2_{Y \sim Z D,X}$ (%)	$R^2_{D \sim Z X}$ (%)	
<i>Panel A: Firm + State-Year + Industry-Year FEs</i>											
Log(Patents)	0.12	0.03	3.65	0.41	6.25	3211	2.79	0.17	2.17	0.00	
Log(CW Patents)	0.12	0.04	3.16	0.31	5.42	3211	1.70	0.17	0.98	0.00	
<i>Panel B: Firm + Year FEs</i>											
Log(Patents)	0.13	0.03	4.82	0.56	7.23	4128	2.52	0.04	1.31	0.00	
Log(CW Patents)	0.14	0.03	4.27	0.44	6.42	4128	1.50	0.04	0.37	0.00	

Notes: Each row reports the sensitivity of the Interlock coefficient in the corresponding OLS specification to unobserved confounding, following the framework of Cinelli and Hazlett (2020). Panel A reports the specifications with firm, state-year and industry-year fixed effects; Panel B those with firm and year fixed effects. Estimate, Std. err., t -value and df are obtained by re-estimating that specification by OLS with its own fixed effects entered as dummy variables, on the same estimation sample as the corresponding baseline regression. The point estimates coincide with those reported in the baseline OLS tables, but the standard errors are homoskedastic and are *not* the firm-clustered standard errors reported there. $R^2_{Y \sim D|X}$ is the partial R^2 of the treatment with the outcome. RV is the robustness value: the minimum strength of association, as a share of residual variance, that an unobserved confounder would need with both the treatment and the outcome to reduce the estimate to zero. The last four columns bound the strength of a confounder Z assumed to be as strong as the named observed covariate ($k_D = k_Y = 1$). $R^2_{Y \sim Z|D,X}$ is the share of the residual variance of the outcome explained by Z , and $R^2_{D \sim Z|X}$ is the share of the residual variance of the treatment explained by Z . A benchmark whose implied values fall below RV cannot on its own overturn the estimated effect.

Figure 2: Event Studies - Patent Counts



Notes: Each panel reports staggered-adoption event study estimates of the effect of first adopting an interlock on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) framework, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. Event time is defined relative to the first interlock ($k = 0$) between any two firms. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Logged outcomes use $\log(1+x)$. All specifications include calendar-year fixed effects and the baseline dyad-level controls as described in the text. *Patents* and *Citation-Weighted Patents* are defined at the dyad level.

steady increase in innovation following a new interlock rather than a rapid jolt immediately following the event. This is also consistent with the view that innovation and production take some time to adjust following an event that may impact them. Further, note that the magnitudes of the coefficients in these event studies match those from the OLS regressions: Figure 2 shows that, a few years after the formation of a new interlocking directorate, interlocked firms issue on average seven to thirteen percent more patents than non-interlocked firms, which is similar to the range reported in Table 3. Despite these analyses being fundamentally quite different, they lead to the estimation of the same sign and similar range of magnitudes.

5 Mechanisms

5.1 Patent Citations

Why does innovation increase after the formation of an interlock? The question may not only be interesting per se, but may also carry policy implications. For example, an increase in patents may be consistent with a parallel increase in the transmission of know-how between the two competitors (Stein, 2008). This may lead either side of the dyad to innovate and formulate new patents that may otherwise never see the light of day. If that is the case, the lax approach on the part of U.S. authorities may have been beneficial to American innovation over the years. On the other hand, the same increase in patents could also be explained by a more strategic or defensive approach. According to this alternative theory, firms use their board links to learn about the interlocked partners' innovation portfolios and use that to issue patents that may guard their own innovation.

In this section, I probe these mechanisms. In Subsection 5.1.1, I explore whether the increase in patents after an interlocking event is mirrored by an increase in cross-citations between the two firms. This will give us a sense of whether the new interlocked firms show greater attention to the partner's patents to begin with. In Appendix A.6.1, I explore whether these dynamics differ between origin and destination firms. Finally, in Subsection 5.2, I study which patents get issued after the formation of an interlock, including a look at continuing patents. If this business practice is consistent with a theory of positive know-how use, we should see interlocking firms making more innovative patents.

For example, we should expect that a newly interlocked firm starts innovating in areas or patent classes that the interlocked partner is already familiar with. On the other hand, if a more strategic channel dominates, firms should be less likely to enter the partner’s technological space and more likely to cite the partner’s prior patents in ways that are more consistent with protecting their existing product market positions.

5.1.1 Cross Citations


All the event studies in this and subsequent subsections follow the same dyad-level design illustrated in Equation 4. Figure 3 reports the event-study dynamics for total cross citations between two firms. We can see that, after the first interlock, cross citations rise steadily and become statistically different from zero, reaching roughly 2 percentage points in magnitude by the end of the timeframe.¹⁴

5.2 Innovation or Strategic Patenting?

At the beginning of this paper, I argued that connections between competing firms could either promote or dampen innovation. In the last few pages, we showed that the net effect is positive, at least for corporate connections at the board level. However, it is difficult to judge from patents and citation counts alone whether these results are consistent with technological learning or not. As a matter of fact, firms could use the same information obtained from newly formed interlock links in two contrasting ways. They could either use this knowledge to start developing a more innovative patent portfolio, or use it to defend their existing products without producing innovation. The former would be related to the IP leakage mechanism in [Cabezon and Hoberg \(2026\)](#), which shows that dense overlapping-director networks among competitors accelerate the propagation of technologies across rivals. The latter would, for example, involve the use of strategic patenting (?). To understand which of these two mechanisms dominates, I first run a series of event studies that look at whether the patents and citations of newly interlocked dyads differ from the control group formed by never-interlocked dyads. Then, I analyze

¹⁴In Figure A.6.1 of Appendix A.6, I show the same cross-citation event study split by origin (firms that appoint the interlocked director first) and destination firms (firms that appoint the interlocked director at a later time). The effect is symmetric for origin and destination firms.

Figure 3: Event Studies - Intensive Margin of Patent Citations



Notes: This figure reports staggered-adoption event study estimates of the effect of first adopting an interlock on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) framework, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. Event time is defined relative to the first interlock ($k = 0$) between any two firms. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Logged outcomes use $\log(1 + x)$. All specifications include calendar-year fixed effects and the baseline dyad-level controls as described in the text. *Cross-Citations* is the share of a firm-year's patent citations that go to (or come from) firms with which the focal firm shares at least one director in year t .

whether there are differences in how interlocked firms develop continuing patents.

5.2.1 Event Studies


I start with new technological classes. CPC classes provide a classification of the technological content of patents, which allows me to examine whether interlocking directorates broaden the range of partner technologies they draw from for their own inventions. If board connections do indeed decrease informational barriers to accessing a partner's knowledge, as [Cohen and Levinthal \(1989\)](#) and [Stein \(2008\)](#) may for example suggest, then one possible consequence is greater exploration of that partner's technological portfolio. Examining CPC classes therefore complements citation counts by revealing the technological breadth of the knowledge being cited. This captures one potential margin of learning, while recognizing that firms may also learn within familiar technological areas or innovate by recombining existing ideas, as emphasized by [Weitzman \(1998\)](#).

Therefore, for each unordered dyad $\{f, j\}$, I track the history of partner patents and partner technology classes cited by the pair up to year $t - 1$. I say that a citation in year t is classified as a new partner class citation if the cited patent belongs to at least one Cooperative Patent Classification (CPC) technology group that had not previously appeared among the dyad's cross-citations. One may be interested in analyzing these measures to understand whether interlocks facilitate exploration of previously unvisited parts of the partner's knowledge base. [Figure 4](#) shows the results using this outcome measure. The results are flat: after two firms become interlocked for the first time, their patent citations do not start moving into new parts of the partner's technological space. The additional citation activity appears to remain within technology areas that were already familiar to the pair.

Next, I separate technologically proximate citations from non-technologically proximate ones. A citation is defined as technologically close when the citing patent and the cited patent share at least one technology group from the CPC classification system. This gives me a dyad-year share of citations to technologically close partner patents. I then compute the share of dyad-year citations directed to technologically and non-technologically close partner patents. Technological proximity allows me to analyze how closely the partner knowledge being cited relates to the technology embodied in the citing

invention, which may be useful for several reasons. For instance, citations to technologically proximate partner patents may reflect firms building on knowledge closely related to their own inventions and drawing from it, while citations to more distant patents capture connections across technological domains, which may possibly indicate a different use of the partner's portfolio. It is important to highlight again that technological learning does not occur exclusively between similar technologies, since combining domains may itself generate innovation (Weitzman, 1998), but this exercise is still useful in helping to characterize the technological relevance of increased cross citations between two interlocked partners.

Figure 4: Event Study - Citations of New CPC Classes in Partner Patents



Notes: This figure reports staggered-adoption event study estimates of the effect of first adopting an interlock on the stated outcome. The reported estimates are from the Callaway and Sant'Anna (2021) framework, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. Event time is defined relative to the first interlock ($k = 0$) between any two firms. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level.

Figure 5 shows the results using these outcomes. In particular, Figure 5a tells us that, after first interlock formation, the two firms do not start citing each other more within overlapping technological areas. Instead, as Figure 5b points out, they start citing each

other more outside their overlapping technology space.¹⁵ Finally, I look at the age of the cited partner patents as a way to distinguish attention to the partner’s recent inventions from attention to its established stock of patented knowledge. Caballero and Jaffe (1993) explicitly model these competing forces of diffusion and technological obsolescence, and Mehta, Rysman and Simcoe (2010) examine how they are reflected in patent citation age profiles. If board connections accelerate access to a partner’s recent technological developments, one possible consequence is greater attention to newer partner patents. Increased citations to older patents would instead indicate greater use or scrutiny of knowledge that has been publicly available for longer. Here, I define older partner patents as those issued at least ten years before the recorded citation, and newer patents as those cited in the ten prior years. Figure 6 shows again that board interlocks are associated with an increase in citations to older partner patents.


