CREATIVE CONSTRUCTION: KNOWLEDGE SHARING AND COOPERATION BETWEEN FIRMS*

Evgenii FADEEV †

June 12, 2024 Click here for the latest version

Abstract

Knowledge flows among firms are often assumed to take the form of involuntary spillovers, which are commonly measured by patent citations. I show that citations are highly concentrated and primarily come from business partners. I provide empirical evidence suggesting that citations measure intentional knowledge transfers between collaborating firms rather than involuntary spillovers. The concentration of citations has increased since 2000, indicating more selective knowledge sharing among partners. I develop a theory explaining how market forces affect knowledge flows. Technologies consist of multiple knowledge components, each of which can be either patented or kept secret. Patenting facilitates the division of labor among business partners but discloses information to competitors. Higher competition leads to greater reliance on secrecy, which impedes contractible knowledge sharing with partners and makes it more selective. Using data on trade secret litigation and import competition from China, I provide evidence supporting the predictions of the model. JEL Codes: L24, 033, 034

^{*}The previous title of the paper was "Creative Construction: Knowledge Sharing in Production Networks." I am immensely grateful to Pol Antràs, Oliver Hart, Josh Lerner, and Andrei Shleifer for their guidance and support. I have been fortunate to benefit from insightful discussions with Ashish Arora, Sharon Belenzon, Jordi Blanes i Vidal, Maxim Boycko, Ronnie Chatterji, Gabriel Chodorow-Reich, Wesley Cohen, Emmanuel Farhi, Edward Glaeser, Daniel Gross, Gordon Hanson, Sharique Hasan, Elhanan Helpman, Katrin Hussinger, Sergei Izmalkov, Myrto Kaloutpsidi, William Kerr, Kristina Komissarova, Josh Krieger, Robin Lee, Marc Melitz, Tom Nicholas, Ariel Pakes, Peter Thompson, Gaétan de Rassenfosse, Jesse Shapiro, Scott Stern, Ludwig Straub, Hosny Zoabi. I thank the participants in the NBER SI Innovation Session, REER Conference, DISS Workshop, and NBER Productivity Conference. I also thank Lisa Ouellette, John Squires, and Christopher Tehrani for discussions about the legal issues surrounding patents and trade secrets.

[†]Duke University, Fuqua School of Business, email: evgenii.fadeev@duke.edu.

1 Introduction

Knowledge flows among firms are often assumed to take the form of involuntary spillovers "in the air" (Marshall 1920) rather than intentional knowledge sharing.¹ This paper shows evidence suggesting that firms have significant control over the knowledge they generate, selectively sharing it with a limited set of business partners, such as suppliers and customers. Intentional knowledge sharing has substantially declined over time due to changes in how firms protect their intellectual property and collaborate with partners.

A prevailing measure of knowledge spillovers is patent citations.² The existence of patents is often justified by the claim that they promote knowledge diffusion through the disclosure of inventions.³ I show that the distribution of citations across firms is highly concentrated, raising a question about the role of patents in the diffusion of knowledge. For example, IBM's patent in Table 1 is heavily cited, but almost all of its citations come from IBM's *input supplier*, Amkor Technology Inc. Figure 1 shows that citations are highly concentrated in general: the most cited patents granted in the U.S. between 1980 and 2000 received around 50% of citations from one firm only, and this concentration increased to 77% in 2014.

Patent Number	Assignee	Total Number of Citations	% of Citations from Amkor Technology Inc		
5877043	IBM	$218 \\ top \ 0.005\%$	94%		

Table 1: Example of a patent with a high concentration of citations

I collect data on various inter-firm relations between publicly listed U.S. companies. I show that business partners—such as suppliers, customers, or firms with research collaboration—account for approximately 76% of inter-firm citations. Changes in citation patterns among partners explain around 84% of the rise in the concentration since 2000. In particular, the average number of partners citing a typical firm declined by 42%, even though the overall number of partners for a firm increased. These changes are explained by increasing differences across firms in their citation probabilities rather than shifts in the distribution of pattern counts or the composition of partners. For example, a rise of superstar firms with large patent portfolios cannot explain the increasing concentration.

¹The existence of knowledge spillovers is key in models of economic growth (Romer 1990; Grossman and Helpman 1991; Aghion and Howitt 1992) and a common justification for R&D subsidies (Bloom et al. 2019).

²For example, patent citations are used to estimate growth models (Caballero and Jaffe 1993; Eeckhout and Jovanovic 2002; Akcigit and Kerr 2018), evaluate the localization of spillovers in space (Jaffe et al. 1993; Thompson and Fox-Kean 2005; Singh and Marx 2013), and provide policy recommendations (Liu and Ma 2023).

³"[T]he patent system represents a carefully crafted bargain that encourages both the creation and the public disclosure of new and useful advances in technology, in return for an exclusive monopoly for a limited period of time" (*Pfaff v. Wells Elecs., Inc.*, 525 U.S. 55, 63, 1998); see also Mazzoleni and Nelson (1998).



Figure 1: Concentration of Citations

This figure shows the average concentration of citations for the most cited patents between 1976 and 2014. In each grant year and technology class for the period of 1976–2014, I track citations within a five-year window for the top 1% of the most cited patents. For each cited patent, the concentration is defined as the share of citations coming from the most citing firm. The technological classes are defined at the group level in the cooperative patent classification system. To construct the aggregate measure, I take the average concentration across patents within each class and then the average across classes weighted by the number of patents.

In addition, citations are specific to firms citing each other rather than individual inventors. I show that an inventor who heavily cites a particular patent in one company significantly decreases her citations to this patent once she moves to another firm. This evidence suggests that citations should be explained based on firm relationships rather than human capital factors. Furthermore, the results are not driven by citations from patent examiners or attorneys and are robust to various specifications.

This evidence is puzzling because, according to the U.S. Supreme Court, one of the primary functions of the patent system is the broad dissemination of knowledge.⁴ This presumption is based on the patent system's disclosure requirement, by which inventors receive patent rights in exchange for publicly disclosing their knowledge. This arrangement assumes that inventors reveal information that would otherwise remain secret, enabling subsequent inventors to learn technologies by studying patent applications (Landes and Posner 2003). Therefore, patent

⁴Kewanee Oil Co. v. Bicron Corp., 416 U.S. 470, 481 (1974): "When a patent is granted and the information contained in it is circulated to the general public and those especially skilled in the trade, such additions to the general store of knowledge are of such importance to the public weal that the Federal Government is willing to pay the high price of 17 years of exclusive use for its disclosure, which disclosure, it is assumed, will stimulate ideas and the eventual development of further significant advances in the art." See also Bonito Boats, Inc. v. Thunder Craft Boats, Inc., 489 U.S. 141, 150-51 (1989): "the ultimate goal of the patent system is to bring new designs and technologies into the public domain through disclosure."

applications are expected to diffuse knowledge similarly to scientific publications. In the same way, models of endogenous economic growth often assume that patented knowledge is protected from imitation but serves as a free input for future research, meaning that follow-on inventors are not expected to be business partners of patent owners (e.g., Romer 1990; Kortum 1997).⁵

I argue that firms might control the diffusion of their knowledge, at least as measured by citations. They selectively share knowledge with a few partners, thereby providing them an advantage in creating follow-on innovations. According to U.S. patent law, these partners are required to cite the original patent when they file for patents on the follow-on innovations. As a result, citations are concentrated because only a few partners of a patent owner gain access to the knowledge embodied in the patent.

I develop a theory of knowledge disclosure and sharing that can rationalize citation patterns and explain how firms control the diffusion of their knowledge. This theory is based on three main assumptions. To illustrate how these assumptions help explain the citation patterns, I contrast them with the standard assumptions in models of economic growth and firm dynamics that are based on involuntary knowledge spillovers (e.g., see reviews in Jones 2005; Aghion et al. 2014; Buera and Lucas 2018).

First, I assume that the efficient allocation of resources might require the division of labor across different firms. For instance, a firm conducting research on a new technology might outsource further development and manufacturing to another company, as in the case of IBM and Amkor Technology.⁶ This division of labor creates a network of business partners who need to share knowledge. In contrast, standard growth models often assume that all stages of research, development, and production are done within one firm. Thus, these models typically focus on competition among firms rather than cooperation. Since most inter-firm citations occur between partners, a theory explaining these citations should account for inter-firm partnerships.

Second, patents are primarily seen as protection against imitation. This protection is crucial for R&D incentives if knowledge can be reverse engineered. However, if knowledge is not reverse engineerable, firms can also protect their intellectual property through secrecy. Surveys consistently show that firms rank secrecy higher than patents for intellectual property protection (Hall et al. 2014). Legal scholars argue that viewing patents merely as protection from imitation is incomplete, and "a more illuminating way of thinking about the patent system is as a response

⁵"The owner of a design has property rights over its use in the production of a new producer durable but not over its use in research. If an inventor has a patented design for widgets, no one can make or sell widgets without the agreement of the inventor. On the other hand, other inventors are free to spend time studying the patent application for the widget and learn knowledge that helps in the design of a wodget. The inventor of the widget has no ability to stop the inventor of a wodget from learning from the design of a widget" (Romer 1990, p. 84).

⁶IBM's patent in Table 1 covers the technology related to semiconductor packaging. Amkor was IBM's assembly manufacturer in this area (https://ir.amkor.com/news-releases/news-release-details/amkor-enters-long-term-semiconductor-assembly-and-test-services).

to economic problems inherent in trade secrecy" (Landes and Posner 2003, p. 294).

I argue that the citation patterns are better understood through the lens of secrecy. Without patent rights, some knowledge might remain secret, impeding the division of labor across firms due to the risk of knowledge expropriation (Arrow 1962). For instance, without formal protection, once a research firm reveals its secrets to a manufacturer, the manufacturer might refuse to pay. Patents enable contractual knowledge transfers and the division of labor across firms by establishing and protecting intellectual property rights, thereby preventing misappropriation by partners (Arora et al. 2001). If a firm does not need to share knowledge with others, it may rely entirely on secrecy to prevent competitors from accessing and improving upon that knowledge, potentially replacing the firm.

The final assumption can explain the puzzle of disclosure ineffectiveness described earlier. Suppose technologies consist of multiple components: formally, knowledge corresponds to a vector of productivity levels in a multi-input production function. Each component can be either patented or kept secret and is only valuable with other components. If knowledge components are sufficiently complementary to each other, patenting some components while keeping others secret can create a dependency barrier: partners cannot fully benefit from the secrets without access to the patented knowledge, and competitors face significant difficulties in improving the technology without observing the secret knowledge. Thus, this strategy can minimize the expropriation by partners while concealing much of the knowledge from competitors.

The standard growth models often assume that knowledge about a particular technology can be summarized in a one-dimensional metric, such as product quality or total factor productivity. The multi-dimensionality of knowledge implies that the same technology can be protected through a combination of different forms of intellectual property protection, such as patenting and secrecy. This idea is supported by firm surveys (Cohen et al. 2000; Mezzanotti and Simcoe 2023), management studies (Amara et al. 2008), patent policy discussions (Anton et al. 2006), and legal research (Roin 2005; Fromer 2008; Jorda 2008). In Section 4, I discuss examples of how firms combine patenting with secrecy in the production of mRNA-based vaccines, ammonia, semiconductors, and self-driving cars.

In the model, information disclosure links competition (creative destruction) and cooperation (creative construction) among firms. Firms choose the optimal patent-secrecy combination to balance the benefits of patent rights, which facilitate the division of labor among partners, against the costs of disclosing information to competitors. As competition intensifies and the costs of disclosure rise, firms increasingly rely on secrecy. However, this reliance on secrecy hinders knowledge sharing among partners and results in more concentrated knowledge flows.

I test two predictions from the model. First, the model predicts that bundling patents with trade secrets leads to more concentrated knowledge flows. Testing this connection is challenging because trade secrets are not observable. I use trade secret litigation data to find patents that were likely to be bundled with trade secrets. For example, the legal case *Waymo v. Uber* was about the misappropriation of trade secrets related to the light detection and ranging (LiDAR) technology for self-driving cars.⁷ However, the same lawsuit also had claims regarding patent infringement, and in the legal complaint, Waymo described the complementarity between their patents and secrets. I show that patents which were involved in both patent and trade secret litigation have more concentrated citations relative to similar patents within the same firm that were involved in patent litigation only.

Second, the model predicts that higher Schumpeterian competition leads to more concentrated knowledge flows. By Schumpeterian competition, I mean the competition that reduces incentives to innovate due to lower profit rents after R&D investments (Schumpeter 1942; Shapiro 2012).⁸ Autor et al. (2020) show that import competition from China reduced profitability, patenting, and R&D investments of U.S. companies, suggesting that it increased Schumpeterian competition in the U.S. Using the same empirical strategy, I show that technologies more exposed to import competition from China experienced higher growth in the concentration of citations.

Literature

This paper is related to the literature on innovation and knowledge diffusion. Involuntary knowledge spillovers are fundamental in models of economic growth (Jones 2005; Aghion et al. 2014; Buera and Lucas 2018), provide a rationale for R&D subsidies (Bloom et al. 2019), and can explain the agglomeration of economic activity (Marshall 1920; Saxenian 1996).

Patent citations have been the prevailing measure of knowledge spillovers since the seminal work of Jaffe et al. (1993). Several surveys of inventors and firms confirm that citations are correlated with knowledge flows (e.g., Jaffe et al. 2000; Duguet and MacGarvie 2005). While these flows might represent intentional knowledge sharing, patent citations are commonly interpreted as unintended knowledge spillovers, for example, in the growth literature (e.g., Caballero and Jaffe 1993; Akcigit and Kerr 2018) and urban studies (e.g., Ellison et al. 2010; Singh and Marx 2013). The evidence in this paper suggests that firms might control the diffusion of their knowledge, and citations reflect intentional knowledge transfers rather than involuntary spillovers.

This paper is related to the literature on markets for technologies (e.g., Gans and Stern 2000; Arora and Gambardella 2010; Hegde and Luo 2017; Arqué-Castells and Spulber 2022).

⁷ Waymo LLC v. Uber Technologies, Inc., No. 17-2235 (Fed. Cir. 2017).

⁸Schumpeterian competition contrasts with ex-ante competition occurring before investments in R&D (Arrow 1962). Ex-ante competition positively affects innovation incentives (Shapiro 2012; Whinston 2012).

In my model, the benefits of patent protection and licensing are weighed against the cost of information disclosure to competitors. Therefore, firms control knowledge diffusion not only through licensing contracts but also through information disclosure and trade secret sharing.

This paper is also related to the literature on competition and innovation (e.g., Aghion et al. 2005; Shapiro 2012). The standard growth models assume knowledge flows in the form of involuntary spillovers (e.g., Aghion et al. 2014). In these models, the distribution of knowledge spillovers across firms is not affected by changes in market conditions. I show that import competition from China changed the distribution of patent citations. I rationalize this finding through increased reliance on secrecy and more selective knowledge sharing among partners.

Akcigit and Ates (2022) argue that a decline in knowledge flows among competitors might be responsible for the recent macro trends in the U.S. economy, such as the rising market concentration and declining business dynamism. The evidence in this paper is consistent with the decline in knowledge diffusion, especially among business partners. An increasing reliance on secrecy can explain the decline in knowledge flows to competitors and partners.

The firms identified as suppliers and customers account for the most citations among partners. The importance of vertical knowledge flows is highlighted in the historical studies on technological change (e.g., Rosenberg 1963) and in the literature on R&D cooperation (e.g., Cassiman and Veughelers 2002), general purpose technologies (e.g., Helpman 1998; Bresnahan 2010), and foreign direct investment (e.g., Alfaro-Ureña et al. 2022; Bai et al. 2022).

Finally, the evidence in this paper gives a new interpretation to the spatial localization of patent citations (Jaffe et al. 1993). This finding is commonly interpreted as evidence of the importance of proximity and serendipitous face-to-face interactions in knowledge diffusion (Carlino and Kerr 2015). However, the results in this paper suggest that such diffusion is either not serendipitous or that the localization of citations is a consequence of the co-location of business partners for reasons unrelated to spillovers (Ellison et al. 2010).

This paper is organized as follows. Section 2 documents that patent citations are highly concentrated. Section 3 shows that citations primarily come from business partners. Section 4 develops and tests the model of optimal information disclosure and knowledge sharing among partners. Section 5 concludes the paper.

2 Concentration of Patent Citations

This section documents that patent citations are highly concentrated. Section 2.1 provides background on the U.S. patent system and describes the data. Section 2.2 shows that, even for the most cited patents, the majority of citations come from one firm only, and this concentration has significantly increased since 2000. Section 2.3 shows that the concentration of citations

should be explained based on the difference across firms in their citation behavior rather than differences in the number of patents. In Section 2.4, I use the movement of inventors across firms to show that the concentration is primarily driven by firms rather than inventors. Section 2.5 documents multiple additional facts and robustness checks. Finally, Section 2.6 extends the concentration measure to the firm level.

For statistics on the concentration of citations, I also report the counterfactual statistics that would be observed if citations were random. Since I observe the universe of all patents and citations in the USPTO, there is no sampling uncertainty (Abadie et al. 2020). Instead, I assume that the randomness of citations provides a basis for inference. The details are given in Section 2.3.

2.1 Data and Background

Patents consist of two parts: a written description of an invention, including citations to prior art (patents, publications, etc.), and claims defining the boundaries of intellectual property rights. To be patentable, an invention must be patent-eligible, useful, novel, and non-obvious. Additionally, the text of the application should satisfy the disclosure requirements.⁹ Patent examiners use references to prior art to check whether the invention is novel and non-obvious. In the U.S., applicants have a "duty of candor" to disclose relevant prior art that they are aware of, and failure to do so can lead to patent invalidation. Prior art is generally used to strengthen, narrow, or reject certain claims. Therefore, citations serve the legal function of delimiting intellectual property rights on an invention.

I use the data on utility patents granted by the U.S. Patent and Trademark Office (USPTO) for the period from 1976 to 2019. Most granted patents contain information about assignees (patent owners). I clean assignee names to group patents by firms, individual inventors, universities, and other organizations. Autor et al. (2020) provide a matching of patent assignees to names of publicly traded firms in the Compustat data set. I extend their matching for the additional years of 2015–2019 and for private firms. Details are given in Appendix A.1.

2.2 High Concentration of Patent Citations

For each year and technological class from 1976 to 2014, I track citations within a five-year window for the top 1% of the most cited granted patents. The sample includes patents assigned to publicly traded companies, private firms, and non-corporate entities. In Section 2.5, I show

⁹These requirements are governed by the U.S. Code, Title 35, sections 101, 102, 103, and 112. For a review, see Scotchmer (2004), ch. 3.

that the results are robust to other thresholds, time frames, and sample selections.

