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

Filippo Mezzanotti*

First Draft: January 2015

This Draft: January 30, 2017

Abstract

The recent spike in patent litigation raises concerns about the ability of the current intellectual property system to effectively promote innovation. Using a difference-in-difference design around the 2006 Supreme Court decision "eBay vs. MercExchange," I examine how patent enforcement can reduce the negative effects of litigation on firms' innovation. This ruling was intended to curb abusive patent lawsuits by providing more flexibility in the way courts remedy patent violations. I estimate the causal impact of the decision by comparing firms that were differentially affected by the shock, measured by exogenous firm-level exposure to patent litigation before 2006. Across a large sample of innovative firms, the decision led to an increase in the quality and quantity of patenting and, for public firms, in R&D investment. In terms of channels, I show that patent litigation reduces investment in innovation by lowering the returns from R&D and by exacerbating financing constraints. This result confirms that patent litigation plays an important role in hindering innovation, and more generally that legal risk has real impact on firms' activity.

JEL CODES: G30, G38, K11, O34

*Kellogg School of Management, Northwestern University. I am indebted to my advisors Josh Lerner, David Scharfstein, Andrei Shleifer, and Jeremy Stein, for their guidance and encouragement. I also thank Lauren Cohen, Jiashuo Feng, Samuel G. Hanson, Victoria Ivashina, Xavier Jaravel, Louis Kaplow, James Lee, Simone Lenzu, Christopher Malloy, Diana Moreira, Gianpaolo Parise, Andrea Passalacqua, Tom Powers, Giovanni Reggiani, Jonathan Rhinesmith, Stanislav Sokolinski, Adi Sunderam, Haris Tabakovic, Edoardo Teso, John Zhou and the seminar participants at Harvard University, Columbia University, Yale, Cornell, University of Rochester, FED Board, Texas A&M, Bocconi, Cornerstone for helpful comments. I thank the Harvard University, Harvard Business School and Kellogg School of Management for financial support. All errors are my own. Email: filippo.mezzanotti@kellogg.northwestern.edu.

1 Introduction

The main goal of the patent system is to protect intellectual property and thus to spur innovation and 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. Over the last thirty years, lawsuits involving patents more than tripled (Figure 1) and their estimated cost surpassed $300 billion (Bessen et al., 2015).[1] Furthermore, a large share of this increase can be explained by a surge in lawsuits involving patent assertion entities, also known colloquially as "patent trolls" (Cohen et al., 2014). This rise in litigation may reduce the incentives of firms to invest in R&D, and therefore curb innovation and growth (Bessen and Meurer, 2008b; Boldrin and Levine, 2002; Jaffe and Lerner, 2011).

In this paper, I show that patent litigation has a real impact on innovation, and that improvements in patent enforcement can have positive effects on corporate R&D. To examine this issue, I develop a new research design that exploits a landmark legal decision, the 2006 Supreme Court decision "eBay vs. MercExchange." The ruling was intended to curb abusive patent lawsuits by providing more flexibility in the way courts remedy patent violations. In particular, this decision ended the practice of granting a permanent injunction almost automatically after a patent infringement, giving courts more power to decide when an injunction is appropriate. An injunctive order forces firms to shut down any operation related to the violated technology regardless of the nature and magnitude of the infringement, and therefore poses a major risk for companies accused of violating a patent.

Previous research has suggested that injunctions play a central role in fostering abusive litigation. Specifically, a "near-mandatory" injunction can increase the extent to which companies can be held-up by a plaintiff, even when lawsuits are based on frivolous claims or minor violations (Lemley and Shapiro, 2006; Shapiro, 2010, 2016a). For instance, patent assertion entities frequently used the threat of injunction "as a bargaining tool to charge exorbitant fees to companies that seek to buy licenses to practice the patent" (Court, 2006). Moreover, this issue is exacerbated by the large legal uncertainty about patent validity and infringement that characterizes intellectual property (Lemley and Shapiro, 2005).

A prominent example of these practices is the lawsuit in which Research in Motion (RIM), BlackBerry's producer, was found to have violated five patents owned by NTP. During the court trial, RIM unsuccessfully tried to prove that NTP patents were not valid, and, after the verdict, requested a re-examination of the patents by the U.S. Patent Office (USPTO). At that point however, NTP had the power to request an injunction, which would have led to a shutdown of the BlackBerry system. RIM was therefore forced to pay a

[1] This estimate refers only to public firms sued by Non-practising entities and it is constructed using an event study methodology.

record settlement of \$612.5 million to NTP before the re-examination was finished in order to avoid the high risk of receiving an injunction. Unfortunately, a few years after the settlement, the USPTO deemed invalid most of the claims contained in the NTP patents. Many similar cases can be identified in the literature (Jaffe and Lerner, 2011), demonstrating how injunction in intellectual property disputes can come at a large cost for firms.

This Supreme Court decision had a first-order impact on the way intellectual property is enforced.[2] For instance, Chien and Lemley (2012) found that the likelihood of obtaining an injunction decreased by at least 25%, and this drop was larger for plaintiffs more likely to be involved in abusive lawsuits like patent assertion entities (PAE). These companies, known for their aggressive patent assertion strategies, are accused to act like "patent trolls" exploiting the injunction threat in litigation.[3] Confirming the importance of the ruling, I identify a set of potential PAE that are public by looking at non-practicing entities. I find that these firms experienced large negative returns around the time of the decision, with average cumulative returns of about -10%.

However, the overall effect of the decision on innovation may be ambiguous. On the one hand, the reduction in injunction rate may suggest that the new policy was successful in fixing part of the distortions of the patent system, making investments in innovation easier and more valuable. This view was shared by many scholars and practitioners. According to the American Innovators Alliance, an association representing some 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."[4] On the other hand, a reduction in injunction may instead lead to lower deterrence, therefore encouraging more violations and lowering the incentive to invest in this area (e.g. Epstein, 2008). This effect is potentially small, because injunction rates remained high and the decline was mostly concentrated in cases more likely related to abusive litigation (Seaman, 2016). However, the actual size of this channel is not determined.

Given this ambiguity, I turn to data to examine the real impact of the decision on firms. This exercise has two objectives: first, it can help us to evaluate the quality and effects of this change, which has important implications for policy. Second, this test can also provide broader insights on the role of intellectual property litigation in shaping firms' incentives to innovate. In particular, this analysis can be used to explore the channels through which this legal risk affects corporate activity.

I estimate the causal impact of the decision using a difference-in-difference design that exploits hetero-

[2] This is confirmed by a large literature in law, e.g. Bessen and Meurer (2008a); Holte (2015); Shapiro (2010); Tang (2006); Venkatesan (2009).

[3] As discussed later in the paper, Gupta and Kesan (2015); Seaman (2016); Grumbles III et al. (2009) find results consistent with Chien and Lemley (2012).

[4] American Innovators Alliance represents large tech firms, such as Microsoft, Micron, Oracle and Intel. The sentences are taken from the "amicus curiae" submitted for the Supreme Court case.

geneity in the intensity of the treatment. In particular, I use variation in firm exposure to patent litigation in 2006 to identify companies that are more likely to be affected by the decision. The intuition for this is simple: while 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. Indeed, patent lawsuits are not equally spread across industries but tend to be more prevalent in certain areas. This concentration, which is persistent over time, reflects in part the focus of patent assertion entities in some specific technologies (Feng and Jaravel, 2015).

In particular, my measure of litigation exposure captures variation in the risk of litigation that exogenously affects a company because it innovates in specific technology fields. Looking across more than four-hundred technology classes defined by the U.S. patent office, I identify the fields where each company operates. Moreover, I create a measure of patent litigation intensity for each of these fields using litigation data from WestLaw (Thomson Reuters). The final firm-level measure of exposure to patent litigation is simply a weighted-average of litigation intensity across all the technology classes, where the weights are the share of patents developed in each class by the company.

Implementing this estimator on a sample of almost twenty thousand innovative firms, I find that the ruling had a positive effect on innovation. Firms that were more exposed to litigation before the decision increased patenting more after the decision. These effects are not only statistically significant, but also economically relevant. For example, examining two firms one standard deviation apart in exposure to litigation, I find that the more exposed company increased patent applications 3% more than the less exposed one, which on average corresponds to almost one extra patent in the two years after the shock. Similarly, the same firm was 2% more likely to patent something, which is equivalent to a 5% jump in the probability of patenting in the sample.

These results are robust to several tests. First, I provide evidence in favor of the parallel trend assumption by showing that differential exposure to litigation does not predict differential behavior before the decision, both looking at measures of quantity and quality of innovation. Second, I implement a battery of placebo tests and a randomized permutation test (Chetty et al. 2009) to further reject that my results could be capturing other spurious factors unrelated to litigation exposure. Furthermore, results are not driven by other sources of heterogeneity that may affect innovation and might be correlated with litigation exposure. In particular, I construct an industry classification based on the major technology of operation (Hall et al., 2001), and I show that the results are not driven by industry effects. Similarly, results are stable when controlling for the location of the firm, the size of the pre-shock patent portfolio, the average quality of the output before the decision or the young age of the firm. Finally, I closely replicate the results with alternative measures of litigation exposure.

Next, I find that the decision also positively affected the quality of innovation. I find that after the decision firms are more likely to develop a potential "breakthrough innovation" (Kerr, 2010), defined as a patent that is at the top of the citation distribution within the same patent class and year group. These results may suggest that better enforcement made firms more prone to take on riskier projects. Since returns from innovation tend to be skewed (Pakes, 1986), this should have positive effects on the ability of a firm to grow and compete. Moreover, patenting behavior also increased when considering the number of patents weighted by citations received, which excludes that the increase in quantity came at a detriment to quality.

A limitation of my main sample is that I observe only patenting instead of the actual investment in R&D. Without further information, it would be challenging to understand whether these results stem from a true change in innovation or instead come from shifts in the incentives to patent. To rule out this alternative explanation, I focus on a sub-sample of public firms that were active in innovation around the decision. Using the same methodology as before, I find that the ruling also increased R&D intensity. When comparing two firms that are one standard deviation apart in terms of exposure to patent litigation, I find that the more exposed firm experienced a relative increase of 8% in R&D over assets.

In line with the intentions of the Supreme Court, the ruling successfully improved companies' ability to innovate. Following the decision, firms increased their investments in R&D and their patenting. Furthermore, the shock positively impacted the quality of output. These results suggest that better patent enforcement can increase the certainty with which intellectual property are protected and reduce the risks of hazardous or abusive litigation, thereby increasing the incentive of firms to invest in R&D. Therefore, this evidence confirms that patent litigation can significantly hinder firms' ability to innovate. The large size of my effects confirms that these distortions can be substantial.

Finally, I examine the mechanisms through which patent litigation affects corporate innovation. Intuitively, intellectual property litigation can affect the incentives to invest by reducing the returns to innovation. The prospect of future litigation lowers the NPV of potential investments, as firms internalize both the monetary and non-monetary costs that this would entail. Clearly, this decline in profitability will be larger for projects in highly litigated areas (Lerner, 1995). I explore the within-firm propensity to work in projects with high risk of litigation to provide evidence in line with this idea. I show that, after the decision, firms reshuffled their internal resources towards projects in higher litigation areas. This effect is driven by firms entering in new technology fields where litigation risks are high.

I also find that patent litigation affects R&D because it exacerbates the financing problems of innovation (Brown et al., 2009; Hall and Lerner, 2010). Companies operating in high litigation environments may be forced to devote a larger share of resources to monitoring and defensive activities (Cohen et al., 2014). Similarly, they are more likely to pay higher settlements and licensing.[5] In the presence of financial frictions, the increase in costs due to patent litigation reduces the amount of internal resources available and therefore forces firms to cut down on investments. Consistent with this implication, I find that firms likely to be financially constrained before the decision increased R&D intensity more in its aftermath. These findings establish the important role played by financial constraints in explaining the negative effects of patent litigation.

A large literature in finance suggests that more and better innovation can increase firms' valuation (Kogan et al., 2012). In line with this research, I find that the decision had a positive effect on stock market returns. By looking at abnormal returns on the day that the decision was made public, I demonstrate that firms that are characterized by high exposure had larger returns, and that this effect does not disappear over the following days. For instance, looking across the top quartile of treatment, I find that value-weighted returns of the two groups are identical before the announcement, but that the more exposed group outperformed the other group by 80 bps on the day of the announcement.

By showing that changes in patent enforcement can have sizable effects on corporate innovation, this paper contributes to the literature that examines how property rights and legal institutions shape economic incentives (Acharya et al., 2011; Claessens and Laeven, 2003; Demirgüç-Kunt and Maksimovic, 1998; King and Levine, 1993; La Porta et al., 1997; Lerner and Schoar, 2005).[6] Previous research has demonstrated that secure property rights favor a more efficient allocation of resources and foster growth, but in many cases good enforcement is as important as good rules in determining economic outcomes (Djankov et al., 2003; Iverson, 2014; Ponticelli, 2013). 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 lawsuits are therefore frequent (Lanjouw and Lerner, 1998). My paper confirms that enforcement can be very important and suggests that, in a manner similar to other interventions (Acharya and Subramanian, 2009; Mann, 2013), a fine-tuning of patent law can have substantial effects on fostering corporate innovation.

My analysis also provides new evidence about the real costs of patent litigation, which is central in today's policy debate (White House 2013). While the idea that litigation could harm innovation is commonly accepted, direct evidence that supports this claim is relatively sparse. In this direction, Smeets (2014) shows that firms decrease R&D intensity after being litigated. My results are consistent with his work and extend his idea by showing that high litigation may harm innovation even among companies that do not directly engage in it. In addition, Tucker (2015) shows that a high level of litigation at the industry level may reduce VC investments. Furthermore, my work provides new insights on the operation of non-practicing entities

[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] This topic is related to the finance literature focusing on the relationship between litigation and corporate policies (Arena and Julio, 2014; Beatty et al., 2008; Kim and Skinner, 2012; Haslem, 2005; Rogers and Van Buskirk, 2009).

and contributes to the growing literature on this topic (Cohen et al., 2014, Feng and Jaravel, 2015; Tucker, 2014).

This paper also adds to the body of empirical work that analyzes the relationship between property rights and innovation (e.g. Galasso and Schankerman, 2015; Lerner, 2009; Sakakibara and Branstetter, 2001; Williams, 2015). In particular, I show that the decision "eBay vs. MercExchange" which can be considered a restriction of certain aspects of intellectual property rights,[7] had beneficial effects on innovation incentives.[8] Historically, automatic injunction was introduced in patent cases because intellectual property law was derived as an extension of standard property law. This strict property rule works well only when ownership rights are clear and easy to identify, like for tangible assets (Calabresi and Melamed, 1972). Instead, patents are different because the exact boundaries of these assets are very hard to define. In such cases, a strict property rule may favor extractive behavior by opportunistic parties rather than incentivize investments, and a hybrid system that provides more flexibility may be superior (Kaplow and Shavell, 1996). Overall, this paper argues that patents are different from other assets, and therefore the design of patent enforcement should take into account these differences (Schwartzstein and Shleifer 2013).

The paper is organized as follows. In Section (2), I provide more background information about the Supreme Court decision, also discussing its potential effects on corporate innovation. In Section (3), I present the data used in the paper and I discuss in detail how I construct my measure of exposure to patent litigation at firm level. In Section (4), I present the main results of my analysis. In Section (5), I discuss and test different channels through which patent litigation can affect innovation. In Section (6) I look at the stock market reaction around the decision. Lastly, Section (7) discusses policy implications and avenues for future research.

2 The "eBay vs. MercExchange" case

This section provides background information on the Supreme Court decision "eBay vs. MercExchange" and its consequences. First, I analyze the importance of injunction on the pre eBay world, in particular in relationship with litigation. Second, we discuss how the ruling could affect innovation, therefore setting the foundation for our hypothesis and research design. Lastly, I provide some preliminary and novel evidence of the importance of the ruling for patent enforcement. In particular, I show that the decision had negative effects on the stock market valuation of patent assertion entities.

[7] During the argument, Justice Scalia said "We're talking about a property right here, and a property right is the exclusive right to exclude others."

[8] However, this is not in contrast with Williams (2015). In fact, most of the previous literature studies strength of property rights looking at length and breadth of patents coverage, rather than the right to exclude after court.

2.1 The role of injunction and the 2006 decision

With the 2006 "eBay vs. MercExchange" decision, the Supreme Court revisited the norms regulating the issuance of permanent injunction in cases involving intellectual property.[9] Injunction is a legal remedy that can be requested by a plaintiff after a violation. If granted by a court, injunction forces the infringer to stop using any technology covered by the contested patents, irrespective of the magnitude of the infringement. Historically, it has been extensively used in patent cases (Chien and Lemley, 2012). Before 2006, a plaintiff that was able to prove a violation had the automatic right of obtaining an injunction. In other words, the norm was that "a permanent injunction should be issued when infringement was proven" (Court, 2006). Exceptions to this rule were quite uncommon and mostly due to reasons of public interest.

