http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–29
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Roadblock to Innovation: The Role of Patent Litigation in Corporate R&D

Filippo Mezzanotti[a]

[a] Kellogg School of Management, Northwestern University, Evanston, Illinois 60208
Contact: filippo.mezzanotti@kellogg.northwestern.edu, https://orcid.org/0000-0001-9720-1906 (FM)

Received: September 11, 2019
Revised: April 7, 2020; June 19, 2020
Accepted: July 10, 2020
Published Online in Articles in Advance: February 2, 2021

https://doi.org/10.1287/mnsc.2020.3816

Copyright: © 2021 INFORMS

Abstract. I examine how patent enforcement affects corporate research and development (R&D), exploiting the legal changes induced by the Supreme Court decision *eBay v. MercExchange*. This ruling increased courts' flexibility in remedying patent cases and effectively lowered the potential costs of patent litigation for defendants. For identification, I compare innovative activity across firms differentially exposed to patent litigation before the ruling. Across several measures, I find that the decision led to a general increase in innovation. This result confirms that the changes in enforcement induced by the ruling reduced some of the distortions caused by patent litigation. Exploring the channels, I show that patent litigation negatively affects investment because it lowers the returns from R&D and exacerbates its financing constraints.

History: Accepted by Gustavo Manso, finance.
Supplemental Material: Data and the online appendices are available at https://doi.org/10.1287/mnsc.2020.3816.

Keywords: patent litigation • innovation • corporate investments

1. Introduction

The main goal of the patent system is to protect innovation and, thus, to spur growth. Whether this goal is achieved depends on how patents are defined and protected, which itself depends on how the legal system resolves intellectual-property disputes. Indeed, the courts appear to have played an increasingly important role in the patent system (Cohen et al. 2019). Over the last 30 years, lawsuits involving patents more than tripled (Figure 1), and their estimated cost surpassed $300 billion (Bessen et al. 2018). In part, this shift in patent enforcement has been connected with the increase in patent-holders' rights over the past decades (Jaffe and Lerner 2011).[1]

In principle, increasing the rights of patent-holders should have positive effects on innovation. However, this simple logic may fail for several reasons. For instance, stronger rights for patent-holders may actually have detrimental effects on innovation when we account for the role of competition (Arrow 1962, Aghion and Howitt 1992) or we consider innovation as a cumulative process (Scotchmer 1991, Bessen and Maskin 2009). Furthermore, the idea that stronger patent rights will increase innovation also hinges on the assumption that patent enforcement is sufficiently effective in resolving legal disputes. In practice, the patent system is characterized by a lot of uncertainty (Lemley and Shapiro 2005), and enforcement tends to be a noisy process (Jaffe and Lerner 2011). In this context, strong patent rights may actually increase the cost that frivolous patent litigation can impose on innovative firms, which, in turn, may curb innovation and growth (Boldrin and Levine 2002, Bessen and Meurer 2008b). Therefore, to understand the connection between patent protection and innovation, we need to also consider how changes in patent rights affect enforcement. This issue is particularly crucial today, given the increasing importance of intangible assets to the modern corporation (e.g., Corrado and Hulten 2010, Eisfeldt and Papanikolaou 2014, and Peters and Taylor 2017).

Consistent with this hypothesis, this paper shows that a reduction of patent holders' rights—which lowered the costs of patent litigation for defendants—may have positive effect on corporate innovation. To examine this issue, I develop a new research design that exploits a landmark Supreme Court decision, *eBay v. MercExchange L.L.C.*, 547 U.S. 388 (2006). The ruling ended the practice of automatically granting a permanent injunction after a patent violation. The issuance of an injunction forces the defendant firm to shut down any operation related to the violated technology, regardless of the nature or magnitude of the infringement. Given this large operational risk, the presence of automatic injunction gave patent holders a very strong bargaining position in intellectual property (IP) disputes and negotiations (Section 2.1). With the 2006 decision, courts were allowed to decide on a case-by-case basis whether an injunction was appropriate, therefore increasing court flexibility to

Figure 1. (Color online) Number of Cases Involving Patents

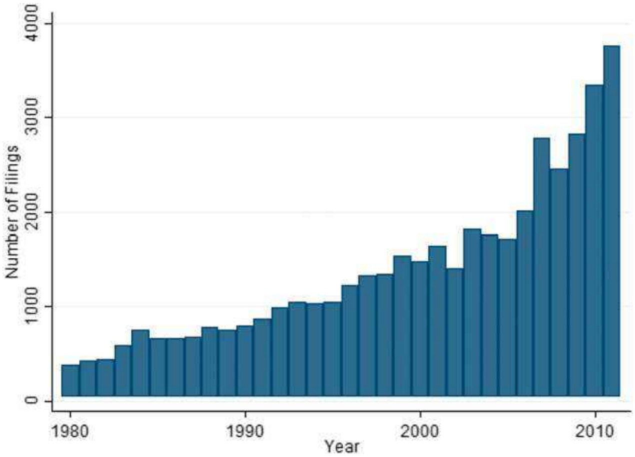

Notes. This plot reports the number of filings involving patents of any type per year of filing, between 1980 and 2012. The data come from WestLaw-ThomsonReuters, which collected filings information from public records. Data are plotted at the docket-number level; therefore, they do not account for the fact that each case can involve multiple defendants. More information on the data is available in Section 3.

remedy patent disputes and lowering the potential cost for defendants.

A prominent example that demonstrates the important role of automatic injunction in patent litigation before 2006 is the lawsuit initiated by NTP Inc. against Research in Motion (RIM), the maker of BlackBerry, in the early 2000s. The district court sided with NTP and found RIM guilty of infringing on a few of NTP's patents. With the objective of avoiding an injunction, RIM started to negotiate with NTP. Even if the infringement covered only a small fraction of the portfolio of patents used to run the BlackBerry system, an injunction order would likely have led to a shutdown of the whole system. Leveraging on its ability to obtain an injunction, NTP was able to negotiate a record settlement of more than $610 million, approximately half of RIM's revenues in the previous year. Interestingly, years later, some of the claims contained in NTP patents were deemed invalid, as RIM had argued initially. According to several experts, the presence of an automatic injunction played a fundamental role in the decision to settle early and the size of the transfer.[2]

In general, the overall impact of the *eBay v. MercExchange* decision is ex ante ambiguous. On the one hand, removing the automatic injunction may lower deterrence against violations (e.g., Epstein 2008), therefore reducing the incentives to innovate. On the other hand, this increase in court flexibility may actually positively affect innovation by reducing the costs that patent litigation may impose on innovative firms. In fact, the presence of a "near-mandatory" injunction increases the extent to which companies can be held up by a plaintiff (Lemley and Shapiro 2006, Shapiro 2010, 2016a), as is clear from the RIM–NTP case. In general, leveraging on the high degree of uncertainty in the patent system (Lemley and Shapiro 2005), the injunction threat may help plaintiffs to obtain large settlements, even if accusations are based on frivolous claims or minor violations.[3] In this context, the new rules should reduce the hold-up costs of litigation, therefore positively affecting the incentives and financial ability to do research and development (R&D) investments.[4]

The post-eBay changes in patent enforcement suggests that the impact that the decision had on litigation costs may have been substantial (e.g., Tang 2006, Bessen and Meurer 2008a, Venkatesan 2009, Shapiro 2010, and Holte 2015). In aggregate, Chien and Lemley (2012) found that the likelihood of obtaining an injunction after the Supreme Court ruling decreased by at least 25%. However, injunctions are still granted in the majority of cases, suggesting that this remedy is still available to firms in case of a violation. Furthermore, the decline in injunction rates was larger for cases where the firms involved were not competitors or when the case involved a patent assertion entity (PAE) (Seaman 2016). In line with this result, I find that public PAEs experienced large negative returns around the time of the court decision, with average cumulative returns of about −10%.

Given the theoretical ambiguity, this paper empirically studies the effect of the ruling by estimating a difference-in-difference model that exploits variation in firm exposure to patent litigation in 2006 to identify companies that are more likely to be affected by the eBay ruling. The intuition for this choice is simple: Although the shock potentially touched every firm, companies that operate in areas where patent litigation is more intense should be relatively more affected by the decision. I construct this measure of exposure as a weighted average of litigation activity across all the technology areas in which a firm operates, where the weights are given by the share of patents in each U.S. Patent and Trademark Office (USPTO) technology class for the firm. Therefore, this measure captures exposure to litigation coming from the area in which the company operates, and it is orthogonal to endogenous decisions of the firm to engage in patent litigation. As a validation, I show that heterogeneity in exposure to the shock predicts variation in abnormal returns the day in which the decision was made public. In particular, firms more exposed to litigation outperformed less exposed firms when the ruling was released.

As a first step, I use this model to examine how the decision affected patent applications for a sample of almost 20,000 innovative firms. Firms that were more exposed to litigation before the decision increased patenting more after the decision. In particular, a one-standard-deviation increase in exposure leads to a 3%

higher application rate—which translates into almost one extra patent in the two years after the shock—and a 2% increase in the probability to patent anything. As discussed in the paper, these results are not driven by differential trends across heterogeneously exposed firms, and they are robust to control for industry trends—measured by the main technology area of the firm (Hall et al. 2001)—as well as other confounding factors. Furthermore, I also argue that my findings cannot be explained by other contemporaneous legal changes.

To better characterize the effect of the decision on innovation, I then extend this analysis in different directions. First, I find that the change in enforcement led to a shift in patent quality. Although increasing their patenting relatively more, firms more exposed to patent litigation did not lower the average quality of their output. Instead, they became relatively more likely to develop a potential "breakthrough innovation" (Kerr 2010, Lin et al. 2020), defined as a patent that is at the top of the citation distribution within the same patent class and year group. Second, using different metrics (Abrams et al. 2013, Srinivasan 2018), I show that the shock also decreased the share of defensive patents for highly affected firms. Third, I show that the ruling had a positive effect on R&D investment for public innovative firms.

Altogether, these results confirm that the positive effect on patenting did not simply reflect an increase in defensive activity (Hall and Ziedonis 2001) or a shift in the incentives to file for a patent. Instead, these results are consistent with the idea that the change in enforcement caused by the ruling positively affected innovation. This suggests that the reduction in plaintiff bargaining power reduced some of the distortions caused by the litigation environment. In line with this hypothesis, I also show that the R&D effect was more pronounced for firms that were more likely to be involved in litigation as a defendant.

Finally, I examine how an improvement in enforcement rules may affect the process of innovation. First, enforcement seems to influence innovation because it determines the relative returns of different R&D projects. Consistent with this hypothesis, I find that after the Supreme Court ruling, firms marginally reshuffled their internal resources toward projects that are in higher litigation areas, at least at the extensive margin. Second, enforcement rules seem to also affect R&D because they exacerbate the financing problems of innovation (Brown et al. 2009, Hall and Lerner 2010). In fact, companies operating in high-litigation environments are forced to devote a larger share of resources for defensive activities (Cohen et al. 2016b, 2019) and spend more money on settlements or licensing.[5] In the presence of financial frictions, this increase in costs may negatively impact the ability to undertake investments. Consistent with this implication, firms that were more likely to be financially constrained before the decision increased R&D intensity in its aftermath. These findings establish the important role played by financial constraints in explaining the negative effects of patent litigation.

This analysis provides several contributions to the literature. First, the paper shows that the impact of a change in patent rights on innovation crucially depends on how this shift affects patent enforcement. Previous work has shown that the strength of patent system has a limited direct effect on innovation (e.g., Sakakibara and Branstetter 2001, Lerner 2002, Moser 2005, Lerner 2009, and Moser 2013). At the same time, stronger patents may negatively affect innovation indirectly, for instance, by reducing knowledge diffusion (e.g., Murray and Stern 2007, Galasso and Schankerman 2015, and Williams 2016). However, less attention has been devoted to understanding the relationship between property rights, enforcement, and innovation decisions. The eBay ruling represents a shift of patent enforcement toward principles of "proportionality." In particular, it gives courts more flexibility to balance the interests of competing parties and therefore reduces the prerogative of a patent owner. This paper shows that in the current patent system, this type of intervention may have beneficial effects for innovation by reducing distortions caused by hold-up in litigation (Shapiro 2016a).[6] Therefore, although this paper does not directly help settle the broad debate on the optimal strength of patents, it provides novel insights about the role of patent enforcement in this area.

This discussion is not surprising within the context of the law and economics literature. In general, a strict property rule—as the mandatory injunction policy in place before the ruling—works well when ownership rights are clear and easy to identify, as with tangible assets (Calabresi and Melamed 1972). If the boundaries of the assets are hard to define, like for patents (Lemley and Shapiro 2005), a strict property rule may fail to provide the best incentives, and it may be inferior to a hybrid system characterized by more flexibility (Kaplow and Shavell 1996).

Furthermore, these results provide new evidence on the real costs of patent litigation—which is central in today's research (Hall and Harhoff 2012) and policy debate (Executive Office of the President 2013). Previous work in this area—in particular, Smeets (2014) and Cohen et al. (2019)—has shown that innovation activity declines when a firm is directly targeted by patent lawsuits. Although my results are consistent with these papers, this work also extends this literature in several directions. First, changes to the patent-litigation environment—and therefore not only the direct involvement in a lawsuit—significantly affect the quantity and the direction of innovation. One implication of this result is that focusing only on the direct costs of patent litigation (e.g., Bessen et al. 2018) will underestimate the real impact. Second, my results expand their analyses by examining the channels through which litigation may affect R&D decisions. Lastly, in this area, this paper complements the contemporaneous work by Appel et al. (2017), which examines the detrimental effects of patent litigation on business creation and start-up activity.[7]

More broadly, this paper shows that changes in patent enforcement can have a significant impact on the incentives of firms to invest in R&D, therefore contributing to the finance literature studying the relationship between legal institutions and economic activity (King and Levine 1993, La Porta et al. 1997, DemirgüçKunt and Maksimovic 1998, Claessens and Laeven 2003, Lerner and Schoar 2005, Acharya et al. 2011, Ferreira et al. 2018, Hochberg et al. 2018). Previous research has demonstrated that secure property rights favor a more efficient allocation of resources and fosters growth, but in many cases, good enforcement is just as important as good rules in determining economic outcomes (e.g., Djankov et al. 2003 and Ponticelli 2016). The role of enforcement is particularly important in intellectual property because the exact boundaries of patents are hard to define (Lemley and Shapiro 2005), and, therefore, lawsuits are frequent (Lanjouw and Lerner 1998). This paper highlights the role of enforcement in innovation and suggests that, similar to other interventions (Acharya and Subramanian 2009, Hsu et al. 2014, Lin et al. 2020, Mann 2018, Moshirian et al. 2020), a fine-tuning of patent law can have substantial effects on fostering corporate innovation.[8]

The paper is organized as follows. In Section 2, I provide more background information about the Supreme Court decision and discuss its potential effects on corporate innovation. In Section 3, I present the data used in the paper, while in Section 4, I present the identification and discuss in detail the measurement of exposure to patent litigation at the firm level. In Section 5, I present the main results of my analysis. In Section 6, I discuss and test different channels through which patent litigation can affect innovation. Lastly, Section 7 concludes.

2. The *eBay v. MercExchange* Case

This section provides background information on the Supreme Court decision *eBay v. MercExchange* and discusses its possible effects on innovation. First, I analyze the importance of injunction on the preruling world and provide some background on the decision. Second, I discuss how the ruling could affect innovation, therefore setting the foundation for the hypothesis and research design. Lastly, I provide some preliminary and novel evidence of the importance of the ruling for patent enforcement.

2.1. The Role of Injunction and the 2006 Decision

With the 2006 *eBay v. MercExchange* decision, the Supreme Court revisited the norms regulating the issuance of permanent injunction in cases involving intellectual property.[9] Injunction is a remedy that can be requested by a plaintiff. If granted by a court, an injunction forces the infringer to stop using any technology covered by the contested patents, irrespective of the magnitude of the infringement. Before 2006, a plaintiff that was able to prove a violation had essentially the automatic right to obtain a permanent injunction. In other words, the norm was that "a permanent injunction should be issued when infringement was proven" (*eBay v. MercExchange L.L.C.*). Exceptions to this rule were quite uncommon and mostly due to reasons of public interest.