There are two reasons to focus on the set of the most cited patents. First, the value of patents is highly skewed, with many being of no value to the firm, so the empirical literature on innovation is often focused on the most cited patents (e.g., Aghion et al. 2023). This focus is justified by the positive correlation between the number of patent citations and the firm's stock market valuation (Hall et al. 2005; Kogan et al. 2017). Second, by construction these patents are expected to generate most follow-on innovations and knowledge flows, providing a lower bound on the concentration measure. Taking a five-year window controls for the truncation bias that older patents have more time to accumulate citations. Comparing patents within each technological class controls for differences across classes in citation patterns (Lerner and Seru 2022). To classify technologies, I use the group level of the Cooperative Patent Classification (CPC) system, which has 672 groups.

For each patent among the most cited ones, I compute the distribution of citations across different organizations. The variable $n_{k,i}$ denotes the number of citations from organization *i* to patent *k*. Organizations are mostly firms but also include non-corporate entities, such as government agencies and universities. I exclude citations from individual inventors and patents with missing assignee information. The concentration measure for patent *k* is the share of citations coming from the most citing organization:¹⁰

$$\mathcal{C}_k = \max_i \left\{ \frac{n_{k,i}}{n_k} \right\} \tag{2.1}$$

where n_k is the total number of citations patent k receives. The most citing organizations are predominantly corporations, so I will use the terms "firms" and "organizations" interchangeably. To construct an aggregate measure, for each year I take the average of patents' concentration measures within each technological class and then the average across technological classes weighted by the number of patents in a class. Appendix A.2 provides more details.

Figure 1 on page 2 shows the resulting aggregate concentration measure. On average, a patent (among the most cited ones) granted between 1976 and 2000 received around 50% of citations from one firm only. This concentration has significantly increased since 2000: a patent granted in 2014 received around 77% of citations from one firm only.

 $^{^{10}}$ In Section 2.6, I also consider the Herfindahl-Hirschman Index (HHI) for the distribution of citations across citing firms.

2.3 Control for Specialization and Number of Patents

The concentration of citations could be explained by technological specialization. Suppose all firms have equal access to the knowledge disclosed in a patent. However, this knowledge is valuable only within a narrow technological space, and only a few firms specialize (patent) in this space. According to this explanation, Amkor is responsible for 94% of citations to IBM's patent in Table 1 because Amkor has more patents than other companies in a technology that benefits from IBM's patent.

I show that the concentration of citations is observed even within narrow technological classes. For instance, contrary to the "specialization" hypothesis, Amkor has fewer patents in its technological space than other companies. Citations are concentrated because a large share of Amkor's patents cite IBM's patent, while other companies do not.

Consider patent k that can receive citations from M firms. Firm i has $N_i > 0$ patents and makes $n_{k,i} \ge 0$ citations. The number of citations can be decomposed into intensive and extensive margins

$$n_{k,i} = p_{k,i} \cdot N_i \tag{2.2}$$

where $p_{k,i} = n_{k,i}/N_i$ is the share of patents in firm *i* that make citations to patent *k*. The concentration measure for patent *k* is defined as

$$C_{k} = \max_{i} \left\{ \frac{n_{k,i}}{\sum_{j=1}^{M} n_{k,j}} \right\} = \max_{i} \left\{ \frac{p_{k,i} \cdot N_{i}}{\sum_{j=1}^{M} p_{k,j} \cdot N_{j}} \right\}$$

The "specialization" explanation assumes that firms do not differ in their citation behavior $(p_{k,i} = p_{k,j} \text{ for } i \neq j)$, but one firm dominates others in terms of the number of patents $(N_i \gg N_j \text{ for } j \neq i)$.

To evaluate this theory, one needs to find all patents that could have potentially made citations to patent k. I divide all granted patents into disjoint groups based on common characteristics. Then, I find all patents that share similar characteristics with patents that *actually* make citations to patent k. For characteristics of patents, I choose an application year and a detailed technological class (main subgroup level in CPC). Below I also describe a robustness check based on textual similarity of patents using BERT model.¹¹

For example, patents making citations to IBM's patent in Table 1 (k = 5877043) are divided in 68 disjoint groups based on the application year and technological class. Most citations (16 out of 218, all from Amkor) come from patents in technological class H01L23 and with application year 2003. Overall, there are 1465 patents with such characteristics, and only 20 of

¹¹Bidirectional Encoder Representations from Transformers (BERT) is a family of language models developed by Google in 2018 (Devlin et al. 2019).

them are assigned to Amkor. Many companies have more patents than Amkor in this class and year, for instance, Intel and Micron Technology have 122 and 101 patents, respectively.

To evaluate the role of "specialization" theory in the concentration of citations, for each patent I do the decomposition from (2.2) within all possible groups of citing patents:

$$n_{k,i}(g) = p_{k,i}(g) \cdot N_i(g)$$

where $n_{k,i}(g)$ is the number of citations from firm *i* to patent *k* from patents within group *g*. For example, for IBM's patent the number of citations from patents in technological class H01L23, with application year 2003, and assigned to Amkor is $n_{k,i}(g) = 16$, where k = 5877043, i = Amkor, and $g = \{H01L23, 2003\}$. In this example, $N_i(g) = 20$ and $p_{k,i}(g) = \frac{16}{20}$ for i = Amkor, and $p_{k,j}(g) = 0$ for $j \neq \text{Amkor}$.

Next, I do two types of Monte-Carlo simulations. First, I randomize $n_{k,i}(g)$ citations across all patents with the same characteristics g. This exercise equates citation rates across firms that have patents with characteristics g ($p_{k,i}(g) = p_{k,j}(g)$ for $i \neq j$). Second, I allocate $n_{k,i}(g)$ citations randomly across firms assuming that all firms with patents in g have the same number of patents. This exercise equates both the citation rates and the number of patents across firms ($p_{k,i}(g) = p_{k,j}(g)$ and $N_i(g) = N_j(g)$ for $i \neq j$). I do these Monte-Carlo simulations for the most cited patents, and then I recompute the aggregate concentration measure from Section 2.

Figure 2 shows the results of Monte-Carlo simulations. The upper solid line shows the actual aggregate concentration (\mathcal{AC}_t) . The dashed and the dotted lines in the middle show the 95th quantile and the median of the concentration in which citation probabilities are equalized across firms (denote by \mathcal{RC}_t the median). Finally, the crossed line at the bottom shows the median concentration with equalized citation rates and equalized patenting across firms (\mathcal{PC}_t).

The variable \mathcal{PC}_t shows a part of the concentration that is driven by the number of firms. For example, with M firms the concentration cannot be lower than 1/M. The measure \mathcal{PC}_t has declined over time, indicating an increase in the number of firms filing patents within the same technological class over the same time period.

The difference $\mathcal{RC}_t - \mathcal{PC}_t$ quantifies the impact of variations in the number of patents among firms, conditional that each firm has at least one patent in a given (application year, technological class) pair. These variations in patenting rates account for only a minor portion of the overall concentration in citations.

Finally, the difference $\mathcal{AC}_t - \mathcal{RC}_t$ shows the role of variations in citation intensities across firms. This difference has significantly increased over time, indicating growing disparities in citation behavior among firms. The difference in citation intensities across firms explains around



Figure 2: Decomposition of the Concentration of Citations

This figure compares the actual concentration of citations (the upper solid line) with the counterfactual ones in which citation rates are equalized across firms (the dashed and the dotted lines in the middle), and in which both citation rates and the number of patents are equalized across firms (the crossed line at the bottom). These counterfactual concentration measures are constructed using Monte-Carlo simulations in which citations are allocated randomly across observationally similar patents (equalized citation rates) and across firms with observationally similar patents (equalized citation rates are given in Appendix A.4.

38% of the concentration in 1980, 50% in 1990, and 79% in 2014.¹² Although the "specialization" theory provides a reasonable approximation to citations until the 1990s, it is inconsistent with the observed changes in citation patterns.

The decomposition in Figure 2 might underestimate the role of specialization in patenting if technological classes are not granular enough to capture this specialization. I use the technological classification which is more granular than the one commonly used in the literature.¹³

As a robustness check, I also use a natural language processing model called BERT to find textual similarity between patents, in addition to considering application years and technological classes. For each cited patent, I measure a similarity between patents citing it. I then select all patents that exhibit at least as much similarity to the citing patents as the citing ones do among themselves. The details are given in Appendix A.4. Figure B1 in Appendix B shows that the results are robust with this more restrictive specification.

Another important patent characteristic that might affect citations is a location of inventors.

¹²The difference in citation intensities explains x% of the concentration where $x = 100 \cdot \frac{AC_t - \mathcal{R}C_t}{\mathcal{A}C_t - \mathcal{P}C_t}$.

¹³The main subgroup level in CPC has 7137 detailed categories while the literature (e.g., Jaffe et al. 1993 and Bell et al. 2019) often considers technologies to be similar if they come from the same 3-digit USPC or NBER sub-class classifications, which have 876 and 445 categories, respectively.

In Section 2.4, I show that the concentration is high even within the same inventor who is located in the same geographical area and works across multiple companies.

The decomposition in Figure 2 might also overestimate the role of specialization in patenting because it excludes citations from never-citing technologies. For example, IBM's patent in Table 1 receives most citations from IBM's supplier, Amkor Technology. Therefore, the citations are allocated randomly across the patents that are similar to Amkor's patents and are likely to represent non-competing technologies to IBM. This randomization exercise does not take into account many patents from IBM's competitors that could have made citations to it.

2.4 Movement of Inventors: Concentration is Explained by Firm Relationships

One of the theories for the agglomeration of economic activity is based on serendipitous face-toface interactions between inventors (Marshall 1920; Saxenian 1996). According to this theory, patents are surrounded by tacit knowledge, and interactions between inventors facilitate the diffusion of this knowledge (Carlino and Kerr 2015). Patent citations are used to measure knowledge spillovers from these interactions and to evaluate their geographical localization (e.g., Jaffe et al. 1993; Atkin et al. 2022).

If citations measure spillovers from inventors' interactions, then the concentration of citations can be explained by limited communication between inventors from different firms. For instance, Amkor might be responsible for 94% of citations to IBM's patent in Table 1 because inventors from IBM exclusively communicate with inventors at Amkor. The central question is whether firms control this communication or if it occurs independently of firm relationships.

I separate the roles of firms and inventors by tracing citations of inventors who filed similar patents in multiple companies. I define patents to be similar if they have close application years, the same narrow technological class, and the same geographical location of an inventor-mover. I show that inventors significantly change their citation probabilities to a particular patent once they move to another company and that the concentration of citations is primarily explained by firm-specific factors. This evidence suggests that firms might have significant control over the diffusion of their knowledge, at least for knowledge flows measured by patent citations.¹⁴

Formally, consider the following statistical framework. Suppose inventor ℓ worked in two companies, *i* and *j*, and created $N_i^{\ell}(g)$ and $N_j^{\ell}(g)$ patents with characteristics *g*, respectively. Assume that each patent in firm *i* (*j*) makes an independent citation to patent *k* with probability

¹⁴For instance, consider an inventor who has learned tacit knowledge associated with a particular technology and then moves to a new firm. This movement does not necessarily imply that the new employer can access the tacit knowledge. The inventor might be bound by a non-disclosure agreement, or the new employer may lack a licensing agreement on patents that complement the tacit knowledge.

 $p_{k,i}^{\ell}(g) \ (p_{k,j}^{\ell}(g))$. Then the expected number of citations from inventor ℓ in firm i to patent k is

$$n_{k,i}^{\ell}(g) = p_{k,i}^{\ell}(g) \cdot N_i^{\ell}(g),$$

and the goal is to test whether $p_{k,i}^{\ell}(g) = p_{k,j}^{\ell}(g)$. To do this, I compare the actual concentration of citations within an inventor with the counterfactual one where citation probabilities are equalized across companies $(p_{k,i}^{\ell}(g) = p_{k,j}^{\ell}(g))$.

For example, during 2008 to 2017 inventor Stefan G. Schreck from California created 16 patents in the technological class A61F2 while working in Endologix Inc. In 15 out of 16 patents, he made a citation to patent 5690642 assigned to Cook Incorporated. He also applied for 9 patents with similar characteristics in another company, Edwards Lifesciences Corporation, but made zero citations to patent 5690642. If 15 citations were allocated randomly across 16 + 9 = 25 patents, the expected share of citations from Endologix Inc. would be $\frac{9.5}{15} = 0.63$ and the 95th quantile would be $\frac{12}{15} = 0.8$. However, the actual share ($\frac{15}{15} = 1$) is significantly higher. Notice that for an inventor who worked in two companies the concentration (the share of citations from the most citing firm) cannot be less than 50%.

I compute the average concentration of citations across firms within all inventors who moved between companies, and do the decomposition from Section 2.3. The details are given in Appendix A.5. Figure 3 shows the actual average concentration within an inventor (\mathcal{AC}_t^w) , the 95th quantile and the median of the concentration with equalized citation rates across firms $(\mathcal{RC}_t^w(q95) \text{ and } \mathcal{RC}_t^w)$, and the median concentration in which both citation rates and patenting are equalized across firms (\mathcal{PC}_t^w) .¹⁵ The actual average concentration within an inventor is significantly higher relative to the what we would expect if citation probabilities were equalized across firms $(\mathcal{AC}_t^w > \mathcal{RC}_t^w(q95))$. The average decomposition over all years

$$\overline{\underline{\mathcal{AC}}^{w} - \overline{\mathcal{PC}}^{w}}_{22.1\%} = \overline{\underline{\mathcal{AC}}^{w} - \overline{\mathcal{RC}}^{w}}_{17.0\%} + \overline{\underline{\mathcal{RC}}^{w} - \overline{\mathcal{PC}}^{w}}_{5.1\%}$$

shows that the concentration is primarily explained by the differences across firms in citation probabilities rather than by the variance in the number of patents. The difference $\mathcal{AC}_t^w - \mathcal{RC}_t^w$ is stable over time, and there was a slight increase in the difference $\mathcal{RC}_t^w - \mathcal{PC}_t^w$, meaning that the dispersion in the number of patents across firms within an inventor has slightly increased by the end of the period.

$$\frac{1}{4} \cdot 1 + \frac{1}{4} \cdot 0.5 + \frac{1}{4} \cdot 0.5 + \frac{1}{4} \cdot 1 = 0.75$$

¹⁵The variable \mathcal{PC}_t^w is greater than 50% because the majority of inventor-movers worked in two firms only. For example, with two citations randomly allocated across two firms the expected concentration measure is



Figure 3: Decomposition of the Concentration Within Inventors Who Moved Across Firms

This figure shows the same decomposition from Figure 2 for the concentration of citations across firms within inventors who patented in multiple companies. The solid line shows the actual aggregate within-inventor concentration of citations across firms. The dashed and the dotted lines show the 95th quantile and the median of the same measure in the Monte-Carlo simulations where citations rate are equalized across firms within an inventor. The crossed line at the bottom shows the concentration (median) where both the citation rates and the number of patents are equalized across firms within an inventor. The averages are almost the same as the medians. To increase the sample size I consider citations from all years, not only 5-year window. The graph is taken until 2009 to ensure that patents have enough time to accumulate citations from inventors-movers. More details are given in Appendix A.5.

This evidence should be interpreted with caution because inventors-movers might differ in their citation rates from inventors who always work in one company. My conjecture is that nonmovers would have higher concentration of citations across firms if they were randomly moved to another company. Below I argue that citations are correlated with access to trade secrets. Based on this interpretation, the conjecture is that inventors who do move between companies are less bound by contractual obligations, such as confidentiality agreements and non-compete clauses, resulting in a less concentrated distribution of citations. An interesting area for future research is to study the movement of inventors caused by exogenous shocks to firms, for example, natural disasters (Barrot and Sauvagnat 2016) or financial constraints (Chodorow-Reich 2014).

2.5 Robustness and Additional Results

Figures B2 and B3 in Appendix B provide additional results on the concentration. Figure B2 shows that the increase in the concentration is primarily driven by changes within technological classes rather than the rise of technologies with high concentrations of citations. An increase in

the concentration after 2000 is observed in 87% of classes. Panel (a) of Figure B3 shows that the average number of citations has significantly increased over time: the most cited patents granted in 2014 received 11 times more citations within five years from the grant day than patents granted in 1976. Panel (b) shows that more cited patents have a higher concentration of citations, and that this relationship is driven by patents granted after 2000.¹⁶ Therefore, the increase in the concentration of citations is not driven by the decline in the number of citations.

Figures B4 – B6 provide multiple robustness checks. In particular, the results are robust to different thresholds for the most cited patents (top 5% and 10%). The concentration is not driven by superstar firms in patenting, and it is robust to the exclusion of firms' self-citations to themselves. The concentration is also robust when I group citations of patents from the same within-country family (continuations, continuations-in-part, and divisionals) as a single citation, indicating that its rise is not driven by increasing patent families. Citation patterns might be affected by patent examiners (Alcácer et al. 2009). I show that the concentration of citations from patent examiners is around two times lower than the concentration based on citations from non-examiners. Citations might also be affected by patent attorneys. In Section 2.4, I use the movement of inventors across companies to show that the concentration of citations is driven by firms rather than inventors. Using the same technique, I show that the concentration of citations is not driven by patent attorneys either. Appendix A.3 describes more robustness checks, including controls for outliers, different samples, and weighting schemes.

2.6 Concentration at the Firm Level

In Section 2.2, the concentration of citations for the most cited patents is defined as a share of citations coming from the most citing firm. This definition is easy to interpret and allows for detailed controls at the patent level. However, it omits citations to patents that are not among the most cited ones and citations from firms other than the top citing entity.

This section extends the concentration measure in two directions. First, I consider citations to all patents of a given firm, for example, all citations to IBM's patents granted in a particular year rather than to each individual patent. Defining the concentration measure at the firm level will be helpful in Section 3 because the data on relationships between firms are not available at the patent level. Second, I compute the concentration using the Herfindahl-Hirschman Index (HHI) rather the share of citations from the most citing firm.

For each firm in a year, I define the concentration measure as the Herfindahl-Hirschman Index (HHI) of all citations to the firm's patents granted in this year within a five-year period.

 $^{^{16}}$ The relationship holds for patents with more than seven citations. For patents with fewer than seven citations, the concentration is high due to a low number of citations. See Section 2.3 for more details.

The aggregate concentration in year t is defined as

$$\mathcal{H}_t = \sum_k s_{k,t} \mathcal{H}_{k,t} \tag{2.3}$$

where the summation goes across all firms receiving citations to their patents granted in year t, $s_{k,t}$ is the share of citations firm k receives in year t in the total number of citations firms receive in this year, and $\mathcal{H}_{k,t}$ is the HHI concentration of citations for firm k.

Figure B7 in Appendix B shows that the firm-level concentration follows the same pattern as in Figure 1: a stable or slightly declining concentration of citations until 2000 and a significant rise after. This result is robust to different sample selections: I consider citations among all corporate patents, excluding self-citations, and between U.S. publicly listed companies only. The results are also robust if the concentration is defined as the share of citations from the most citing firm, as in Section 2.2. The additive structure of the HHI measure will be useful for evaluating the role of partners in the concentration in Section 3.