5.2.2 Continuing Patents

Inventors may file continuing applications, that is, patent child applications that extend the examination of the parent application. Continuations (CONs) pursue claims based on the same disclosure as the parent, continuations-in-part (CIPs) add new claims, and divisionals (DIVs) separate inventions that were jointly disclosed in the original application. By allowing applicants to postpone or revise claims, continuing applications help applicants deal with legal, technological, and commercial uncertainty. Nevertheless, this flexibility may also promote strategic patenting behavior and has been associated with greater litigation and overlapping patent rights that can impede competitor firms (Berger, Blind and Thumm, 2012, Hegde, Mowery and Graham, 2009, Righi and Simcoe, 2023, Righi, Cannito and Vladasel, 2023).

CIPs are the continuation type most closely associated with additional inventive ac-


¹⁵It is important to distinguish this exercise from the one reported in Figure 4, as the two sets of outcomes capture different margins. For instance, a citation can fall outside the dyad’s overlapping technology space while still pointing to a partner technology class that the dyad had cited previously. Suppose that, before year t , the dyad had already cited partner patents in CPC class H01, which covers electric elements. If, in year t , one firm cites an H01 partner patent from one of its own patents classified in G06, which covers computing and calculating technologies, the citation would not be technologically close, even though H01 is not a newly cited partner class. The patterns in Figures 4 and 5 are therefore consistent, as interlocks may increase citations across technologically distant areas without necessarily expanding the set of partner technology classes that firms cite.

Figure 5: Event Studies - Technological Proximity



Notes: Each panel reports staggered-adoption event study estimates of the effect of first adopting an interlock on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) framework, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. Event time is defined relative to the first interlock ($k = 0$) between any two firms. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Logged outcomes use $\log(1+x)$. All specifications include calendar-year fixed effects and the baseline dyad-level controls as described in the text.

Figure 6: Event Studies - Patent Age



Notes: Each panel reports staggered-adoption event study estimates of the effect of first adopting an interlock on the stated outcome. I use the [Callaway and Sant’Anna \(2021\)](#) estimator, mapping cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Logged outcomes use $\log(1 + x)$. A patent is “old” if it is at least ten years old at the time of citation, and “new” otherwise. All specifications include calendar-year fixed effects and the baseline dyad-level controls described in the text.

tivity, as patents in this category necessarily must contain additional technical disclosure relative to their parent application (see [USPTO \(2024\)](#) §§ 201.06–201.08). For instance, [Hegde, Mowery and Graham \(2009\)](#) show that CIPs are disproportionately used by R&D-intensive firms, particularly in chemical and biotechnological technologies. Of course, one should be careful in stating that CIPs represent innovations. A CIP establishes that the applicant has added new matter to the parent application, but that does not indicate that such new matter represents a major technological advance. Nevertheless, among the continuation types, CIPs offer the strongest signal of incremental innovation with respect to the parent application.

On the other hand, in a CON application, the applicant is legally prohibited from adding new matter ([USPTO, 2024](#)). In fact, the USPTO defines it as another application for inventions already disclosed in the parent. What can change through a continuation application is the set of claims, and therefore the legal scope of protection sought over that pre-existing disclosure. This is the reason why recent empirical work in Economics views CONs through the lens of strategic patenting behavior ([Righi and Simcoe, 2023](#), [Righi,](#)

Cannito and Vladasel, 2023). Recent work also documents that continuation applications are followed by a 37%-47% increase in the probability of litigation, with stronger effects when the continuation's claims differ substantially from those of its parents (Righi, 2026).

Finally, DIV applications are somewhat ambiguous for several reasons. To begin, DIVs cannot add new matter, and they only carve an independent or distinct invention out of subject matter that was already disclosed in the parent. But, unlike CONs, DIVs are often generated by the USPTO examination process itself. This happens when the examiner tells the applicant that the application contains multiple independent or distinct inventions and that they must be prosecuted separately. DIVs cannot, therefore, be interpreted as evidence of active strategic patenting behavior. Nevertheless, Hegde, Mowery and Graham (2009) report that both CON and DIV applications tend to be associated with less important patents and interpret their use as consistent with strategic patenting behavior. Both because of this reason and because of transparency, I also report the results concerning DIVs alongside CONs and CIPs in the regressions below.¹⁶

Therefore, one may want to look at whether the composition of these patents changes. For example, does the number of CONs, CIPs, and DIVs as a share of continuing patents change differently for interlocked firms? Table 5 illustrates the answers to this question. First, Columns (1) and (2) show that interlocked firms file more continuing patents applications. Then, columns (3) and (5) show that, after absorbing firm, industry-year, and state-year effects, the share of CONs is 6.3 percentage points higher for interlocked firms, while the share of CIPs is 4.5 percentage points lower. With only firm and year fixed effects, the CON share remains positive and statistically significant at the 10% level in column (4), while the CIP estimate is close to zero in column (6). The DIV estimates are negative in columns (7) and (8), but only the latter is statistically significant at the 10% level. According to the commonly held views about these types of continuing patents in the literature, therefore, it seems that the patenting behaviour of interlocked companies is a little more consistent with strategic patenting compared to their non-interlocked counterparts.

¹⁶Note that these categories are not mutually exclusive: while the vast majority of continuing applications are associated with only one type, some continuing applications may be linked to more than one continuing application type.

Table 5: Composition of Continuing Applications and Section 8 Interlock Exposure

	(1) Log (Continuing)	(2) Log (Continuing)	(3) CON / Continuing	(4) CON / Continuing	(5) CIP / Continuing	(6) CIP / Continuing	(7) DIV / Continuing	(8) DIV / Continuing
Interlock	0.072* (0.038)	0.095** (0.039)	0.063*** (0.023)	0.036* (0.021)	-0.045* (0.025)	-0.000 (0.021)	-0.035 (0.023)	-0.039* (0.021)
N	4,300	4,542	1,896	2,171	1,943	2,220	1,941	2,217
R-squared	0.901	0.880	0.635	0.510	0.570	0.436	0.622	0.480
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one Section 8-violating horizontal interlock in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the firm's total director fixed effects, and the corresponding sets of fixed effects as reported in the table. The two Jaccard technology and product-market spillover covariates are forced into every specification on top of the LASSO selection.

6 Robustness to Alternative Specifications

I assess the robustness of my findings with a battery of robustness exercises reported in Appendix A.5. To begin, I report the results that I obtain if, instead of running OLS, I use a Poisson specification. This would be my preferred baseline specification if it were not for the fact that sensitivity analyses for Poisson regressions are not available, and thus cannot assess robustness to omitted variable bias the same way I do with OLS. In any case, Table A.5.1 shows that the Poisson specification fully agrees with OLS.

In Appendix A.5.2, I report the results I obtain when I use the average AKM-estimated ability in a certain firm as a control for board member ability as opposed to using the total firm-level sum of the estimated fixed effects, as I do in my baseline. In other words, in this alternative specification, I take the firm-level sum of the estimated AKM fixed effects, I divide them by the number of board members at the firm, and I use this figure as a control variable. Table A.5.2 reports these results, which confirm the estimates we obtain in the baseline table of results.

One may also be interested in seeing how these results change if I were to add common ownership as a control variable in my regressions. Indeed, if a firm shares common owners with its competitors on top of shared board members, one may for example presume that these results are actually driven by common owners rather than shared directors. I

test this hypothesis in Appendix A.5.3 by adding average levels of common ownership at the firm level. I report the results in Table A.5.3 and find that the results are unchanged.

Next, I try broadening the treatment definition from Section 8—violating horizontal interlocks to all interlocks, including vertical and other non-horizontal board links, which assesses whether the results are tied exclusively to the legal screen emphasized in the main analysis or instead reflect a broader relationship between board connectedness and innovation. These analyses can be found in Table A.5.4. The results for the raw patent counts are confirmed, although statistical significance grows thinner since, in this alternative specification, the vast majority of firms in the sample become interlocked and there is less variation provided by the control group. Finally, in Appendix A.5.5 I show the results I obtain when I broaden my definition of competitors using the Hoberg–Phillips text-based network industry classification (Hoberg and Phillips, 2010, 2016) and find similar results.

In Appendix A.7, I assess the robustness of my event studies by reporting their placebo versions. In each event study, I assign each treated pair a fake treatment date and then re-estimate the same staggered adoption design. I find that these placebo event studies do not reproduce the post-treatment dynamic effects illustrated in the main text. Thus, this strengthens the interpretation that the observed increases in cross citations and patents is tied to actual interlock formation, and that the event study results are not generated by generic pre-existing trends around eventually treated firms.

7 Conclusion

This paper studies whether shared board members between competitors reshape firm-level innovation and, if so, through which channels. I find that firms engaged in this practice produce more patents, including citation-weighted patent counts, but do not display corresponding increases in R&D expenditure or market-based patent values. Therefore, interlocks do not primarily operate by inducing firms to spend more on formal research, but they appear to alter how firms convert existing innovative effort into observable patent output. The sensitivity analyses indicate that these results are very robust to residual omitted variable bias and that overturning the baseline patenting estimates would require unobservables to be implausibly strong relative to the rich set of observ-

ables already included in the specifications.

