



Essays in Corporate Finance

Citation

Mezzanotti, Filippo. 2016. Essays in Corporate Finance. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:33493570>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays in Corporate Finance

A dissertation presented

by

Filippo Mezzanotti

to

The Committee for the PhD in Business Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Business Economics

Harvard University

Cambridge, Massachusetts

April 2016

©2016 Filippo Mezzanotti

All rights reserved.

Dissertation Advisors:

Professor Josh Lerner

Professor Jeremy Stein

Author:

Filippo Mezzanotti

Essays in Corporate Finance

Abstract

Macroeconomic and institutional shocks are important drivers of firms' activities. In chapter one, I examine the role of patent litigation in affecting companies' innovation. Studying a landmark Supreme Court decision, I show that an improvement in patent enforcement positively affects the innovation activity of corporations. In chapter two, I study the role of private equity in period of large financial turmoil. In the context of the 2008 crisis in United Kingdom, I show that private equity backed companies experienced a lower decline in investment than a control group of similar companies that were not related to private equity. This effect is explained by the ability of private equity to relax the financing constraints of the portfolio companies when access to credit markets is limited. In chapter three, I explore the role of sovereign securities held by banks in the propagation of a financial shock to the economy. Using detailed loan level data matching firms and banks in Italy, the paper finds that the shock to banks' sovereign portfolio caused by the Greek bailout (2010) was passed on to firms through a contraction in credit. The effects of this shock were particularly disruptive for smaller companies.

Contents

Abstract	iii
Acknowledgments	xiii
Introduction	1
1. Roadblock to Innovation: The Role of Patent Litigation in Corporate R&D	3
1.1. Introduction	3
1.2. The “eBay vs. MercExchange” case	11
1.2.1. The Supreme Court decision	12
1.2.2. Injunction and innovation	16
1.2.3. The decision and non-practicing entities	17
1.3. Data	21
1.3.1. Firm level data	21
1.3.2. Measuring exposure of litigation at firm level	23
1.4. The effect of the Supreme Court decision on innovation	28
1.4.1. Empirical framework	29
1.4.2. The effect of the decision on innovation output	30
1.4.3. Identification assumptions	33
1.4.4. Robustness and other results	38

1.4.5.	Timing of the effects	40
1.4.6.	Evidence on patent quality	42
1.4.7.	Evidence from public firms	45
1.5.	How does litigation exposure affect innovation?	48
1.5.1.	Litigation lowers innovation returns: evidence from the composition of innovation	48
1.5.2.	Litigation exacerbates financial constraints	51
1.6.	Supreme Court decision and stock prices	57
1.7.	Conclusion	59
2.	Private Equity, Financial Strategy, and the Crisis¹	62
2.1.	Introduction	62
2.2.	Data	68
2.2.1.	Sample construction	68
2.2.2.	Control group	70
2.2.3.	Other data	74
2.3.	Empirical Strategy	75
2.4.	Investment Analysis	78
2.4.1.	Main Results	78
2.4.2.	Robustness	83
2.4.3.	Heterogeneity: investments and financial constraints	86
2.5.	Performance Analysis	91
2.5.1.	Change in performance	91
2.5.2.	Exit Analysis	93
2.6.	Conclusion	97

¹Co-authored with Shai Bernstein and Josh Lerner.

3. Sovereign debt exposure and the bank lending channel: impact on credit supply and the real economy²	99
3.1. Introduction	99
3.2. The onset of the sovereign crisis	106
3.3. Data	109
3.4. The bank lending channel	114
3.4.1. Identification strategy	114
3.4.2. The bank lending channel: main results and robustness	118
3.5. Credit supply and corporate behavior	137
3.5.1. Supply shock and credit market's access	139
3.5.2. Information frictions and market segmentation	142
3.5.3. Real effects on investments and employment	144
3.6. Conclusion	149
 Bibliography	 151
 A. Appendix to Chapter 1	 160
A.1. Background information on “eBay versus MercExchange”	160
A.2. The timing of the decision	161
A.3. Data	162
A.3.1. Samples	162
A.3.2. Variables definition	165
A.3.3. Stock Market data	168
A.4. Additional Figures	169
A.5. Additional Tables	172

²Co-authored with Margherita Bottero and Simone Lenzu.

B. Appendix to Chapter 2	182
B.1. Data and variable construction	182
B.2. Additional Figures	184
B.3. Additional Tables	185
C. Appendix Chapter 3	194
C.1. A Simple model of loan supply	194
C.2. Sovereign Holdings and Banks: A Cross-Country Comparison	195
C.3. Correlation Between Greek and Italian Sovereign Yields	198
C.4. Semi-parametric estimation of the bank lending channel	200
C.5. Estimates of effect of the sovereign crisis on aggregates bank credit via the lending channel	203
C.6. Data selection and other information on data construction	204
C.7. Variables description	205
C.8. Additional Figures	208
C.9. Additional Tables	210

List of Tables

1.1. Summary Statistics	22
1.2. Effect of the policy change on patenting: main results	32
1.3. Timing of the effects	41
1.4. Evidence on patent quality	43
1.5. Effect of the decision on public firms	46
1.6. Evidence on Patent Mix	50
1.7. Effect of the decision across firm size	54
1.8. Effect of the decision across measures of financial constraint	56
1.9. Stock Market returns and Litigation Exposure	59
2.1. Summary Statistics: Level	71
2.2. Summary Statistics: Changes	73
2.3. Investment and funding policies	80
2.4. Investment heterogeneity and financial constraints	87
2.5. Investment heterogeneity across fund age measures	90
2.6. Accounting Performance	92
2.7. Exit Analysis	95
3.1. Summary Statistics - Loan-level and Firm-level variables	112
3.2. The Bank Lending Channel	120

3.3. Bank Lending Channel: Total Credit and Credit Lines	122
3.4. The Bank Lending Channel: alternative measures	126
3.5. The Bank Lending Channel w/ relationship controls	128
3.6. The Bank Lending Channel: Lender-Borrower Sorting and Single Lender Firms	132
3.7. Transmission Mechanism of The Bank Lending Channel: Capital and Funding Channel	136
3.8. The Firm Borrowing Channel	141
3.9. Real Effects: Sovereign Exposure on Investments and Employment	145
3.10. Real Effects: Size and Dependence on External Finance	148
A.1. Distribution of Litigation	172
A.2. Effect of the policy change on patenting: Poisson model	173
A.3. Stock Market returns and Litigation Exposure	174
A.4. Robustness: Leave-out-Industry	175
A.5. Effect of the policy change on patenting: alternative <i>Exposure</i> measure . . .	176
A.6. Robustness: differential linear effect before and after the shock	177
A.7. Timing of the effects across technologies	178
A.8. Robustness: differential linear effect before and after the shock	179
A.9. Evidence on Patent Mix: pre-trend analysis	180
A.10. Heterogeneity of the effects: robustness	181
B.1. Summary Statistics alternative matching: Level	185
B.2. Summary Statistics alternative matching: Changes	186
B.3. Robustness: alternative matching model	187
B.4. Robustness: no MBO	188
B.5. Robustness: only 2007-2008	189

B.6. Robustness: only no exit companies	190
B.7. Robustness: Industry time-varying effects	191
B.8. Exit across more vs. less equity injection	192
B.9. Activity across more vs. less equity injection	193
C.1. International Comparison of Banking Systems	199
C.2. Correlation between the Italian and Greek spread	200
C.3. Distribution of Firms Across Industries	205
C.4. Banks Characteristics and Sovereign Holdings	210
C.5. Banks Characteristics and Credit Supply	211
C.6. Pre-Trending and Placebo Regressions	212
C.7. The Bank Lending Channel: Firms' Heterogeneity	213
C.8. The Bank Lending Channel: Different Sovereign Holdings	214
C.9. Transmission Mechanism of The Bank Lending Channel: Foreign and BCC .	215
C.10.Fixed Effects and Demand-Side Shocks	216
C.11.Real Effect: Sovereign Exposure on Investment and Employment - Pre-Trending	217
C.12.The Firm Borrowing Channel: Adding controls	218
C.13.Real Effects: Adding Firm-Controls	219
C.14.Real Effects across External Dependence: Adding Controls	220
C.15.Real Effects: Labor Efficiency	221

List of Figures

1.1. Number of cases involving patents	12
1.2. NPEs Stock Returns around the decision	20
1.3. Distribution of firm industries for more exposed firms vs. whole sample . . .	24
1.4. Effect of litigation over time	34
1.5. Placebo test over time	36
1.6. Permutation Test: distribution of test statistic	38
1.7. Effect of litigation on R&D intensity over time	47
1.8. Stock Market Reaction: High vs. Low exposure	58
2.1. Industry Distribution	70
2.2. Investment: PE vs. non-PE	79
2.3. Equity Contribution: PE vs. non-PE	82
2.4. Debt: PE vs. non-PE	83
3.1. The burst of the sovereign crisis	108
3.2. The bank lending channel	123
3.3. Pre-trending test	124
3.4. Sovereign holdings and credit supply dynamics	130
A.1. R&D and Patenting by Corporations	169
A.2. Number of papers citing “eBay vs. MercExchange”	169

A.3. Persistence of patent litigation over time	170
A.4. Effect of litigation over time	171
A.5. Returns NPEs-alternative specifications	172
B.1. Private Equity Activity	184
C.1. Holding of Domestic Sovereign by Domestic Banks	198
C.2. The bank lending channel	202
C.3. 10-year spread of European bonds over German bonds	208
C.4. Investors pool: European Bank and Insurance CEO conference	209

Acknowledgments

I would like to thank my advisors - Professors Josh Lerner, David Scharfstein, Andrei Shleifer, and Jeremy Stein - for their encouragement and guidance during the past five years. Working with them has been an incredibly stimulating experience. I have also benefited a great deal from other faculty members at Harvard that over the years and in different roles have helped me during my PhD. In particular, I want to thank Professors Victoria Ivashina, Lauren Cohen, Adi Sunderam, Samuel Hanson, Alberto Alesina, and Gary Chamberlain.

I am also indebted to my co-authors, who have shared with me the stresses and joys of academic research. In particular, I want to thank Shai Bernstein, who co-authored the second chapter of this dissertation with Josh Lerner and I. In addition, I would like to thank Simone Lenzu and Margherita Bottero, my co-authors on the third chapter of this dissertation. Lastly, I would also like to thank James Lee, James Feigenbaum, and Giovanni Reggiani, who are my co-authors on other projects.

Graduate school would not have been as fun and rewarding without many of my classmates and colleagues. On top of the previously mentioned people, I want to thank Tommaso Denti, Xavier Jaravel, Divya Kirti, Ludovica GAZZE', Gordon Liao, Guilherme Lichand, Luca Maini, Diana Moreira, Vijay Narasiman, Gianpaolo Parise, Andrea Passalacqua, Daniel Pollmann, Tom Powers, Jonathan Rhinesmith, Rajesh Vijayaraghavan, Edoardo Teso, Yao Zeng, John Zhou, and many others. Outside Cambridge, a special thanks also go to my friends Nicola,

Diego, Raffaele, Vincenzo, Lorenzo, and Corrado, and to my “adopted” American family, the whole Gamez-Djokic clan.

This dissertation is dedicated to my family, in particular my fiancé Monica, my parents Alfio and Paola and my brother Francesco. Any achievement would have been impossible without their love and support. I hope that I can give them back as much as they have given me every single day. One last acknowledgment also goes to my grandparents, whose lives have always been for me a source of inspiration and motivation.

Introduction

This dissertation examines how various corporate activities – like innovation or investment decisions – are affected by the macroeconomic and institutional environment where companies operate.

In the first chapter of my dissertation, I examine the importance of patent litigation in affecting companies' incentive to innovate. This paper has been motivated by the recent spike in patent litigation, which has raised 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. Then, I show that patent litigation reduces investment in innovation by lowering the returns from R&D and by exacerbating financing constraints. This evidence confirms that patent litigation plays an important role in hindering innovation, and therefore that adjustments in the enforcement of patent law can have sizable effects on R&D.

The second chapter (with Shai Bernstein and Josh Lerner) studies how the corporate behavior of private equity backed companies was affected by the 2008 financial crisis. To investigate this issue, the paper uses a sample of about seven hundred British firms targeted by a private equity deal before the crisis and it compares them to a carefully selected control group. On average, during the financial crisis PE-backed companies experienced a lower decline in investment than the control group. This effect is related to the ability of the private equity sponsors to relax the financing constraints of the portfolio companies when access to credit markets is limited. Consistent with this idea, PE-backed companies also experienced a larger increase in overall debt and equity contributions during the financial crisis. Furthermore, the relative increase in investment is particularly strong for companies more likely to be financially constrained and when the private equity sponsor is more likely to have resources to help the portfolio company. However, the higher investments do not translate into better performance. These results shed new light on the role of the private equity industry during periods of financial turmoil.

Lastly, the third chapter (with Margherita Bottero and Simone Lenzu) examines the transmission of a bank balance sheet shock to corporate credit and its effects on real activity in the context of the European sovereign crisis. Using detailed loan level data matching firms and banks in Italy, the paper finds that the shock to banks' sovereign portfolio caused by the Greek bailout (2010) was passed on to firms through a contraction in credit. The contraction of the credit supply was similar in size for both large and small firms. However, it led to a reduction in investments and employment only for the smaller firms, especially those relying heavily on external financing. These effects were further exacerbated by the geographical segmentation of the credit market. Investigating the heterogeneity of the bank lending channel across financial intermediaries, we found a sharper tightening of credit supply among banks closer to the regulatory capital constraint.

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

1.1. 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 twenty years, lawsuits involving patents more than tripled (Figure 1.1) 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 & Meurer, 2008b; Boldrin & Levine, 2002; Jaffe & Lerner, 2011).

In this paper, I show that patent litigation has a real impact on innovation, and that im-

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

provements 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 & Shapiro, 2006; Shapiro, 2010). 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 & 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 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 & Lerner, 2011), demonstrating how injunction in intellectual property disputes can come at a large cost of firms.

This Supreme Court decision had a first-order impact on the way intellectual property is enforced.² For instance, Chien & 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. 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%.

Many scholars and practitioners considered the ruling an important step in fixing the distortions of the patent system, making investments in innovation easier and more valuable. 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.”³ However, the actual effectiveness of this ruling crucially depends on the ability of courts to handle these cases. Criticizing the position of the Supreme Court, other parties claimed that removing the automatic injunction would simply reduce deterrence against real violations (e.g. Epstein, 2008). In their view, the new system would just encourage more violations and therefore lower the incentives to invest. Given this potential competing effect, I turn to data to understand the impact of the decision.

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

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

³American 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.

exploits heterogeneity 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 & 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, 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 even patent something, and this corresponds to a 5% jump over 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 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 detriment of 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 increased R&D over assets by 8% more.

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 & 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.⁴ In the presence of financial frictions, the increase in costs due to

⁴Litigation claims “whether meritorious or not, (...) could require expensive changes in our methods of doing

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.*, 2011b; Claessens & Laeven, 2003; Demirgüç-Kunt & Maksimovic, 1998; King & Levine, 1993; La Porta *et al.*, 1997; Lerner & Schoar, 2005).⁵ 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; Iversen, 2014; Ponticelli, 2013). The role of enforcement is particularly important in intellectual property because the exact boundaries of patents are hard to define (Lemley & Shapiro, 2005) and lawsuits are therefore frequent (Lanjouw & Lerner, 1998). My paper confirms

business, or could require to enter into costly royalty or licensing agreements" (eBay 2006 10-K).

⁵This topic is related to the finance literature focusing on the relationship between litigation and corporate policies (Arena & Julio, 2014; Beatty *et al.*, 2008; Kim & Skinner, 2012; Haslem, 2005; Rogers & Van Buskirk, 2009).

that enforcement can be very important and suggests that, in a manner similar to other interventions (Acharya & 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 and contributes to the growing literature on this topic (Cohen *et al.*, 2014, Feng & Jaravel, 2015; Tucker, 2014).

This paper also adds to the body of empirical work that analyzes the relationship between property rights and innovation (Lerner, 2009; Sakakibara & 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,⁶ had beneficial effects on innovation incentives.⁷ 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 & 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

⁶During the argument, Justice Scalia said “We’re talking about a property right here, and a property right is the exclusive right to exclude others.”

⁷However, 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.

investments, and a hybrid system that provides more flexibility may be superior (Kaplow & 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 & Shleifer 2013).

The paper is organized as follows. In Section (1.2), I provide more background information about the Supreme Court decision, also discussing its potential effects on corporate innovation. In Section (1.3), 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 (1.4), I present the main results of my analysis. In Section (1.5), I discuss and test different channels through which patent litigation can affect innovation. In Section (1.6) I look at the stock market reaction around the decision. Lastly, Section (1.7) discusses policy implications and avenues for future research.


1.2. 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 content of the ruling in more detail in order to set the foundation for the hypothesis and the research design. Second, I discuss how the shock could have affected the trajectory of innovation. Lastly, I show that the decision had negative effects on the stock market valuation of patent assertion entities, which confirms that the ruling had first-order implications for patent enforcement.

1.2.1. The Supreme Court 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.⁸ 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 & Lemley, 2012).

Figure 1.1.: 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 (1.3).

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

⁸I provide some background legal information about the “eBay vs. MercExchange” case in Appendix (A.1).

uncommon and mostly due to public interest reasons. This is not surprising because patent law was derived from property law, where injunction is the standard method used to remedy violations.

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 & 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 case discussed in the introduction. First, although the dispute involved only a few patents, issuing an injunction would have forced the firm 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 it took a lengthy re-examination process that lasted several years to make this discovery. Altogether, NTP's ability to leverage on the near-mandatory injunction was the main driver to obtain the large settlement.⁹

⁹In 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."

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. Indeed, a large number of lawsuits today are initiated for purely opportunistic reasons (Cohen *et al.*, 2014; Jaffe & Lerner, 2011). These opportunistic players can leverage on strong injunction to hold-up defendants and obtain large transfers from them (Lemley & Shapiro, 2006; Shapiro, 2010).¹⁰ 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.

Consistent with this idea, patent assertion entities (PAE) were actively using the threat of permanent injunction as a way to scare counterparties and therefore obtain larger settlements (Lemley & Shapiro, 2006). Patent assertion entities are companies that hold large patent portfolios and are extremely active in intellectual property litigation. In particular, these companies, which are sometimes referred to as “patent trolls” in the popular press, tend to aggressively assert patents against cash-rich firms with the objective of obtaining large settlements.¹¹ As a result, the important role of patent assertion entities help explaining why permanent injunction was so relevant in fostering the costs and risks related to abusive lawsuits.

The idea that strong injunction can foster abusive litigation 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 a discussion of the case for a law journal, Wesenberg & O’Rourke

¹⁰This is a classical “hold-up” problem (Williamson 1975, 1983): at the time of litigation, the target company already made an irreversible investments and an injunction would either force the company to shut down a product line or modify it to comply with the court order.

¹¹While 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 & Jaravel, 2015; Tucker, 2015, 2014).

(2006) suggest:

“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

Previous research confirms that the decision had a large impact on the way intellectual property is enforced (Bessen & Meurer 2008a; Holte 2015; Shapiro 2010; Tang 2006; Venkatesan 2009). For instance, Chien & Lemley (2012) find that the ruling reduced the likelihood of obtaining an injunction by about 25%. The same study also shows that parties more likely to be involved in abusive litigation were more affected by the shock. This is consistent with the idea that the ruling was intended to curb the costs of abusive litigation. Further-

¹²In 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).

more, the literature in law discussing the case confirms that the content of the decision was unexpected. In Appendix (A.2), I argue that this also is consistent with a wide range of qualitative evidence collected around the decision from news and other public records.

1.2.2. Injunction and innovation

The intent of the Supreme Court was to introduce new norms that would improve the status of patent enforcement, by curbing abusive lawsuits and reducing the overall uncertainty in the patent system. If the ruling was successful, better enforcement should translate into more and better innovation by U.S. corporations. In particular, more clarity and less risk in the intellectual property market should increase the attractiveness of R&D investments. Furthermore, reducing the cost of opportunistic litigation could free-up resources that could be employed for innovative activities.

Many industry experts, from the software and computer industry in particular, supported the decision of the Supreme Court and claimed that strong injunction was hampering their ability to innovate. For instance, according to the Computer & Communication Industry Association, automatic injunction did “produce anti-competitive behavior, foster more litigation, and undermine innovation.”¹³ The motivations for these negative effects are clear in a comment from the American Innovators Alliance, an association representing the interests of some large high-tech companies. The group claims 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

¹³The 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.

¹⁴American 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.

However, some scholars have criticized this position, arguing that the ruling could actually end up reducing innovation (Epstein, 2008; Holte, 2015). One of the concerns was that eliminating automatic injunction would lower deterrence against real violations. According to this argument, if courts are unable to identify abusive lawsuits, more firms may decide to infringe because they anticipate that injunction may not be issued if they get caught. In turn, more violations should translate into lower returns for innovative firms, therefore reducing the overall incentives to engage in R&D.

Given these competing channels, I turn to data to determine the actual effects of the Supreme Court decision on innovation. In fact, the previous discussion suggests that theory is ambiguous about the overall effectiveness of the ruling. Without further restrictions on courts' ability to deal with patent cases, both channels may be important. Therefore, I take a data-driven approach by simply examining how R&D activity by innovative firms reacted to the Supreme Court decision. In particular, as I explain in detail below, I exploit heterogeneity in the intensity of the treatment to estimate the causal impact of the decision in innovation.

This analysis can help understanding whether the Supreme Court decision was effective in improving the status of innovation, with crucial implications for guiding future policy work in this area. However, this test can also provide broader insights on the role of intellectual property litigation in shaping firms' incentives to innovate. In fact, companies should react positively to the decision only if patent litigation before 2006 was an obstacle to their activity. Building on this intuition, my analysis can be used to quantify this distortion and explore the mechanisms through which this operates.

1.2.3. The decision and non-practicing entities

In his concurring opinion, Justice Kennedy identified patent assertion entities (PAE) as one of the main players taking advantage of almost-automatic injunction (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 & 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.

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 before 2006 (Chien & Lemley, 2012). Furthermore, the elimination of automatic injunction 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.¹⁵ 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.¹⁶

Studying the returns of these companies around the decision, I identify four important stylized facts.¹⁷ 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 t-statistics that range between 4.08 and 4.75. Second, firms suffered negative returns also in the days before the decision (Figure 1.2).¹⁸ While the largest one-day drop is experienced the day of the Supreme Court decision, stocks also lost value in the week before the ruling, and in particular, over 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 to the benchmarks (Table A.3). Finally, these negative effects do not revert back in the days following the decision.¹⁹

¹⁵The 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/>).


¹⁶The majority of the companies appear in both list - six - and only one company is only listed by Patent-Freedom. The companies are Acacia Technologies, Asure Software, Rambus, Tessera Technologies, VirnetX Holding Corp., Universal Display, Document Security Systems, Pendrell, ParkerVision, Unwired Planet, Interdigital, Spherix.

¹⁷More information about the analysis can be found in Appendix (A.3.3). 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.

¹⁸In Figure (A.5) I replicate the same results under alternative models as robustness.

¹⁹These results are qualitative identical when I use value-weighted measures.

Figure 1.2.: 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, 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 (1.2). More info on this analysis is in Appendix (A.3.3). The straight red line correspond to the trading day right before the decision.

In summary, these facts confirm that public NPEs suffered a great deal around the Supreme Court decision. In particular, the shock led to a large drop in market value, which was not reverted in the following weeks. The results are robust to the removal of each of the of the NPEs considered in the sample.²⁰ 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.

²⁰For 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.

1.3. Data

1.3.1. 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,²¹ 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 2014,²² 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.

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). Using this data, I link one or more identifiers in the patent data to one Compustat identifier using a patent level matching. Since patent numbers are very easy to match, this approach greatly reduces the possibility of errors and missing information. After applying the standard filters, I have a sample of more than one thousand non-financial public companies that are active in innovation around the decision. Lastly, I match these firms to CRSP using the standard Compustat-CRSP bridge file. In the Appendix (A.3) I provide more details on the data construction and matching.

²¹Data can be found: <http://funginstitute.berkeley.edu/tools-and-data>.

²²The 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..

Table 1.1.: Summary Statistics

(a) Full sample			
	Obs.	Mean	S.D.
$\#Patent_{jt}$	32,118	31.310	273.499
$1\{Patent_{jt} = Top^{10\%}\}$	32,118	0.338	0.473
$1\{Patent_{jt} = Top^{25\%}\}$	32,118	0.526	0.499
<i>Citation Weighted Pat_{jt}</i>	32,118	31.618	245.204
<i>Exposure_j</i>	32,118	0.769	0.775
<i>Exposure_j^{OLD}</i>	32,118	0.680	0.562
<i>Average Citation Pre</i>	32,118	1.188	1.806
$1\{Years\ first\ Patent \leq 3\}$	32,118	0.195	0.456
<i>Size Pre Portfolio</i>	32,118	16.580	128.962
(b) Public Firms			
	Obs.	Mean	S.D.
$\#Patent_{jt}$	2,032	93.122	443.469
<i>R&D/Asset</i>	2,032	0.030	0.039
<i>Exposure_j</i>	2,032	0.945	0.818
<i>Exposure_j^{OLD}</i>	2,032	0.789	0.575
<i>Average Citation Pre</i>	2,032	1.469	2.008
$1\{Years\ first\ Patent \leq 3\}$	2,032	0.040	0.196
<i>Size Pre Portfolio</i>	2,032	83.049	358.87

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.3). The variable construction is described in detail in the Appendix (A.3) - for outcomes - and in the Section (1.3) for the measures of exposure.