The availability of a quasi-automatic injunction grants a lot of power to plaintiffs in IP negotiations (Hall and Ziedonis 2001). Experts have criticized this feature of the law, arguing that the presence of automatic injunction exacerbates the hold-up problem during the negotiation between firms involved in litigation (Shapiro 2016b). In particular, the view was that the ability to leverage on an injunction threat may allow the plaintiff to obtain transfer of resources that exceed the value of the disputed technology, either before or during a formal court proceeding. In the context of intellectual property, the hold-up problem is particularly concerning because technologies tend to be characterized by high complementarity. Therefore, even an injunction granted for a relatively small violation can deeply impair a company's operations. Furthermore, the high uncertainty characterizing the patent system may exacerbate this issue. In fact, the discovery process in IP can be long and costly, and cases of involuntary infringements or false positives in court decisions tend to be common (Lemley and Shapiro 2005). In this context, even when a lawsuit is based on relatively weak claims, the threat of injunction may force the defendant into costly settlements to avoid an uncertain court procedure.[10]

The RIM versus NTP case discussed in the introduction represents a very clear example of how an injunction can magnify the cost of patent litigation. First, although the dispute involved only a few patents, the settlement was more than $600 million, almost half of RIM's previous-year revenue. This high settlement is explained by the fact that a likely injunction would have forced RIM to completely block its Blackberry sales, increasing the chance of bankruptcy for the firm. Second, RIM was forced to settle, despite the fact that most of the NTP claims were eventually found to be invalid. This invalidity was impossible to prove in court at the time of the litigation. Altogether, NTP's ability to leverage on the near-mandatory injunction was the main driver to obtain the large settlement.

This argument—which links some of the distortion in the litigation market to the presence of automatic injunction—was prevalent among academics (e.g., Bessen and Meurer 2008a), practitioners, and legal experts. For example, the American Innovators Alliance, an association of high-tech companies, claimed that, because of automatic injunction, "money that could go to productive investments is instead diverted to legal fees and settlement payments," therefore having "profound implications for technological innovation in the United States."[11] This view was also expressed by the Supreme Court in the motivation of the ruling. For instance, Justice Kennedy wrote that the threat of injunction has been extensively used "as a bargaining tool to charge exorbitant fees to companies that seek to buy licenses to practice the patent." This quote is in line with the evidence that parties accused of aggressively asserting their patents—for instance, the patent assertion entities discussed later—were actively using the threat of permanent injunction as a way to scare counterparties and, therefore, obtain larger settlements (Lemley and Shapiro 2006).

The ruling *eBay v. MercExchange*, which was made public on May 15, 2006, changed the landscape by shifting negotiation power away from plaintiffs (e.g., Tang 2006; Bessen and Meurer 2008a; Venkatesan 2009; Shapiro 2010, 2016a). Specifically, the decision stated clearly that the issuance of an injunction should not happen automatically. Instead, courts should decide on a case-by-case basis, using a four-factor test balancing "the hardships between plaintiff and defendant" (*eBay v. MercExchange L.L.C.*). In practice, the eBay case started a new hybrid system, in which monetary damages could be used instead of an injunction to remedy violations. In other words, the court recognized that a "damages award is sometimes sufficient to maintain incentives while preventing patentees from amassing disproportionate rewards, significantly injuring the public, and stifling innovation" (Carrier 2011). In the context of the policy debate, the decision was perceived as an attempt to remove some of the distortions that characterized the system, however, leaving injunction as an option when this is the only way to remedy a violation.

One important step to understand the effectiveness of this decision is to examine its impact on patent enforcement, in particular, regarding the use of injunction. Quantifying this effect is challenging, because of the obvious selection issue. In fact, the decision did not only affect how courts will make decisions, but it also changed the balance of costs and benefits to file a lawsuit. Despite this limitation, three stylized facts are identified in the literature. First, the ruling in aggregate substantially reduced the likelihood of obtaining an injunction. For instance, Chien and Lemley (2012) find that the likelihood of obtaining an injunction declined by about 25%.[12] Because firms appear now less likely to seek an injunction in the first place (Gupta and Kesan 2015), this estimate may be considered a lower bound of the actual effect. Second, despite this decline, injunction is still a valuable tool for companies seeking protection from patent violations. In fact, injunctive reliefs are still granted in the majority of cases. Third, the drop in the injunction rate is mostly driven by cases that are more likely to be motivated by the strategic and opportunistic reasons. For instance, Seaman (2016) finds that injunction rates decline across all categories of plaintiffs, but this reduction is much larger when the two parties are not or when the case involves a nonpracticing entity. In principle, this evidence is consistent with the idea that the policy was able to reduce the risk of litigation, without, however, completely removing injunction as a remedy against real violations.

Therefore, the eBay decision led to a significant shift of bargaining power from plaintiff to defendants in both court cases and out-of-court negotiations (Shapiro 2016b). Consistent with the reduction in plaintiff's power, firms heavily involved in litigation activity have tried to change their strategies to try to limit the impact of this change on their bargaining power. For instance, Chien and Lemley (2012) find that after the eBay ruling, companies started to become more active in bringing claims in front of the International Trade Commission (ITC), which—in some cases—could still issue injunctions. At the same time, Cohen et al. (2016a) discuss how firms, to make their claims more credible, increased their likelihood of filling a lawsuit rather than simply issuing demand letters.[13] The previous discussion in the literature—also consistent with a revealed preference argument—suggests that these changes could not undo the shift caused by eBay. Our evidence on nonpracticing entity (NPE) returns will corroborate these claims. However, our analyses—looking at the effect of the reform net of any strategic response—will be able to shed light on this issue.

2.2. Hypothesis Development

Given the discussion from the previous section, it is clear how the shock led to a significant reduction in the bargaining power of plaintiffs in litigation. However, the court ruling's effect on innovation is less obvious ex ante. For a firm, this shock to enforcement can have two effects. On the one hand, limiting the ability of a plaintiff to hold up a defendant in litigation implies a reduction in the cost of litigation (Shapiro 2016a, b).[14]

In turn, this shift should increase firms' incentive to innovate and also allow a company to transfer more resources from the defensive activity into R&D.[15] On the other hand, the shock can also decrease the ability of a company to deter possible violations (Epstein 2008, Holte 2015) and, as a result, lower appropriability. Even if injunction were still available to firms facing infringements, the level of ex ante deterrence perceived by the firms may still be lower than before. This alternative channel would lower the returns from innovation and, therefore, induce firms to invest less.

In light of this discussion, it is clear that the overall effect is ex ante ambiguous. The evidence on the changes in injunction rates discussed before provides some guidance for thinking about this problem. However, it is not sufficient to evaluate the full impact of the decision. In this context, the only way to evaluate the policy is to test how firms actually responded to the change in incentives. This analysis will inform us about how a change in patent enforcement—which effectively reduced the rights of patent owners—would affect the innovation activity. Furthermore, this type of analysis can provide important insights on how the risk of patent litigation can distort firm innovation.

Following most of the empirical literature in this area, the paper will start by measuring innovation activity at the firm level using patent-application counts. Patents provide a multidimensional measure of innovation activity in both private and public firms and can be used to identify shifts in the quality of innovation (Lerner and Seru 2017). However, a drawback in using patents is that this measure does not allow us to directly distinguish the impact of the shock on innovation from a shift in the incentives to file patents.[16] In particular, one concern is that the count of patent applications may confound an increase in defensive activity with an increase in innovation. This result is particularly relevant in this case, because previous research looking at the semiconductor industry in the 1990s found that an increase in hold-up can increase firms' incentive to patent for defensive reasons (Hall and Ziedonis 2001). Therefore, looking at patent counts in isolation may be problematic.

Because of these concerns, the paper will also explore innovation across several other dimensions. First, for the sample of public firms, I will explore the effect of the decision on R&D expenditure as well. Second, this study will examine the effect of the decision on the quality of patenting. Along this dimension, a standard model of innovation and patenting should provide different predictions depending on the firms' motives. In particular, if a higher propensity to patent caused the increase in patenting, I expect the quality of the output to decrease after the ruling, because the marginal project should be worse than the average patent. The same result should not be true if the increase in patenting is caused by more genuine innovation. Third, given the specific concerns, I will directly examine the impact on strategic patenting. If the increase in patenting were to be explained by more defensive activity, then the share of strategic patenting should actually increase. The same should not be the case otherwise. Although none of these tests may be perfect, taken together, they provide a great overview of the overall effect on innovation. In the last part, I will try to exploit this setting to provide valuable insights into channels through which patent enforcement can foster or hinder innovation

2.3. The Economic Importance of the Decision: The Case of NPE

Previous discussion presents the shock as a reduction in the bargaining power of plaintiffs in litigation. In this section, I provide evidence consistent with this hypothesis by examining the stock returns of a set of public PAEs. In particular, I find that the ruling led to a drop of about 10% in the stock price of these companies.

Following the literature (Tucker 2014, Feng and Jaravel 2015, Kiebzak et al. 2016, Cohen et al. 2019), I identify PAEs looking at nonpracticing entities, which are companies that generate most of their revenue through licensing and settlement fees, rather than from manufacturing.[17] This setting is a useful laboratory for several reasons. First, previous research has confirmed that NPEs extensively used injunction threats when negotiating licensing agreements or settlements (Chien and Lemley 2012). Therefore, the ruling should have somehow harmed their business model. Second, unlike for other companies, automatic injunction does not constitute a major risk for these firms because they generally do not directly use intellectual property to develop products or sell services.

Therefore, if the ruling had a big impact on patent enforcement, I expect NPEs to be negatively affected by the decision. In particular, I test this hypothesis by looking at the stock market returns of public NPEs around the time of the ruling. The main challenge in this type of analysis is that most NPEs are private. For instance, "Intellectual Ventures"—allegedly the largest NPE today—is a private firm. I start by combining two lists of NPEs, provided respectively by PatentFreedom, one of the most important firms in assessing NPE risk and now owned by RPX, and by EnvisionIP, a law firm involved in strategic IP consulting.[18] Then, I identify the firms in these lists for which returns information is available in the Center for Research in Security Prices (CRSP) around the date of the event. This analysis yields a final list of 10 companies.[19]

Studying the returns of these companies around the decision, I identify four important stylized facts.[20] First, on the day of the decision, these firms experienced a drop in stock price of 3.3%–3.8%, depending on whether I look at raw returns or abnormal returns. These effects are highly significant, with the Sharpe ratios ranging between 4.08 and 4.75. Second, firms suffered negative returns also in the couple of days before the decision (Figure 2). Although the largest one-day drop occurred the day of the Supreme Court ruling, these stocks also lost value in the three days before it. Examining the abnormal returns with respect to the S&P500, the firms lost 6.3% on average the week before the ruling. One explanation for this result is that investors, anticipating the arrival of news regarding the case, started to require a premium to hold these stocks until the day of the decision. Third, I find that the drop is not capturing a negative trend in the data. When I consider a month or two months before the ruling—excluding the five trading days before it—I find no out-performance of this group of firms with respect the benchmarks (Table A.6 in the online appendix). Finally, these negative effects do not revert back in the days following the decision. Even if the largest negative returns are experienced in the day in which the news became public, the portfolio continues to experience negative returns for the following month, reaching the bottom in mid-June. Overall, these results hold when using alternative models (Figure A.4 in the online appendix).

Figure 2. (Color online) NPEs Stock Returns Around the Decision

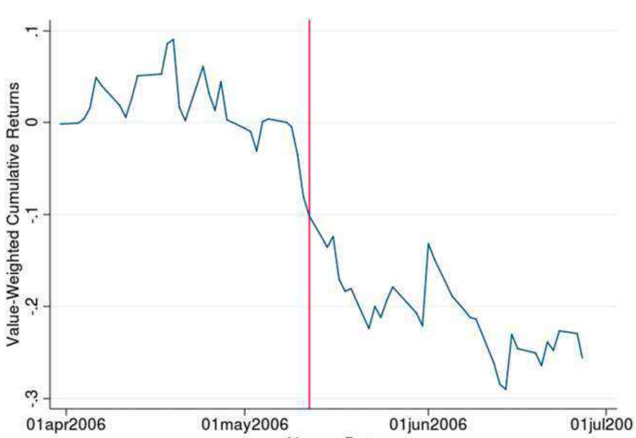

Notes. This figure plots the average cumulative returns for the sample of NPEs identified in the paper. The sample of firms used are 10 companies that are (a) identified as NPEs and (b) public at the time of the Supreme Court decision. The companies are Acacia Technologies, Asure Software (formerly Forgent Network), Rambus, Tessera Technologies, Universal Display, Document Security Systems, ParkerVision, Unwired Planet (formerly Openwave), Interdigital, and Spherix. Information on the sample constructions is provided in Section 2. More information on this analysis is in Online Appendix A.3.3. The straight vertical line corresponds to the trading day right before the decision.

In summary, these facts confirm that public NPEs suffered a great deal around the Supreme Court decision.

In particular, the shock led to a large drop in market value, which did not revert back in the weeks that followed. The results are robust to the removal of each of the NPEs considered in the sample.[21] Overall, this evidence demonstrates that the ruling was a critical event in patent enforcement and greatly affected the players in this market. Furthermore, these results confirm that the decision was not completely anticipated by market participants.

3. Data

To estimate the impact of the *eBay v. MercExchange* Supreme Court decision on corporate innovation, I compare innovative activity across firms that were differentially affected by the decision. In the first part of the paper, I proxy innovation with counts of granted patents, where the timing is defined based on the application date. This measure allows me to observe innovation for a large sample of both public and private companies. The data come from the Fung Institute (University of California Berkeley) patent data set, which is an updated version of the Harvard Business School Patent Network Database (Li et al. 2014) used extensively in literature. These data contain complete information on all patents granted between 1975 and 2014 and contain a new disambiguate assignee ID, which I use to identify a firm across different patents.[22] In most of the analyses, I focus on a sample of more than 16,000 firms active in patenting around the time of decision.

I also supplement the patent data with balance-sheet information from Compustat. I match Compustat to patent information using a procedure that takes advantage of the recent data from Kogan et al. (2012). In short, I link one or more identifiers in the patent data to one Compustat identifier using a patent-level matching. Because patent numbers are easy to match, this approach greatly reduces the probability of errors and missing information. After applying the standard filters,[23] I am left with a sample of more than 1,000 public companies active in innovation around the decision and with R&D information at the quarterly level. Lastly, I match these firms to CRSP using the standard Compustat-CRSP bridge file. In Online Appendix A.3, I provide more details on the data construction and matching.

As stated earlier, the main measure of innovation activity employed in the paper is based on the simple count of patents applied for by a firm in a specific period. I focus on the application date because this is closer to the time of the actual invention. When I focus on public firms, I supplement patent-based innovation measures with R&D intensity data, constructed as quarterly R&D expenses scaled by total assets of the firm (e.g., Mann 2018). In the end, patent data are also used to construct a variety of measures of patent quality, which are discussed in the paper as they are used.

Furthermore, I use patent data to generate firm-level control variables. For every firm, I construct an industry classification based on the major (large) technology class in which the firm patents during the four years around the time of the decision (Hall et al. 2001). I use the addresses reported in the patent application to identify the state of location of the firm. In addition, I construct a proxy for firm age by looking at the time at which a firm first applied for a patent ("start-up"),[24] and a proxy of patent portfolio size by counting the number of patent applications in the two years before the estimation window.

Table 1 reports the summary statistics of the main variables used. On average, the firms in the sample applied for almost 10 (granted) patents per year over the window considered. These numbers are large, but they are justified by the fact that I focus most of the analysis on a subset of firms that are highly active in patenting around the time of the decision. In terms of citations, they receive an average of one citation per patent, where the number of citations is adjusted for technology-class and year. As expected, innovative public firms appear to patent more than the average firm in the full data set—around 50 patents per year—and they have, on average, quarterly R&D expenses of roughly 3% of their assets.

Table 1. Summary Statistics

Variable	Observations	Mean	Standard deviation
Panel A: Full sample			
$\#Patent_{jt}$	32,128	20.28	164.41
$1\{Patent_{jt} = Top^{10\%}\}$	32,128	0.30	0.46
$1\{Patent_{jt} = Top^{25\%}\}$	32,128	0.48	0.50
$Exposure_j$	32,128	0.77	0.79
$Exposure_j^{OLD}$	32,128	0.68	0.56
$AverageCitationPre$	32,128	1.19	1.81
$1\{YearsfirstPatent \leq 3\}$	32,128	0.29	0.46
$SizePrePortfolio$	32,128	18.97	146.74
Panel B: Public firms			
$\#Patent_{jt}$	2,034	101.97	463.87
$R\&D/Asset$	2,034	0.03	0.04
$Exposure_j$	2,034	0.93	0.81
$Exposure_j^{OLD}$	2,034	0.77	0.55
$AverageCitationPre$	2,034	1.50	2.06
$1\{YearsfirstPatent \leq 3\}$	2,034	0.04	0.19
$SizePrePortfolio$	2,034	90.75	375.16

Notes. These two panels report the summary statistics for the two main samples used in the main analyses. Therefore, a period t is defined as a two-year window either before or after the ruling. In panel A, I present the summary statistics for the variables that are used for the first set of analyses, where I employ both private and public firms active in innovation around the time of the decision. In particular, I use the sample that is used in the regressions, which is the sample of firms that applied to at least one granted patent in the two years before and in the year after the time of decision. In panel B, instead, I report summary statistics for the sample that is used in the second part of the analysis, which focuses on public firms that patented around the decision. More information on the sample construction is available in Online Appendix A.3. The variable construction is described in detail in Online Appendix A.3 for outcomes and in Section 4 for the measures of exposure.