Event studies show that when two firms interlock for the first time, they begin citing each other more often in their patents. Nevertheless, the composition of those citations is not consistent with a story of technological learning or exploratory knowledge recombination: post-interlock cross citations are not disproportionately directed toward newer partner patents or technologically proximate inventions, as one would probably expect under technological learning. Instead, the stronger responses are concentrated in citations to partner patents that are older and not technologically close to the citing patent. In other words, interlocks do seem to transmit information across connected competitors, but the evidence presented in this paper suggests that this information is used primarily to sharpen firms' awareness of each other's patent boundaries and to support more strategic or defensive patenting behavior rather than to facilitate transmission of frontier technical knowledge.

The results concerning continuing patent applications reinforce this interpretation. On average, conditional on applying for a continuing patent, I find that interlocked firms shift the composition of these continuing patents more toward continuation patents – which are commonly associated with strategic patenting behavior–, while continuation-in-part patents –which usually reflect true innovation, as they add new technological specifications to the original claims– decrease as a share.

These results may carry some welfare and policy implications. If interlocks promoted knowledge diffusion and transmission of know-how between firms, then an increase in patenting could be interpreted as at least partly innovation-enhancing. On the other hand, the results presented here provide evidence that the increase in post-interlock patent activity mainly reflects strategic patenting, which may for instance raise legal congestion in overlapping technology space. From an antitrust perspective, this may imply that the costs of interlocking directorates extend beyond the standard concerns about softened product-market competition or explicit coordination ([Herrera-Caicedo, Jeffers and Prager, 2026](#)): shared board members may also alter how firms defend incumbent positions, monitor rivals' intellectual-property portfolios, and shape the legal environment in which innovation happens in the first place. With that being said, to calculate the net welfare effects of this practice we would also need to take into account the effects on the consumers, which is beyond the scope of this paper.

The findings presented here reinforce the traditional view that interlocking directorates are organizational channels through which firms can reallocate effort and adjust strategic behavior (Allen, 1974, Bunting, 1976, Burt, 1983, Pfeffer, 1972, Pfeffer and Salancik, 1978, Stearns and Mizruchi, 1993) and offer a new mechanism through which this practice can alter innovation outcomes. This perspective also opens several avenues for future work. One of them, as I mentioned above, would be to undertake an assessment of the net welfare effect of this practice, which would need to take the effects on consumers into account in a structural model. Another would be to embed interlock formation in a structural network framework that permits counterfactual policy analysis of alternative enforcement regimes. Finally, future research could analyze whether interlocked firms experience lower litigation risks as a result of their new patenting behavior.

Understanding the effects of interlocking directorates requires looking not only at prices, output, and market structure, but also at how board connections reshape the legal and strategic margins of innovation. There is still a lot we do not understand about interlocking directorates, but the possibilities for future research seem ripe and promising.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis.** 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333.
- Allen, Michael Patrick.** 1974. "The Structure of Interorganizational Elite Cooptation: Interlocking Corporate Directorates." *American Sociological Review*, 39(3): 393–406.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber.** 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy*, 113(1): 151–184.
- Andrews, M. J., L. Gill, T. Schank, and R. Upward.** 2008. "High wage workers and low wage firms: negative assortative matching or limited mobility bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 171(3): 673–697.
- Andrews, M.J., L. Gill, T. Schank, and R. Upward.** 2012. "High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias." *Economics Letters*, 117(3): 824–827.
- Antón, Miguel, Florian Ederer, Mireia Giné, and Martin Schmalz.** 2023a. "Common Ownership, Competition, and Top Management Incentives." *Journal of Political Economy*, 131(5): 1294–1355.
- Antón, Miguel, Florian Ederer, Mireia Giné, and Bruno Pellegrino.** 2023b. "Mergers and Acquisitions under Common Ownership." *AEA Papers and Proceedings*, 113: 294–98.
- Antón, Miguel, Florian Ederer, Mireia Giné, and Martin Schmalz.** 2025. "Innovation: The Bright Side of Common Ownership?" *Management Science*, 71(5): 3713–3733.
- Azar, José, Martin C. Schmalz, and Isabel Tecu.** 2018. "Anticompetitive Effects of Common Ownership." *The Journal of Finance*, 73(4): 1513–1565.
- Azar, José.** 2022. "Common Shareholders and Interlocking Directors: The Relation Between Two Corporate Networks." *Journal of Competition Law & Economics*, 18(1): 75–98.

- Baltrunaite, Audinga, and Egle Karmaziene.** 2024. “Fast tracks to boardrooms: Director supply and board appointments.” *Journal of Corporate Finance*, 88: 102642.
- Barone, Guglielmo, Fabiano Schivardi, and Enrico Sette.** 2025. “Interlocking Directorates and Competition in Banking.” *The Journal of Finance*, 80(4): 1963–2016.
- Battaggion, Maria Rosa, and Vittoria Cerasi.** 2020. “Strategic interlocking directorates.” *Journal of Economic Behavior & Organization*, 178: 85–101.
- Berger, Florian, Knut Blind, and Nikolaus Thumm.** 2012. “Filing behaviour regarding essential patents in industry standards.” *Research Policy*, 41(1): 216–225.
- Bertrand, Marianne, and Antoinette Schoar.** 2003. “Managing with Style: The Effect of Managers on Firm Policies*.” *The Quarterly Journal of Economics*, 118(4): 1169–1208.
- Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler.** 2023. “How Much Should We Trust Estimates of Firm Effects and Worker Sorting?” *Journal of Labor Economics*, 41(2): 291–322.
- Bunting, David.** 1976. “Corporate Interlocking. Part III. Interlocks and Return on Investment.” *Directors & Boards (fall)*:4-11.
- Burt, Ronald S.** 1983. *Corporate Profits and Cooptation : Networks of Market Constraints and Directorate Ties in the American Economy*. New York: Academic Press.
- Caballero, Ricardo J., and Adam B. Jaffe.** 1993. “How High Are the Giants’ Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth.” *NBER Macroeconomics Annual*, 8: 15–74.
- Cabazon, Felipe, and Gerard Hoberg.** 2026. “Leaky Director Networks and Innovation Herding.” *The Review of Financial Studies*, 39(1): 158–197.
- Callaway, Brantly, and Pedro H.C. Sant’Anna.** 2021. “Difference-in-Differences with multiple time periods.” *Journal of Econometrics*, 225(2): 200–230.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline.** 2018. “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics*, 36(S1): S13 – S70.

- Card, David, Jörg Heining, and Patrick Kline.** 2013. "Workplace heterogeneity and the rise of West German wage inequality." *The Quarterly Journal of Economics*, 128(3): 967–1015.
- Cavaco, Sandra, Patricia Crifo, Antoine Reberlioux, and Gwenael Roudaut.** 2017. "Independent directors: Less informed but better selected than affiliated board members?" *Journal of Corporate Finance*, 43: 106–121.
- Cinelli, Carlos, and Chad Hazlett.** 2020. "Making Sense of Sensitivity: Extending Omitted Variable Bias." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1): 39–67.
- Cohen, Wesley M., and Daniel A. Levinthal.** 1989. "Innovation and Learning: The Two Faces of R&D." *The Economic Journal*, 99(397): 569–596.
- Cohen, Wesley M., Richard R. Nelson, and John P. Walsh.** 2000. "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)." National Bureau of Economic Research, Inc NBER Working Papers 7552.
- De Loecker, Jan, and Frederic Warzynski.** 2012. "Markups and Firm-Level Export Status." *American Economic Review*, 102(6): 2437–71.
- Diegert, Paul, Matthew A. Masten, and Alexandre Poirier.** 2025. "Assessing Omitted Variable Bias when the Controls are Endogenous."
- Dooley, Peter C.** 1969. "The Interlocking Directorate." *The American Economic Review*, 59(3): 314–323.
- Ederer, Florian, and Bruno Pellegrino.** 2026. "A Tale of Two Networks: Common Ownership and Product Market Rivalry." *The Review of Economic Studies*, 93(3): 1746–1788.
- Eeckhout, Jan, and Philipp Kircher.** 2011. "Identifying Sorting-In Theory." *The Review of Economic Studies*, 78(3): 872–906.
- Eldar, Ofer, and Jillian Grennan.** 2021. "Common Ownership and Entrepreneurship." *AEA Papers and Proceedings*, 111: 582–86.
- Federico, Giulio, Gregor Langus, and Tommaso Valletti.** 2017. "A simple model of mergers and innovation." *Economics Letters*, 157(C): 136–140.

- Federico, Giulio, Gregor Langus, and Tommaso Valletti.** 2018. "Horizontal mergers and product innovation." *International Journal of Industrial Organization*, 59: 1–23.
- Ferris, Stephen P., Murali Jagannathan, and A. C. Pritchard.** 2003. "Too Busy to Mind the Business? Monitoring by Directors with Multiple Board Appointments." *The Journal of Finance*, 58(3): 1087–1111.
- Graham, John R., Si Li, and Jiaping Qiu.** 2012. "Managerial Attributes and Executive Compensation." *The Review of Financial Studies*, 25(1): 144–186.
- Gualdani, Cristina.** 2021. "An econometric model of network formation with an application to board interlocks between firms." *Journal of Econometrics*, 224(2): 345–370.
- Guiso, Luigi, Fabiano Schivardi, and Luana Zaccaria.** 2026. "Changing the board game: Horizontal spillovers of gender quotas." *Journal of Financial Economics*, 183: 104326.
- Hall, Bronwyn H, Adam Jaffe, and Manuel Trajtenberg.** 2001. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." *NBER Working Paper*, , (8498).
- Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg.** 2005. "Market Value and Patent Citations." *The RAND Journal of Economics*, 36(1): 16–38.
- Hall, Bronwyn H., and Rosemarie Ham Ziedonis.** 2001. "The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995." *The RAND Journal of Economics*, 32(1): 101–128.
- Hegde, Deepak, David C. Mowery, and Stuart J. H. Graham.** 2009. "Pioneering Inventors or Thicket Builders: Which U.S. Firms Use Continuations in Patenting?" *Management Science*, 55(7): 1214–1226.
- Herrera-Caicedo, Alejandro, Jessica Jeffers, and Elena Prager.** 2026. "Collusion through Common Leadership." *Journal of Political Economy*, 0(ja).
- Hoberg, Gerard, and Gordon Phillips.** 2010. "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis." *The Review of Financial Studies*, 23(10): 3773–3811.