The main measure of innovation activity employed in the paper is based on the simple count of granted patents applied by a firm in a specific period.²³ 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

²³If patents are assigned to more than one assignee, then I equally divide the patent count across firms.

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 1.1 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.

1.3.2. 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

Figure 1.3.: Distribution of firm industries for more exposed 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.3). The first chart is constructed using only firms in the top 50% of litigation exposure, where litigation is measured using $Exposure_j$. This is constructed 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.

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 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 set $t(j) = \{\sigma_i^j\}_{i=1}^T$; (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:

$$Exposure_j = \sum_{i=1}^T \sigma_i^j p_i \tag{1.1}$$

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

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 2006.²⁴

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 & Schankerman, 2001). Using the online tool LitAlert,²⁵ I searched for all

²⁴For 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.

²⁵http://intranetsolutions.westlaw.com/practicepages/template/ip_litalert.asp?rs=IPP2.0&vr=1.0

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 1.1) 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 than 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:

$$p_i = \frac{\sum_{c \in \text{cases}} \#Patents_c^i}{\sum_{i \in \text{Tech.Classes}} \sum_{c \in \text{cases}} \#Patents_c^i} \quad (1.2)$$

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.3). 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.²⁶

I estimate $Exposure_j$ combining these two measures as in equation (1.1). My favorite measure uses litigation data from 2000 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 1.1). Furthermore, the distribution of the score is skewed and there are fewer firms with high scores. 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 1.3). 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,

²⁶For 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 & Jaravel, 2015).

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.

1.4. 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.

1.4.1. 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:

$$y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt} \quad (1.3)$$

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, which are interacted with time dummies to allow them to have a differential effects before and after the decision. Consistent with the previous discussion, $\beta > 0$ is consistent with the fact 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 2006.²⁷ 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 correlation in the data. In any case, any analysis in the paper is conducted by clustering standard errors at firm level. Furthermore, I also show that using the full panel does not affect my results.

1.4.2. 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

²⁷In 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.

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.²⁸ 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 (1.6) and I report these results in columns (1) and (4) of Table (1.2). 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. Comparing again two firms one standard deviation apart, the estimates suggest that the more exposed firm increased patent applications by almost one extra patent. Similarly, repeating the same comparison at the extensive margin, I find that the more exposed firm increased its likelihood to apply to one

²⁸In 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.

Table 1.2.: Effect of the policy change on patenting: main results

	(1)	(2)	(3)	(4)	(5)	(6)
OLS						
$Post \cdot Exposure_j$	0.041*** (0.009)	0.037*** (0.011)	0.035*** (0.011)	0.011*** (0.002)	0.028*** (0.002)	0.028*** (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	0.319	0.359	0.394
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 1.6), where I estimate the effect of the decision on quantity of innovation. In particular, the outcomes are: (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 are in Appendix (A.3). Standard errors are clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

granted patent in two years after the ruling by between 2%-5% more. As presented below, the magnitude of these results remain stable across different specifications.

1.4.3. 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:

$$y_{jt} = \alpha_j + \alpha_t + \sum_{\tau=-8}^8 \beta_{T-\tau} Exposure_j + \epsilon_{jt} \quad (1.4)$$

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 (1.4). 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.

Figure 1.4.: Effect of litigation over time



This Figure plots the β_t from equation (1.4). 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.

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.²⁹ While less flexible than the previous specification, this approach allows me to obtain more precise

²⁹I essentially estimate $y_{jt} = \alpha_j + \alpha_t + \beta^{PRE} R_j \cdot Pre + \beta^{POST} R_j \cdot Post + \epsilon_{jt}$

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 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.4) 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 (1.6) but center the analysis in a quarter where there is no change in patent law.³⁰ 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 2004Q1.³¹


The results are reported graphically in Figure (1.5), where I estimate the main specification

³⁰In order to do so, I reconstruct the outcomes and regressors as if the shock happened right after the quarter of interest.

³¹Clearly, 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 (1.2). As expected, the convergence is smooth and the effects turns positive and significant at 95% only at the end of 2005.

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.

Figure 1.5.: 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 (1.6), 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.

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 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.³² 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 1.6). Also this test confirms the quality of my empirical framework.

³²For 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).

Figure 1.6.: 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.

1.4.4. 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.4). 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 1.2). 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.³³ 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 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;³⁴ 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 (1.6), 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.2). The model is estimated using quarterly ob-

³³The z-score on the difference is small, around 0.29.

³⁴For consistency with the rest of the measures, I look at the patents applied between four and two years before the decision.

servations 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 (1.6) 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.

Lastly, I explore the robustness of my results using an alternative proxy for litigation exposure, which is based only on patent litigation carried out by non-practicing entities (NPE). As previously discussed, litigation initiated by non-practicing entities is generally considered a particularly important source of risk for companies (Cohen *et al.*, 2014). Therefore, I would expect my main results to be similar or unchanged using my measure of litigation exposure constructed using NPE lawsuits only. In an unreported Table, I confirm my previous results using a measure of exposure to litigation that exploits only lawsuits that involve patents owned by NPEs.

1.4.5. 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 time-varying fixed effects.

Table 1.3.: Timing of the effects

	$\ln(Patents_{jt})$					
	1 Year After		2 Years After		3 Years After	
$Post \cdot Exposure_j^{LARGE}$	0.036*** (0.011)		0.050*** (0.012)		0.058*** (0.0124)	
$Post \cdot Exposure_j$		0.029*** (0.008)		0.041*** (0.009)		0.0482*** (0.009)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	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 1.6. 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.3). Standard errors are clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

There are two main results in Table (1.3). 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 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.

1.4.6. 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.³⁵ 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 & Lerner, 2000). Here, I construct different measures based on citations in order to capture different aspects of quality (Appendix A.3).

³⁵Since 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 & Kerr, 2010). See also Lerner & Seru (2015).

Table 1.4.: Evidence on patent quality

(a) Panel A						
OLS	(1)	(2)	(3)	(4)	(5)	(6)
	$1\{Patent_{jt} = Top^{10\%}\}$			$1\{Patent_{jt} = Top^{25\%}\}$		
<i>Post · Exposure_j</i>	0.010** (0.005)	0.016*** (0.006)	0.016*** (0.006)	0.018*** (0.006)	0.022*** (0.007)	0.021*** (0.007)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Indu. × Time F.E.</i>		Y	Y		Y	Y
<i>Other Controls_{jt}</i>			Y			Y
<i>R²</i>	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			
OLS	(1)	(2)	(3)
	$ln(Citation\ Weighted\ Pat_{jt})$		
<i>Post · Exposure_j</i>	0.034* (0.018)	0.052** (0.023)	0.049** (0.023)
<i>Firm F.E.</i>	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y
<i>Indu. × Time F.E.</i>		Y	Y
<i>Other Controls_{jt}</i>			Y
<i>R²</i>	0.001	0.002	0.002
Observations	22,673	22,673	22,673

This Table reports the estimate of the linear difference-in-difference specification (equation 1.6), 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.3). Standard errors are clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

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 & 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.³⁶ I find that citation-weighted patents increased as well after the decision (Table 1.4).³⁷ 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 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 (1.4), 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

³⁶In this case, however, this adjustment does not play a major role and results with standard citations are very similar.

³⁷Since 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 .

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 controls. Furthermore, as previously discussed, these outcomes do not show any evidence of pre-trending before the decision (Table A.6; Figure A.4).

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.³⁸

1.4.7. 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 (1.3).

³⁸In 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.

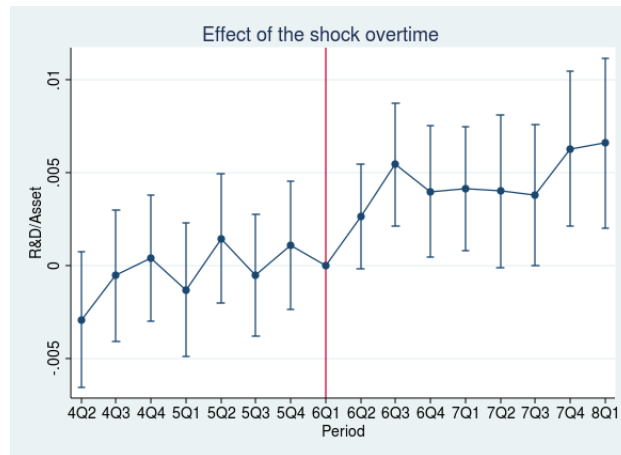
Table 1.5.: Effect of the decision on public firms

	(1)	(2)	(3)	(4)	(5)	(6)
OLS		$\ln(Patents_{jt})$			$R\&D_{jt}/Asset_{jt}$	
$Post \cdot Exposure_j$	0.064* (0.0329)	0.102** (0.0448)	0.0923** (0.0452)	0.003*** (0.001)	0.004*** (0.002)	0.004*** (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	2,032	2,032	2,032	2,032	2,032	2,032

This Table reports the estimate of the linear difference-in-difference specification (equation 1.6), 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.3). The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the five years before the decision and patent litigation at technology class since 2000. The data set is a balanced two-period panel, where each period collapses firm information in the two years before and two years after the Supreme Court decision. The sample is a set of non-financial, US located public firms that applied to at least one patent in the two years before and one after (see appendix A.3). I always control for firm fixed-effects and time effects. In Columns (2) and (5) I augment this with industry-time fixed effect. 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.3). Standard errors are clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

With this sample at hand, I estimate the same equation (1.6), looking both at patenting and R&D intensity. Results are reported in Table (1.5). 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 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 1.7 and Table A.8). As suggested by Figure (1.7), 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.

Figure 1.7.: Effect of litigation on R&D intensity over time



This Figure plots the β_t from equation (1.4) 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.

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.

1.5. 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.

1.5.1. 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 & 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:

$$y_{jtr} = \alpha_{jr} + \alpha_{jt} + \beta 1\{Risk_r\} \cdot Post \tag{1.5}$$

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.

Table 1.6.: 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 (0.025)	-0.027 (0.022)	0.270*** (0.005)	0.171*** (0.005)
<i>Split</i>	10%	25%	10%	25%
<i>Firm</i> \times <i>Time F.E.</i>	Y	Y	Y	Y
<i>Firm</i> \times <i>Risk F.E.</i>	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 (1.5), 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. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Results are reported in Table (1.6). 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.³⁹

1.5.2. 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

³⁹One 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.

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 & 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 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 (1.6) 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:

$$y_{jt} = \alpha_j + \alpha_t + \beta_1(Exposure_j \cdot FinCon_j \cdot Post) + \beta_2(FinCon_j \cdot Post) + \beta_3(Exposure_j \cdot Post) + \gamma X_{jt} + \epsilon_{jt} \quad (1.6)$$

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 & 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.

Table 1.7.: 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	<i>Small</i>	<i>Large</i>	<i>All</i>	<i>Small</i>	<i>Large</i>	<i>All</i>
<i>Post</i> · <i>Exposure_j</i>	0.004** (0.002)	-0.001 (0.01)	0.004** (0.002)	0.005** (0.002)	-0.001 (0.001)	0.005*** (0.002)
<i>Post</i> · <i>Small_j</i>			-0.004** (0.002)			-0.002 (0.002)
<i>Post</i> · <i>Exposure_j</i> · · <i>Small_j</i>			0.004** (0.002)			0.003* (0.002)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Indu.</i> × <i>Time F.E.</i>				Y	Y	Y
<i>Other Controls_{jt}</i>				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	<i>Small</i>	<i>Large</i>	<i>All</i>	<i>Small</i>	<i>Large</i>	<i>All</i>
<i>Post</i> · <i>Exposure_j</i>	0.003** (0.001)	-0.001 (0.001)	0.003** (0.002)	0.006*** (0.002)	-0.001 (0.001)	0.005*** (0.002)
<i>Post</i> · <i>Small_j</i>			-0.003** (0.002)			0.003 (0.002)
<i>Post</i> · <i>Exposure_j</i> · · <i>Small_j</i>			0.004** (0.002)			0.003 (0.002)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Indu.</i> × <i>Time F.E.</i>				Y	Y	Y
<i>Other Controls_{jt}</i>				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 1.6), 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. Firm controls are the same as in previous tables. More info on the variables are provided in the Appendix (A.3). Standard errors are clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

The results are reported in Tables (1.7) and (1.8). 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.⁴⁰

As a robustness, I show that, in my case, the results are not simply capturing heterogeneity across firms in the growth (Farre-Mensa & Ljungqvist, 2015). To rule this out, I augment equation (1.6) 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 β_1 is estimated more precisely when I add the growth controls. In an unreported Table, I find the same when looking at asset growth. Overall, my analysis is not just capturing a spurious correlation of these measure of financial constraint with different growth trajectories.

⁴⁰On 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 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.

Table 1.8.: Effect of the decision across measures of financial constraint

(a) Heterogeneity by dividend payers						
	(1)	(2)	(3)	(4)	(5)	(6)
	$R\&D_{jt}/Asset_{jt}$					
	<i>No Div.</i>	<i>Div.</i>	<i>All</i>	<i>No Div.</i>	<i>Div.</i>	<i>All</i>
<i>Post · Exposure_j</i>	0.004** (0.001)	-0.002 (0.002)	0.004** (0.001)	0.006*** (0.002)	-0.003 (0.003)	0.005*** (0.002)
<i>Post · 1{Div_j = 0}</i>			-0.005*** (0.002)			-0.004** (0.002)
<i>Post · Exposure_j · 1{Div_j = 0}</i>			0.005** (0.002)			0.005* (0.003)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Indu. × Time F.E.</i>				Y	Y	Y
<i>Other Controls_{jt}</i>				Y	Y	Y
<i>R²</i>	0.019	0.024	0.019	0.104	0.087	0.069
<i>Observations</i>	1,322	712	2,034	1,322	712	2,034

(b) Heterogeneity by rating status						
	(1)	(2)	(3)	(4)	(5)	(6)
	$R\&D_{jt}/Asset_{jt}$					
	<i>No Rating</i>	<i>Rating</i>	<i>All</i>	<i>No Rating</i>	<i>Rating</i>	<i>All</i>
<i>Post · Exposure_j</i>	0.003** (0.001)	-0.001 (0.001)	0.003** (0.001)	0.005*** (0.002)	-0.001 (0.001)	0.005*** (0.001)
<i>Post · 1{Rat_j = NO}</i>			-0.002** (0.001)			-0.001 (0.001)
<i>Post · Exposure_j · 1{Rat_j = NO}</i>			0.004** (0.002)			0.002 (0.002)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Indu. × Time F.E.</i>				Y	Y	Y
<i>Other Controls_{jt}</i>				Y	Y	Y
<i>R²</i>	0.014	0.013	0.014	0.090	0.092	0.066
<i>Observations</i>	698	1,336	2,034	698	1,336	2,034

These Tables report the estimate of the linear difference-in-difference specification (equation 1.6), 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, winsorized at 1%. The variable $Exposure_j$ captures the exposure of firm j to patent litigation, using patent application in the 5 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 the full interaction with the whole sample. I always have firm and time fixed effects, but in Columns (4)-(6) I add controls interacted with time dummies. These controls are the same as in previous tables (see Appendix A.3). Standard errors are clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

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.

1.6. 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.⁴¹

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 1.2.2), 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

⁴¹See section A.3.3 in the Appendix for more info on data construction and analysis.

ruling had positively affected the 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 (1.8), which plots the cumulative value-weighted returns of high and low exposure firms.⁴² 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.

Figure 1.8.: Stock Market Reaction: High vs. Low exposure



The Figure 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.

I can confirm all these results in a regression framework (Table 1.9), where I run cross-sectional value-weighted regression between firm returns and ex-ante exposure to litigation.

⁴²High risk firms are firms above the top 25% of litigation exposure.

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.⁴³

Table 1.9.: Stock Market returns and Litigation Exposure

	<i>Event Day</i>		<i>Event [-1; +1]</i>		<i>Event [0; +5]</i>		<i>Event [-5; -1]</i>	
	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>
<i>Exposure_j</i>	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)
Obs.	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 β 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.3). *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

1.7. 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

⁴³In 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 errors are very similar to the robust one.

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), this result suggests that the shock 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 argue that patent litigation reduces innovation because it lowers the returns from performing R&D activities. Consistent with this idea, I find that 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 & 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 not 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 address the various problems in the patent system today.

2. Private Equity, Financial Strategy, and the Crisis¹

2.1. Introduction

Private equity (PE) is a special form of investment characterized by extensive financial restructuring and a strong involvement of external investors. The general consensus is that private equity has positive effects on companies' performance. Starting with the seminal work of Kaplan (1989), research has linked investments by private equity to various improvements in operational and financial performance (e.g., Cumming *et al.*, 2007; Bernstein & Sheen, 2013; Bernstein *et al.*, 2015; Davis *et al.*, 2014), both in the United States and in Europe (Bergström *et al.*, 2007; Boucly *et al.*, 2011; Harris *et al.*, 2005).

At the same time, private equity appears to be tightly connected to credit cycles. In aggregate, periods characterized by booming financial markets also experience stronger private equity fundraising, more leverage and higher valuation for private equity deals (Figures B.1). In line with these patterns, credit conditions are identified as one of the main determinants of leverage in buyouts (Axelson *et al.*, 2013) and PE fundraising (Kaplan & Stromberg, 2009). These results confirm the importance of macroeconomic cycles for PE activity. However, we

¹Co-authored with Shai Bernstein and Josh Lerner.

know much less about the role of a strong private equity sector in affecting the real economy during credit cycles.

On the one hand, private equity investments may exacerbate the negative effects of shocks to the financial sector, aggravating the boom and bust dynamic of the economy. In line with this idea, the Bank of England suggests that buyouts should be monitored for macro-prudential reasons, because “the increased indebtedness of such companies poses risk to the stability of the financial system” (Bank of England Quarterly Bulletin, 2013Q1).² However, differences in leverage may not be the only reason for the higher fragility of private equity companies. First, pressure to complete deals during boom times may lead to the selection of lower-quality firms and a worse structuring of deals (Kaplan & Stein, 1993), leaving PE-backed companies more exposed to changes in underlying economic conditions. Second, booms in fundraising during pre-crisis periods may reduce the ability and incentives of private equity groups to monitor ex-post their portfolio companies.

On the other hand, private equity companies may be particularly resilient to large financial shock, and can therefore play a stabilizing role during bad times. In fact, PE firms tend to have superior operational expertise (Bernstein & Sheen, 2013; Davis *et al.*, 2014; Hotchkiss *et al.*, n.d.), which is crucial for weathering periods of crisis. Furthermore, PE-backed companies may be better positioned to obtain external funding when financial markets are dysfunctional. First, PE firms have strong ties with the banking industry (Ivashina & Kovner, 2011) and can therefore use these relationships to access credit markets during periods of crisis. Second, since private equity sponsors are generally deep-pocket investors, PE-backed companies may also have an easier time obtaining liquidity from shareholders when financial markets are barred.

This paper tries to resolve this tension in the literature and understand to what extent private

²This idea is consistent with part of the literature in the area showing how PE deals tend to be over-leveraged during credit booms (Axelson *et al.*, 2013).

equity contributed to the fragility of the economy in the United Kingdom (UK) during the recent financial crisis. The 2008 financial crisis represents an unprecedented shock to the banking sector that affected credit markets all over the world. In this setting, the UK is a perfect environment in which to study the performance of private equity companies. First, the UK had the largest private equity market as a share of GDP before the crisis (Blundell-Wignall, 2007) and one of the largest in absolute value. Second, the UK provides unique data on income and financial information for almost every active company, whether public or private (Brav, 2009; Michaely & Roberts, 2012).

In order to examine this issue, we study how PE-backed companies differed from non-PE companies in terms of investment, financing and performance in the wake of the financial crisis. Evidence that PE-backed companies experienced lower investment and worse performance than similar non-PE companies would suggest that intense private equity activity may indeed increase the sensitivity of the economy to credit cycles. Conversely, results in the opposite direction would limit these concerns and highlight the positive role of private equity in overcoming times of financial turmoil.

Our main analysis focuses on a large sample of about seven hundred companies that were backed by private equity before the financial crisis. Using a difference-in-difference design between 2003 and 2011, we study how the corporate policies of PE-backed companies were affected by the onset of the crisis relative to a control group. We generated this control group by identifying for each target company up to five firms with similar observable characteristics in 2007. Specifically, our baseline model looks at companies that were operating in the same industry and had similar size and profitability (Boucly *et al.*, 2011).

In our final sample, PE-backed companies look very similar to the control group across a variety of observable measures. This is true both for firm characteristics measured in 2007 and for the growth rates of these characteristics entering the financial crisis. We find not only that companies are similar in terms of size and profitability, but also that PE-backed

companies are characterized by comparable investment rates and growth, and by only slightly higher leverage after the matching. The small difference in leverage is somewhat surprising, and it may be related to the overall credit market conditions in 2007. Therefore, when interpreting our results, we need to remember that the effect of being a PE-backed company is estimated while keeping constant pre-shock balance sheet characteristics.

We start by showing that PE-backed firms decreased investments relatively less than non PE companies during the financial crisis. This effect is both statistically and economically significant. In terms of assets, the investment of PE-backed companies decreased about 5% less than the control group. Looking at the timing of the effects, the two groups present similar investment dynamics in the pre-crisis period, but the investment rate of the PE group substantially diverges from the control group in 2008. This result is robust when controlling for heterogeneity in other firm characteristics in the pre-crisis periods.

We argue that higher investments by PE-backed companies reflect the ability of the private equity investors to relax the financial constraints on their portfolio companies during the crisis. In particular, the PE-backed companies may leverage the reputational and financial capital of their sponsors to raise external funding when credit markets are not functioning. Consistent with this story, we find that PE-backed firms experienced larger increases in both overall debt and equity contributions during the financial crisis. As a result, leverage stays roughly constant. As before, these effects appear promptly in 2008 and do not revert in the following years. The effect on equity injection is particularly strong – both statistically and economically – across different specifications, and it accounts for a 2% difference in equity contribution relative to assets across the two groups.

Furthermore, PE-backed companies may cut investments less because they have better operational expertise. A more competent management may be in a better position to find resources within the company by reorganizing operations. In line with this story, we also find that PE-backed companies partially finance the expansion in long-term assets with a

reduction in working capital.

These results are robust to a battery of tests. First, we show that they are not driven by non-parallel trends in the pre-crisis period. Second, we examine the robustness of our matching by augmenting our baseline model with an extra dimension: leverage. We demonstrate that adding leverage as a matching variable does not affect the overall quality of the matching, but it reduces the number of companies matched. However, the main results are not affected by this change. Third, we show that our main results are generally not affected when we exclude companies whose private equity deals were categorized as management buyouts (MBOs), which are deals characterized by lower engagement of the private equity firms. Fourth, we argue that our results do not simply reflect the difference in attrition between PE and non-PE companies. Lastly, we show that our analysis is very similar if we control for time-varying industry shocks around the crisis.

The idea that private equity can help relax the financial constraints of portfolio companies is also consistent with two other findings. First, the positive effect on investment is particularly large among companies that were ex-ante more likely to be financially constrained. We find this result across various proxies like size, industry dependence on external finance (Rajan & Zingales, 1998) and pre-crisis leverage. Second, the increase in investment is larger when the private equity sponsor has more resources available to help the portfolio company. To explore this dimension, we exploit the heterogeneity across private equity firms in the age of the last fund raised before the crisis. Our intuition said that firms with younger funds should have more resources – both financial and operational – to invest. Our results confirm this hypothesis.

These results show that private equity companies do not appear to be relatively more sensitive to the onset of the financial crisis. Instead, during a period in which capital formation dropped dramatically, PE-backed companies invested more aggressively than peer companies. This ability to maintain a high level of investment is related to the superior access of

PE-backed firms to funding, potentially both debt and equity.

In the last part of the paper, we examine the performance of PE-backed companies during the financial crisis. First, looking at average performance, we show that PE-backed companies did not underperform the control groups. Across various measures, we find that performance changed similarly across the two groups. In particular, PE-backed companies do not become relatively less profitable and do not shrink their workforce or revenue more. The main parameter of interest is not only non-significant, but also small in magnitude.

As an alternative measure of performance, we also examine the exit patterns of PE-backed companies relative to the control group. We find that PE-backed companies are more likely to experience economic or financial distress after the financial crisis. However, we also find that these companies have a higher probability of being the target of a potentially profitable M&A transaction. Significantly, this is not driven by M&A transactions related to distress. The same patterns hold when we control for pre-crisis company characteristics. Overall, this suggests that the average performance does not significantly change for PE-backed companies, but the variance in their potential outcomes is higher.

Our results show that the investment policies of PE-backed companies are more resilient to large financial shocks than similar non-PE firms. In particular, focusing on the recent financial crisis, we find that PE-backed companies in the UK decreased investments significantly less than the control group of companies with similar balance sheet and income characteristics. We explain this result with the ability of private equity sponsors to relax the financial constraints on portfolio companies. Consistent with this idea, we find that PE-backed companies also experience a larger increase in equity contribution and debt, while keeping leverage roughly constant. Furthermore, we show that the positive investment effects are larger among companies that are ex-ante more sensitive to funding conditions. However, the greater investments do not translate into stronger long-term prospects for the firms. While PE-backed companies do not on average outperform the control group, they have a

higher likelihood of reaching extreme outcomes – like entering distress or being sold – during the financial crisis. In other words, the volatility of the final performance is higher for this group.