4. Empirical Setting

4.1. The Framework

The objective of my study is to examine how the Supreme Court decision *eBay v. MercExchange* affected the innovation of corporations. In principle, every firm patenting in the United State has been affected by this legal change, and, therefore, there is no straightforward control group in this experiment. However, the shock should not have affected every firm in the same way. In particular, firm exposure to patent litigation should represent an important factor in determining whether the ruling was significant for a company. Firms operating in technology areas where patent litigation was nonexistent should be essentially unaffected by the decision. For the same reason, the ruling was, instead, very salient for firms that innovate in high-litigation technologies.

Following this logic, the paper exploits variation in the intensity of the treatment—measured by the extent to which a firm was exposed to patent litigation at the time of the Supreme Court decision—to identify the effects of the decision. In this framework, firms with little or no exposure to litigation, which supposedly were not affected by the shock, provide a counterfactual for firms that were, instead, highly exposed to litigation. The key advantage of this approach is that it does not impose any restriction on the effect of the ruling on firms. In fact, firms more exposed to patent litigation will benefit from the decision because of a reduction in litigation distortions, but will also be hurt because of the potential reduction in deterrence. The estimates will provide evidence regarding the overall net effect of the different channels.

This design is equivalent to a difference-in-difference model, where I study how innovation changed as a function of the exposure to the shock. In other words, I estimate:

$$y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j * Post) + \gamma_t X_j + \epsilon_{jt}, \qquad (1)$$

where y_{jt} is an outcome of firm j at time t, $Post = 1\{time > decision\}$, (α_t, α_j) is a set of firm and time fixed effects, and $Exposure_j$ is the index of exposure to litigation, which is discussed in next section. For robustness, I can augment the specification with a matrix of controls X_j. As I discuss later, the controls are a set of firm-level characteristic measured before the decision, which are interacted with time dummies to allow them to have differential effects before and after the decision (Angrist and Pischke 2008, Gormley and Matsa 2014).

When it is not specified otherwise, I estimate this equation over a four-year window, considering the two years before and after the announcement of the Supreme Court decision on May 15, 2006.[25] To facilitate the interpretation of the different margins and in line with the literature (Bertrand et al. 2004), I run my main results collapsing the data in one observation before and one after the decision. For instance, when looking at patenting, the outcome is the total number of applications in the windows before and after the shock.[26] This choice does not affect in any way my results—indeed, I will also show the results using the full panel dimension when presenting the estimate quarter-by-quarter—and it has the advantage of allowing me to focus on intensive and extensive margin separately. For consistency, the standard errors will be clustered at firm-level across all the analyses.[27]

4.2. Measuring Exposure of Litigation at the Firm-Level

Although most of the previous literature has studied the determinants of litigation at the patent-level (e.g., Lanjouw and Schankerman 2001), a crucial component of my identification relies on measuring exposure to litigation at firm-level. Intuitively, a firm is more exposed to patent litigation if it operates in a technology area where patent litigation is more intense. For instance, companies that operate in software or drugs, where IP lawsuits are more frequent, will be more concerned with patent litigation than companies doing mechanical research, where litigation is much less intense. Therefore, one potential approach is to measure a firm's exposure to litigation by looking at the absolute level of litigation in the area in which it undertakes R&D. This approach takes advantage of two features of the patent system. First, there is a lot of variation across technology fields in the intensity of patent litigation. This is true both across major technology areas—for instance, between "Communications & Computer" and "Chemicals"—and within the major technology fields. Second, many companies operate across different technology fields (Bessen and Hunt 2007). This fact implies that even firms that operate in relatively safer areas may still be influenced by litigation because of a subset of the patent portfolio.

Formalizing this intuition, I can express the exposure to patent litigation of an individual firm j as a function of two quantities: (1) the technology fields i in which the firm j operates, defined by a vector $t(j) = [\sigma_i^j]_{i=1}^T$; and (2) the distribution of the patent litigation activity across different technology fields i, which is defined by a vector $p = [p_i]_{i=1}^T$. In particular, I can define $t(j)$ as a vector whose entries σ_i^j are the share of firm j patents across the different technology fields i. Therefore, firm j exposure to litigation $Exposure_j$ can be constructed by weighting the litigation risk in each technology field by the share of activity that firm j has in each of these fields. This is:

$$Exposure_j = \sum_{i=1}^{T} \sigma_i^j p_i, \tag{2}$$

with $Exposure_j \in [\min(p), \max(p)]$.

Although the variable $Exposure_j$ is intrinsically unobservable, its components—$t(j)$ and p—can be constructed from the data. First, I use patent data to measure $t(j)$, the technology space where the company operates. The USPTO categorizes each patent across more than 400 technology classes, therefore providing a very precise and narrow definition of technology. Using this classification, I construct σ_i^j as the share of granted patents of firm j in technology class i that were applied for before 2006.[28]

Second, I estimate the distribution of patent litigation across technology fields—the vector p—using litigation data from WestLaw, a subsidiary of Thomson Reuters (e.g., Lanjouw and Schankerman 2001 and Lerner 2010).[29] Using all filings involving IP between 1980 and 2006, I extract all the patents that were asserted by the plaintiff and then use this information to construct a proxy for p. After some preliminary data cleaning, this corresponds to more than 30,000 fillings, with the number of cases per year increasing over time (Figure 1). Similar to Kiebzak et al. (2016), I reshape the data at defendant–plaintiff–patent level to make cases comparable across filings.[30] I then measure the size of the litigation in each of the USPTO technology classes by computing the number of patents in a specific class involved in the litigation, normalized by the total number of patents litigated. In other words, my index is:

$$p_i = \frac{100\sum_{c\in cases}\#Patents_c^i}{\sum_{i\in Tech.Classse}100\sum_{c\in cases}\#Patents_c^i}, \tag{3}$$

where i defines one of the USPTO technology classes, and c is a specific filing. In line with the previous evidence, patent litigation is not equally spread across technology classes, but, rather, tends to be more concentrated. For instance, the top 50 technology classes in terms of litigation account for half of the patent level litigation (Table A.7 in the online appendix).[31]

I estimate $Exposure_j$ by combining these two measures as in Equation (2). My preferred measure uses litigation data and patents in the five years before the Supreme Court decision. As expected, the distribution of the score is skewed, and some areas, such as "Drug" and "Computer and Communication," have a larger share of highly exposed firms (Figure 3). However, even within this major industry, there is a relatively large variation in litigation exposure. Indeed, I will show that my results are similar, even when I control for this main industry effect.

This way of measuring exposure to patent litigation has three important advantages. First, this score can be calculated for every firm that is active in patenting using existing data, and its computation is relatively simple. Second, the measure is exogenous to firm *j*'s strategies in litigation. Unlike other approaches, this measure does not depend on the actions that firms take regarding litigation, but only on the area in which a firm operates. Third, the fact that litigation across technologies is highly persistent over time (Figure A.5 in the online appendix) suggests that this measure does not simply reflect some heterogeneity in technology shocks in the years before the Supreme Court decision, but, rather, some structural characteristic of the field.

However, the validity of my approach relies on the assumption that the score captures relevant heterogeneity in exposure to litigation across firms. In general, the fact that I will find that exposure predicts a change in R&D behavior around the shock already suggests that this approach is capturing information that is relevant to understand the impact of the decision.[32] However, to provide more direct evidence in this direction, I examine the stock market reaction of innovative firms around the announcement. Given the previous NPE result (Section 2.3), it is clear that the stock market thought that the patent ruling represented significant news for innovative firms. Therefore, we should expect to find a systematic correlation between this measure and stock returns of innovative firms around the event.

I explore this issue by measuring returns and abnormal returns around the Supreme Court announcement and then correlating these measures with the litigation exposure score (Katz et al. 2017). The main result of the analysis can be synthesized by Figure 4, which plots the cumulative value-weighted returns of high- and low-exposure firms, where the split is made at the top 25% of the litigation distribution. I find that the two groups behave in the same way in the days before the ruling. However, the day of the decision, the high-risk group outperforms the low-risk group by almost 1%. This initial outperformance does not revert immediately afterward. The same results hold in a formal regression analysis (Table A.1 in the online appendix). Online Appendix A.4.1 contains a more detailed discussion of these results. These analyses suggest that our score captures significant exposure to the ruling. Furthermore, the effect appears to be positive for innovative companies.[33] As I discuss later, this effect is consistent with the rest of evidence presented in the paper.

5. The Effect of the Supreme Court Decision on Innovation

This section contains the main results of the analysis. I start by showing that the Supreme Court decision positively affected the ability of companies to patent new technologies. Next, I discuss the main identification assumption—in particular, the parallel-trend assumption—and I provide further evidence that confirms the quality of my model. Lastly, I examine

Figure 3. (Color online) Technology Distribution for the Top 50% Riskier Firms vs. Whole Sample

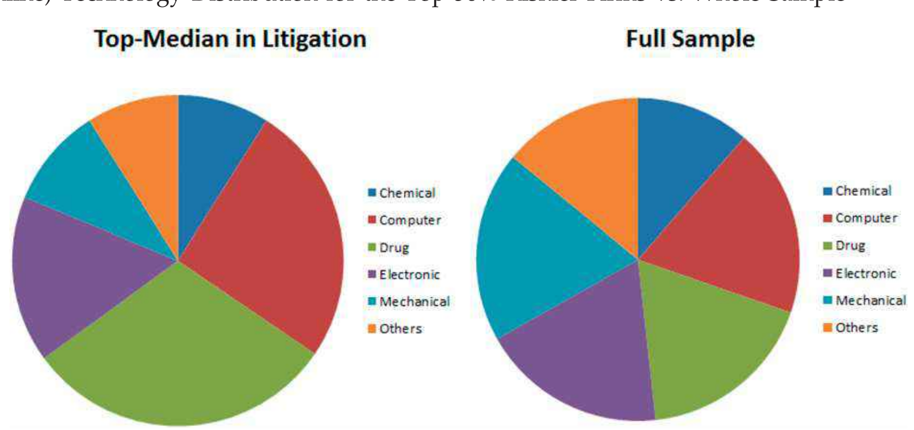

Notes. This figure reports the pie chart of the patents by industry, across the full sample and the sample of firms that are more exposed to litigation. Industries are identified based on patent applications across macro-technology areas (Hall et al. 2001), and the construction is discussed in detail in Online Appendix A.3. The first chart is constructed using only firms in the top 50% of litigation exposure, where litigation is measured using $Exposure_j$. This is constructed by using litigation in the five years before the decision and using patents since 2000. The second chart is instead constructed by using the full sample. Furthermore, the sample that was used to construct this plot is the sample of innovative firms that applied for at least one patent in the two years before or two years after the decision.

Figure 4. (Color online) Stock Market Reaction: High vs. Low Exposure

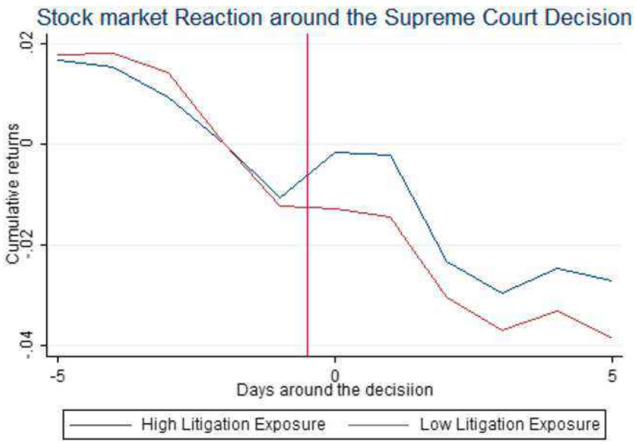

Notes. This figure plots the value-weighted cumulative returns across high- and low-exposure firms. High-litigation firms are firms in the top 25% of the litigation distribution. Cumulative returns are normalized to zero for both groups two days before the decision. The straight vertical line is plotted between the day before and the day of the decision (which is defined to be zero in calendar time). The value-weights are based on the market value of traded stocks, and they are kept fixed five days before the decision.

the effect of the decision on the quality of innovation and on R&D intensity for public firms.

5.1. The Effect on Innovation Output

I begin my analysis by exploring how the decision affected innovation output, measured by the count of granted patents using the application date as time reference of the patent. In particular, I construct two outcomes using this data. First, I look at $\ln(pat_{jt})$, which is the logarithm of the total number of patents that firm j applied for during time t (intensive margin). In order to keep the panel balanced and, therefore, estimate a purely intensive margin, I estimate the model using every firm in the patent data that applied for at least one patent before and after the shock.[34] Second, I examine an alternative outcome variable: a dummy equal to one when the firm has applied for any subsequently granted patent in the period, $1\{Patent_{jt} > 0\}$ (extensive margin). In this case, the sample contains every firm that applied for at least one patent in the five years before the Supreme Court decision, which is a minimal requirement to construct the measure of litigation exposure.

Table 2 starts presenting the results by estimating the baseline version of Equation (1). Looking at both the intensive (column (1)) and extensive (column (4)) margins, I find that firms that were operating in more litigious areas increased their patent application relatively more. These effects are not only statistically significant, but also economically relevant. A one-standard-deviation increase in the exposure to litigation leads to a relative increase in patent applications of 3%. Comparing these estimates to the patenting baseline, this effect corresponds to an increase of almost one additional patent for innovative firms (0.7). Similarly, a one-standard-deviation increase also implies a 0.8% increase in the probability of patenting, which is a 2% increase relative to the baseline probability over the whole period. This result suggests that removing the threat of automatic injunction did not discourage firms from filing patent applications. If anything, those firms that were more likely to be affected by the new rules saw a relative increase in patenting activity. In particular, this evidence seems to be consistent with a "catch-up" mechanism, where firms were able to close part of the gap caused by litigation costs after the change in enforcement.

The causal interpretation of the difference-in-difference approach relies on the parallel-trend assumption. In this case, the assumption requires that the relative behavior of high- and low-exposed firms would have not changed without the Supreme Court ruling. As a first test in this direction, I examine the dynamic of patenting activity in the months before and after the decision. In particular, I use patent data at quarterly frequency, and I estimate the time-varying effect of exposure to litigation on patenting relative to the last quarter before the decision:

$$y_{jt} = \alpha_j + \alpha_t + \sum_{\tau=-8}^{8} \beta_{T-\tau} Exposure_j + \epsilon_{jt}. \tag{4}$$

Consistently with the parallel-trend assumption, I expect to find that: (a) The positive effect appears only in quarters after the Supreme Court decision ($\beta_t > 0$); and (b) before the decision, the changes in patenting behavior are orthogonal to the measure of exposure ($\beta_t = 0$). For completeness, I estimate this equation using a log-plus-one specification, which allows me to look at the effect at both the intensive and extensive margins. These results are presented in Figure 5: Firms characterized by different exposure to litigation did not have a differential pattern of patenting before the Supreme Court decision. The estimated β in this period is always small and statistically non-different from zero. However, after the Supreme Court decision, firms that were more exposed to litigation increased their rate of patenting more. In particular, the effects turn positive already within a few quarters and keep rising afterward. These results are overall comparable with those in the previous regression model, confirming that choice of collapsing the data in the main model does not affect our results.[35] As I discuss later, I confirm the same results looking at other metrics, like innovation quality and R&D investments (Figure A.6 and Table A.5 in the online appendix).