- Hoberg, Gerard, and Gordon Phillips.** 2016. "Text-Based Network Industries and Endogenous Product Differentiation." *Journal of Political Economy*, 124(5): 1423–1465.
- Kline, Patrick.** 2024. "Firm wage effects."
- Kline, Patrick, Raffaele Saggio, and Mikkel Sølvssten.** 2020. "Leave-Out Estimation of Variance Components." *Econometrica*, 88(5): 1859–1898.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman.** 2017. "Technological Innovation, Resource Allocation, and Growth." *The Quarterly Journal of Economics*, 132(2): 665–712.
- Masten, Matthew A., and Alexandre Poirier.** 2026. "The Effect of Omitted Variables on the Sign of Regression Coefficients." *American Economic Review*, 116(7): 2685–2710.
- Masulis, Ronald W., and Shawn Mobbs.** 2014. "Independent director incentives: Where do talented directors spend their limited time and energy?" *Journal of Financial Economics*, 111(2): 406–429.
- Mehta, Aditi, Marc Rysman, and Tim Simcoe.** 2010. "Identifying the age profile of patent citations: new estimates of knowledge diffusion." *Journal of Applied Econometrics*, 25(7): 1179–1204.
- Mills, C. Wright.** 1956. *The Power Elite*. New York: Oxford University Press.
- Mizruchi, Mark S.** 1996. "What Do Interlocks Do? An Analysis, Critique, and Assessment of Research on Interlocking Directorates." *Annual Review of Sociology*, 22: 271–298.
- Nili, Yaron.** 2022. "Horizontal Directors Revisited." *Journal of Competition Law and Economics*, 18(1): 5–28.
- Oster, Emily.** 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business & Economic Statistics*, 37(2): 187–204.
- Palmer, Donald, Roger Friedland, and Jitendra Singh.** 1986. "The Ties That Bind: Organizational and Class Bases of Stability in a Corporate Interlock Network." *American Sociological Review*, 51.

- Pennings, J.M.** 1980. *Interlocking directorates: Origins and consequences of connections among organizations' Board of Directors. The Jossey-Bass social and behavioral science series*, Jossey-Bass.
- Pfeffer, Jeffrey.** 1972. "Size and Composition of Corporate Boards of Directors: The Organization and its Environment." *Administrative Science Quarterly*, 17(2): 218–228.
- Pfeffer, Jeffrey, and Gerald R. Salancik.** 1978. *The External Control of Organizations: A Resource Dependence Perspective*. New York: Harper & Row.
- Poberejsky, Roma.** 2024. "Interlocking Directorates, Competition, and Innovation." *Working paper*.
- Righi, Cesare.** 2026. "Continuation Patents and Litigation." Available at SSRN: <https://ssrn.com/abstract=4027998> or <http://dx.doi.org/10.2139/ssrn.4027998>.
- Righi, Cesare, and Timothy Simcoe.** 2023. "Patenting inventions or inventing patents? Continuation practice at the USPTO." *The RAND Journal of Economics*, 54(3): 416–442.
- Righi, Cesare, Davide Cannito, and Theodor Vladasel.** 2023. "Continuing patent applications at the USPTO." *Research Policy*, 52(4): 104742.
- Selznick, Philip.** 1949. *TVA and the Grass Roots; A Study in the Sociology of Formal Organization*. Berkeley and Los Angeles: University of California Press.
- Stearns, Linda Brewster, and Mark S. Mizruchi.** 1993. "Board Composition and Corporate Financing: The Impact of Financial Institution Representation on Borrowing." *The Academy of Management Journal*, 36(3): 603–618.
- Stein, Jeremy C.** 2008. "Conversations among Competitors." *American Economic Review*, 98(5): 2150–62.
- Stokman, Frans N., Jelle Van Der Knoop, and Frans W. Wasseur.** 1988. "Interlocks in the Netherlands: Stability and careers in the period 1960–1980." *Social Networks*, 10(2): 183–208.

- Thompson, James D., and William J. McEwen.** 1958. "Organizational Goals and Environment: Goal-Setting as an Interaction Process." *American Sociological Review*, 23(1): 23–31.
- USPTO.** 2024. "Manual of Patent Examining Procedure." Ninth Edition, Revision 01.2024 ed., U.S. Department of Commerce.
- Weitzman, Martin L.** 1998. "Recombinant Growth." *The Quarterly Journal of Economics*, 113(2): 331–360.
- Yermack, David.** 2004. "Remuneration, Retention, and Reputation Incentives for Outside Directors." *The Journal of Finance*, 59(5): 2281–2308.
- Zajac, Edward J.** 1988. "Interlocking Directorates as an Interorganizational Strategy: A Test of Critical Assumptions." *Academy of Management Journal*, 31: 428–438.

Appendix

A.1 Section 8 (Clayton Act) Classification of Horizontal Interlocks

This appendix describes, step by step, how I estimate Section 8 illegality for horizontal interlocking directorates. I aim to construct (i) a pair-year indicator for whether an observed competitor interlock is Section 8-violating and (ii) firm-year aggregates that count the number of such violations.

First, the Section 8 screen from the Clayton Act is applied only to interlocks between firms that are classified as competitors in a given year, operationalized as belonging to the same industry. Vertical interlocks are not illegal by construction. Further, I implement the absolute de minimis exemption on competitive sales using the FTC’s annual dollar threshold. Section 8 contains exclusions that are specific to certain entities in some settings, and in practice financial institutions are treated differently. In my implementation, I therefore exclude banks and trust entities by setting $\text{Excluded}_f = 1$.

Then, for each year t , I use the FTC’s annual Section 8 threshold table containing (i) the CSUP (capital, surplus, and undivided profits) coverage threshold, denoted $\overline{\text{CSUP}}_t$, and (ii) the absolute de minimis dollar threshold on competitive sales, denoted $\overline{S}_t$. Figure A.1.1 plots these two threshold series over time (in millions of US dollars) for the years I use in my regressions.

I construct a firm-year proxy for CSUP using Compustat data. Let $\widehat{\text{CSUP}}_{it}$ denote this proxy for firm i in year t . This is defined using a priority rule: (i) use stockholders’ equity if available, else (ii) common equity plus preferred stock when available, else (iii) total assets minus total liabilities. Then, let S_{it} denote firm-year net sales. I define the two threshold indicators:

$$\text{Covered}_{it} = \mathbb{1}\{\widehat{\text{CSUP}}_{it} \geq \overline{\text{CSUP}}_t\}, \quad \text{ExemptAbs}_{it} = \mathbb{1}\{S_{it} < \overline{S}_t\}.$$

A firm-year is eligible for the interlocking directorate restrictions according to Section 8

of the Clayton Act if it is covered and not exempt under the absolute sales cutoff, that is,

$$\text{Eligible}_{it} = \mathbb{1}\{\text{Covered}_{it} = 1 \wedge \text{ExemptAbs}_{it} = 0\}.$$

This is defined only for the firms' observations for which the underlying accounting inputs and thresholds are observed.

Let $\text{Interlock}_{fjt} \in \{0, 1\}$ now indicate whether firms f and j share at least one board member in year t . Consider a horizontal competitor pair-year. The interlock is classified as Section 8-violating if an interlock exists, both firms are eligible, and neither firm is excluded:

$$\text{Sec8Violate}_{fjt} = \mathbb{1}\{\text{Interlock}_{fjt} = 1\} \cdot \mathbb{1}\{\text{Eligible}_{ft} = 1\} \cdot \mathbb{1}\{\text{Eligible}_{jt} = 1\} \quad (\text{A.1.1})$$

$$\cdot \mathbb{1}\{\text{Excluded}_f = 0\} \cdot \mathbb{1}\{\text{Excluded}_j = 0\}, \quad (\text{A.1.2})$$


where Sec8Violate_{fjt} is only defined on the restricted sample of horizontal pairs.

To obtain firm-year counts, I first sum pair-year violations across all competitor pairs before collapsing to a unique firm-year:

$$\text{Sec8Violate}_{jt} = \sum_{j \neq f} \text{Sec8Violate}_{fjt}.$$

This gives me a firm-year panel of Section 8-violating horizontal interlocks. Then, in my baseline regressions, I use an indicator variable Interlock_{jt} that is equal to one if firm j has at least one horizontal interlock at time t , that is, $\text{Sec8Violate}_{jt} \geq 1$.

Figure A.1.1: FTC Annual Section 8 Coverage Thresholds Used in the Section 8 Screen



Notes: Panel (a) plots the FTC's annual threshold for "capital, surplus, and undivided profits" (CSUP), and Panel (b) plots the FTC's annual absolute de-minimis "competitive sales" threshold. Both are shown in millions of dollars for the years used in the implementation window (2000–2024).

A.2 Empirical Facts About Interlocking Directorates

In this section, I report three empirical facts about interlocked directors, firms' likelihood to interlock, and the average persistence of interlocks.