The paper is organized as follows. In Section (2) we present the data used in this study as well as the matching methodology developed to construct the control group. Section (3) then describes the empirical strategy employed in the paper. Sections (4) and (5) present the main results on investment and performance, discussing the possible mechanisms behind our results and presenting a large set of robustness tests. Finally, Section (6) summarizes our results and conclusions.

2.2. Data

2.2.1. Sample construction

We started our data construction by extracting from Capital IQ a sample of more than two thousand UK companies backed by private equity before the financial crisis. We identified private equity deals in Capital IQ by searching for events such as “going private,” “leveraged buyout,” “management buyout,” and “platform.” In so doing, we excluded “growth buyouts,” “venture capital” and “expansion capital” investments, where investors generally buy a stake in the company using little or no leverage. Since we were interested in studying the behavior of private equity-backed firms around the financial crisis, we selected only firms that (i) were operating in United Kingdom; (ii) had been target of a private equity deal by 2007; (iii) did not experience an exit before 2009.³ This gave us an initial sample of 2,107 firms.

We then further filtered our data by keeping only those firms that have reliable financial

³We started from the original Capital IQ data to see whether/when a company had an exit from the PE investor. We then also manually checked whether we saw the company exiting from data completely (out of business) before 2009 or it was engaged in M&A during the same time.

and income information from Amadeus, a Bureau Van Dijk (BvD) commercial data set with balance sheet information on European companies. Amadeus collects data from the “Companies House,” the United Kingdom official national registrar office. As already pointed out by other authors (Brav, 2009; Michaely & Roberts, 2012), the United Kingdom is a perfect setting for studies on private companies. According to current regulations, every registered limited company is required to provide financial and income information annually to the public register.

However, the actual disclosure requirement depends on the size of the firm.⁴ For instance, small and medium-sized firms are allowed to file different abbreviated accounts. Since the amount of information for small firms in Amadeus is extremely limited, we excluded this group from our analysis.⁵ Even if the depth of the information in Amadeus is not comparable to that available for public companies, the reliability of the source and its ability to cover the universe of UK companies is a unique feature for our study. In fact, the largest share of the companies in our initial sample consists of middle-sized private enterprises, for which similar data are not available in the United States.


In order to ensure maximum coverage and reduce concerns about selection, we supplemented Amadeus data with Orbis, which is another data product from BvD. While both Amadeus and Orbis collect information from the “Companies House,” Amadeus generally removes firms from the sample after a few years of inactivity. This is not the case for Orbis. Since the post-financial crisis period has been characterized by an increase in firm exit, using only Amadeus would have generated concerns about selection and undermined the reliability of our results.⁶

⁴For instance, since 2008, a small company must meet at least two of the following criteria: total assets less than £3.26 million, annual turnover less than £6.5 million, average number of employees fewer than 50.

⁵We usually have only total assets, revenue and profit.

⁶Orbis and Amadeus are essentially the same data product. The main two differences are the selection, as discussed above, and the interface used to distribute the data.

Figure 2.1.: Industry Distribution



This Figure reports a pie graph that describe the industry composition of our main sample. In particular, we categorize firms across seven major industries, using SIC categorization.

Starting from the initial Capital IQ sample, only firms meeting the following conditions were selected for our final sample: the firm (i) has been manually matched to Amadeus; (ii) is not a small firm, as defined in Amadeus; (iii) does not operate in the financial, public or regulated sector. This led to an initial sample of 987 unique firms. Once we excluded those firms that did not meet the minimum requirement of data available for the matching,⁷ the sample was composed of 784 firms. As we show in Figure (2.1), the industry composition of our sample is well distributed across the major areas of the economy.

2.2.2. Control group

Private equity-backed companies are clearly not a random sample of the population. For instance, in general they are larger and more leveraged than the average active firm. Therefore,

⁷We require them to have non-missing assets, ROA for 2007, and a two-digit SIC.

the first step in our analysis was to identify a proper control group for the set of PE-backed companies.

Following Boucly *et al.* (2011), we identified a suitable control group by implementing a matching procedure where, for each PE-backed company in our sample, we identified a set of control firms that operate in the same industry and had similar size and profitability in 2007. This procedure involved two steps. First, for all private equity-backed firms in our data, we selected every firm in the Amadeus/Orbis sample that (a) belongs to the same two-digit SIC; (b) has a ROA within a 50% bracket around our PE firm; (c) has total assets within a 50% bracket around our PE firm. Second, when this first step identified more than five firms, we selected the closest five, based on quadratic distance computed on the asset and ROA variables. This procedure is essentially equivalent to that of Boucly *et al.* (2011).⁸

Table 2.1.: Summary Statistics: Level

Level	Matching on ROA-Assets-2 Digit SIC						
	N.	Mean	SD	N.	Mean	SD	Diff.
	PE-Backed Company			Non PE Company			
Assets	712	122.72	635.52	3229	104.44	636.63	0.03
Revenue	678	91.37	218.9	3060	75.43	232.79	15.95
Employment	639	463.54	1093.18	2665	298.68	840.58	164.9***
ROA	712	0.07	0.25	3229	0.06	0.2	0.01
CAPEX % As.	542	0.16	0.29	2511	0.18	0.25	-0.01
Net Contribution % As.	558	-0.02	0.18	2826	0.01	0.15	-0.02***
Tot. Liabilities % As.	658	0.72	0.45	3119	0.68	0.49	0.04
Long Debt % As.	711	0.19	0.28	3226	0.17	0.28	0.02*
Bank Debt % As.	641	0.26	0.41	2801	0.25	0.47	0.01
Debt/EBITDA	622	5.05	30.19	2747	5.01	30.47	0.03
Long Debt/EBITDA	670	1.35	7.47	2839	1.17	7.07	0.18
Bank Debt/EBITDA	661	1.99	17.62	2783	2.56	18.57	-0.56

This Table reports the summary statistics of the companies at 2007 across treated (PE-backed companies) and untreated firms (non PE companies), as well as the mean difference across the two groups. This reports the mean characteristics in level at 2007. More information in the variable definition is available in the Appendix.

⁸The only difference between our methodology and that of Boucly *et al.* (2011) is the fact that we use assets instead of employment to match on the size dimension. The reason for this choice is that employment in Amadeus is a much less populated variable than assets. However, in a robustness test, we added employment as a fourth variable in our matching procedure and we show that this does not greatly affect the results.

Using this methodology, we were able to match 712 of the 784 firms available and we generated a total sample of 3,941 firms. As described later in the paper, we tested the quality and robustness of our matching procedure by augmenting the above model with other matching variables. In particular, we added leverage as a matching variable. We found that this addition does not improve the observable quality of our matching and, even if the size of our final sample is reduced, the main results are generally unaffected. For every firm in the final sample, we extracted from Amadeus/Orbis the full set of income and financial information available for the period 2004-2011. In the Appendix, we provide a full list of the variables used in the paper with a detailed definition. Given the potential errors in data reporting to “Companies House,” we winsorized every ratio at 1%. The analysis was performed using fiscal year company data, with the majority of companies reporting at the end of the year.

In Table (2.1) we compare the characteristics of firms across treatment and control groups in 2007. The two groups are very similar in size, measured by revenue or assets, profitability and investment. The average firm in the sample is a middle-sized firm with around \$130 million in assets and roughly \$90 million in revenue. However, firms in the treatment group are larger when looking at employment. Surprisingly, we find that private equity-backed firms are only weakly more leveraged than the control group. Across different measures, the treatment group tends to be more leveraged than the control group, but this difference is small and significant only when we look at long-term debt. This suggests that differences in leverage across treated and control groups mostly disappear when we compare firms of similar size and profitability within the same industry. This is an important feature of our matching, and it has implications for the interpretation of our results. This similarity is somewhat surprising since involvement of PE generally leads to an increase in leverage. Lastly, private equity-backed firms are characterized by lower net equity contribution.

Table 2.2.: Summary Statistics: Changes

Panel A	Matching on ROA-Assets-2 Digit SIC						Diff.
	N.	Mean	SD	N.	Mean	SD	
Growth rates 1 yr.	PE-Backed Company			Non PE Company			
Growth Asset	539	0.24	0.52	2428	0.28	0.57	-0.04
Growth Revenue	501	0.54	2.41	2197	0.57	2.5	-0.03
Growth Employment	478	0.07	0.35	1916	0.07	0.31	-0.01
Growth CAPEX	377	1.58	7.21	1811	1.65	6.96	-0.07
Growth Contribution	393	0.13	6.29	2023	0.58	7.24	-0.45
Growth Leverage	442	0.01	0.3	2199	0.02	0.35	-0.01
Change rates 1 yr.	PE-Backed Company			Non PE Company			
Change Asset	542	18.12	93.45	2440	15.02	276.55	3.10
Change Revenue	501	8.58	70.52	2197	8.97	62.26	-0.39
Change Employment	478	25.94	196.98	1916	4.03	262.7	21.91
Change Ln(Asset)	539	0.18	0.73	2428	0.2	0.48	-0.02
Change Ln(Revenue)	501	0.12	0.87	2197	0.14	0.85	-0.02
Change Ln(Employment)	478	0.04	0.48	1916	0.03	0.41	0.01
Change CAPEX	377	0.01	0.39	1811	0.01	0.36	0.01
Change Contribution	394	0.01	0.2	2026	0.02	0.21	-0.01
Change Leverage	443	0	0.31	2219	-0.01	0.26	0.01
Panel B							
Growth rates 2 yr.	PE-Backed Company			Non PE Company			
Growth Asset	391	0.57	1.57	1720	0.79	2.21	-0.22
Growth Revenue	351	0.89	4.93	1530	1.13	5.05	-0.24
Growth Employment	553	0.14	0.59	2280	0.16	0.59	-0.02
Growth CAPEX	253	2.34	7.68	1232	2.73	7.72	-0.39
Growth Contribution	263	2.66	15.23	1379	2.52	12.69	0.15
Growth Leverage	312	0.05	0.44	1532	0.04	0.5	0.01
Change rates 2 yr.	PE-Backed Company			Non PE Company			
Change Asset	393	30.54	179.02	1729	28	411.24	2.54
Change Revenue	351	13.34	102.84	1530	17.03	70.04	-3.69
Change Employment	337	10.44	595.74	1340	3.31	452.01	7.13
Change Ln(Asset)	391	0.28	0.55	1720	0.39	0.93	-0.11*
Change Ln(Revenue)	351	0.16	0.95	1530	0.26	1.02	-0.10
Change Ln(Employment)	337	0.02	0.59	1340	0.05	0.56	-0.04
Change CAPEX	253	0.08	0.34	1232	0.11	0.36	-0.02
Change Contribution	263	0.05	0.23	1379	0.05	0.21	-0.01
Change Leverage	312	0.01	0.33	1548	-0.02	0.35	0.03

This Table reports the summary statistics of the companies at 2007 across treated (PE-backed companies) and untreated firms (non PE companies), as well as the mean difference across the two groups. Panel A reports the one-year trend in the characteristics. Panel B reports the two-year trend of the characteristics. More information in the variable definition is available in the Appendix.

This analysis suggests that the observable characteristics of the treatment and control groups are relatively close entering in the financial crisis. Next, we repeat the same comparison examining the trends in the main observable variables (Table 2.2). In particular, we compare the one-year growth rate and changes⁹ for the main firm characteristics considered so far. We find that differences in growth rates between the two groups are always non-significantly different than zero. In other words, pre-crisis growth rates across the two groups are not statistically different. As we discuss in the next section, these results are reassuring for our methodology. In particular, they are consistent with the assumption of parallel trends between treated and control groups, which is the main identification assumption in our difference-in-difference design.

2.2.3. Other data

In the end, we supplemented the data from Amadeus/Orbis with information from other sources.

First, we identified potential shareholder exit in the post-crisis period. We started by constructing two different variables that identify whether a firm went out of business for reasons related to economic or financial distress. In particular, we generated a dummy “Out of Business,” which is equal to one if the information on the total assets of the firm goes missing in Amadeus/Orbis by 2011. Information on total assets is always required by reporting rules, and therefore when it disappears from the data, it usually suggests that the firm went out of business. However, this exit may be caused by very heterogeneous reasons – bankruptcy, a merger, or other – and therefore the interpretation of results using this variable may be difficult.

⁹Growth rates are simple growth rate between 2006 and 2007, winsorized at 1%. Changes instead are computed as a simple difference between 2007 and 2006.

Therefore, we further refined this measure by generating a dummy – “Bad Exit” – which identifies companies that went out of business unambiguously because of distress. We generated this using the firm status history, available only through Orbis. The data provider collects information from the “Companies House” and then assigns to the each firm a status, such as active, dissolved, dormant, in liquidation, or others, which may change over time. Then, we defined a company as in a “bad exit” if two criteria were met: (a) it was not active by 2011; (b) before disappearing from the data, it was defaulting in some payments, involved in liquidation or in insolvency proceedings.

Similarly, we used Capital IQ to identify potential profitable exits by looking at firms involved in M&A transactions after 2008. Since M&A transactions may also arise because of distress, we provided an alternative measure by excluding companies that were involved in M&A but were also identified in the same period as going out of business. In the analysis, we provide more information useful for the interpretation of the results.

Lastly, we also collected information on private equity firms’ history. When available, we collected the list of investors for the firm’s last private equity deals using Capital IQ. We then manually searched the private equity investors in Thomson ONE and collected information about the fundraising history. We used this information to construct variables that capture the quality and the strength of the private equity fund at the time of the crisis.

2.3. Empirical Strategy

To understand how the crisis affected the financial and investment policies of PE companies, we developed a difference-in-difference design where we compare our PE-backed companies to a control group of non-PE companies around the financial crisis. We estimate this model

using the constructed panel data over the period 2003-2011,¹⁰ which is a symmetric window around the 2008 shock. The choice of 2008 as the first year of the crisis is in line with the official Bank of England view and is consistent with a large body of empirical evidence on the crisis.¹¹

In practice, we estimate the following equation:

$$y_{it} = \alpha_i + \alpha_t + \beta(PEfirm_i \cdot Crisis) + \gamma X_{it} + \epsilon_{it} \quad (2.1)$$

where y_{it} is an outcome of firm i at time t , $Post = 1\{time > decision\}$, (α_j, α_t) are a set of firm and time fixed effects, $PEfirm_i$ is a dummy for the firms that are backed by private equity and $Crisis$ is a dummy for the crisis period (2008-2011). We augment our specification with a set of time-varying firm covariates X_{it} . To avoid endogeneity issues with controls (Angrist & Pischke, 2008; Gormley & Matsa, 2014), these variables are measured in 2007, right before the post-period window, and then interacted with the crisis dummy to allow them to have a differential impact before and after the shock.¹² Standard errors are clustered at firm level (Bertrand *et al.*, 2004).

The causal interpretation of these results crucially depends on the parallel trend assumption. In particular, we need to assume that PE-backed companies would have experienced the same change in behavior as non PE-backed companies in the absence of the financial crisis.¹³ The parallel trend assumption is intrinsically untestable, since we cannot observe the true counterfactual in the absence of the shock. However, we can still strengthen the interpretation of our analysis by providing evidence consistent with this assumption.

¹⁰To make the sample as balanced as possible in the pre-shock period, we made sure that any given control firm does not enter the data before the treatment firm.

¹¹For instance, lending to UK businesses started declining by the end of 2007, and kept this negative trend for all of 2008 and most of 2009. See the various editions of the Bank of England “Trends in Lending” publication for summaries of evidence on this issue. In particular, see the post-2009 issues.

¹²When constructing the controls, to minimize the loss of observations because of some idiosyncratic missing values in 2007, we use the 2006 value if the 2007 one is missing.

¹³For instance, it would be problematic if treated firms differed from untreated firms along some characteristics that would be affected by the financial shock independently from the status of private equity company.

First, it is important to recognize that our treatment and controls groups are very similar, at least in terms of observable characteristics. By construction, both groups have the same industry distribution, and there is therefore no difference between them in terms of exposure to industry shock during the crisis. Furthermore, as we discussed before, size, profitability, investment and leverage are very similar across these groups. Even more importantly for the parallel trend assumption, the PE and non-PE companies have similar growth rates in the years leading up to the crisis.

Building on this result, we can also formally examine the time-varying behavior of the treatment effects for the main outcomes in our analysis. With this test, we can rule out the possibility that a positive (or negative) effect for PE-backed firms merely reflects the higher (or lower) growth rate of this group of firms before the crisis. In particular, we can estimate:

$$y_{it} = \alpha_i + \alpha_t + \beta_t PEfirm_i + \epsilon_{it} \quad (2.2)$$

where we estimate a different β for every year between 2004-2011, keeping as reference year the last year before the crisis, 2007.

If our parameter β in the standard equation is correctly capturing the causal effect of the crisis on private equity firms – rather than a differential trend between the two groups – then we expect the effect of private equity to appear only after the crisis. In the next section, we will show that our main results are indeed not driven by the lack of parallel trends.¹⁴

To further strengthen our analysis, we take two extra steps. First, we augment our specifications with controls that capture the heterogeneity across firms in important balance sheet characteristics before the crisis. In particular, we control for firm size (log of revenue), growth of revenue, cash flow level (cash-flow over assets), profitability (ROA), leverage (long-term debt over assets). This enables us to rule out the possibility that our results could be driven

¹⁴A potential drawback for this approach is clearly the loss of statistical power: by increasing the number of parameters to be estimated, we decrease the level of precision. On the other hand, this specification is very flexible and can be used to gain insight into the timing of the effects.

by the presence of some unbalanced observable characteristics across treatment and control groups before 2008.

Second, as a robustness test for our main results, we add a full set of time-varying industry fixed effects, which can account for changes in industry demands and other industry characteristics around the financial crisis. In particular, we interact two digit industry fixed effects with the post dummy.

2.4. Investment Analysis

2.4.1. Main Results


In this part of the paper, we examine whether companies backed by private equity firms by 2007 were more or less affected by the financial crisis in terms of investments. While overall investments dropped significantly in United Kingdom during the crisis period,¹⁵ it is important to understand whether private equity played any role in exacerbating the effects of this negative financial shock.

We start our analysis by studying the change in investment policies between PE and non-PE-backed firms.¹⁶ In Table (2.3), we find that PE-backed companies decreased investments less than non-PE companies around the financial crisis. This effect is not only statistically significant, but also large in economic magnitude. In terms of assets, the PE firms saw their investment increasing almost 5% more than non-PE companies in the post-crisis period. The results are unchanged – both in terms of size and precision – when we add the standard set of firm-level controls.

¹⁵For instance, World Bank data shows that gross capital formation in UK between 2007 and 2009 dropped by about 20%, both in absolute value and as a share of GDP.

¹⁶In particular, we proxy investment by looking at the change in assets, accounting for depreciation, scaled by firm assets.

Figure 2.2.: Investment: PE vs. non-PE



This Figure reports the time-varying effect of being a PE-backed company on investment. Investment is measured by the standard CAPEX proxy – change in asset netting depreciation, scaled by asset. More info on this measure is available in the paper and in the Appendix. Specifically, this Figure reports the β_t of the following equation: $y_{it} = \alpha_i + \alpha_t + \beta_t PE firm_i + \epsilon_{it}$. As explained in the paper, the year 2007 is used as base period and normalized to zero. The central dot reports the point estimate while the straight line represents the 90% confidence interval. The confidence interval is constructed using standard errors clustered at firm level.

In Figure (2.2), we show that this result is not driven by non-parallel trends in the investment behavior of PE companies before the crisis. PE-backed companies did not behave differently from normal companies in the years before the financial crisis. However, they did change their investment behavior as the financial crisis hit. Most of the shift happens in the first and most important year of the crisis, 2008. After this shock, the positive effect on investment does not disappear and stay roughly at the same level in the coming years.

Overall, this suggests that companies financed by private equity seem to be more resilient to financial crisis shock. The next step is to understand the mechanism behind this result.

Table 2.3.: Investment and funding policies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	<i>Capex/Asset</i>	<i>Capex/Asset</i>	<i>Debt/Asset</i>	<i>Debt/Asset</i>	<i>Ln(Debt)</i>	<i>Ln(Debt)</i>	<i>NetContr./Asset</i>	<i>NetContr./Asset</i>	<i>Wkcap/Asset</i>	<i>Wkcap/Asset</i>
<i>PEFirm_i · Crisis</i>	0.050*** (0.012)	0.047*** (0.012)	0.005 (0.014)	0.013 (0.014)	0.056* (0.033)	0.061** (0.031)	0.022*** (0.007)	0.017** (0.007)	-0.010* (0.005)	-0.009* (0.005)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>		Y		Y		Y		Y		Y
<i>Observations</i>	20315	18874	24403	20453	24192	20342	22449	19260	24172	20327
<i>Clusters</i>	3525	3121	3914	3189	3903	3184	3842	3142	3926	3201
<i>Adj. R²</i>	0.141	0.147	0.003	0.016	0.028	0.037	0.036	0.051	0.002	0.006

This Table reports the main results of the paper, where we estimate the standard difference-in-difference fixed effect model on investment and funding variables. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In columns (1) and (2) the outcome is CAPEX scaled by asset; in columns (3) and (4) is total leverage; in columns (5) and (6) is the natural log of total debt; in columns (7) and (8) is the net equity contribution over asset; in columns (9) and (10) is working capital over asset. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

The first hypothesis is that private equity firms have superior operational ability (Bernstein & Sheen 2013; Davis *et al.* 2014), and therefore that they are better positioned to find resources within the firm when financial markets freeze. We find partial evidence consistent with this hypothesis when we look at working capital. In particular, PE-backed companies seem to cut working capital more extensively than do those in the control group during the financial crisis. Relative to assets, the decline in working capital for PE firms is about 1% higher during the crisis. While the effect is not negligible, the size is small with respect to the investment effect.


Alternatively, private equity firms may help their portfolio companies to maintain high investments by relaxing their financial constraints. This can happen in two ways. First, private equity firms tend to be deep-pocket investors and may therefore be in a better position to inject liquidity into the companies if access to financial markets is barred. Second, private equity firms have strong ties with banks (Ivashina & Kovner, 2011) and should therefore find it easier to access credit markets during periods of turmoil. We find evidence that is generally consistent with both these channels (Table 2.3).

First, net equity contributions increase more for PE-backed firms than for the control group around the crisis.¹⁷ In terms of assets, equity contributions during the financial crisis are 2% higher for PE-backed companies relative to non-PE firms. This suggests that PE funds were eager to support the operations of their portfolio companies by injecting equity into the firms. Similar to investment, this effect is not driven by non-parallel trends before the crisis and this shift in financial policy appears as early as in 2008 (Figure 2.3).

At the same time, PE-backed firms also experience a relative increase in debt financing.

¹⁷Notice that we define equity contribution by looking at the changes in equity that are not explained by profit (see Appendix). Therefore, we cannot distinguish whether positive effects are due to paying down more capital or paying out less capital.


Figure 2.3.: Equity Contribution: PE vs. non-PE



This Figure reports the time-varying effect of being a PE-backed company on equity contribution. The measure of equity contribution is just the difference in equity, netting out the profit and scaled by asset. More info on this measure is available in the paper and in the Appendix. Specifically, this Figure reports the β_t of the following equation: $y_{it} = \alpha_i + \alpha_t + \beta_t PEfirm_i + \epsilon_{it}$. As explained in the paper, the year 2007 is used as base period and normalized to zero. The central dot reports the point estimate while the straight line represents the 90% confidence interval. The confidence interval is constructed using standard errors clustered at firm level.

While on average debt declines during the financial crisis, for PE-backed firms this decline is 5% lower. The result is similar when adding controls and it is not driven by non-parallel trends in debt financing around the time of the crisis (Figure 2.4). Furthermore, while overall debt increases, PE companies do not become more leveraged. The coefficient on this outcome is indeed positive, but it is non-significant and small in size. This result should reflect the joint increase in equity injections and debt.

Figure 2.4.: Debt: PE vs. non-PE



This Figure reports the time-varying effect of being a PE-backed company on the log of the total liabilities. Total liability is just the overall sum of all the company liability. More info on this measure is available in the paper and in the Appendix. Specifically, this Figure reports the β_t of the following equation: $y_{it} = \alpha_i + \alpha_t + \beta_t PEfirm_i + \epsilon_{it}$. As explained in the paper, the year 2007 is used as base period and normalized to zero. The central dot reports the point estimate while the straight line represents the 90% confidence interval. The confidence interval is constructed using standard errors clustered at firm level.

These results suggest that private equity companies can help stabilize the economy during financial turmoil, allowing firms to invest more even when credit markets are frozen and economic uncertainty is high. In particular, private equity firms can take advantage of their financial and reputational capital to raise external finance – potentially both debt and equity – and help the funding of their portfolio companies. At the same time, there is some evidence that PE funds may also benefit from possessing important operational expertise, in particular in managing working capital.

2.4.2. Robustness

In this section, we further examine the robustness of the previous results. We find that our results, in particular on investment and equity contribution, are very robust and stable

across all the tests.