Injunction is one of the major risks faced by companies active in intellectual property. Since technologies tend to be highly complementary, an injunction granted for a relatively small violation can deeply impair a company's operations. A company receiving an injunction may be forced to shut down a line of business or change the way it produces or markets a product. Furthermore, the uncertainty characterizing the patent system complicates things even more. Since the boundaries of intellectual property are generally unclear (Lemley and Shapiro, 2005), innovative firms may involuntarily infringe existing patents. This can be particularly problematic because in some cases the patent office may issue dubious or overlapping patents, where the same technology is assigned to multiple companies in different times. Because of this uncertainty, companies have a hard time navigating the patent system and predicting the outcomes of patent disputes, making injunction even more salient.

These two factors are very clear in the RIM vs. NTP case discussed in the introduction. First, although the dispute involved only a few patents, issuing an injunction would have forced the RIM to shutdown the whole Blackberry system. This is because patents are highly complementary, and therefore blocking only a few of them may limit the use of the whole patent portfolio. Second, RIM was forced to settle despite the fact that most of NTP claims were eventually found to be invalid. This was impossible to prove in court and this evidence was confirmed only after a lengthy re-examination process that lasted several years. Altogether, NTP's ability to leverage on the near-mandatory injunction was the main driver to obtain the large settlement.[10]

As illustrated in the RIM case, an automatic injunction can foster abusive lawsuits that use the threat of litigation as a way to make money out of innovative firms (Shapiro, 2016b). Consistent with this idea, patent assertion entities (PAE) were actively using the threat of permanent injunction as a way to scare

[9] I provide some background legal information about the "eBay vs. MercExchange" case in Appendix (A.2).

[10] In an interview for the National Law Journal (March 13th 2006, Volume 27, Issue 77), patent litigator David Clonts of Akin Gump Strauss Hauer & Feld's, states 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."

counterparties and therefore obtain larger settlements (Lemley and Shapiro, 2006). Patent assertion entities are companies that hold large patent portfolios and are extremely active in intellectual property litigation and they are accused of aggressively asserting patents with the objective of obtaining large settlements.[11] In these cases, since patent uncertainty is high and the downside costs are large, the threat of injunction can loom large and have important consequences even when lawsuits are based on minor violations.

This idea was recognized by the Supreme Court and well understood by practitioners. When expressing his motivation for the decision, Justice Kennedy argued that 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." Similarly, in an analysis of the case Wesenberg and O'Rourke (2006) reports:

> "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."

The ruling "eBay vs. MercExchange" which was made public on May 15th, 2006, dramatically changed this landscape by granting courts more discretion to decide when it is appropriate to issue an injunction. In particular, the decision stated clearly that the issuance of injunction should not happen automatically. Instead, courts should decide on a case-by-case basis, balancing "the hardships between plaintiff and defendant" (Court, 2006). In other words, the Court recognizes that an automatic injunction can actually be less efficient than a hybrid system, where monetary damages can be sometimes sufficient to remedy a violation.[12] The content of the ruling was generally not anticipated (see Appendix A.3) and it had a deep impact on the way companies think about intellectual property protection. In the next section, I discuss previous evidence on the importance of the law and I also examine how we expect the ruling to affect innovation by firms.

2.2 Hypothesis Development

The intent of the Supreme Court was to introduce new norms that would improve the status of patent enforcement, by curbing the cost of abusive lawsuits and reducing the overall uncertainty in the patent system. However, simple theory has ambiguous predictions over the effectiveness of this intervention.

[11] While this is not a completely new phenomenon (Moser, 2013), it has increased dramatically over the last twenty years and has recently attracted the attention of policy makers and academics (Feng and Jaravel, 2015; Tucker, 2015, 2014).

[12] In other words, "damages award is sometimes sufficient to maintain incentives while preventing patentees from amassing disproportionate rewards, significantly injuring the public, and stifling innovation" (Carrier, 2011).

On the one hand, the presence of automatic injunction can create distortions for firms that are active in R&D, increasing the risks for companies exposed to patent litigation and forcing them to invest resources for defensive reasons. Therefore, we may expect this decision to positively impact innovation by U.S. corporations. From this perspective, removing automatic injunction should lead to better enforcement and fewer risks in the intellectual property market, which can increase the attractiveness of investing in this area. Furthermore, reducing the cost of opportunistic litigation could free-up resources that could be employed for innovative activities. This positive view on eBay was widely held by many academics (Bessen and Meurer, 2008a) and industry experts. For instance, according to the Computer & Communication Industry Association, automatic injunction did "produce anti-competitive behavior, foster more litigation, and undermine innovation."[13] The motivation for these effects were well explained by American Innovators Alliance, another association representing the interests of large high-tech companies, which claimed that because of 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."[14]

On the other hand, however, some scholars have criticized this decision, arguing that the ruling could end up negatively affecting innovation (Epstein, 2008; Holte, 2015). Their idea was that, if courts are unable to identify abusive litigation, eliminating automatic injunction just lowers deterrence against real violations. As a result, more firms may decide to infringe because they anticipate that injunction may not be issued if they get caught, and this lowers returns for innovative firms and reduces the overall incentives to engage in R&D. For this concern to be relevant, we need that post-eBay firms experiencing real violations would have a much harder time to obtain an injunction.

Recent research in this area can guide us in interpreting the importance of these channels. First of all, various papers have confirmed that the decision had a large impact on the way intellectual property is enforced (Bessen and Meurer 2008a; Shapiro 2010, 2016a; Tang 2006; Venkatesan 2009). Empirically, Chien and Lemley (2012) find that the ruling reduced the likelihood of obtaining an injunction by about 25% [15], while Gupta and Kesan (2015) finds that the ruling also decreased the rate at which injunction is sought. Therefore, both qualitative and quantitative evidence suggests that the decision had a deep impact on this area and this effectiveness was clear already in the few months after the ruling. However, this body of evidence also confirms that injunction is still a valuable tool that companies can use against real violations. First, despite the decline in court acceptance rate, injunction is still granted in the majority of cases.[16] Second, the

[13] The quote is from the "amicus curiae" submitted by the Computer & Communication Industry Association (CCIA) for the Supreme Court case. The CCIA is a Washington based advocacy organization that represents the interests of the computer, internet and information technology industry.

[14] 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.

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

[16] Looking at the results from Chien and Lemley (2012), the rate at which injunction is granted goes from almost 100% to

drop in injunction rate was mostly driven by cases where the plaintiff is more likely to be involved in abusive litigation (Seaman, 2016), like cases initiated by parties that are not competitors of the plaintiff.

Altogether, previous literature suggests that the ruling effectively reduced litigation risks - in particular for the most problematic cases - without removing the protection provided by injunction when a real violation occurs. This evidence reinforces the positive interpretation of the eBay decision, but it cannot shed conclusive lights on the real effects of the Supreme Court decision. Therefore, I turn to data to determine the actual overall effects of the Supreme Court decision on innovation. In particular - as I discuss in the next sections - I examine how R&D activity by innovative firms reacted to the decision, by exploiting heterogeneity across companies in the intensity of the treatment. This analysis has two objectives: first, this exercise can help us to evaluate the quality of this policy change. Second, this test can also provide broader insights on the role of intellectual property litigation in shaping firms' incentives to innovate. In fact, the ruling should only positively affect companies if litigation at 2006 was an important obstacle for companies interested in investing in innovation. Building on this intuition, my analysis can be used to quantify this distortion and explore the mechanisms through which this operates.

2.3 The economic importance of the decision: the case of NPE

In his concurring opinion, Justice Kennedy identified patent assertion entities (PAE) as one of the main players taking advantage of almost-automatic injunction policy (Court, 2006). In this section, I provide further evidence on the importance of the decision by studying the stock market returns of a set of public patent assertion entities. In particular, in line with previous literature, I identify PAE by looking at non-practicing entities. Consistent with the importance of the decision, I find that the ruling led to a drop of about 10% in the stock price of these companies.

In general, it is complicated to identify patent assertion entities in the data. One approach taken in the literature (Cohen et al., 2014; Feng and Jaravel, 2015; Tucker, 2015, 2014) has been to identify PAE by looking at non-practicing entities (NPE). As the name suggests, these company generate most of their revenue by licensing and settlement fees rather than from manufacturing, and therefore they are more likely to aggressively assert patents in courts. Clearly, not every NPE can be accused to act like a "patent troll." For instance, universities and other research institutions are categorized in this way. By the same token, not all the abusive behavior is specific of NPEs.

NPEs are a useful laboratory to test whether the decision had a first order impact on the enforcement of patents. Previous research has confirmed NPEs extensively used injunction threats when bargaining licensing agreements or settlements (Chien and Lemley, 2012). Furthermore, the elimination of automatic injunction

around 70%.

is unambiguously a bad news for these firms. First, automatic injunction reinforces the bargaining position of patent holders and therefore it is advantageous for NPEs when they negotiate the license of one of their patents. Second, differently from other companies, automatic injunction does not constitute a major risk for these firms because they generally do not directly engage in manufacturing.

Therefore, if the ruling had a big impact on patent enforcement, I expect NPE to be negatively affected by the decision. In particular, I test whether the market value of public NPEs declined around the ruling of the Supreme Court. The main challenge in this type of analysis is that most NPEs are private. For instance, "Intellectual Ventures" - allegedly the largest NPE nowadays - 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.[17] Then, I identify the firms in these lists that have returns information in CRSP around the date of the event. This analysis yields a final list of ten companies.[18]

Studying the returns of these companies around the decision, I identify four important stylized facts.[19] 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 3).[20] While the largest one-day drop is experienced the day of the Supreme Court decision, stocks also lost value also the three days before it. Examining the abnormal returns with respect to the S&P500, the firms lost 6.3% ($t = -4.53$) on average the week before the ruling. One explanation for this result is that investors, anticipating the arrival of news regarding the case, have started to require a premium to hold these stocks 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.2). Finally, these negative effects do not revert back in the days following the decision.[21]

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 was not reverted in the following

[17]The first providers publish a list of top NPEs active in USA at 2014 (https://www.patentfreedom.com/about-npes/holdings/), where companies are selected based on number of patents held. The second instead published a study on stock returns on NPEs in 2013, where they used 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/).

[18]The majority of the companies appear in both list - 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, Spherix.

[19]More information about the analysis can be found in Appendix (A.4.3). One caveat of the data set is that it is compiled based on a recent list; therefore I may miss a NPE that was active and public in 2006, but defunct today. While I cannot exclude this, I could not find any example of this phenomenon in the data.

[20]In Figure (A.4) I replicate the same results under alternative models as robustness.

[21]These results are qualitative identical when I use value-weighted measures.

weeks. The results are robust to the removal of each of the of the NPEs considered in the sample.[22] 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

3.1 Firm level data

To estimate the impact of the "eBay vs. 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 look at patenting behavior of companies as a proxy of innovation. This allows me to measure innovation for a large sample of both public and private companies. The data comes from the Fung Institute (University of California at Berkeley) patent data set,[23] which is an updated version of the Harvard Business School Patent Network Database (Li et al., 2014) extensively used in literature. This data contains full information on all patents granted between 1975 and 2014,[24] with a new disambiguate assignee identifier that I use to identify a firm across different patents. Overall, the first part of the paper focuses on a sample of more than 16 thousand firms that are active in patenting around the decision, meaning that they patent at least once both before and after the decision. However, I also examine an alternative and more extensive sample to gauge the robustness of the results.

In the second part of this work, I 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. Since patent numbers are easy to match, this approach greatly reduces the probability of errors and missing information. After applying the standard filters,[25] I have a sample of more than one thousand public companies that are active in innovation around the decision and with R&D information at quarterly level. Lastly, I match these firms to CRSP using the standard Compustat-CRSP bridge file. In the Appendix (A.4) I provide more details on the data construction and matching.

The main measure of innovation activity employed in the paper is based on the simple count of granted

[22]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 any case, the result is 1% significant.

[23]Data can be found: http://funginstitute.berkeley.edu/tools-and-data.

[24]The bulk of my analysis is run with applications made by the end of 2008, therefore allowing more than the five years recommended by Dass et al. (2015) to eliminate risk of truncation bias..

[25]I consider firms in non-financial and non-regulated industries, headquartered in USA, not involved in financial restructuring and with information reported in the quarterly Compustat data. More details in the Appendix Appendix (A.4).

patents applied by a firm in a specific period.[26] I focus on application because they are closer to the time of the actual investment. 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. R&D expenses are adjusted to take into account the acquisition of in process R&D during the quarter (Mann, 2013). In the end, patent data are also used to construct a variety of citations-based measures of patent quality.

Furthermore, I use patent data to generate firm-level control variables. For every firm, both public and private, I construct an industry classification based on the major (large) technology class in which the firm patents in the four-year around the decision (Hall et al., 2001). I use addresses reported in 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, 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 to roughly 31 patents in the two-year window considered. These numbers are large and 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 decision. Looking at citations, they receive an average of one citation per patent, where number of citations is adjusted for technology-class and year. As expected, public firms appear to patent more than the average firm in the full data set and they have on average quarterly R&D expenses of roughly 3% of their assets.

3.2 Measuring exposure of litigation at firm level

A crucial component of my identification strategy relies on measuring firm exposure to patent litigation. While true litigation risk is unobservable, I can use heterogeneity in the intensity of patent litigation across different technology fields to construct a firm-level measure of patent litigation. In this section, I first discuss in detail how I construct this measure, using data from patent lawsuits and patent applications. Second, I highlight the advantages of this approach and discuss the possible shortcomings.

Intuitively, a firm is more exposed to patent litigation if its R&D is focused in technology fields where patent litigation is more intense. For instance, companies that operate in software or drugs, where intellectual property lawsuits are more frequent, will be more concerned with patent assertion than companies doing mechanical research, where instead litigation is much less intense. This approach takes advantage of two main 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

[26] If patents are assigned to more than one assignee, then I equally divide the patent count across firms.

in different technology fields.

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 $t(j) = [\sigma_i^j]_{i=1}^T$; (2) the distribution of patent litigation risk 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. Clearly, in this case I would have that σ_i^j is between zero and one and that $\sum_{i=1}^T \sigma_i^j = 1$.

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:

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

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

While the variable $Exposure_j$ is intrinsically unobservable, its components - $t(j)$ and p - can be constructed from data. First, I use patent data to measure $t(j)$, the technology space where the company operates. I identify different technology fields using the U.S. Patent Office (USPTO) classification in technology classes. In particular, the USPTO categorizes each patent across more than 400 technology classes, which provide a very precise and narrow definition of technology. Then, for each firm, I define σ_i^j as the share of granted patents of firm j in the technology class i that were applied before 2006.[27]

Second, I estimate the distribution of patent litigation across technology fields - the vector p - using litigation data from WestLaw, a subsidiary of Thomson Reuters. Westlaw is one of the primary provider of legal data in United States and use public records to develop a complete overview of lawsuits in United States. The same data, also known as Derwent LitAlert data, were previously used by other empirical work on patent litigation (e.g. Lerner, 2006, Lanjouw and Schankerman, 2001). Using the online tool LitAlert,[28] I searched for all the litigation involving patents between 1980 and 2006.

From each filing, I extract all the patents that were asserted by the plaintiff and then use this information to construct a proxy for p. After cleaning the raw data, I have more than thirty thousand cases filed until 2006. In line with the previous literature, the number of cases increases over time (Figure 1) and more than tripled between the beginning of the 1980s and the most recent data. Then, I use an approach similar to Tucker (2015) to adjust the data and make cases comparable across filings. First of all, each filing may contain multiple defendants. Firm A suing both firm B and C in the same filing should have more weight

[27]For instance, if a company operates in four technology classes with 2 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 technologies where the company patented something.

[28]http://intranetsolutions.westlaw.com/practicepages/template/ip__litalert.asp?rs=IPP2.0&vr=1.0

that Firm A suing only firm D. Secondly, 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 deal with this, I reshape the data at single defendant-plaintiff-patent level.

Then, I measure the size of litigation in each of the USPTO technology classes by computing the number of patents in a specific class involved in litigation, scaled by the total number of patents litigated. In other words, my index is the share of total patents litigated within each technology class:

$$p_i = \frac{\sum_{c \in cases} \#Patents_c^i}{\sum_{i \in Tech.Classes} \sum_{c \in cases} \#Patents_c^i} \tag{2}$$

where i defines one of the USPTO technology classes and c is a specific filing.

There are two important features of this measure for my identification. First, patent litigation is not equally spread across technology classes, but rather tend to be more concentrated in some technology classes. Using the index between 1980 and 2006 as a benchmark, I find that the top 50 technology classes in terms of litigation accounts for half of the patent level litigation. Similarly, around 10% of technological classes have no litigation in this period (Table A.1). This heterogeneity gives me the cross-sectional variation that I will exploit in my analysis.

Second, this measure is highly persistent over time. In other words, technology fields where litigation is high in the years immediately before the Supreme Court decisions are also intensively litigated when looking at litigation data over the previous decades. For instance, the score constructed using data between 1980 and 2000 has an 84% correlation with the same score constructed with lawsuits between 2000 and 2006 (Figure A.2). As I discuss more when considering the identification assumptions of my model, this result is reassuring because it suggests that the cross-sectional distribution of patent litigation across technology classes does not simply reflect some heterogeneity in technology shocks in the years before the Supreme Court decision, but rather some structural characteristics of the field.[29]

I estimate $Exposure_j$ combining these two measures as in equation (1). My favorite measure uses litigation data and patents in the five years before Supreme Court decision. This index incorporates the most recent data on both patent litigation and firm activity, therefore reducing potential measurement error. However, for robustness, I also estimate my results using an alternative measure $Exposure_j^{LONG}$, which is constructed using litigation data from 1980 and looking at patents applied over the ten years before the decision. Results are stable across the two indexes.