Table 2. Effect of the Policy Change on Patenting: Main Results

	(1)	(2)	(3)	(4)	(5)	(6)
OLS		$ln(Patents_{jt})$			$1\{Patent_{jt} > 0\}$	
$Post \cdot Exposure_j$	0.040***	0.036***	0.034***	0.010***	0.027***	0.027***
	(0.008)	(0.011)	(0.011)	(0.002)	(0.002)	(0.002)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	N	N	Y	N	N
Indu. × *Time F.E.*		Y	Y		Y	Y
$Controls_j$ × *Time F.E.*			Y			Y
R^2	0.005	0.007	0.033	0.216	0.282	0.290
Observations	32,128	32,128	32,128	155,876	155,876	155,876

Notes. This table reports the estimate of the linear difference-in-difference specification (Equation (1)), where I estimate the effect of the decision on quantity of innovation. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j{\cdot}Post) + \gamma X_{jt} + {}_{jt}$, where y_{jt} is (a) the (natural) logarithm of granted patent that firm j applied for during period t for columns (1)–(3); and (b) a dummy equal to one if the firm j applied to at least one patent in period t. The data set is a balanced two-period panel. Each period collapses firm information in the two years before and two years after the Supreme Court decision. The sample depends on the outcome: When looking at the intensive margin (columns (1)–(3)), I use every firm that applied to at least one patent in the two years before and in the year after the decision; when I look at the extensive margin (columns (4)–(6)), I use the sample of every firm with at least one patent in the five year before the decision, which is the minimal requirement to construct the measure of exposure. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. In columns (1) and (4), I control for firm fixed effects and time effects. In columns (2) and (5), I add industry-time fixed effects to the equation. Industries are constructed based on the macro technology area where the company patented the most over the four years before the decision, where macro technology classes are constructed as in Hall et al. (2001). In columns (3) and (6), I further augment the specification using location dummies of the firm (constructed by using the model location reported in patent data), the size of the portfolio before the estimation period, the start-up status (looking at whether a firm applied for the first patent ever within the previous three years), and average quality of the patent portfolio in the pre period, measured by average citations. All these controls are interacted with a time dummy to allow the variable to have a differential effect before and after. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.
$^*p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The analysis of the pretrend provides a first glimpse into the timing of the effect, which is discussed more extensively in Online Appendix A.4.2. Overall, there are three results to highlight. First, the effect is increasing over time, consistent with the idea that incentive changes will affect output with a lag. Relative to one year after the decision, the effect over two years increases by 38% and over three years by 50% (Table A.3 in the online appendix). Second, although, on average, I find some response within one year, this short-term response is driven by companies operating in "Computer and Communications" technologies (Hall et al. 2001) (Table A.4 in the online appendix). This result is reassuring because these technologies tend to have a much faster R&D cycle than other areas—for instance, drugs—and, therefore, they should respond first. Third, the effect does not appear to be short-lived. In Figure A.1 in the online appendix, I replicate the same dynamic plot using data on patents up to 2010, showing that the effects discussed before are still economically and statistically significant essentially five years after the decision.

Figure 5. (Color online) Effect of Litigation over Time

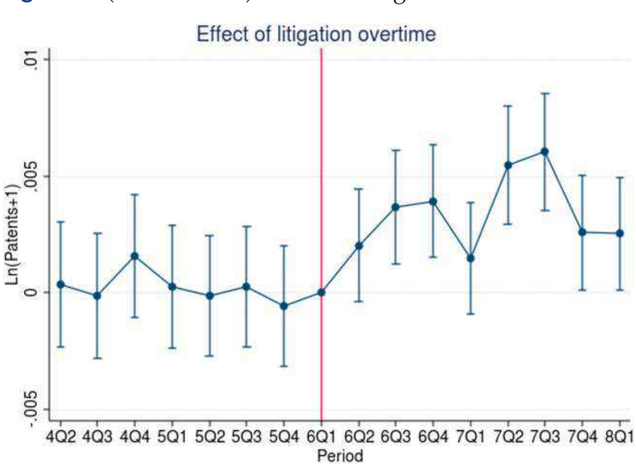

Notes. This figure plots the β_t from Equation (4). The straight vertical line corresponds to the last period of the predecision period. Every β_t is plotted with the corresponding CI at 5%. Every period is label with the corresponding quarter. Notice that quarters are in "event time" not calendar time: In fact, I set the end of the first quarter artificially to be the one ending in May 15th (the other quarters are constructed relative to this). The data used correspond at the two years before and after the decision, in event time. The sample used corresponds to the one of the extensive margin. Standard errors are clustered at the firm-level.

Another way to provide evidence consistent with the parallel-trend assumption is by using placebo tests, where I replicate my analysis in periods where there is no change in the rules.[36] The idea is, again, to show that, absent a large shock to enforcement like eBay, exposure to litigation does not predict differential changes in patenting. In order to avoid arbitrarily choosing a period in which to run the placebo, I consider as the fictional shock period every quarter in the closest two years before the shock and such that the postperiod does not overlap with the posttreatment period (2002Q2–2004Q1).[37] Figure 6 presents

Figure 6. (Color online) Placebo Tests over Time

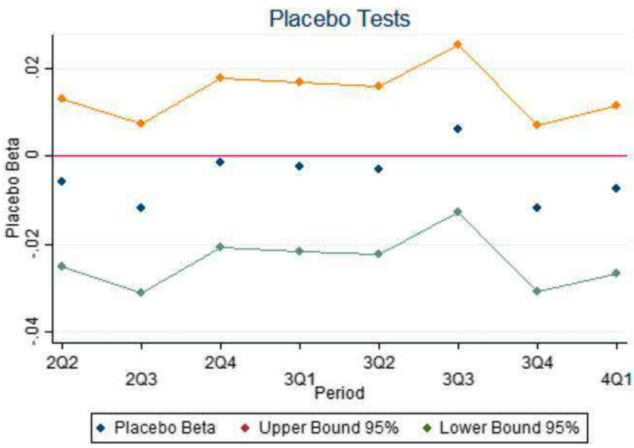

Notes. This figure presents the results from a set of placebo tests. In particular, I construct a series of placebo samples, centered around fictional shocks in the two years that are completely outside the period after the decision. The date in the *x* axis are the quarter around which the analysis is centered. In each case, I reconstruct the data around this placebo shock, both the outcomes and the measures of exposure $Exposure_j$. Then, I run the standard regression. The figure plots the β from Equation (1), as well as the 95% confidence intervals, estimated over different samples. For clarity, I estimate the simple equation without further controls and looking at the intensive margin. Notice that quarters are in "event time," not calendar time: In fact, I set the end of the first quarter artificially to be ending on May 15 (the other quarters are constructed accordingly). Data used correspond to the two years before and after the decision, in event time. Standard errors are clustered at the firm level.

the results of this analysis by plotting the β from the intensive margin regression and its 95% confidence interval for each quarter considered as the fictional shock. As expected, the coefficient is never positive and significant. In other words, in periods where there is no major shift in the patent-enforcement law, I do not find that firms operating in high-litigation fields increase innovation more than firms in low-litigation fields. If anything, the coefficient actually tends to be negative in sign, but the size is always small and never statistically different from zero. As I discuss later, the same result holds for R&D spending.

This detailed discussion of pretrend and placebo tests provides strong evidence in favor of the parallel-trend assumption. However, these tests cannot fully exclude the presence of a shock contemporaneous to the ruling that was correlated positively (negatively) with exposure to litigation and positively (negatively) affects innovation. In this context, the main concern is the presence of technology shocks in highly litigated areas. To exclude this possibility, I augment my model with industries by time fixed effects, where industries are defined based on the main technology category within which they fall, as previously discussed (Hall et al. 2001). This set of controls removes from the data any technology trend, comparing patenting by firms with different levels of exposure to litigation within the same industry. The results—reported in columns (2) and (5) of Table (2)—show a relatively small change with respect to our main findings. In particular, at the intensive margin, the change in the estimated coefficient is minimal, and this difference is not statistically different from zero.[38] At the extensive margin, the inclusion of this control instead significantly increases the magnitude of the effects. Therefore, although industry dynamics could be important in explaining patenting behavior around this period, they do not seem to drive my results.

However, the presence of contemporaneous shocks may also materialize in other aspects of economic activity. For instance, local economic factors may be a concern to the extent that there is geographical clustering of firms that are similarly exposed to litigation. To deal with these concerns, in columns (3) and (6) of Table (2), I augment the previous specification with an extra set of firm-level controls in every case interacted with time dummies. To directly tackle the concern of geographical clustering, I consider among my controls the location of the firm's R&D facilities—based on state of operation where I find the most patents before the decision. Furthermore, I also add controls for the size of the portfolio of the firm, measured by the number of patents published in the years before the decision, but outside the estimation window;[39] quality of portfolio, measured by the average number of citations before the decision; and a dummy for firms that patented for the first time in the three years before the decision. In general, the addition of these extra controls does not significantly change my estimates. In Figure A.2 in the online appendix, I also show that the pretrend analysis is qualitatively identical when I also add these controls interacted with postdummy.[40]

In terms of contemporaneous shocks, in Online Appendix A.2, I also argue that our results are not explained by the presence of other legal changes around the same time as eBay. In general, the federal law was stable during the period analyzed.[41] Furthermore, the other Supreme Court decisions involving innovation are unlikely to have an effect in our setting. In fact, the scope and importance of other rulings were generally more limited than eBay. On top of this, the timing of our effects are also inconsistent with the impact of other important Supreme Court decisions. Online Appendix A.2 contains a thorough discussions on these issues. As a final step, I show that the results are also robust to three extra tests. First, I implement a simple permutation test (Fisher 1922, Chetty et al. 2009). As discussed in Online Appendix A.4.3, this test allows me to provide inference based on weaker assumptions than the standard linear model and to rule out that my identification strategy is somehow mechanically capturing other spurious firm characteristics. The results are reassuring, as I find that the p-value constructed based on the random permutation test is similar to the standard one and lower than 1% (Figure A.3 in the online appendix). Second, these results are identical when using a fixed-effects Poisson model instead of the linear specification (Table A.9 in the online appendix). Third, in Table A.8 in the online appendix, I can replicate the results estimated by Equation (1) using an alternative measure of patent litigation exposure $Exposure_j^{LARGE}$. As discussed before, this measure uses data on patents applied for by the firm in the 10 years before the shock and patent-litigation data since 1980.

5.2. Evidence on Patent Quality

The previous results show that firms that were more exposed to litigation appeared to have responded to the Supreme Court decision by increasing the rate of patenting. At face value, these results seem to support the idea that removing automatic injunction did not cause a collapse in innovation activity, but, instead, had a positive effect. However, more work is required to interpret these results as evidence on firms' innovation. First, an increase in patenting may simply represent an increase in the propensity to patent, rather than a real increase in innovation. Second, patent applications may increase because firms feel the need to increase their defensive patents (Hall and Ziedonis 2001). This result could be consistent with the idea that eBay reduced protection against real violations.

As a first step, I start exploring measures of innovation quality. In order to do so, I use the same empirical model as before, but focus on a set of quality metrics that are constructed based on patent citations.[42] Previous research has shown that forward citations are correlated with the quality of the underlying patent and its economic value (Kortum and Lerner 2000, Hall et al. 2005). Here, I construct different measures based on citations in order to capture different aspects of quality (Online Appendix A.3).

First, I examine how the average quality of the patents—measured by average scaled citations—changes around the time of the decision. Because comparing the number of citations across technologies and over time can be challenging (Lerner and Seru 2017), I adjust my baseline citations by scaling them by the average number of citations received by other assigned patents in the same technology class and year.[43] Across the three specifications, I find no change on the average patent in a firm's portfolio (columns (1)–(3), Table 3). In general, the coefficient is positive, but never statistically different from zero. This result confirms that the new marginal patents applied for after the decision were not of worse quality than those before.[44] Therefore, this evidence rejects the

Table 3. Evidence on Patent Quality

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
OLS	*AverageScaledCitations*			$1\{Patent_{jt} = Top^{10\%}\}$			$1\{Patent_{jt} = Top^{25\%}\}$		
$Post \cdot Exposure_j$	0.013	0.014	0.016	0.010**	0.016***	0.016***	0.018***	0.022***	0.021***
	(0.023)	(0.032)	(0.032)	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
Firm F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time F.E.	Y	N	N	Y	N	N	Y	N	N
Indu. × *Time F.E.*		Y	Y		Y	Y		Y	Y
$Controls_j$ × *Time F.E.*			Y			Y			Y
R^2	0.001	0.001	0.004	0.001	0.001	0.005	0.001	0.001	0.004
Observations	32,128	32,128	32,128	32,128	32,128	32,128	32,128	32,128	32,128

Notes. Tables 3 and 4 report the estimate of the linear difference-in-difference specification (Equation (1)), where I estimate the effect of the decision on the quality of innovation. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + {}_{jt}$, where y_{jt} is a measure of patent quality. In particular, in this table, I consider three outcomes: (a) the average number of scaled citations received by firms j in period t; (b) a dummy that is equal to one if firm j has published in period t at least one patent that is in the top 10% of the distribution of citations among patents granted in the same year in the same technology class; and (c) similar dummy, but constructed considering the top 25% of the distribution. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. I always control for firm fixed effects and time effects. Furthermore, I always augment the specification with industry-time fixed effect, which are constructed based on the macro technology area where the company patented the most over the four years before the decision (Hall et al. 2001). Lastly, I further augment every specification with location dummies of the firm, the size of the portfolio before the estimation period, and the start-up status. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

Q:69

Table 4. Evidence on Patent Quality

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
OLS		*ShareDefensivePatent*			*ShareBM Pat.*			*ShareBM Pat.Broad*	
Post · *Exposure*$_j$	–0.008**	–0.011**	–0.010**	–0.004***	–0.006***	–0.006***	–0.002	–0.004*	–0.004*
	(0.004)	(0.005)	(0.005)	(0.001)	(0.001)	(0.001)	–(0.001)	(0.002)	(0.002)
Firm F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time F.E.	Y	N	N	Y	N	N	Y	N	N
Indu. × *Time F.E.*		Y	Y		Y	Y		Y	Y
Controls$_j$ × *Time F.E.*			Y			Y			Y
R^2	0.001	0.001	0.005	0.002	0.003	0.007	0.001	0.003	0.006
Observations	32,128	32,128	32,128	32,128	32,128	32,128	32,128	32,128	32,128

Notes. Tables 3 and 4 report the estimate of the linear difference-in-difference specification (Equation (1)), where I estimate the effect of the decision on the quality of innovation. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j{\cdot}Post) + \gamma X_{jt} + {}_{jt}$, where y_{jt} is a measure of patent quality. In this table, I look at other measures of defensive patenting. In particular, I have three outcomes: (a) the share of defensive patents, where the defensive patents are patents are in the top 25% in terms of dispersion across technology (originality) among patents of same technology class and year, despite being in the bottom three quartiles in terms of citations for the same group; (b) the share of business method (BM) patents, which are defined as patents in technology class 705; and (c) alternative share of BM patents with broader definition, where BM patents are defined in Hall (2003, table 3) (i.e., technology classes: 84; 119; 379; 434; 472; 380; 382; 395; 700; 701; 702; 703; 704; 705; 706; 707; 709; 710; 711; 712; 713; 714; 715; 717; and 902). As before, the data set is a balanced two-period panel where I employ every firm that published at least one patent in the two year before and in the year after the decision. The variable *Exposure*$_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. I always control for firm fixed effects and time effects. Furthermore, I always augment the specification with industry-time fixed effect, which are constructed based on the macro technology area where the company patented the most over the four years before the decision (Hall et al. 2001). Lastly, I further augment every specification with location dummies of the firm, the size of the portfolio before the estimation period, and the start-up status. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

hypothesis that an increase in the innovation output was driven by a surge in low-quality patents.

Second, because the returns on innovation are highly skewed (Pakes 1986), it may be useful to also examine whether the decision increased companies' ability to develop breakthrough innovation (Kerr 2010). In principle, we may expect a positive effect on this dimension because of a scale effect: Firms now may have more resources to do R&D—consistent with previous findings—and, therefore, they are more likely to engage in at least one high-quality project. This mechanism may be magnified if firms conducting breakthrough projects are even more likely to be targeted by litigation. For instance, high-quality research may require combining knowledge across different areas, and this type of activity may expose firms to a higher change of infringement. In order to look at this, I examine the probability that a company applies for a patent that is at the top 10% (or top 25%) of the citation distribution in the relevant reference group. In line with previous literature, the reference group is composed of assigned patents that are in the same USPTO technology class and were developed in the same year (e.g., Lin et al. 2020), but I also show similar results with alternative benchmarks.