First, we reproduced our main results using a different matching procedure. In particular, we added leverage as an additional variable used to match treatment and control. Specifically, we looked for potential matching companies between those UK firms that are in the same two-digit industry and within a 50% bandwidth of assets, ROA and leverage, measured by total debt over assets. After we recreated this alternative sample, we found that the final quality of the matching is similar to the baseline one (Table B.1), but the number of firms that we were able to match drops by about 10%. In Table (B.3), we show that in general results are similar to our baseline model. The main difference is that the effect on total debt becomes statistically weaker.

Second, we also drop management buyout (MBO) from our main sample. Eliminating MBO leads to a substantial drop in total observations, since these firms constitute about 25% of deals. In principle, MBOs are characterized by lower engagement of private equity firms. If their inclusion completely drove the results, the interpretation and generalization of the analysis would be more difficult. However, as we show in Table (B.4), this is not the case. Results are generally similar both in size and statistical significance. The only difference is for working capital, which is still negative but not significant. Since MBOs do not appear to drive the main results and they can help with estimation, we include these observations for the rest of the paper.

Third, we show that our results are robust to concerns related to attrition. As usual with panel data, endogenous exit of firms from the data may bias results. This may be particularly problematic here since, as we discuss later, PE companies are more likely to enter into distress or be targeted by M&A transactions. Therefore, it is important to understand how much our results may be reflecting potential differences in survivorship across the treatment group.

To start, it is important to notice that this concern is unlikely to drive our results given the

timing of the effects. As we explained before, we constructed our sample to be conditional on PE firms' not going out of business before 2009. On top of this, we have found that most of the shift in investment and financing policies arose in 2008 – the first year of the financial crisis – when no firms were found to go out of business. This is clear from the figures where we show the time-varying effect, but we can also show this more formally by estimating our standard model using data from 2007 and 2008 only. As we show in Table (B.5), for every outcome but working capital, the results are confirmed. For working capital, the effect is still negative but slightly above the significance level.

Furthermore, we can show that results are also similar when we exclude firms that exit the data before the end of the sample. In particular, in Table (B.6) we take a conservative approach and drop every firm belonging to a matching where either the PE-backed company or one firm in its control group exits the database before 2011.¹⁸ Overall, this leads to a substantial drop in observations. In any case, we are able to replicate our previous results, in particular for the investment, equity contribution and working capital outcomes, which are still strongly significant and similar in size. Only the effect on overall debt becomes non-significant, while being still similar in magnitude.

In the end, we show that our results are robust to changes in industry dynamic. One concern is that PE-backed firms may be more or less sensitive than the control group to change in demand contemporaneous to the shock. In principle, this should not be a problem because treatment and control distribution across industries is equivalent. However, to strengthen our inference we augmented our analysis with a full set of industry fixed effects interacted with the crisis dummy. This set of fixed effects can non-parametrically control for changes in demand and other industry characteristics. As we show in Table (B.7), despite the large number of fixed effects that the model introduces, the main results are unchanged.

¹⁸For every PE, there are up to 5 matched control firms. There are 310 companies that exit before 2011, which corresponds to 245 groups of companies that are dropped.

2.4.3. Heterogeneity: investments and financial constraints

So far, we have shown that PE-backed companies were actually better than non-PE companies at keeping investment high during the financial crisis. This divergence appears in the first year of the financial crisis and does not revert over the next few years. In particular, PE sponsors help their portfolio companies to fund their activities by injecting more equity and facilitating access to credit markets. Overall, these results are consistent with the idea that private equity can play an important role during financial turmoil by relaxing the financial constraints faced by their portfolio companies. In this section, we provide more evidence consistent with this hypothesis by studying the heterogeneity of the investment results.

We start by showing that the increase in investment for private equity-backed companies was even stronger for companies more likely to be financially constrained at the time of the crisis. We identify heterogeneity in the likelihood of being financially constrained in various ways.

Table 2.4.: Investment heterogeneity and financial constraints

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
OLS	<i>Capex/Asset</i>							
<i>PEFirm_i · Crisis</i>	0.063*** (0.015)	0.059*** (0.015)	0.065*** (0.015)	0.062*** (0.015)	0.032** (0.013)	0.028** (0.013)	0.027* (0.016)	0.031* (0.016)
<i>Large_i · Crisis</i>	0.024** (0.010)	0.014 (0.013)	0.004 (0.010)	-0.016 (0.015)				
<i>PEFirm_i · Large_i · ·Crisis</i>	-0.060*** (0.023)	-0.059*** (0.023)	-0.045* (0.023)	-0.046** (0.022)				
<i>HighDep_i · Crisis</i>					-0.039*** (0.010)	-0.036*** (0.010)		
<i>PEFirm_i · HighDep_i · ·Crisis</i>					0.047* (0.025)	0.048** (0.024)		
<i>HighLev_i · Crisis</i>							-0.049*** (0.010)	-0.047*** (0.010)
<i>PEFirm_i · HighLev_i · ·Crisis</i>							0.050** (0.023)	0.034 (0.023)
<i>Measured based on</i>	<i>Employment</i>		<i>Revenue</i>					
<i>Firm & Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>		Y		Y		Y		Y
<i>Observations</i>	18563	17390	20066	18835	20315	18874	20066	18835
<i>Clusters</i>	3160	2845	3416	3109	3525	3121	3416	3109
<i>Adj. R²</i>	0.146	0.151	0.141	0.147	0.142	0.147	0.141	0.147

This Table reports the heterogeneity of result on investment. In particular, I estimate the standard difference-in-difference fixed effect model on investment, adding a further level of interaction with various proxy of likelihood of financial constraint at 2007. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007 and the financial constraint dummy. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). We use various ways to identify firms more likely to be financially constrained: size -both in terms of employment and revenue-, firms highly dependent on external finance (RZ index) and leverage at 2007. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

First, we study how the effect on investment differs between large and small firms (Table 2.4). Consistent with the idea that small companies are more likely to be financially constrained, recent studies have found that small businesses are more sensitive to shocks in the credit

markets (Petersen & Rajan 1994; Chodorow-Reich 2014). In our sample, we identify large firms by looking at the top quartile of revenue or employment at 2007, the last year in our pre-shock period. Across both measures, we find that the positive effect on investment is stronger for small companies. Furthermore, when we compare the magnitude of the interactions, we find that that among large companies, the effect of being funded by private equity is essentially zero and therefore most of the average effect is driven by smaller companies.

Second, we find similar results when we look across companies that operate in industries that are differentially dependent on external finance (Rajan & Zingales, 1998), identified using the standard Rajan-Zingales index. Firms more dependent on external finance should be more affected by turmoil in credit markets. Therefore, if private equity provides some relief to financial stress, companies in industries characterized by larger Rajan-Zingales indices should benefit more. Consistent with this idea, we find that the positive effect of being backed by private equity is even larger for more dependent firms. This interaction effect is relatively large in size and robust to the addition of the standard controls.

Third, we find similar results when comparing firms that are more or less leveraged entering the crisis. In general, firms with higher leverage are characterized by less financial flexibility and therefore should have a harder time adjusting to a shock to credit markets. As a consequence, these companies should suffer more during the financial crisis. Comparing companies based on 2007 leverage, we find that higher pre-crisis leverage always predicts lower investment. However, when looking at firms in the top quartile of leverage relative to the rest of the population, we find that highly-leveraged PE-backed companies increased investment more than their non-PE counterparts. In terms of size, being backed by PE seem to counter-balance the negative effect of being highly leveraged on investments. However, this effect is statistically significant only without controls.¹⁹

¹⁹Clearly, leverage at 2007 is endogenous to many firm characteristics, in particular debt capacity. If anything, firms that are beforehand more able to deal with a negative credit shock should ex-ante leverage more. Therefore, it is reasonable to think that our results are actually characterized by a downward bias.

These results show that the positive effect of private equity on investments is even stronger when looking at companies that are ex-ante more likely to be financially constrained. In unreported regressions, we do not find the same pattern when looking at measures of financial policy. Therefore, the investment of companies more likely to be financially constrained increased more because they were more sensitive to any shock to funding, and not because they received more financing. However, it may still be the case that funding variables are characterized by more noise than the investment measure.

To further strengthen this claim, we turn to study how our effects differ across PE funds characterized by differing age. If private equity relaxes the financial stress faced by companies, PE firms that raised the last fund closer to the financial crisis should be better positioned to help their portfolio companies. This should be the case because these firms should have more resources – both financial and operational – available for this purpose. For every PE-backed company, we used Capital IQ to identify the investors in the deal and then looked at the year in which these investors raised a fund before the financial crisis. Since variation at investor level is only available for the PE-backed firms, this analysis was only performed on this sample of firms.²⁰ Given the previous hypothesis, we expected the effects on investments to be even stronger when looking at PE firms that had raised funds more recently before the financial crisis.

²⁰Furthermore, this information was not found for every company, either because we could not use Capital IQ to identify the PE fund backing the deal or because we were not able to find fundraising information for the PE firm.

Table 2.5.: Investment heterogeneity across fund age measures

	(1)	(2)	(3)	(4)
OLS				
		<i>Capex/Asset</i>		
<i>PEFirm_i · YearLastFund</i>	0.002 (0.002)	0.003* (0.002)		
<i>PEFirm_i · 1{Fund98 – 08}</i>			0.096** (0.049)	0.104** (0.046)
<i>Firm F.E.</i>	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y
<i>Firm Controls_{it}</i>		Y		Y
<i>Observations</i>	2272	2217	2272	2217
<i>Clusters</i>	397	396	397	396
<i>Adj.R²</i>	0.102	0.177	0.103	0.178

This Table reports an heterogeneity to the main result on investment. In particular, I estimate the standard difference-in-difference fixed effect model on investment, adding a further level of interaction with measures of age of the PE fund backing the company. For this reason, these analyses are run using only the set of PE-backed companies. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007 and the fund information. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). We look at fund age using the year of the last pre 2008 fundraising of the PE firm and a dummy if this fundraising was between 1998 and 2008. This info is not available for every PE-backed firm. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

In line with this hypothesis, in Table (2.5) we find stronger effects when PE investors have raised a fund more recently, suggesting that availability of resources is an important dimension in explaining our results. We look at this effect both with a continuous measure – year in which the last fund was raised – and a dummy, where we compare companies whose investors raised any fund in the previous ten years to the rest of the population. The effects are consistent across the two measures, albeit more precisely estimated with the dummy specification. To gauge the size of the effects, we find that companies whose investors raised at least one fund in the previous ten years increased investment almost 10% more than the

other group. In this case as well, we do not find the same heterogeneity results when looking at funding, like equity contribution or debt.

2.5. Performance Analysis

2.5.1. Change in performance

In the first part of the paper, we have shown that PE-backed companies cut investment relatively less than non-PE companies in the 2008 financial crisis. Furthermore, we have linked this behavior to the superior ability of these companies to access funding in a time when access to the credit market was barred. The next step is to understand to what extent higher investment increases the long-term prospects of PE-backed companies. In this section, we examine this question by looking at various measures of company performance.

In particular, we start by examining accounting measures of firm performance around the crisis period. Using the same specification as before, we test how the performance of PE-backed companies changes relative to non-PE companies around 2008. In Table (2.6), we look at five metrics, which are ROA, EBITDA scaled by revenue or employment, employment and revenue. In general, we find no evidence that PE-backed firms underperformed the control group. Our main parameter of interest is always non-significant and relatively small. If anything, the sign of the coefficient in the baseline specification tends to be positive but non-significant.²¹

Similar to the results for investment and financing policies, the performance results are robust

²¹We find a negative effect only when we look at EBITDA/Asset. However, this effect seems to be mostly driven by the fact that PE-firms are increasing more in size. Consistent with this, when looking at the effect over time, we find that the effect builds up only in 2010 and 2011, as the size effects kicks in.

Table 2.6.: Accounting Performance

OLS	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	<i>EBITDA/Emp.</i>	<i>EBITDA/Rev.</i>	<i>Ln(Rev.)</i>	<i>Ln(Emp.)</i>	<i>ROA</i>	<i>ROA</i>	<i>ROA</i>	<i>ROA</i>	<i>ROA</i>	<i>ROA</i>
<i>PEFirm_i · Crisis</i>	0.692 (3.281)	-0.166 (2.850)	0.009 (0.025)	-0.006 (0.025)	0.013 (0.048)	-0.023 (0.047)	-0.013 (0.026)	-0.004 (0.026)	-0.005 (0.008)	-0.001 (0.008)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	20200	18476	21582	19944	23654	20463	21024	18773	24554	20601
<i>Clusters</i>	3439	3034	3639	3201	3902	3205	3570	3041	3941	3206
<i>Adj.R²</i>	0.006	0.013	0.129	0.135	0.326	0.399	0.001	0.007	0.005	0.033

This Table reports the main results of the paper, where we estimate the standard difference-in-difference fixed effect model on performance measures. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In columns (1) and (2) the outcome is EBITDA scaled by asset; in columns (3) and (4) is total EBITDA scaled by revenue; in columns (5) and (6) is the natural log of revenue; in columns (7) and (8) is the log of employment; in columns (9) and (10) is ROA. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

to various concerns raised before. We find a similar null result across all the outcomes when we use alternative matching procedures or when we exclude MBOs. By the same token, the results are close in magnitude and non-significant when we account for the possible survivorship bias or when we control for time-varying change in industry characteristics.²²

This analysis suggests that the mean performance of PE-backed companies is not differentially affected by the financial shock. In other words, higher investment does not automatically translate into higher average performance in this setting. However, these accounting measures of performance may fail to fully capture underlying change in asset quality and company value around the crisis. For this reason, in the next section we consider alternative measures of performance, based on exit patterns in the post-crisis period. In particular, we compare the relative likelihood of PE-backed companies to enter into distress or to be the target of an M&A deal. Rather than measuring the average performance, these variables can provide insight into the extreme outcomes – positive or negative – that a firm may face.

2.5.2. Exit Analysis

Even assuming that the financial crisis did not induce any differential effect on the average firm performance, PE-backed companies may still outperform or underperform the control group when looking at the extreme outcomes of the distribution. For instance, strong investments may not affect the average profitability, but they may increase the company's probability of becoming the target of an acquisition. Similarly, it may affect the likelihood of negative tail events, like distress. In this section, we jointly study how the private equity status affected the probability of both entering into distress and experiencing a profitable shareholders' exit in the post-crisis period.

As we discussed in the Data section, cases of potential financial or economic distress were

²²See previous robustness Tables.

identified using companies that exited from the Amadeus/Orbis data (marked “Out of Business”) and, among these, the exits that happened after signs of financial issues in the previous years (“Bad exit”). Here we instead look at potentially profitable exits focusing on company sales through M&A transactions (“M&A”), with specific interest in the transactions that do not appear to be related to company distress (“M&A no distress”).

This analysis studies how post-crisis exit patterns differ across PE and non PE-backed companies. One problem with this question is that variation is only cross-sectional and therefore the difference-in-difference design is not suitable for these tests. To make PE and non-PE-related companies comparable in the cross-section, we use only within-industry variation (two-digit SIC) and we control for a battery of firm-level covariates. Even with these adjustments in mind, a causal interpretation of these results is based on a much stronger assumption than the previous set of results. In particular, we estimate the following equation:

$$Exit_i = \alpha_{ind(i)} + \beta PEfirm_i + \gamma X_i^{PRE} + \epsilon_i$$

where $Exit_i$ is a firm-level dummy that identifies the exit activity, $PEfirm_i$ is the usual treatment variable identifying PE-backed companies, $\alpha_{ind(i)}$ is a set of industry-level fixed effects at two digit SIC level, and X_i^{PRE} are a set of firm-level characteristics measured before the crisis (2007). Since in this case the outcome is only cross-sectional and discrete – whether the firm took part in such an event during the crisis – we estimate this model using a conditional logit model, where we condition on two-digit SIC industry. In this case, since observations are at the firm level, we cluster our errors at the two-digit SIC industry level. To facilitate interpretation, all the results are presented as marginal effects at the mean.

In Table (2.7), we show the main results of this section. We find that PE-backed firms are more likely to experience both good and bad exits. In terms of direction and statistical significance, the results are similar with and without controls, but the magnitude is usually smaller when controls are added. Comparing the most conservative estimates, PE-backed

companies are about 9% more likely to be engaged in an M&A than are non-PE firms within the same industry. Similarly, this group is also about 19% more likely to enter into financial or economic distress.

Table 2.7.: Exit Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Marginal Eff.	1{ <i>M&A</i> }		1{ <i>M&A No Dist.</i> }		1{ <i>OutBusiness</i> }		1{ <i>BadExit</i> }	
<i>PEFirm_i</i>	0.382***	0.117***	0.373***	0.088**	0.248***	0.254***	0.188***	0.199***
	(0.013)	(0.045)	(0.016)	(0.042)	(0.045)	(0.090)	(0.068)	(0.062)
<i>Firm Controls_{it}</i>		Y		Y		Y		Y
<i>Observations</i>	3005	3005	3005	3005	2859	2859	2859	2859
<i>Clusters</i>	40	40	40	40	37	37	37	37

This Table reports the marginal value (at the mean) of a conditional logit model where we study the effect of being a PE-backed company on various outcomes. The conditioning in the model is given by two digit industry. Odd columns do not have any additional controls while even columns have firm level controls at 2007. In Columns (1) and (2) the outcomes is a dummy if the company was the target of an M&A activity in the post-crisis period; in Columns (3) and (4) is instead a dummy if the company was still target of an M&A activity and the company does not exit from the data in the same time frame; in columns (5) and (6) is the dummy equal to one if the company exit the data set in the post period; lastly in columns (7) and (8) is a dummy if the company exit the data and it reported some financial difficulties before the exit. See the Appendix and the paper for more info on the variables. Standard errors are clustered at industry level. The number of observations corresponds to the number of observations that were effectively used by the logit in the estimation, which are observations that present some within-industry (condition) variation in the outcomes. Cluster at industry as well. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

These results show that the shareholder exit patterns are very different for PE-backed companies. In line with the previous analysis, the joint behavior of the two outcomes excludes the possibility that PE-backed companies simply under- or over-performed the control group in the post-crisis period. While it is true that PE-backed companies are more likely to enter into distress, they also appear to be more likely to be targeted by a potentially profitable M&A. Therefore, this result is consistent with the idea that the larger investments did not lead to higher mean performance, but that they did increase the range of outcomes. In other words, the mean is not affected but the overall variance increases.

The higher variance in outcomes for PE-backed companies can be explained in several ways. First, it can be justified by the stronger risk-propensity of PE investors. Stronger risk-taking may lead to more aggressive investment policies, but it can also explain the increase in the outcomes' variance. Second, PE firms may simply select the best performers in their portfolios and then invest massively in them, while liquidating the other businesses. These two channels can be rationalized by the compensation structure of private equity funds, which may benefit from an increase in returns' skewness. Lastly, these patterns may simply reflect the private equity cycle. As private equity funds have a fixed length, portfolio companies purchased in the years before the financial crisis generally had to exit a few years into the crisis. At that point, PE funds may have just chosen the most efficient exit on a case-by-case basis, whether it was a sale or a liquidation. Clearly, this last explanation is not orthogonal to the previous two.

To conclude the section, we provide some suggestive evidence that these exit patterns may be related to fund selection. In line with this hypothesis, we should expect to find that companies selected by private equity funds actually experience a lower likelihood of distress and potentially better performance. One potential proxy of investors' selection is the ex-post level of net equity contribution that a company receives. Comparing companies in the top quartile of 2008-2011 equity contribution to the rest of the population, we show in Table (B.8) that PE-backed companies receiving larger contributions actually seem to experience a lower likelihood of distress. This finding is particularly large in magnitude when we focus on exits that are linked to financial distress. Not surprisingly, when we look at firm outcomes in the standard difference-in-difference model, we find that this same set of firms outperforms the average PE-backed firm in terms of investment rates and ROA, at least in the baseline model (Table B.9).²³ However, since our sorting measure is measured ex-post, these results should be interpreted cautiously because they may be capturing only reverse causality.

²³By construction, the sorting variable also predicts stronger equity injection.

2.6. Conclusion

In this paper, we have studied how PE-backed companies responded to the large turmoil in the financial markets caused by the 2008 crisis. In particular, we have examined how the investment behavior of companies backed by private equity firms reacted in terms of investment, funding and performance. One of the main objectives of this analysis is to understand to what extent the presence of a large number of private equity companies can be considered a source of fragility for a developed economy. Concerns in this direction have been raised recently by the Bank of England (Bank of England Quarterly Bulletin, 2013Q1). Furthermore, this analysis can improve our understanding of the connection between economic cycles and the private equity industry.

In our study, we examined the corporate activity of almost seven hundred companies that were backed by private equity before 2007 in the United Kingdom. This is a particularly suitable setting because the UK represents one of the most developed private equity markets in the world. Furthermore, for almost every active company in the UK – both private and public – the UK company registrar provides a rich set of financial and income information. To study how the financial crisis affected PE-backed companies, we matched these firms to a set of non-PE firms operating in the same industry and characterized by similar size and profitability in 2007. Indeed, we find that PE-backed companies and the control group share very similar observable characteristics, both in 2007 levels and in terms of growth rates during the pre-crisis period. This is not only true for size and profitability, but also for investment levels, revenue growth, and leverage.

We find that PE-backed companies actually decreased investments relatively less than the control group during the financial crisis. The overall effect is large in magnitude, and we estimate that PE-backed companies experienced a 5% lower decline in investment over asset. This result can be explained with the ability of PE-backed companies to leverage the financial

and reputational capital of their private equity sponsors to raise funding when credit markets are dysfunctional. Consistent with this idea, we find that PE-backed companies also experienced a larger increase in overall debt and equity contributions during the financial crisis. Furthermore, we show that the positive investment effects of private equity are particularly large in companies more likely to be financially constrained at the time of the crisis. Lastly, we also see stronger investment effects when the private equity sponsor was more likely to have fresh resources – both financial and operational – to invest in the portfolio company. However, these stronger investments do not automatically translate into better performance. Instead, we find that PE-backed companies experienced more extreme exit outcomes during the post-crisis period. Therefore, while mean performance appears unaffected, its variance increases.

Overall, these results suggest that private equity institutions can actually play a positive, stabilizing role in the economy. Relative to companies with similar observable characteristics at the onset of the crisis, PE-backed companies kept investment higher and did not systematically underperform. These effects are related to the ability of the private equity sponsors to compensate for the tightening in credit during financial turmoil.

3. Sovereign debt exposure and the bank lending channel: impact on credit supply and the real economy¹

3.1. Introduction

Financial intermediaries play a fundamental role in enhancing economic growth (Schumpeter 1934; King & Levine, 1993), lending to firms and households and reallocating capital to the highest value use. But loans are not the only assets held by banks. A large fraction of their portfolio is composed by securities, real properties, and equity holdings. While there are complementarities between these different investments, swings in the prices and riskiness of these assets may lead to adjustments in banks' credit supply, with potential adverse effects on the real economy.

This paper studies how shocks to the security portfolio of banks affect credit supply and businesses' real activity, with a particular emphasis on the impact on small firms. Using

¹Co-authored with Margherita Bottero and Simone Lenzu.

the 2010 Greek bailout as a natural experiment, we estimate how tensions originating in the sovereign market can lead financial intermediaries more exposed to government securities to tighten their credit supply. This contraction affected the corporate sector causing a reduction in real economic activity. Our results suggest that sovereign securities can play an important role in the propagation of financial shocks to the real economy (Gennaioli *et al.* 2013).

Our data includes all Italian financial intermediaries and a large, representative sample of non-financial firms operating in the time window around the first Greek bailout in April 2010. To assess the impact of the tensions in sovereign markets on credit supply, and then on real outcomes, we analyze over 500,000 bank-firm credit relationships and compare the change in credit supply around the Greek bailout request, across banks differentially exposed to sovereign securities.² To capture the exposure of financial institutions to the sovereign shock, we construct a bank-level proxy that exploits the cross-sectional heterogeneity in banks' holdings of Italian sovereigns measured *before* the Greek bailout. We argue that, in the pre-bailout period, sovereign holdings were a function of bank-specific characteristics and portfolio strategies, but they were unrelated to the post-bailout events. We take advantage of the widespread presence of firms that established simultaneous lending relationships with different financial institutions to control for changes in observable and unobservable firm's characteristics, such as credit demand and creditworthiness (Gan 2007a,b; Khwaja & Mian 2008). In practice, we run a within-firm difference-in-difference regression which compares changes in credit supply to the same borrower from lenders that differed in their relative exposure to Italian government debt.