Fact #1: Interlocked directors are more likely to be CEOs or Presidents. Figure [A.2.1](#) summarizes how interlocked directors differ from non-interlocked directors along a range of observable characteristics. Each point in the graph is an estimate that comes from a separate regression from the others, and it reflects both direct selection into interlocks and sorting across firms with different baseline propensities to interlock. I document that interlock status at the director-level is positively correlated with corporate authority. For instance, CEOs and Presidents are more likely to hold at least one concurrent board seat at another firm. The same graph also shows that interlocked directors have slightly higher compensation and larger professional networks. This is consistent with the idea that interlocked individuals are more productive, more in demand across firms, more influential, and therefore more likely to sit on the boards of multiple firms.

Fact #2: Firms prefer to interlock with similar firms. Table [A.2.2](#) reports coefficients from logit regressions and characterizes interlock formation at the firm-pair level. The propensity to form an interlock between any two firms can be strongly predicted by proximity and similarity. For instance, geographic distance is negative and statistically significant across specifications, which implies that more distant dyads have lower odds of interlock formation ([Azar, 2022](#)). I find that being in the same industry also matters, because the results illustrate that interlocks form disproportionately within sectors: being close in the product market space predicts interlock formation, as a high Hoberg–Phillips ([Hoberg and Phillips, 2010, 2016](#)) similarity indicator enters positively and is statistically significant. In the more restrictive specifications, several differences in board composition enter with negative point estimates, including a statistically significant negative coefficient on differences in gender composition, which suggests that interlocks form more often between firms whose boards resemble each other from a gender perspective. This is again consistent with a view of interlock formation as a non-random matching process in which firms connect to geographically proximate and compositionally similar companies.

Fact #3: Horizontal interlocks are typically persistent. Finally, in Figure [A.2.2](#), I take a look at the average length of a horizontal interlock. In particular, I report the


Table A.2.1: Director-Level Summary Statistics, 2000-2024

	(1) All	(2) Independent	(3) Executive
<i>Panel A: Director-Level Statistics</i>			
Tenure (years)	5.98	6.01	5.56
Age at Appointment	60.27	60.42	58.41
Age at Exit	65.25	65.43	62.97
Female (%)	17.23	18.19	5.05
<i>Panel B: Board Connectivity and Committees</i>			
Concurrent Seats (avg.)	1.80	1.82	1.44
Committee Roles (% any of A/C/N)	91.92	95.76	43.17
N (director-firm spells)	154,871	143,555	11,316
N (unique directors)	94,461	86,184	8,277

Notes: Columns in Panel A report means across director-firm spells (a unique director-board appointment), while columns in Panel B report means across the panel. “Independent” denotes non-executive directors, while “Executive” denotes inside directors. *Tenure (years)* is the length of the employment spell (last minus first appointment year + 1). *Age at Appointment* and *Age at Exit* are computed at the first and last year of the appointment spell, respectively. *Female (%)* is the share of spells held by female directors. *Concurrent Seats (avg.)* is the mean number of simultaneous board seats the director holds in a given year within a spell. *Committee Roles (% any of A/C/N)* illustrates the share of spells in which the director served on at least one Audit, Compensation/Remuneration, or Nomination/Governance committee at any point during a spell. *N (director-firm spells)* counts unique spells; *N (unique directors)* counts distinct directors. All the variables are constructed from BoardEx.

observed spell length in years of this practice. I document that around 30% of horizontal interlocks are observed for only one year, although 68.3% of those end in the dyad's last observed year, so it is unclear how many actually dissolved. Around 52% last three or more years. Roughly 1 in 10 lasts at least eight years. Therefore, when two boards connect through an interlock, they usually do so for an extended period of time, either with just one interlocked director or several interlocked board members.

Figure A.2.1: Director-Level Correlates of Being Interlocked




Notes: Each point is the OLS coefficient from a separate regression. Each regression uses an indicator for whether a director is interlocked as the outcome variable, and a single director-level characteristic as the explanatory variable. I use the total concurrent number of boards a director works for and year fixed effects as control variables. Formally, for each director d and characteristic k , the regression I estimate is $\mathbb{1}\{\text{Interlocked}_d\} = \alpha_k + \beta_k X_{d,k} + \gamma_t + \omega \log(\text{Number of Boards}) + \varepsilon_{d,k}$. The plotted β_k therefore summarizes the correlation between characteristic $X_{d,k}$ and interlocked status. The whiskers around the dots report the corresponding 95% confidence intervals.

Table A.2.2: Interlock Formation Between Firm Pairs

	(1) Prob(Interlock)	(2) Prob(Interlock)	(3) Prob(Interlock)	(4) Prob(Interlock)
Distance	-0.327*** (0.017)	-0.308*** (0.018)	-0.449*** (0.111)	-0.452*** (0.112)
Size Difference	-0.178*** (0.031)	-0.086*** (0.031)	-0.206** (0.104)	-0.029 (0.131)
R&D Difference	-0.225*** (0.022)	-0.213*** (0.022)	-0.064 (0.068)	-0.020 (0.070)
High Product Similarity	1.451*** (0.197)	1.250*** (0.186)	0.522 (1.203)	0.399 (1.107)
Same Industry		0.920*** (0.051)	0.572 (0.556)	0.546 (0.570)
Nationality Difference			-0.437 (0.454)	-0.465 (0.452)
Gender Difference			-2.012** (0.978)	-2.012** (0.967)
Busy Directors Difference			-0.504 (0.592)	-0.486 (0.594)
Network Size Difference			-0.000*** (0.000)	-0.000*** (0.000)
Time Spent on Board Difference			-0.026 (0.029)	-0.024 (0.029)
Revenue Difference				-0.016 (0.106)
COGS Difference				-0.225** (0.102)
Constant	-17.332*** (0.143)	-18.370*** (0.139)	-13.874*** (2.140)	-13.754*** (1.980)
N	2,205,283	2,205,283	151,408	151,401
Year FE	YES	YES	YES	YES

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Entries are logit coefficients, with standard errors in parentheses. All specifications include year fixed effects and two-way clustered standard errors. *High Product Similarity* is an indicator equal to one if the Hoberg–Phillips product-market similarity score for the firm pair falls in the top percentile of the corresponding year-specific distribution. All variables labeled *Difference* are constructed as absolute differences between the two firms in the dyad. *Size Difference* is the absolute difference in the number of employees. *Nationality Difference* is the absolute difference in the nationality composition of the two boards, as reported by BoardEx. *Gender Difference* is the absolute difference in the share of female directors across the two boards. *Busy Directors Difference* is the absolute difference in the share of busy directors across the two firms. *Network Size Difference* is the absolute difference in average director network size across the two boards. *Time Spent on Board Difference* is the absolute difference in average board tenure across the two firms, based on BoardEx records.

Figure A.2.2: Spell Lengths of Horizontal Interlocking Directorates




Notes: This line graph reports the observed spell length (in years) for each horizontal interlocking directorate in the data, with the exclusion of those formed before the beginning of the panel. A spell length is defined as the difference between the year in which two firms first interlock and the year in which they dissolve that link. For each spell length, I report the percentage of interlocks that last for at least that period of time.

A.3 AKM Model

A.3.1 AKM Model Diagnostics

Figure A.3.1: Diagnostics of the AKM Model



Notes: Panel A splits the estimated AKM board member and firm fixed effects into deciles and then, for each of the one hundred cells, reports their mean residuals (in % of log compensation). Panel B presents an event-study mobility test that categorizes companies by quartiles of their mean log wages and then classifies each job change by the quartile of the origin and destination workplaces. Both tests are from Card, Heining and Kline (2013).


In this Appendix section, I report common tests that have been developed to test whether the assumptions of the baseline AKM model are respected. To begin, Figure A.3.1 reports graphical tests from Card, Heining and Kline (2013). In particular, Figure A.3.1a evaluates the additively separable assumption of the canonical AKM model to exclude match effects between the workers and the firms. In this test, both the board member and the firm fixed effects are split into deciles. Then, for each of the one hundred different cells, I compute the average size of the residuals as a percentage of the logarithm of the compensation for the board members in that cell. A flat distribution would be evidence against economically important match effects and is consistent with the additively separable structure of director and firm fixed effects. The results from Figure A.3.1a show that, in my estimated board member model, the distribution of these residuals is extremely flat and does not show evidence of systematic patterns across deciles.

Figure A.3.1b presents a mobility diagnostic for the same model. The figure follows mean log compensation in event time for directors who move between boards, grouping moves by the quartile of the mean coworker compensation at the origin and destination boards. The exogenous mobility assumption implies that, absent match-specific or transitory shocks driving moves across firms, compensation profiles should be stable before the transition and exhibit a discrete change at the move. The sign and magnitude of this change should also vary systematically with the direction of mobility. This is what Figure A.3.1b shows. Directors who move from low-pay to higher-pay companies experience upward shifts in compensation at the time of the move, with larger gains for moves to higher destination quartiles, whereas directors who move from high-pay boards to lower-pay boards display corresponding declines. Moves within the same quartile generate modest changes. These visual results are consistent with the interpretation that the estimated board component captures persistent differences in pay levels across boards.

A.3.2 Firm-Level Results

First, I investigate whether the firm-level AKM aggregates are systematically related to innovation outcomes. If the AKM director fixed effects capture a relevant dimension of board member ability, then firms with boards composed of directors with higher estimated fixed effects should also exhibit stronger innovative performance. Or, equivalently, better board members should select into better, more innovative firms. To evaluate whether this pattern holds in the data, I plot a set of binscatter relationships between the firm-level AKM ability measures and several innovation outcomes. Figure A.3.2 shows that the relationship is indeed strongly positive throughout: firms with higher estimated board ability spend more on R&D, produce more patents, and generate patents that are both more highly cited and more valuable in market terms.