In Section 3, I evaluate the role of business partners in the sample of inter-firm citations among publicly traded companies after 2001. Figures B8 and B9 in Appendix B provide more details on the rise in the concentration of citations in this sample. Figure B8 decomposes the increase in the concentration from 2001 to 2014 based on the methodology of Melitz and Polanec (2015). The rise in the concentration is primarily explained by the increasing concentration of citations within firms rather than the re-allocation of citations across firms or the exit and entry of new firms out of and into patenting. Figure B9 shows the decomposition of the concentration has increased due to changes in citation probabilities across firms rather than shifts in patent counts.

3 The Role of Business Partners

This section evaluates the role of business partners in the concentration of citations. Section 3.1 describes the data. Section 3.2 shows that business partners account for most citations between firms. The concentration of citations has increased due to changes within firms in how they receive citations from partners. Section 3.3 discusses the results and provides some robustness checks.

3.1 Data

I use three data sets to find the types of relationships between cited and citing firms. First, I use the FactSet Revere Supply Chain Relationships data set, which is based on public sources such as filings with the U.S. Securities and Exchange Commission (SEC), investor presentations, and press releases to collect data on business relations between firms. The data list partners such as suppliers and customers, firms with licensing agreements, research collaborations, joint product offerings, and firms with ownership stakes (e.g., joint ventures). Second, I use the Compustat Segments to collect data on supplier–customer relationships between firms. Finally, I use the USPTO data on patent reassignment to find firms that have traded patents with each other. The data sets cover the period from 2003 to 2022.¹⁷ I match all data sets with patents using company names. Details are given in Appendix A.1.

3.2 Business Partners in Patent Citations

To document the role of partners in patent citations, I focus on inter-firm citations between U.S. publicly listed firms for patents granted after 2001. The sample selection is driven by data limitations.¹⁸ I discuss how these limitations might affect the results at the end of this section. I consider patents granted until 2014 and trace citations within a five-year window from the grant date. The large role of partners is robust to alternative time windows.

I compute the distribution of citations across firms based on their relationship with a cited firm. I define firms as business partners if they had at least one of the following relationships between 2003 and 2022:¹⁹ suppliers or customers, research collaborations, licensing agreements, joint product offering, patent reassignment, joint ownership stakes (e.g., a joint venture), and partners with an uncertain relationship.²⁰

Panel (a) of Figure 4 shows that 76% of citations occur between business partners. This share is 72% if I consider citations within a 10-year window. Panel (b) shows the distribution of citations among partners based on a relationship between firms. If two firms have multiple relationships, I divide citations evenly across them. Firms with a supplier–customer relation and research collaboration account for the majority of citations: 39% and 27%, respectively. Other types of partners account for 34% of citations and include firms that signed licensing agreements, reassigned patents, had joint ownership stakes, and had joint product offerings.

 $^{^{17}}$ FactSet covers most relationships between firms, though its coverage only started from 2003. In general, Compustat Segments and USPTO reassignment data provide coverage since 1976 and 1968, respectively.

¹⁸The data on relationships cover the period from 2003 to 2022. I include the years 2001 and 2002 to analyze the rise in the concentration since 2001, but the results are similar without these years. The coverage of relationships is limited for private and foreign firms.

¹⁹Around 75% of citations between partners occur during the period of a reported relationship. The data are based on information disclosed by firms in public sources, but the disclosed dates often do not correspond to the actual dates of a contractual relationship between companies. Firms might delay the disclosure of a relationship relative to its actual start and might stop reporting it before it actually ends. Therefore, 75% provides a lower bound on the share of citations occurring during a business relationship.

 $^{^{20}}$ Table B1 in Appendix B provides the full list of all business relationships. FactSet defines some partners without clarifying the nature of the relationship. These account for only 0.7% of citations between partners.



Figure 4: Distribution of Citations Based on Relationship Between Firms

This figure shows the distribution of patent citations across different types of relationships between cited and citing firms. Panel (a) shows the share of citations coming from business partners. Panel (b) shows the distribution of citations coming from business partners across different types of partners.

As I discuss in Section 3.3, the distribution of citations across partnerships might be biased because types of relationships are self-reported. The correct interpretation of panel (b) is that partners responsible for most citations had *at least* a supplier–customer relationship or research collaboration. It is reasonable to assume that they also had licensing agreements.

To evaluate the importance of partners in the concentration of citations, I decompose the HHI concentration within a firm (2.3) into citations from partners and other firms:

$$\mathcal{H}_{k,t} = \sum_{i} \left(\frac{n_{k,i,t}}{n_{k,t}}\right)^2 = \mathcal{H}_{k,t}^{\mathcal{P}} + \mathcal{H}_{k,t}^{\mathcal{O}} = \underbrace{\sum_{i \in \mathcal{P}_k} \left(\frac{n_{k,i,t}}{n_{k,t}}\right)^2}_{\text{Partners },\mathcal{H}_{k,t}^{\mathcal{P}}} + \underbrace{\sum_{i \in \mathcal{O}_k} \left(\frac{n_{k,i,t}}{n_{k,t}}\right)^2}_{\text{Other },\mathcal{H}_{k,t}^{\mathcal{O}}}$$
(3.1)

where $n_{k,i,t}$ is the number of citations from firm *i* to firm *k* for patents granted in year *t*, $n_{k,t}$ is the total number of citations firm *k* receives to patents granted in year *t*, $i \in \mathcal{P}_k$ means that firm *i* is a business partner to firm *k*, and $i \in \mathcal{O}_k$ means that firm *i* is not a (revealed) business partner to firm *k*.

Figure 5 shows the decomposition into partners and other firms from equation (3.1). On average, partners explain around 66% of the concentration in citations: 59% in 2001 and 75% in 2014. Changes in citations among partners account for 84% of the rise in the concentration. Suppliers, customers, and firms with research collaboration account for 60% of the rise in the concentration. Figure B10 in Appendix B offers more details.

To provide further evidence, I also analyze citations by industries of the cited and citing



Figure 5: Decomposition of the Concentration Based on Relationship Between Firms

This figure decomposes the concentration into the roles of partners and other firms; see equation (3.1).

firms. Figure B11 in Appendix B shows that firms within the same industry account for only 8% of the increase in the concentration of citations. Therefore, the rise in the concentration is largely driven by changes in citations among firms that are unlikely to be direct competitors.²¹

Figure 6 provides more details on the rise in the concentration of citations among partners. Panel (a) shows that, from 2001 to 2014, the average number of citing partners for a typical firm declined by 42%, even though the number of (reported) partners increased over time. Panel (b) shows that citations became more skewed within citing partners: over time, the actual concentration of citations significantly deviated from the counterfactual one in which citations are allocated uniformly across citing partners. In Figure B12 from Appendix B, I do an exercise similar to the one from Section 2.3 to show that changes in partners' composition and patenting activity cannot explain the rise in the concentration of citations among partners.²²

To sum up, most inter-firm citations among U.S. publicly traded companies occur between

 $^{^{21}}$ In general, defining competitors requires a definition of the market, and industry affiliation might be an imperfect proxy for competitors. Nevertheless, firms from the same industry are more likely to compete with each other than with firms from different industries. Industries are defined at the four-digit level in Standard Industry Classification.

²²Consider patents assigned to firm A and citing a patent of firm B. I find all patents with the same characteristics as the citing ones (see Section 2.3) and assigned to firms with the same relationship to firm B as the relationship between A and B during a comparable period. For instance, if A is a supplier to B in 2005, I find other suppliers to B reported from 2003 to 2007 that had patents with the same application year and narrow technological class as patents from firm A. Then, I randomize citations to equalize citation rates and patenting across partners as in Section 2.3. Citations are highly concentrated even among partners filing patents from the same narrow classes and application years and with the same relationship to the patent owner. Moreover, the concentration primarily increased due to changes in citation probabilities among partners.



Figure 6: Rise in the Concentration of Citations Among Partners

This figure shows the roles of extensive and intensive margins in the concentration of citations within partners. Panel (a) compares the average number of citing partners relative to the total number of partners for a typical firm. For each firm, the number of partners is weighted by the number of citations this firm receives. The total number of partners of a firm in a given year is defined using partners reported within five years of this year. Panel (b) shows the average concentration of citations (HHI) from business partners. The dotted line shows the counterfactual concentration, in which citations are distributed across firms as uniformly as possible to minimize the HHI, given the existing number of citing firms. Both graphs are constructed based on a sample of patents with a unique assignee. Patents with more than one assignee account for less than 1% of citations.

business partners. Moreover, the rise in the concentration of citations is primarily explained by the decline in citations among partners. Below, I discuss data limitations and additional robustness checks. In Section 4, I develop and test a theory that can rationalize the evidence on patent citations.

3.3 Discussion

In this section, I discuss data limitations and whether strategic withholding of citations can explain the high share of partners in citations.

Data Limitations

The data on inter-firm relations have some limitations because they are primarily based on the disclosure of relationships by publicly traded companies. Therefore, the data do not cover relations between all firms. As a robustness check, I focus on a sample of citations in which at least one of the firms is a U.S. publicly traded company, but not necessarily both. Figure B13 in Appendix B shows that partners explain 51% of citations, and the distribution within partners is similar to Panel (b) in Figure 4. The share of partners is lower in this sample, but it might reflect the incompleteness of the data for private and foreign firms rather than the absence of a relationship. Moreover, Figure B7 in Appendix B shows that the concentration dynamics are consistent whether considering all firms or just U.S. publicly traded companies. Therefore, I expect a minimal bias regarding the role of partners in the rise of the concentration when focusing on the sample of U.S. publicly traded companies.

Another limitation of the data is related to firms' incentives to disclose relationships with other firms. Although regulation SFAS No. 131 requires publicly traded companies to report the identity of any customer representing more than 10% of their total sales, smaller customers or other types of relationships are self-reported. Firms' incentives to disclose a relationship with another firm might differ across types of relationships. For example, firms might have stronger incentives to conceal the identity of research collaborators relative to established input suppliers. In addition, the data lack details on contracting between firms, and certain types of relationships are not mutually exclusive. For instance, firms with a research collaboration might not report that they also have licensing agreements as part of the collaboration. Therefore, the comparison of citation patterns based on the type of partnership should be interpreted with caution.

Strategic Citations

The large share of partners in patent citations could be explained by strategic withholding of citations. U.S. patent applicants must disclose all relevant prior art that they are aware of. However, firms might potentially withhold citations to make patent claims broader (Lampe 2012). If firms follow this strategy, then partners are less likely to withhold citations because it is harder for them to prove they were unaware of the prior art.

There are several arguments for why strategic withholding of citations is unlikely to explain the citation patterns. First, withholding of citations can make the whole patent unenforceable, not just specific claims (Ouellette and Masur 2023). Therefore, this practice bears considerable risks, especially for patent attorneys. Second, Kuhn et al. (2023) provide empirical evidence and legal arguments that firms are unlikely to practice the strategic withholding of citations.

Finally, I test this hypothesis using the movement of patent attorneys across business partners of a given firm. Similar to the exercise on the movement of inventors, I find attorneys who filed similar patents in multiple companies. In addition, I restrict the sample to citations among partners and attorneys moving from one partner of a given firm to another. If citations were strategic, such attorneys would be expected to cite a patent from all partners of the patent owner. In contrast, Figure B14 shows that even in this sample, citations are highly firm-specific: an attorney who heavily cites a particular patent in one firm does not continue citing it in another firm, despite both firms being partners to the patent owner.

4 Patenting, Secrecy, and Knowledge Sharing

Several surveys of inventors and firms show that patent citations are correlated with knowledge flows (e.g., Jaffe et al. 2000; Duguet and MacGarvie 2005). In the previous sections, I show that citations predominantly come from a few business partners of patent owners. I provide various tests to determine whether the citation patterns can be interpreted as involuntary spillovers. I also provide robustness checks for many issues with patent data raised in the literature. Overall, the evidence suggests that knowledge flows measured by citations are intentional rather than involuntary spillovers. Furthermore, citations have become more concentrated over time, suggesting a decline in knowledge sharing among collaborating firms.

This interpretation of citations is not entirely new to the empirical literature on knowledge spillovers, as it is commonly assumed that self-citations do not represent knowledge spillovers (e.g., Jaffe et al. 1993). However, given the significant role of partners in patent citations, it would be more accurate to differentiate inter-firm citations and self-citations as *intentional* knowledge flows between and within firm boundaries rather than categorizing them as spillovers and non-spillovers.

It is important to note that patent citations do not capture all knowledge flows. Thus, the evidence in this paper does not imply the general absence of knowledge spillovers. However, understanding the citation patterns is important for at least two reasons. First, citations are one of the primary measures of knowledge spillovers in the literature.²³ Second, the U.S. Supreme Court views the patent system as a central mechanism for knowledge dissemination.²⁴

In this section, I develop a theory that can rationalize the citation patterns and explain how firms control the diffusion of their knowledge. The idea is that firms choose a combination of patenting and secrecy to balance the benefits of patent rights in knowledge sharing with partners against the cost of knowledge disclosure to competitors. Section 4.1 discusses several examples of how firms combine patenting with secrecy. Section 4.2 develops a formal model. Sections 4.3 and 4.4 test two predictions from the model. Finally, Section 4.5 discusses the implications for the innovation literature.

4.1 Combining Patenting and Secrecy: Examples

There are multiple examples on how firms combine patenting and secrecy to protect the same technology. For instance, patents on biological products, including mRNA-based COVID-19 vaccines, do not disclose manufacturing details (Price and Rai 2016; Price et al. 2020). The undisclosed knowledge is not necessarily tacit: the manufacturing details are codified and

²³See footnote $\frac{2}{2}$ and the literature review on page 5.

 $^{^{24}}$ See footnote 4.

submitted to the Food and Drug Administration but not disclosed to the public. Patent applications on these products are filed very early in the technology development. At this stage, many details about the technology are unknown, and firms are not required to disclose them later (Ouellette 2016).

Arora (1997) describes how dyestuff producers in the 20th century patented individual chemical compounds but kept the knowledge of how to combine them secret. Similarly, "[the] Haber-Bosch process for ammonia, a truly significant process innovation, was protected by more than 200 patents that covered the apparatus, temperatures, and pressures, but avoided particulars about the catalysts employed or their preparation. The catalyst was critical to the successful operation of the process, and keeping it secret significantly increased the expense and time for firms trying to circumvent the Haber-Bosch patent" (Arora 1997, p. 394).

Lim (2009) describes the development of the copper interconnect technology for semiconductor chips. IBM, the prominent leader in this technology, disclosed "general ideas relating to the damascene process and copper technology" but kept secret "the type of material it uses for the barrier layer and the chemical composition of its electroplating bath" and other "process-specific knowledge" (Lim 2009, p. 1261).

In Section 4.3, I use trade secret litigation to find patents that were likely to be bundled with trade secrets. For instance, the legal case *Waymo LLC v. Uber Technologies, Inc.* was about misappropriation of the trade secrets related to the LiDAR technology for self-driving cars.²⁵ However, the same lawsuit also had claims regarding patent infringement for three patents.²⁶ In the complaint, Waymo describes how these patents and trade secrets are complementary to each other: "The Replicated Board reflects Waymo's highly confidential proprietary LiDAR technology and Waymo trade secrets. Moreover, the Replicated Board is specifically designed to be used in conjunction with many other Waymo trade secrets and in the context of overall LiDAR systems covered by Waymo patents." Similar to the biological products, the undisclosed knowledge is not necessarily tacit: "Waymo's trade secret information further includes the resolution profile that is achieved through its proprietary positioning and orientation of laser diodes and optical elements in its current generation LiDAR system."

These examples highlight that technologies often consist of multiple pieces of knowledge, some of which are patented while others are kept secret. Therefore, replicating and building upon such technologies requires access not only to patented knowledge but also to trade secrets.

 $^{^{25}\,}Waymo\ LLC\ v.$ Uber Technologies, Inc., No. 17-2235 (Fed. Cir. 2017).

²⁶The patent numbers are 8836922, 9368936, and 9086273.

4.2 Model of Information Disclosure and Knowledge Sharing

I start with a simple model of licensing between a technology provider (supplier) and adopters (customers). I show how the model can be mapped to patent citations. Then, I provide various extensions, including a more general production structure and other types of partnerships among firms. Finally, I discuss assumptions and predictions of the model.

A product can be produced using a constant returns to scale production function $F(z_1x_1, z_2x_2)$, where x_1 and x_2 are two inputs, z_1 and z_2 are input productivity levels (knowledge components). A marginal cost of producing one unit of the product is denoted by $c(z_1, z_2)$:

$$c(z_1, z_2) = \min_{x_1, x_2} \{ w_1 x_1 + w_2 x_2 \}$$

s.t. $F(z_1 x_1, z_2 x_2) = 1$

where w_i is a price of input $i \in \{1, 2\}$.

There are *n* markets, and each market has one manufacturer who can produce and sell the product in this market. Denote by $v_j(c_j)$ the profit of producer *j* as a function of its marginal costs, c_j . Profits $v_j(c_j)$ are decreasing in c_j . The marginal cost c_j depends on producer *j*'s access to knowledge about input productivity levels, $c_j = c(z_{1,j}, z_{2,j})$. Producer *j* always has free access to a generic technology with $z_{1,j} = z_{2,j} = \underline{z}$. However, it can also buy a superior technology from technology providers. To simplify notation, I will denote $v_j(c(z_{1,j}, z_{2,j}))$ as $v_j(z_{1,j}, z_{2,j})$.

An incumbent technology provider knows a technology with $z_1 = z_2 = \overline{z} > \underline{z}$. The knowledge for each component $i \in \{1, 2\}$ of the technology can be patented or kept secret. If it is patented, the incumbent can license it to the manufacturers (customers). If it is kept secret, once the incumbent shares it with the customers, they can use it without compensating the incumbent.

Patenting also makes knowledge publicly available. There are n_e entrants who can improve productivity. Entrant k can discover a new pair $(z_{1,k}^e, z_{2,k}^e)$:

$$z_{i,k}^e = \psi_k \cdot Z_i, \ i \in \{1, 2\}$$

where $Z_i = \overline{z}$ if component *i* is patented, $Z_i = \underline{z} < \overline{z}$ if it is kept secret, and ψ_k is a random draw from a continuous distribution $\Psi(\cdot)$ with the support on $[0, \infty)$. The random variables ψ_k are independent across entrants.