In the main sample, the average litigation exposure score is 0.76 and the standard deviation is similar (Table 1). Furthermore, the distribution of the score is skewed and there are fewer firms with high scores.

[29]For instance, the activity of patent-assertion entities, which explains a large share of lawsuits before the decision (Cohen et al., 2014), tends to be highly concentrated in specific fields (Feng and Jaravel, 2015).

This is the result of two things. First, some areas, such as “Drug” and “Computer and Communication,” have a larger share of highly exposed firms (Figure 4). Second, even within this major industry there is a relatively large variation in litigation exposure. Consistent with this, later on the paper I show that even adding a full set of industry-time fixed effects my results have similar size and they are still significant.

This way of measuring exposure to patent litigation has two important advantages. First, this score can be constructed for every firm that is active in patenting using existing data and its development is relatively simple, intuitive and transparent. Second, the measure is exogenous to the firm j strategies in litigation. Differently from 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. This is an important since the decision of a firm to engage in litigation may be function of other unobservable firm characteristics, which may then be correlated with investment opportunities.

At the same time, this approach assumes that court litigation is a good proxy of the effective litigation level for a technology class. In principle, this can be problematic because many disputes involving intellectual property do not end up in court, and the decision to file a lawsuit is clearly non-random. However, my approach only requires that the cross-sectional distribution in patent litigation based on lawsuits is representative of the overall, true status of litigation in the intellectual property market. In particular, I do not impose any condition on the homogeneity in the quality of litigation happening inside and outside court. Furthermore, even the presence of some heterogeneity in case selection across technology classes is not a major problem for my identification. As long as this is not systematically correlated with contemporaneous shock to the productive function of innovation, this selection should only result in higher measurement error and therefore lead to a larger attenuation bias.

4 The effect of the Supreme Court decision on innovation

This section contains the main results of the analysis. I start by presenting the empirical framework that I developed to evaluate the effect of the Supreme Court decision on innovation. Then I show that the Supreme Court decision positively affected the ability of companies to patent new technologies. Following this, 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 the effect of the decision on the quality of innovation and on R&D intensity for public firms.

4.1 Empirical framework

The objective of my study is to examine how the Supreme Court decision "eBay vs. MercExchange" affected the innovation output of corporations. In particular, in the first part of the paper I explore measures of innovation based on patent counts. In principle, every firm is affected by the legal change, and therefore there is no natural control group in this experiment. However, the shock should not affect every firm in the same way, and therefore I can exploit variation in the intensity of the treatment to identify the causal effect of the decision.

In particular, I use firm-level variation in exposure to patent litigation before the decision to identify companies that are more or less affected by the decision. As I discussed before, the Supreme Court decision only changed the way injunction is issued in patent lawsuits. Therefore, companies should care about the ruling only to the extent that they are concerned with patent litigation. In fact, a firm that operates in an area where patent lawsuits are rare or inexistent should be unaffected by "eBay vs. MercExchange."

Building on this intuition, I examine how innovation activity changed after the ruling across companies that were differentially exposed to patent litigation. In this framework, firms with little or no exposure to litigation, which supposedly were not affected by the shock, provides a counterfactual for firms that were instead highly exposed to litigation. 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. This means I estimate:

$$y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt} \tag{3}$$

where y_{jt} is an outcome of firm j at time t, $Post = 1\{time > decision\}$, (α_j, α_t) are a set of firm and time fixed effects and $Exposure_j$ is the index of exposure to litigation, as previously discussed. For robustness, I can augment the specification with a matrix of controls X_{jt}. As I discuss later, the controls are a set of firm-level characteristic measured at the time of the decision -and therefore they pre-determied with respect to it - which are interacted with time dummies to allow them to have a differential effects before and after the decision (Angrist and Pischke, 2008; Gormley and Matsa, 2014). Consistent with the previous discussion, $\beta > 0$ confirms that the Supreme Court decision had a positive effect on innovation.

When it is not specified otherwise, I estimate this equation in a four-year window, considering the two years before and after the announcement of the Supreme Court decision on May 15th 2006.[30] Following the literature in applied econometrics (Bertrand et al., 2004), I run my main results collapsing in one observation before and after the decision. This specification provides inference that is robust to concerns of serial

[30]In tables and figures 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 15th. The other quarters are then constructed consistent with this.

correlation in the data. In any case, any analysis in the paper is conducted by clustering standard errors at firm level. Later in the paper, I discuss different robustness to this main set-up.

4.2 The effect of the decision on innovation output

I start by studying how the decision affected the output of innovation, by looking at whether firms more exposed to the shock started applying to more patents after the Supreme Court decision.

In particular, I consider two outcomes. First, I look at $ln(pat_{jt})$, which is the logarithm of the patent applications that firm j filed to during time t (intensive margin). In order to keep the panel balanced, I estimate the model using every firm in the patent data that applied to at least one patent before and after the shock.[31] This corresponds to a sample of slightly more than sixteen thousand firms. Second, I examine an alternative outcome variable: a dummy equal to one when the firm has applied to any granted patent in the period, $1\{Patent_{jt} > 0\}$ (extensive margin). In this case, the sample contains every firm that applied to at least one patent in the five years before the Supreme Court decision. This is a minimal requirement to construct the measure of litigation exposure. As expected, this sample is much larger than the first one, and it contains around seventy-seven thousand firms.

I estimate the simple version of equation (3) and I report these results in columns (1) and (4) of Table (2). I find that firms more affected by the shock applied to more patents than firms that were less affected. This is true looking at the intensive and extensive margin, as previously defined. This suggests that, consistent with the intention of the Supreme Court, the output of innovation increased with the decision. These effects are not only statistically significant, but also economically relevant. Comparing two firms that are one standard deviation apart in exposure to litigation, the more exposed company increased patent applications by 3% more and it was almost 2% more likely to apply to something.

To better appreciate the economic magnitude of these effects, I can compare these reduced form estimates with the baseline patenting activity. A one standard deviation increase in the exposure to the decision leads to an increase by almost one extra patent for innovative firms. Similarly, repeating the same comparison at the extensive margin, the same change increased the likelihood of apply to one granted patent in two years after the ruling by between 2%-5% more, depending on the specification. While the magnitude is not huge - a fact that would probably be concerning in this context -, these results are economically relevant and they confirm that the effects of the decision were sizable.

[31] In particular, in the reported table, I require the firm j to have applied to at least one granted patent in the two year 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 are going to 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.

4.3 Identification assumptions

The previous analysis suggests that the decision may have had a positive effect on innovation. However, before I can interpret these effects causally, I need to provide a more thorough discussion of the identification assumption of the model.

In particular, the causal interpretation of the difference-in-difference approach relies on the parallel trend assumption. In a discrete treatment setting, this assumption requires that the relative dynamic of both treatment and control would have been the same in the absence of the shock. In this case, this requires that the relative behavior of high and low exposed firms would have not changed without the Supreme Court ruling. For instance, this assumption would be violated if litigation exposure were just a proxy of higher growth in innovation. While the evidence on the persistence of litigation over time would be at odds with this specific hypothesis, I can explore the dynamic of innovation by firms before the decision to rule this out and to provide more general evidence consistent with the parallel trend assumption.

I estimate the time-varying effect of exposure to litigation on patenting relative to the last period before the decision. To do this, I use patent data at quarterly frequency and I estimate the following equation:

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

Consistent with the parallel trend assumption, I would expect to find that: (a) the positive effect only appears in quarters after the Supreme Court decision ($\beta_t > 0$); (b) before the decision, the changes in patenting behavior are orthogonal to the measure of exposure ($\beta_t = 0$). I present the results of this test in the Figure (5). Firms characterized by different exposure to litigation did not have differential pattern in patenting before the Supreme Court decision. The estimated β in this period is always small in size and statistically non-different from zero. However, after the Supreme Court decision firms that were more exposed to litigation increased their patenting more. In particular, the effects turn positive already within a few quarters and keep rising afterwards.

An alternative approach is to test for pre-trending assuming that the relationship between exposure to litigation and patenting were linear, both before and after the shock.[32] While less flexible than the previous specification, this approach allows me to obtain more precise estimates of the trends. As expected, exposure to litigation does not predict differential before the decision, but only afterwards (Table A.6). The estimate of the effects of exposure to litigation before the decision is not only non-significant, but also small in size and of the opposite sign than a violation of this assumption would predict.

The same pattern is confirmed by looking to quality metrics of innovation, which are discussed later in

[32] I essentially estimate $y_{jt} = \alpha_j + \alpha_t + \beta^{PRE} R_j \cdot Pre + \beta^{POST} R_j \cdot Post + \epsilon_{jt}$

the paper. While not every outcome is positively affected by the decision, in every case I find that before the Supreme Court decision, the measure of litigation exposure does not predict differential growth rates. This is true both in a non-parametric test (Figure A.3) and when assuming linearity of the treatment effect (A.6). All in all, these analyses confirm that the positive effect of the ruling on innovation does not reflect differential trends in the data.

In order to provide further evidence in favor of the quality of the setting, I implement a battery of placebo tests. Since the effects of my analysis are contingent on the Supreme Court ruling, I should find no effect of exposure to litigation if I replicate my analysis in periods where there is no change in rules. In line with this intuition, I estimate the same model in equation (3) but center the analysis in a quarter where there is no change in patent law.[33] In order to avoid to arbitrarily choose a period where to run the placebo, I estimate a battery of placebos. In particular, I center my analyses in every quarter in the closest two years before the shock and such that the post-period does not overlaps with the post-treatment period. This is to say, I look at every quarter between 2002Q2 and 2004Q1.[34]

The results are reported graphically in Figure (6), where I estimate the main specification looking at the intensive margin. In particular, I plot the β and its 95% confidence interval for each quarter of the period considered. As expected in a placebo test, the coefficient is never positive and significant. In other words, in periods where there is no major shift in 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 size is always very small and never statistically different from zero. Altogether, the results from the placebo test seems to support the quality of my empirical setting.

Lastly, I develop a permutation test (Chetty et al. 2009; Fisher 1922), where I compare the t-statistic from my analysis to a non-parametric distribution of statistics that I obtain by randomly assigning technology classes to firm. The objective of this test is twofold. First, this methodology allows me to provide inference based on weaker assumptions than the standard linear model. Second, this test can be used to evaluate whether my analysis is capturing some other spurious firm characteristic that is different than litigation exposure but somehow correlated with it. For instance, this analysis allows me to reject that my results are somehow driven mechanically by the way the exposure index is constructed.

The intuition for this test is simple: if my results correctly capture the exposure to litigation through the technology fields, I would expect to find no results when technology exposure is randomly assigned. Rather

[33] In order to do so, I reconstruct the outcomes and regressors as if the shock happened right after the quarter of interest.

[34] Clearly, after 2004Q1, the post period of the placebo analysis would overlaps 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 closest I go to 2006Q1, the more my analysis would look like the main results. Consistent with this, I find that post 2004Q1 the β starts converging towards 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.

than a one-to-one comparison, I implement this test by constructing a full distribution of test statistics obtained in this way. If my model is correctly capturing true exposure to patent litigation, I would expect the true statistic to be on the top percentiles of this distribution.

In short, the procedure is implemented in the following way: I start by re-assigning randomly the technology classes in which a company operates for every firm in the sample.[35] Then, based on this, I reconstruct the exposure index $Exposure_j$ and I run the main specification presented before. I repeat this procedure for a thousand random iterations and then I plot the non-parametric distribution of the t-statistic I obtain from this. I compute the p-value of my true model by looking at the percentile in which my true t-statistic is within the constructed distribution. As expected, I find that the p-value constructed based on the random permutation test is similar to the standard one, and lower than 1% (Figure 7). Also this test confirms the quality of my empirical framework.

4.4 Robustness and other results

The previous section helped me to rule out the presence of non-parallel trends across firms differentially exposed to the shock, as well as other confounding factors in the analysis. However, I cannot exclude with those analyses the presence of a shock contemporaneous to the ruling that was correlated positively (negatively) with the exposure to litigation and affecting positively (negatively) innovation. The main candidate for an omitted variable is an industry level shock. For instance, during the same months of the Supreme Court decision, there could have been a positive productivity shock to a high litigation industry like computer.

I exclude that this could be driving force behind my results in two ways. First, I show that my results still hold when I exclude, one at the time, all the major industries in my sample (Table A.3). As previously discussed, I categorize every firm in an industry based on the major technology in which the company patented. In line with previous literature in innovation, I use large technology grouping from Hall et al. (2001). Second, I replicate the analyses exploiting only within-industry variation. I implement this by augmenting my model with industry by time fixed effects (Table 2). This set of controls removes from the data any industry trend, comparing patenting by firms with different exposure to litigation within the same industry. When I compare these estimates to those obtained in the simple model, I find no difference in the statistical significance and magnitude of the results.[36] Therefore, while industry dynamics could be important in explaining patenting behavior around this period, they do not seem to drive my results.

In the same Table, I augment the previous specification with another set of controls. In particular, I add a set of dummies that non-parametrically control for the location of the firm R&D facilities - based on state

[35] For instance, a firm that has obtained two patents in class 131 (Tobacco) and three in 428 (Stock material or miscellaneous articles) can be assigned to have two patents in class 432 (Heating) and three in 125 (Stone working).

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

of operation where I find more patents before the decision. Similarly, I control for the size of the portfolio of the firm, measured by the count of patents published in the years before the decision, but outside the estimation window;[37] quality of portfolio, measured by the average number of citation before the decision; and a dummy for firms that patented for the first time in the three years before the decision. In Table (3), I show that adding these controls does not systematically affect the results, as both magnitude and statistical significance remain very similar.

As expected, these results replicate when using a Poisson model, instead of the linear specification. In particular, I estimate this using a fixed-effects Poisson model, where I allow errors to be clustered within firms (Table A.4). The model is estimated using quarterly observations for the same time period and firms used in the other analyses. Since the log-linear specification employed is simply a log-transformation of a Poisson model, the coefficient of the two models are directly comparable. As expected, I find that the coefficients do not change in magnitude.

Furthermore, I can replicate the results estimated by equation (3) using an alternative measure of patent litigation exposure $Exposure_j^{LARGE}$. As discussed before, this measure uses patent data applied by the firm in the ten years before the shock and patent litigation data since 1980. Results are reported in Table (A.5) and they show essentially no change in our interpretation of the results. More broadly, results are stable when using alternative sub-periods of the data in estimating patent exposure. This is not surprising, since both the technology focus at firm level and cross-sectional distribution of litigation intensity are very stable over time.

4.5 Timing of the effects

I examine the timing of the effects by studying how they change when considered only one, two or three years after the decision. In order to do so, I repeat the same estimation as before, keeping the pre-period fixed and moving the post-period accordingly. Since I am interested in the change in magnitude of the coefficient across different specifications, I estimate this model without industry time trends.

There are two mains results in Table (3). First, I find that the effect is increasing over time. Relatively to one year after the decision, the effect over two years increases by 38% and over three years it increases by 50%. This is consistent with the idea that changes in the production function of innovation will reflect in the output with a lag. Second, the effect is already positive and significant after one year. On the one hand, this quick response allows me to relax the identification assumptions. In fact, the lack of pre-trending and the fast response confirm that this shift in innovation output is the result of a shock that happened in the early summer of 2006, the time of the Supreme Court decision. However, on the other hand, this result

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

may suggest that, at least partially, the increase in patent application stems from a shift in the incentives for patents rather than a true change in the innovation.

To shed light on this result, I explore the heterogeneity of the results across industries. In Table (A.7), I show that the whole positive result in the first year is driven by companies whose main industry is "Computer and Communications" (Hall et al., 2001). For this area, the R&D cycle is faster than the other technologies and therefore it is not surprising that these companies can react quicker to a change in incentives. However, the difference between this industry and the rest of the sample fades away over time. This confirms that the larger one-year effect does not reflect that this industry was, all else equal, more impacted by the ruling, but rather a different timing of R&D cycle.

4.6 Evidence on patent quality

The results so far suggest that the Supreme Court decision had a positive impact on the firms' ability to patent. In this Section, I show that also quality of innovation changes. In order to do so, I use the same empirical model as before, but focus on set of quality metrics that are constructed on patent citations.[38] Previous research has shown that forward citations are correlated with the quality of the underlying patent and its economic value (Hall et al., 2005, Kortum and Lerner, 2000). Here, I construct different measures based on citations in order to capture different aspects of quality (Appendix A.4).

First, I examine whether the increase in quantity of innovation came at complete detriment of quality. In order to do this, I examine the number of citation-weighted patents filed by companies around the decision (Mann, 2013). This measure combines both information about the quantity and quality of the innovation output. Since comparing number of citations across technologies and over time can be challenging (Lerner and Seru, 2015), 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.[39] I find that citation-weighted patents increased as well after the decision (Table 4).[40] This result rejects the hypothesis that an increase in the innovation output was reached simply by lowering the quality of R&D.