The choice of this setting reflects three considerations. First, the Greek events fundamentally and unexpectedly changed the risk perception for sovereign assets. As argued by Giordano

²In our database, one observation corresponds to the difference between the average log credit granted by bank b to firm j during the period 2010:Q2-2011:Q1 and the average log credit granted by the same bank to the same firm during the period 2009:Q1-2010:Q1. This results in a total of almost 5 million observations in the quarterly firm-bank panel data set.

et al. (2013), the sequence of events which culminated into Greece’s bailout request was a “wake-up call” for investors that prompted them to discriminate among sovereigns based on the quality of their fundamentals, which were largely ignored until then. This change in attitude boosted the volatility of government bond yields in peripheral European countries, and widened their spreads vis-a-vis the German Bund. As sovereign securities represent a large fraction of banks’ assets, typically the second largest after loans, the shock had a sizeable impact on their activity. Second, our measure of banks’ exposure exploits only pre-crisis variation in sovereign holdings, allowing us to tackle endogeneity concerns related to banks’ portfolio adjustments. In fact, as the situation began to deteriorate in spring 2010, intermediaries started to (endogenously) adjust their exposure to government securities (Acharya *et al.* 2011a). Finally, our event window excludes periods characterized by unconventional measures adopted by the ECB to counteract the dysfunctions in sovereign markets, which might confound the empirical analysis.³

We document that the negative shock to the sovereign bond market had a negative, causal effect on lending. When we compare lending to the same firm by two banks which are one standard deviation apart in terms of pre-crisis sovereign exposure, we find that the more exposed financial intermediary reduced its credit supply by 7% more relative to the other. Not only banks more exposed to the shock cut lending more intensively, but they were also more likely to terminate ongoing credit relationships. Our results are robust to alternative measures of sovereign exposure, other potential confounding factors, such as bank-firm relation specific characteristics, and are not driven by non-parallel trends in lending growth before the shock. Moreover, we can exclude any systematic sorting of firms

³In particular, in our main specifications we exclude from our estimation period the second half of 2011 when the ECB re-activated its Securities Markets Programme with the objective of restoring the smooth functioning of the European sovereign market. The reactivation of the SMP was followed by two other measures, two longer-term refinancing operations (LTROs) with an extended maturity of 3 years announced in December 2011, and the OMT, announced in July 2012 (see Casiraghi *et al.* 2013). See also Carpinelli & Crosignani (2015) for a study of the effect of the the LTRO on bank credit supply.

experiencing high idiosyncratic shocks with banks more exposed to sovereigns.

The sovereign shock affected lending because it unexpectedly increased the riskiness of bank assets, inducing banks to adjust their lending behavior. Sovereigns that were considered to have little or no risk before the Greek bailout, started carrying a non-trivial amount of credit risk in the aftermath of the bailout. Therefore, banks concerned with the need to increase their capitalization or to raise funding might have preemptively tightened credit supply in order to adjust the riskiness of their assets. In line with this hypothesis, we find that the shock lead to particularly severe credit tightening by banks closer to the regulatory capital requirement. Lastly, we show that the reduction in lending was not caused by concerns related to a possible reduction in the collateral value of sovereign assets in the funding market and unlikely related to changes in soft regulatory pressure.

Looking at the real economy, the credit supply shock had a sizable impact on firms' activity, by impairing their access to credit. In particular, we show that credit market imperfections prevented firms from fully undoing the credit shock. On average, firms were unable to compensate the reduction in credit from more exposed lenders by expanding existing credit relationships or establishing new ones with less exposed financial intermediaries. In fact, we find that the lenders' exposure at the onset of the sovereign crisis is highly predictive of the change in a firm's *total* bank credit spanning the burst of the sovereign crisis. To one standard deviation increase in lenders' average holdings of Italian sovereign securities before the sovereign shock corresponds a reduction of 5% in the firm's total bank borrowing, compared to its pre-crisis amount. Conducting a simple counterfactual exercise, our estimates suggest that the lending channel can account for a drop of almost 2% in aggregate lending.

Next, we assess what kind of credit market frictions prevent firms from smoothing out the bank lending shock. First, in line with information asymmetries and agency costs theories (Stiglitz & Weiss 1981), we find a sharper tightening of credit supply to smaller companies, holding constant credit demand (Petersen & Rajan 1994). Second, firms operating in geo-

geographical areas with higher concentrations of exposed banks suffered the highest reduction of bank credit following the sovereign shock. This effect confirms that the geographical segmentation of credit markets can amplify credit contractions by affecting firms ability to smooth out credit supply shocks (Petersen & Rajan 2002; Guiso *et al.* 2004; Degryse & Ongena 2005). As expected, geographical segmentation is particularly problematic for smaller firms.

Finally, we look at the dynamics of corporate investments and employment between 2009 and 2011. The key problem in the real-effects literature is the identification of credit constrained firms (Fazzari *et al.* 1988; Kaplan & Zingales 1997). Our approach has the advantage of directly observing which firms experience a restricted access to credit using lenders' exposure to sovereigns, which - as we show - is orthogonal to firms' characteristics and investment opportunities. We find that smaller firms whose lenders were more exposed to Italian public debt cut investments more than other less exposed firms, both at the intensive and extensive margin. For every percentage point of lenders' exposure to Italian sovereigns, small firms decreased investments by 0.3 percent over the period studied, while we find no effect on the investment rate of larger firms. Looking at employment, the effect of the sovereign shock is mostly concentrated among small firms that are highly reliant on external finance (Rajan & Zingales, 1998).⁴ All in all, our analysis suggests that the real costs of a credit shock are highly heterogeneous across firms. In particular, more sensitive firms, such as smaller firms, pay a disproportionately larger price during periods of financial turmoil.

This paper contributes to the literature assessing the real effects of shocks to the balance sheet of financial intermediaries (Kashyap *et al.* 1992; Greenstone *et al.* 2014; Jiménez *et al.* 2014), potentially contributing to the debate over optimal regulation of bank's activity (Hanson *et al.* 2011). We emphasize that the effects of a credit shock are highly heterogeneous across

⁴These findings are in line with those in Duygan-Bump *et al.* (2015), who find that workers in small firms were more likely to become unemployed during the 2007–2009 recession than comparable workers in large firms, but for firms operating in industries highly dependent on external finance.

borrowers, with small firms being disproportionately more affected (Khwaja & Mian 2008; Chodorow-Reich 2014), and that the geographical segmentation of credit markets exacerbates this size effect.⁵

By documenting the economic costs of a financial system that is encumbered by public debt, this work also contributes to the strand of literature investigating the nexus between sovereigns, financial intermediation and the real economy. Among these works, a group of papers tries to explain the increase of sovereign holdings in banks' balance sheets during the recent sovereign crisis (Acharya & Steffen 2013; Battistini *et al.* 2013; Angelini *et al.* 2014), while others examine how this dynamic in government debt holdings may crowd out private credit from banks' balance sheet (Ahtik & Albertazzi 2014; Becker & Ivashina 2014). Closer to our work, Bofondi *et al.* (2013) and Acharya *et al.* (2015) investigate the economic costs of the sovereign debt crisis via the bank lending channel. Bofondi *et al.* (2013) exploit the heterogeneous exposure to country risk of Italian versus foreign banks to identify the causal effect of the European sovereign debt crisis on credit supply using Credit Registry from Bank of Italy.⁶ Acharya *et al.* (2015) uses a sample of European firms active in the syndicated loan market to show that companies borrowing from banks incorporated in distressed countries were more prone to reduce investments and job creation during the sovereign crisis.⁷ Our work is consistent with these papers, but it differs from them in two important respects. First, our estimation strategy relies on a direct, bank-specific measure of sovereign exposure, which allows us to directly disentangle country-specific risk from a bank-specific lending channel. Second, our paper extends the findings in Acharya *et al.* (2015) by looking at a larger, representative sample of firms and at the population of banks

⁵Using Italian Credit Register data, Albertazzi & Bottero (2013) examines the lending channel in the context of the 2008 credit crunch, and Cingano *et al.* (2014) identifies the effects of the financial crisis on both bank lending and investment policies of firms.

⁶Other related papers are Correa *et al.* 2014, De Marco 2014, and Popov & Van Horen 2013, that study the effect of the sovereign crisis using bank-level data.

⁷Similar to Acharya *et al.* (2015), Balduzzi *et al.* (2014) uses survey data to document the real effects of the sovereign and financial crisis in Italy.

active in Italy. This is crucial because it allows us to highlight the large heterogeneity in the response to the sovereign shock, both at firm and bank level, and to provide new insights into on its transmission channels. Finally, our paper confirms and provides causal evidence of the findings in Gennaioli *et al.* (2013) which, looking at a panel dataset in a large number of countries, document a large, negative and statistically significant correlation between a bank's bond holdings and its subsequent lending during periods of sovereign defaults.

From a policy perspective, this paper shows that, in a banking system characterized by large and concentrated holdings of the sovereign assets, tensions in the public debt market can have sizable, negative effects on the real economy. While the interpretation of this result is simple, actual policy recommendations are more complicated and they need to take into account other important factors, such as the potential costs of banks' portfolio adjustments. Furthermore, this paper highlights that the costs of a credit contraction are very asymmetric, with smaller firms bearing most of the downstream, negative effects. This has important implications for the design of public policies aiming to shield real activity from financial market turmoils.

The remainder of the paper is organized as follows. In Section 3.2, we review some background information about the European sovereign debt crisis. In Section 3.3, we describe the data used in the paper. In Section 3.4, we introduce and discuss our identification strategy, provide evidence of the presence of the bank lending channel, and document its heterogeneous effects across firms and banks. In Section 3.5, we show that the sovereign shock impaired firms' access bank credit, consequently affecting their investments and employment. Section 3.6 concludes.

3.2. The onset of the sovereign crisis

The key event for our study is the bailout request advanced by Greece in April 2010. The bailout prompted a reassessment of the default risk of a number of countries of the European Economic and Monetary Union, arguably the first since the adoption of the single currency, which directly affected financial intermediaries that invested in sovereign securities of distressed countries.

Until late 2009 neither financial markets nor the media appeared to be particularly concerned with the sustainability of sovereign debt in peripheral European countries. For over ten years since the introduction of the Euro, the yields of 10-year bonds issued by European countries had been low and stable.⁸ However, after the parliamentary elections held in Greece in October 2009, the newly elected government acknowledged accounting budget misreporting in previous years and a larger-than-expected fiscal deficit in the coming year, generating concerns about the health of the Greek economy and the solvency of its sovereign debt.⁹ The situation became even more dramatic in the spring 2010. On April 23, the Greek government requested an EU/IMF bailout package to cover its financial needs for the remaining part of the year. A few days later, on April 27, Standard & Poor's downgraded Greece's sovereign debt rating to BB+ ("junk bond") and, in response to these events, yields on Greek government bonds rose sharply, barring the country's access to capital markets.

The Greek crisis can be regarded as a *wake-up call* on sovereign risk (Goldstein 1998, Giordano *et al.* 2013), which increased investors' sensitivity to country-specific macroeconomic fundamentals and prompted them to re-assess the default risk of other euro-area countries.¹⁰

⁸The convergence began right after the institution of the European monetary union and was interrupted only by a short-lived increase of the interest rates of peripheral countries in the second half of 2008 driven by banks bailouts in Ireland and Greece (Acharya *et al.* 2011a).

⁹After a series of upward deficit revisions, the Greek government estimated to deficit at 12.7 percent of GDP for 2009, up from the 7.7 percent in 2008. We refer to Lane (2012) for a detailed description of the European sovereign crisis.


¹⁰The literature distinguishes between three types of international transmission of area-specific shocks: "wake-

Shortly after the events in Greece, investors began to be concerned with the solvency and liquidity of the public debt issued by other European countries, starting with Ireland and Portugal and spreading soon afterwards to Spain and Italy (Angelini *et al.* 2014). The shift in the perceived riskiness of sovereign debt after the “wake-up call” is confirmed by a number of other indicators. In Italy, the spread between the BTP and the German Bund (henceforth BTP-Bund spread) increased from 85 bps at the end of the first quarter of 2010 to almost 160 bps in the third quarter of the same year (Figure 3.2a). To put the economic magnitude of this event into perspective, this jump corresponds to an increase of about 1.7 standard deviations of the spread series since 2005. Similarly, the CDS on Italian bonds with 5 years of maturity doubled soon after the bailout. As documented by Abbassi *et al.* (2014), Italian banks also paid a higher cost in the interbank market lending after the Greek bailout, although volumes were not significantly affected.¹¹


up call contagion”, “pure contagion”, and “shift contagion”. See Pericoli & Sbracia (2003) for a review. Giordano *et al.* 2013 show that the Greek crisis was a wake-up call for investors who largely ignored macroeconomic fundamentals for peripheral EMU countries before the end of 2009. By contrast they find no evidence of other forms of contagion. This narrative is in line with a large body of evidence (Acharya & Steffen 2013; Augustin *et al.* 2014; Abbassi *et al.* 2014), which finds that the Greek bailout represented a shock to sovereign bonds’ risk. Benzoni *et al.* (2014) present a theoretical model that shows that unexpected events as the Greek bailout may trigger a widespread increase in uncertainty with negative repercussions for other countries.

¹¹Furthermore, the correlation between the Italian and Greek spread, that was essentially one before the October 2009, reflecting the fact that the spreads with Germany were negligible for all euro area countries, dropped substantially in the months immediately before the bailout, driven by the rising yields on the Greek bond, and then increased again, as tensions spilled-over to other European countries driving upwards the yields on their sovereign securities (Table C.2 in Appendix C.3).

Figure 3.1.: The burst of the sovereign crisis



(a) 10y Spread Italy and Media Coverage



(b) Bank's CDS Spread

Figure 3.2a shows, on the left-hand axis (solid line), the dynamics of the spread between the yield of 10-year Italian zero-coupon bonds and that of 10-year zero-coupon bonds issued by Germany. Data from ECB. On the right-hand axis (dashed line) Figure 3.2a displays the frequency of Google searches of key words "Euro Crisis" using Google Trends. The y-axis reports the Google searches in every week between the beginning of 2008:Q1 and the end of 2011:Q2 as a fraction of the total Google searches of the key words over the same period. Source: <http://www.google.com/trends/>. Sources: ECB, Google.

Figure 3.2b reports the time series of the average of the CDS spreads on unsecured senior debt of the top 5 Italian banks (solid line). Data are taken from Markit database and include only the CDS issued in Euro. Sources: Markit.

The sudden change in the expected risk profile of government securities had a direct negative effect on the balance sheet of banks holding these assets. In September 2010, participants to the European Bank and Insurance CEO conference envisaged fears originating from sovereign markets to be the biggest threat to bank share prices (Figure C.5a).¹² The same survey reveals that investors were ranking banks in countries mostly affected by the sovereign shock - Italy, Greece and Portugal - among the financial institutions with the lowest expected performance (Figure C.5b). This negative outlook is confirmed when looking at CDS on bonds issued by large Italian banks (Figure 3.2b). In particular, looking at the quarter before and after the Greek bailout, the average spread of the top-five Italian banks rose by 78 percent, over 2.3 times its standard deviation. By any metric, the Greek bailout represented the breaking point of the recent events related to the EU sovereign crisis.

3.3. Data

The building block of our database is Bank of Italy's Credit Register, which contains detailed information on the credit relationships entertained by intermediaries operating in Italy.¹³ Quarterly accounting information - balance sheet, income statement, and a detailed picture on sovereign bond holdings - for each bank are collected from Bank of Italy Supervisory records. Balance sheet, income statement, industry, and location of firms appearing in the Credit Register are drawn from the Cerved database and the CADS database (sub-sample of firms), two proprietary firm-level panel database owned by Cerved Group S.p.A.. Appendix C.6 discusses the data construction process and variables used in this work.

We restrict our analysis to a two-year window around the Greek bailout (May 2010), which

¹²“Competing in the age of austerity”, Bank of America Merrill Lynch Banking & Insurance CEO Conference, London, 29 September 2010. Source: <http://ftalphaville.ft.com/2010/10/04/359726/european-bank-watch-past-present-and-future/>

¹³For a performing credit relationship the reporting is mandatory only if it exceeds the threshold of 30,000 euro. Positions in default need to be reported irrespective of their amount.

we split into a pre-crisis period - from 2009:Q2 to 2010:Q1 - and a post-crisis period - from 2010:Q2 to 2011:Q1. After applying standard filters and consolidating bank balance sheet items, our final database includes 539 different bank holding companies, 302, 538 non-financial firms, and 538, 348 unique firm-bank credit relationships, for a total of more than 4.5 million observations between 2009:Q2 and 2011:Q1. 188, 381 firms established simultaneous multiple lending relationships.

This sample is representative of the universe of Italian banks and firms and, therefore, it is well suited for our purpose of investigating the real effects of credit supply shocks across heterogeneous types of firms. In the CADS sample, firms have a median fixed assets of 996 thousand Euros, and median revenues of 6,383 thousand Euros. Half of this sample has less than 36 employees, while firms belonging to the top 1 percent of the distribution have more than 773 employees. Therefore, differently from other works exploring the real effects of credit shocks (e.g. Chodorow-Reich, 2014), our definition of small firms encompasses companies at the bottom of the overall size distribution, whose main way to access debt financing is through the local banking sector (Petersen & Rajan, 1994).

Our main dependent variable is the percentage change in average outstanding loans between the pre- and post-shock period for every firm-bank credit relation in our data set (Khwaja & Mian 2008). More precisely, we collapse the quarterly amount of credit granted to firm j by bank b to two a pre-shock average (2009:Q2-2010:Q1) and the post-shock average (2010:Q2-2011:Q1), and we calculate the percentage change as the log-difference between the two averages ($\Delta \ln(\text{Loans}_{bj})$). The stock of Italian government bonds at the end of 2010:Q1 scaled by risk-weighted assets ($\text{Sovereigns}_{2010Q1}$) provides us with a bank-specific measure of financial institutions' exposure to the sovereign shock.¹⁴ Our empirical models include the following set of bank-level controls: bank profitability, size, capitalization, retail funding,

¹⁴In a robustness exercise, we use alternative measures of exposure to the crisis (like total sovereign holdings, and sovereign holdings of “peripheral” European countries) and a different scaling variable (Tier1 and total assets rather than risk-weighted assets). We show that our empirical results are unchanged.

interbank funding, liquidity, quality of lending portfolio, and status of the bank as a mutual bank (BCC).¹⁵ Moreover, we control for the length of the lending relationship between a borrower and each of its lenders, and for the contribution of each lender to the total bank debt of the borrower. All bank-specific and relationship-specific controls are measured at the end of the first quarter of 2010, i.e. the last quarter of the pre-crisis period. Information on firm's industry and geographical location are drawn from the Cerved database. From the CADS database we obtain information on firms' assets, revenues, fixed assets, and employment for a subset of the firms the Credit Registry (about 35,000 firms), for which we are able to measure investments and changes in employment. To limit the influence of outliers in the growth rate of credit, employment, and investment, we winsorize the top and bottom 1% of the distribution of these variables.

¹⁵Mutual banks are characterized by a different statutory objective than other banks and potentially a different lending policy. The presence of mutual banks is not an Italian *unicum*, but they are widely found across Europe and the US.

Table 3.1.: Summary Statistics - Loan-level and Firm-level variables

Panel (a): Relationship-specific Variables							
	Obs.	Mean	SD	SD ^{Within}	SD ^{Between}	1 st Perc	99 th Perc
$\Delta\ln(\text{Loans})$	424191	0.011	0.567	0.366	0.478	-1.771	1.702
$\Delta\ln(\text{Credit Lines})$	424191	0.005	0.519	0.336	0.463	-1.704	1.546
$\Delta\ln(\text{Tot Credit})$	424191	0.011	0.514	0.317	0.450	-1.572	1.564
Cut Credit	464435	0.050	0.219	0.142	0.194	0.000	1.000
Lenght Relationship _{2010Q1}	424191	10.372	5.555	3.258	4.780	1.000	17.000
Share Relationship _{2010Q1}	424191	33.124	21.076	16.459	13.247	2.932	90.312

Panel (b): Firm-specific Variables						
	Obs.	Mean	SD	1 st Perc	99 th Perc	
Fixed Assets ₂₀₀₉	35165	3,443.703	7,771.823	7.00	53,724.00	
Employees ₂₀₀₉	35165	69.01	118.00	2	773	
Revenues ₂₀₀₉	35165	12,252.73	21,406.13	225.00	148,025.00	
1{Invest < 0}	35165	0.44	0.50	0.000	1.000	
% Δ Invest	35165	0.03	0.56	-1.87	2.18	
1{ Δ Empl < 0}	16695	0.43	.50	0.000	1.000	
% Δ Empl	16695	0.00	0.25	-0.99	0.93	
RZ Index	35165	-1.22	2.88	-14.96	0.91	
Sovereigns Province _{2010Q1}	35165	0.09	.02	0.06	0.16	

Panel (c): Bank-specific Variables						
	Obs.	Mean	SD	1 st Perc	99 th Perc	
Sovereigns _{2010Q1}	539	0.251	0.364	0.000	0.892	
Sovereigns ^{TIER1} _{2010Q1}	539	1.423	1.192	0.000	5.643	
Sovereigns _{2010Q1} /Total Sovereigns _{2010Q1}	539	0.989	0.0571	0.068	1.000	
ROA _{2010Q1}	539	0.004	0.002	-.0313	0.006	
Size _{2010Q1}	539	5.160	1.623	2.768	11.194	
Tier1 _{2010Q1}	539	0.172	0.119	0.062	0.593	
Low Capital Ratio _{2010Q1}	539	0.186	0.389	0.000	1.000	
Deposits _{2010Q1}	539	0.812	0.465	0.096	1.995	
Liquidity _{2010Q1}	539	0.014	0.019	0.001	0.035	
Net Interbank Debt _{2010Q1}	539	-0.093	0.242	-0.672	0.258	
Bad Loans _{2010Q1}	539	0.038	0.033	0.002	0.183	
BCC	539	0.758	0.428	0.000	1.000	
Total Sovereigns _{2010Q1}	539	0.254	0.369	0.000	1.320	
Total Sovereigns GIPSI _{2010Q1}	539	0.252	0.364	0.000	1.320	
Total Sovereigns GIPS _{2010Q1}	539	0.001	0.006	0.000	0.032	

This table reports the summary statistics of the relationship-specific, firm-specific and bank-specific variables used in our analysis, using firms with multiple lending relationship in the specific data set. See Appendix ((C.7)) for a detailed description of the variables.

Table 3.1 displays the summary statistics of the variables of interest.¹⁶ On average, credit relationships displayed a weakly positive growth rate between the post- and pre-crisis period (Table 3.1, Panel a). On the extensive margin, about 5% of credit relationships that existed before the onset of the sovereign crisis were terminated between 2010:Q1 and 2011:Q2. However, these changes mask a very large of heterogeneity in cross section. Indeed, the standard deviation of $\Delta \ln(\text{Loans}_{bj})$ is about 74 times higher than its mean, and the interquartile range spans from -19 percent to +14 percent. Moreover, we observe high variation in the percentage change of loans both across firms and within-firm. A high degree of heterogeneity also characterizes the distribution of corporate investments and employment (Table 3.1, Panel b). Over the period 2009-2011 firms increase their investments by 3 percent on average, with a standard deviation about 20 time higher. The average workforce does not change between 2009 and 2011 for the firms in our sample, but the standard deviation of the percentage change in employment is very high (25 percent).

Moving to our main treatment variable, in 2010:Q1 the average exposure of Italian banks to sovereign was 25 percent, with a standard deviation of about 36 percent (Table 3.1, Panel c). These statistics are indicative of both a high average exposure and of a high dispersion in the exposure to the sovereign shock across the financial intermediaries in our sample. Italian sovereign debt amounts, on average, to almost 99 percent of banks' sovereign portfolio. This is true also when we look at the banking groups with the most diversified portfolio of sovereign bonds. A bank holding company located at the first percentile of the distribution allocated 58 percent of his sovereign portfolio to Italian government bonds in 2010:Q1. While these numbers might seem surprising, in Appendix C.2 we show that the strong home-bias of financial institutions in our sample is not a peculiarity of the Italian banking system, but

¹⁶We present further details about the distribution of the most important variables in Appendix B. Appendix C.7 reports a description of the variables.

rather a common feature of many European countries like Germany, France and Spain.

3.4. The bank lending channel

3.4.1. Identification strategy

In this section, we present the identification strategy employed to investigate how the sudden and sharp shock to sovereign securities triggered by the Greek bailout was transmitted to the credit supply through the bank balance sheet channel. To gain intuition on the main challenges of this analysis, we present a simple model of lending under costly external financing in the spirit of Khwaja & Mian (2008).¹⁷

When external financing is costly due, for example, to asymmetric information in the wholesale market (Stein 1998), a shock to the market value of sovereigns held in bank's portfolio increases bank's cost of funding and translates into a tightening of credit supply (bank lending channel). The sovereign shock, however, may also lead to a simultaneous reduction of firm's investment opportunities, lowering firms' credit demand (e.g. Bocola 2013). Formally, consider a simplified environment where a bank b lends to firm j using short-term funding backed by sovereign securities or issuing bonds. Following a reduction of liquidation value of sovereign assets from γ to τ ($0 \leq \tau \leq \gamma$), the change in bank credit supplied by bank b to firm j can be decomposed as follows:

$$\Delta L_{bj} = \frac{1}{\alpha_B + \alpha_L} (\Delta\theta_j + \Delta\theta + \alpha_B \tau G_b) = \beta_0 + \beta_1 G_b + \rho_j \quad (3.1)$$

where the parameter $\alpha_L > 0$ measures the concavity of the demand for loans with respect to loan size, $\alpha_B > 0$ captures the cost of external finance, $\Delta\theta$ and $\Delta\theta_j$ capture, respectively, aggregate and idiosyncratic changes in the bank's marginal return on the loan. Thus, we

¹⁷Appendix C.1 provides a complete description of the model and its solution.

can decompose the change in loans from $t = 1$ to $t = 2$ can be decomposed in three factors: (i) an idiosyncratic firm-specific shock to the profitability or productivity of the borrower ($\rho_j = \frac{1}{\alpha_B + \alpha_L} \Delta\theta_j$); (ii) an economy-wide shock ($\beta_0 = \frac{1}{\alpha_B + \alpha_L} \Delta\theta$); (iii) and the sensitivity of credit supply to the balance sheet shock triggered by bank b holdings of sovereigns ($\beta_1 G_b = \frac{\alpha_B \tau}{\alpha_B + \alpha_L} G_b$). The last term captures the bank lending channel, that is the reduction of credit supply imputable to the transmission of the bank-specific shock driven by direct holdings of distressed sovereigns.