The competition between the incumbent and the entrants consists of four stages. First, upon receiving ψ_k , entrant k can observe whether $c(z_{1,k}^e, z_{2,k}^e) < c(\overline{z}, \overline{z})$. It decides whether to patent $z_{1,k}^e$ and $z_{2,k}^e$ or not. Patenting involves a small fixed cost $\delta > 0$. Second, both the incumbent and the entrants can observe $c(z_{1,k}^e, z_{2,k}^e)$ for $k \in \{1, 2, \ldots, n_e\}$, and they decide whether to enter



Figure 7: Timing in the Model

the market for technologies, which requires paying a small fixed cost $\varepsilon > 0$. Third, providers who have paid ε decide whether to share their secrets with each customer, given that trade secret sharing is non-contractual. Finally, these providers make simultaneous licensing offers on their patented knowledge to each customer. Each customer requires only one technology from one provider, and knowledge components from different providers are incompatible. In the case of indifference, I assume a customer accepts the offer from the most efficient provider. This provider is unique with probability one by continuity of $\Psi(\cdot)$.

I assume that fixed costs δ and ε are sufficiently small and ignore them in profit equations. The fixed cost of patenting (δ) ensures that inefficient firms do not patent. The fixed cost of entry into the market (ε) simplifies the equilibrium because only the most efficient firm will pay it and set monopoly prices.

The summary of the timing in this game is the following (see also Figure 7):

- t1. For each knowledge component, the incumbent decides whether to patent it or keep it secret;
- t2. The entrants receive independent productivity draws, $\psi_k \sim \Psi(\cdot)$ for $k \in \{1, 2, \ldots, n_e\}$;
- t3. Each entrant observes whether the marginal cost of its technology is lower than the incumbent's one. The entrants decide whether to patent their knowledge;
- t4. All technology providers observe marginal costs of each other. They decide whether to enter the market for technologies;
- t5. Providers that have entered make separate decisions for each customer on whether to share their secrets with that particular customer;
- t6. Providers that have entered make simultaneous licensing offers on their patented knowledge.

To simplify the equilibrium, I make the following assumptions:

Assumption 1. The cost function is symmetric, $c(z_1, z_2) = c(z_2, z_1)$

This assumption simplifies the analysis because only the number of knowledge components that are patented or kept secret will matter in equilibrium, not their identity.

Assumption 2. A profit function of customer $j \in \{1, 2, ..., n\}$ satisfies the following conditions:

- 1. Gains from licensing: $v_j(\overline{z},\overline{z}) v_j(\underline{z},\underline{z}) \ge \max\{v_j(\overline{z},\overline{z}) v_j(\underline{z},\overline{z}), v_j(\overline{z},\underline{z}) v_j(\underline{z},\underline{z})\}$
- 2. Imperfect Substitution: $v_j(\overline{z}, \overline{z}) v_j(\underline{z}, \overline{z}) > 0$

The first condition specifies that the incumbent prefers licensing both knowledge components rather than just one. Suppose both components are patented, and the incumbent makes a takeit-or-leave-it offer to customer j. In that case, it will set the licensing fee to $v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \underline{z})$. The third condition states that knowledge components are not perfect substitutes. It is still valuable to license one component even if customer j already has access to another.

Equilibrium

A perfect subgame equilibrium of this game is the following. Since each customer requires only one technology, in the last two stages, the most efficient provider wins the competition and serves all customers. Anticipating this outcome, only the most efficient provider pays the entry fixed cost and makes take-it-or-leave-it licensing offers to the customers in the final stage. In the third stage, entrants have no incentives to conceal their knowledge because there is no future entry. For entrants with $c(z_{1,k}^e, z_{2,k}^e) < c(\overline{z}, \overline{z})$, the probability of having the most efficient technology is positive. Since the cost of patenting δ is sufficiently small, these entrants patent both components of their knowledge.

In the first stage, the incumbent faces the following trade-off. Complete secrecy eliminates licensing opportunities: once the incumbent shares its private knowledge to a customer, the customer can use the technology without paying for it. Therefore, this strategy is not optimal. In the extensions, I show how this strategy might be optimal with richer production structure.

Patenting both knowledge components discloses all information to the entrants, but, conditional on surviving the competition, the incumbent can extract maximum surplus from the customers. The incumbent's expected profits are

$$\Pi_{\text{pat, pat}} = \left(\sum_{j=1}^{n} \left[v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \underline{z}) \right] \right) \cdot \left[\mathbb{P} \left(\psi_k \leqslant 1 \right) \right]^{n_e}$$

where $v_j(\overline{z},\overline{z}) - v_j(\underline{z},\underline{z})$ is a licensing fee to customer j, and $[\mathbb{P}(\psi \leq 1)]^{n_e} = [\Psi(1)]^{n_e}$ is the probability that all entrants have marginal costs higher than the incumbent.²⁷

²⁷This probability is equal to $\mathbb{P}(\min_k \{c(\psi_k \cdot \overline{z}, \psi_k \cdot \overline{z})\} \ge c(\overline{z}, \overline{z}))$ and simplifies to $[\mathbb{P}(\psi_k \le 1)]^{n_e}$ due to the assumption on the constant returns to scale in the production function and independence of ψ_k across entrants.

Suppose the incumbent patents one knowledge component while keeping another secret. If it survives the competition, it has to decide whether to share this secret with each customer. The following lemma shows that the incumbent shares trade secrets with a customer if the patented and the secret knowledge are complementary in the customer's profits:

Lemma 1. The incumbent firm relying on a patent-secrecy combination (partial disclosure) shares its trade secrets with customer j if and only if $v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \overline{z}) \ge v_j(\overline{z}, \underline{z}) - v_j(\underline{z}, \underline{z})$.

Proof. If the incumbent does not share its secrets with customer j, it can charge $v_j(\overline{z}, \underline{z}) - v_j(\underline{z}, \underline{z})$ for its patented knowledge. If it shares the secrets, it can charge $v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \overline{z})$. Therefore, the incumbent shares its secrets with customer j if $v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \overline{z}) \ge v_j(\overline{z}, \underline{z}) - v_j(\underline{z}, \underline{z})$. \Box

The complementarity of knowledge components in profits of customer j depends on the technological complementarity between inputs in the production function and market conditions in market j, defined by the curvature of $v_j(\cdot)$. Formally, for small productivity changes around $(\underline{z}, \underline{z})$, the knowledge components are complements (substitutes) for customer j if the following expression evaluated at $(\underline{z}, \underline{z})$ is positive (negative):

$$\frac{\partial^2 v_j}{\partial z_1 \partial z_2} = v_j''(c) \frac{\partial c}{\partial z_1} \frac{\partial c}{\partial z_2} + v_j'(c) \frac{\partial^2 c}{\partial z_1 \partial z_2}$$

For example, suppose the production function is Cobb-Douglas, $F(z_1x_1, z_2x_2) = (z_1x_1)^{\frac{1}{2}} (z_2x_2)^{\frac{1}{2}}$, and input prices are one, $w_1 = w_2 = 1$. Suppose customer j faces demand $Q_j(p) = A_j - p$ in its market, where p is a product price, and A_j is some constant. The monopoly profit of customer jis $v_j(c_j) = (A_j - c_j)^2/4$, where $c_j(z_{1,j}, z_{2,j}) = (z_{1,j}z_{2,j})^{-\frac{1}{2}}$. The knowledge components are complements if $A_j \cdot \underline{z} \leq 2$ and substitutes otherwise.

Denote by \mathcal{N}_c the set of partners for whom the knowledge components are complements, and by \mathcal{N}_s the set of partners for whom they are substitutes. I assume that these two sets are non-empty:

Assumption 3. There is at least one partner for whom the incumbent's knowledge components are complements $(\mathcal{N}_c \neq \emptyset)$, and at least one partner for whom they are substitutes $(\mathcal{N}_s \neq \emptyset)$.

The incumbent's expected profits under partial disclosure are

$$\Pi_{\text{pat, sec}} = \left(\underbrace{\sum_{j \in \mathcal{N}_s} \left[v_j(\overline{z}, \underline{z}) - v_j(\underline{z}, \underline{z}) \right]}_{j \in \mathcal{N}_c} + \underbrace{\sum_{j \in \mathcal{N}_c} \left[v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \overline{z}) \right]}_{j \in \mathcal{N}_c} \right) \cdot \left[\mathbb{P} \left(\psi \leqslant \frac{c(\overline{z}, \underline{z})}{c(\overline{z}, \overline{z})} \right) \right]^{n_e}$$

By Assumption 2, the licensing profits are lower under the patent-secrecy strategy relative to full patenting. However, the probability of surviving the competition goes up because less knowledge is disclosed to the entrants.

In the rest of this section, I assume that Assumptions 1-3 are satisfied. Under the full patenting, all n customers get access to both knowledge components of the incumbent's technology. Under the partial disclosure, the customers for whom the knowledge components are substitutes get access to the patented component, but the incumbent does not share the secret knowledge with them. This discussion proves the following proposition:

Proposition 1. Bundling patents with secrets leads to more concentrated knowledge flows relative to full patenting.

Proof. Follows from Lemma 1 and the discussion in the text.

The intuition behind this proposition is the following. Reliance on secrecy impedes contractible knowledge sharing and affects how the incumbent sells its technology to the customers. The incumbent shares its secrets with a customer only if the customer's access to these secrets raises the value of the patented knowledge. As a result, secrecy creates inequality among customers: some access patented and secret knowledge, while others access patented knowledge only. Even though the incumbent sells the same technology, customers are different due to variations in their market conditions, defined by the curvature of $v_j(\cdot)$. In contrast, if both knowledge components are patented, the variation in market conditions affects only the price of the technology but not how the incumbent sells it.

The reliance on secrecy is endogenous. In particular, it depends on the degree of competition measured by the number of entrants, n_e . The next proposition summarizes the relationship between the competition and knowledge flows:

Proposition 2. *Higher competition (number of entrants) leads to more concentrated knowledge flows.*

Proof. If $n_e = 0$, then $\Pi_{\text{pat, pat}} \ge \Pi_{\text{pat, sec}}$. If $n_e \to \infty$, $\Pi_{\text{pat, pat}}/\Pi_{\text{pat, sec}} \to 0$ for the non-degenerate cumulative distribution function $\Psi(\cdot)$. Since $\Pi_{\text{pat, pat}}/\Pi_{\text{pat, sec}}$ is continuous and monotonically decreasing in n_e , there is n_e^* such that for $n_e \ge n_e^*$ partial disclosure dominates full patenting. Then, the result follows from Proposition 1.

The intuition behind this proposition is the following. A higher number of competing entrants raises the cost of knowledge disclosure. Thus, the incumbent is more likely to rely on secrecy, which leads to more concentrated knowledge flows by Proposition 1.

Mapping Theory To Patent Citations

Consider the following stylized mapping of the model to citations. Suppose the final products require two tasks: upstream and downstream. Technology providers (suppliers) handle the upstream task, while manufacturers (customers) handle the downstream task. Each provider has one patent containing information about one or two knowledge components. Each manufacturer patents and cites the provider's patent with probability \underline{q} if it gains access to one knowledge component, and with probability $\overline{q} > q$ if it gains access to both components.

If the incumbent patents both knowledge components, the expected number of citations to its patent from customers is $n \cdot \overline{q}$. If it patents only one knowledge component, the expected number of citations is $n_s \cdot \underline{q} + n_c \cdot \overline{q}$, where $n_c = |\mathcal{N}_c|$ is the number of partners for whom the upstream knowledge components are complements, and $n_s = |\mathcal{N}_s|$ is the number of partners for whom they are substitutes. In the latter case, citations are expected to be more concentrated because customers cite with different probabilities. For example, if $\overline{q} = 1 > 0 = \underline{q}$, the concentration of citations is $\frac{1}{n_s+n_c}$ under full patenting and $\frac{1}{n_c} > \frac{1}{n_s+n_c}$ under partial secrecy.

In the model, all entrants with $c(z_{1,j}^e, z_{2,j}^e) < c(\overline{z}, \overline{z})$ patent their knowledge. Since they rely on the incumbent's knowledge, assume that they cite its patent. Suppose an increase in the number of entrants from n_e to n'_e leads to the shift from complete to partial disclosure, which leads to more concentrated citations among partners. However, the effect on citations from the entrants is ambiguous. In the original equilibrium, the expected number of citations from the entrants is $n_e \cdot \mathbb{P}(\psi > 1)$. In the new equilibrium, this number is $n'_e \cdot \mathbb{P}\left(\psi > \frac{c(\overline{z}, z)}{c(\overline{z}, \overline{z})}\right)$. While there are more entrants, each of them cite with lower probability.

The described mapping of the model to citations cannot fully explain the data. However, it might broadly account for the trends documented in Section 3. Since 2000, the concentration of citations has significantly increased, primarily due to a decline in citations from business partners. The changes in citations among firms within the same industry (a proxy for competitors) explain only a small portion of this trend. In the model, higher competition from potential entrants leads to the shift from full patenting to patent-secrecy combination. This change increases the concentration of citations among partners. However, it has less impact on competing entrants because the higher number of entrants is offset by a lower probability of successfully replacing the incumbent firm.

Extensions

In Appendix C, I consider various extensions to the model. First, I incorporate investments in R&D. This extension clarifies that competition in this model reduces incentives to innovate (Schumpeter 1942). Second, I consider a model with a continuum of knowledge components, accounting for tacit and reverse engineerable knowledge. The presence of tacit or reverse engineerable knowledge limits a firm's choice of intellectual property protection strategies. However, knowledge sharing with partners remains unchanged, conditional on this choice. Third, I consider an extension where the incumbent can manufacture the product. In this

case, it might be optimal to rely entirely on secrecy. Finally, I consider more general relationships between partners, such as the incumbent sharing knowledge with its supplier. If the supplier serves the incumbent and other customers, the incumbent can use patent rights to extract surplus from sales to others.

Discussion

Several caveats are worth discussing. First, the model does not capture all the benefits and costs of patenting described in the literature (e.g., Hall et al. 2014). I focus on the benefit of patents in facilitating knowledge sharing among partners and the cost of knowledge disclosure, both of which are relevant for knowledge diffusion. Other aspects of patenting decisions are beyond the scope of this paper.

Second, I make a simplifying assumption that patents provide complete protection from imitation, whereas secrets are not protected once shared with a partner. In general, patent rights do not guarantee this level of protection (Lemley and Shapiro 2005), and firms can benefit from secret sharing even without property rights (e.g., Anton and Yao 1994). However, the main results will hold as long as patent rights allow firms to extract more surplus in knowledge sharing with partners than without these rights.

Third, in the model, the incumbent shares its secrets with a partner only if doing so increases the value of its patented knowledge. Thus, partners are differentiated based on the market conditions they face, which in turn affect the complementarity between secrets and patents. Partners might also be differentiated along other dimensions. For instance, if a firm is concerned with trade secret leakage, it might prefer working only with long-term partners and rely on relational contracting. As long as secrecy leads to more selective knowledge sharing compared to patenting, the main predictions from this model remain qualitatively the same.

Fourth, I assume that patent rights facilitate knowledge sharing because an alternative to patenting is secrecy. If knowledge has already been disclosed, patent rights could impede followon innovations (Murray and Stern 2007; Galasso and Schankerman 2015). For example, consider a technology licensing from firm A to firm B. If B's access to the technology cannibalizes A's profits, A will only license the technology if B's profit increase exceeds A's losses. However, if knowledge is disclosed and A does not have patent rights, B can use the technology even if it reduces the total surplus.

Finally, Ouellette (2012) and Hegde et al. (2023) show that patents disclose some valuable information for others. In the model, I assume that patenting reveals valuable information but not necessarily all information about a technology, consistent with legal studies (e.g., Fromer 2008; Ouellette 2016). The value of disclosed knowledge depends on its complementarity with undisclosed information, being useless only if they are perfect complements.

4.3 Evidence: Secrecy and Knowledge Flows

This section tests Proposition 1, which states that bundling patents with secrets leads to more concentrated knowledge flows. Testing the connection between secrecy and patent citations is challenging for at least for two reasons. First, trade secrets are not observable. I suggest using trade secret litigation to make progress in this measurement problem. Specifically, I use Lex Machina data to find patents involved in federal trade secret litigation.²⁸ These patents are likely to be a part of a broader technology that also involves trade secrets. For instance, Section 4.1 describes the complementarity between patents and trade secrets from the case *Waymo LLC v. Uber Technologies, Inc.* Details on the data are given in Appendix A.1.

The second challenge in testing the connection between secrecy and patent citations is identification. Patents bundled with secrets and involved in trade secret litigation are not random. For instance, citations to these patents might differ from citations to other patents due to a publicity effect of litigation. Furthermore, the intellectual property strategies of firms engaged in litigation could differ from those of other companies. To partially address this concern, I find control patents involved in patent infringement litigation but *without* trade secret claims, and I require both treatment and control patents to be from the same firm. In various specifications, I also require patents to share similar characteristics, such as grant years, technological classes, and a number of citations. Nevertheless, the comparison of citation patterns to these patents should be interpreted with caution.

For each patent involved in trade secret litigation, I find control patents which were involved in patent infringement litigation but without trade secret claims. I compute the difference in the concentration of citations between treatment and control patents using two measures: Herfindahl-Hirschman Index and the share of citations from the most citing firm ("Top Share"). Then, I take the average of this difference across patents. I test whether the average difference in the concentration between patents with and without trade secret claims significantly deviates from the difference we would expect if citations were random, controlling for application years, locations of inventors, and technological classes of patents (see Section 2.3). The details are given in Appendix A.6.

Figure 8 gives a visual test against the null hypothesis. The vertical line shows the actual difference in the concentration, and the histogram shows the distribution of this difference if citations were random within the same application years, locations of inventors, and technological classes. For the HHI measure, the difference in concentration is 0.085 which corresponds to approximately 16% higher concentration of citations for patents bundled with

²⁸The misappropriation of trade secrets can be litigated in both state and federal courts. However, the infringement of patents is litigated in federal courts. Therefore, there is no loss of generality in a focus on federal litigation.



Figure 8: The Difference in the Concentration of Citations Between Patents With and Without Trade Secrets

These figures show the average difference in the concentration of citations between patents involved in litigation with and without trade secret claims. The vertical lines show the actual difference in the concentration. The histograms show the distribution of the difference when citations are random within (time, location, technology) triple (see Section 2.3 for the details). Panel (a) shows the results when the concentration is measured as the Herfindahl-Hirschman Index, and Panel (b) shows the results when the concentration is measured as the share of citations from the most citing firm.

trade secrets. For the "Top Share" measure, the difference in concentration is 0.06 which corresponds to approximately 9.2% higher concentration of citations. Table B2 in Appendix B shows the results based on different criteria for selecting control patents, that is whether treatment and control patents have the same grant year, receive a similar number of citations, belong to the same technological class, or are assigned to plaintiffs rather than defendants.²⁹ Figure 8 shows the results for the most restrictive set of controls (columns 4 and 8 in Table B2). All specifications show a positive and significant difference in concentration between patents involved in litigation with and without trade secret claims.