Second, I test whether the decision was able to increase companies' ability to develop breakthrough innovation (Kerr, 2010). Since the returns of innovation are highly skewed (Pakes, 1986), these patents can be very relevant for both firm value and welfare. In order to look at this, I examine the probability that a

[38]Since patent citations increase over time, their measure is sensitive to the date at which the patent was granted, relatively to the last date in 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 to 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 and there is large serial correlation in citations (Akcigit and Kerr, 2010). See also Lerner and Seru (2015).

[39]In this case, however, this adjustment does not play a major role and results with standard citations are very similar.

[40]Since some firms receive no citations, the panel is not perfectly balanced. However, since data are collapsed in two period (pre and post decision), this feature does not affect the estimated β.

company applies for patents that are at the top - in particular top 10% and 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 the same USPTO technology class and were developed in the same year. Once I identified these exceptional patents, I generate a dummy equal to one if the company has applied to any of these top patents.

As reported in Table (4), I find that the Supreme Court decision increased the probability of applying for a potential breakthrough patent. In particular, firms that were more exposed to litigation appeared to be more likely to patent a technology which is in the top 10% or 25% of the quality distribution after the decision. In economic terms, an exposure one standard deviation higher before the decision led to a higher increase in the probability of patenting something in the top 10% of the quality distribution by about 1%. This corresponds to a 3% increase with respect to the baseline probability. The results hold when controlling for usual set of covariates. Furthermore, as previously discussed, these outcomes do not show any evidence of pre-trending before the decision (Table A.6; Figure A.3).

Overall, these results suggest that the improvement in patent enforcement was able to increase the amount of innovation produced by companies, but also to positively affect the quality of the output.[41]

4.7 Evidence from public firms

So far, the analysis showed that the Supreme Court decision had positive effects on one specific proxy for innovation, firms patenting. I now study a smaller set of firms – innovative firms that were public at the time of the decision – and show that for this group the decision also led to an increase in R&D intensity.

If the decision really affected the innovation incentives, I expect to find an increase in both the input and output of innovation. While 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 one thousand firms that are active in patenting around the decision. These are identified matching patent data to Compustat using data from Kogan et al. (2012), as discussed in section (3).

With this sample at hand, I estimate the same equation (3), looking both at patenting and R&D intensity. Results are reported in Table (5). First, I find that quantity of patenting increases also for this smaller sample. Second, I find that the decision also led to an increase in R&D intensity, here measured by R&D over asset. Comparing two firms one standard deviation apart in litigation exposure, R&D investment increased for the

[41] In addition, in an unreported Table, I find that firms also increased the quantity of "good patents", which are patents that received at least one citation after the decision.

more exposed firm by about 1% per year. This corresponds to an 8% increase with respect to the sample average of R&D intensity. These results are not driven by failure of non-parallel trend assumption (Figure 8 and Table A.8). As suggested by Figure (8), firms started increasing their R&D spending already within one year from the decision and this does not revert back in the following one. If anything, the R&D of more exposed firms keeps rising also in this second period.

This confirms that the Supreme Court decision "eBay versus MercExchange" had a significant, positive effect on innovation. Both patenting and R&D, at least for public firms, increased. The decision, which reduced the potential cost faced by firms in case of litigation, was successful in freeing up resources for innovation.

5 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 margin. Furthermore, this change in enforcement also positively affected patent quality, fostering the development of potential breakthrough patents. Lastly, also R&D investment increased. 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.

5.1 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. Since direct involvement in patent litigation can be extremely 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.

This channel has two predictions regarding what should happen when the burden of patent litigation is exogenously reduced. 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 area where patent litigation is more intense should become relatively more valuable. This reshuffle should happen in every firm, irrespective of whether they are more or less exposed to litigation. In other words, every firm should perceive the investment in more risky patents to be more valuable.

In order to provide evidence in favor of this idea, I study whether firms experienced a relatively higher increase in risky patents after the decision. In order to focus on the within-firm resource allocation, I sort patents applied by each firm across two categories - risky and non-risky – 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. Since 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:

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

where α_{jt} is a set of firm-time fixed effects, α_{jr} is a set of fixed effects at firm-group level, $1\{Risk_r\}$ is a dummy for more risky groups. As mentioned before, I group patents in two classes, such that $r = \{high\,risk; low\,risk\}$. 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 non-risky patents within the firm portfolio, which is $\beta > 0$.

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 patent applications that firm j filed to during time t in the class of risk r. To obtain a purely intensive margin, I estimate this regression with a subset of firms – around 3,000 – that are simultaneously active in both risk classes, around the decision time. Second, I look at the extensive margin where I have y_{jtr} 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, since I consider every firm that has applied to at least one patent in the ten years before the decision. As usual, the analysis is collapsed before and after the decision to provide more conservative inference (Bertrand et al., 2004) and standard errors are clustered at firm level.

Results are reported in Table (4). When I look at the intensive margin, I find no relative effect on risky patents: within firm, patents belonging to more intensively litigated patent classes do not appear to increase more. Estimates are not only non-significant but also small. 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. Results are similar whether risky patents are defined by looking at the top 10% or the top 25%. Furthermore, in Table (A.9) I find 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. While 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.[42]

5.2 Litigation exacerbates financial constraints

The previous results confirm that patent litigation, lowering the returns on innovation, reduces firms' incentives to invest in it. In this section, I argue that this is not the only channel in place. Instead, operating in a high-litigation environment can also hinder innovation by reducing the amount of resources available for R&D. Given the frictions in the financing of innovation, this reduction in internal resources can translate into lower investment.

The idea that exposure to litigation can deplete corporate resources is supported by both previous research and anecdotal evidence. Firms in sectors where litigation is more intense are more likely to pay large settlements or overpaying for licensing agreements. This happens because companies want to avoid the escalation of legal conflicts to courts or just limit its negative consequences, like in the BlackBerry case previously discussed. 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. These views can be often identified in public records: 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 in deterring attacks (Cohen et al., 2014).

It is worthy to point out that intense litigation does not only affect monetary resources. In line with the above comment, the time of management and R&D specialists is another dimension of this issue. In companies exposed to litigation the management has to invest extra time and effort around intellectual property issues, in order to avoid incurring potential violations or attracting the interest of patent assertion entities. Overall these concepts are well summarized by a quote from a VC surveyed by Feldman (2014): "when companies spend money trying to protect their intellectual property position, they are not expanding; and when companies spend time thinking about patent demands, they are not inventing."

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. Within this framework, the non-monetary aspect of this reduction

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

in resources does nothing but aggravating the overall issue.

To test whether the theory is true in the data, I examine the heterogeneity of the decision effects across firms characterized by 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 reduction in litigation cost for companies facing more financial frictions.

In order to study this, I modify the standard model described by equation (3) 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:

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

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). 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 dividends. Also in this case, I define a company as non-dividend payers if she pays no cash dividends in the three years before the decision.

The results are reported in Tables (7) and (8). The decision led to an increase in R&D intensity only for firms that are 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 non-significant 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 than non-financially constrained firms in terms of patent applications.[43]

[43] On the one hand, this is consistent with the presence of two distinct channels. Independent from the financial situation, every firms should patent more after the decision because innovation becomes more profitable or less risky. Therefore, I should not find

As a robustness, I show that, in my case, the results are not simply capturing heterogeneity across firms in the 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 decisions to my treatment. In particular, in Table (A.10) I report the results looking at revenue growth. I find that, if anything, the main coefficient $\beta_1$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.

6 Supreme Court decision and stock prices

In the end, I show that the decision had positive effects on the stock returns of innovative firms. Implementing an event study around the decision, I find evidence that firms more exposed to litigation experienced higher excess returns around the decision.[44]

Previous research in finance has shown that innovation can positively affect the stock market valuation of firms (Kogan et al., 2012). If this is the case, an improvement in the enforcement of patents should positively affect the stock prices of innovative firms. This should be particularly the case for companies for which this dimension is particularly relevant, such as firms that operate in areas where patent litigation is intense. The problem with this type of analysis is that investors may not be immediately aware of the positive effects of the decision on innovation. In particular, with respect to the effect on non-practicing entities (NPE), the consequences of the decision on standard innovative firms may be harder to identify. First, while the decision had an unambiguous, negative effect on NPEs, the impact of the ruling on innovation is less clear (section 2.2), in particular without a deep understanding of the patent litigation market. Second, a change in patent enforcement is clearly more salient for non-practicing entities.

In order to study this question, I measure returns and abnormal returns around the announcement and I correlate these measures with the measure of litigation exposure. If the ruling had positively affected the

that only financially constrained firms increase patent applications after the shock. On the other hand, this is puzzling because I would still expect firms more likely to be financially constraint to respond relatively more in terms of patent applications. A tentative explanation for this null result is that the effect of financial constraints is harder to be detected with this outcome because patent applications respond for every firm. Therefore, the treatment effect can be expected to be smaller and harder to identify. Furthermore, more financially constrain firms may invest less in R&D as they operate in more intensively litigated area, but this lower investment does not need to fully translates into lower quantity of output. For instance, companies may still produce innovation, but focus on less expensive or ambitious areas. In this case, as the amount of internal resources increase, company may channel the extra resources both to increase the output and change the type or quality of the projects undertaken.

[44] See section A.4.3 in the Appendix for more info on data construction and analysis.

value of innovative firms, the out-performance around the decision should be larger in companies that operate in technology fields where litigation is more intense.

The main result of the analysis can be synthesized by Figure (9), which plots the cumulative value-weighted returns of high and low exposure firms.[45] The two groups almost overlap in the days before the decision. However, the day of the decision, the high-risk group outperforms the low-risk group by almost 1%. This out-performance does not revert right after, and the two groups seem to have similar paths in the following days.

I can confirm all these results in a regression framework (Table 9), where I run cross-sectional value-weighted regression between firm returns and ex-ante exposure to litigation. As usual, I focus on the sample of innovative public firms for which I find return information on CRSP around the decision. I find similar results when looking at raw returns or abnormal returns at the time of the news release. The effect is still positive and significant when looking at the end of the week. Furthermore, the formal test also rejects that this result could be driven by differential trends in returns before the decision. For instance, in the week before the decision differential exposure to litigation does not seem to predict differential returns.[46]

7 Conclusion

This paper examines how patent litigation affects innovation using the 2006 Supreme Court decision "eBay versus MercExchange" as an exogenous shock to patent enforcement. The evidence provided suggests that this intervention had a positive effect on innovation. 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 sub-sample 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.

The decision also impacted quality of innovation. Firms more exposed to patent litigation increased the likelihood of patenting breakthrough technology. Moreover, I found evidence for an increase in patenting even when examining patents weighted by citations received, which excludes that the increase in quantity came at total detriment of quality. I interpret these results as suggesting that the decision made firms more able to take risky projects. Given that the returns of innovation are highly skewed (Pakes, 1986), the shock

[45] High risk firms are firms above the top 25% of litigation exposure.

[46] In unreported regression, I replicate this result estimating the standard errors clustering them at major technology level, finding almost identical results. Since the number of major technology is small, there may be a concern that this result is indeed driven by the small number of cluster. Therefore, I also check the size of the standard errors when clustering at SIC two-digit level, as reported in Compustat. The results are again very similar. Overall, the clustered standard errros are very similar to the robust one.

had positive effects on the ability of a firm to grow and compete.

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 towards 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 the recent policy interventions, such as the American Innovation Act (2011). This analysis is crucial for guiding future policy work and it can also provide a nice laboratory to gain better insights on the mechanisms through which abusive litigation hinders innovation. In addition, more work can be done to examine the role of patent litigation in start-up companies. The nature of my identification strategy focuses on established firms and therefore the results do directly apply to start-up companies. However, the evidence on the importance of financial frictions to determine the cost of patent litigation may suggest that start-up companies should be even more affected.

The results presented in this paper support the idea that patent litigation can significantly affect companies' innovation. As a result, policies that mitigate the overhang of litigation can have beneficial effects on technology advancement. In particular, improvements in the quality of patent enforcement, which reduces the legal uncertainty around patents and limits abusive behaviors in this market, can increase firms' ability and incentives to invest in R&D. Recent efforts in the U.S., such as the American Innovation Act (2011), 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.

References

Acharya, V. V., Y. Amihud, and L. Litov (2011). Creditor rights and corporate risk-taking. *Journal of Financial Economics 102*(1), 150–166.

Acharya, V. V. and K. V. Subramanian (2009). Bankruptcy codes and innovation. *Review of Financial Studies*, hhp019.

Akcigit, U. and W. R. Kerr (2010). Growth through heterogeneous innovations. Technical report, National Bureau of Economic Research.

Almeida, H., M. Campello, and M. S. Weisbach (2004). The cash flow sensitivity of cash. *The Journal of Finance 59*(4), 1777–1804.

Angrist, J. D. and J.-S. Pischke (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Arena, M. P. and B. Julio (2014). The effects of securities class action litigation on corporate liquidity and investment policy. *Journal of Financial and Quantitative Analysis*.

Beatty, A., J. Weber, and J. J. Yu (2008). Conservatism and debt. *Journal of accounting and economics 45*(2), 154–174.

Bernstein, S. (2015). Does going public affect innovation? *The Journal of Finance*.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics 119*(1), 249–275.

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

Bessen, J. and M. J. Meurer (2008b). *Patent failure: How judges, bureaucrats, and lawyers put innovators at risk*. Princeton University Press.

Bessen, J. and M. J. Meurer (2013). Direct costs from npe disputes, the. *Cornell L. Rev. 99*, 387.

Bessen, J. E., P. Neuhäusler, J. L. Turner, and J. W. Williams (2015). Trends in private patent costs and rents for publicly-traded united states firms. *Boston Univ. School of Law, Public Law Research Paper*.

Boldrin, M. and D. Levine (2002). The case against intellectual property. *American Economic Review*, 209–212.

Brown, J. R., S. M. Fazzari, and B. C. Petersen (2009). Financing innovation and growth: Cash flow, external equity, and the 1990s r&d boom. *The Journal of Finance 64*(1), 151–185.

Calabresi, G. and A. D. Melamed (1972). Property rules, liability rules, and inalienability: one view of the cathedral. *Harvard Law Review*, 1089–1128.

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

Chetty, R., A. Looney, and K. Kroft (2009). Salience and taxation: Theory and evidence. *American Economic Review 99*(4), 1145–1177.

Chien, C. V. and M. A. Lemley (2012). Patent holdup, the itc, and the public interest. *Cornell Law Review*.

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

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

Cohen, L., U. Gurun, and S. D. Kominers (2014). Patent trolls: Evidence from targeted firms. *HBS Working Paper*.

Court, S. (1908). Continental paper bag co. v. eastern paper bag co.

Court, S. (2006). ebay inc. v. mercexchange, ll.

Court, V. D. (2003). Mercexchange, llc v. ebay, inc.

Dass, N., V. K. Nanda, and S. C. Xiao (2015). Truncation bias in patent data: Does it explain why stock-liquidity seemingly reduces innovation? *SSRN 2584221*.

Demirgüç-Kunt, A. and V. Maksimovic (1998). Law, finance, and firm growth. *Journal of finance*, 2107–2137.

Djankov, S., R. La Porta, F. Lopez-de Silanes, and A. Shleifer (2003). Courts. *The Quarterly Journal of Economics*, 453–517.

Epstein, R. A. (2008). The property rights movement and intellectual property. *Regulation 30*(4), 2007–2008.

Farre-Mensa, J. and A. Ljungqvist (2015). Do measures of financial constraints measure financial constraints? *Review of Financial Studies*.

Fazzari, S. M., R. G. Hubbard, B. C. Petersen, A. S. Blinder, and J. M. Poterba (1988). Financing constraints and corporate investment. *Brookings Papers on Economic Activity*, 141–206.

Feldman, R. (2014). Patent demands & startup companies: The view from the venture capital community. *Yale Journal of Law & Technology 16*(2), 236.

Feng, J. and X. Jaravel (2015). Patent trolls and the patent examination process.

Fisher, R. A. (1922). On the interpretation of $\chi 2$ from contingency tables, and the calculation of p. *Journal of the Royal Statistical Society*, 87–94.

Galasso, A. and M. Schankerman (2015). Patents and cumulative innovation: Causal evidence from the courts. *The Quarterly Journal of Economics 130*(1), 317–369.

Gormley, T. A. and D. A. Matsa (2014). Common errors: How to (and not to) control for unobserved heterogeneity. *Review of Financial Studies 27*(2), 617–661.

Grumbles III, E., R. C. Hughey, and S. Perera (2009). the three year anniversary of ebay v. mercexchange: a statistical analysis of permanent injunctions. *Intellectual Property Today*.

Gupta, K. and J. P. Kesan (2015). Studying the impact of ebay on injunctive relief in patent cases. *Available at SSRN 2629399*.

Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *RAND Journal of economics*, 16–38.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). The nber patent citation data file: Lessons, insights and methodological tools. Technical report, National Bureau of Economic Research.

Hall, B. H. and J. Lerner (2010). The financing of r&d and innovation. *Handbook of the Economics of Innovation 1*, 609–639.

Haslem, B. (2005). Managerial opportunism during corporate litigation. *The Journal of Finance 60*(4), 2013–2041.

Helm, J. S. (2006). Why pharmaceutical firms support patent trolls: The disparate impact of ebay v. mercexchange on innovation. *Mich. Telecomm. & Tech. L. Rev. 13*, 331.