Equation (3.2) highlights a first, crucial difficulty in the estimation of the bank lending channel. Since the firm-specific demand shift – ρ_j – is unobservable, an OLS estimator capturing the effect of the sovereign exposure on lending will be biased if firm demand is correlated to bank-exposure to sovereigns. In particular, if banks with high sovereign exposure systematically lend to firms with negative demand shock ($\text{corr}(G_b, \Delta\theta_j) < 0$), we would have a downward bias in our estimates. This would falsely lead us to attribute the drop in credit to a credit supply shift, even when the overall effect is entirely driven by a reduction of firms credit demand.¹⁸

Following Khwaja & Mian (2008) we circumvent the omitted variable problem by focusing on firms with multiple lending relationships, and adding firm fixed effects (ρ_j) in a first-differenced model to control for observable and unobservable idiosyncratic changes in demand-side factors which might be systematically correlated with lenders’ exposure to the crisis.¹⁹ We bring Equation (3.1) to the data using bank b pre-bailout (2010:Q1) holdings of Italian sovereign securities to measure bank’s exposure to the shock (Sovereigns_{2010}), and

¹⁸This systematic sorting between highly exposed banks and “weak” firms is not only a theoretical possibility but can be empirically relevant. Consider the case of poor areas within a country: in these areas banks may end up holding more sovereign asset on average because of lower investment opportunities. At the same time, they will lend to local firms, which may be weaker and therefore more sensitive to sovereign shocks.

¹⁹Studying a number of countries, Ongena & Smith (2000) and Detragiache *et al.* (2000) report that firms borrowing from one bank is the exception more than the rule. In the United States 55.5 percent of small and medium firms have more than one bank, and the median number of credit relationships established by them is two.

compare credit supply in the four quarters *before* April 2010, when tensions on the Italian sovereign debt grew unexpectedly, to credit supply in the four quarter *after* it:

$$\Delta \ln(\text{loan}_{bj}) = \beta_0^I + \beta_1^I \text{Sovereigns}_{b,2010Q1} + \Gamma^I \cdot X_{b,2010Q1} + \rho_j + \epsilon_{bj} \quad (3.2)$$

where ρ_j is a firm fixed effect in a first-differenced data and $X_{b,2010Q1}$ is a set controls measured at the pre-bailout period. To reduce concerns related to serial correlation of the errors (Bertrand *et al.* (2004)) and to average out any seasonality effect (Duchin *et al.* 2010), we collapse the quarterly data into a pre-shock and a post-shock average, and calculate the before-to-after percentage change in loans granted from bank b to firm j as the log difference of the two averages ($\Delta \ln(\text{Loans}_{bj})$). The empirical model in Equation (3.2) is equivalent to a Difference-in-Difference specification, where intermediaries with lower exposure to Italian debt are used as the control group for banks with higher exposure. A negative and statistically significant value of the coefficient $\hat{\beta}_1$ indicates the presence of the lending channel triggered by banks' sovereign holdings.²⁰

In an ideal experiment, the amount of pre-bailout sovereigns would be randomly assigned across banks. In reality, the holdings of sovereign securities in the portfolio of financial institutions and their credit supply are a function of a set of bank-specific characteristics and asset management strategies (Gennaioli *et al.* 2013). In fact, at the onset of the sovereign crisis, bigger, less capitalized and less liquid banks held a smaller share of assets invested in sovereign securities, while financial institutions with more more deposits and less reliant on interbank funding held more (Table C.4 in Appendix C.9). The same characteristics are correlated with changes in banks' propensity to lend.²¹ To isolate the effect of sovereign

²⁰Differently from previous studies (Khwaja & Mian 2008), our estimates of β_1 should not be interpreted as the elasticity of credit supply with respect to the shock to market value or liquidity of sovereign holdings. Instead, β_1 measures the percentage change of a lender's credit supply with respect to the other lenders of the same firm, which results from his exposure to sovereign securities.

²¹Comparing the bank-level change in total corporate loans between our pre- and post-shock period, Table C.5 (Appendix C.9) shows that bigger and more profitable banks cut lending more than smaller and less profitable financial institutions. The same is true for banks with more stable sources of funding, those that

exposure on credit supply from other bank-characteristics which may be correlated with the overall sovereign holding and have a direct impact on lending behavior, we augment Model (3.2) with the set of bank-specific controls ($X_{b,2010Q1}$), measured at the end of the last quarter before the shock (2010:Q1). In particular, we include bank profitability, size, capitalization, funding (retail deposit and wholesale), liquidity, quality of lending portfolio, and status of the bank as a mutual bank.

The coefficient β_{11}^I estimated in equation (3.2) captures the pure *intensive margin* of the bank lending channel, as the left-hand side variable measures the change in the stock of loans granted by lender b to borrower j focusing only on relationships existing in both the pre- and the post-shock period. Using a similar specification we can estimate the lending channel along the *extensive margin* by testing whether banks more exposed to the shock were also more likely to terminate ongoing lending relationships after the onset of the sovereign crisis:

$$\text{Cut Credit}_{bj} = \beta_0^E + \beta_1^E \text{Sovereigns}_{b,2010Q1} + \Gamma \cdot X_{b,2010Q1} + \rho_j + \epsilon_{bj} \quad (3.3)$$

where Cut Credit_{bj} is binary variable taking value one whenever the credit relation between bank b and firm j in place at the end of 2010:Q1 was terminated within the following year, after the onset of the sovereign crisis. The coefficient β_1^E tests whether, after the sovereign shock, a firm is more likely to interrupt a lending relationship with those lenders more exposed to government securities after the sovereign shock.

The validity of this identification strategy relies on the following conditions. First, financial institutions should not have anticipated the imminent transmission of the sovereign crisis to the Italian debt and adjusted their sovereign portfolio beforehand. If the shock to Italian sovereigns were expected before the downgrade of Greece, holdings at 2010:Q1 might reflect

have higher liquidity holdings, and those that participate in the wholesale market as net borrowers. On the contrary, more capitalized financial institutions increased their corporate lending relative to less capitalized banks before-to-after the Greek bailout.

strategic adjustments undertaken in expectation of the imminent crisis, and confound the interpretation of our results. The stylized facts presented in Section 3.2 suggest that this was not the case. Before the downgrade of Greek debt both financial markets and media were not expecting such a sharp contagion of the sovereign crisis to Italy and other European countries. Moreover, the origin of the tensions on Italian sovereigns can be traced back to large government deficits and high public debt rather than imputed to a structural weakness of its banking system (Acharya *et al.* 2011a; Lane 2012; Angelini *et al.* 2014).

Second, our quasi-experimental design also requires the parallel trend assumption to hold. In other words, we need to assume that, in the absence of the sovereign crisis, financial institutions with higher sovereign holdings (the treated group) would have displayed a credit supply trend comparable to banks with lower holdings (the control group). While the parallel trend assumption is fundamentally untestable due to the absence of an observable counterfactual, the next section presents and discusses extensive evidence that supports it.

3.4.2. The bank lending channel: main results and robustness

We bring to the data the econometric model described by Equations (3.2) and (3.3). Results are reported in Table 3.2. According to our estimates, higher exposure of a financial intermediary to Italian sovereign debt had a negative and statistically significant effect on its credit supply (Columns (1) and (2)).²² This change in bank lending is a pure supply shock, which is orthogonal to idiosyncratic demand-side shocks such as changes in credit demand or borrower’s riskiness. The lending channel is not only statistically, but also economically relevant. On average, comparing the change in credit supply of two lenders whose exposure to distressed sovereigns is one standard deviation apart (0.36), the lender with the higher

²²While an observed reduction of bank credit from bank b to firm j could be the result of an increase in the interest rate charged by b , previous works suggest that banks are reluctant to increase borrowing costs to avoid financing more risky borrowers (adverse selection) or to discourage firms to take more risk (moral hazard) (Stiglitz & Weiss 1981).

holdings of sovereigns tightens its credit supply 7% more relative to the other. All the other coefficients in our regression are in line with the predictions of economic theory.²³

Similar effects are detected along the extensive margin (Columns (3) and (4)). Comparing two lenders of the same firm that are one standard deviation apart in terms of exposure to sovereigns, the lender with the highest holding interrupts existing lending relationship with a 2% higher probability relative to the other bank. This is a sizable increase since in our sample the unconditional probability of not renewing in the post-shock period a loan that existed in the pre-shock period is about 5%.

²³Clustering standard errors at bank level, only size, Tier1, BCC, share of lending relationship and length of lending relationship remain statistically significant. The negative relationship between the time-length of a credit relationship and its percentage change before-to-after the sovereign shock is likely capturing the credit-cycle effect. The older a credit relationship is, the more likely it is that the credit will be rolled-over or renegotiated in the post-shock period. Since firms experience a tightening of credit standards in the post-shock period, the amount of new credit will be likely reduced by all lenders.

Table 3.2.: The Bank Lending Channel

<i>Dep. Var :</i>	$\Delta\ln(\text{Loans})$		Cut Credit	
	(1)	(2)	(3)	(4)
Sovereigns _{2010Q1}	-0.176*** (0.057)	-0.176*** (0.022)	0.053*** (0.025)	0.053*** (0.007)
ROA _{2010Q1}	0.021 (0.031)	0.021** (0.010)	-1.432 (1.466)	-1.432** (0.609)
Size _{2010Q1}	0.012** (0.005)	0.012*** (0.001)	0.001 (0.002)	0.001*** (0.000)
Tier1 _{2010Q1}	0.728*** (0.202)	0.728*** (0.061)	-0.076* (0.042)	-0.076*** (0.018)
Deposits _{2010Q1}	0.164*** (0.044)	0.164*** (0.012)	-0.007 (0.011)	-0.007 (0.004)
Liquidity _{2010Q1}	2.393 (2.159)	2.393*** (0.651)	-0.531 (0.843)	-0.531** (0.226)
Net Interbank Debt _{2010Q1}	0.158*** (0.057)	0.158*** (0.022)	0.033 (0.031)	0.033*** (0.007)
Bad Loans _{2010Q1}	-0.231 (0.160)	-0.231*** (0.062)	0.036 (0.081)	0.036 (0.023)
BCC	0.067** (0.026)	0.067*** (0.006)	-0.007 (0.006)	-0.007*** (0.002)
Firm FE	Y	Y	Y	Y
Sample	Multiple	Multiple	Multiple	Multiple
Cluster	Bank	Firm	Bank	Firm
Adj. R ²	0.599	0.599	0.585	0.585
Observations	424191	424191	464435	464435

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table examines the bank lending channel. It reports the estimates obtained from model (3.2) (Columns (1) and (2)) and model (3.3) (Columns (3) and (4)) on the sample of firms with multiple lending relationships appearing in the Credit Register. The outcome variable in Columns (1) and (2) ($\Delta\ln(\text{Loans})$) is the log-difference in average loans granted by bank b to firm j between (2010:Q2-2011:Q1) and (2009:Q2-2010:Q1). The outcome variable in Columns (3) and (4) (Credit Cut) is a dummy equal to one if bank b granted a loan to firm j before the sovereign shock (2009:Q2-2010:Q1) but did not renew the loan to firm j after the shock (2010:Q2-2011:Q1). The main independent variable is the stock of Italian sovereigns held by the lender at the end of 2010:Q1 scaled by RWA (Sovereigns_{2010Q1}). All regressions include a set of bank-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at firm or bank level depending on the specification. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Before we move forward, we present two additional empirical findings. First, we show that results are similar when we look at revolving credit lines, rather than term loans.²⁴ Results are reported in Column (1) of Table 3.3. Interestingly, the bank lending channel effect on credit lines appears smaller in magnitude. This is in line with recent research showing a particular sticky response of revolving credit lines to supply-side shocks because of “ever-greening” practices conducted by banks (Albertazzi & Marchetti 2010). The same results hold when looking at the total bank credit, i.e revolving credit lines plus term loans (Column (2)). Secondly, we show that banks do not discriminate borrowers based on risk or size. In particular, we augment equation (3.2) with an interaction term between sovereign exposure and firm size. As measure of firm size we use a dummy equal to one if the firm revenues in 2009 are below the median of the cross sectional distribution of revenues at the time.²⁵ Regressions are estimated on the CADS sub-sample. Results are reported in Table C.7 in Appendix C.9. We find that the direct effect of the shock is negative and significant, while the interactions are non-significant and small, providing evidence that the credit supply shocks was homogeneous across different types of borrowers.

²⁴The data from Italian Credit Register are in line with the patterns documented in Sufi 2009, as more than 90 percent (85% in the US) of the firms in our sample have at least one line of credit available.

²⁵The results are unchanged if we use log revenues as a proxy of size.

Table 3.3.: Bank Lending Channel: Total Credit and Credit Lines

<i>Dep. Var :</i>	$\Delta\ln(\text{Credit Lines})$	$\Delta\ln(\text{Tot. Bank Credit})$
	(1)	(2)
Sovereigns _{2010Q1}	-0.240** (0.107)	-0.270*** (0.099)
Bank Controls _{2010Q1}	Y	Y
Relationship Controls _{2010Q1}	Y	Y
Firm FE	Y	Y
Sample	Multiple	Multiple
Cluster	Bank	Bank
Adj. R ²	0.596	0.642
Observations	424191	424191

Sample: Firms with multiple lending relationships appearing in the Credit Register

In this table we explore the bank lending channel for alternative measures of change in bank credit. The table reports the estimates obtained from model (3.2) on the sample of firms with multiple lending relationships. We consider two alternative measures of bank credit. In Column (1) we focus only on the percentage change in the credit lines granted by bank b to firm j before-to-after the Greek bailout ($\Delta\ln(\text{Credit Lines})$). In Column (2) we consider the percentage change total amount of bank credit, including credit lines and term loans ($\Delta\ln(\text{Tot. Bank Credit})$). The main independent variable is the exposure of the lender to Italian sovereigns (Sovereigns_{2010Q1}). All regressions include a set of bank-specific and relationship-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

3.4.2.1. Identifying assumptions

As we pointed out in the previous section, the causal interpretation of our findings relies on the parallel trend assumption. We now provide evidence in line with this hypothesis by showing that differential exposure to sovereigns does not predict differential lending before the Greek bailout. Instead, we find a negative negative, significant effect of sovereign holdings on lending starting immediately after the bailout.

Figure 3.2.: The bank lending channel




Figure 3.2 illustrates the bank lending channel semi-parametrically by comparing lending to firms from banks with high holdings of Italian sovereign bonds, the most exposed to the sovereign shock, and banks with lower holdings. See appendix C.4 for a detailed description of the procedure used to construct this figure.

We start by showing these patterns hold at the aggregate level. First, we sort banks into a “*High Sovereign*” group and a “*Low Sovereign*” group based their (conditional) holdings of Italian sovereigns in the last quarter before the shock (2010:Q1). Second, for every quarter, we extrapolate variation in credit supply of bank b to firm j which cannot be explained by bank-specific characteristics. Then, we aggregate the residuals of corporate loans granted by “High sovereign” and those granted by “Low Sovereign” banks, and plot them over time in Figure 3.2. The two time series are normalized such that aggregate lending is zero in 2010:Q1 for each group.²⁶ Before the sovereign shock, the aggregate credit provided by the institutions with high and low holdings displays a very similar dynamic. However, after

²⁶Appendix C.4 provides a detailed description of the procedure followed to conduct the semi-parametric test.

April 2010 the two groups start diverging. More exposed intermediaries cut lending more extensively, while credit supply of less exposed bank does not react. These patterns present a first evidence in favor of the parallel trend assumption.²⁷

Figure 3.3.: Pre-trending test




Figure 3.3 presents a graphical test of the pre-trending assumption behind our identification strategy. It plots the coefficient capturing the correlation of sovereign bonds holdings in quarter 2010:Q1 ($Sovereigns_{b,2010Q1}$) - our proxy of the bank balance sheet shock - and the growth rate of credit between quarter t and quarter 2010:Q1 ($\ln(Loans_{ib,t}) - \ln(Loans_{ib,2010Q1})$) - the last quarter before the sovereign shock. Quarter t is reported on the x-axis. All regressions are run on the sample of firms who established multiple lending relationships, and include bank-level and relationship-level controls measured in 2010:Q1 and firm fixed-effects. To facilitate the comparison across periods, coefficients are reported as Z-scores. 90% confidence intervals are displayed. Standard errors clustered bank level.

To address the concern that more exposed banks might have experience more severe reduction in credit demand we turn to the micro-data, and provide a parametric test of the parallel trend assumption using our main specification. Figure 3.3 plots the coefficient

²⁷Further analysis shows no systematic match between banks more exposed to sovereigns and firms operating in particular industries or located in particular geographic regions. Tables available upon request.

β_1^I over time, capturing the relation between sovereign bonds holdings in quarter 2010:Q1 ($\text{Sovereigns}_{b,2010Q1}$) - our proxy of the bank balance sheet shock - and the growth rate of credit between quarter t and quarter 2010:Q1 ($\ln(\text{Loans}_{ib,t}) - \ln(\text{Loans}_{ib,2010Q1})$).²⁸ On the one hand, before the Greek bailout, we find a no significant difference in credit supply comparing banks who were differently exposed to sovereigns at the onset of the crisis. On the other hand, the graph displays a significant and long-lasting effect of the balance shock immediately after the Greek bailout and for the following quarters. All in all, the empirical evidence suggests that banks with lower sovereign holdings represent a valid control group for more exposed intermediaries, providing strong support for the identifying assumptions behind our empirical strategy.

3.4.2.2. Robustness tests

We conduct several econometric tests to evaluate the robustness of our results. First, our findings are robust to different assumptions about the variance-covariance structure of the errors and to an alternative measures of exposure to the sovereign crisis. Columns (1) and (3) and Columns (2) and (4) in Table 3.2 estimate our baseline regressions clustering standard errors at bank and firm level respectively. In both cases, our estimates are highly statistically significant. If anything, clustering errors at bank level produces standard errors that are substantially larger. Opting for the most conservative and conceptually correct choice, we maintain the bank-level clustering structure for the remaining part of the within-firm analysis.²⁹

²⁸Quarter t is reported on the x-axis. To facilitate the comparison across periods, coefficients are reported as Z-scores. Corroborating these results, Table C.6 in Appendix C.9 replicates the regression presented in Table 3.5 using as left-hand side variable the log difference in loans granted by bank b to firm j between 2010:Q1 (the last accounting period *before* the crisis) and quarter t ($\Delta \ln(\text{Loans})_{2010Q1-2009Q2}$). As in Model (3.2), all regressions include bank-level and relationship-level controls measured in 2010:Q1 and firm fixed-effects.

²⁹When we employ the full panel dimension of our data - like in Figure 3.3 - we can also adjust errors at both bank and quarter level (Petersen, 2009) without affecting the statistical significance of our inference.

Table 3.4.: The Bank Lending Channel: alternative measures

<i>Dep. Var :</i>	$\Delta \ln(\text{Loans})$		Cut Credit		$\Delta \ln(\text{Loans})$		Cut Credit	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sovereigns ^{TOT ASSETS} _{2010Q1}	-0.277*** (0.095)	-0.277*** (0.037)	0.076 (0.047)	0.076*** (0.013)				
Sovereigns ^{TIER1} _{2010Q1}					-0.016*** (0.005)	-0.016*** (0.002)	0.006* (0.003)	0.006*** (0.001)
Bank Controls _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Firm FE	(0.031)	(0.010)	(1.500)	(0.613)	(0.031)	(0.010)	(1.530)	(0.617)
Sample	Y	Y	Y	Y	Y	Y	Y	Y
Cluster	Multiple	Multiple	Multiple	Multiple	Multiple	Multiple	Multiple	Multiple
Adj. R ²	Bank	Firm	Bank	Firm	Bank	Firm	Bank	Firm
Observations	0.582	0.582	0.585	0.585	0.582	0.582	0.585	0.585
	424191	424191	464435	464435	424191	424191	464435	464435

This table examines the bank lending channel using the stock of Italian sovereigns held by the lender at the end of 2010:Q1 by Tier1 (Sovereigns^{TOT ASSETS}_{2010Q1}) and by Tier1 (Sovereigns^{TIER1}_{2010Q1}) as a measure of lenders' exposure to the sovereign shock. It reports the estimates obtained from model (3.2) and model (3.3) on the sample of firms with multiple lending relationships. The outcome variable in Columns (1), (2), (5) and (6) is the log-difference in average loans granted by bank *b* to firm *j* between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis. The outcome variable in Columns (3), (4), (7) and (8) is a dummy equal to one if bank *b* granted a loan to firm *j* before the onset of the sovereign crisis (2009:Q2-2010:Q1) but did not renew the loan to firm *j* after the crisis (2010:Q2-2011:Q1). All regressions include a set of bank-specific controls are measured at the end 2010:Q1 and firm fixed effects. Standard Errors are clustered at firm or bank level depending on the specification. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Second, we show that our results are not affected by our definition of bank's sovereign exposure. Table 3.4 shows that the effects of the balance sheet shock on credit supply is significant and similar in economic magnitudes if we scale banks' exposure to Italian sovereigns by alternative proxies of size - Total Assets and Tier1 capital - rather than risk-weighted assets. Moreover, in line with our economic intuition, we find a negative effect on lending associated to holdings of sovereign securities issued by EU other countries experiencing tensions on their sovereign market (GIPS) (Table C.8 in Appendix C.9, Column (2)), but a positive effect associated to holdings of German sovereigns (Column (4)). Both coefficients, however, are imprecisely estimated due to the strong home-bias characterizing the sovereign portfolio of Italian banks. The effect is both statistically and economically significant when we pool together all government debt issued by the GIPS plus Italy (GIPSI), or when we use the whole sovereign portfolio (respectively Columns (3) and (5)).

Third, we show that our results are not driven by difference in the nature of the credit relationship established by more and less exposed intermediaries. Because economies scale exist in information production, and information is durable and not easily transferable, firms with strong and long-lasting lending relationships are expected to be rationed less than others (Hoshi *et al.* 1990b,a, 1991; Petersen & Rajan 1994). Thus, our results would be biased upwards if banks with higher sovereign holdings systematically establish "weak" credit relationships with their borrowers. We augment our regressions with a set of relationship-specific variables that capture the pre-shock length and strength of the lending relationship between bank b and firm j the coefficient of the lending channel increases (Table 3.5). Our findings are unaffected. Because these controls allow us to obtain more precise estimates, we are going to include them in all our regressions.

Table 3.5.: The Bank Lending Channel w/ relationship controls

<i>Dep. Var :</i>	$\Delta\ln(\text{Loans})$	
	(1)	(2)
Sovereigns _{2010Q1}	-0.285*** (0.105)	-0.285*** (0.022)
ROA _{2010Q1}	0.055 (0.035)	0.055*** (0.009)
Size _{2010Q1}	0.017*** (0.006)	0.017*** (0.001)
Tier1 _{2010Q1}	0.996*** (0.218)	0.996*** (0.059)
Deposits _{2010Q1}	0.155*** (0.058)	0.155*** (0.012)
Liquidity _{2010Q1}	2.696 (3.812)	2.696*** (0.632)
Net Interbank Debt _{2010Q1}	0.044 (0.122)	0.044** (0.021)
Bad Loans _{2010Q1}	0.225 (0.292)	0.225*** (0.062)
BCC	0.106*** (0.028)	0.106*** (0.006)
<i>Lenght Relationship</i> _{2010Q1}	-0.023*** (0.000)	-0.023*** (0.000)
<i>Share Relationship</i> _{2010Q1}	-0.025*** (0.002)	-0.025*** (0.002)
Firm FE	Y	Y
Sample	Multiple	Multiple
Cluster	Bank	Firm
Adj. R ²	0.599	0.599
Observations	424191	424191

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table examines the bank lending channel. It reports the estimates obtained from model (3.2) estimated on the sample of firms with multiple lending relationships appearing in the Credit Register. The outcome variable in Columns (1) and (2) ($\Delta\ln(\text{Loans})$) is the log-difference in average loans granted by bank b to firm j between (2010:Q2-2011:Q1) and (2009:Q2-2010:Q1). The main independent variables is the stock of Italian sovereigns held by the lender at the end of 2010:Q1 scaled by RWA (Sovereigns_{2010Q1}). All regressions include a set of bank-specific and relationship-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at firm or bank level depending on the specification. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Lastly, we run a battery of placebo tests to rule out the possibility that the results presented in this paper reflect a “structural” negative correlation between holding of sovereigns in period $t - 1$ and future credit supply. Using our preferred specification (Equation 3.2) to control for credit demand dynamics, Figure 3.4 plots the coefficient capturing the correlation of sovereign bonds holdings in quarter $t - 1$ ($\text{Sovereigns}_{b,t-1}$) and the average growth rate between of credit $\Delta \ln(\text{Loans}_{ib,t})$, where $\Delta \ln(\text{Loans}_{ib,t}) = \ln(0.25 * \sum_{\tau=0}^3 \text{Loans}_{ib,t+\tau}) - \ln(0.25 * \sum_{\tau=-1}^{-4} \text{Loans}_{ib,t+\tau})$. All regressions include bank-level and relationship-level controls measured in time $t - 1$ and firm fixed-effects. To facilitate the comparison across periods, coefficients are reported as Z-scores. We find zero or positive correlation between sovereign holdings in $t - 1$ and changes in credit supply before the Greek default, when no tensions were present in sovereign markets.³⁰ It is only after the events in Greece that bank’s holdings of sovereign securities predict a subsequent credit tightening.³¹

³⁰If anything, these results suggest that, in normal times, financial institutions use government bonds as a storage of liquidity in expectation of future investments opportunities (Gennaioli *et al.* 2014).