²⁹Controls for grant years and technological classes ensure that patents represent similar technologies. A control for the number of citations ensures that there are no mechanical differences in the concentration. Requiring patents to be assigned to a plaintiff increases a probability that patents are bundled with trade secrets involved in litigation. For example, if firm A shares secrets with firm B under some contractual arrangement (e.g., an acquisition), and firm B patents these secrets, then firm A might sue firm B for misappropriation of the trade secrets. However, in this situation, patents are not bundled with secrets. Requiring patents to be assigned to a plaintiff eliminates such cases.

4.4 Evidence: Competition and Knowledge Flows

This section tests Proposition 2, which states that higher Schumpeterian competition leads to more concentrated knowledge flows. Autor et al. (2020) show that import competition from China reduced profitability, patenting, and R&D investments of U.S. companies, suggesting that it increased Schumpeterian competition in the U.S. Moreover, the rise in trade with China occurred around 2000, which coincided with the start of the increase in the concentration of patent citations. Using the empirical strategy and data from Autor et al. (2020), I study the relationship between import competition from China and citation patterns between firms.

Following Autor et al. (2020), I define the measure of trade exposure at the four-digit Standard Industry Classification (SIC) over the two subperiods, 1991 to 1999 and 1999 to 2007,

$$\Delta IP_{i1} = \frac{M_{i,1999} - M_{i,1991}}{Y_{i,91} + M_{i,91} - E_{i,91}} \text{ and } \Delta IP_{i2} = \frac{M_{i,2007} - M_{i,1999}}{Y_{i,91} + M_{i,91} - E_{i,91}}$$
(4.1)

where $M_{i,t}$ is the U.S. imports from China for industry *i* and year $t \in \{1991, 1999, 2007\}$, and $Y_{i,91} + M_{i,91} - E_{i,91}$ is the absorption at the start of the period (industry shipments plus imports minus exports). For each patent, I calculate the import penetration for its technological class using the mapping of four-digit SIC industries to technological classes implied from patents owned by publicly traded firms as in Autor et al. (2020).

Panel (a) of Figure 9 shows the concentration measure from Figure 1 for different technological classes based on their exposure to import competition from China, ΔIP_{i2} . All classes experienced an increase in the concentration of citations after 2000. However, the dynamics are different based on the exposure to import competition. For classes with the least exposure (below first quartile Q_{25}), the concentration of citations was initially high (around 55%) and stable prior to 2000, and then it increased up to 75%. For classes with the most exposure (above the third quartile Q_{75}), the concentration was initially lower (around 50%), decreased to around 40% near 2000, and then it also increased up to 75%. Thus, technological classes with the most exposure to China shock experienced faster growth in the concentration.

I study the change in the concentration of citations around 2000 in the regression specifications. For each technological class, I define the average concentration measure among patents granted in the following seven-year periods: 1977–1983, 1984–1990, 1991–1997, 1998–2004, 2005–2011.³⁰ Appendix A.7 provides more details. The concentration measure for technological class j and for the period starting from t is denoted by $C_{j,t}$. I define the following growth measures

 $\Delta y_{j1} = 100 \cdot \ln \left(\mathcal{C}_{j,1998} / \mathcal{C}_{j,1991} \right)$ and $\Delta y_{j2} = 100 \cdot \ln \left(\mathcal{C}_{j,2005} / \mathcal{C}_{j,1998} \right)$

 $^{^{30}}$ The periods are chosen to map the data from Autor et al. (2020).



(a) The concentration for technological classes with different exposure to import competition.

(b) Changes in the concentration and in the import penetration from China.

Figure 9: Trade with China and Concentration of Citations

These figures show the relationship between the concentration of citations and the exposure to import competition from China. Panel (a) shows the concentration measure from Section 2 (Figure 1) for different technological classes divided into quartiles based on their exposure to import competition from China, ΔIP_{i2} in (4.1). The concentration measure for year t shows the average concentration in years [t - 4, t]. Panel (b) shows the binned scatter plot of the change in the concentration of citations and the change in the import exposure from China. The specification is weighted by the number of Compustat-matched U.S.-inventor patents in a technology class.

Panel (b) of Figure 9 shows that technological classes more exposed to trade with China $(\Delta I P_{j\tau})$ experienced a higher growth in the concentration of citations $(y_{j\tau})$.

I estimate the following specification

$$\Delta y_{j\tau} = \beta \Delta I P_{j\tau} + \gamma X_{j0} + \varepsilon_{j\tau} \tag{4.2}$$

where $\tau \in \{1,2\}$ and X_{j0} is the set of controls. To control for the aggregate trend in the concentration of citations, I include time fixed effects. Since the concentration measure depends on the total number of citations, I also include the change in the average number of citations for each technological class. Moreover, I include two lags of the outcome variable to control for technology-specific trends prior to China shock. I also include fixed effects for 11 manufacturing sectors and for 6 main NBER technological categories. Finally, I control for the rising importance of software inventions (Chattergoon and Kerr 2021). Specifically, for each technological class I include a dummy variable indicating whether it has more than 50% of software subclasses (Graham and Vishnubhakat 2013).

Panel A in Table 2 shows the results of simple OLS regressions, and Panel B shows the results for the specification in which changes in US import exposure $(\Delta IP_{j\tau})$ are instrumented

	(1)	(2)	(3)	(4)	(5)	(6)			
Panel A: OLS									
Δ Tech Class Exposure	2.06	1.52	1.51	1.77	1.41	1.41			
to Chinese Imports	(0.44)	(0.42)	(0.42)	(0.37)	(0.42)	(0.42)			
Panel B: 2SLS									
Δ Tech Class Exposure	2.31	1.57	1.55	1.92	1.71	1.70			
to Chinese Imports	(0.49)	(0.50)	(0.50)	(0.44)	(0.61)	(0.62)			
Time FE Δ Citations 2 Lags of outcomes 11 sectors, 6 Tech Software Patents		Yes	Yes Yes	Yes Yes Yes	Yes Yes Yes Yes	Yes Yes Yes Yes Yes			

Table 2: Trade with China and Increase in the Concentration of Citations

Notes: This table shows the estimated coefficient β for the specification in (4.2). Panel A shows the results for simple OLS regressions. Panel B shows the results for the specification in which the import penetration from China is instrumented with Chinese exports to non-U.S. high-income markets (Autor et al. 2020). Regressions consider the effect of higher growth in import penetration from China on the increase in the concentration of citations at the technology class level. Industry exposure to Chinese competition is mapped to technology class exposure using the mapping implied by the U.S. publicly traded firms in Compustat as in Autor et al. (2020). Controls include time fixed effects, a change in the average number of citations for a technology class, 2 lags of the outcome variable, fixed effects for 11 manufacturing sectors and for 6 main NBER technology categories, and a dummy for software technological classes. I define the software classes as classes where more than 50% of software subclaclasses according to Graham and Vishnubhakat (2013). All specifications are weighted by the number of Compustat-matched U.S.-inventor patents in a technology class. Standard errors are clustered at the technology class level.

by changes in Chinese exports to non-U.S. high-income countries (Autor et al. 2013; Autor et al. 2020). All specifications show a positive and significant relationship between the changes in the import competition from China and the growth in the concentration of citations.

Discussion and Robustness Checks

Trade with China has had a profound effect on various aspects of the U.S. economy, from innovation (Autor et al. 2020; Hoberg et al. 2021) and labor markets (Autor et al. 2013) to the structure of supply chains (Antràs et al. 2017). The evidence in this section suggests that it might have also affected incentives for knowledge sharing among business partners. I provide additional discussion and robustness checks for this result below.

The instrument from Autor et al. (2013) is designed to isolate changes in the trade with
China unrelated to U.S. demand and technological shocks. Table B3 in Appendix B shows the results from two additional placebo exercises. First, I show that the relationship between China shock and the rise in the concentration is insignificant for patents assigned to non-corporate entities (e.g., universities and government agencies). Therefore, the effect of trade competition with China is specific to the corporate sector, and the results are unlikely to be driven by the correlation between general technological changes and globalization. Second, I regress lagged outcome variables (pre 1991) on future changes in imports from China. The coefficients are insignificant, indicating that the main results are unlikely to be driven by contemporaneous changes in the technological opportunities and trade.

Using the methodology outlined in Section 2.3, I also study whether China shock influenced the rise in the concentration of citations through changes in firm patenting behavior (N) or in firm citation rates (p).³¹ Figure B15 in Appendix B shows that 84% of the rise in the concentration of citations in the technologies most exposed to trade with China is attributed to changes in the citation rates. In contrast, for the least exposed technologies this share is 45%. Overall, China shock changed the way firms cite each other. This effect is distinct from the decline in patenting documented in Autor et al. (2020). For instance, firms' exit from patenting cannot fully explain the rise in the concentration of citations.

4.5 Implications for the Innovation Literature

In Section 4.2, I develop a model of endogenous knowledge flows among firms. Knowledge flows to competitors are controlled through incomplete information disclosure in patent applications; knowledge sharing among partners is managed through licensing and trade secret sharing. In this section, I discuss the model's implications for the innovation literature.

The assumption on the accumulation of knowledge lies at the heart of most growth models. The endogenous growth literature (e.g., Romer 1990; Grossman and Helpman 1991; Aghion and Howitt 1992) assumes that the productivity growth is independent of the current productivity. Semi-endogenous growth models (e.g., Jones 1995; Kortum 1997; Segerstrom 1998) assume that the higher the productivity, the harder it is to improve upon it. Bloom et al. (2020) provide empirical evidence supporting the latter assumption.

One way to model productivity dynamics is to assume that $z^e = \psi \cdot (\overline{z})^{\beta}$, where \overline{z} is the current productivity level, z^e is the productivity of an entrant who can replace the incumbent firm, and ψ is an independent draw from a continuous distribution on \mathbb{R}_+ (Buera and Oberfield 2020). In this specification, β measures the degree of knowledge diffusion and is usually exogenously assumed. If $\beta = 1$, then the probability of positive productivity growth $\mathbb{P}[z^e > z]$ is independent

 $^{^{31}}$ In addition to application years and technological classes, I also control for the location of inventors to take into account the geographical specialization of firms (see Appendix A.4).

of the current productivity level; if $\beta < 1$, then this probability negatively depends on the current productivity level (ideas are harder to find).

The choice of information disclosure can endogenize β . Consider the model from Section 4.2 with Cobb–Douglas function, $y = (z_1 x_1)^{\frac{1}{2}} (z_2 x_2)^{\frac{1}{2}}$, and input prices equal to 1, $w_1 = w_2 = 1$. Suppose that $\overline{z} > \underline{z} = 1$ and there is one entrant $(n_e = 1)$ with total productivity $z^e = (z_1^e)^{\frac{1}{2}} (z_2^e)^{\frac{1}{2}}$. Under full patenting, the productivity dynamic is $z^e = \psi \cdot \overline{z}$, and the probability of positive productivity growth $\mathbb{P}(z^e/\overline{z} = \psi > 1)$ is independent of the current productivity. Under partial disclosure, the productivity dynamic is $z^e = \psi \cdot \overline{z}^{1/2}$, and the probability of positive productivity growth $\mathbb{P}(z^e/\overline{z} = \psi/\overline{z}^{1/2} > 1)$ negatively depends on the current productivity level. Under complete secrecy, this probability is even lower, $\mathbb{P}(z^e/\overline{z} = \psi/\overline{z} > 1)$. In other words, ideas are harder to find if less information is disclosed to the public.

In the previous example, β can take only values in $\{0, \frac{1}{2}, 1\}$ depending on the information disclosure. The model can be extended to make $\beta \in [0, 1]$. Suppose the production function is $y = \exp\left(\int_0^1 \ln(z_i x_i) di\right)$ with a continuum of inputs in [0, 1]. The technology corresponds to a continuum of productivity levels $\{z_i\}_{i \in [0,1]}$. As in Section 4.2, each component can be patented or kept secret. Suppose share β of components is patented and revealed to the public, while share $1 - \beta$ is kept secret. Then, the productivity dynamic is $z^e = \psi \cdot (\overline{z})^{\beta}$. The parameter β can be determined by exogenous factors: for instance, some knowledge components can be reverse engineered, while others are tacit and unpatentable. However, as long as there are knowledge components for which a firm can choose between patenting and secrecy, β is endogenously determined. Appendix C.2 provides more details.

One implication of a firm's ability to control knowledge flows is that the degree of ideas getting harder to find might respond to changes in market conditions. For example, Proposition 2 shows that higher Schumpeterian competition can lead to greater reliance on secrecy. Despite a higher number of competing firms, business dynamism can decrease once firms rely more on secrecy. Formally, suppose if there are n_e entrants, the firm fully discloses information in patents. The probability of the incumbent's replacement is $1 - [\mathbb{P}(\psi \leq 1)]^{n_e}$. Suppose the firm shifts to partial disclosure if there are $n'_e > n_e$ entrants. The probability of the incumbent's replacement is $1 - [\mathbb{P}(\psi \leq 1)]^{n_e}$. Suppose the firm shifts to partial disclosure if there are $n'_e > n_e$ entrants. The probability of the incumbent's replacement is $1 - [\mathbb{P}(\psi \leq 1)]^{n_e}$. Although there are more entrants, $n'_e > n_e$, partial disclosure makes it harder for each to improve upon the incumbent's technology, $\mathbb{P}(\psi > \overline{z}^{1/2}) < \mathbb{P}(\psi > 1)$.

Akcigit and Ates (2022) argue that a decline in knowledge flows among competitors can explain several trends in the U.S. economy, including the decline in business dynamism. They model this decline through a decrease in the exogenous knowledge diffusion parameter responsible for the probability of a lagging competitor catching up with a frontier firm. Their quantitative results "stress the importance of future research to understand the underlying reasons for slower knowledge diffusion" (Akcigit and Ates 2022, p. 2064). I show how a competitive threat from entrants can lead to a decrease in knowledge flows to competitors if the frontier firm can choose the degree of information disclosure to the public.

The citation patterns documented in Sections 2 and 3 show that business partners, especially vertical ones, account for a significant share of knowledge flows. These partners are likely to develop complementary technologies. The importance of complementary innovations is highlighted in the literature on general purpose technologies (e.g., Helpman 1998; Bresnahan 2010) and economic history (e.g., Rosenberg 1963; Rosenberg 1979): "The growing productivity of industrial economies is the complex outcome of large numbers of interlocking, mutually reinforcing technologies, the individual components of which are of very limited economic consequence by themselves. The smallest relevant unit of observation is seldom a single innovation but, more typically, an interrelated clustering of innovations" (Rosenberg 1979, p. 28–29). In this paper, I show how competition can affect knowledge sharing among partners, linking creative destruction and construction forces.

5 Conclusion

This paper provides evidence that firms have significant control over the diffusion of the knowledge they generate. They combine patents and trade secrets, sharing secrets with selected partners. Knowledge flows are shaped by the firm's incentives for cooperation with other companies and the risks associated with disclosing knowledge to rivals. As a result, knowledge flows depend on the details of the economic environment, such as the degree of competition. Further research on how firms manage their intellectual property and share knowledge with partners might provide deeper insights into the process of economic growth.

References

- Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2020). Sampling-Based Versus Design-Based Uncertainty in Regression Analysis. *Econometrica*, 88(1):265–296.
- Aghion, P., Akcigit, U., and Howitt, P. (2014). What Do We Learn from Schumpeterian Growth Theory, *Handbook of Economic Growth*, 2B:515–563.
- Aghion, P., Bergeaud, A., and Van Reenen, J. (2023). The Impact of Regulation on Innovation. American Economic Review, Fortcoming.
- Aghion, P., Bloom, N., Blundell, R., Griffith, R., and Howitt, P. (2005). Competition and Innovation: An Inverted-U Relationship. *Quarterly Journal of Economics*, 120(2):701–728.
- Aghion, P. and Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–51.
- Akcigit, U. and Ates, S. T. (2022). What Happened to U.S. Business Dynamism? Journal of Political Economy, 131(8):2059–2124.
- Akcigit, U. and Kerr, W. (2018). Growth Through Heterogenous Innovations. Journal of Political Economy, 126(4):1374–1443.
- Alcácer, J., Gittelman, M., and Sampat, B. (2009). Applicant and Examier Citations in U.S. Patents: An Overview and Analysis. *Research Policy*, 38:415–427.
- Alfaro-Ureña, A., Manelici, I., and Vasquez, J. P. (2022). The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages. *Quarterly Journal of Economics*, 137(3):1495–1552.
- Amara, N., Landry, R., and Traoré, N. (2008). Managing the Protection of Innovations in Knowledge-Intensive Business Services. *Research Policy*, 37:1530–1547.
- Anton, J. J., Greene, H., and Yao, D. A. (2006). Policy Implications of Weak Patent Rights. Innovation Policy and the Economy, Volume 6, edited by Adam B. Jaffe, Josh Lerner, and Scott Stern, pages 1–26.
- Anton, J. J. and Yao, D. A. (1994). Expropriation and Inventions: Appropriable Rents in the Absence of Property Rights. *American Economic Review*, 84(1):190–209.
- Antràs, P., Fort, T. C., and Tintelnot, F. (2017). The Margins of Global Sourcing: Theory and Evidence from US Firms. American Economic Review, 107(9):2514–2564.

- Arora, A. (1997). Patents, Licensing and Market Structure in the Chemical Industry. Research Policy, 26(4–5):391–403.
- Arora, A., Fosfuri, A., and Gambardella, A. (2001). Markets for Technology. The Economics of Innovation and Corporate Strategy. The MIT Press.
- Arora, A. and Gambardella, A. (2010). Ideas for Rent: An Overview of Markets for Technology. Industrial and Corporate Change, 19(3):775–803.
- Arqué-Castells, P. and Spulber, D. F. (2022). Measuring the Private and Social Returns to R&D: Unintended Spillovers versus Technology Markets. *Journal of Political Economy*, Forthcoming.
- Arrow, K. (1962). Economic Welfare and the Allocation of Resources for Invention. in R. Nelson, eds., The Rate and Direction of Inventive Activity. Princeton University Press, Princeton, NJ, pages 609–625.
- Atkin, D., Chen, K., and Popov, A. (2022). The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley. *Working Paper*.
- Autor, D., Dorn, D., Hanson, G. H., Pisano, G., and Shu, P. (2020). Foreign Competition and Domestic Innovation: Evidence from US Patents. *American Economic Review: Insights*, 2(3):357–374.
- Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6):2121– 2168.
- Bai, J., Barwick, P., Cao, S., and Li, S. (2022). Quid Pro Quo, Knowledge Spillover, and Industrial Quality Upgrades: Evidence from the Chinese Auto Industry. *Working Paper*.
- Barrot, J.-N. and Sauvagnat, J. (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *Quarterly Journal of Economics*, 131(3):1543–1592.
- Bell, A., Chetty, R., Jaravel, X., Petkova, N., and van Reenen, J. (2019). Who Becomes an Inventor in America? The Importance of Exposure to Innovation. *Quarterly Journal of Economics*, 134(2):647–713.
- Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4):1104–1144.