Holte, R. T. (2015). The misinterpretation of ebay v. mercexchange and why: An analysis of the case history, precedent, and parties. *Chapman Law Review*.

House, W. (2013). Patent assertion and u.s. innovation.

Iverson, B. C. (2014). Get in line: Chapter 11 restructuring in crowded bankruptcy courts. *Available at SSRN 2156045*.

Jaffe, A. B. and J. Lerner (2011). *Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it.* Princeton University Press.

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

Kashyap, A. K. and O. A. Lamont (1994). Credit conditions and the cyclical behavior of inventories. *Quarterly Journal of Economics 109*(3).

Kerr, W. R. (2010). Breakthrough inventions and migrating clusters of innovation. *Journal of Urban Economics 67*(1), 46–60.

Kim, I. and D. J. Skinner (2012). Measuring securities litigation risk. *Journal of Accounting and Economics 53*(1), 290–310.

King, R. G. and R. Levine (1993). Finance and growth: Schumpeter might be right. *The Quarterly Journal of Economics*, 717–737.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2012). Technological innovation, resource allocation, and growth. Technical report, National Bureau of Economic Research.

Kortum, S. and J. Lerner (2000). Assessing the contribution of venture capital to innovation. *RAND journal of Economics*, 674–692.

La Porta, R., F. Lopez-de Silanes, A. Shleifer, and R. W. Vishny (1997). Legal determinants of external finance. *Journal of Finance*, 1131–1150.

Lanjouw, J. and M. Schankerman (2001). Characteristics of patent litigation: a window on competition. *RAND journal of economics*, 129–151.

Lanjouw, J. O. and J. Lerner (1998). The enforcement of intellectual property rights: A survey of the empirical literature. *Annales d'Économie et de Statistique*, 223–246.

Lemley, M. A. and C. Shapiro (2005). Probabilistic patents. *Journal of Economic Perspectives*, 75–98.

Lemley, M. A. and C. Shapiro (2006). Patent holdup and royalty stacking. *Tex. L. Rev. 85*, 2163.

Lerner, J. (1995). Patenting in the shadow of competitors. *Journal of Law and Economics 38*(2), 463–95.

Lerner, J. (2006). Trolls on state street?: The litigation of financial patents, 1976-2005. *Harvard Business School manuscript*.

Lerner, J. (2009). *The empirical impact of intellectual property rights on innovation: Puzzles and clues.*

Lerner, J. and A. Schoar (2005). Does legal enforcement affect financial transactions? the contractual channel in private equity. *The Quarterly Journal of Economics*, 223–246.

Lerner, J. and A. Seru (2015). The use and misuse of patent data: Issues for corporate finance and beyond.

Li, G.-C., R. Lai, A. D'Amour, D. M. Doolin, Y. Sun, V. I. Torvik, Z. Y. Amy, and L. Fleming (2014). Disambiguation and co-authorship networks of the us patent inventor database (1975–2010). *Research Policy 43*(6), 941–955.

Mann, W. (2013). Creditor rights and innovation: Evidence from patent collateral. *Available at SSRN 2356015*.

Moser, P. (2013). Patents and innovation: evidence from economic history. *The Journal of Economic Perspectives*, 23–44.

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

Ponticelli, J. (2013). Court enforcement and firm productivity: evidence from a bankruptcy reform in brazil. *Chicago Booth Research Paper* (14-08).

Rogers, J. L. and A. Van Buskirk (2009). Shareholder litigation and changes in disclosure behavior. *Journal of Accounting and Economics 47*(1), 136–156.

Sakakibara, M. and L. Branstetter (2001). *Do stronger patents induce more innovation? Evidence from the 1988 Japanese patent law reforms.*

Schwartzstein, J. and A. Shleifer (2013). An activity-generating theory of regulation. *Journal of Law and Economics 56*, 1–38.

Seaman, C. B. (2016). Permanent injunctions in patent litigation after ebay: An empirical study. *Washington and Lee University School of Law.*

Shapiro, C. (2010). Injunctions, hold-up, and patent royalties. *American Law and Economics Review.*

Shapiro, C. (2016a). Patent remedies. *The American Economic Review 106*(5), 198–202.

Shapiro, C. (2016b). Property rules vs. liability rules for patent infringement. *Liability Rules for Patent Infringement (May 4, 2016).*

Smeets, R. (2014). Does patent litigation reduce corporate r&d? an analysis of us public firms. *Working Paper*.

Tang, Y. H. (2006). Future of patent enforcement after ebay v. mercexchange, the. *Harv. JL & Tech. 20*, 235.

Tucker, C. (2014). Patent trolls and technology diffusion: The case of medical imaging.

Tucker, C. (2015). The effect of patent litigation and patent assertion entities on entrepreneurial activity. *Research Policy*.

Venkatesan, J. (2009). Compulsory licensing on nonpracticing patentees after ebay v. mercexchange. *Va. JL & Tech. 14*, 26.

Wesenberg, E. and P. O'Rourke (2006). *The Toll on the Troll*.

Williams, H. L. (2015). Intellectual property rights and innovation: Evidence from health care markets. In *Innovation Policy and the Economy, Volume 16*. University of Chicago Press.

Figure 1: Number of cases involving patents

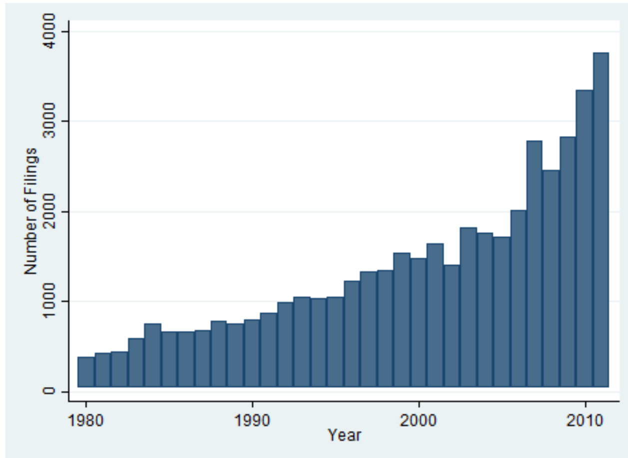

This plot reports the number of filings involving patents of any type per year of filing, between 1980 and 2012. The data comes from WestLaw-ThomsonReuters, which collected filings information from public records. Data are plotted at docket-number level, therefore they do not account for the fact that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

Figure 2: R&D and Patenting by Corporations

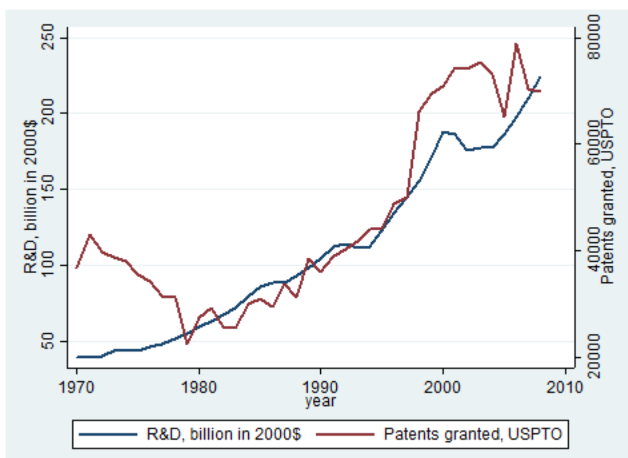

This figure reports the time series plot of R&D expenditure by corporations and patents granted to corporations by year since 1970 until 2008. Data on R&D expenditure is collected from the National Science Foundation (NSF), Division of Science Resources Statistics, in the national patterns in R&D 2008. The R&D expenditure is expressed in billion of 2000$. Patents data are instead from USPTO aggregate statistics that can be found at http://www.uspto.gov/web/offices/ac/ido/oeip/taf/h_at.htm. The series report the raw data, no adjustments have been made.

Figure 3: NPEs Stock Returns around the decision

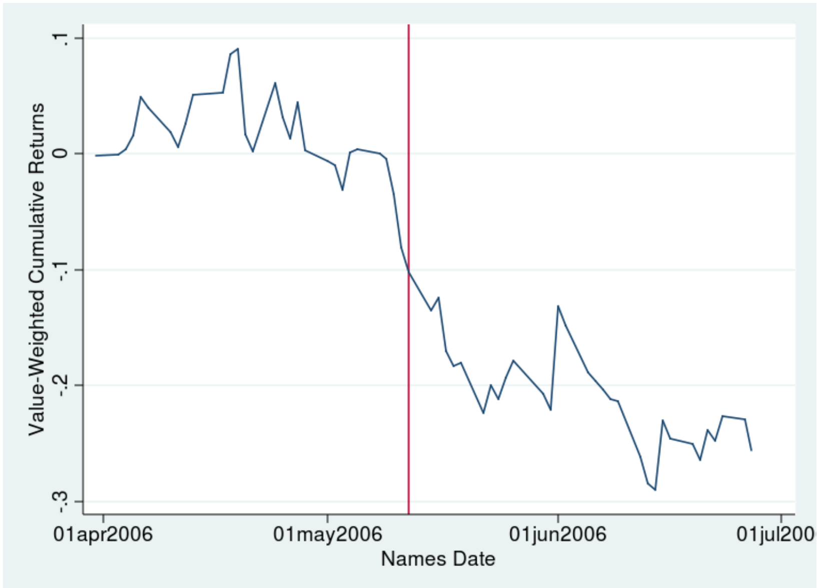

This Figure plots the average cumulative returns, for the sample of NPEs identified in the paper. The sample of firms used are 10 companies that are (a) Identified as NPEs; (b) Public at the time of the Supreme Court decision. The companies are Acacia Technologies, Asure Software (formerly Forgent Network), Rambus, Tessera Technologies, Universal Display, Document Security Systems, ParkerVision, Unwired Planet (formerly Openwave), Interdigital, Spherix. Information on the sample constructions are provided in Section (2). More info on this analysis is in Appendix (A.4.3). The straight red line correspond to the trading day right before the decision.

Figure 4: Distribution of firm industries for the top 50% more risky firms vs. whole sample

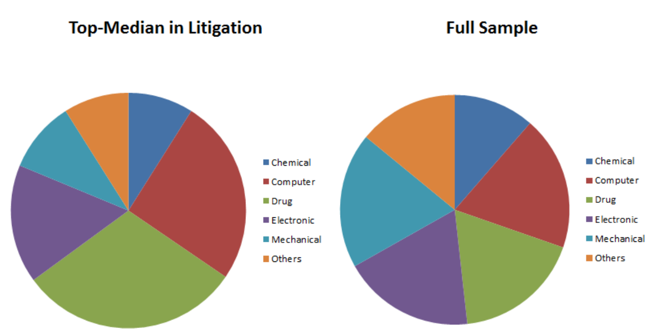

This Figure reports the pie chart of the patents by industry, across the full sample and the sample of firms that are more exposed to litigation. Industries are identified based on patent applications across macro-technology area (Hall et al., 2001) and the construction is discussed in detail in Appendix (A.4). 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 using litigation in the five years before the decision, and using patents since 2000. The second chart is instead constructed using the full-sample. Furthermore, the sample that was used to construct this plot is the sample of innovative firms that applied to at least one patents in the two years before or two years after the decision.

Figure 5: Effect of litigation over time

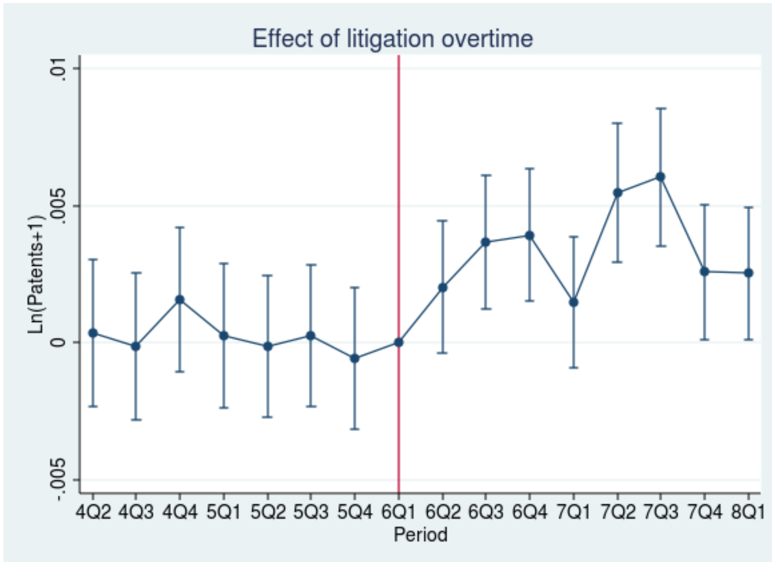

This Figure plots the β_t from equation (4). The red vertical line correspond to the last period of the pre-decision period. Every β_t is plotted with the correspondent CI at 5%. Every period is label with the corresponding quarter. Notice that quarters are in "event time" not calendar time: in fact, 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). Data used corresponds at the two years before and after the decision, in event time. The sample used correspond to the one of the extensive margin. Standard errors are clustered at firm-level.

Figure 6: Placebo test over time

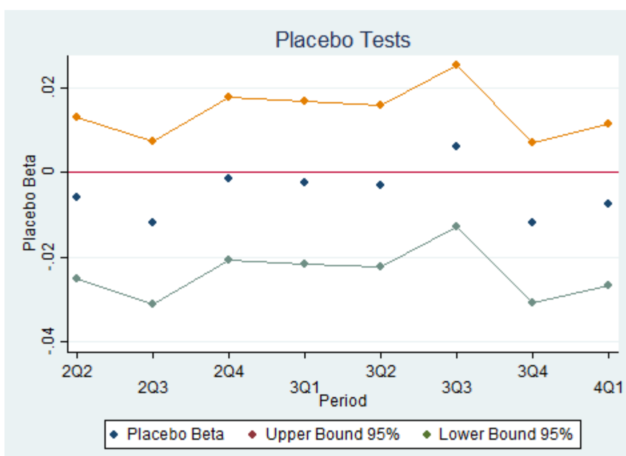

This represent the results from a set of placebo tests. In particular, in this Figure I construct a series of placebo samples, centered around fictional shocks in the two years that are completely outside the two years that are completely outside the period after the decision. The date in the x-axis is 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 (3), 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 the one ending in May 15th (the other quarters are constructed accordingly). Data used corresponds at the two years before and after the decision, in event time. Standard errors are clustered at firm level.

Figure 7: Permutation Test: distribution of test statistic

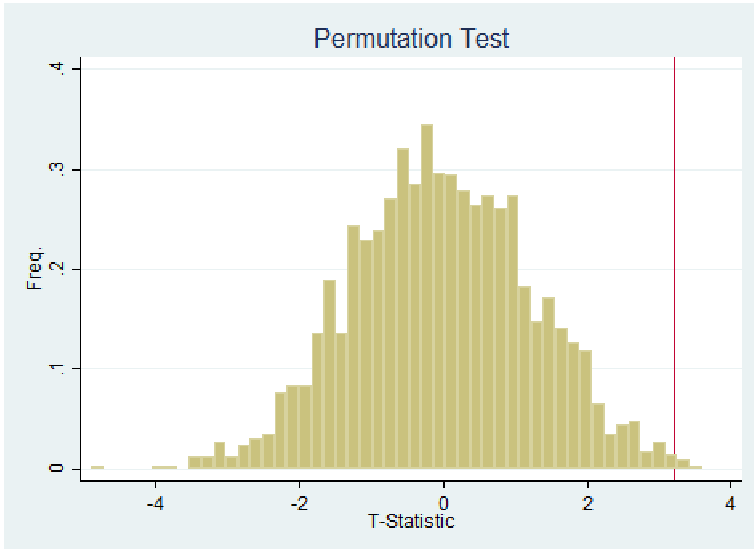

The Figure reports the results of the permutation test, where I compare the value of the t-statistic on the "true results" - which is reported by the straight red line - which the distribution of statistics that are constructed randomly assigning industries to firms. For every iteration of the procedure, I randomly assign technology class to firms. Then, I run the standard regression and store the t-statistic. Finally, after one thousand iterations, I plot them in a histogram as above. As mentioned, I plot the true estimates in the red line, which in this case belongs to the top 1% of the distribution of coefficients.

Figure 8: Effect of litigation on R&D intensity over time

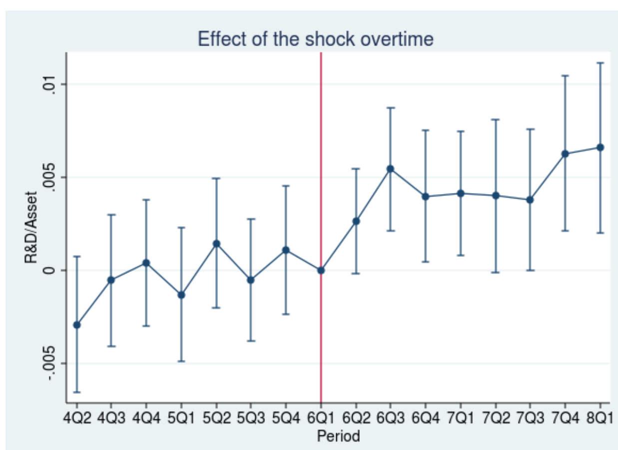

This Figure plots the β_t from equation (4) with the standard controls, where the outcome is R&D over asset. The red vertical line correspond to the last period of the pre-decision period. Every β_t is plotted with the correspondent CI at 5%. Every period is label with the corresponding quarter. Notice 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). Data used corresponds at the two years before and after the decision, in event time. The sample is the standard Compustat sample of innovative firms used.