³¹As discussed in Section 3.2, the Greek sovereign default triggered a change in the relationship between bank’s sovereign holdings and credit supply. After the burst of the sovereign crisis the correlation between government bonds and lending remains negative less precisely estimated as the more and more periods used to calculate the before-to-after change in credit supply are periods characterized by sovereign distress.

Figure 3.4.: Sovereign holdings and credit supply dynamics




Figure 3.4 investigates the relationship between the stock of sovereigns in bank’s portfolio and credit supply dynamics. It plots the coefficient capturing the correlation of sovereign bonds holdings in quarter $t - 1$ ($\text{Sovereigns}_{b,t-1}$) and the average growth rate of credit between $\Delta \ln(\text{Loans}_{ib,t})$, where $\Delta \ln(\text{Loans}_{ib,t}) = \ln(0.25 * \sum_{\tau=0}^3 \text{Loans}_{ib,t+\tau}) - \ln(0.25 * \sum_{\tau=-1}^{-4} \text{Loans}_{ib,t+\tau})$. Quarter t is reported on the x-axis. All regressions are run on the sample of firms who established multiple lending relationships, and include bank-level controls and relationship-level measure in quarter t and firm fixed effects. To facilitate the comparison across periods, coefficients are reported as Z-scores. 90% confidence intervals are displayed. Standard errors clustered bank level.

3.4.2.3. Correlated demand-side shocks

Restricting the analysis to the sample of firms with multiple lending relationships allowed us to include firm fixed effects to control for idiosyncratic demand shocks. Were banks more exposed to the sovereign crisis leading to firms experiencing particularly severe demand shocks? To answer this question, we proceed in steps. First, we apprise the information content of the estimated firm(-time) fixed effects and then we test for systematic sorting of exposed lenders and borrowers with poor investment opportunities and credit merit. Lastly,

we assess whether how our estimates of the lending channel change on the sample of single-lender firms.

Other studies employing a similar within-firm identification strategy have treated the estimated fixed effects as nuisance parameters (Gan 2007b; Khwaja & Mian 2008; Jiménez & Ongena 2012; Cingano *et al.* 2014; Jiménez *et al.* 2014). However, to the extent that they proxy real demand-side shocks, the estimated fixed effects may convey useful information on the transmission of the sovereign shock to the real economy. We take a step in this direction by conducting a two-stage analysis in which the fixed effects $\hat{\rho}_j$ are treated as the dependent variable in a second stage regression. A more negative $\hat{\rho}_j$ parameter should indicate more severe idiosyncratic shocks hitting firm j after the onset of sovereign crisis. If $\hat{\rho}_j$ captures changes in demand-side factors which affects firms' activity between the pre- and post-shock period, we should observe a significant correlation between the firm's fixed-effect estimates and proxies of riskiness and demand for credit.

Table 3.6.: The Bank Lending Channel: Lender-Borrower Sorting and Single Lender Firms

<i>Dep. Var :</i>	$\Delta \ln(\text{Loans})$		
	(1)	(2)	(3)
<i>Sovereigns</i> _{2010Q1}	-0.285*** (0.105)	-0.347*** (0.129)	-0.367*** (0.138)
Bank Controls _{2010Q1}	Y	Y	Y
Relationship Controls _{2010Q1}	Y	Y	Y
Firm FE	Y	N	N
Sample	Multiple	Multiple	Single
Cluster	Bank	Bank	Bank
Adj. R ²	0.599	0.086	0.086
Observations	424191	424191	114157

Sample: Firms with multiple and single lending relationships appearing in the Credit Register

This table reports the result for our baseline specification, presented in equation (3.2), estimated on different samples of firms and with/without firm fixed effects appearing in the Credit Register. The outcome variable is the log-difference in average loans granted by bank b to firm j between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis ($\Delta \ln(\text{Loans})$). Column (1) reports our baseline regression on the sample of firms with multiple lending relationships, including firm fixed effects. Columns (2) replicates column (1) without including firm fixed effects. Column (3) replicates column (2) on the sample of firms which established only one lending relationship. The main independent variable is the exposure of the lender to Italian sovereigns (*Sovereigns*_{2010Q1}). All regressions include a set of bank-specific and relationship-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Using firms in the CADS data set, we find a strong, positive correlation between the fixed effects and a firm's growth in revenues and assets between 2009 and 2010 (Table C.10, Appendix C.9). Similarly, the $\hat{\rho}_j$ are positively correlated with the credit score of the borrower after the onset of the crisis. While the estimates $\hat{\rho}_j$ are likely noisy, these findings corroborate the hypothesis that the firm fixed-effect estimated by model (3.2) capture and control for relevant information about changes firms' credit demand and creditworthiness.³²

³²In unreported regressions we show that there is positive and significant correlation of between $\hat{\rho}_j$ and each of the right-hand side variables presented Table C.10 individually. Also, we find a significant and sizable correlation with asset and revenues growth measured over over a longer period (2009-2011).

In the next step, we run our baseline model (3.2) without including firm fixed effects on the multi-lender firms and on the full sample of firms (including both firms with multiple and single lending relationships) to test possible biases driven by correlated demand-side shocks. Results are presented in Table 3.6. For exposition purposes, Column (1) reports the coefficient capturing the bank lending channel estimated by the baseline model of equation (3.2) (cfr. Table 3.5). Column (2) shows the estimates of the OLS regression using the sample of firms with multiple lending relationships, but without including firm fixed effects. While the magnitude of $\hat{\beta}_1^I$ increases when we do not control for firm fixed effects, statistical tests on the equality of these two coefficients cannot reject the null hypothesis, suggesting no systematic sorting of borrowers experiencing idiosyncratic shocks and banks more exposed to sovereign holdings.³³

Finally, we can estimate the baseline model without firm fixed effects on the sample of firms which established only one lending relationship to assess whether the lending channel confined to multi-lender firms (Column (3)).³⁴ The magnitude of the coefficient capturing the bank lending channel increases even further compared to the estimate on the multi-lender firms of Columns (1) and (2) but, also in this case, this difference is not statistically significant.

3.4.2.4. Transmission mechanism of the sovereign shock

We argued that the sovereign shock reduced banks' credit supply because it unexpectedly increased the riskiness of a large fraction of their assets that were considered safe until the Greek crisis. In this Section, we provide further evidence in favor of this hypothesis.

³³While not affecting the consistency of the OLS estimation, the results in Table 3.6 suggest that the inclusion of firm demand-side controls reduce the residual variance and significantly increases the precision of the estimates (standard errors fall by 20%, the adjusted R^2 increases by 50%).

³⁴The sample of firms with one lending relationships is a random sample of 70% of firms appearing in the Bank of Italy's Credit Register which established only one lending relationship over the period of interest and satisfy the filters described in Appendix C.6.

The literature proposes two main channels of transmission for the shock via direct holdings of distressed securities (Panetta *et al.* 2011). First, the sovereign shock may affect banks' lending by affecting its capital position. Even if these assets are not mark-to-market, regulators may exercise hard or soft pressure on banks with high holdings of distress sovereigns, for instance requiring banks to deleverage in response to the reduction of the market value of their sovereign portfolio. Furthermore, concerns about future health of the bank might have induced bank managers and shareholders to reduce the value at risk by shrinking the loan portfolio.³⁵ In both cases, if this *capital channel* is economically relevant, we should observe a stronger credit contraction for intermediaries that have a weaker balance sheet - such as a low bank capital.

An alternative transmission mechanism is the *collateral channel*. In this case, banks would be affected by the sovereign shock because the value of collateral available for interbank transactions gets reduced. This might be particularly relevant because a substantial part of inter-bank lending is collateralized by sovereign securities. In this case, banks that obtain a larger share of funding from interbank markets should be more affected by the sovereign shock and tighten credit supply more than others.

To shed some light on this issue, we test for heterogeneous responses to the shock across banks that are different in terms of capitalization and funding structure. To proxy the weakness of the balance sheet, we use the ratio of non-performing loans to RWA ($\text{Bad Loans}_{2010Q1}$) and the bank's Tier1 ratio (Tier1_{2010Q1}). To account for the non-linearity of the effect, we also test whether the impact of the shock was more severe for the subset of banks closer to the regulatory threshold defining a the dummy variable - $\text{Low Capital Ratio}_{2010Q1}$ - which takes value one when a bank's Tier1 ratio is below 10%.³⁶ To distinguish banks with differ-

³⁵This might have been the case, for instance, if excess exposure to risky sovereigns would affect banks' ability to access funding by questioning the long-term solvency of the institution. Since external investors do not observe the direct exposure of a specific intermediary, the shock may not necessarily affect immediately bank funding costs.

³⁶During the period of our analysis the regulatory threshold was set at 8%.

ent funding structure, we use the deposit-to-RWA ratio (Deposits_{2010Q1}) and the interbank borrowing-to-RWA ratio ($\text{Net Interbank Borrowing}_{2010Q1}$). Results are reported in Table 3.7. The data reject the collateral channel, while we find an economic and statistically significant effect of the capital channel. Importantly, a comparison of columns (1) and (5) reveals that the capital channel is highly non-linear, as distance from the regulatory threshold seems to matter more than capitalization per se. In fact, our results suggest that banks close to the capital regulatory threshold tighten their credit supply twice as much in response to the shock.

In Appendix (C.9) we push the analysis one step further, trying to discriminate between a capital channel triggered by a change in banks' perception of the riskiness of their balance sheet, and a capital channel resulting from "soft regulatory pressure." We identify two categories of banks which are differentially exposed to regulatory pressure: mutual banks and subsidiaries foreign banks. Mutual banks ("Banche di Credito Cooperativo", or BCC) are smaller and adopt a business model based on local fund-raising and local lending. Moreover, differently from the largest commercial banks, BCC are usually scrutinized by local branches of the national regulator. Similarly, subsidiaries of international banks may be affected by different incentives and/or receive different pressures when dealing with national regulators with respect to their national competitors. Comparing these two types of banks to the other financial intermediaries we do not find any significant difference in the magnitude of the lending channel (Table C.9). While not conclusive, these results are against "soft regulation" hypothesis and, indirectly, suggestive of the change in bank's risk management as the main driver of our results.

Table 3.7.: Transmission Mechanism of The Bank Lending Channel: Capital and Funding Channel

<i>Dep. Var.</i>	$\Delta \ln(\text{Loans})$						
	(1)	(2)	(3)	(4)	(6)	(7)	
Sovereigns _{2010Q1}	-0.174** (0.083)	-0.227** (0.103)	-0.304*** (0.099)	-0.392*** (0.141)	-0.275*** (0.103)	-0.275*** (0.103)	
Sovereigns _{2010Q1} inter. with: Low Capital Ratio _{2010Q1}	-0.235* (0.137)				-0.507*** (0.173)	-0.507*** (0.173)	
Tier1 _{2010Q1}		0.349 (0.366)				-0.508 (0.548)	
<i>Bad Loans</i> _{2010Q1}				-1.541 (1.131)	-1.763 (1.123)	-1.763 (1.123)	
Net Borrower Interbank _{2010Q1}			-0.139 (0.156)		0.128 (0.155)	0.128 (0.155)	
Deposits _{2010Q1}				0.101 (0.069)	0.156 (0.117)	0.156 (0.117)	
<i>Top5</i> _{2010Q1}					0.134 (0.339)	0.134 (0.339)	
Bank & Relationship Contr. _{2010Q1}	Y	Y	Y	Y	Y	Y	
Firm FE	Y	Y	Y	Y	Y	Y	
Sample	Multiple	Multiple	Multiple	Multiple	Multiple	Multiple	
Cluster	Bank	Bank	Bank	Bank	Bank	Bank	
Adj. R ²	0.599	0.599	0.600	0.599	0.600	0.600	
Observations	424191	424191	424191	424191	424191	424191	

This table investigates the channels of transmission of the sovereign shock through banks' balance sheet. It reports the estimates obtained from model (3.2) on the sample of firms with multiple lending relationships appearing on the Credit Register. We interact exposure to the sovereign shock with a set of bank characteristics which are proxies for alternative balance sheet channels of transmission. The outcome variable is the log-difference in average loans granted by bank b to firm j between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis ($\Delta \ln(\text{Loans})$). The main independent variable is the exposure of the lender to Italian sovereigns (Sovereigns_{2010Q1}), and its interactions with different proxies of the transmission channels. All regressions include a set of bank-specific and relationship-specific controls as measured at the end 2010:Q1. The interaction variables include: bad loans over RWA, a dummy equal one if the financial institution is a net borrower in interbank markets, Tier1 ratio, deposit over RWA, close capital (dummy equal 1 if Tier1 ratio of the bank is between 8 and 10 percent), Top 5 (a dummy equal 1 if the bank is one of the biggest five Italian banks). Every specification contains firm fixed effects. Standard Errors are clustered at bank level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

3.5. Credit supply and corporate behavior

The previous results confirm the presence of a sizable contraction in credit triggered by the turmoil in sovereign markets. We find that the shock to sovereign assets impaired the ability of banks to provide credit to firms. The next step is to evaluate whether this event had actual consequences on firm behavior. Lower credit can impair companies' ability to invest if firms cannot compensate lower credit from exposed lenders with more borrowing from other banks or with other sources of funding. Indeed, our results demonstrate that credit market frictions prevented firms to fully smooth out the bank lending channel, with sizable effects on firms investment rates and employment, especially among small firms.³⁷

As a first step, we construct a measure of firm-level exposure to the sovereign shock by computing the average exposure of its lenders. Lenders' exposures are weighted by the share of total bank loans the firm received by the bank before the Greek bailout. Formally, let \mathcal{B}_j be the set of all lenders to firm j in 2010:Q1. Then, we construct firm j 's average exposure as

$$\text{Sovereigns}_{j,2010Q1}^{AVE} = \frac{1}{B} \sum_{b \in \mathcal{B}_j} \omega_{bj} \text{Sovereigns}_{b,2010Q1}$$

where $\text{Sovereigns}_{b,2010Q1}$ is the stock of Italian sovereign over RWA held by lender b in 2010:Q1.³⁸ We then study how different firm-level outcomes (y_j) are affected by the firm specific exposure to the sovereign shock:

$$y_j = \alpha_0 + \alpha_1 \text{Sovereigns}_{j,2010Q1}^{AVE} + \Gamma \cdot X_{j,2010Q1}^{AVE} + \Lambda \cdot Z_{j,2010Q1}^{AVE} + u_j \quad (3.4)$$

The intuition behind this test is the following. If the sovereign shock had no effects on firms' operations, then the lenders' exposure to sovereign securities should not predict any change

³⁷Even if the firms are able to completely undo the bank lending channel by borrowing from banks less exposed to the shock or resorting to other forms of financing, the sovereign crisis might still propagate to the real economy through other channels. See for example Bocola (2013) or Neri & Ropele (2013).

³⁸In our data set, on average, the exposure of firms to the sovereign crisis is 22 percent (mean of $\text{Sovereigns}_{j,2010Q1}^{AVE}$), with a standard deviation of 23 percent (standard deviation of $\text{Sovereigns}_{j,2010Q1}^{AVE}$).

in outcomes ($\hat{\alpha}_1 = 0$). This would suggest that, despite the credit tightening, firms were able to take actions to effectively neutralize the negative effects of this drop in lending. If the opposite were true, the exposure of firms' lenders before the shock would still predict changes in y_j .

Similar to the within-firm model presented before, the identification of α_1 requires the orthogonality between the banks' exposure to sovereign securities and firms' unobservables that might affect credit demand and investment opportunities. Specifically, the concern is that banks more exposed to the sovereign shock may be matched with firms with worse investment opportunities and therefore these firms may have responded to the sovereign shock regardless of lenders' supply of credit. For instance, this sorting may be induced by geographic or industry clustering.

However, unlike model (3.2), we cannot directly control for unobservable demand-side shocks performing a within-firm estimation, since the unit of observation of this analysis is the firm rather than its (multiple) lending relationships. However, the previous findings suggest that systematic sorting between highly exposed banks and low quality firms is not a first-order problem in our setting. In particular, we have shown that the loan-level estimates with and without the firm-fixed are not statistically different, and therefore that the bias induced by demand is either nonexistent or relatively small. Nevertheless, we can still improve our identification in several ways.

First, we add a full set of province and industry fixed effects to every specification. Industry fixed effects control for specialization of exposed lenders in industries suffering more severe contractions of economic activity. Province fixed effects control for spatial clustering of banks and borrowers. These fixed-effects should take care of two important sources of matching between companies and banks. Secondly, we rely on the discussion in Section (3.4) and use the estimates of the firm fixed effects from model (3.2) - $\hat{\rho}_j$ - as a proxy of changes in demand-side factors (Albertazzi & Bottero 2013; Cingano *et al.* 2014). Given the importance

of demand-side factors in the context of the sovereign crisis, this choice allows us to improve the precision of our estimators. Finally, as a robustness exercise, we add a full set of firm level variables which can control for heterogeneity in characteristics across firms before the shock. All in all, this strategy gives us confidence in interpreting our estimates as causal. We estimate the following model

$$y_j = \alpha_0 + \alpha_1 \text{Sovereigns}_{j,2010Q1}^{AVE} + \hat{\rho}_j + \Gamma \cdot X_{j,2010Q1}^{AVE} + \Lambda \cdot Z_{j,2010Q1}^{AVE} + \tau_{province} + \tau_{industry} + u_j \quad (3.5)$$

where y_j is a firm-level outcome, $\text{Sovereigns}_{j,2010Q1}^{AVE}$ the previously described measure of exposure, $\tau_{province}$ and $\tau_{industry}$ are a set of province fixed effects and industry fixed effects (SIC 2-digits), and $\hat{\rho}_j$ are the set of firm-fixed effects estimated in the within-firm specification. We also control for a weighted average of lender-specific ($X_{j,2010Q1}^{AVE}$) and relationship-specific characteristics ($Z_{j,2010Q1}^{AVE}$) which might simultaneously affect firm j 's lenders exposure to the shock and their credit supply decisions.³⁹

3.5.1. Supply shock and credit market's access

The shock to sovereign holdings triggered a decline in credit supply via the bank lending channel. However, firms may have been able to limit the economic impact of the shock by borrowing from alternative, less exposed financial intermediaries. To investigate this issue, we estimate model (3.5) looking at the total change in bank credit before to after the shock. Using the previous notation, our dependent variable is $\Delta \ln(\text{Tot Loans}_j) = \sum_{b \in \mathcal{B}_j} \ln(\text{Loans}_{bj,POST}) - \sum_{b \in \mathcal{B}_j} \ln(\text{Loans}_{bj,PRE})$.

In Table (3.8), we show that firms have been unable to fully undo the decline in credit, as the average exposure their lenders at the onset of the sovereign shock is predictive of

³⁹The weighted averages of bank-specific and relationship-specific variables are constructed similar to $\text{Sovereigns}_{j,2010Q1}^{AVE}$. The only exception is the dummy for mutual bank which is equal to one if the major bank is a mutual bank. Standard errors are clustered at the level of the major bank.

the change in total bank credit. On average, to a one standard deviation (0.23) increase in lenders' average holdings of Italian sovereign securities corresponds a reduction of 5% of total bank borrowing before-to-after the sovereign shock.⁴⁰ As a robustness exercise, we estimate model (3.5) without controlling for firm-specific demand shocks (Column 2). Our estimate of the shock propagation does not significantly change, providing further evidence against the assortative matching hypothesis, but the inclusion of $\hat{\rho}$ improves the precision of our model. Similar results, but smaller in magnitude, hold if we restrict the analysis to the CADS sample (Columns 3 and 4). As we are going to show, this should not be surprising once we take into account the larger size of the firms in the CADS subsample.⁴¹ Lastly, results are also stable when the firm-level controls are added (Table C.12).

This effect is not only relevant from a microeconomic point of view, but has also sizable macroeconomic implications. As we discuss in Appendix, we conduct a simple exercise where we compare the actual bank credit available to firms after the crisis to a counterfactual amount of credit constructed under the assumption that the shock had no effect on credit supply. We then aggregate this micro estimates in order to isolate the size the decline in aggregate bank credit which can be attributed to the transmission of sovereign shock via the lending channel. We find that the initial sovereign shock led to an aggregate decline in corporate lending of about 2% within the first year following the Greek bailout request.

⁴⁰The coefficient is very similar ($\hat{\alpha}_1 = -0.219^{***}$) if we do not control for industry and province fixed effects. In unreported regressions, we also show that these results are similar when we measure firms' exposure to the sovereign crisis by looking only at the exposure of the major lender in 2010:Q1.

⁴¹The average size firms in the CADS sample is higher than the average size firm in the Credit Registry, both in terms of assets, revenues, and employment. As we are going to show, this selection on size can rationalize the smaller magnitude of the estimates using the CADS sample.

Table 3.8.: The Firm Borrowing Channel

<i>Dep. Var :</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\Delta \ln(\text{Tot Loans})$							
Sovereigns ^{AVE} _{2010Q1}	-0.207*** (0.021)	-0.238*** (0.065)	-0.121*** (0.033)	-0.111* (0.064)	-0.122** (0.051)	0.054 (0.096)	-0.057 (0.047)	-0.128 (0.129)
Sovereigns ^{AVE} _{2010Q1} x Sovereigns Prov _{2010Q1}					-0.748** (0.354)	-1.667* (0.947)		0.616 (1.304)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉							-0.111** (0.052)	0.307* (0.180)
Sovereigns ^{AVE} _{2010Q1} x Sovereigns Prov _{2010Q1}								-3.864** (1.724)
Small Firm ₂₀₀₉							0.011** (0.005)	-0.009 (0.023)
Small Firm ₂₀₀₉ x Sovereigns Prov _{2010Q1}								0.273 (0.239)
Estimated Firm FE ($\hat{\rho}_j$)	0.914*** (0.003)		0.782*** (0.005)		0.914*** (0.003)	0.783*** (0.005)	0.781*** (0.005)	0.781*** (0.005)
Bank & Relationship Contr. ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Sample	CR	CR	CADS	CADS	GR	CADS	CADS	CADS
Cluster	Lead	Lead	Lead	Lead	Lead	Lead	Lead	Lead
	Bank	Bank	Bank	Bank	Bank	Bank	Bank	Bank
Adj. R ²	0.685	0.089	0.644	0.057	0.685	0.627	0.627	0.627
Observations	188381	188381	35123	35123	188381	35123	35123	35123

This table examines the firm borrowing channel. It reports the estimates obtained from model (3.5) on the sample of firms with multiple lending relationships appearing in the Credit Register and on the subs-sample appearing in the CADS database. The outcome variable is the log-difference in average total bank loans granted to firm j between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis ($\Delta \ln(\text{Tot Loans})$). The main independent variable is the weighted average of the exposure to Italian sovereign scaled by RWA of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). In Column (1) we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated from the baseline regression of the bank lending channel (equation (3.2)). Column (2) presents the results of the same econometric model estimated in Column (1), but without controlling for demand-side shocks ($\hat{\rho}_j$). Columns (3) and (4) replicate the regressions in Column (1) and (2) on the CADS sample. Columns (5)-(9) investigate the heterogeneity of the effect. The interaction variables include: a dummy equal to one if firm j 's revenue in 2009 is below the median across firms at the time (Small Firm₂₀₀₉), and the average exposure of banks operating in the same province of the firm (Sovereigns Province_{2010Q1}). All regressions include a set of weighted averaged bank-specific and relationship-specific controls measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1. Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

3.5.2. Information frictions and market segmentation

Firms were unable to compensate lower credit from exposed banks by borrowing more from other financial intermediaries. In this Section, we try to understand why this was the case by exploring different channels that may have prevented firms to smooth the credit shock.

Frictions related to information asymmetries and agency costs can impair the efficient allocation of credit. With this respect, the literature puts special emphasis on firm size. Smaller companies are generally characterized by larger information asymmetries and lower transparency, making them more vulnerable to credit supply shocks (e.g. Gertler & Gilchrist 1994). We test this heterogeneity adding an interaction between sovereign exposure and a dummy Small Firm_{2009} , which identifies firms whose revenue was below the median of the revenue in 2009.⁴² In Table (3.8), we find that most of the negative effect on credit is actually driven by smaller companies. In particular, the interaction between the size dummy and the shock is negative, large and strongly statistically significant. The economic magnitude of these effects is large: a firm in the bottom half of the size distribution experienced almost a 3% higher decline in credit relative to an equivalent larger firm at the average level of the sovereign shock. As we discussed before (Section 3.4), this result is not driven by the fact that banks were strategically cutting lending to smaller companies.