Figure A.3.2: Correlations Between Firm-Level AKM Director Estimates and Outcomes



Notes: These binscatter plots show the correlation between firm-year aggregates of estimated board-member fixed effects and firm-year outcomes. The horizontal axis is partitioned into 100 quantile-spaced bins, and each blue dot reports the mean of the x-variable and the mean of the y-variable within a bin. The red line reports the corresponding linear fit.

A.4 Sensitivity Analyses

The modern literature on sensitivity analyses in Economics begins, to my knowledge, with [Altonji, Elder and Taber \(2005\)](#). Studying the effect of Catholic high school attendance on subsequent college outcomes, they develop a framework for assessing the importance of omitted variable bias in settings where identification does not rest on a clear exogenous shock or a valid instrument. Their idea is to compare selection on unobservables with selection on observables after conditioning on a rich set of covariates.

To formalize this idea, let $\hat{\beta}^S$ denote the coefficient on the interlock variable from a short regression of Y_{it} on the treatment only, and let $\hat{\beta}^M$ denote the coefficient from a medium regression that adds the observed controls W_{it} . Let R_S^2 and R_M^2 be the corresponding R^2 values. In [Altonji, Elder and Taber \(2005\)](#), the relevant statistic is the ratio

$$\text{AET} = \frac{\hat{\beta}^M}{\hat{\beta}^S - \hat{\beta}^M},$$

which compares the coefficient that remains after controlling for observables to the amount by which the coefficient shrinks once observables are added. A larger value means that selection on unobservables would have to be large, relative to selection on observables, to drive the coefficient of interest to zero. While this approach is useful and intuitive, it does not scale coefficient movements by the explanatory power of the added controls. [Oster \(2019\)](#) builds on and sharpens this idea by combining the movement in the treatment coefficient with the movement in R^2 . Under proportional selection, and given a hypothesized upper bound $R_{\max}$ for the R^2 of the infeasible long regression, Oster's bias-adjusted coefficient is

$$\beta^*(\delta, R_{\max}) = \hat{\beta}^M - \delta(\hat{\beta}^S - \hat{\beta}^M) \frac{R_{\max} - R_M^2}{R_M^2 - R_S^2},$$

where δ measures the strength of selection on unobservables relative to selection on observables. Setting $\delta = 1$ imposes the benchmark assumption that selection on unobservables is as strong as selection on observables. Solving $\beta^*(\delta, R_{\max}) = 0$ yields the Oster breakdown value

$$\delta_0 = \frac{\hat{\beta}^M(R_M^2 - R_S^2)}{(\hat{\beta}^S - \hat{\beta}^M)(R_{\max} - R_M^2)},$$

that is, the value of δ required to drive the adjusted coefficient to zero.

While the Oster δ measures the amount of selection needed to make the coefficient equal zero, [Masten and Poirier \(2026\)](#) develop a test to quantify the selection required to overturn its sign. As they explain in their paper, these thresholds can differ because the coefficient compatible with Oster's assumptions can admit a coefficient with the opposite sign while still ruling out a coefficient of exactly zero. Interestingly, in some cases, less selection may actually needed to reverse the sign than to explain the estimate away.

Let W_0 denote the fixed effects, W_1 the observed covariates, and W_2 the omitted controls. Let β^{long} denote the coefficient on the treatment in the long regression that includes both observed controls W_1 and omitted controls W_2 , and let β^{med} be the coefficient from the regression that includes only W_1 . Consistently with Oster's framework, I set the hypothetical long regression explanatory power to $R_{\text{long}}^2 = \min\{1.3 R_{\text{med}}^2, 1\}$. I additionally restrict the magnitude of the omitted variable bias to $|\beta_{\text{long}} - \hat{\beta}_{\text{med}}| \leq 2|\hat{\beta}_{\text{med}}|$, which is a bias bound of twice the baseline coefficient and an established choice proposed and examined by [Masten and Poirier \(2026\)](#).

Let $\mathcal{B}(\delta, R_{\text{long}}^2)$ denote the set of long-regression coefficients compatible with Oster's model and the observed regression moments, and define $s = \text{sign}(\hat{\beta}^{\text{med}})$. The reported sign-change breakdown point is

$$\hat{\delta}_{\text{sign}}^{\text{bp}} = \inf \left\{ |\delta| : \begin{array}{l} \text{there exists } b \in \mathcal{B}(\delta, R_{\text{long}}^2) \text{ such that } sb \leq 0, \\ |b - \hat{\beta}^{\text{med}}| \leq 2|\hat{\beta}^{\text{med}}| \end{array} \right\}.$$

This implies that the statistic measures the smallest absolute value of Oster's selection parameter consistent with a zero coefficient or a coefficient of the opposite sign, subject to the specified explanatory power and bias bound. Larger values indicate greater sign robustness under these assumptions.

This analysis is closely related to the work of [Diegert, Masten and Poirier \(2025\)](#) (DMP henceforth): both approaches assess how much selection on unobservables is needed to overturn regression estimates and address different limitations of earlier sensitivity analyses. While [Masten and Poirier \(2026\)](#) distinguish explaining away a coefficient from reversing its sign within Oster's model, DMP develop a different sensitivity framework that allows observed controls to be correlated with omitted variables and uses separate measures of relative selection in the treatment and outcome equations. The [Masten and](#)

Poirier (2026) statistic I report here retains Oster’s selection parameter and assumptions, and that is why I prefer it to the type of statistic discussed in DMP.¹⁷

Another approach to assess sensitivity to omitted variable bias is provided by Cinelli and Hazlett (2020). They depart from the proportional selection logic underlying the approaches discussed above and instead parameterizes the strength of unobserved confounding directly through partial R^2 measures. To formalize their idea, consider the population regression

$$Y = \tau D + X'\beta + \gamma U + \varepsilon, \quad (\text{A.4.1})$$

where D is the treatment or regressor of interest, X is the vector of observed controls, and U is an omitted confounder. Cinelli and Hazlett (2020) show that the magnitude of omitted variable bias can be illustrated in terms of two quantities: $R^2_{D \sim U|X}$ and $R^2_{Y \sim U|D,X}$. The first term measures the proportion of the residual variation in D , after conditioning on X , explained by the omitted confounder U . The second one denotes the proportion of the residual variation in Y , after conditioning on both D and X , explained by U . These two partial R^2 statistics therefore capture the treatment side and outcome side strength of the omitted confounding, respectively. The omitted variable bias in the estimated treatment coefficient can be written as a direct function of these two quantities: let $\widehat{\tau}_{\text{res}}$ denote the coefficient obtained from the regression omitting U , and let $\widehat{\tau}$ denote the coefficient that would be obtained if U were observed and included. Then, the absolute bias satisfies

$$|\widehat{\tau}_{\text{res}} - \widehat{\tau}| = \text{se}(\widehat{\tau}_{\text{res}}) \sqrt{\text{df} \frac{R^2_{Y \sim U|D,X} R^2_{D \sim U|X}}{1 - R^2_{D \sim U|X}}}.$$

Thus, any hypothesized pair $(R^2_{D \sim U|X}, R^2_{Y \sim U|D,X})$ implies a corresponding degree of omitted variable bias and, consequently, a bias-adjusted estimate of the treatment effect.

For a given proportional reduction q in the estimated coefficient, the robustness value from Cinelli and Hazlett (2020) is the minimum common partial R^2 that an omitted confounder would need to have with both the treatment and the outcome to generate such a reduction. In the commonly considered case with $q = 1$, the robustness value measures the minimum strength required from confounding variables to reduce the estimated treat-

¹⁷Both approaches are implemented in the Stata package `reg sensitivity`. The calculations reported here use its `oster` option, which implements the Masten and Poirier (2026) analysis within Oster’s model, rather than its `dmp` option.

ment effect to zero. Analogously, RV_α measures the minimum confounding strength required to make the estimate statistically insignificant at a specified significance level. It is important to highlight that the restriction that the two partial R^2 values are equal is used only to construct this scalar summary. In practice, the underlying sensitivity analysis allows the treatment side and outcome side strength of confounding to differ.

Table A.4.1: Sensitivity Analyses on the Effects of Horizontal Interlocks on Innovation Inputs, Outputs and Values

	(1) Log (Patents)	(2) Log (Patents)	(3) Log (CW Patents)	(4) Log (CW Patents)
Oster- β	[0.048]	[0.048]	[0.032]	[0.032]
Oster- δ	{1.684}	{1.542}	{1.364}	{1.279}
MP Sign Breakdown	<10.886>	<9.581>	<13.889>	<11.386>
N	4,575	4,810	4,575	4,810
R-squared	0.933	0.918	0.927	0.909
Firm FE	YES	YES	YES	YES
Year FE	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO

Notes: Each column reports the sensitivity of the *Interlock* coefficient of the corresponding OLS specification to unobserved confounding. Oster- β is the Oster (2019) bias-adjusted coefficient at $\delta = 1$ and a maximum R^2 of 1.3 times the observed one, capped at 1. Oster- δ is the value of δ that would set that coefficient to zero at the same maximum R^2 . MP Sign Breakdown is the [Masten and Poirier \(2026\)](#) breakdown point beyond which the sign of the effect can no longer be determined.