As reported in Table 3, I find that, around the time of the ruling, firms more exposed to litigation were relatively more likely to apply for a breakthrough patent. The result holds when looking at both the top 10% and 25% of the quality distribution after the decision. In economic terms, an increase by one standard deviation in the index led to a 1% increase in the probably of applying for a patent in the top 10% of the distribution, which represents more than a 3% jump from the baseline probability. The results are qualitatively similar across the various specifications. In Table A.11 in the online appendix, I find that the same results hold when I construct two alternative versions of the reference groups used to compute the quality threshold. In particular, rather than benchmarking citations within both technology class and year, I also construct a version of the data where top citation patents are identified only looking at the technology class (odd columns) and the year (even columns). The results using these alternative measures are still positive and significant at the conventional level.

Overall, firms more exposed to patent litigation did not lower the average quality of their patents, and they were more likely to develop breakthrough technologies. Furthermore, these findings—both average and extreme outcomes—do not appear to be driven by any pretrend before the decision (Table A.5 and Figure A.6 in the online appendix). Overall, this evidence appears more in line with the hypothesis that the increase in patenting was caused by more innovative activity.

5.3. Evidence on Strategic Patenting

As the next step in the analysis, I explore the incidence of strategic patents around the time of the decision. I define as strategic those patents whose value does not rest on the intrinsic technology covered by the IP, but, instead, comes from the ability to use it for litigation purposes, either offensive or defensive. If the positive effect on patenting is explained by an increase in defensive activity, then we should expect to find this increase to be explained by strategic patenting. In this section, using two alternative measures of strategic patenting, I will show that this is not the case.

To start measuring strategic patents, I count the number of patents that are low quality—measured by forward citations—but whose patent claim spans a very large set of different technologies, measured by originality (Hall et al. 2001). The intuition behind this measure is simple: The value of a defensive patent does not rest on the quality of the innovation covered by the patent, but, rather, on its ability to be used in court. Consistent with this argument, Abrams et al. (2013) find that patents with a high strategic value are actually characterized by lower quality, measured by forward citations. Instead, patents are more valuable for court cases when they are characterized by high originality (Hall et al. 2001), which is a measure used in the literature to identify patent claims that span a large set of different technologies.[45] In practice, my outcome is the share of granted patent applications that are in the top 25% in terms of originality among patents in the same technology class and year, but also are in the bottom three quartiles in terms of citations for the same group.[46] In these tables, I refer to this variable as the share of defensive patents.

The results are reported in Table 4. I find that firms more exposed to litigation actually experienced a reduction in the share of defensive patents around the decision time (columns (1), (2), and (3)). The results are consistent in size and significance across the different specifications, but they are larger when I add controls. Looking at the most saturated model (column (3)), the estimates show that a one-standard-deviation increase in exposure to litigation translates to a reduction in the percentage of defensive patents by roughly 1%, which corresponds to a 5% reduction in defensive patents relative to the average for the period.[47]

To validate the previous results, I provide an alternative definition of strategic patents based on business-methods patents. With this approach, rather than trying to identify strategic or defensive patents across the whole sample, I focus on a specific set of technologies—business-method patents—where the strategic value of patents is generally considered to be one of the key determinants of firms' patenting behavior (Srinivasan 2018).[48] In line with the previous analyses, in columns (4), (5), and (6) of Table 4, I examine how the share of business-method patents changed around the decision. Across all the specifications, I find that companies that were more exposed to litigation applied for a lower share of business-method patents after the decision time. In these results, business-method patents are defined as simply patents in the official business-method technology class—technology class 705. For further robustness, in the last three columns, I consider an alternative definition of business-method patents, as discussed in Hall (2003).[49] Also in this case, I find a negative relationship between firm exposure and the change in the share of business-method patents. However, the effects appear to be statistically weaker, particularly in the baseline specification.

Overall, this evidence on strategic patents appears at odds with the hypothesis that the change in enforcement rules led to more patenting because it increased firms' defensive activity. In fact, firms that were more exposed to litigation appear to lower their efforts on strategic patents, suggesting that they perceived a decline in the need for building up a defensive portfolio. Together with the results on the quality of the innovation output, these analyses support the interpretation of the patenting results as an increase in innovation. Furthermore, this result also helps in squaring this paper with the previous evidence on the relationship between patenting and injunction risk from Hall and Ziedonis (2001), which documented that an increase in injunction risk in the semiconductor industry led to an increase in defensive patenting. In the same way, I show that a decrease in injunction risk overall led to a decline in the intensity of defensive patenting.[50]

5.4. R&D Investment for Public Firms

As a further step to nail down the effects of the decision, I next look at the effects of the ruling on R&D investment. This aspect is important for two reasons. First, looking at R&D investment provides more insight on how the ruling affected firms' activity in innovation. Second, consistent with the discussion in the previous section, evidence on R&D investment would help to interpret the patenting results. In fact, if the decision really increased innovation activity, we would expect a response at both the input and output ends of innovation. The same would not hold if the motives behind the increase in patenting were simply an increase in the propensity to patent or enhance defensive activity. One important constraint in this analysis is data. Although previous analysis has the advantage of focusing on a very large, heterogeneous set of firms, the amount of information that is available is limited to patent data. Looking at public firms, I can instead observe the total amount of monetary resources that a company has devoted to R&D. In particular, I focus on a set of around 1,000 firms active in patenting around the time of the decision (Section 3).

Using this sample, I estimate the same Equation (1) looking at R&D investment. As a preliminary check, in columns (1), (2), and (3) of Table 5, I find that patent counts positively respond to the ruling for this group of firms as well. In magnitude, these effects are larger than those estimated using the full sample, in principle, suggesting that public firms were particularly affected by the decision.[51] This result would be consistent with the idea that public firms are perceived as very profitable targets for strategic litigation.[52] Then, in columns (4), (5), and (6), I show that the same result holds for R&D investment. Specifically, firms that were more exposed to litigation at the time of the Supreme Court decision experienced a larger increase in R&D spending around the same period. A one-standard-deviation increase in litigation exposure leads to an 8% increase in R&D intensity relative to the baseline model. These results are essentially unchanged when I add the usual controls (columns (5) and (6)).

Also in this case, the results do not appear to be driven by a failure of the parallel-trend assumption. This is true both when looking at the nonparametric test—where I plot the quarter specific coefficient across time (Figure 7)—and when I assume linear trends in the model Table A.13 in the online appendix. Similar to our previous results, I find that the differential among firms appears within one year from the decision, and it does not seem to close in the following quarters. If anything, the gap seems to increase over time. Lastly, the same type of placebo analysis discussed before also works for R&D, as reported in Figure A.7 in the online appendix. Overall, this result confirms that the decision had a significant impact on the R&D decisions of innovative firms.

5.5. Plaintiff vs. Defendant: Heterogeneous Effect of the Ruling

As a final robustness test to the main mechanism, I examine how the firms' reactions differ depending on whether a company was more or less likely to be active as a plaintiff in litigation around the time of the ruling. This test is motivated by the simple intuition that a company that expects to be more active as a plaintiff rather than a defendant should benefit far less from the enforcement changes. In fact, even if permanent injunction were still available to plaintiffs,

Table 5. Effect of the Decision on Public Firms

	(1)	(2)	(3)	(4)	(5)	(6)
OLS		$ln(Patents_{jt})$			$R\&D_{jt}/Asset_{jt}$	
$Post \cdot Exposure_j$	0.062*	0.092**	0.096**	0.002**	0.004***	0.004***
	(0.033)	(0.044)	(0.045)	(0.001)	(0.001)	(0.001)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	N	N	Y	N	N
Indu. × *Time F.E.*		Y	Y		Y	Y
$Controls_j$ × *Time*			Y			Y
R^2	0.007	0.017	0.081	0.010	0.018	0.063
Observations	2,034	2,034	2,034	2,034	2,034	2,034

Notes. This table reports the estimate of the linear difference-in-difference specification (Equation (1)), where I estimate the effect of the decision on patenting and R&D intensity. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j{\cdot}Post) + \gamma X_{jt} + {}_{jt}$, where y_{jt} is (a) the (natural) logarithm of granted patents that firm j applied during period t for columns (1)–(3); and (b) *R&D/Asset* is the average over the period of the quarterly R&D expenses scaled by total assets for columns (4)–(6). Outcomes are winsorized at 1%, and the exact construction of the variables is discussed in the paper and in Online Appendix A.3. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. The data set is a balanced two-period panel, where each period collapses firm information in the two years before and two years after the Supreme Court decision. The sample is a set of nonfinancial, U.S.-located public firms that applied to at least one patent in the two years before and one after (see Online Appendix A.3). I always control for firm fixed effects and time effects. In columns (2) and (5), I augment this with industry-time fixed effect. Industries are constructed based on the macro technology area where the company patented the most over the four years before the decision, where macro technology classes are constructed as in Hall et al. (2001). In columns (3) and (6), I further augment the specification using location dummies of the firm (constructed using the model location reported in patent data), the size of the portfolio before the estimation period, quality of the patent portfolio before the decision (measured by average citations), and the "start-up" status (looking at whether a firm applied for the first patent ever within the previous three years), which would be more correct to refer as firm age in this sample. All these controls are interacted with a time dummy to allow the variable to have a differential effect before and after. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.

*$p < 0.10$; **$p < 0.05$; ***$p < 0.01$.

its strength as a bargaining tool was weakened by the decision. This analysis is, however, constrained by a few theoretical and empirical issues. First, the data do not allow us to observe whether a firm will be more likely to be involved as a plaintiff or a defendant in the future: All we can observe is whether a firm was involved in a lawsuit in the past. Second, the decision to enter into a formal lawsuit is clearly endogenous (Cohen et al. 2019). As pointed out above, the number of lawsuits that end up in court are just a fraction of the overall litigation activity, and the decision to go to court rather than settle is clearly a function of the expected costs and benefits of the different actions. Third, data do not allow us to understand the real motives of the dispute (i.e., strategic versus defensive lawsuits).

With these caveats in mind, I use the lawsuits data at the firm level to explore this dimension. Starting with the lawsuits files from Westlaw discussed earlier, I obtain the list of all the firms that were involved in litigation—either as a defendant or a plaintiff—and I name-match them to the patent data with the help of a research assistant. Information on this matching is available in Online Appendix A.3.4.[53] Using this information for the period 2001–2005, I define a firm as being more likely to be the plaintiff if it appeared more times in the filings as a plaintiff than as defendant.

In Table 6, I start exploring these analyses for R&D intensity. First, I look at the effect of exposure to litigation separately for companies that I identify as more likely to be a plaintiff (column (2)) versus the rest, which I define by complementarity as the group of firms more likely to be a defendant (column (1)).

Figure 7. (Color online) Effect of Litigation on R&D Intensity over Time

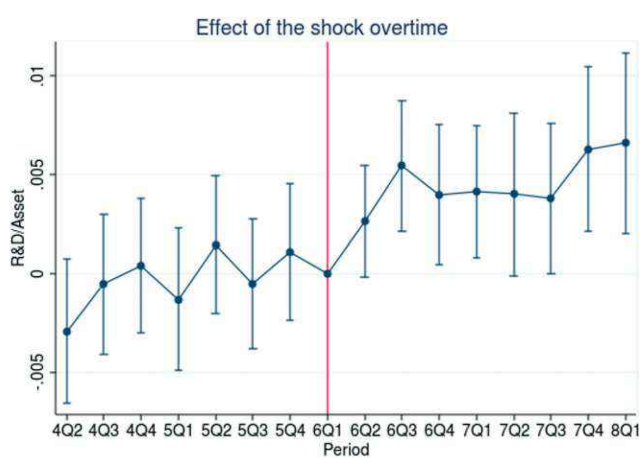

Notes. This figure plots β_t from Equation (4) with the standard controls, where the outcome is R&D over asset. The straight vertical line corresponds to the last period of the predecision period. Every β_t is plotted with the corresponding CI at 5%. Every period is label with the corresponding quarter. Notice that quarters are in "event time," not calendar time: I set the second quarter to be the first quarter ending after the Supreme Court decision (and the others are defined relative to this quarter). The data used correspond to the two years before and after the decision. The sample is the standard Compustat sample of innovative firms used in the paper.

Table 6. Effect Across Defendant Status

	(1)	(2)	(3)	(4)	(5)	(6)
	$R\&D_{jt}/Asset_{jt}$					
Public firms	*LikelyDef.*	*LikelyPlaint.*	*All*	*LikelyDef.*	*LikelyPlaint.*	*All*
$Post \cdot Exposure_j$	0.003**	−0.002	0.003***	0.005***	−0.002	0.005***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)
$Post \cdot LikelyPlaint.$			0.004***			0.003**
			(0.002)			(0.002)
$Post \cdot Exposure_j \cdot LikelyPlaint.$			−0.005*			−0.004*
			(0.003)			(0.003)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	N	N	N
Indu. × *Time F.E.*				Y	Y	Y
$Controls_j \times Time$				Y	Y	Y
R^2	0.016	0.013	0.016	0.079	0.273	0.068
Observations	1,642	392	2,034	1,642	392	2,034

Notes. This table reports the estimate of the linear difference-in-difference specification (Equation (1)), where I allow the effect of the exposure to the decision to be heterogeneous across firm likelihood of being a plaintiff at the time of the decision. The outcome is always *R&D/Asset*, which is the average over the period of the quarterly R&D expenses scaled by total assets. As usual, this is winsorized at 1%. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation by technology class since 2000. The dummy *Likely Plaintiff* is equal to one if the firm has been involved in more lawsuits as a plaintiff rather than a defendant. I first report the regressions as split between likely plaintiff and the complementary group—which I call likely defendant—and then I report the fully interacted regression in the whole sample. I always control for firm and time fixed effects, but in columns (4)–(6), I add extra controls interacted with time dummies. As in the previous analyses, I control for industry; location dummies of the firm; the size of the portfolio before the estimation period; a dummy for "start-up," which in this context is firms that published the first patent in the three years before the decision; and average quality of the patent portfolio in the pre period, measured by average citations. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.
$^*p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Consistently with the hypothesis presented before, I find that the effect is mostly driven by those firms that were more likely to be on the defensive side of a lawsuit. In fact, within this sample the effect of exposure is significant and positive, while for firms more likely to be a plaintiff the result is null. In column (3), I pull together the sample and formally test for the difference in the effects across the two groups. As expected, I confirm that the difference between the two groups is statistically significant. In the remaining columns (columns (4)–(6)), I show that this result is qualitatively identical when I add the usual set of controls.

As the final step, in Table A.14 in the online appendix, I use the same type of analysis, but for patenting. In the two panels, I explore this separately for my full sample and for public firms only. The results across the two data sets are very consistent. As before, I find that firms more exposed to patent litigation responded positively to the shock, but only in the sample of firms that were more likely to be defendants. In particular, the effect for firms more likely to be a plaintiff is always small—if not negative—and highly insignificant statistically. However, unlike before, these differences are not statistically significant at the conventional level. Although these results are not completely in line with the R&D estimates, they seem to confirm that the subset of firms that were more likely to be involved in a lawsuit as plaintiffs did not benefit much from the new rules.

Overall, these results provide a final robustness test for the mechanism of the paper, which suggests that the new rules regarding injunctions had an impact on innovation activity by affecting the balance of enforcement in patent litigation. Consistent with this mechanism, firms that were more likely to be a defendant appeared to have responded more positively than those that were more likely to be a plaintiff.

6. How Does Litigation Exposure Affect Innovation?

In the previous sections, I showed that the Supreme Court decision led to an increase in patenting, both at the intensive and extensive margins. Furthermore, this change in enforcement also positively affected patent quality and R&D investments. Overall, this evidence suggests that patent litigation had real distortive effects on firms' ability to innovate in 2006, and the decision was able to reduce some of this burden faced by innovative firms.

In this section, I explore why patent litigation affects innovation by firms.

6.1. Litigation Lowers Innovation Returns: Evidence from the Composition of Innovation

Firms exposed to litigation may reduce innovation for different reasons. The most intuitive channel is that patent litigation lowers the returns from investing in innovation. Because direct involvement in patent litigation can be very expensive (Bessen and Meurer 2013), firms will take into account this risk when assessing whether to invest in a project. As a result, when the risk of patent litigation is too high, firms may choose to forgo some good investment opportunities. If this channel is quantitatively important, I should expect to find two results in the data. First, firms operating in more intensively litigated areas should be more positively affected. This is what I found in the main results. Second, within a firm, projects in an area where patent litigation is more intense should become relatively more valuable. This reshuffle should happen in every firm, irrespective of whether it is more or less exposed to litigation. In other words, every firm should perceive the investment in riskier patents to be more valuable.