Figure 9: Stock Market Reaction: High vs. Low exposure

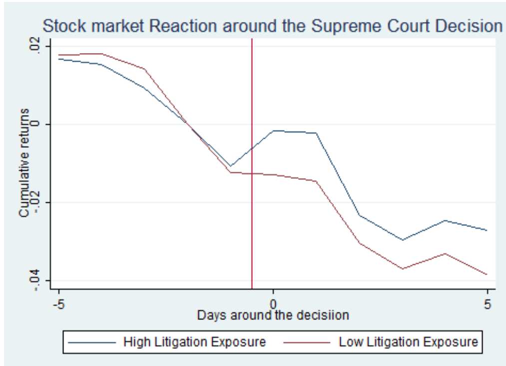

The Figures plots the value-weighted cumulative returns across high and low exposure firms. High litigation firms are firms at the top 25% of the litigation distribution. Cumulative returns are normalized to zero for both groups two days before the decision. The straight red line is plot between the 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.

Table 1: Summary Statistics

(a) Full sample

	Obs.	Mean	S.D.
$\#Patent_{jt}$	32,118	20.28	164.34
$1\{Patent_{jt} = Top^{10\%}\}$	32,118	0.30	0.46
$1\{Patent_{jt} = Top^{25\%}\}$	32,118	0.48	0.50
$Citation\,Weighted\,Pat_{jt}$	32,118	20.69	151.85
$Exposure_j$	32,118	0.77	0.79
$Exposure_j^{OLD}$	32,118	0.68	0.562
$Average\,Citation\,Pre$	32,118	1.19	0.46
$1\{Years\,first\,Patent \leq 3\}$	32,118	0.3	0.46
$Size\,Pre\,Portfolio$	32,118	18.97	146.74

(b) Public Firms

	Obs.	Mean	S.D.
$\#Patent_{jt}$	2,032	101.93	463.62
$R\&D/Asset$	2,032	0.03	0.04
$Exposure_j$	2,032	0.92	0.79
$Exposure_j^{OLD}$	2,032	0.77	0.55
$Average\,Citation\,Pre$	2,032	1.50	2.06
$1\{Years\,first\,Patent \leq 3\}$	2,032	0.04	0.19
$Size\,Pre\,Portfolio$	2,032	90.72	375.06

The two tables report summary statistics for the two main samples used in the analyses. In the first panel, I present the summary statistics for the variables that are used for the first set of analysis, where I employ both private and public firms, which are active in innovative activity around 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 decision. In the second panel, instead, I report summary statistics for the sample that is used in the second part of the analysis, which focus on public firms that patented around the decision. More info on the sample construction is available in the Appendix (A.4) . The variable construction is described in detail in the Appendix (A.4) - for outcomes - and in the Section (3) for the measures of exposure.

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***	0.035***	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	Y	Y	Y	Y	Y
$Indu. \times Time\,F.E.$		Y	Y		Y	Y
$Other\,Controls_{jt}$			Y			Y
R^2	0.005	0.007	0.033	0.217	0.283	0.290
Observations	32,118	32,118	32,118	155,866	155,866	155,866

This Table reports the estimate of the linear difference-in-difference specification (equation 3), where I estimate the effect of the decision on quantity of innovation. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt}$, where y_{jt} is: (a) the (natural) logarithm of granted patent that firm j applied during period t for Columns (1)-(3); (2) 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 year 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 Column (2) and (5), I add industry-time fixed effect to the equation. Industry 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 modal location reported in patent data), the size of the portfolio before the estimation period, the start-up status (looking at whether the 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 info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 3: Timing of the effects

	$ln(Patents_{jt})$					
	1 Year After		2 Years After		3 Years After	
$Post \cdot Exposure_j^{LARGE}$	0.036***		0.050***		0.058***	
	(0.012)		(0.012)		(0.0124)	
$Post \cdot Exposure_j$		0.029***		0.040***		0.047***
		(0.008)		(0.008)		(0.009)
$Firm\, F.E.$	Y	Y	Y	Y	Y	Y
$Time\, F.E.$	Y	Y	Y	Y	Y	Y
R^2	0.209	0.209	0.005	0.006	0.113	0.113
Observations	32,118	32,118	32,118	32,118	32,118	32,118

In this table I report the estimation of the equation 3. The data set is constituted by a balanced two-period panel.The first period is fixed to the two year before the decision, while the second period depends on the specification and in particular it moves from 1 to 3 years after. The outcome is always the (natural) logarithm of granted patent that firm j applied during period t. In this case, I use every firm that applied to 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. Similarly, the variable $Exposure_j^{LARGE}$captures the exposure of firm j to patent litigation, using patent application in the ten years before the decision and patent litigation at technology class since 1980. More info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 4: Evidence on patent quality

(a) Panel A

	(1)	(2)	(3)	(4)	(5)	(6)
OLS	$1\{Patent_{jt} = Top^{10\%}\}$			$1\{Patent_{jt} = Top^{25\%}\}$		
$Post \cdot Exposure_j$	0.010**	0.016***	0.016***	0.018***	0.022***	0.021***
	(0.005)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)
$Firm\,F.E.$	Y	Y	Y	Y	Y	Y
$Time\,F.E.$	Y	Y	Y	Y	Y	Y
$Indu. \times Time\,F.E.$		Y	Y		Y	Y
$Other\,Controls_{jt}$			Y			Y
R^2	0.001	0.001	0.005	0.001	0.001	0.004
Observations	32,118	32,118	32,118	32,118	32,118	32,118

(b) Panel B

	(1)	(2)	(3)
OLS	$ln(Citation\,Weighted\,Pat_{jt})$		
$Post \cdot Exposure_j$	0.032*	0.048**	0.047**
	(0.018)	(0.023)	(0.023)
$Firm\,F.E.$	Y	Y	Y
$Time\,F.E.$	Y	Y	Y
$Indu. \times Time\,F.E.$		Y	Y
$Other\,Controls_{jt}$			Y
R^2	0.001	0.001	0.002
Observations	22,673	22,673	22,673

This Table reports the estimate of the linear difference-in-difference specification (equation 3), where I estimate the effect of the decision on quality of innovation. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt}$, where y_{jt} is a proxy of average quality of innovation in the two years before or after the decision. In particular, in panel A, I consider two outcomes: (a) a dummy which 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 (within 3 years) among patents granted in the same year in the same technology class; (b) similar dummy, but constructed considering the top 25% of the distribution. In panel B instead I look at the logarithm of the citation weighted patents over the same period. Here, citations are scaled by the average number of citations received by patents in the same technology class and year, in order to account for time-varying patterns in patent citations. 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 (constructed using the modal location reported in patent data), the size of the portfolio before the estimation period and the start-up status (looking at whether the a firm applied for the first patent ever within the previous three years). More info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 5: Effect of the decision on public firms

	(1)	(2)	(3)	(4)	(5)	(6)
OLS		$ln(Patents_{jt})$			$R\&D_{jt}/Asset_{jt}$	
$Post \cdot Exposure_j$	0.063*	0.093**	0.099**	0.003**	0.004***	0.004***
	(0.033)	(0.045)	(0.045)	(0.001)	(0.002)	(0.002)
$Firm\, F.E.$	Y	Y	Y	Y	Y	Y
$Time\, F.E.$	Y	Y	Y	Y	Y	Y
$Indu. \times Time\, F.E.$		Y	Y		Y	Y
$Other\, Controls_{jt}$			Y			Y
R^2	0.007	0.017	0.078	0.010	0.018	0.063
Observations	2,034	2,034	2,034	2,034	2,034	2,034

This Table reports the estimate of the linear difference-in-difference specification (equation 3), where I estimate the effect of the decision on patenting and R&D intensity. In particular, I estimate $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt}$, where y_{jt} is: (a) the (natural) logarithm of granted patent that firm j applied during period t for Columns (1)-(3); (2) $R\&D/Asset$ 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 Appendix (A.4). 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 non-financial, US located public firms that applied to at least one patent in the two years before and one after (see appendix A.4). I always control for firm fixed-effects and time effects. In Columns (2) and (5) I augment this with industry-time fixed effect. Industry 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 modal 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 the 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 info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 6: Evidence on Patent Mix

	(1)	(3)	(5)	(7)
	Extensive Margin		Intensive Margin	
	$ln(Patents_{jtr})$		$1\{Patents_{jtr} > 0)$	
$Post \cdot 1\{Risk_r\}$	0.012	-0.027	0.270***	0.171***
	(0.025)	(0.022)	(0.005)	(0.005)
$Split$	10%	25%	10%	25%
$Firm \times Time\, F.E.$	Y	Y	Y	Y
$Firm \times Risk\, F.E.$	Y	Y	Y	Y
R^2	0.909	0.909	0.829	0.811
Sample	2,785	3,893	54,844	54,844
Observations	11,140	15,572	219,376	219,376

This Table estimates equation (5), which is: $y_{jtr} = \alpha_{jr} + \alpha_{jt} + \beta 1\{Risk_r\} \cdot Post$, where where α_{jt} is a set of firm-time fixed effects, α_{jr} is a set of fixed effects at firm-group level, $1\{Risk_r\}$ 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 $r = \{high\, risk; low\, risk\}$. 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 time. I consider two outcomes: in columns (1)-(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. Since 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. This leads to a sample of around 3,000 firms depending on the split. Then, in columns (3)-(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 ten years before the decision. Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 7: Effect of the decision across firm size

(a) Heterogeneity by size: revenue

	(1)	(2)	(3)	(4)	(5)	(6)
			$R\&D_{jt}/Asset_{jt}$			
Median Revenue	*Small*	*Large*	*All*	*Small*	*Large*	*All*
$Post \cdot Exposure_j$	0.004**	-0.001	0.004**	0.006***	-0.001	0.005***
	(0.002)	(0.01)	(0.002)	(0.002)	(0.001)	(0.002)
$Post \cdot Small_j$			-0.004**			-0.002
			(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	Y	Y	Y
$Indu. \times Time\, F.E.$				Y	Y	Y
$Other\, Controls_{jt}$				Y	Y	Y
R^2	0.016	0.007	0.022	0.120	0.076	0.072
Observations	956	1,078	2,034	956	1,078	2,034

(b) Heterogeneity by size: employment

	(1)	(2)	(3)	(4)	(5)	(6)
			$R\&D_{jt}/Asset_{jt}$			
Median Employment	*Small*	*Large*	*All*	*Small*	*Large*	*All*
$Post \cdot Exposure_j$	0.003**	-0.001	-0.001	0.006***	-0.001	0.002**
	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
$Post \cdot Small_j$			-0.003**			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	Y	Y	Y
$Indu. \times Time\, F.E.$				Y	Y	Y
$Other\, Controls_{jt}$				Y	Y	Y
R^2	0.015	0.003	0.021	0.111	0.059	0.072
Observations	969	1,065	2,034	969	1,065	2,034

These Tables report the estimate of the linear difference-in-difference specification (equation 3), 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 application 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 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 info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 8: Effect of the decision across measures of financial constraint

(a) Heterogeneity by dividend payers

	(1)	(2)	(3)	(4)	(5)	(6)
			$R\&D_{jt}/Asset_{jt}$			
	No Dividend	*Dividend*	*All*	*No Dividend*	*Dividend*	*All*
$Post \cdot Exposure_j$	0.004**	-0.002	0.005**	0.006***	-0.003	0.005***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.003)	(0.002)
$Post \cdot 1\{Div_j = 0\}$			-0.005***			-0.004**
			(0.002)			(0.002)
$Post \cdot Exposure_j \cdot 1\{Div_j = 0\}$			0.005**			0.005*
			(0.002)			(0.003)
$Firm\,F.E.$	Y	Y	Y	Y	Y	Y
$Time\,F.E.$	Y	Y	Y	Y	Y	Y
$Indu. \times Time\,F.E.$				Y	Y	Y
$Other\,Controls_{jt}$				Y	Y	Y
R^2	0.019	0.024	0.019	0.104	0.087	0.069
Observations	1,322	712	2,034	1,322	712	2,034

(b) Heterogeneity by rating status

	(1)	(2)	(3)	(4)	(5)	(6)
			$R\&D_{jt}/Asset_{jt}$			
	No Rating	*Rating*	*All*	*No Rating*	*Rating*	*All*
$Post \cdot Exposure_j$	0.003**	-0.001	0.003**	0.005***	-0.001	0.005***
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
$Post \cdot 1\{Rating_j = NO\}$			-0.002**			-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	Y	Y	Y
$Indu. \times Time\,F.E.$				Y	Y	Y
$Other\,Controls_{jt}$				Y	Y	Y
R^2	0.014	0.013	0.014	0.090	0.092	0.066
Observations	698	1,336	2,034	698	1,336	2,034

These Tables report the estimate of the linear difference-in-difference specification (equation 3), 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. Panel (a) 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. In Panel (b), I do the same 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 info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table 9: Stock Market returns and Litigation Exposure

	Event Day		*Event* $[-1;+1]$		*Event* $[0;+5]$		*Event* $[-5;-1]$	
	Ret.	*A.Ret.*	*Ret.*	*A.Ret.*	*Ret.*	*A.Ret.*	*Ret.*	*A.Ret.*
$Exposure_j$	0.012***	0.013***	0.013***	0.011***	0.008**	0.006*	0.001	-0.001
	(0.004)	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)	(0.003)
Observations	986	986	986	986	986	986	986	986

The table reports cross-sectional value-weighted regressions between litigation exposure and returns. Returns are measured either raw or as abnormal returns, where this is constructed as $r_j - \beta_j r^{S\&P500}$, where $\beta$are estimated by firm regressions between one months and twelve months before the decision. Furthermore, returns are measured over different windows, which are reported in the header of the table. Returns are 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. The weights are given by the firm market value of equity seven days before the decision. Standard errors are robust to heteroskedasticity. More info on the variables are provided in the Appendix (A.4). All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

A Appendix

A.1 Other Figures and Tables

Figure A.1: Number of papers citing “eBay vs. MercExchange

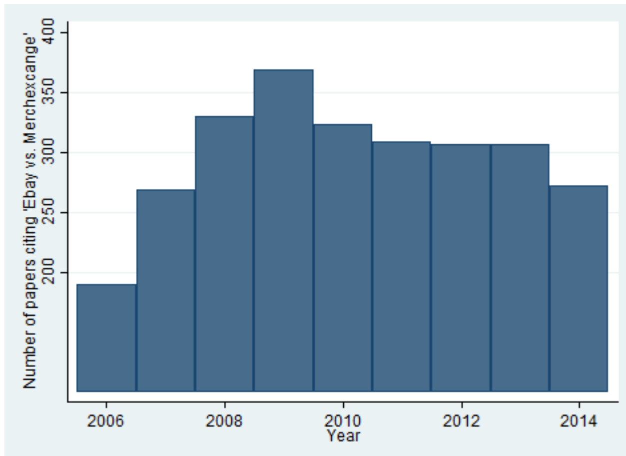

This Figure reports the number of papers citing the case “eBay vs. MercExchange” between 2006 and 2014. The total number of papers is about 2673. The search has been performed using Google Scholar on September 2015. In particular, I have search the key work “eBay vs. MercExchange” and extracted the data by year, as organized by Google.

Figure A.2: Persistence of patent litigation over time

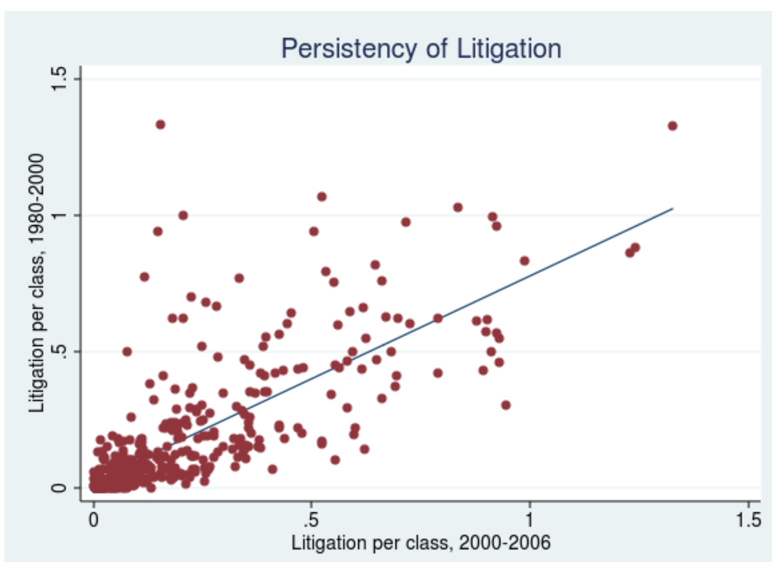

This figure provides a scatter plot of the size of litigation technology class level, as measured by equation (2), measured over two samples. In the vertical axis, I measure it using lawsuits between 1980 and 2000. In the horizontal axis, I use data between 2000 and 2006 (excluded). More information for the construction of this measure is provided in Section (3). For the clarity of the figure, I used every technology class with score p_c lower than 1.5. The blue line in the figure is the linear fit of the data, which has a coefficient of 1.05 in this case.