Relative to [Oster \(2019\)](#) and [Masten and Poirier \(2026\)](#), [Cinelli and Hazlett \(2020\)](#) characterize confounding directly through the explanatory power that an omitted variable would need to have for the residualized treatment and residualized outcome. This allows the econometrician to express their sensitivity parameters to be benchmarked directly against the partial R^2 values associated with observed covariates, which allows for easier interpretation. The framework is nevertheless closely connected to the approaches discussed above. For example, [Oster \(2019\)](#)'s $R_{\max}$ can be mapped into the outcome side

of the [Cinelli and Hazlett \(2020\)](#) sensitivity parameter according to

$$R_{Y \sim U|D,X}^2 = \frac{R_{\max} - R_{Y \sim D,X}^2}{1 - R_{Y \sim D,X}^2}.$$

The principal conceptual difference lies in the treatment side. While proportional selection approaches impose structure linking selection on unobservables to selection on observables, [Cinelli and Hazlett \(2020\)](#) leave $R_{D \sim U|X}^2$ as a sensitivity parameter to be calibrated directly. Hence, their method does not require the researcher to posit a particular proportional relationship between observed and unobserved selection, like the [Oster \(2019\)](#) test does.

The [Cinelli and Hazlett \(2020\)](#) analysis should therefore be viewed as complementary to the [Oster \(2019\)](#) and [Masten and Poirier \(2026\)](#) exercises, and reporting all these types of sensitivity analysis provides robustness assessments under distinct, but closely related, ways of formalizing the potential importance of unobserved confounding. Accordingly, in the main part of the paper, I report the sensitivity analysis test from [Cinelli and Hazlett \(2020\)](#). In this Appendix, Table [A.4.1](#) reports the sensitivity results for innovation outputs using the [Oster \(2019\)](#) and [Masten and Poirier \(2026\)](#) tests. It indicates that these estimates are unlikely to be driven by omitted variable bias. The corresponding Oster β values are 0.048 for the raw patent counts for both specifications, implying that the coefficients remain positive even under the benchmark assumption that selection on unobservables is as strong as selection on observables. The Oster- δ values, 1.684 and 1.542, imply that selection on unobservables would need to be roughly 1.6 times as strong as selection on observables to drive the corresponding coefficients to zero.¹⁸ The MP sign breakdown values, 10.886 and 9.581, point in the same direction: selection on unobservables would need to be approximately 10.9 times as strong as selection on observables, as measured by Oster's relative selection parameter, to overturn the conclusion that the effect is positive. The results regarding the citation-weighted counts are similar.

¹⁸For reference, [Oster \(2019\)](#) reports that only 45% of the non-randomized papers from the top 5 journals in Economics have an Oster delta greater than 1.

A.5 Alternative Specification

In this section of the appendix, I test the robustness of my baseline results to alternative specifications, from changing the estimator to trying a different definition of market competition and competing firms.

A.5.1 Poisson Regressions

A Poisson regression is usually the estimator of choice when dealing with patent counts, which inevitably contain a lot of zeros (Hausman, Hall and Griliches, 1984, Silva and Tenreyro, 2006). However, as I explain in the main text, I must use OLS because it is the only estimator that can be paired with common sensitivity analysis to test for omitted variable bias. Nevertheless, in this subsection, I still report Poisson as a robustness check and, from Table A.5.1, I find that the results are the same as OLS.

Table A.5.1: Poisson Estimates of the Effects of Horizontal Interlocks on Innovation Inputs, Outputs and Values

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	R&D	R&D	Patents	Patents	CW Patents	CW Patents	Avg. Value	Avg. Value
Interlock	-0.039 (0.032)	-0.010 (0.029)	0.119** (0.049)	0.183*** (0.062)	0.237*** (0.061)	0.207*** (0.063)	-0.080 (0.065)	-0.023 (0.058)
N	2,827	3,085	3,454	3,779	3,454	3,779	3,172	3,501
Pseudo R-squared	0.990	0.986	0.972	0.959	0.970	0.958	0.899	0.853
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Estimates are Poisson pseudo-maximum likelihood (ppmlhdf), with the dependent variables in levels. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one Section 8-violating horizontal interlock in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

A.5.2 Results with Average AKM Ability

In my baseline specification, after estimating my AKM model, I use aggregated firm-level estimated director fixed effects as a control variable. In this subsection, I test whether the results change if I were to use average estimated AKM effects instead. Table A.5.2 reports these results.

Table A.5.2: Effects of Horizontal Interlocks on Innovation Inputs, Outputs and Values

	(1) Log (R&D)	(2) Log (R&D)	(3) Log (Patents)	(4) Log (Patents)	(5) Log (CW Patents)	(6) Log (CW Patents)	(7) Log (Avg. Value)	(8) Log (Avg. Value)
Interlock	0.017 (0.027)	0.008 (0.024)	0.115*** (0.042)	0.134*** (0.044)	0.119** (0.049)	0.143*** (0.052)	0.001 (0.051)	0.005 (0.045)
N	4,283	4,528	4,575	4,810	4,575	4,810	4,246	4,461
R-squared	0.990	0.986	0.933	0.918	0.927	0.909	0.840	0.787
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one Section 8-violating horizontal interlock in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

A.5.3 Results with Common Ownership as a Control Variable

In this subsection, I try adding common ownership to my vector of control variables and show the results of this exercise in Table A.5.3. Once again, the results remain unchanged.

Table A.5.3: Effects of Horizontal Interlocks on Innovation Inputs, Outputs and Values


	(1) Log (R&D)	(2) Log (R&D)	(3) Log (Patents)	(4) Log (Patents)	(5) Log (CW Patents)	(6) Log (CW Patents)	(7) Log (Avg. Value)	(8) Log (Avg. Value)
Interlock	0.024 (0.026)	0.016 (0.023)	0.110** (0.043)	0.129*** (0.044)	0.206*** (0.057)	0.196*** (0.055)	-0.004 (0.052)	-0.012 (0.046)
N	4,353	4,602	4,374	4,637	2,615	2,766	4,372	4,635
R-squared	0.990	0.986	0.934	0.919	0.943	0.929	0.840	0.792
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one Section 8-violating horizontal interlock in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

A.5.4 All Interlocks

As another robustness exercise, I relax the baseline treatment definition and replace Section 8-violating horizontal interlocks with a broader indicator for *all* interlocks. Under this alternative specification, a firm is treated in year t if it shares at least one director with any other firm, regardless of whether the linked firms are horizontal rivals or are connected through vertical or not. Figure A.5.1 reports the statistics for this type of interlock.

Figure A.5.1: Firm-Level Statistics About All Interlocking Directorates



Notes: Panel (a) shows the average number of horizontal and vertical interlocking directorates (blue line, left axis) across all firms in the sample by year, and the percentage of firms that have at least one interlocking directorate by year (red line, right axis). Panel (b) shows the distribution of the number of interlocking directorates by firm in 2024, along with the percentage of firms for each bracket.

Table A.5.4: Effects of Section 8 Interlocks on Innovation Inputs, Outputs and Values

	(1) Log (R&D)	(2) Log (R&D)	(3) Log (Patents)	(4) Log (Patents)	(5) Log (CW Patents)	(6) Log (CW Patents)	(7) Log (Avg. Value)	(8) Log (Avg. Value)
Interlock	0.030 (0.044)	0.012 (0.041)	0.130** (0.063)	0.123** (0.060)	0.126* (0.072)	0.131** (0.066)	0.087 (0.058)	0.064 (0.051)
N	4,283	4,528	4,575	4,810	4,575	4,810	4,246	4,461
R-squared	0.990	0.986	0.932	0.918	0.926	0.909	0.840	0.787
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one interlock (horizontal or vertical) in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

A.5.5 Results with Hoberg–Phillips-based Definition of Competition

Here, I broaden the baseline definition of competitors by combining NAICS-2 classifications (TNIC-2) with the Hoberg–Phillips (Hoberg and Phillips, 2016, 2010) text-based network industry classification (TNIC-2). I now say that two firms are classified as competitors if they share the same two-digit NAICS code or are linked in the official TNIC-2 network in the corresponding year. I then define the interlock indicator as equal to one when a firm shares at least one director with a competitor under this combined definition. The results, show in Table A.5.5, illustrate that the coefficients are robust to this alternative specification.

Table A.5.5: Effects of Product-Market (TNIC-2) Interlocks on Innovation Inputs, Outputs and Values

	(1) Log (R&D)	(2) Log (R&D)	(3) Log (Patents)	(4) Log (Patents)	(5) Log (CW Patents)	(6) Log (CW Patents)	(7) Log (Avg. Value)	(8) Log (Avg. Value)
Interlock	-0.012 (0.026)	0.000 (0.022)	0.106** (0.041)	0.100** (0.041)	0.103** (0.048)	0.101** (0.047)	0.023 (0.049)	0.011 (0.042)
N	4,283	4,528	4,575	4,810	4,575	4,810	4,246	4,461
R-squared	0.990	0.986	0.933	0.918	0.926	0.909	0.840	0.787
Firm FE	YES	YES	YES	YES	YES	YES	YES	YES
Year FE	NO	YES	NO	YES	NO	YES	NO	YES
Industry-Year FE	YES	NO	YES	NO	YES	NO	YES	NO
State-Year FE	YES	NO	YES	NO	YES	NO	YES	NO


Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one interlock in year t with a firm that is either in the same NAICS-2 category or with a firm that is a Hoberg–Phillips TNIC-2 competitor in year t , where both firms meet the Section 8 size thresholds and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

A.6 Event Studies

A.6.1 Citations by Origin and Destination

In this Section, I inspect whether the citation outcomes that I report in the main text differ depending on the direction of director movement within the dyad. In particular, I define the origin firm as the firm in which the shared director was first appointed chronologically, and the destination firm as the company in the dyad that the director joins at a later date. For simplicity, I exclude dyads in which directors were appointed in the same year. The results illustrate that the dynamic pattern is symmetric across origin and destination firms. The two panels of Figure [A.6.1](#) show that, after the first interlock, citations increase on both sides of the dyad. The effect only grows gradually over time, which mirrors the same dynamic effects illustrated in Figure [3](#). The magnitude is very slightly larger on the origin side at longer horizons, but the two profiles are remarkably similar overall. Interlocks therefore seem to facilitate knowledge flows in both directions across the connected pair.