In order to provide evidence in favor of this idea, I study whether firms experienced a relatively increase in risky patents after the decision. In order to focus on within-firm resource allocation, I sort patents applied for by each firm across two categories—risky and nonrisky—depending on whether they belong to one of the USPTO technology classes in the top 10% (or 25%) of litigation. This reshape of the data implies that each firm has two observations per period. Because I am interested in the within-firm allocation, I can now test whether risky patents increased relatively more after the decision conditional on a full set of firm-by-time fixed effects. In practice, I estimate the following equation:

$$y_{jtr} = \alpha_{jr} + \alpha_{jt} + \beta 1\{Risk_r\} * Post, \tag{5}$$

where α_{jt} is a set of firm-time fixed effects, α_{jr} is a set of fixed effects at the firm-group level, and $1\{Risk_r\}$ is a dummy for riskier groups. As mentioned above, I group patents in two classes, such that $r = \{highrisk; lowrisk\}$. If the return channel is the driving force behind the response of innovation to the ruling, I would expect risky patents to grow substantially more than nonrisky patents within the firm portfolio, which is $\beta > 0$. Standard errors are clustered at firm-level.

In this analysis, I consider two outcomes: First, I explore the intensive margin of the effect by looking at $\ln(pat_{jtr})$, which is the logarithm of the grant patent that firm j applied for during time t in the class of risk r. To obtain a purely intensive margin, I estimate this regression with a subset of firms that are simultaneously active in both risk classes around the decision time. Second, I look at the extensive margin with y_{jtr} equal to $1\{Pat_{jtr} > 0\}$, which is a dummy equal to one if the firm j applies for any subsequently granted patent in risk-group r at time t. In this case, my sample is much larger, because I consider every firm that has applied for at least one patent in the 10 years before the decision.

Results are reported in Table 7. When I look at the intensive margin, I find that patents belonging to

Table 7. Evidence on Patent Mix

	(1)	(2)	(3)	(4)
	Extensive margin		Intensive margin	
Variable	$\ln(Patents_{jtr})$		$1\{Patents_{jtr} > 0\}$	
Post · $1\{Risk_r\}$	0.005 (0.020)	−0.027* (0.015)	0.280*** (0.003)	0.170*** (0.004)
Split	10%	25%	10%	25%
Firm × Time F.E.	Y	Y	Y	Y
Firm × Risk F.E.	Y	Y	Y	Y
R^2	0.981	0.976	0.834	0.810
Observations	8,712	15,572	219,616	219,616

Notes. This table estimates Equation (5), which is: $y_{jtr} = \alpha_{jr} + \alpha_{jt} + \beta 1\{Risk_r\} \cdot Post$, where α_{jt} is a set of firm-time fixed effects, α_{jr} is a set of fixed effects at the firm-group level, and $1\{Risk_r\}$ is a dummy for more risky groups. Data are reshaped for this analysis at the firm-time-riskiness group level. In practice, I group patent within a firm in a certain time across two classes depending on the level of riskiness r, such that $r = \{highrisk;lowrisk\}$. Patents are assigned to one of the two groups based on the intensity of litigation of their technology class after 2000 based on the WestLaw Litigation data. In particular, I split the data across both 10% and 25%. Furthermore, data are collapsed before and after the decision: Therefore, every firm is in the data exactly four times. I consider two outcomes: In columns (1) and (2), I use $\ln(pat_{jtr})$, which is the logarithm of the patent applications that firm j filed to during time t in the class of risk r. Because this should capture the intensive margin of the treatment, I use all firms that are simultaneously active in both risk classes, around the decision time. Then, in columns (3) and (4), I have y_{jtr} to be equal to $1\{Pat_{jtr} > 0)\}$, which is a dummy equal to one if the firm j applies to any granted patent in risk-group r at time t. In this case, my sample is much larger, and I consider every firm that has applied to at least one patent in the 10 years before the decision. Standard errors are clustered at firm level. All regressions include a constant.

$^*p < 0.10$; $^{***}p < 0.01$.

more intensively litigated patent classes do not appear to increase relatively more within a firm's portfolio. Estimates are very small and noisy.[54] On the other hand, I find that firms are more likely to patent in a risky class in the two years after the decision, rather than in the two before. The results are similar whether risky patents are defined as being in the top 10% or the top 25%. Lastly, in Table A.12 in the online appendix, I use the full panel dimension—without collapsing pre and post—and I estimate the effect differentially for pre and post relative to the quarter before the decision. This analysis shows that this effect is not driven by differential trends in patenting before the decision.

At least partially, these results are consistent with the return channel: The decision also shifted the patenting behavior of firms across classes, in particular, by making companies more likely to patent in a more risky area after the decision. Although a similar effect is not identified at the intensive margin, these results are in line with the reshuffle idea that should occur if the decision were to increase the perceived returns of R&D investment.[55]

6.2. Litigation Exacerbates Financial Constraints

Operating in a high-litigation environment can also hinder innovation by reducing the amount of resources available for R&D. The idea that exposure to litigation can deplete corporate resources is intuitive. Firms in sectors where litigation is more intense are more likely to pay large settlements or overpays for licensing agreements. This happens because companies want to avoid the escalation of legal conflicts to courts or just limit their negative consequences, as in the BlackBerry case discussed previously. Furthermore, ex ante, these companies may be forced to devote larger resources to monitor potential threats and modify their products to minimize the risk of litigation. For example, eBay in the 2006 10-K recognizes that litigation claims, "whether meritorious or not, are time-consuming and costly to resolve and could require expensive changes in our methods of doing business, or could require to enter into costly royalty or licensing agreements." In response to this, companies may invest more intensively in defensive tools, such as a large legal department within the company, which seems to have some effects on deterring attacks (Cohen et al. 2019).

If the financing of innovation were frictionless, this shift of monetary resources should not affect firms' ability to invest in good projects. In reality, firms face constraints in funding innovation (Brown et al. 2009, Hall and Lerner 2010), and therefore a reduction in internal resources has an impact on firms' ability to innovate. When this is the case, intense patent litigation exacerbates this financing problem, and, therefore, it increases the inefficiency in funding R&D. To test whether this theory is true in the data, I examine the heterogeneity of the decision effects across firms characterized by a differential likelihood of being financially constrained. If this channel is relevant, I expect companies that are more likely to be financially constrained to react more positively to the shock. In other words, this story would predict a higher elasticity between investment in R&D and a reduction in litigation costs for companies facing more financial frictions.

In order to study this, I modify the standard model described by Equation (1) by adding an interaction with a dummy $FinCon_j$, which is equal to one for firms that are more likely to be financially constrained. More specifically, I estimate:

$$y_{jt} = \alpha_j + \alpha_t + \beta_1(Exposure_j * FinCon_j * Post) + \beta_2(FinCon_j * Post) + \beta_3(Exposure_j * Post) + \gamma X_{jt} + \epsilon_{jt}. \tag{6}$$

Furthermore, I separately study the behavior of the two groups of firms. In line with previous discussion, I would expect $\beta_1 > 0$.

Following the relevant literature in finance, I identify firms that are more likely to be financially constrained in three different ways. First, I study the differential behavior of small versus large firms. Previous research has found that smaller firms tend to have a harder time accessing external funding (Fazzari et al. 1988, Chodorow-Reich 2014, Bottero et al. 2016). In my setting, I focus on smaller firms within the public-firm sample. In particular, I construct two definitions of small firms, looking at whether they are below the median of employment or revenue. Second, I identify firms with no rating on public debt as companies that are more likely to be financially constrained (Kashyap and Lamont 1994, Almeida et al. 2004). More specifically, I look at firms with no rating reported in the three years before the Supreme Court decision. Lastly, I examine the heterogeneity across firms that pay and do not pay cash dividends, looking at the three years before the decision.

The results are reported in Tables 8, 9, and 10. The decision led to an increase in R&D intensity only for firms more likely to be financially constrained. When splitting the sample across the two groups, I systematically find that the coefficient is positive and significant for the financially constrained group, while nonsignificant and small for the other group. When using the full sample, more financially constrained firms increase R&D intensity more. This is true across all the measures, although it is not statistically significant in some cases. Lastly, I find that more financially constrained firms did not respond more

Table 8. Effect of the Decision Across Firm Size

	(1)	(2)	(3)	(4)	(5)	(6)
	$R\&D_{jt}/Asset_{jt}$					
Variable	*Small*	*Large*	*All*	*Small*	*Large*	*All*
	Panel A: Heterogeneity by size: Revenue					
Median revenue						
$Post \cdot Exposure_j$	0.004**	–0.001	–0.001	0.005***	–0.001	0.002*
	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
$Post \cdot Small_j$			–0.004**			–0.003*
			(0.002)			(0.002)
$Post \cdot Exposure_j \cdot Small_j$			0.004**			0.003*
			(0.002)			(0.002)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	N	N	N
Indu. × Time F.E.				Y	Y	Y
$Controls_j$ *× Time F.E.*				Y	Y	Y
R^2	0.017	0.007	0.022	0.120	0.076	0.072
Observations	956	1,078	2,034	956	1,078	2,034
	Panel B: Heterogeneity by size: Employment					
Median employment						
$Post \cdot Exposure_j$	0.003**	–0.001	–0.001	0.005***	–0.001	0.002**
	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
$Post \cdot Small_j$			–0.003**			–0.001
			(0.002)			(0.002)
$Post \cdot Exposure_j \cdot Small_j$			0.004**			0.003
			(0.002)			(0.002)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	N	N	N
Indu. × Time F.E.				Y	Y	Y
$Controls_j$ *× Time F.E.*				Y	Y	Y
R^2	0.016	0.003	0.021	0.112	0.058	0.073
Observations	969	1,065	2,034	969	1,065	2,034

Notes. This table reports the estimate of the linear difference-in-difference specification (Equation (1)), where I allow the effect of the exposure to the decision to be heterogeneous across firm size. The outcome is always *R&D/Asset*, which is the average over the period of the quarterly R&D expenses scaled by total assets. As usual, this is winsorized at 1%. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent applications in the five years before the decision and patent litigation at technology class since 2000. Panel A reports the result measuring size based on revenue before the decision, and, in particular, I divide the sample above and below the median. In Panel B, I do the same, but using employment as sorting variables. I first report the regressions as split between large and small firms, and then I report the fully interacted regression in the whole sample. I always control for firm and time fixed effects, but in columns (4)–(6), I add extra controls interacted with time dummies. As in the previous analyses, I control for industry; location of the firm; the size of the portfolio before the estimation period; a dummy for "start-up", which in this context is firms that published the first patent in the three years before the decision; and average quality of the patent portfolio in the pre period, measured by average citations. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.

$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

than nonfinancially constrained firms in terms of patent applications.[56]

As a robustness, I show that, in my case, the results are not simply capturing heterogeneity across firms in growth (Farre-Mensa and Ljungqvist 2015). To rule this out, I augment Equation (6) by fully interacting measures of firm growth in the two years before the decision to my treatment. In particular, in Table A.15 in the online appendix, I report the results looking at revenue growth. I find that, if anything, the main coefficient β_1 is estimated more precisely when I add the growth controls. In an unreported table, I find the same when looking at asset growth. Overall, my analysis is not just capturing a spurious correlation of these measure of financial constraint with different growth trajectories.

These results suggest that a decline in R&D returns is not the only channel through which patent litigation may affect innovation. Instead, financial constraint is an important dimension to consider when evaluating the effect of operating in area where litigation is intense.

7. Conclusion

This paper examines how patent rights affect innovation using the 2006 Supreme Court decision in *eBay v. MercExchange* as an exogenous shock to patent enforcement. The evidence provided suggests that this intervention had a positive effect on innovation.

Table 9. Effect of the Decision Across Measures of Financial Constraint: Heterogeneity by Dividend Payers

	(1)	(2)	(3)	(4)	(5)	(6)
	$R\&D_{jt}/Asset_{jt}$					
Variable	*NoDividend*	*Dividend*	*All*	*NoDividend*	*Dividend*	*All*
$Post \cdot Exposure_j$	0.003**	−0.002	−0.002	0.006***	−0.002	0.001
	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)
$Post \cdot 1\{Div_j = 0\}$			−0.005***			−0.004**
			(0.002)			(0.002)
$Post \cdot Exposure_j \cdot 1\{Div_j = 0\}$			0.005**			0.005*
$Post \cdot Exposure_j \cdot 1\{Div_j = 0\}$			(0.003)			(0.003)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	N	N	N
Indu. × *Time F.E.*				Y	Y	Y
$Controls_j$ × *Time F.E.*				Y	Y	Y
R^2	0.019	0.016	0.019	0.105	0.092	0.069
Observations	1,322	712	2,034	1,322	712	2,034

Notes. Tables 9 and 10 report the estimate of the linear difference-in-difference specification (Equation (1)), where I allow the effect of the exposure to the decision to be heterogeneous across firms characterized by different rating status or dividend policies. The outcome is always *R&D/Asset*, which is the average over the period of the quarterly R&D expenses scaled by total assets. As usual, this is winsorized at 1%. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. This table reports the result dividing the sample across firms that paid positive cash dividends in any quarters in the three years before the decision and firms that did not. I first report the regressions as split between the two groups, and then I report the fully interacted regression in the whole sample. I always control for firm and time fixed effects, but in columns (4)–(6), I add extra controls interacted with time dummies. As in the previous analyses, I control for industry; location dummies of the firm; the size of the portfolio before the estimation period; a dummy for "start-up," which in this context is firms that published the first patent in the three years before the decision; and average quality of the patent portfolio in the pre period, measured by average citations. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.
$^*p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Firms that were more exposed to the change in rules—companies operating in areas where patents were more intensively litigated—increased innovation output more after the decision. Similarly, for a subsample of public firms, I found that R&D intensity was positively affected. This is consistent with the idea that patent litigation may have negative, distortive effects on firm investment in innovation. The effects were large in magnitude, suggesting that these distortions can be substantial. Although the average quality of the patents did not change, firms more exposed to patent litigation increased the likelihood of patenting breakthrough technology. Similarly, firms exposed to the shock saw a lower increase in the share of defensive patents. Overall, these results are consistent with the idea that patent litigation may have negative, distortive effects on firm investment in innovation.

Furthermore, I investigate the specific channels through which patent litigation reduced innovation. First, I show that patent litigation reduces innovation because it lowers the returns from performing R&D activities. Consistent with this idea, firms partially reshuffled their portfolios toward patents with higher risk of lawsuits after the decision. Second, I explore whether patent litigation also reduces investment in R&D because it diminishes the amount of internal resources available for productive activities, therefore exacerbating the financing problem of innovation (Brown et al. 2009, Hall and Lerner 2010). In line with this hypothesis, I find that the increase in R&D is mostly concentrated in firms that are more likely to be financially constrained.

There are several avenues for future research in this area. A primary question is to examine the effectiveness of recent policy interventions, such as the America Invents Act (2011). In addition, more work can be done to examine the role of patent litigation in start-up. The nature of my identification strategy focuses on established firms, and, therefore, the results do not directly apply to start-up companies. However, there are good reasons to think ex ante that the results on this set of companies should not be reversed. First, the importance of financial constraints in explaining the results suggests that the litigation channel may have been relevant also for start-ups, because these firms are generally more financially constrained than established companies. Second, aggregate evidence is also consistent with the fact that eBay did not significantly harm start-up investments. For instance, Mezzanotti and Simcoe (2019) suggest that venture capital investments and aggregate innovation did not slow down during this period (and, if anything, grew at a faster rate).