- Bloom, N., Van Reenen, J., and Williams, H. (2019). A Toolkit of Policies to Promote Innovation. *Journal of Economic Perspectives*, 33(3):163–184.
- Bresnahan, T. (2010). General Purpose Technologies. *Handbook of the Economics of Innovation*, 2:761–791.
- Buera, F. and Lucas, R. (2018). Idea Flows and Economic Growth. Annual Review of Economics, 10:315–45.
- Buera, F. J. and Oberfield, E. (2020). The Global Diffusion of Ideas. *Econometrica*, 88(1):83–114.
- Caballero, R. and Jaffe, A. B. (1993). How High are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth. NBER Working Paper No. 4370.
- Carlino, G. and Kerr, W. (2015). Aglomeration and Innovation. *Handbook of Regional and Urban Economics*, 5:349–404.
- Cassiman, B. and Veughelers, R. (2002). R&D Cooperation and Spillovers: Some Empirical Evidence from Belgium. *American Economic Reivew*, 92(4):1169–1184.
- Chattergoon, B. and Kerr, W. R. (2021). Winner Takes All? Tech Clusters, Population Centers, and the Spatial Transformation of U.S. Invention. *Research Policy*, Forthcoming.
- Chodorow-Reich, G. (2014). The Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008-9 Financial Crisis. *Quarterly Journal of Economics*, 129(1):1–59.
- Cohen, W. M., Nelson, R. R., and Walsh, J. P. (2000). Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). National Bureau of Economic Research Working Paper 7552.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv, Available at https://arxiv.org/abs/1810.04805.
- Duguet, E. and MacGarvie, M. (2005). How Well Do Patent Citations Measure Flows of Technology? Evidence from French Innovation Surveys. *Economics of innovation and New Technology*, 14(5):375–393.
- Eeckhout, J. and Jovanovic, B. (2002). Knowledge Spillovers and Inequality. American Economic Review, 92(5):1290–1307.

- Ellison, G., Glaeser, E. L., and Kerr, W. (2010). What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns. *American Economic Review*, 100(3):1195–1213.
- Fromer, J. C. (2008). Patent Disclosure. *Iowa Law Review*, 94(2):539–606.
- Galasso, A. and Schankerman, M. (2015). Patents and Cumulative Innovation: Causal Evidence From the Courts. *Quaterly Journal of Economics*, 130(1):317–369.
- Gans, J. S. and Stern, S. (2000). Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction. *Journal of Economics and Management Strategy*, 9(4):485–511.
- Graham, S. and Vishnubhakat, S. (2013). Of Smart Phone Wars and Software Patents. *Journal* of *Economic Perspectives*, 27(1):67–86.
- Graham, S. J. H., Marco, A. C., and Myers, A. F. (2018). Patent Transactions in the Marketplace: Lessons from the USPTO Patent Assignment Dataset. *Journal of Economics* and Management Strategy, 27(3):343–371.
- Grossman, G. M. and Helpman, E. (1991). Quality Ladders in the Theory of Growth. *Review* of *Economic Studies*, 58(1):43–61.
- Hall, B. H., Helmers, C., Rogers, M., and Sena, V. (2014). The Choice between Formal and Informal Intellectual Property: A Review. *Journal of Economic Literature*, 52(2):375–423.
- Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2005). Market Value and Patent Citations. *RAND Journal of Economics*, 36(1):16–38.
- Hegde, D., Herkenhoff, K., and Zhu, C. (2023). Patent Publication and Innovation. Journal of Political Economy, 131(7):1845–1903.
- Hegde, D. and Luo, H. (2017). Patent Publication and the Market for Ideas. *Management Science*, 64(2):652–672.
- Helpman, E. (1998). General Purpose Technologies and Economic Growth. The MIT Press.
- Hoberg, G., Li, Y., and Phillips, G. M. (2021). U.S. China Innovation Competition. Working Paper.
- Jaffe, A. B., Trajtenberg, M., and Fogarty, M. S. (2000). Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors. *American Economic Review: Papers and Proceedings*, 90(2):215–218.

- Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *Quarterly Journal of Economics*, 108(3):577–598.
- Jones, C. I. (1995). R&D Based Models of Economic Growth. *Journal of Political Economy*, 103(4):759–784.
- Jones, C. I. (2005). Growth and Ideas. Handbook of Economic Growth, Volume 1B, edited by Philippe Aghion and Steven N. Durlauf.
- Jorda, K. F. (2008). Patent and Trade Secret Complementariness: An Unsuspected Synergy. Washburn Law Journal, 48(1):1–32.
- Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological Innovation, Resource Allocation, and Growth. *Quarterly Journal of Economics*, 132(2):665–712.
- Kortum, S. (1997). Research, Patenting, and Technological Change. *Econometrica*, 65(6):1389–1419.
- Kuhn, J., Younge, K., and Marco, A. (2020). Patent Citations Reexamined. RAND Journal of Economics, 51(1):109–132.
- Kuhn, J., Younge, K., and Marco, A. (2023). Strategic Citation: A Reassessment. Review of Economics and Statistics, 105(2):458–466.
- Lampe, R. (2012). Strategic Citation. Review of Economics and Statistics, 94(1):320–333.
- Landes, W. M. and Posner, R. A. (2003). *The Economic Structure of Intellectual Property Law*. Harvard University Press.
- Lemley, M. A. and Shapiro, C. (2005). Probabilistic Patents. *Journal of Economic Perspectives*, 19(2):75–98.
- Lerner, J. and Seru, A. (2022). The Use and Misuse of Patent Data: Issues for Corporate Finance and Beyond. *Review of Financial Studies*, 35(6):2667–2704.
- Lim, K. (2009). The Many Faces of Absorptive Capacity: Spillovers of Copper Interconnect Technology for Semiconductor Chips. Industrial and Corporate Change, 18(6):1249–1284.
- Liu, E. and Ma, S. (2023). Innovation Networks and R&D Allocation. Working Paper.
- Marshall, A. (1920). Principles of Economics. London: MacMillan.
- Mazzoleni, R. and Nelson, R. R. (1998). The Benefits and Costs of Strong Patent Protection: a Contribution to the Current Debate. *Research Policy*, 27(3):273–284.

- Melitz, M. and Polanec, S. (2015). Dynamic Olley-Pakes productivity decomposition with entry and exit. *The RAND Journal of Economics*, 46(2):362–375.
- Mezzanotti, F. and Simcoe, T. (2023). Innovation and Appropriability: Revisiting the Role of Intellectual Property. *Working Paper*.
- Murray, F. and Stern, S. (2007). Do Formal Intellectual Property Rights Hinder the Free Flow of Scientific Knowledge? An Empirical Test of the Anti-Commons Hypothesis. *Journal of Economic Behavior & Organization*, 63:648–687.
- Ouellette, L. L. (2012). Do Patents Disclose Useful Information. Harvard Journal of Law and Technology, 25(2):545–607.
- Ouellette, L. L. (2016). Pierson, Peer Review, and Patent Law. Vanderlbilt Law Review, 69(6):1825–1848.
- Ouellette, L. L. and Masur, J. S. (2023). *Patent Law: Cases, Problems, and Materials.* Independently published.
- Price, N. and Rai, A. K. (2016). Manufacturing Barriers to Biologics Competition and Innovation. *Iowa Law Review*, 101(3):1023–1063.
- Price, N., Rai, A. K., and Minssen, T. (2020). Knowledge transfer for large-scale vaccine manufacturing. *Science*, 369(6506):912–914.
- Roin, B. N. (2005). The Disclosure Function of the Patent System (Or Lack Thereof). Harvard Law Review, 118(6):2007–2028.
- Romer, P. (1990). Endogenous Technical Change. *Journal of Political Economy*, 98(5):S71–S102.
- Rosenberg, N. (1963). Technological Change in the Machine Tool Industry, 1840–1940. The Journal of Economic History, 23(4):414–443.
- Rosenberg, N. (1979). Technological Interdependence in the American Economy. *Technology* and *Culture*, 20(1):25–50.
- Saxenian, A. (1996). Regional Advantage: Culture and Competition in Silicon Valley and Route 128. Harvard University Press.
- Schumpeter, J. (1942). Capitalism, Socialism and Democracy. New York: Harpaer & Brothers.
- Scotchmer, S. (2004). Innovation and Incentives. The MIT Press.

- Segerstrom, P. S. (1998). Endogenous Growth without Scale Effects. *American Economic Review*, 88(5):1290–1310.
- Shapiro, C. (2012). Competition and Innovation: Did Arrow Hit the Bull's Eye? The Rate and Direction of Inventive Activity Revisited, edited by Josh Lerner and Scott Stern.
- Singh, J. and Marx, M. (2013). Geographical Constraints on Knowledge Spillovers: Political Borders vs. Spatial Proximity. *Management Science*, 59(9):2056–2078.
- Thompson, P. and Fox-Kean, M. (2005). Patent Citations and the Geography of Knowledge Spillovers: A Reassessment. *American Economic Review*, 95(1):450–460.
- Whinston, M. D. (2012). Comment on "Competition and Innovation: Did Arrow Hit the Bull's Eye?". The Rate and Direction of Inventive Activity Revisited, edited by Josh Lerner and Scott Stern.

Appendix. For Online Publication Table of Contents

A	Data Appendix								
	A.1 Data Details	47							
	A.2 Details on the Concentration of Patent Citations	48							
	A.3 Robustness to Section 2: Concentration of Citations	50							
	A.4 Details on Monte-Carlo Simulations	51							
	A.5 Movement of Inventors and Citation Patterns	53							
	A.6 Details on Trade Secret Litigation	54							
	A.7 Details on Import Competition from China	55							
в	Figures and Tables	56							
С	Theory Appendix	72							
	C.1 Investments in R&D	72							
	C.2 A Continuum of Knowledge Components	72							
	C.3 Production by the Incumbent	73							
	C.4 A Generalization of the Types of Relationships	74							

A Data Appendix

A.1 Data Details

Patents. The main source of patent data is PatentsView. Autor et al. (2020) provide a matching of patent assignees to Computat firm names for publicly traded firms. I use their existing matching of assignee names to Computat firms for the period 1976–2014 to extend it for years up to 2019. For the rest of the patents, I follow the procedure outline in Autor et al. (2020) for cleaning and standardizing firm names (e.g., replace "Incorporated" with "INC"). Finally, I matched around 100 thousand patents manually for the largest assignees.

FactSet Revere. I use two data sets from FactSet. FactSet Revere Company provides basic information on companies, including their names. FactSet Revere Supply Chain Relationships provides information on business relationships between firms. I discuss the types of relationships below. The following quote from FactSet's manual describes how they collect data on business relationships:

"FactSet analysts systematically collect companies' relationship information exclusively from primary public sources such as SEC 10-K annual filings, investor presentations, and press releases, and classify them through normalized relationship types. Company information is fully reviewed annually, and changes based on corporate actions are monitored daily. The result is a comprehensive, detailed and up-to-date dataset of material intercompany relationships."

FactSet also provides information on the names of subsidiaries in some business relationships. I match all these names, including subsidiaries, with the names of assignees listed in patents, using the same matching procedure as I did for patents.

FactSet records licensing and supplier-customer relationships in a duplicate manner. Specifically, a firm receiving a licensing is also recorder as a customer, and a firm providing the intellectual property is also listed as a supplier. For all firm pairs involved in licensing agreements, I exclude the recordings on supplier-customer connections.

Compustat Segment. I take the supplier-customer data from Compustat Segments data set. For publicly traded firms, the data list names of the main customers, which are mostly other firms but can also be government agencies. Regulation SFAS No. 131 requires publicly traded firms to report the identity of any customer representing more than 10% of their total sales. Using Compustat Segments, Barrot and Sauvagnat (2016) constructed a data set of suppliers and customers for the U.S. publicly traded firms for the period 1976–2013. I extend their data up to 2022 using name matching and manual inspection.

USPTO Patent Reassignment. The data on patent reassignment is described in Graham et al. (2018). I leave only reassignment of patents between companies. Formally, I leave only transactions with 'convey_ty' equal to 'assignment'. I clean firm names in the same way as for patents. Then I match these data to patent assignees.

All three data sets—FactSet, Segments, and Reassignment—provide information on the date of a transaction or a relationship. These dates are self-reported, so they might not correspond to the dates of actual relationships. For each firm-pair, I find the minimum and the maximum of years in which the relationship between firms was active. I leave only relationships active between 2003 (the first year in FactSet) and 2022.

I group certain relationships into more aggregated groups. Table B1 in Appendix B provides a summary of all relationships.

Lex Machina. Lex Machina offers comprehensive data on federal litigation involving patents and trade secrets. Each case entry in the database includes details such as the names of plaintiffs, defendants, and any third parties involved. In cases related to patent litigation, the database also lists the patents at issue. Additionally, the data indicate whether a given case has overlapping claims with trade secret litigation.

In my analysis, I identified 1,092 cases that featured a total of 2,541 patents and were involved in both patent and trade secret litigation. These cases were filed from 2001 to 2021.

China Shock. Autor et al. (2020) provides a measure of the exposure to import competition from China at the main group level in the US Patent Classification (USPC). I use these data in Section 4.4.

A.2 Details on the Concentration of Patent Citations

The concentration measure from Figure 1 is constructed in the following way. First, I identify the top 1% of the most cited patents within each grant year and technological class. Second, for these patents I compute the share of citations coming from the most citing firm. Finally, I aggregate these measures within and across technological classes.

The first step is to identify the top 1% of the most cited patents. Denote $y_{km} = 1$ if patent $m \in \mathcal{P}$ makes a citation to patent $k \in \mathcal{P}$ and $y_{km} = 0$ otherwise, where \mathcal{P} is the set of all granted patents. Each patent has an assignee (owner) or, in rare cases (around 3%), multiple assignees. For the majority of patents, the assignee is a corporate firm but it can also include universities, government agencies, and individual inventors. In the second step, I compute the distribution of citations across different organizations, so I exclude citations from individual inventors and patents with missing assignee information.³² Each patent $k \in \mathcal{P}$ has a grant year

³²Formally, I set $y_{km} = 0$ where patent *m* belongs to an individual inventor and does not have have an assignee information.

 t_k^g and a primary technological class c_k . I define the technological classes at the group level in the Cooperative Patent Classification (primary class is a class listed first in the patent file). Denote the set of all groups in Cooperative Patent Classification (CPC) system by CPC. For each patent $k \in \mathcal{P}$, I compute the number of citations within a 5-year window from a grant day

$$n_k = \sum_{m \in \mathcal{T}_k} y_{km} \text{ where } \mathcal{T}_k = \{ m \in \mathcal{P} : \text{Grant Date}_m - \text{Grant Date}_k \leq 5 \cdot 365 \text{ Days} \}$$
(A.1)

For each grant year t and technological class c, I define the set of all granted patents receiving at least one citation

$$\Omega_{t,c} = \{k \in \mathcal{P} : t_k^g = t \text{ and } c_k = c \text{ and } n_k > 0\}$$

and take the top 1% of patents in terms of the number of citations within this set. Denote it by $\Omega_{t,c}^{top}$, and define the "Main" sample as top patents for all years and technology classes:

$$Main = \left\{\Omega_{t,c}^{top}\right\}_{c \in \mathcal{CPC} \text{ and } t=1976\dots 2014}$$
(A.2)

In the second step, for each patent in the *Main* sample I compute the share of citations coming from the most citing organization. To account for patents with multiple assignees, I define a weighted citation as $y_{km}^w = y_{km}/F_m$ where F_m is the number of assignees for patent m. Define the number of citations to patent $k \in \mathcal{P}$ from organization i as

$$n_{k,i} = \sum_{m \in i, m \in \mathcal{T}_k} y_{km}^w$$

where $m \in i$ means that organization *i* is an assignee for patent $m \in \mathcal{P}$. Then, the concentration measure for patent $k \in \mathcal{P}$ is

$$\mathcal{C}_k = \max_i \left\{ \frac{n_{k,i}}{n_k} \right\}$$

In the third step, I aggregate these measures. Specifically, within each grant year (t) and technological class (c) I compute a simple average across patents³³

$$\mathcal{C}(t,c) = \frac{1}{|\Omega_{t,c}^{top}|} \sum_{k \in \Omega_{t,c}^{top}} \mathcal{C}_k$$
(A.3)

where $|\Omega_{t,c}^{top}|$ is the number of patents in $\Omega_{t,c}^{top}$. Then I aggregate across technological classes using the weighted average of $\mathcal{C}(t,c)$ where weights are defined by the number of patents in each

 $^{^{33}{\}rm The}$ results are robust if instead of a simple average I use a citation-weighted average, or median instead of an average.

 $\Omega_{t,c}^{top} \neq \emptyset$

$$C(t) = \sum_{c \in CPC} \frac{|\Omega_{t,c}^{top}|}{\sum_{c \in CPC} |\Omega_{t,c}^{top}|} C(t,c)$$
(A.4)

The variable C(t) for $t = 1976 \dots 2014$ is shown in Figure 1.

A.3 Robustness to Section 2: Concentration of Citations

Panel (a) in Figure B4 shows that the concentration of citations is similar if we restrict the sample to corporate patents only. I also consider different thresholds for the most cited patents: top 5% and 10%. Finally, I exclude the sample of citing patents that are assigned to superstar firms. Specifically, in each year and group level in Cooperative Patent Classification I find top 1% of firms in terms of the number of patents, and exclude their patents from the sample of citing patents. Panel (b) shows that the results are robust if one uses the citation-weighted average or the median instead of the average to aggregate concentration measures within technological classes. Panel (c) shows that the results are the same when I exclude self-citations of firms to itself, so the concentration is driven by citations between firms rather than self-citations. Kuhn et al. (2020) argue that the quality of citations as a measure of knowledge flows has declined over time due to a small number of patents responsible for a large share of backward citations. Panel (c) shows that the results on the concentration are robust when I exclude top 1% of patents in terms of the number of backward citations. Finally, I exclude citations between patents sharing a common law firm to ensure that the concentration is not driven by lawyers citing themselves. I also group citations from patents from the same within-country family (continuations, continuations-in-part, divisionals) as a single citation. This ensures that the rise in concentration is not driven by increasing patent families. Finally, for the period after 2001 I separate citations made by patent examiners and non-examiners. Figure B5 shows the concentration based on citations from patent examiners is around two times lower than the concentration based on citations from non-examiners.³⁴

I also check whether citations are not driven by lawyers. Specifically, I track citations of lawyers who worked in multiple firms similar to the movement of inventors in Section A.5. The only difference is that there is no data for the location of lawyers, so I consider patents which are filed by the same lawyer in at least two companies, have a similar application period, and are classified to the same main subgroup in Cooperative Patent Classification system. Figure B6 shows that the actual concentration of citations across firms is around 95% within lawyers who represented similar companies. It is significantly higher relative to the 95th quantile of the concentration measure where citation rates are equated across companies within a lawyer.