Figure A.6.1: Event Studies - Citations by Origin and Destination



Notes: Each panel reports staggered-adoption event study estimates of the effect of first adopting an interlock on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) framework, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. Event time is defined relative to the first interlock ($k = 0$) between any two firms. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Logged outcomes use $\log(1 + x)$. All specifications include calendar-year fixed effects and the baseline dyad-level controls as described in the text. The outcome is *Cross-Citations*, the share of a firm-year's patent citations that go to (or come from) firms with which the focal firm shares at least one director in year t , split between origin and destination firms as defined in the text.


A.7 Robustness Checks

A.7.1 Placebo Event Studies

To assess whether the patterns in the staggered event studies illustrated in the main paper could be generated mechanically by generic trends around future interlocks rather than by interlocks themselves, I estimate placebo versions of these event studies.


For the dyad-level citation outcomes, I begin from the same firm-pair-year panel used in the main event studies and construct pseudo-treatment dates by shifting each treated dyad's true first-interlock year backward by six years. That is, if G_{ij} denotes the first year in which dyad (i, j) becomes interlocked, I define the placebo cohort as $G_{ij}^P = G_{ij} - 6$ for treated dyads, while never-treated dyads remain untreated. This is the same placebo event study framework discussed in [Borusyak, Jaravel and Spiess \(2024\)](#) and a common way to generate placebo event studies. I then re-estimate the same staggered-adoption [Callaway and Sant'Anna \(2021\)](#) design as in the main text, replacing the true treatment date with G_{ij}^P and keeping $k = -1$ as the omitted period.

Figure A.7.1: Placebo Event Studies - Patent Counts




Notes: This figure reports staggered-adoption event study estimates of the effect of the placebo adoption date on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) estimator, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Event time is defined relative to the placebo treatment date ($k = 0$); $k = -1$ is the omitted baseline.

Figure A.7.2: Placebo Event Studies - Citations




Notes: This figure reports staggered-adoption event study estimates of the effect of the placebo adoption date on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) estimator, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Event time is defined relative to the placebo treatment date ($k = 0$); $k = -1$ is the omitted baseline.

Figure A.7.3: Placebo Event Studies – New Classes




Notes: This figure reports staggered-adoption event study estimates of the effect of the placebo adoption date on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) estimator, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Event time is defined relative to the placebo treatment date ($k = 0$); $k = -1$ is the omitted baseline.

Figure A.7.4: Placebo Event Studies - Technological Proximity



Notes: This figure reports staggered-adoption event study estimates of the effect of the placebo adoption date on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) estimator, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Event time is defined relative to the placebo treatment date ($k = 0$); $k = -1$ is the omitted baseline.

Figure A.7.5: Placebo Event Studies - Patent Age



Notes: This figure reports staggered-adoption event study estimates of the effect of the placebo adoption date on the stated outcome. The reported estimates are from the [Callaway and Sant'Anna \(2021\)](#) estimator, which maps cohort- and time-specific $ATT(g, t)$ to event time $k = t - g$, with $k = -1$ normalized to zero. The dots are the estimates from such regressions. The darker red band shows the 90% confidence intervals, whereas the lighter red band shows the 95% confidence intervals. The standard errors are clustered at the firm-pair level. Event time is defined relative to the placebo treatment date ($k = 0$); $k = -1$ is the omitted baseline.

A.8 Markups

A.8.1 Implementation Details for Markup Estimation

I build a firm–year panel and retain observations with non-missing gross revenue, materials expenditure, and the book value of plant, property, and equipment. Inputs are constructed as logs of gross output, materials, and capital and log employment as a labor quantity. I estimate a translog gross-output production function separately within 2-digit industries:

$$\ln Q_{it} = \beta_0 + \beta_M \ln M_{it} + \beta_L \ln L_{it} + \beta_K \ln K_{it} + \beta_{MM} (\ln M_{it})^2 + \beta_{ML} \ln M_{it} \ln L_{it} + \beta_{MK} \ln M_{it} \ln K_{it} + \dots + \omega_{it} + \varepsilon_{it}. \quad (\text{A.8.1})$$

To address simultaneity and selection, I use the [Olley and Pakes \(1996\)](#) investment proxy estimator with the [Akerberg, Caves and Frazer \(2015\)](#) correction to estimate translog coefficients and a firm–year transitory error estimate $\hat{\varepsilon}_{it}$ from the first-stage residual. The mapping from elasticities and expenditure shares to markups follows the [De Loecker and Warzynski \(2012\)](#) framework.

Given the estimated translog, the materials elasticity varies by firm and year:

$$\theta_{M,it} = \hat{\beta}_M + 2\hat{\beta}_{MM} m_{it} + \hat{\beta}_{ML} l_{it} + \hat{\beta}_{MK} k_{it}.$$

The materials revenue share is

$$\alpha_{M,it} = \frac{\text{COGS}_{it}}{\widetilde{R}_{it}}, \quad \widetilde{R}_{it} = \frac{\text{Sales}_{it}}{\exp(\hat{\varepsilon}_{it})},$$

i.e., I scale revenue by $\exp(\hat{\varepsilon}_{it})$ to adjust for the estimated transitory error component. Firm–year markups are computed as

$$\mu_{it} = \frac{\theta_{M,it}}{\alpha_{M,it}},$$

as implied by cost minimization with competitive input markets ([De Loecker and Warzynski, 2012](#)). For each firm–year, I retain $(\mu_{it}, \alpha_{M,it}, \hat{\varepsilon}_{it})$ for subsequent analysis. The dependent variable in the markup regressions is $\ln(1 + \mu_{it})$.

A.8.2 Markup Results

Table A.8.1 presents the results of running Equations 1 and 2 on the markup numbers estimated using the procedure illustrated in this section. In each specification, the coefficient on firm-level markups is null.

Table A.8.1: Effects of Horizontal Interlocks on Markups


	(1) Log(Markup)	(2) Log(Markup)
Interlock	0.011 (0.008)	0.009 (0.008)
N	4,636	4,877
R-squared	0.895	0.880
Firm FE	YES	YES
Year FE	NO	YES
Industry-Year FE	YES	NO
State-Year FE	YES	NO

Notes: *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors (clustered by firm) are reported in parentheses. The unit of observation is the firm-year. *Interlock* is an indicator equal to one if the firm has at least one Section 8-violating horizontal interlock in year t and zero otherwise. All specifications include a set of firm-level controls chosen by double LASSO, the technology and product-market spillover covariates entered on top of that selection, and the corresponding sets of fixed effects as reported in the table.

A.9 Results by Industry

Figure A.9.1 reports industry-specific estimates. I estimate industry-specific effects using a pooled OLS specification that interacts the interlock indicator with industry indicators. For each industry, the reported marginal effect combines the coefficient on the interlock indicator with the corresponding interaction coefficient.

Figure A.9.1: Effects of Section 8 Interlocks on Innovation by Industry



Notes: These graphs report the average marginal effect of Interlock_{jt} estimated from Regression 1 augmented with a full set of interactions between Interlock_{jt} and industry indicators. Each dot denotes the estimated effect for a different industry, defined by NAICS-2 sectors (with codes 31–33, 44–45 and 48–49 collapsed into single sectors). Only industries containing at least one firm whose interlock status changes over the sample period are shown. The horizontal navy blue bars denote the corresponding 95% confidence intervals, based on standard errors clustered at the firm level.

Bibliography of the Appendix

- Ackerberg, Daniel A., Kevin Caves, and Garth Frazer.** 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica*, 83(6): 2411–2451.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess.** 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *The Review of Economic Studies*, 91(6): 3253–3285.
- Callaway, Brantly, and Pedro H.C. Sant'Anna.** 2021. "Difference-in-Differences with multiple time periods." *Journal of Econometrics*, 225(2): 200–230.
- De Loecker, Jan, and Frederic Warzynski.** 2012. "Markups and Firm-Level Export Status." *American Economic Review*, 102(6): 2437–71.
- Hausman, Jerry, Bronwyn Hall, and Zvi Griliches.** 1984. "Econometric Models for Count Data with an Application to the Patents-RD Relationship." *Econometrica*, 52(4): 909–38.
- Hoberg, Gerard, and Gordon Phillips.** 2010. "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis." *The Review of Financial Studies*, 23(10): 3773–3811.
- Hoberg, Gerard, and Gordon Phillips.** 2016. "Text-Based Network Industries and Endogenous Product Differentiation." *Journal of Political Economy*, 124(5): 1423–1465.
- Masten, Matthew A., and Alexandre Poirier.** 2026. "The Effect of Omitted Variables on the Sign of Regression Coefficients." *American Economic Review*, 116(7): 2685–2710.
- Olley, G. Steven, and Ariel Pakes.** 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6): 1263–1297.
- Silva, J. M. C. Santos, and Silvana Tenreyro.** 2006. "The Log of Gravity." *The Review of Economics and Statistics*, 88(4): 641–658.