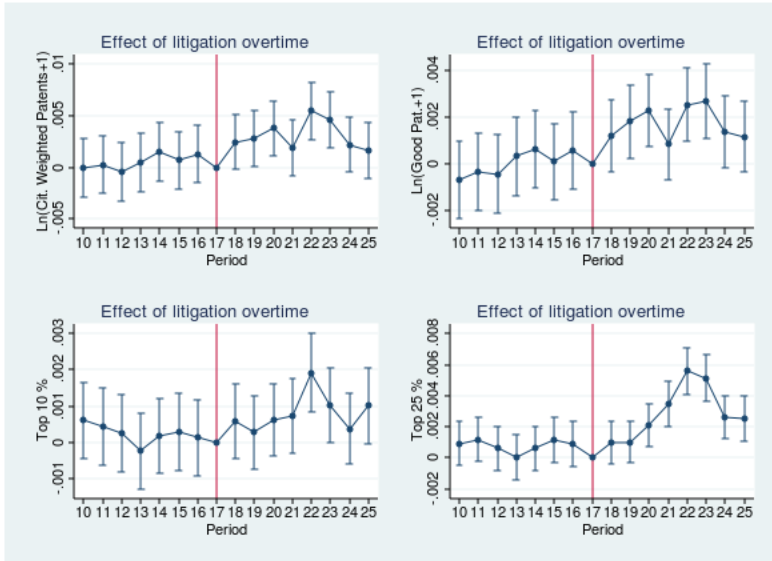

Figure A.3: Effect of litigation over time

This Figure plots the β_t from equation 4, using the usual sample. With respect to the other Figure (A.3), this is identical but with different outcomes. In order, I consider as outcomes the number of good patents, citation-weighted patents and dummies for firm patenting at top of the distribution (10% and 25%). More detailed description of the outcomes are in Appendix (A.4). The red vertical line correspond to the last period of the pre-decision period. 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). Data used corresponds at the two years before and after the decision.

Figure A.4: Returns NPEs-alternative specifications

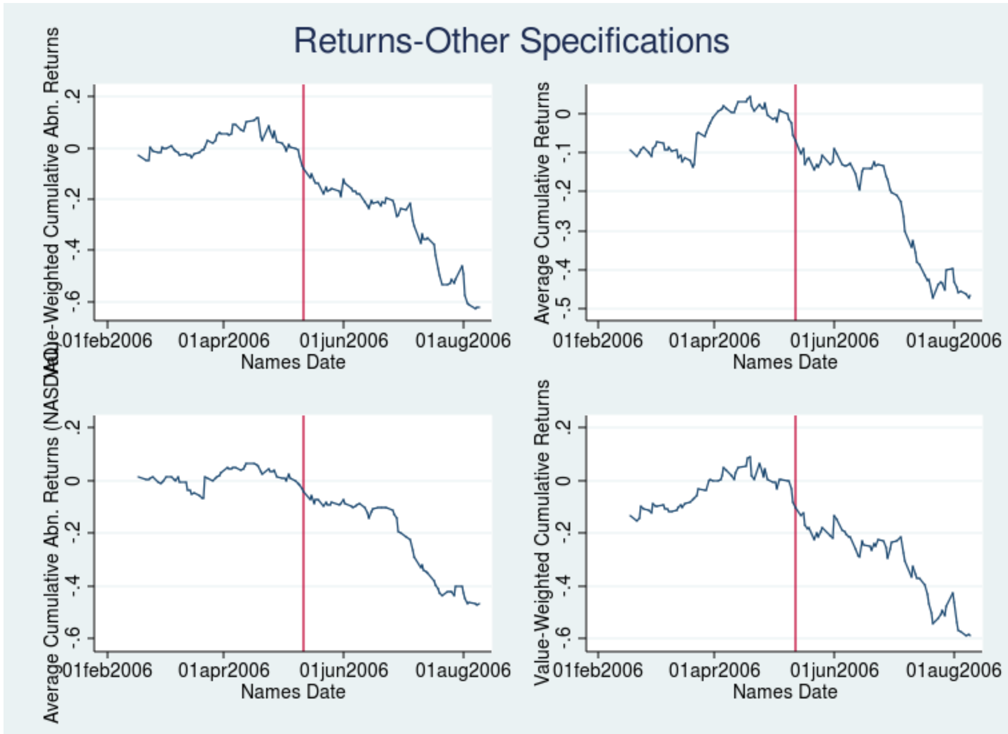

These Figures plot the average cumulative abnormal returns, for the sample of NPEs identified in the paper. The sample of firms used are 10 companies that are (a) Identified as NPEs; (b) Public at the time of the Supreme Court decision. The companies are Acacia Technologies, Asure Software, Rambus, Tessera Technologies, VirnetX Holding Corp., Universal Display, Document Security Systems, Pendrell, ParkerVision, Unwired Planet, Interdigital, Spherix. Information on the sample constructions are provided in Section (2.2). Abnormal returns are constructed with respect to the S&P500, as discussed in the Appendix (A.4.3). The straight red line correspond to the trading day right before the decision.

Table A.1: Distribution of p_i

	Obs.	Mean	SE	1%	10%	25%	50%	75%	90%	99%
$p_i^{1980-2006}$	438	0.22	0.35	0	0.01	0.03	0.1	0.27	0.60	1.77
$p_i^{2000-2006}$	438	0.22	0.43	0	0	0.02	0.08	0.24	0.61	2.30

This Table reports construction of technology-class size of patent litigation, as it is described in Section (1), and in particular by equation (2).

Table A.2: Stock Market returns and Litigation Exposure

	Event Day		*Event* $[-1;+1]$		*Event* $[-5;-1]$		*Event* $[-20;-5]$		*Event* $[-40;-5]$	
	Ret.	*A.Ret.*	*Ret.*	*A.Ret.*	*Ret.*	*A.Ret.*	*Ret.*	*A.Ret.*	*Ret.*	*A.Ret.*
Mean	-0.034***	-0.038***	-0.036**	-0.033**	-0.076***	-0.064***	-0.026	-0.071	0.129*	0.063
	(0.008)	(0.008)	(0.012)	(0.013)	(0.015)	(0.014)	(0.045)	(0.046)	(0.057)	(0.054)
Observations	10	10	10	10	10	10	10	10	10	10

This Table reports the average returns -either raw or abnormal-over a specific time span for the set of NPEs considered in Section (2), and a t-test for the difference from zero of the average. Standard errors are robust to heteroskedasticity. Abnormal returns refer to abnormal returns with respect to the S&P 500. More info on the test is available in Appendix (A.4.3). *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.3: Robustness: Leave-out-Industry

	$ln(Patents_{jt})$											
Leave out:	*Chemical*		*Computer*		*Drug*		*Electronic*		*Mechanical*		*Others*	
$Post \cdot Exposure_j$	0.042***	0.039***	0.038***	0.033***	0.060***	0.031**	0.043***	0.035***	0.034***	0.033***	0.039***	0.036***
	(0.009)	(0.012)	(0.009)	(0.012)	(0.016)	(0.017)	(0.009)	(0.011)	(0.009)	(0.011)	(0.009)	(0.011)
$Firm\,F.E.$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Time\,F.E.$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Indu. \times Time\,F.E.$		Y		Y		Y		Y		Y		Y
$Other\,Controls_{jt}$		Y		Y		Y		Y		Y		Y
Observations	28,408	28,408	25,938	25,938	26,448	26,448	25,982	25,982	26,204	26,204	27,610	27,610

In this table we report the estimation of the equation 3. The data set is constituted by a balanced two-period panel. The first and second period are the collapse of firm information in the two years before and two years after the Supreme Court decision. The outcome is the (natural) logarithm of granted patent that firm j applied during period t. In this case, I use 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. In every column I drop one industry, as it is reported in the header of the column. Industries are constructed based on the technology class of publication over the previous four years, as previously defined (Hall et al. (2001). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.4: Effect of the policy change on patenting: Poisson model

	(1)	(2)	(3)
Poisson		$\#Patents_{jt}$	
$Post \cdot Exposure_j$	0.043***	0.047**	0.040**
	(0.015)	(0.021)	(0.020)
$Firm\,F.E.$	Y	Y	Y
$Time\,F.E.$	Y	Y	Y
$Indu. \times Time\,F.E.$		Y	Y
$Other\,Controls_{jt}$			Y
Observations	256,944	256,944	256,944

This Table reports the estimate of the standard difference-in-difference specification (equation 3) using an equivalent fixed-effect Poisson model. The properties of the Poisson model implies that the parameter β on the main variable of interest $Post \cdot Exposure_j$ can be interpreted as a semi-elasticity, similarly to the log-linear difference-in-difference model previously estimated. In this model, the outcome is the number of granted patent applications made by firm j in period t. The data set is a balanced quarterly panel over the same set of innovative firms employed before. 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 1990. In every specification, I essentially control for both firm and quarter fixed effects. In Column (2), I add industry-time fixed effect to the equation. Industry 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), I further augment the specification using location dummies of the firm (constructed using the modal location reported in patent data), the size of the portfolio before the estimation period, the start-up status (looking at whether the 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 the decision. More info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.5: Effect of the policy change on patenting: robustness with alternative *Exposure* measure

	(1)	(2)	(3)	(4)	(5)	(6)
OLS		$ln(Patents_{jt})$			$1\{Patent_{jt} > 0\}$	
$Post \cdot Exposure_j^{LARGE}$	0.050***	0.049***	0.047***	0.014***	0.037***	0.037***
	(0.012)	(0.016)	(0.015)	(0.003)	(0.003)	(0.003)
$Firm\,F.E.$	Y	Y	Y	Y	Y	Y
$Time\,F.E.$	Y	Y	Y	Y	Y	Y
$Indu. \times Time\,F.E.$		Y	Y		Y	Y
$Other\,Controls_{jt}$			Y			Y
R^2	0.005	0.007	0.033	0.319	0.359	0.394
Observations	32,118	32,118	32,118	155,866	155,866	155,866

In this Table I replicate the estimates from Table (2), using an alternative measure to firm exposure to litigation. In particular, I estimate equation (3), which is $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt}$, where y_{jt} is: (a) the (natural) logarithm of granted patent that firm j applied during period t for Columns (1)-(3); (2) a dummy equal to one if the firm j applied to at least one patent in period t. The variable $Exposure_j^{LARGE}$captures the exposure of firm j to patent litigation, using patent application in the ten years before the decision and patent litigation at technology class since 1980. 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 published at least one patent in the two year 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. In Columns (1) and (4), I control for firm fixed-effects and time effects. In Column (2) and (5), I add industry-time fixed effect to the equation. Industry 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 modal location reported in patent data), the size of the portfolio before the estimation period, the start-up status (looking at whether the 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 info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.6: Robustness: differential linear effect before and after the shock

	(1) $ln(Patents_{jt}+1)$	(2) $Patent^{scaled}_{jt}$	(3) $1\{Patent_{jt}=Top^{10\%}\}$	(4) $1\{Patent_{jt}=Top^{25\%}\}$	(5) $ln(Citation\,Weighted\,Pat_{jt}+1)$
$Post \cdot Exposure_j$	0.012***	0.010***	0.004*	0.010***	0.010*
	(0.004)	(0.003)	(0.002)	(0.003)	(0.005)
$Pre \cdot Exposure_j$	0.004	0.002	0.001	0.001	0.004
	(0.004)	(0.003)	(0.002)	(0.003)	(0.005)
$Firm\,F.E.$	Y	Y	Y	Y	Y
$Time\,F.E.$	Y	Y	Y	Y	Y
R^2	0.009	0.016	0.001	0.002	0.002
Observations	256,944	256,944	256,944	256,944	256,944

In the table I report the estimation of an equation where I use the data as a panel and I estimate the same specification as equation (3), but where I interact the risk exposure measure with both a dummy for after the decision (equal to one for quarters after May 15th 2006) and a dummy for the quarters before 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. Therefore, the interaction compares corporate behavior before and after the decision to the behavior in the quarter that concluded with the decision. The data set is constituted by a panel with eight quarter and it is balanced in any specification. In any case, I use every firm that published at least one patent in the two year before and in the year after the decision. Column (1) has the (natural) logarithm plus one of granted patent that firm j applied during period t. Column (2) has granted patent that firm j applied during period t, scaled by the total number of patents in the two years before the decision. Columns (3) and (4) have the dummy which is equal to one whether the firm j applied during period t at least to one patent that is in the top 10% or 25% of the matched patents (same year and same technology class). Column (5) has the (natural) logarithm plus one of granted patent that firm j applied during period t weighted by citations received in the first three years of life. Citations are scaled by the average number of citations received by the patents in the same technology class and year. More info on the data are available in the Appendix (A.4).All regressions include a constant. Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.7: Timing of the effects

	$ln(Patents_{jt})$					
	1 Year After		2 Years After		3 Years After	
$Post \cdot Exposure_j$	0.001		0.035***		0.051***	
	(0.001)		(0.010)		(0.010)	
$Post \cdot Exposure_j \cdot Computer$	0.094***		0.026		-0.003	
	(0.020)		(0.020)		(0.022)	
$Post \cdot Exposure_J^{LARGE}$		0.004		0.052***		0.072***
		(0.013)		(0.014)		(0.010)
$Post \cdot Exposure_j^{LARGE} \cdot Computer$		0.085***		0.016		-0.008
		(0.027)		(0.029)		(0.030)
$Firm\, F.E.$	Y	Y	Y	Y	Y	Y
$Time\, F.E.$	Y	Y	Y	Y	Y	Y
R^2	0.215	0.215	0.006	0.006	0.117	0.117
Observations	32,118	32,118	32,118	32,118	32,118	32,118

In this table I report the estimation of the equation 3, where I interact the shock measure with a dummy for firms that are in the Computer industry, as defined in Appendix A.4. The data set is constituted by a balanced two-period panel. The first period is fixed to the two year before the decision, while the second period depends on the specification and in particular it moves from 1 to 3 years after. The outcome is always the (natural) logarithm of granted patent that firm j applied during period t. In this case, I use 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. Similarly, the variable $Exposure_j^{LARGE}$captures the exposure of firm j to patent litigation, using patent application in the ten years before the decision and patent litigation at technology class since 1980. Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.8: Robustness: differential linear effect before and after the shock

	(1)	(2)	(3)	(4)
		$R\&D_{jt}/Asset_{jt}$		
$Post \cdot Exposure_j^{LARGE}$	0.004**	0.006**		
	(0.002)	(0.003)		
$Pre \cdot Exposure_j^{LARGE}$	0.001	0.001		
	(0.002)	(0.002)		
$Post \cdot Exposure_j$			0.003**	0.004**
			(0.001)	(0.001)
$Pre \cdot Exposure_j$			0.002	-0.001
			(0.001)	(0.002)
$Firm\&Time\,F.E.$	Y	Y	Y	Y
$Indu. \times Time\,F.E.$		Y		Y
$Other\,Controls_{jt}$		Y		Y
R^2	0.011	0.025	0.005	0.016
Observations	16,272	16,272	16,272	16,272

In the table I report the estimation of an equation where I use the data as a panel and I estimate the same specification as equation (3), but where I interact the risk exposure measure with both a dummy for after the decision (equal to one for quarters after May 15th 2006) and a dummy for the quarters before the decision. Therefore, the interaction compares corporate behavior before and after the decision to the behavior in the quarter that concluded with the decision. The data set is constituted by a panel with eight quarter and it is balanced in any specification. 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 and the variable $Exposure_j^{LARGE}$captures the exposure of firm j to patent litigation, using patent application in the ten years before the decision and patent litigation at technology class since 1980. In any case, I use every firm that published at least one patent in the two year before and in the one year after the decision. The table has have $R\&D/Asset$, measured at quarterly frequency. The even columns are augmented with industry, as constructed in the Appendix, interacted with time dummies (per quarter). More info on the data are available in the Appendix (A.4). All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.9: Evidence on Patent Mix: pre-trend analysis

	(1)	(2)	(3)	(4)
	Extensive Margin		Intensive Margin	
	$ln(Patents_{jtr})$		$1\{Patents_{jtr} > 0)$	
$Post \cdot 1\{Risk_r\}$	-0.019	-0.009	0.040***	0.027***
	(0.017)	(0.015)	(0.002)	(0.002)
$Pre \cdot 1\{Risk_r\}$	-0.016	0.0015	-0.002	0.001
	(0.018)	(0.015)	(0.003)	(0.002)
$Split$	10%	25%	10%	25%
$Firm \times Time\, F.E.$	Y	Y	Y	Y
$Firm \times Risk\, F.E.$	Y	Y	Y	Y
R^2	0.924	0.913	0.688	0.687
Sample	2,785	3,893	54,844	54,844
Observations	89,120	124,576	1,755,008	1,755,008