The results presented in this paper support the idea that patent litigation can significantly affect

Table 10. Effect of the Decision Across Measures of Financial Constraint: Heterogeneity by Rating Status

	(1)	(2)	(3)	(4)	(5)	(6)
	$R\&D_{jt}/Asset_{jt}$					
Variable	*NoRating*	*Rating*	*All*	*NoRating*	*Rating*	*All*
$Post \cdot Exposure_j$	0.003**	−0.001	−0.001	0.005***	−0.001	0.002*
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
$Post \cdot 1\{Rating_j = NO\}$			−0.003**			−0.001
			(0.001)			(0.001)
$Post \cdot Exposure_j \cdot 1\{Rating_j = NO\}$			0.004**			0.002
			(0.002)			(0.002)
Firm F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	N	N	N
Indu. × Time F.E.				Y	Y	Y
$Controls_j$ *× Time F.E.*				Y	Y	Y
R^2	0.014	0.012	0.014	0.092	0.089	0.066
Observations	1,336	698	2,034	1,336	698	2,034

Notes. Tables 9 and 10 report the estimate of the linear difference-in-difference specification (Equation (1)), where I allow the effect of the exposure to the decision to be heterogeneous across firms characterized by different rating status or dividend policies. The outcome is always *R&D/Asset*, which is the average over the period of the quarterly R&D expenses scaled by total assets. As usual, this is winsorized at 1%. The variable $Exposure_j$ captures the exposure of firm *j* to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. In this table, I do the same as in Table 8A, but sorting based on whether the firm has any rating reported in Compustat in the three years before, looking at S&P Domestic Long Term Issuer Credit Rating. I first report the regressions as split between the two groups, and then I report the fully interacted regression in the whole sample. I always control for firm and time fixed effects, but in columns (4)–(6), I add extra controls interacted with time dummies. As in the previous analyses, I control for industry; location dummies of the firm; the size of the portfolio before the estimation period; a dummy for "start-up," which in this context is firms that published the first patent in the three years before the decision; and average quality of the patent portfolio in the pre period, measured by average citations. More information on the variables is provided in Online Appendix A.3. Standard errors are clustered at firm level. All regressions include a constant.
$^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

companies' innovation. As a result, policies that mitigate the overhang of litigation can have beneficial effects on technology advancement.[57] In particular, improvements in the quality of patent enforcement, which reduce the legal uncertainty around patents and limit abusive behaviors in this market, can increase firms' ability and incentives to invest in R&D. Recent efforts in the United States, such as the America Invents Act (2011) or *Alice Corp. v. CLS Bank International*, 573 U.S. 208 (2014), have started to take steps in this direction. However, more comprehensive policy work needs to be done to further addresses the various problems in the patent system today.

Acknowledgments

The author thanks Pere Arque-Castells, Lauren Cohen, Jiashuo Feng, Samuel G. Hanson, Victoria Ivashina, Xavier Jaravel, Louis Kaplow, James Lee, Simone Lenzu, Josh Lerner, Christopher Malloy, Diana Moreira, Gianpaolo Parise, Andrea Passalacqua, Tom Powers, Giovanni Reggiani, Jonathan Rhinesmith, David Scharfstein, Andrei Shleifer, Tim Simcoe, Stanislav Sokolinski, Luke Stein, Jeremy Stein, Adi Sunderam, Haris Tabakovic, Edoardo Teso, Rosemarie Ziedonis, and John Zhou; and the seminar participants at the Boston University Technology Initiative conference, American Finance Association (AFA) 2018, Northwestern Law, European Finance Association (EFA) 2017, Harvard University, Columbia University, Yale University, Cornell University, the University of Rochester, Federal Reserve Bank (FED) Board, Texas A&M University, Bocconi, and Cornerstone for helpful comments. The author also thanks Davin Wang and Matthew Nicholas Nicholson for excellent research assistance. All errors are the author's own.

Endnotes

[1] A short discussion regarding the increase in patent-holders' rights and the related research in this area can be found in Lerner (2002).

[2] In an interview with the *National Law Journal* (March 13, 2006, Volume 27, Issue 77), patent litigator David Clonts of Akin Gump Strauss Hauer & Feld, stated that "If BlackBerry knew it could successfully defend against an injunction and instead have a trial on money damages, the settlement value would have been a tenth of what it was."

[3] To quote the Supreme Court majority opinion, the threat of injunction was frequently used "as a bargaining tool to charge exorbitant fees to companies that seek to buy licenses to practice the patent."

[4] This view was shared by many scholars and practitioners. According to the American Innovators Alliance, an association representing large high-tech companies, because of high injunction risk, "money that could go to productive investments is instead diverted to legal fees and settlement payments," leading to "... less innovation." The sentences are taken from the "amicus curiae" submitted for the Supreme Court case.

[5] Litigation claims, "whether meritorious or not, (...) could require expensive changes in our methods of doing business, or could require to enter into costly royalty or licensing agreements" (eBay 2006 10-K).

[6] One limitation of this study—which is driven by the methodology used —is that its results cannot be directly generalized to start-ups. Although more research is definitely needed to explore this dimension,

in the conclusion (Section 7), I discuss how evidence from this study and other works (e.g., Mezzanotti and Simcoe 2019) may suggest that results for that sample should not be reversed. Furthermore, it is important to highlight how established firms undertake the large majority of R&D investment in the United States.

[7] On top of looking at different dimensions of firm activity (innovation versus business creation), Appel et al. (2017) also differ from this work across (at least) two important dimensions. First, the papers look at very different types of policy interventions. Second, the two works focus on different populations (established versus new firms), which are likely to be affected in different ways by litigation.

[8] This analysis is also related to the body of work in finance that focuses on the effect of litigation risk—mostly shareholder litigation—on corporate policies (Haslem 2005, Rogers and Van Buskirk 2009, Kim and Skinner 2012, Appel 2019, Arena and Julio 2015, Lin et al. 2020).

[9] I provide some background legal information about the *eBay v. MercExchange* case in Online Appendix A.1. More discussion on the background of the case and its policy implication can also be found in Mezzanotti and Simcoe (2019).

[10] One interesting quote can be found in the analysis of the case in Wesenberg and O'Rourke (2006): "In determining whether to settle a case, a market participant must consider many factors, including (1) the expense of litigation, (2) the potential exposure, and (3) the threat of an injunction forcing the company to either terminate a product or excise a component or part from a larger product, at potential prohibition, cost or delay. Oftentimes, it is this final threat of injunctive relief that forces the market participant to settle. As a practical matter, certainty trumps justice and accused defendants agree to pay an exorbitant license fee for a questionable patent and continue to operate rather than risk discontinuing a product or operations altogether."

[11] American Innovators Alliance is a lobby group that represents large tech firms, such as Microsoft, Micron, Oracle, and Intel. The sentences are taken from the amicus curiae that the group submitted for the Supreme Court case. Similar quotes can be found in the amicus curiae submitted by the Computer & Communication Industry Association (CCIA): For instance, they claim that automatic injunction did "produce anti-competitive behavior, foster more litigation, and undermine innovation."

[12] Similar results are also provided in an earlier empirical analysis in Grumbles et al. (2009).

[13] This hypothesis may explain why the number of lawsuits kept increasing after the eBay ruling. Another proposed explanation is that defendants after eBay were less concerned with going to court and, therefore, became less likely to settle ex ante.

[14] As discussed in Shapiro (2016a), injunction could lead to an excessive compensation of the patent-holder both during ex post (litigation) and ex ante negotiation. Because both negotiations are strictly tied in practice, the language of the paper refers to both of them as being part of the "litigation channel."

[15] Importantly, this effect does not necessarily imply that the number of lawsuits should go down, as discussed before. In fact, the number of lawsuits is an equilibrium outcome and, therefore, also depends on the willingness of the accused firm to settle versus go to court and on the need of the accusing firm to bring the counterparty to court to make the threat credible. These two effects can actually increase the observed number of lawsuits. Therefore, after eBay, companies interested in strategic litigation may need to go to court more to prove the seriousness of their intentions (Cohen et al. 2016a). At the same time, defendants may be less concerned about a court case, therefore making them more willing to go to court rather than settle.

[16] Although an increase in patenting may still have some positive effect—for example, Hegde and Luo (2018) for the role of disclosure—these benefits are likely to be much lower than those caused by an actual increase in innovation.

[17] Not every NPE can be accused of acting like a "patent troll." For instance, universities and other research institutions are in the NPE category. By the same token, not all the abusive behavior is specific to NPEs.

[18] The first firm published a list of top NPEs active in the United States in 2014 (https://www.patentfreedom.com/about-npes/holdings/), where companies were selected based on the number of patents held. The second, instead, published a study on stock returns on NPEs in 2013, using both public and private information for compiling a list of NPEs that are publicly traded (http://patentvue.com/2013/04/15/508-publicly-traded-patent-holding-companies-yield-impressive-returns/).

[19] The majority of the companies appear in both lists—six—and only one company is only listed by PatentFreedom. The companies are Acacia Technologies, Asure Software (formerly Forgent Network), Rambus, Tessera Technologies, Universal Display, Document Security Systems, ParkerVision, Unwired Planet (formerly Openwave), Interdigital, and Spherix.

[20] More information about the analysis can be found in Online Appendix A.3.3. One caveat of the data set is that it is compiled based on a recent list; therefore, I may have missed an NPE that was active and public in 2006, but is defunct today. Although I cannot exclude this possibility, I could not find any example of this phenomenon in the data.

[21] For instance, the average return the day of the decision is −3.4%. When dropping one company at the time, I get results between −2.97% and −3.75%. In all cases, the result is 1% significant.

[22] Almost all the analyses are run with applications made by the end of 2008, therefore allowing more than the five years recommended by Dass et al. (2017) to eliminate risk of truncation bias.

[23] I consider firms in nonfinancial and nonregulated industries, headquartered in the United States, not involved in financial restructuring and with information reported in the quarterly Compustat data. More details are available in Online Appendix A.3.

[24] This term "start-up" is likely not particularly accurate to describe the variable. To make sure readers will understand what this captures, I provide a clear definition of this variable (and the others) in each table.

[25] Dates are usually reported in terms of quarters (e.g., 2006Q1): These quarters are constructed in event time, where I artificially set the end of the first quarter of the year at May 15. The other quarters are then constructed consistent with this.

[26] In the text, I will discuss case by case the definition of the outcome used.

[27] As expected based on the literature (Bertrand et al. 2004), clustering the errors does not have a material impact when the sample is collapsed in two periods (pre and post), but it is important when presenting the result as dynamic effect and, therefore, using the panel at quarterly frequency.

[28] For instance, if a company operates in four technology classes with two patents granted in the each of these classes, then the vector $t(j)$ will be equal to zero for every technology class where there were no patents and equal to 0.25 for the four technology classes where the company patented something.

[29] One of the advantages of WestLaw is that these data go back to 1980, unlike other sources. For instance, RPX—another leading data source used for research in this area (Cohen et al. 2019)—generally provides data on litigation starting from 2005. The same data are also

known as Derwent LitAlert data. The data were accessed through the online tool LitAlert

[30] First, each filing may contain multiple defendants. Firm A suing firms B and C in the same filing should carry more weight that firm A suing only firm D. Second, each filing may contain more than one patent, because in the same case, the plaintiff may sue the defendant over multiple technologies. In order to address this, I reshape the data at the single defendant–plaintiff–patent level.

[31] Qualitative inspection also provides supportive evidence, as I find technology classes that are expected to have high level of litigation to actually be on the top of the ranking. For instance, at the top of the ranking, I find the two main classes for drugs (514 and 424) and two of the prominent classes for communications technologies (379 and 340), plus business method (705) and one electronic class covering illumination technologies (362).

[32] To the extent that $exposure_i$ captures just a very noisy measure of litigation, I expect measurement error to bias the results toward zero. Therefore, at the extreme, measurement error should push us toward finding no effect.

[33] As discussed in detail in Online Appendix A.4.1, I repeat the same analysis on the day in which the Supreme Court decided to hear the case (November 28, 2005) and at the oral hearing (March 29, 2006). Although the decision to hear the case is not associated with any significant response, the oral argument is associated with negative outperformance of firms more exposed to litigation. The magnitude of this negative effect is around one-half of the size of the positive outperformance at announcement. This effect seems to be consistent with the idea that eBay did not come out as a clear winner from the oral argument, therefore confirming that the outcome of the decision was unexpected.

[34] In particular, in the reported table, I require the firm j to have applied for at least one granted patent in the two years before and in the one year after. This choice is motivated by the fact that I want the sample in this table to be equivalent to the one I use in one of the next sections, where I estimate the same equation over different periods, from one to three years after. Results are unchanged if I consider the set of firms with at least one patent in the two years before and one in the two years after.

[35] As an alternative to the previous analysis, I estimate the trends in the model by assuming that the relationship between exposure to litigation and patenting is linear. I essentially estimate $y_{jt} = \alpha_j + \alpha_t + \beta^{PRE} R_j * PRE + \beta^{POST} R_j * Post + \epsilon_{jt}$. The excluded period in this case is the quarter right before the decision, and, therefore, the two coefficients should be interpreted as changes relative to that quarter. Although this approach is less flexible than the previous specification, it allows me to obtain more precise estimates of the trends and therefore to rule out that the lack of a pre-trend may have been due to a lack of power. As expected, exposure to litigation does not predict differential behavior before the decision, but only after (columns (1) and (2), Table A.5 in the online appendix).

[36] In other words, I estimate the same model in Equation (1), but center the analysis in a quarter where there is no change in patent law. In order to do so, I reconstruct the outcomes and regressors as if the shock occurred right after the quarter of interest.

[37] Clearly, after 2004Q1, the post period of the placebo analysis would overlap with the post-treatment period. Because of this, a similar placebo centered after 2004Q1 would not be a true placebo, because the estimated parameters would capture part of the treatment effects. Furthermore, the closer I come to 2006Q1, the more my analysis would look like the main results. Consistent with this, I find that post 2004Q1 the β starts converging toward the main results in Table 2. As expected, the convergence is smooth and the effects turns positive and significant at 95% only at the end of 2005.

[38] The z-score on the difference is small, around 0.29.

[39] For consistency with the rest of the measures, I look at the patents applied for between four and two years before the decision.

[40] This figure is constructed by using the longer sample, but it is qualitatively identical using the same period as the main analyses. It is interesting to notice that the dynamic of the effect with and without controls are qualitative identical. The main difference is that the result with controls is characterized by slightly larger magnitude and more precise estimates.

[41] One potential exception was the Medicare reform of 2003, which may have influenced the pattern of innovation, at least for healthcare related technologies. However, as discussed in Online Appendix A.2, the timing of this reform would be inconsistent with our results, as the impact of Medicare reform appeared already much earlier than 2006 (e.g., Krieger et al. 2018).

[42] Because patent citations increase over time, their measure is sensitive to the date on which the patent was granted, relative to the last date on which the data were updated. To avoid this truncation problem, I look at citations in the three years after the grant of the patents. This approach is consistent with other works in this area (e.g., Bernstein 2015), and it reflects the fact that patents tend to receive most of their citations early in their life, with strong serial correlation in citations afterward (Akcigit and Kerr 2018).

[43] Adjusting citations may rise concerns related to "bad controls," because the adjustment factor may be itself affected by the policy. As shown in Table A.10 in the online appendix, the results using unadjusted citations (raw average citations) still confirm that the policy did not lead to a decline in average quality. In this case, I actually find a positive and significant effect after I add controls.

[44] In Table A.5 in the online appendix, where I analyze the pre-trend on this variable, I actually find that in this panel specification, there is some weak, positive effect on the average citations.

[45] In line with the rest of the literature, originality is measured as one minus the Herfindahl index of technology class dispersion of citations made by the patent to other patents. In other words, this is a measure of the dispersion of the patent's references across the different technologies.

[46] I cannot specify a general threshold in terms of citations, because the threshold depends on the comparable set of patents, which are patents in the same technology and year. However, the bottom three quartiles in terms of citations capture the bulk of patents of very low quality. The median patent in this group has zero citations, and the average is only 0.4.

[47] For the average firm in the sample, the share of defensive/strategic patents is in fact 21%. For the median firm, the same value is about 8%.

[48] This approach builds on the recent work by Srinivasan (2018), who shows that the development of business method patents appear to be mostly related to strategic considerations.

[49] The list of these other technology classes is in Hall (2003, table 3). In particular, these are technology classes: 84; 119; 379; 434; 472; 380; 382; 395; 700; 701; 702; 703; 704; 705; 706; 707; 709; 710; 711; 712; 713; 714; 715; 717; and 902.

[50] One difference with Hall and Ziedonis (2001) is that the overall patenting activity did not decrease. This difference can probably be explained by two important differences between this experiment and the setting in Hall and Ziedonis (2001). First, the two papers study completely different time periods, with a very different litigation landscape: In particular, Hall and Ziedonis (2001) focus on the 1980s. Second, Hall and Ziedonis (2001) focus on the semiconductor industry. This is an industry characterized by very specific business conditions that are hard to generalize outside the specific context.

Interestingly, despite these big differences, both papers find consistent results on the direction of the elasticity between injunction use and defensive patenting.

[51] However, I want to be cautious in this interpretation. First, public firms operate to a much larger scale than the average company in the full sample. This difference in scale may be important to compare adjustments across groups, therefore affecting the relative magnitude of the effects. Second, this comparison only accounts for intensive margin of adjustment, and I know from the previous results that extensive margin also plays an important role in the full sample.

[52] For instance, Cohen et al. (2019) suggest that a large part of strategic litigation is driven by the presence of "extra cash" on targets' balance sheets, therefore linking aggressive behavior of NPEs to the ability to extract large profits through settlements. Although Cohen et al. (2019) examine variation within public firms, it is reasonable to assume that public firms in general can be perceived as being more profitable targets for litigation.