³⁴The USPTO started to separate examiner and non-examiner citations only around 2001.

A.4 Details on Monte-Carlo Simulations

The details on the Monte-Carlo simulations are the following. First, I divide all granted patents into disjoint groups based on common observational characteristics. Then, for each cited patent I find all patents sharing the same observational characteristics as the citing patents. Second, for each patent I randomize citations across similar patents to equalize citation rates across firms. Third, I randomize citations to equalize both citation rates and patenting across firms. For each patent, I compute the concentration measure on the simulated sample. I repeat this procedure 300 times to construct the distribution of concentration measures. Finally, I aggregate different moments of this distribution in a way similar to the actual concentration measure.

A.4.1 Step 1: Find Patents with Similar Characteristics

In Section 2.3, I use application years and technological classes for patent characteristics. For technological classes, I use the main subgroup level in CPC. Denote by t_k^a and \tilde{c}_k an application year and a technological class of patent k. For each patent, I divide its citing patents based on their characteristics $g = (t^a, \tilde{c})$. Then, I find all patents, citing and non-citing, with the same characteristics.

As a robustness, in Section 2.3 I also group patents based on their textual similarity of abstracts. Specifically, as in the main analysis, for each cited patent I find all patents with the same characteristics $g = (t^a, \tilde{c})$ as the citing ones. Then, for all these patents, citing and non-citing, I compute the vector embedding of their abstracts using BERT model developed by Google. I use the version of BERT model called "all-MiniLM-L6-v2". The embedding is a vector that provides a mapping from text to a numerical representation. Then, I compute pairwise cosine similarities between patents that actually make citations. I find the minimum of this similarity. Next, for each non-citing patent I compute all pairwise cosine similarities with all citing patents. I leave a non-citing patent in the sample only if its similarity with at least one citing patent is greater or equal to the minimum similarity among citing patents. In this exercise, patent characteristics are specific to a cited patent.

In Section 2.4, I use a geographical location of the majority of inventors for patent characteristics. Denote the location for patent k by ℓ_k . I define the geographical location at the state level if an inventor is located in the U.S., and at the country level if an inventor is located outside the U.S. For example, the location for an inventor living in Cambridge, MA, USA is (USA, MA), and for an inventor living in Berlin, Germany is Germany. If a patent has several inventors in different locations, I define the location for a patent based on the location of the majority of inventors. In the case of a tie, I take the location based on the alphabetical order. In this case, the set of patent characteristics is $g = (t^a, \tilde{c}, \ell)$.

A.4.2 Step 2: Equalize Citation Rates

Denote the number of citations to patent k from patents with characteristics g by $n_k(g)$. For each g, I equate citation rates across firms. Formally, for each patent characteristic g I randomize $n_k(g)$ citations across all patents that have characteristics g and satisfy sample selection constraints from Section A.2.³⁵ As a result of this randomization, every patent can make a citation with the same probability. Denote the total number of such patents, citing and non-citing, by $N_k(g)$. Then every patent makes a citation with probability

$$p_k(g) = \frac{n_k(g)}{N_k(g)}$$

A.4.3 Step 3: Equalize Citation Rates and Patenting

This Monte-Carlo exercise is similar to the previous one except the details on the randomization of citations. To equate citation rates and patenting, I assume that $n_k(h)$ citations are allocated randomly to firms with the same probability. In other words, I assume that all firms have the same number of patents. Formally, denote by $\mathcal{F}_k(g)$ the set of firms that have at least one patent with characteristic g. Each citation out of $n_k(g)$ is randomly allocated to firm $j \in \mathcal{F}_k(g)$ with probability

$$\frac{1}{|\mathcal{F}_k(g)|}$$

where $|\mathcal{F}_k(g)|$ is the number of firms in $\mathcal{F}_k(g)$.

A.4.4 Step 4: Aggregation

I repeat the randomization procedures 300 times, and each time I compute the counterfactual concentration of citations for patent k. For an exercise with equal citation rates, denote the concentration measure for patent k in round s by $\mathcal{RC}_{k,s}$. For an exercise with both equal citation rates and the number of patents, denote the concentration measure for patent k in round s by $\mathcal{PC}_{k,s}$. I compute the median and the 95th quantile based on the distribution $\{\mathcal{RC}_{k,s}\}_{s=1}^{300}$ and $\{\mathcal{PC}_{k,s}\}_{s=1}^{300}$. These moments are denoted by $\mathcal{RC}_k(q)$ and $\mathcal{PC}_k(q)$, where q denotes quantile. Then I aggregate these measures across patents in the same way as with the actual concentration, see equations (A.3) and (A.4).

³⁵Citations should be within a 5-year window from a grant day of a cited patent. I also exclude citations from patents assigned to individual inventors or with missing assignee information

A.5 Movement of Inventors and Citation Patterns

To distinguish whether the concentration of citations is driven by firms or inventors, I track citations of inventors who worked for multiple companies. I compute the concentration measure similar to the one in Section 2 but within inventors-movers, and then I do the decomposition of the concentration measure similar to the one in Section 2.3. This exercise follows the same procedure as the Monte-Carlo exercise in Section A.4 except that the sample is restricted to inventors who worked in multiple companies, and citations are randomized within an inventor.

To increase the sample size, I consider all citations rather than the ones within a five year window. As a result, I consider the trend in citation patterns for cited granted patents until 2009, so that they have 10 years to accumulate citations. I also focus on the sample of patents granted to publicly listed firms in Compustat.³⁶ Moreover, I exclude patents assigned to multiple companies because it is impossible to distinguish which company an inventor represents.

For each patent, I compute the distribution of citations across inventors. I leave only patents that received at least 20 citations from one inventor. The results are robust to other thresholds. This is done in order to ensure greater variability in the concentration measure. For example, if an inventor cited a patent only one time, then this patent would always receive a citation from one firm only, and the within-inventor concentration measure would always be 100%. For each citing patent, I find all patents that were filed by the same inventor in the same U.S state or foreign country and the same main subgroup category in Cooperative Patent Classification system.³⁷ Patents should also be applied in the same time period: I find all patents applied in the period $[t_j^a, t_j^a + 2]$ where t_j^a is the application year of the citing patent. Then, I equate citation rates across firms by randomizing citations within each inventor across all these patents with similar characteristics: citing ones and control patents that are observationally similar to the citing ones. I remove citing patents where no inventor worked in at least two companies and filed for similar patents. To equate citation rates and patenting, I randomize citations across firms that had observationally similar patents filed by the same inventor. The procedures are the same as in Section A.4.

The final data set is the following. Each cited patent has at least one citing inventor who filed similar patents in multiple firms. I compute the actual concentration of citations within each of these citing inventors (if there are many). Next, I compute the same concentration in Monte-Carlo simulations where citations are allocated randomly. For each cited patent, I take

³⁶Matching of patents to Compustat firms is cleaner in a sense that I use the data from Autor et al. (2020) to control for potential subsidiary-parent relationship. If a patent is granted to a subsidiary of a certain firm, I match it to the parent company. Therefore, when the same inventor has patents in two firms in Compustat, these firms are more likely to represent different organizations relative to cases where the same inventor has patents in two private firms or foreign firms not listed in the U.S.

 $^{^{37}}$ If there are several inventors in the citing patent, I do this procedure for each of them.

the average of the concentration measures across all citing inventors-movers. This gives withininventor actual and counterfactual average concentrations of citations for each cited patent. Then, I aggregate within and across technological classes in a way similar to Section 2. Figure 3 shows the results. The average within-inventor concentration measure is significantly higher relative to the 95th quantile of the same measure in Monte-Carlo simulations. This means that citations are driven by firms rather than inventors: inventors tend to cite different patents in various companies despite doing similar technologies.

A.6 Details on Trade Secret Litigation

Using Lex Machina data, I identify 2541 patents that were involved in both patent and trade secret litigation (for cases filed between 2001 and 2021). I exclude patents granted after 2014 to leave 5 years for accumulation of citations.

For each patent involved in trade secret litigation ("treated" patents), I find control patents that have only been involved in patent litigation, without any trade secret claims. I use four criteria, each with progressively stricter conditions, to select control and "treated" patents. First, control patents should be assigned to the same firm and have the same grant year as the "treated" patent. Second, in addition to the first condition, control patents should have approximately the same number of citations as the "treated" patent: between 0.8 and 1.2 of the number of citations that the "treated" patent receives. Third, control patents should come the same CPC group as the "treated" patent. Finally, I focus on "treated" patents assigned to plaintiffs only.

To increase the sample size, I trace citations from all years, rather than limiting citations to a 5-year window. The reason for this approach is to ensure that patents accumulate enough citations for the computation of the concentration measure. Notice that patents involved in in trade secret litigation are not necessarily the most cited ones. Since both "treated" and control patents have the same grant years, the results are not biased due to truncation of the citation data. I compute the concentration of citations using two measures: Herfindahl-Hirschman Index (HHI) and the share of citations coming from the most citing firm ("Top Share"). For each "treated" patent, I compute the difference between its concentration and the average concentration of citations for its control patents. Then I take the average of these differences across all "treated" patents.

Next, I compute the counterfactual distribution of the difference in the concentration if citations were random. Specifically, for both "treated" and control patents I find all patents sharing the same application year, technological class (main subgroup CPC), and location of inventors. Then, I randomize citations to equalize citation rates across patents. The details are given in A.4.

To explain the importance of such randomization, consider the following example. Suppose a patent involved in trade secret litigation (the "treated" patent) receives all of its citations from one firm. Therefore, the concentration is equal to 1. All patents of this firm come from the same application year, technological class, and location of inventors. This firm is the only one who has patents with such characteristics. Therefore, even under randomization of citations the concentration for the "treated" patent would be equal to 1.

Suppose the "treated" patent has one control patent. The control patent receives an equal number of citations from two firms. Therefore, the concentration is equal to 0.5. Suppose that, as with the "treated" patent, these two citing firms specialize in their respective technologies: one firm is a patent monopolist in technology A and another firm is a patent monopolist in technology B. There are no other firms who have patents in these technologies. Therefore, even under randomization of citations the concentration for the control patent would be equal to 0.5.

The actual difference in the concentration is 1 - 0.5 = 0.5. However, this difference is driven by specialization of citing firms and would be observed even under random citations. I show that the actual difference in the concentration of citations between patents with and without trade secret claims is significantly higher relative to the difference explained by observable patent characteristics.

A.7 Details on Import Competition from China

Section 4.4 estimates how import competition from China affected citations patterns. For this exercise, I follow the methodology in Autor et al. (2020) (Appendix B.3) for the analysis at the technological class level. Specifically, I do the following steps.

First, I take the set of the top 1% of the most cited patents (*Main* sample defined in appendix A.2). For the specification with corporate patents, I leave only patents assigned to corporate firms (both public and private). Denote by f_k the assignee of patent k and by t_k^g the grant year of patent k. I group patents into five 7-year periods based on the grant year. Formally, I define the following sets

$$S_{1977} = \{k \in Main : f_k \in \text{Corporate}, t_k^g \in [1977, 1983]\}$$

$$S_{1984} = \{k \in Main : f_k \in \text{Corporate}, t_k^g \in [1984, 1990]\}$$

$$S_{1991} = \{k \in Main : f_k \in \text{Corporate}, t_k^g \in [1991, 1997]\}$$

$$S_{1998} = \{k \in Main : f_k \in \text{Corporate}, t_k^g \in [1998, 2004]\}$$

$$S_{2005} = \{k \in Main : f_k \in \text{Corporate}, t_k^g \in [2005, 2011]\}$$

where the sample Main is defined in Section A.2.

Second, for each set S_t and each technology class I compute the aggregated concentration measure. I take the simple average across concentration measures. Autor et al. (2020) provides a mapping between USPC technological classes and SIC industries. Moreover, there exists a matching from USPC to NBER technology categories that will be used as controls. Therefore, for technology classes I use the USPC system. Denote the aggregate concentration measure for technology class j and set S_t by $C_{j,t}$.

Third, given the constructed $C_{j,t}$ the analysis proceeds as described in Section 4.4. Data construction with non-corporate patents is the same except that in the first step I leave only non-corporate patents from the *Main* sample.



B Figures and Tables

Figure B1: Decomposition of the Concentration of Citations Based on the Textual Similarity Between Patents

This figure does a decomposition similar to Figure 2 using textual similarity of patents. Specifically, in addition to application years and technological classes, I also control textual similarity of abstracts between non-citing patents and citing patents. I use BERT model to measure textual similarity. The details are given in Appendix A.4.



Figure B2: Concentration of Citations Within Technological Classes

Panel (a) shows the aggregate concentration of citations that is driven by changes within technological classes. In the baseline specification, I aggregate concentration measures across classes by taking an average weighted by the number of patents in a class. The dotted red line shows the concentration in which the average across classes is unweighted. In Panel (b), for each technological class (a group category in CPC) I compute the ratio of the average concentration between 2009 and 2014 to the average concentration between 1976 and 2000. Panel (b) shows the distribution of the increase in the concentration measure across classes.



Figure B3: Number of Citations and Concentration

Panel (a) shows the average number of citations by years. Panel (b) shows the relationship between the average number of citations and the concentration. The dotted line shows the relationship based on patents granted in all years. The lines with circles and asterisks show the results for patents granted before and after 2000, respectively.



(a) Corporate patents, top 1%, top 10%, exclude superstar firms in patenting



(b) Alternative aggregation: citation-weighted average and median within classes.



of backward citations.

(d) Exclude Citations from common lawyers and group patents from one family.

Figure B4: Robustness for the Concentration of Patent Citations

These figures show robustness exercises for the concentration measure in Figure 1. Panel (a) shows that the concentration of citations in the sample of corporate patents only. I also consider different thresholds for the most cited patents: top 5% and 10%. Finally, I exclude citations from superstar firms in terms of the number of patents. Specifically, in each year and group level in Cooperative Patent Classification I find top 1% of firms in terms of the number of patents, and exclude their patents from the sample of citing patents. Panel (b) shows the results if one uses the citation-weighted average or the median instead of the average to aggregate concentration measures within technological classes. Panel (c) shows the concentration in the sample without self-citations of firms to themselves. It also shows the concentration in the sample without top 1% of patents in terms of the number of backward citations. Figure (d) shows the results in the sample without citations between patents sharing a common law firm. I also group citations from patents from the same within-country family (continuations, continuations-in-part, divisionals) as a single citation.



Figure B5: Concentration of Citations: Examiners vs. Non-examiners

This figure shows the concentration of citations across firms in which I separate citations from examiners and non-examiners. The USPTO started to distinguish citations from examiners in 2001. The dashed lines show 95th quantiles of the same measures in Monte-Carlo simulations in which citation probabilities are equalized across firms within the same (application year, technological class, location of inventors), see details in Appendix A.4.



Figure B6: Concentration of Citations for Lawyers Who Represented Several Firms

This figure shows the concentration of citations across firms within lawyers who represented multiple companies. I compare patents with similar characteristics: the same main subgroup level in Cooperative Patent Classification and application time (within 2 years from the citing application year). The solid line shows the actual aggregate within-lawyer concentration of citations across firms measured by the share of citations coming from one firm only. The dashed line shows the 95th quantile of the same measure in Monte-Carlo simulations where citations are allocated randomly within a law firm.



Figure B7: Concentration of Citations at the Firm Level

Panel (a) shows the results for the concentration defined by the Herfindahl-Hirschman Index of citations across firms, see equation (2.3) on page 16. The solid line shows the baseline concentration using all citations. The dashed line shows the concentration without self-citations. The dotted line shows the concentration based on citations between US publicly listed firms. Panel (b) provides a robustness check for the firm-level concentration defined as the share of citations from the most citing firm. For each firm in a year, I define the concentration based on citations within five years to the firm's patents granted in this year. The aggregate measure is defined as the average concentration across firms weighted by the number of citations.



Figure B8: Concentration Over Time: Within-Firm, Between-Firm, Entry, and Exit

This figure decomposes the rise in the concentration of citations, defined in equation (2.3) on page 16, from 2001 to 2014 into within-firm, between-firm, and entry/exit components (Melitz and Polanec 2015). For each component, I plot how the average concentration would evolve if it were driven by this component only. For years t and t + s, the decomposition considers changes within and between firms that received citations both in t and t + s ("surviving firm"). Additionally, it accounts for the impact of firms that stopped receiving citations by t + s ("exiting firm") and those that began receiving them in t + s ("entering firms"). Around 97.5% of the rise in the concentration from 2001 to 2014 is explained by changes in the concentration within surviving firms. If the concentration within these firms stayed constant, the average concentration to firms with low concentration of citations. On average, entrants have higher concentration relative to surviving firms, and they explain 45% of the rise in the aggregate concentration. Exiting firms also have higher concentration relative to surviving firms, and they explain 45% of the rise in the aggregate concentration.



Figure B9: Decomposition of the Firm-Level Concentration of Citations

This figure shows a decomposition similar to Figure 2 for the firm-level concentration of citations. There are two main differences from the computation in Figure 2. First, I allocate citations randomly within the set of publicly traded companies, rather than across all firms. The random allocation of citations follows the procedure described in Appendix A.4. Second, I aggregate citations at the firm level, as defined in Section 2.6.



Figure B10: Role of Partners in the Rise of the Concentration

This figure shows the share of the rise in the concentration of citations that can be explained by business partners. Panel (a) shows the aggregate role of partners. Formally, $\frac{\mathcal{H}_{2014}^{\mathcal{P}} - \mathcal{H}_{2001}^{\mathcal{P}}}{\mathcal{H}_{2014} - \mathcal{H}_{2001}} \approx 0.835$, where \mathcal{H}_t is the aggregate concentration from equation (2.3), and $\mathcal{H}_t^{\mathcal{P}}$ is the aggregate concentration from partners based on equation (3.1). Panel (b) shows a similar decomposition within partners across different types of relationships.



Figure B11: Decomposition of the Concentration Based on Industries

This figure decomposes the concentration into the roles of firms from the same industry as a cited firm and firms from other industries. The decomposition is similar to the one in equation (3.1): instead of partners and other firms I define cited-citing firm pairs as coming from the same 4-digit Standard Industry Classification Industry or from different industries.



Figure B12: Decomposition of the Concentration within Partners

This figure shows a decomposition similar to Figure 2 for the firm-level concentration of citations among business partners. Consider patents assigned to firm A and citing a patent of firm B. I find all patents with the same characteristics as the citing ones (see Section 2.3) and assigned to firms with the same relationship to firm B as the relationship between A and B during a comparable period. For instance, if A is a supplier to B in 2005, I find other suppliers to B reported from 2003 to 2007 that had patents with the same application year and narrow technological class as patents from firm A. I allocate citations randomly following the procedure described in Appendix A.4. If the rise in the concentration of citations among partners were driven by changes in the composition of partners or their patenting activity, the actual concentration would be similar to the concentration with random citations. However, the actual concentration has significantly deviated from the random one. Section 2.3 provides more details on this approach.