This Table provides a study of pre-trending for results reported in Table (4). In order to do so, I estimate the same specification as before without collapsing the data and estimating a differential effect for the treatment before and after the decision. Data are at quarterly level in a four year around the decision, for a total of 16 periods. In practice, I estimate: $y_{jtr} = \alpha_{jr} + \alpha_{jt} + \beta^{POST} 1\{Risk_r\} \cdot Post + \beta^{PRE} 1\{Risk_r\} \cdot Pre$, where where α_{jt} is a set of firm-time fixed effects, α_{jr} is a set of fixed effects at firm-group level, $1\{Risk_r\}$ is a dummy for more risky groups. Here $Post$ identifies quarters after the decision and Pre those before the decision. The quarter of the decision - 20061- is the reference period for interpreting the coefficients. 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 $r = \{high\, risk; low\, risk\}$. 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%. I consider two outcomes: in columns (1)-(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. Since 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. This leads to a sample of around 3,000 firms depending on the split. Then, in columns (3)-(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 ten years before the decision. Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.10: Heterogeneity of the effects: growth

	(1)	(2)	(3)	(4)
		$R\&D_{jt}/Asset_{jt}$		
	$Standard$	$Controls\,Growth$	$Standard$	$Controls\,Growth$
$Post \cdot Exposure_j \cdot Small_j^{REVENUE}$	0.004**	0.004**	0.003*	0.003**
	(0.002)	(0.002)	(0.002)	(0.002)
$Post \cdot Exposure_j \cdot Small_j^{EMP}$	0.004**	0.004**	0.003	0.003*
	(0.002)	(0.002)	(0.002)	(0.002)
$Post \cdot Exposure_j \cdot 1\{Dividend_j = 0\}$	0.005**	0.008***	0.005*	0.008***
	(0.002)	(0.002)	(0.003)	(0.002)
$Post \cdot Exposure_j \cdot 1\{Rating_j = NO\}$	0.004**	0.003**	0.002	0.003
	(0.002)	(0.002)	(0.002)	(0.002)
$Firm\,F.E.\ \&\ Time\,F.E.$	Y	Y	Y	Y
$Indu. \times Time\,F.E.$			Y	Y
$Other\,Controls_{jt}$			Y	Y
Observations	2,034	2,034	2,034	2,034

These Tables report the estimate of the coefficient β_1 of the following equation:

$$y_{jt} = \alpha_j + \alpha_t + \beta_1(Exposure_j \cdot FinCon_j \cdot Post) + \beta_2(FinCon_j \cdot Post) + \beta_3(Exposure_j \cdot Post)$$

$$+\beta_4(Growth_j \cdot Post) + \beta_5(Growth_j \cdot FinCon_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt}$$

across different specifications. First, different rows measure financial constraint in different ways. In particular, I use: (1) size, as measure by firm below the median revenue in my sample; (2) size, as measure by firm below the median employment in my sample; (3) dividend, where I look at firms that paid no cash dividends in any quarters in the three years before the decision ; (4) rating, where I sort based on whether the firm has any rating. Second, in columns (1)-(3), I report the standard results I have already reported, and in columns (2)-(4) I introduce a fully interacted control for firm growth over the pre-period. This measure is the simple growth of revenue over the two years of pre period. Even if not reported, all the regressions are estimated as fully interacted. 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. 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 info on the variables are provided in the Appendix (A.4). Standard errors are clustered at firm level. All regressions include a constant.*** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

A.2 Background information on "eBay versus MercExchange"

The main object of the dispute in the 2006 "eBay vs. MercExchange" case was a patent on the popular "Buy It Now" function on the eBay platform.[47] In the early 2000s, MercExchange accused eBay of infringing some of the company's online auction patents. In 2003, the Virginia Circuit Court agreed with these accusations, but then decided to reject MercExchange's request to issue an injunction on eBay's technologies (Court, 2003). However, this decision was subsequently reversed by the Court of Appeals, which clearly stated that the issuance of a permanent injunction, "absent exceptional circumstances," was a general rule in the U.S. intellectual property enforcement system. In 2005, eBay decided to petition this decision in front of the Supreme Court, which agreed to discuss the case in the following year. As I discuss in the paper, the final ruling of the Court rejects the idea that injunction should always be issued in normal cases after a patent violation.

In fact, in line the opinion of the Court of Appeals, injunction was issued almost automatically after a violation was proved before 2006. This idea dates back to a 1908 Supreme Court case between Continental Paper Bag and Eastern Paper Bag (Court, 1908). In this case, where the Supreme Court clearly states that "exclusion may be said to have been of the very essence of the right conferred by the patent, as it is the privilege of any owner of property to use or not use it, without question of motive." In practice, the only cases where a firm would not receive injunction is when the firm could successfully argue in favor of the public interest of its products. Otherwise, the ruling granted to patent owner full ability to exclude others from using the technologies covered by the patent.

A.3 The timing of the decision

The ruling should affect the trajectory of innovation only if it was not anticipated by agents in this market.[48] In this section I argue that this was the case using different pieces of qualitative evidence from news sources and other public records. This is consistent with the large body of research in law that discusses the decision (e.g. Bessen and Meurer 2008a; Holte 2015; Shapiro 2010; Tang 2006; Venkatesan 2009)

As a first step, I review the news about the case and I do not find any evidence that the content of the ruling was anticipated. This is the case both when looking at news published in the weeks before the decision and right after it. If anything, the news the day of the decision appeared to find the ruling surprising.[49]

[47] The patent is USPTO number 5,845,265.

[48] The activities of the Supreme Court are planned in advanced and therefore the general public knew that the case "eBay vs. MercExchange" was under review. At the time, there was not an exact calendar, but generally there was agreement that the final decision was going to be taken most likely by the end of June. See, for instance, the article "Supreme Court to Take Another Look at "Automatic" Injunctions for Prevailing Patent Owners in Infringement Cases", appeared online December 12th 2005 in the Mondaq Business Briefing, a news provider for legal expert.

[49] MarketWatch defined the ruling "a surprising turn" of the Court. MarketWatch was accessed through Bloomberg "Supreme Court Rules for eBay in Patent Case: Expert Lawyer Calls Decision Surprising," May 15th 2006

Furthermore, looking across different parties that had some interest in the case, I find that the opinions about the case were generally divergent. First, the Justices appeared to be divided during the oral hearing. For instance, Patently-O, one of the most reputable patent law blogs, claimed that "based on oral arguments, pundits see a potential split decision in the eBay v. MercExchange injunction case."[50] In particular, Justice Scalia appeared to be in favor of considering an injunction as automatic after an infringement: "we're talking about a property right here, and a property right is the exclusive right to exclude others." Interestingly, in the end the Supreme Court made the decision unanimously. Second, in the weeks before the decisions, the government took a clear stand against eBay, and therefore in disagreement with what the Court later decided. In particular, on March 10th – two weeks before the oral argument of the case - the Office of the Solicitor General (OSG), on behalf of both the Federal Trade Commission and the antitrust division of the Justice Department, asked to confirm the injunction to eBay.[51] While the opinions of the OSG are in no way binding for the Supreme Court, they have an impact on the public perception of these issues. Third, even the business community was split on this issue. While large drug companies were opposing any change in the way injunction was issued, large companies in high tech – such as Intel, Cisco, Hewlett-Packard, Microsoft– were explicitly supportive of eBay.[52]

A.4 Data

A.4.1 Samples

This is a more detailed clarification on Section (3) in the paper, where I discuss the data and variable construction.

Firm level data comes from two sources. Patent data comes from the Fung Institute (University of California at Berkeley),[53] and they are an updated version of the Harvard Business School Patent Network Database (Li et al., 2014). In particular, I use all the assigned granted patents in the data, which does not have missing information on the grant date, application date, assignee ID and technology class. It is worthy to point out that all analyses are carried based on application date, since I am interested in capturing firm behavior. The data were download in August 2014 and they contain all patents that were granted before 29th April 2014. For the full sample, I define a firm based on the assignee identifiers in the data for the analysis using the full set of innovative firms. To evaluate quality of the data, I compare them to the aggregate

[50] See the article "eBay v. MercExchange Oral Arguments," from March 31st 2006, which can be found at the following address http://patentlyo.com/patent/2006/03/ebay_v_mercexch_3.html

[51] See for reference, Washington Post article on this issue "Government Sides Against eBay in Patent Dispute," March 11th 2006. A copy is available online at the following link: http://www.washingtonpost.com/wp-dyn/content/article/2006/03/10/AR2006031001918.html

[52] Helm (2006) argues that these difference stems from the different use of injunction for large firms across these industries.

[53] Data can be found: http://funginstitute.berkeley.edu/tools-and-data (downloaded in August 2014)

statistics that USPTO provides online.[54] In particular, I compare patents granted to corporations according to USPTO aggregate data to the data used in this paper. I find that the two series almost overlap across the whole period and strongly co-move over 2002-2008 period. In the year, where they differ the most, the difference is only about of 2%.[55]

Most of the analyses in the paper are carried using a sample of innovative firms, which are firms active in patenting before and after the decision. In particular, I define innovative firms as firms that applied to at least one patent in the two years before the decision and one in the year after this. The advantage of this is to have a sample that is the same when analyzing different sample period (one, two or three years after the decision). Furthermore, notice that this sample is intrinsically balanced when I do the analysis considering data collapsed before and after the decision. When I am interested in the intensive margin, I need to consider a larger set of firms. In particular, I take firms that have at least one patent in the four years before the decision, but not necessarily anything afterwards.

In the second part of this work, I supplement patent level data with information on R&D at firm level for the subset of innovative firms, for which public information are available. Data on firms' financials come from Compustat quarterly data. This allows me to construct pre and post period windows that are exactly around the Supreme Court decision. In order to add patent information to Compustat data, I use the data provided in Kogan et al. (2012). I construct a bridge file which is based on patent ID: this approach does not have the concerns of a matching performed by name. Essentially, I match the two data sets based on USPTO patent ID - as defined in the Fung Institute data - and then I use this match to bridge the assignee IDs to the ID used in Kogan et al. (2012). Since the assignee ID in the patent data is based on name disambiguity, one firm in patent data may correspond to more than one company ID in Compustat: therefore, the analysis at firm level use the more aggregate Compustat ID. Furthermore, in about 90% of the cases, the company ID in the patent data corresponds to only one Compustat ID. In the remaining cases, I use the Compustat ID that received more unique matching over the period considered. An hand-check of the data supports the quality of this choice.

In order to end up in the final sample, I apply some of the common filters. In particular, I focus on non-financial companies and non-government related companies, with the headquarter in USA. Furthermore, I exclude firms that do not have a balance reporting around the decision. In particular, while some data entry may be missing for random reasons, total assets and revenue should always be populated. I therefore eliminate those firms that do not have balance reporting on these variables in the four-year symmetric window around

[54] Table can be found here: http://www.uspto.gov/web/offices/ac/ido/oeip/taf/h_at.htm#PartA2_1

[55] This difference can stem from two things. First, it is not super clear how USPTO categorize companies, so there may be some discrepancy in this dimension Second, I would expect that the bulk of difference is probably made up of patents that had missing info in the micro data, such as missing date or technology.

the shock. Lastly, I want to exclude those companies that may be under financial distress or restructuring: in order to do this, I exclude companies that systematically report negative equity over the usual period. Furthermore, I focus only on the same of innovative firms, similarly to the previous analysis. All in all, I have a sample of more than one thousand firms. Using the Compustat IDs, I then also match the firm to stock returns information from CRSP.

In the end, as discussed in Section (3), I use patent lawsuits data from public filings to construct the measure of litigation size at technology class level. The data are collected from WestLaw, a subsidiary of Thomson Reuters. Westlaw is one of the primary provider of legal data in United States and use public records to develop a complete overview of lawsuits in United States. The same data, also known as Derwent LitAlert data, were previously used by other empirical work on patent litigation (e.g. Lerner 2006; Lanjouw and Schankerman 2001).

Using the online tool LitAlert, I searched for all the litigation involving patents between 1980 and 2006.[56] Every filing should report the date of the filing, the plaintiffs, the defendants and information on the intellectual property that is used to go to court. As a preliminary step, I eliminate the few filings with missing information about the date. To avoid issues with duplicates, I keep only one case in situations where multiple observations share the same entries for plaintiff, defendant, filing data and patents. Since I am interested in utility patents, at this point I keep only filings that report at least one utility patent. As discussed in Section (3), I make filings comparable across each other by reshaping the data at plaintiff-defendant-patent level. Then, I match patents with their technology fields and I aggregate them at technology class level over different periods of time. I then use equation (2) to construct the final score at technology class level.

A.4.2 Variables definition

In the analysis involving the full sample of innovative firms, I use various outcomes.

For measuring intensity of innovation, I look at two measures. First, I look at the logarithm of the patents produced by the firm j at time t, $ln(pat_{jt})$, which is consistent to an intensive margin of our treatment. Using this outcome, I consider the sample of every firm, either private or public, which applied to at least one patent before and after the shock, as previously discussed. In this sample, there are slightly more than 16 thousand firms that satisfy this condition. Consistent with the literature, I count patents weighting them based on the number of assignee to which the patent is granted. In particular, I weight assignee equally. However, results are completely unaffected when I use a normal patent count, where I count patents as one even when assigned to multiple parties. Second, in order to estimate something closer to an extensive margin of the treatment, I consider an alternative outcome variable, which is a dummy equal to one when the firm has applied to any

[56] http://intranetsolutions.westlaw.com/practicepages/template/ip__litalert.asp?rs=IPP2.0&vr=1.0

granted patent in the period, $1\{Patent_{jt} > 0\}$. In order to measure exposure to litigation for a firm, this has to have to at least one (applied) patent before the shock. Because of this, as I discussed before, I consider the set of firms that has at least one patent in the four years before the Supreme Court decision, for a total of around 77 thousand firms.[57] Results do not change if I shrink the window by looking only at three years before the decision or I increase it to six years before the decision.

For measuring quality of innovation, I construct few metrics based on patent citations. Following the literature in this area, I count patent citations at a fixed window -3 years - after the granting (e.g. Bernstein, 2015). I then construct various outcomes based on this. I consider three main outcomes. First, I construct a patent count measure, where patents are weighted by the number of citations received. In particular, I adjust citations by scaling them by the average number of citations that other assigned patent in the same technology classes and applied in the same year received. In my case, results with scaled and unscaled citations do not differ much. Second, I look at good patents, which are simply patents that received at least one citation. This result is not reported as an outcome in this updated version of the paper, but only to check pre-trending. Lastly, I look at the probability that a company applies to patents that are at the top of the citation distribution in the relevant reference group as a proxy for breakthrough innovation. The reference group is composed by assigned patents that are the same USPTO technology class and were developed by the company in the same year, based on application date. I then look at whether the company has applied to any patent which is on the top 10% and 25% of the distribution of citations.

I also use the patent data to construct a set of other controls, which are used along the paper. I construct a new measure of industry of the firm, which is based on patent application, rather than self-reporting industry. The main advantage of this measure is that I can use it both across public and private firms. Firm j is assigned to a certain industry by looking at the major industry in which the firm has applied to the highest number of patents. In line with the literature, major industries are defined as in the Appendix (1) of Hall et al. (2001). I use patents in the four years around the decision for the analysis. Similarly, I define a measure of location of the firm based on patent data. In particular, I assigned to firm j the location c if location c is the modal location for the patents applied in the four years before the Supreme Court decision. An extra code is used for firms for which no state location can be determined. I also construct a measure of size of the portfolio of the firm in question. I do so, looking at patents that were filled in the two years before the estimation window. Clearly, I cannot use patenting before the decision inside the estimation window because it would be collinear with the outcomes.

In the second part of the paper, I then use a set of balance sheets variables. All balance sheet ratios

[57] The outcome variable is constructed looking only at the two years before and after the shock, as in the intensive margin measure.

are winsorize at 1% to ensure that results are not driven by outliers. My main measure of R&D intensity is R&D/Asset. R&D expenditure is measured using quarterly Compustat data (variable xrdq) and it is adjusted for acquisition of in process R&D expenses (variable rdipq), as in Mann (2013). Notice than the adjustment does not produce first-order effects in the outcome, as the share of firm-quarter with non-zero in process R&D expenses is, as expected, very small. The quarterly data are augmented, if necessary, with yearly data. These data are consistently adjusted at quarterly level assuming equal R&D across quarters within the fiscal year. Lastly, in line with the literature, missing R&D data is replaced with zero.

A.4.3 Stock Market data

When dealing with stock market data, I usually report the results both as raw returns and abnormal returns.

Raw returns are simply computed based on the standard stock returns. Abnormal returns are instead constructed relative to a benchmark, which is usually either the S&P500 or the NASDAQ. The S&P500 returns are also obtained from CRSP, while the NASDAQ data are obtained online from Yahoo Finance. In order to construct abnormal returns, I compute the predicted returns estimating the β of each stock using daily returns between 343 trading days before (January 1st 2005) and 30 trading days before the events. Conditional on providing a sufficiently large window to estimate the β precisely, results are not affected by the choice of the estimation window. When considering cumulative returns, I compute them as simple sum of the returns. Furthermore, when I use value-weights, I compute the weights based on equity capitalization seven trading days before the decision and keep them constant throughout.[58]

When I test the returns of NPEs around the event I report t-statistic, that tests the difference of the average returns from zero. This is constructed based on heteroskedasticity robust errors, and the estimation is implemented for simplicity using least-squares.

[58] I use the stock price and the number of shares provided by CRSP to compute the market value of equity.