[53] I thank Matthew Nicholas Nicholson for the excellent support on the name matching.

[54] In particular, the effect is nonsignificant when looking at the 10% split, and borderline significant, but negative, when looking at the 25% split. Overall, I interpret this evidence as consistent with a lack of response at this margin. First, the effect is small in size with both outcomes. Second, the borderline negative effect does not appear to be particularly robust. For instance, when we use the full panel dimension to study the pre-trend (Table A.12 in the online appendix), we do not find any significant effect across both outcomes.

[55] One view on this difference is that an intensive margin is harder to trace down empirically. Alternatively, it is possible that firms that already operate across areas with both high and low risk of litigation are endogenously less sensitive to patent litigation. As a result, the positive net present value effect for these firms may be smaller and empirically not relevant.

[56] The presence of multiple (risk versus financial) channels may explain this null effect on patenting. In particular, if the risk channels explain a large portion of the patent changes, it is possible to detect no effect on patenting across financial constrain despite the change in R&D. Alternatively, the null effect may hide a shift in the type of project. For instance, financially constrained firms may have invested less in R&D before, but they still have patented the same and simply shuffled resources toward less expensive projects.

[57] There are also more direct applications of this paper. For instance, there has been some recent attempts to moderate the effect of the eBay case on the U.S. legal system. One example is the recent bill titled "STRONGER Patent Act," which would introduce a presumption of irreparable harm when making injunction in patent cases. For references, see "Congress Shouldn't Overturn eBay Patent Injunction Standard" by Thomas Cotter (2018), https://www.law360.com/articles/1083601/congress-shouldn-t-overturn-ebay-patent-injunction-standard.

References

Abrams DS, Akcigit U, Popadak J (2013) Patent value and citations: Creative destruction or strategic disruption? Technical report, National Bureau of Economic Research, Cambridge, MA.

Acharya VV, Subramanian KV (2009) Bankruptcy codes and innovation. *Rev. Financial Stud.* 22(12):4949–4988.

Acharya VV, Amihud Y, Litov L (2011) Creditor rights and corporate risk-taking. *J. Financial Econom.* 102(1):150–166.

Aghion P, Howitt P (1992) A model of growth through creative destruction. *Econometrica* 60(2):323–351.

Akcigit U, Kerr WR (2018) Growth through heterogeneous innovations. *J. Political Econom.* 126(4):1374–1443.

Almeida H, Campello M, Weisbach MS (2004) The cash flow sensitivity of cash. *J. Finance* 59(4):1777–1804.

Angrist JD, Pischke J-S (2008) *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton University Press, Princeton, NJ).

Appel I (2019) Governance by litigation. *J. Financial Econom.* 133(3):708–725.

Appel I, Farre-Mensa J, Simintzi E (2017) Patent trolls and small business employment. Harvard Business School Working Paper 17-072, Harvard Business School, Boston.

Arena MP, Julio B (2015) The effects of securities class action litigation on corporate liquidity and investment policy. *J. Financial Quant. Anal.* 50(1-2):251–275.

Arrow K (1962) Economic welfare and the allocation of resources for invention in national bureau of economic research. *The Rate and Direction of Inventive Activity: Economic and Social Factors*, *National Bureau of Economic Research Special Conference Series*, vol. 13 (Princeton University Press, Princeton, NJ), 609–626.

Bernstein S (2015) Does going public affect innovation? *J. Finance* 70(4):1365–1403.

Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust differences-in-differences estimates? *Quart J. Econom.* 119(1): 249–275.

Bessen J, Hunt RM (2007) An empirical look at software patents. *J. Econom. Management Strategy* 16(1):157–189.

Bessen J, Maskin E (2009) Sequential innovation, patents, and imitation. *RAND J. Econom.* 40(4):611–635.

Bessen J, Meurer MJ (2008a) Of patents and property. *Regulation* 31:18.

Bessen J, Meurer MJ (2008b) *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk* (Princeton University Press, Princeton, NJ).

Bessen J, Meurer MJ (2013) The direct costs from NPE disputes. *Cornell Law Rev.* 99:387.

Bessen JE, Neuhäusler P, Turner JL, Williams JW (2018) Trends in private patent costs and rents for publicly-traded United States firms. *Internat. Rev. Law Econom.* 56:53–69.

Boldrin M, Levine D (2002) The case against intellectual property. *Amer. Econom. Rev.* 92(2):209–212.

Bottero M, Lenzu S, Mezzanotti F (2020) Sovereign debt exposure and the bank lending channel: Impact on credit supply and the real economy. *J. Internat. Econom.* 103328.

Brown JR, Fazzari SM, Petersen BC (2009) Financing innovation and growth: Cash flow, external equity, and the 1990s R&D boom. *J. Finance* 64(1):151–185.

Calabresi G, Melamed AD (1972) Property rules, liability rules, and inalienability: One view of the cathedral. *Harvard Law Rev.* 85:1089–1128.

Carrier MA (2011) *Innovation for the 21st Century* (Oxford University Press, Oxford, UK).

Chetty R, Looney A, Kroft K (2009) Salience and taxation: Theory and evidence. *Amer. Econom. Rev.* 99(4):1145–1177.

Chien CV, Lemley MA (2012) Patent holdup, the ITC, and the public interest. *Cornell Law Rev.* 98(1):1.

Chodorow-Reich G (2014) The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis. *Quart. J. Econom.* 129(1):1–59.

Claessens S, Laeven L (2003) Financial development, property rights, and growth. *J. Finance* 58(6):2401–2436.

Cohen L, Gurun U, Kominers S (2016a) Shielded innovation. Working paper, Harvard University, Boston.

Cohen L, Gurun UG, Kominers SD (2016b) The growing problem of patent trolling. *Science* 352(6285):521–522.

Cohen L, Gurun UG, Kominers SD (2019) Patent trolls: Evidence from targeted firms. *Management Sci.* 65(12):5461–5486.

Corrado CA, Hulten CR (2010) How do you measure a "technological revolution"? *Amer. Econom. Rev.* 100(2):99–104.

Dass N, Nanda VK, Xiao SC (2017) Truncation bias correction in patent data: Implications for recent research on innovation. *J. Corporate Finance* 44:353–374.

Demirgüç-Kunt A, Maksimovic V (1998) Law, finance, and firm growth. *J. Finance* 53(6):2107–2137.

Djankov S, La Porta R, Lopez-de Silanes F, Shleifer A (2003) Courts. *Quart. J. Econom.* 118(12):453–517.

Eisfeldt AL, Papanikolaou D (2014) The value and ownership of intangible capital. *Amer. Econom. Rev.* 104(5):189–194.

Epstein RA (2008) The property rights movement and intellectual property. *Regulation* 30(4):2007–2008.

Executive Office of the President (2013) Patent assertion and U.S. innovation. Report, Executive Office of the President, Washington, DC.

Fazzari SM, Hubbard RG, Petersen BC, Blinder AS, Poterba JM (1988) Financing constraints and corporate investment. *Brookings Papers Econom. Activity* 1988(1):141–206.

Feng J, Jaravel X (2015) Crafting intellectual property rights: Implications for patent assertion entities, litigation, and innovation. *Amer. Econom. J.: Appl. Econom.:* 12(1):140–181.

Ferreira D, Ferreira MA, Mariano B (2018) Creditor control rights and board independence. *J. Finance* 73(5):2385–2423.

Fisher RA (1922) On the interpretation of χ^2 from contingency tables, and the calculation of p. *J. Roy. Statist. Soc.* 85(1):87–94.

Galasso A, Schankerman M (2015) Patents and cumulative innovation: Causal evidence from the courts. *Quart. J. Econom.* 130(1): 317–369.

Gormley TA, Matsa DA (2014) Common errors: How to (and not to) control for unobserved heterogeneity. *Rev. Financial Stud.* 27(2):617–661.

Grumbles E III, Hughey RC, Perera S (2009) The three year anniversary of eBay v. MercExchange: A statistical analysis of permanent injunctions. *Intellectual Property Today* (November 2), 25–29.

Gupta K, Kesan JP (2015) Studying the impact of eBay on injunctive relief in patent cases. University of Illinois College of Law Legal Studies Research Paper 17-03, University of Illinois, Champaign, IL.

Hall BH (2003) Business method patents, innovation, and policy. Technical report, National Bureau of Economic Research, Cambridge, MA.

Hall BH, Harhoff D (2012) Recent research on the economics of patents. *Annu. Rev. Econom.* 4(1):541–565.

Hall BH, Lerner J (2010) The financing of R&D and innovation. Hall BH, Rosenberg N, eds. *Handbook of the Economics of Innovatio*n, vol. 1 (Elsevier, Amsterdam), 609–639.

Hall BH, Ziedonis RH (2001) The patent paradox revisited: An empirical study of patenting in the US semiconductor industry, 1979-1995. *RAND J. Econom.* 32(1):101–128.

Hall BH, Jaffe A, Trajtenberg M (2005) Market value and patent citations. *RAND J. Econom.* 36(1):16–38.

Hall BH, Jaffe AB, Trajtenberg M (2001) The NBER patent citation data file: Lessons, insights and methodological tools. Technical report, National Bureau of Economic Research, Cambridge, MA.

Haslem B (2005) Managerial opportunism during corporate litigation. *J. Finance* 60(4):2013–2041.

Hegde D, Luo H (2018) Patent publication and the market for ideas. *Management Sci.* 64(2):652–672.

Hochberg YV, Serrano CJ, Ziedonis RH (2018) Patent collateral, investor commitment, and the market for venture lending. *J. Financial Econom.* 130(1):74–94.

Holte RT (2015) The misinterpretation of eBay v. MercExchange and why: An analysis of the case history, precedent, and parties. *Chapman Law Rev.* 18(3):677–733.

Hsu P-H, Tian X, Xu Y (2014) Financial development and innovation: Cross-country evidence. *J. Financial Econom.* 112(1):116–135.

Jaffe AB, Lerner J (2011) *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do About It* (Princeton University Press, Princeton, NJ).

Kaplow L, Shavell S (1996) Property rules versus liability rules: An economic analysis. *Harvard Law Rev.* 109(4):713–790.

Kashyap AK, Lamont OA (1994) Credit conditions and the cyclical behavior of inventories. *Quart. J. Econom.* 109(3):565–592.

Katz DM, Bommarito MJ, Soellinger T, Chen JM (2017) Law on the market? Abnormal stock returns and Supreme Court decision-making. Working paper, Illinois Institute of Technology, Chicago.

Kerr WR (2010) Breakthrough inventions and migrating clusters of innovation. *J. Urban Econom.* 67(1):46–60.

Kiebzak S, Rafert G, Tucker CE (2016) The effect of patent litigation and patent assertion entities on entrepreneurial activity. *Res. Policy* 45(1):218–231.

Kim I, Skinner DJ (2012) Measuring securities litigation risk. *J. Accounting Econom.* 53(1):290–310.

King RG, Levine R (1993) Finance and growth: Schumpeter might be right. *Quart. J. Econom.* 108(3):717–737.

Kogan L, Papanikolaou D, Seru A, Stoffman N (2012) Technological innovation, resource allocation, and growth. *Quart. J. Econom.* 132(2):665–712.

Kortum S, Lerner J (2000) Assessing the contribution of venture capital to innovation. *RAND J. Econom.* 31(4):674–692.

Krieger JL, Li D, Papanikolaou D (2018) Missing novelty in drug development. Technical report, National Bureau of Economic Research, Cambridge, MA.

La Porta R, Lopez-de Silanes F, Shleifer A, Vishny RW (1997) Legal determinants of external finance. *J. Finance* 52(3): 1131–1150.

Lanjouw J, Schankerman M (2001) Characteristics of patent litigation: A window on competition. *RAND J. Econom.* 32(1): 129–151.

Lanjouw JO, Lerner J (1998) The enforcement of intellectual property rights: A survey of the empirical literature. *Ann. Econom. Statist.* (49/50):223–246.

Lemley MA, Shapiro C (2005) Probabilistic patents. *J. Econom. Perspect.* 19(2):75–98.

Lemley MA, Shapiro C (2006) Patent holdup and royalty stacking. *Texas Law Rev.* 85:2163.

Lerner J (2002) Patent protection and innovation over 150 years. Technical report, National Bureau of Economic Research, Cambridge, MA.

Lerner J (2009) The empirical impact of intellectual property rights on innovation: Puzzles and clues. *Amer. Econom. Rev.* 99(2):343–348.

Lerner J (2010) The litigation of financial innovations. *J. Law Econom.* 53:807–831.

Lerner J, Schoar A (2005) Does legal enforcement affect financial transactions? The contractual channel in private equity. *Quart. J. Econom.* 120(1):223–246.

Lerner J, Seru A (2017) The use and misuse of patent data: Issues for corporate finance and beyond. Harvard Business School Working Paper 18-042, Harvard Business School, Boston.

Li G-C, Lai R, D'Amour A, Doolin DM, Sun Y, Torvik VI, Amy ZY, Fleming L (2014) Disambiguation and co-authorship networks of the US patent inventor database (1975–2010). *Res. Policy* 43(6), 941–955.

Lin C, Liu S, Manso G (2020) Shareholder litigation and corporate innovation. Management Science Working paper, University of Hong Kong, Hong Kong.

Mann W (2018) Creditor rights and innovation: Evidence from patent collateral. *J. Financial Econom.* 130(1):25–47.

Mezzanotti F, Simcoe T (2019) Patent policy and American innovation after eBay: An empirical examination. *Res. Policy* 48(5): 1271–1281.

Moser P (2005) How do patent laws influence innovation? Evidence from nineteenth-century world's fairs. *Amer. Econom. Rev.* 95(4): 1214–1236.

Moser P (2013) Patents and innovation: Evidence from economic history. *J. Econom. Perspect.* 27(1):23–44.

Moshirian F, Tian X, Zhang B, Zhang W (2020) Stock market liberalization and innovation. *J. Financial Econom.*, epub ahead of print August 23, https://doi.org/10.1016/j.jfineco.2020.08.018.

Murray F, Stern S (2007) Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anticommons hypothesis. *J. Econom. Behav. Organ.* 63(4):648–687.

Pakes A (1986) Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* 54(4):755–784.

Peters RH, Taylor LA (2017) Intangible capital and the investment-q relation. *J. Financial Econom.* 123(2):251–272.

Ponticelli J (2016) Court enforcement bank loans and firm investment: Evidence from a bankruptcy reform in Brazil (with Leonardo S. Alencar). *Quart. J. Econom.* 131(3):1365–1413.

Rogers JL, Van Buskirk A (2009) Shareholder litigation and changes in disclosure behavior. *J. Accounting Econom.* 47(1):136–156.

Sakakibara M, Branstetter L (2001) Do stronger patents induce more innovation? Evidence from the 1988 Japanese patent law reforms. *RAND J. Econom.* 32(1):77–100.

Scotchmer S (1991) Standing on the shoulders of giants: Cumulative research and the patent law. *J. Econom. Perspect.* 5(1):29–41.

Seaman CB (2016) Permanent injunctions in patent litigation after eBay: An empirical study. *Iowa Law Rev.* 101(5):1949–2019.

Shapiro C (2010) Injunctions, hold-up, and patent royalties. *Amer. Law Econom. Rev.* 12(2):280–318.

Shapiro C (2016a) Patent remedies. *Amer. Econom. Rev.* 106(5): 198–202.

Shapiro C (2016b) Property rules vs. liability rules for patent infringement. Working paper, University of California Berkeley, Berkeley, CA.

Smeets R (2014) Does patent litigation reduce corporate R&D? An analysis of US public firms. Working paper, Rutgers, The State University of New Jersey, Newark, NJ.

Srinivasan S (2018) Patents v. innovation: Evidence from public firms. Working paper, Northwestern University, Evanston, IL.

Tang YH (2006) The future of patent enforcement after eBay v. MercExchange. *Harvard J. Law Tech. 20*(1): 235–252.

Tucker C (2014) Patent trolls and technology diffusion: The case of medical imaging. Ewing Marion Kauffman Foundation Research Paper Series, Kansas City, MO.

Venkatesan J (2009) Compulsory licensing on nonpracticing patentees after eBay v. MercExchange. *Virginia J. Law Tech.* 14:26.

Wesenberg E, O'Rourke P (2006) The toll on the troll. *New Jersey Law J.* (July 24), https://www.law.com/njlawjournal/almID/900005458488/?slreturn=20210012145850.

Williams HL (2016) Intellectual property rights and innovation: Evidence from healthcare markets. *Innovation Policy Econom.* 16:53–87.