Figure B13: Distribution of Citations, Robustness to Figure 4 With a Broader Sample

This figure shows the distribution of patent citations across different types of relationships between cited and citing firms. In Figure 4, I consider citations between publicly traded US companies. In the sample of firms for this figure, only one firm is required to be a publicly traded US company. Panel (a) shows the share of citations coming from business partners. Panel (b) shows the distribution of citations coming from business partners.



Figure B14: Test for Strategic Citations Using Movement of Patent Attorneys

This figure shows the results of a test for the role of strategic citations. I compute the concentration of citations within attorneys who filed similar patents at different business partners of a given firm. The difference from Figure B6 is that I focus on the sample of citations among partners rather than all firms. I compare patents with similar characteristics: the same main subgroup level in Cooperative Patent Classification and application time (within 2 years from the citing application year). The solid line shows the actual aggregate within-lawyer concentration of citations across firms measured by the share of citations coming from one firm only. The dashed line shows the 95th quantile of the same measure in Monte-Carlo simulations where citations are allocated randomly within an attorney.



Figure B15: Trade with China. Actual Concentration vs. Counterfactual One with Equalized Citation Rates

This figure shows the concentration measure from Section 2 (Figure 1) for different technological classes divided into quartiles based on their exposure to import competition from China, ΔIP_{i2} in (4.1). The solid lines show the concentration measures over time for the technologies most (the blue line) and least (the green line) exposed to the competition from China. The dashed lines show the counterfactual concentration measures that are evolved due to changes in patenting across firms holding firms' citation rates fixed at the level of 2000. Formally, for patents granted after 2000 I compute the median concentration measures in the in the Monte-Carlo simulations where citations are allocated randomly across patents sharing the same application years, technological classes, and locations of inventors (see Appendix A.4). I add the difference between the actual and the counterfactual concentration measures in 2000.

Туре	Data	Description from FactSet	Group	
Customer	FactSet	Entities to which the source company sells	Suppliers and	
	Segments	products/services.	Customers	
Supplier	FactSet	Entities from which the source company	Suppliers and	
	Segments	purchases goods or services.	Customers	
Manufacturing	FactSet	Entities who provide paid manufacturing	Suppliers and	
		services to the source company.	Customers	
Marketing	FactSet	Entities who provide paid marketing	Suppliers and	
		and/or branding/advertising services to	Customers	
		the source company.		
Distribution	FactSet	Entities whom the source company	Suppliers and	
		pays to distribute this company's	Customers	
		products/services.		
In-Licensing	FactSet	Entities from whom the source company	Licensing	
		license products, patents, intellectual		
		property, or technology		
Out-Licensing	FactSet	Entities to whom the source company	Licensing	
		licenses products, patents, intellectual		
		property, and technology; also entities		
		where the source company is paid by the		
		target entity, commonly upfront and in		
		periodic future payments.		
Research	FactSet	Entities collaborating with the source	Research	
Collaboration		company for research and development,	Collaboration	
		generally for new product development,		
		common between science companies and		
		between technology companies. This		
		designation is applicable for products in		
		development, not marketed.		

Table B1: Types of Business Partnerships Between Firms

Continued on next page

Type	Data	Description	Group
Equity	FactSet	Entities in which the source company	Ownership
Investment		owns an equity stake. This designation	Stakes
		applies only when the source company	
		owns equity in another company - i.e.	
		working interests, royalties, property, or	
		well claims do not qualify for the Equity	
		Investment designation.	
Investor	FactSet	Entities which own equity in the source	Ownership
		company.	Stakes
Joint	FactSet	Entities where the source company jointly	Ownership
Venture		owns a separate company with one or	Stakes
		more companies.	
Integrated	FactSet	Entities with whom the source	Integrated
Product		company agrees to bundle standalone	Product
Offering		products/services of each company and	Offering
		are then marketed together as one	
		offering. No money is exchanged upfront,	
		and costs, risks, and profits are shared.	
NA	FactSet	Partners with an unknown relationship.	NA
		They are responsible for only 0.7% of	
		citations among partners.	
Patent	USPTO	Entity who transfers its right, title, and	Patent
Reassignor		interest in a patent or patent application	Reassignment
		to an assignee.	
Patent	USPTO	Entity who receives the right, title, and	Patent
Reassignee		interest in a patent or patent application	Reassignment
		from an assignor.	

Table B1 – continued from previous page

Note: This table describes all business partnerships used in Section 3. The description for all relationships except NA and Patent Reassignment comes from FactSet's Data and Methodology Guide. In the analysis of citations between partners, I use the grouping of relationships from the last column.

Concentration Measure	HHI				Top Share				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Δ Concentration	2.11 (0.44)	4.78 (0.59)	6.19 (1.25)	8.49 (1.33)	2.53 (0.62)	5.07 (0.82)	4.53 (1.57)	5.99 (2.14)	
Controls: Same Firm and Time Same N of Citations Same Tech Class Plaintiffs Only	Yes	Yes Yes	Yes Yes Yes	Yes Yes Yes Yes	Yes	Yes Yes	Yes Yes Yes	Yes Yes Yes Yes	

Table B2: Trade Secret Litigation and Concentration of Citations

Notes: This table shows the average difference in the concentration of citations between patents involved in litigation with and without trade secret claims. Columns 1–4 measure the concentration as the Herfindahl-Hirschman Index (HHI), and columns 5–8 — as the share of citations coming from the most citing firm. All measures are multiplied by 100. The numbers in brackets show standard errors of the difference in the concentration of citations when citations are allocated randomly within the same application years, technological classes, and locations of inventors (see Appendix A.4). I use four sets of controls. First, I require control patents (patents involved in patent infringement litigation but without trade secret claims) to be from the same firm and grant year. Second, I require control patents to have approximately the same number of citations as the treatment patents (patents involved in litigation with trade secret claims): between 0.8 and 1.2 of citations. Third, I require patents to be from the same CPC group. Finally, I focus on patents assigned to plaintiffs only.

		Non-Corporate Patents				Lag Outcomes				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	
Δ Tech Class Exposure	-1.68	-1.58	0.90	4.33	5.04	-0.46	-0.47	-1.31	-1.42	
to Chinese Imports	(1.85)	(1.81)	(1.21)	(5.33)	(5.67)	(0.64)	(0.65)	(1.10)	(1.11)	
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	
Δ Citations		Yes	Yes	Yes	Yes		Yes	Yes	Yes	
2 Lags of outcomes			Yes	Yes	Yes					
11 sectors, 6 Tech				Yes	Yes			Yes	Yes	
Software Patents					Yes				Yes	

Table B3: Placebo Tests: Trade with China and Increase in the Concentration of Citations

Notes: This table shows the results for the falsification tests in specification (4.2) and Table 2. Changes in US import exposure are instrumented with Chinese exports to non-U.S. high-income markets (Autor et al. (2020)). In columns (1)-(4), I regress the change in the concentration of citations for non-corporate patents on the changes in import competition from China. In columns (5)-(7), I regress the change in the concentration measure pre-period (pre 1991) on future changes in import exposure. Standard errors are clustered at the technology class level. All specifications are weighted by the number of Compustat-matched U.S.-inventor patents in a technology class.
C Theory Appendix

In this section, I provide various extensions to the model from Section 4.2. First, I incorporate investments in R&D into the baseline model. Second, I consider a continuum of knowledge components, accounting for reverse engineerable and tacit knowledge. Third, I assume that the incumbent can manufacture its technology in a separate market. Finally, I explore more general types of relationships.

C.1 Investments in R&D

Consider an additional stage t0 to the timing in the baseline model. In t0, only the generic technology with $z_1 = z_2 = \underline{z}$ exists and is freely available. The incumbent firm invests c(r) to generate an idea for a superior technology with $z_1 = z_2 = \overline{z}$ with probability r. Next, the game proceeds as before. The incumbent solves the following optimization problem at the investment stage:

$$r \cdot \max \{ \Pi_{\text{pat, pat}}(n_e), \Pi_{\text{pat, sec}}(n_e) \} - c(r)$$

The profits under both intellectual property strategies are decreasing in the number of competitors, n_e . Thus, higher competition leads to lower investments in R&D.

This extension clarifies that I focus on ex-post competition, which occurs after the R&D investments. This type of competition lowers innovation incentives (Schumpeter 1942).

C.2 A Continuum of Knowledge Components

Suppose the production function is Cobb-Douglas, $y = \exp\left(\int_0^1 \ln(z_i x_i) di\right)$, where z_i is a productivity of input $i \in [0, 1]$. The cost function is

$$c(\{z_i\}_i) = \exp\left(-\int_0^1 \ln z_i \mathrm{d}i\right)$$

where I assumed that input prices are $w_i = 1$ for $i \in [0, 1]$. The incumbent knows a technology with $z_i = \overline{z} > \underline{z}$ for $i \in [0, 1]$, where a $z_i = \underline{z}$ is a free generic technology. If partner j has access to knowledge about share α of inputs, it gains access to marginal costs $c = 1/(\overline{z}^{\alpha} \underline{z}^{1-\alpha})$. For notation simplicity, I assume that $\underline{z} = 1$, and partner j's profit is denoted as $v_j(\alpha) = v_j(1/\overline{z}^{\alpha})$.

Suppose that knowledge for share β of inputs is patented and $1 - \beta$ is kept secret. For each partner j, the incumbent decides on the share of secret knowledge to share

$$\max_{\beta_{s}\in[0,1-\beta]}\left\{v_{j}\left(\beta+\beta_{s}\right)-v_{j}\left(\beta_{s}\right)\right\}$$

Similarly to the baseline model, competing entrants can innovate based on the disclosed knowledge. Entrant k receives a random draw ψ and a technology with $z_{i,k}^e = \psi \cdot Z_i$ where $Z_i = \overline{z}$ if knowledge about input *i* is patented and $Z_i = \underline{z}$ if it is secret.

Suppose that knowledge for share β_r of inputs can be reverse enginereed. Thus, it is always patented. Knowledge for share β_t is tacit and cannot be patented. Taking these constraints into account, the incumbent maximizes the following expected profit as a function of the patented knowledge, β :

$$\Pi(\beta) = \max_{\beta \in [0, 1-\beta_r-\beta_t]} \left(\sum_{i=1}^n \max_{\beta_s \in [0, 1-\beta]} \left\{ v_j \left(\beta + \beta_s\right) - v_j \left(\beta_s\right) \right\} \right) \cdot \mathbb{P} \left[\psi \leqslant \overline{z}^{1-\beta} \right]^{n_{\epsilon}}$$

The firm faces the same trade-offs as in the baseline model. Higher reliance on secrecy decreases the likelihood of being replaced but limits the firm's ability to extract surplus from technology buyers.

C.3 Production by the Incumbent

Suppose there is one additional market where the incumbent can manufacture its technology. Therefore, it does not need to transfer knowledge to another firm in this market. The incumbent's profits in this market are denoted by π . The rest of the setup remains the same as in the baseline model. Complete secrecy is not optimal in the baseline model because the incumbent cannot extract surplus from technology buyers and gets zero profits. In this extension, the incumbent can still earn profits even under complete secrecy. Formally, the incumbent's expected profits from different intellectual property protection strategies are:

$$\Pi_{\text{sec, sec}} = \pi \cdot \left[\mathbb{P}\left(\psi \leqslant \frac{c(\underline{z}, \underline{z})}{c(\overline{z}, \overline{z})} \right) \right]^{n_e}$$

$$\Pi_{\text{pat, sec}} = \left(\pi + \sum_{j \in \mathcal{N}_s} \left[v_j(\overline{z}, \underline{z}) - v_j(\underline{z}, \underline{z}) \right] + \sum_{j \in \mathcal{N}_c} \left[v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \overline{z}) \right] \right) \cdot \left[\mathbb{P}\left(\psi \leqslant \frac{c(\overline{z}, \underline{z})}{c(\overline{z}, \overline{z})} \right) \right]^{n_e}$$

$$\Pi_{\text{pat, pat}} = \left(\pi + \sum_{j=1}^n \left[v_j(\overline{z}, \overline{z}) - v_j(\underline{z}, \underline{z}) \right] \right) \cdot \left[\mathbb{P}\left(\psi_k \leqslant 1\right) \right]^{n_e}$$

The first equation shows the profit under complete secrecy, the second one is the profit under the patent-secrecy combination, and the last is the profit under full patenting.

This extension leads to similar predictions as the baseline model. Higher reliance on secrecy results in more selective knowledge sharing. Under complete secrecy, the incumbent has no incentives to share knowledge with any customers. Increased competition leads to greater reliance on secrecy due to the rising costs of information disclosure. Since $\frac{\Pi_{\text{pat, sec}}}{\Pi_{\text{sec, sec}}}$ and $\frac{\Pi_{\text{pat, pat}}}{\Pi_{\text{sec, sec}}}$

converge to zero as the number of competitors n_e goes to infinity, at some point, keeping all knowledge secret becomes the optimal strategy.

C.4 A Generalization of the Types of Relationships

The baseline model considers a simple framework of knowledge transfers from research firms to manufacturers. In this section, I show how this setting can be generalized to other types of relationships. The notation and assumptions follow the baseline model unless otherwise noted.

Assume that there are only two firms: an incumbent and its partner. Later I show how this assumption can be relaxed. The incumbent knows a superior technology which can be used by the partner for two purposes. First, the partner can provide services or products to the incumbent. Second, it can use the incumbent's technology in other activities. For example, if the partner is an input supplier to the incumbent, it can use the technology to serve both the incumbent and potentially other customers. Denote the incumbent's profits as a function of the partner's use of its knowledge by $\pi (z_1^d, z_2^d; z_1^v, z_2^v)$. The vector $(z_1^d, z_2^d) \in \{\underline{z}, \overline{z}\}^2$ represents the productivity levels the partner has in directly serving the incumbent, such as by selling it an input. The vector $(z_1^v, z_2^v) \in \{\underline{z}, \overline{z}\}^2$ represents the partner's productivity levels in other activities. For simplicity, suppose that the partner does not get any surplus by serving the incumbent, but it can generate $v(z_1^v, z_2^v)$ by serving others. This function is increasing in z_1^v and z_2^v . I assume that $\pi(\cdot)$ and $v(\cdot)$ are symmetric in (z_1^d, z_2^d) and (z_1^v, z_2^v) , similarly to Assumption 1 in the baseline model. The baseline model corresponds to a case with $\pi (z_1^d, z_2^d; z_1^v, z_2^v) = 0$ for all productivity levels.

Patents allow the incumbent to control the use of its knowledge. However, this control is broader than in the baseline model. The incumbent controls not only whether the partner has access to the technology but also how it is used. For example, the incumbent can permit the partner to use its knowledge in direct services but restrict its use in services to other companies, implying $z_i^d = \overline{z}$ but $z_i^v = \underline{z}$ for $i \in \{1, 2\}$. Suppose the incumbent always benefits when the partner is more productive in direct activities for the incumbent, $\frac{\partial \pi}{\partial z_i^d} \ge 0$. However, the incumbent's profits might decrease or increase from the partner's other activities. For instance, the incumbent's profits might decrease if its input supplier sells inputs to its competitors. I assume that such competitors cannot use the incumbent's knowledge directly but may affect the incumbent's profits through their use of the partner's services.

In addition to competitors who cannot use the incumbent's knowledge directly, there are competing entrants who can innovate based on the incumbent's knowledge and potentially replace it. Suppose each entrant has their own partner. Similar to the baseline model, the most efficient firm among the incumbent and the entrants wins the competition for final consumers. Due to entry fixed costs, only the most efficient firm enters the market. The incumbent makes a take-it-or-leave-it offer to the partner on the patented knowledge. The incumbent's expected profits from full patenting, patent-secrecy combination, and compete secrecy are the following:

$$\Pi_{\text{pat, pat}} = \left(\max_{z_i \in \{\underline{z}, \overline{z}\}, i=1,2} \left\{ \pi \left(\overline{z}, \overline{z}; z_1, z_2 \right) + \left[v(z_1, z_2) - v(\underline{z}, \underline{z}) \right] \right\} \right) \cdot \left[\mathbb{P} \left(\psi \leqslant 1 \right) \right]^{n_e}$$

$$\Pi_{\text{pat, sec}} = \left(\max_{z_i \in \{\underline{z}, \overline{z}\}, i=1,2} \left\{ \pi \left(\overline{z}, z_2; z_1, z_2 \right) + \left[v(z_1, z_2) - v(\underline{z}, z_2) \right] \right\} \right) \cdot \left[\mathbb{P} \left(\psi \leqslant \frac{c(\overline{z}, \underline{z})}{c(\overline{z}, \overline{z})} \right) \right]^{n_e}$$

$$\Pi_{\text{sec, sec}} = \left(\max_{z_i \in \{\underline{z}, \overline{z}\}, i=1,2} \left\{ \pi \left(z_1, z_2; z_1, z_2 \right) + \left[v(z_1, z_2) - v(z_1, z_2) \right] \right\} \right) \cdot \left[\mathbb{P} \left(\psi \leqslant \frac{c(\underline{z}, \underline{z})}{c(\overline{z}, \overline{z})} \right) \right]^{n_e}$$

The probability terms are the same as in the baseline model: higher reliance on secrecy decreases the probability of being replaced. Conditional on surviving the competition from the entrants, the incumbent's profit consist of direct profits and the surplus extracted from the partner. Under full patenting, the partner always gets access to both knowledge components because it is valuable for the incumbent, $z_1^d = z_2^d = \overline{z}$. However, the incumbent might find it optimal to restrict the use of its technology for other activities, setting $z_1^v = z_2^v = \underline{z}$, if the gains from these activities to the partner are lower than the costs for the incumbent. If one or both of the knowledge components are secret, the incumbent decides whether to share them with the partner, given that it cannot prevent the partner from using them in other activities, such as serving other customers. These restrictions are reflected in the incumbent's choice set. Under full patenting, the incumbent can choose any element from $(z_1^d, z_2^d, z_1^v, z_2^v) \in \{\underline{z}, \overline{z}\}^4$. Suppose only the first component is patented. Then, the incumbent chooses $(z_1^d, z_2^d, z_1^v, z_2^v)$ with the restriction $z_2^v = z_2^d$. If both components are secret, the restrictions are $z_1^d = z_1^v$ and $z_2^d = z_2^v$.

This extension can be generalized to multiple partners. Suppose the incumbent's technology can be used to sell n products. All products are independent in profits and require separate partners. Denote direct and indirect profits for product j by $\pi_j(\cdot)$ and $v_j(\cdot)$, respectively. Then, the analysis is the same as described before.

The predictions of this extension are the same as those of the baseline model. Higher competition n_e increases the costs of disclosure and leads to higher reliance on secrecy, which might lead to more selective knowledge sharing.