Furthermore, the geographical segmentation of the banking sector has been identified as a very relevant friction in the credit markets (Petersen & Rajan 2002; Guiso *et al.* 2004; Degryse & Ongena 2005). If lending markets are geographically segmented, firms' ability to hedge the sovereign shock depends on the presence of other local financial intermediaries less exposed to the sovereign shock. As a result, we expect firms operating in areas with higher concentration of exposed banks to experience the highest reduction of bank credit following the sovereign shock. In order to test this hypothesis, we construct a proxy of geographical

⁴²As discussed in Section (3.3), when we use firm characteristics we only look at firms within the CADS sample, for which we have reliable firm level data.

exposure to the sovereign shock as the average exposure of the financial intermediaries that operate in the province ($\text{Sovereigns}_{j,2010Q1}$), which we interact with the firm-level credit shock ($\text{Sovereigns}_{j,2010Q1}^{AVE}$).⁴³ As reported in Table (3.8), the interaction between the firm-level credit supply shock with the province-level shock is negative and significant, confirming that the geographical segmentation of credit markets affects firms ability to smooth out the bank lending channel.⁴⁴ This effect is large: for any given firm-level shock, a firm located in a province at the 25th percentile of sovereign exposure received 2% more credit than a similarly exposed firm in province located at 75th percentile of the distribution.

Moreover, credit market segmentation should be particularly relevant when information frictions are more important. For example, we expect the burden of credit market segmentation to be disproportionately born to smaller firms, since large firms can more easily tap credit outside the local area. We explore this issue in Table (3.8). We compare the effect of being in a local market characterized by large overall sovereign exposure across large and small firms. In line with the economic predictions, we find that the previous result on geographical segmentation was almost entirely driven by smaller companies. While we cannot completely rule out that geographical segmentation matters also for larger companies, this result suggests that the hurdle imposed on small companies is at least twice the average effect identified before. Overall, all these results are confirmed when we augment the specification with firm-level controls measured before the shock (Table C.12).

⁴³Italian provinces roughly compare to US counties and, as discussed in details by Guiso *et al.* (2013), they constitute a natural geographical unit for small business lending. Firms are distributed across 110 provinces.

⁴⁴The effect is more precisely estimated when we use the full CR sample. For the CADS sample, the same analysis delivers similar but less precise result reflecting the smaller sample of firms in this sample and the bias of CADS towards larger firms.

3.5.3. Real effects on investments and employment

If corporations cannot access alternative sources of financing, the drop in credit might have a sizable impact on firms' real activity and, thus, on the economy as a whole. Previous studies have shown that, when financial markets are imperfect, the availability of external debt financing directly affects firms' investment (Gan 2007a; Gan 2007b; Acharya *et al.* 2015) and employment decisions (Benmelech *et al.* 2011; Bentolila *et al.* 2013; Chodorow-Reich 2014; Greenstone *et al.* 2014).

To shed light on the real effects of the sovereign crisis, we employ the specification in equation (3.5) and regress the average sovereign exposure of a firm's lenders on changes in investments and employment. The key problem in the real-effects literature is the identification of credit constrained firms (Fazzari *et al.* 1988; Kaplan & Zingales 1997). Our approach has the advantage of directly observing which firms experience a restricted access to credit using lenders' exposure to sovereigns, which - we have shown - is orthogonal to firms' characteristics and investment opportunities. Due to data availability, we restrict our investigation to the sub-sample of multi-lender firms appearing in the CADS database for which we have accounting information on assets, revenues, and employment.⁴⁵ We construct two proxies that capture both the intensive and extensive margin dynamics of firms' capital expenditures (Acharya *et al.* 2015). The first is a dummy indicating an decrease in firm's fixed assets between 2009 and 2011 ($1\{\Delta\text{Inv} < 0\}$). The second is the log-difference of fixed assets between 2011 and 2009 ($\%\Delta\text{Inv}$). Similarly, we construct two measures of change in employment to investigate the spillover on labor demand: decrease ($1\{\Delta\text{Empl} < 0\}$) and percentage change

⁴⁵In unreported regressions we show that all previous results on credit remain stable in terms of economic magnitude and statistical significance on this sub-sample of firms. Regressions are available upon request.

Table 3.9.: Real Effects: Sovereign Exposure on Investments and Employment

<i>Dep. Var.</i>	1($\Delta Invest < 0$)	% $\Delta Invest$	1{ $\Delta Empl < 0$ }	% $\Delta Empl$				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sovereigns ^{AVE} _{2010Q1}	0.074 (0.077)	0.119 (0.085)	-0.131 (0.088)	0.044 (0.098)	-0.066 (0.148)	0.110 (0.163)	0.015 (0.068)	-0.021 (0.081)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉		0.335*** (0.096)		-0.308** (0.132)		-0.194 (0.203)		0.075 (0.097)
Small Firm ₂₀₀₉		0.002 (0.010)		0.021 (0.014)		0.026 (0.021)		-0.016 (0.010)
Estimated Firm FE ($\hat{\rho}_j$)	-0.202*** (0.008)	-0.201*** (0.009)	0.253*** (0.010)	0.253*** (0.011)	-0.152*** (0.009)	0.129*** (0.010)	0.093*** (0.007)	0.092*** (0.007)
Bank Controls ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Relationship Controls ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Sample	CADS	CADS	CADS	CADS	CADS	CADS	CADS	CADS
Cluster	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.
Adj. R ²	0.031	0.032	0.033	0.033	0.037	0.050	0.043	0.044
Observations	35165	35165	35165	35165	16672	16672	16672	16672

This table examines the effects of the sovereign crisis on employment transmitted via the lending channel. It reports the estimates obtained from model (3.5) where the dependent variables are two alternative proxies of firm's investments and employment. In Columns (1) and (2) we use a dummy which is equal to 1 if the firm's invested or increased its labor force between 2009 and 2011. In Columns (3) and (4) we look at the growth rate in investments and employment over the same period. The main independent variable is the weighted average of the exposure to Italian sovereigns of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). Columns (1) and (3) show the baseline effect on investments and employment. Columns (2) and (4) investigate the heterogeneity of the effect on investments across firms of different size. Size is constructed as a dummy if firm revenue in 2009 is below the median across firms at the time (Small Firm₂₀₀₉), for the firms for which control and outcomes are available. All regressions include a set of weighted averaged bank-specific and relationship-specific controls measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1, and we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated in the baseline regression of the bank lending channel (equation (3.2)). Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

in employment ($\% \Delta \text{Empl}$).⁴⁶

Table (3.9) presents the results for investments and employment. We find that, on average, the credit shock had little or no effect on investments. In columns (1) and (2), the coefficient associated to lenders' sovereign exposure is negative and, at best, barely significant. The same holds for the measures of employment growth discussed above. However, this small average effect masks a large heterogeneity across firms. Comparing firms above and below the median size, we find that sovereign shock was indeed very relevant for smaller companies (Columns 3 and 4). This is true for changes of investments at the extensive and intensive margin.⁴⁷ Relative to large companies, for every percentage change in the lenders' exposure the shock smaller firms decrease investment 0.3 percent more over the period of interest and they are 0.3% more likely to cut on assets. Robustness tests comparing the change in investments between 2007 and 2009 exclude that these results are driven by pre-trending (Table C.11). Similarly, the results are similar in size and statistical power when adding firm-controls (Table C.13). As we discuss more later in the Section, this is consistent with the fact that small firms were overall more affected by the reduction in bank credit.

The same does not hold for employment, for which we cannot detect any significant effect both on average and across size groups.⁴⁸ This difference between employment and investment may just reflect the smaller sample of firms for which we have employment information. However, an alternative explanation may be the strictness of employment protection law in Italian (Schivardi & Torrini 2008). Since these policies increase the cost of adjusting the workforce, employment should be less responsive than asset to credit shocks. We provide

⁴⁶Information on the workforce is available only for a subsample of firms the firms in CADS. A comparison of the observable characteristics of firms with and without employment information reveals that the latter are smaller firms both in terms of total assets and revenues. Thus, our estimates are likely a lower bound of the true impact of the credit shock on employment.

⁴⁷In unreported regressions we show that we obtain when using a discrete measure of size constructed using firms' assets and when using a continuous measures like asset and revenues.

⁴⁸Since employment is not available for every firm in CADS, we re-define the size dummy for this outcome in the same way as before within this sample.

some suggestive evidence in line with this explanation by showing that the credit shock actually mattered for the subset of firms for which adjustments would be more beneficial: i.e. companies that utilized labor less efficiently. Following the literature on factor utilization, we construct a firm-level measure of labor productivity - the log ratio of value added to employment in 2009 -, and compare the effect of the credit supply shock on employment across firms characterized by different degrees of labor efficiency *within* narrowly defined industries (Hsieh & Klenow 2009). In line with economic intuition, Table C.15 shows that the credit shock led to a significant reduction of employment for firms that over-utilized labor.

Another important source of heterogeneity is the level of dependence to external finance that characterizes a firm. Firms that are more heavily dependent on external finance to run their operations will be – all else equal – more affected by a credit shock. In order to explore this aspect, we rely on underlying technological differences among industries and construct a measure of dependence on external finance following the classification proposed by Rajan & Zingales (1998). More specifically, we use construct the Rajan-Zingales (RZ) index for the firms of US using Compustat, and imputed it to firms in our data set according to their industry.⁴⁹ In line with our hypothesis, we find that firms that are more dependent on external finance cut employment - both at the intensive and extensive margin - relatively more than other firms in response to a shock (Table 3.10). Similar results can be found when looking at investment, at least for the extensive margin.

Consistent with the previous results, the degree of dependence on external finance seems to matter mostly for small firms (Table 3.10). In particular, the decision to cut employment for

⁴⁹The intuition behind the RZ index is the following. For technological reasons some industries rely on external finance more than others. For example some industries operate in larger scale than others, have projects with longer gestation or require continuous investments to keep operating; thus, these industries should suffer more than others an unexpected tightening of credit supply. We refer to Rajan & Zingales (1998) for further details.

Table 3.10.: Real Effects: Size and Dependence on External Finance

<i>Dep. Var :</i>	1($\Delta Invest < 0$)		% $\Delta Invest$		1{ $\Delta Empl < 0$ }		% $\Delta Empl$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sovereigns ^{AVE} _{2010Q1}	0.066 (0.076)	-0.128 (0.089)	-0.173** (0.086)	0.054 (0.099)	-0.015 (0.157)	0.025 (0.173)	-0.031* (0.070)	-0.051 (0.083)
Sovereigns ^{AVE} _{2010Q1} x RZ Index ₂₀₀₉	-0.009 (0.018)	-0.010 (0.027)	-0.036* (0.019)	0.001 (0.023)	0.054** (0.024)	-0.006 (0.043)	-0.041* (0.021)	-0.025 (0.022)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉		0.340*** (0.109)	-0.428*** (0.137)			-0.079 (0.218)		0.044 (0.095)
Sovereigns ^{AVE} _{2010Q1} x RZ Index ₂₀₀₉ x Small Firm ₂₀₀₉		0.002 (0.036)	-0.084* (0.050)			0.102* (0.052)		-0.028 (0.041)
RZ Index ₂₀₀₉ x Small Firm ₂₀₀₉		-0.001 (0.003)	0.010** (0.005)			-0.006 (0.006)		0.001 (0.003)
Small Firm ₂₀₀₉		-0.008 (0.011)	0.035*** (0.013)			0.017 (0.023)		-0.013 (0.011)
Estimated Firm FE ($\hat{\rho}_j$)		-0.203*** (0.008)	0.255*** (0.011)			-0.151*** (0.009)		0.092*** (0.007)
Bank & Relationship Contr. ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Sample	CADS	CADS	CADS	CADS	CADS	CADS	CADS	CADS
Cluster	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.	Lead B.
Adj. R ²	0.031	0.031	0.032	0.033	0.042	0.043	0.044	0.045
Observations	34740	34740	34740	34740	16445	16445	16445	16445

This table examines the effects of the sovereign crisis on corporate investments and employment transmitted via the lending channel. It reports the estimates obtained from model (3.5) on the sample of firms with multiple lending relationships appearing in the CADS database. The dependent variables are two proxies of firm investments and employment. The main independent variable is the weighted average of the exposure to Italian sovereigns of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). We interact with the firm level shock with a proxy of firm's size (Small Firm₂₀₀₉) and dependence on external finance (RZ Index). All regressions include a set of weighted averaged bank-specific and relationship-specific controls are measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1, and we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated in the baseline regression of the bank lending channel (equation 3.2). Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

smaller companies is more sensitive to the credit shock when the company is operating in an industry highly dependent on external finance. By contrast, external dependence does not matter much for larger firms. These findings are in line with those in Duygan-Bump *et al.* (2015) for the US, where the authors find that workers in small firms were more likely to become unemployed during the 2007–2009 recession than comparable workers in large firms, but only if they were employed in industries with high financing needs.

Altogether, our results suggest that the credit tightening caused by the sovereign crisis had a disproportional effect on small companies. In fact, smaller firms experienced a relatively larger drop in bank credit, investment and, to a certain extent, employment. As discussed in Section 3.4, this is not driven by banks cutting credit more extensively to smaller companies. Instead, the results suggest that smaller firms were less able than larger firms to compensate a credit shortage of equal magnitude across different lenders. This is particularly costly also because small businesses have less funding opportunities outside bank credit.⁵⁰ Overall, our results suggest that credit supply shocks are particularly disruptive for small businesses, even when credit markets do not discriminate them (Khwaja & Mian 2008; Chodorow-Reich 2014).

3.6. Conclusion

Using a detailed firm-bank panel data set extracted from the Italian Credit Registry, we document the propagation of sovereign tensions spurred by the Greek bailout to the real economy through a deterioration of banks' balance sheet. Tensions in the sovereign market led to a reduction in the amount of credit supplied by banks. Comparing lending to the

⁵⁰While we do not have a reliable measure of bond market access for the firms in our database, the sporadic use of arms-length financing in Italy and the prevalence of small-size firms in our sample (a firm at the 99th percentile has 773 employees) suggest that credit market frictions might be the most important driving force of our results.

same firm by two banks one-standard deviation apart in terms of sovereign holding, we find that the bank more exposed to sovereigns reduced the loan supply 7% more than the other bank in the year after the Greek bailout. Similarly, more exposed banks are more likely to terminate ongoing lending relationships. This effect is larger for banks that had a weak balance sheet (low capital), especially those for which the regulatory capital constraint was more binding.

Firms were unable to smooth out the shock by compensating lower borrowing from highly impaired banks with more credit from less impaired banks. Furthermore, we show that the geographical segmentation of credit markets helps explaining firms' inability to smooth the credit shock. In fact, firms face a larger decline in overall credit when they operate in areas where banks were on average more exposed to the sovereign market. These effects are mostly driven by small companies that are more reliant on local banks to get credit.

We document the sizable real effects of the sovereign shock on small firms operations. We find that small companies cut investments and employment, especially when relying heavily on external financing, more than larger companies, which appear to be almost not affected. This heterogeneous response does not appear to be supply driven, as we show that banks did cut lending more extensively to small firms. Instead, this seems to reflect a higher sensitivity of small firms to bank credit shocks. Smaller firms are more sensitive to bank credit as they cannot easily replace bank loans with other forms financing, and pay a disproportionately larger price during times of crisis.

Our analysis provides also new empirical evidence on the possible real costs connected with shocks to the security portfolio held by banks. In particular, we highlight the possible risk characterizing a financial system encumbered by large public debt. Shedding light on the mechanisms of propagation of a credit shock, our analysis show a large heterogeneity across companies in the costs of credit shocks to the economy.

Bibliography

- ABBASSI, PURIYA, BRAEUNING, FALK, FECHT, FALKO, & PEYDRO, JOSÉ-LUIS. 2014. *Cross-border liquidity, relationships, and monetary policy: Evidence from the euroarea interbank crisis*. Tech. rept. mimeo. Deutsche Bundesbank.
- ACHARYA, VIRAL, & STEFFEN, SASCHA. 2013. The "greatest" carry trade ever? Understanding eurozone bank risks. *NBER Working paper*.
- ACHARYA, VIRAL, DRECHSLER, ITAMAR, & SCHNABL, PHILIPP. 2011a. A pyrrhic victory? Bank bailouts and sovereign credit risk. *NBER Working paper*.
- ACHARYA, VIRAL, EISERT, TIM, EUFINGER, CHRISTIAN, & HIRSCH, CHRISTIAN. 2015. Real effects of the sovereign debt crises in Europe: evidence from syndicated loans. *Working paper*.
- ACHARYA, VIRAL V, & SUBRAMANIAN, KRISHNAMURTHY V. 2009. Bankruptcy codes and innovation. *Review of Financial Studies*, hhp019.
- ACHARYA, VIRAL V, AMIHUD, YAKOV, & LITOV, LUBOMIR. 2011b. Creditor rights and corporate risk-taking. *Journal of Financial Economics*, **102**(1), 150–166.
- AHTIK, META, & ALBERTAZZI, UGO. 2014. The supply of credit to sovereigns and privates is there crowding out? *Mimeo ECB*.
- AKCIGIT, UFUK, & KERR, WILLIAM R. 2010. *Growth through heterogeneous innovations*. Tech. rept. National Bureau of Economic Research.
- ALBERTAZZI, UGO, & BOTTERO, MARGHERITA. 2013. The Procyclicality of Foreign Bank Lending: Evidence from the Global Financial Crisis. *Journal of International Economics*, **92**(S1), **S22-S35**.
- ALBERTAZZI, UGO, & MARCHETTI, DOMENICO J. 2010. Credit supply, flight to quality and evergreening: an analysis of bank-firm relationships after Lehman. *Bank of Italy Working paper*.
- ALMEIDA, HEITOR, CAMPELLO, MURILLO, & WEISBACH, MICHAEL S. 2004. The cash flow sensitivity of cash. *The Journal of Finance*, **59**(4), 1777–1804.
- ANGELINI, PAOLO, GRANDE, GIUSEPPE, & PANETTA, FABIO. 2014. The negative feedback loop between banks and sovereigns. *Bank of Italy Working Paper*.
- ANGRIST, JOSHUA D, & PISCHKE, JÖRN-STEFFEN. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.
- ARENA, MATTEO P, & JULIO, BRANDON. 2014. The Effects of Securities Class Action Litigation on Corporate Liquidity and Investment Policy. *Journal of Financial and Quantitative Analysis*.
- AUGUSTIN, PATRICK, BOUSTANIFAR, HAMID, BRECKENFELDER, JOHANNES H, & SCHNITZLER, JAN. 2014. Sovereign credit risk and corporate borrowing costs. *Working Paper*.

- AXELSON, ULF, JENKINSON, TIM, STRÖMBERG, PER, & WEISBACH, MICHAEL S. 2013. Borrow cheap, buy high? The determinants of leverage and pricing in buyouts. *The Journal of Finance*, **68**(6), 2223–2267.
- BALDUZZI, PIERLUIGI, BRANCATI, EMANUELE, & SCHIANTARELLI, FABIO. 2014. Financial Markets, Banks' Cost of Funding, and Firm's Decisions: Lessons from Two Crisis. *Working paper*.
- BATTISTINI, NICCOLÒ, PAGANO, MARCO, & SIMONELLI, SAVERIO. 2013. Systemic risk and home bias in the euro area. *European Commission Economic Papers*.
- BEATTY, ANNE, WEBER, JOSEPH, & YU, JEFF JIEWEI. 2008. Conservatism and debt. *Journal of Accounting and Economics*, **45**(2), 154–174.
- BECKER, BO, & IVASHINA, VICTORIA. 2014. Financial Repression in the European Sovereign Debt Crisis. *Swedish House of Finance Research paper*.
- BENMELECH, EFRAIM, BERGMAN, NITTAI K, & SERU, AMIT. 2011. Financing labor.
- BENTOLILA, SAMUEL, JANSEN, MARCEL, JIMÉNEZ, GABRIEL, & RUANO, SONIA. 2013. When credit dries up: Job losses in the great recession.
- BENZONI, LUCA, COLLIN-DUFRESNE, PIERRE, GOLDSTEIN, ROBERT S, & HELWEGE, JEAN. 2014. Modeling credit contagion via the updating of fragile beliefs.
- BERGSTRÖM, CLAS, GRUBB, MIKAEL, & JONSSON, SARA. 2007. The operating impact of buyouts in Sweden: A study of value creation. *The Journal of Private Equity*, **11**(1), 22.
- BERNSTEIN, SHAI. 2015. Does going public affect innovation? *The Journal of Finance*.
- BERNSTEIN, SHAI, & SHEEN, ALBERT. 2013. The operational consequences of private equity buyouts: Evidence from the restaurant industry. *Rock Center for Corporate Governance at Stanford University Working Paper*.
- BERNSTEIN, SHAI, LERNER, JOSH, SORENSEN, MORTEN, & STROMBERG, PER. 2015. Private Equity and Industry Performance. *Management Science*, **51**, 3–44.
- BERTRAND, MARIANNE, DUFLO, ESTHER, & MULLAINATHAN, SENDHIL. 2004. How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics*, **119**(1), 249–275.
- BESSEN, JAMES, & MEURER, MICHAEL J. 2008a. Of patents and property. *Regulation*, **31**, 18.
- BESSEN, JAMES, & MEURER, MICHAEL J. 2013. Direct Costs from NPE Disputes, The. *Cornell L. Rev.*, **99**, 387.
- BESSEN, JAMES, & MEURER, MICHAEL JAMES. 2008b. *Patent failure: How judges, bureaucrats, and lawyers put innovators at risk*. Princeton University Press.
- BESSEN, JAMES E., NEUHÄUSLER, PETER, TURNER, JOHN L., & WILLIAMS, JONATHAN W. 2015. Trends in Private Patent Costs and Rents for Publicly-Traded United States Firms. *Boston Univ. School of Law, Public Law Research Paper*.
- BLUNDELL-WIGNALL, ADRIAN. 2007. The Private Equity Boom. *Financial Market Trends*, **2007**(1), 59–86.
- BOCOLA, LUIGI. 2013. The Pass-Through of Sovereign Risk. *Manuscript, University of Pennsylvania*.

- BOFONDI, MARCELLO, CARPINELLI, LUISA, & SETTE, ENRICO. 2013. Credit supply during a sovereign debt crisis. *Bank of Italy Temi di Discussione (Working Paper) No*, **909**.
- BOLDRIN, MICHELE, & LEVINE, DAVID. 2002. The Case against Intellectual Property. *American Economic Review*, 209–212.
- BOUCLY, QUENTIN, SRAER, DAVID, & THESMAR, DAVID. 2011. Growth IboS. *Journal of Financial Economics*, **102**(2), 432–453.
- BRAV, OMER. 2009. Access to capital, capital structure, and the funding of the firm. *The Journal of Finance*, **64**(1), 263–308.
- BROWN, JAMES R, FAZZARI, STEVEN M, & PETERSEN, BRUCE C. 2009. Financing innovation and growth: Cash flow, external equity, and the 1990s R&D boom. *The Journal of Finance*, **64**(1), 151–185.
- CALABRESI, GUIDO, & MELAMED, A DOUGLAS. 1972. Property rules, liability rules, and inalienability: one view of the cathedral. *Harvard Law Review*, 1089–1128.
- CARPINELLI, LUISA, & CROSIGNANI, MATTEO. 2015. The Effect of Central Bank Liquidity Injections on Bank Credit Supply.
- CARRIER, MICHAEL A. 2011. *Innovation for the 21st Century*. Oxford University Press.
- CASIRAGHI, MARCO, GAIOTTI, EUGENIO, RODANO, MARIA LISA, & SECCHI, ALESSANDRO. 2013. The impact of unconventional monetary policy on the Italian economy during the sovereign debt crisis. *Bank of Italy Occasional Paper*.
- CHETTY, RAJ, LOONEY, ADAM, & KROFT, KORY. 2009. Salience and Taxation: Theory and Evidence. *The American Economic Review*, **99**(4), 1145–1177.
- CHIEN, COLLEEN V, & LEMLEY, MARK A. 2012. Patent holdup, the ITC, and the public interest. *Cornell Law Review*.
- CHODOROW-REICH, GABRIEL. 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.
- CINGANO, FEDERICO, MANARESI, FRANCESCO, & SETTE, ENRICO. 2014. Does credit crunch investments down? New evidence on the real effects of the bank-lending channel. *MoFIR Working paper*.
- CLAESSENS, STIJN, & LAEVEN, LUC. 2003. Financial development, property rights, and growth. *The Journal of Finance*, **58**(6), 2401–2436.
- COHEN, LAUREN, GURUN, UMIT, & KOMINERS, SCOTT DUKE. 2014. Patent Trolls: Evidence from Targeted Firms. *HBS Working Paper*.
- CORREA, RICARDO, LEE, KUAN-HUI, SAPRIZA, HORACIO, & SUAREZ, GUSTAVO A. 2014. Sovereign credit risk, banks' government support, and bank stock returns around the world. *Journal of Money, Credit and Banking*, **46**(s1), 93–121.
- COURT, SUPREME. 1908. *Continental Paper Bag Co. v. Eastern Paper Bag Co.*
- COURT, SUPREME. 2006. *eBay Inc. v. MERCEXCHANGE, LL.*

- COURT, VIRGINIA DISTRICT. 2003. *MercExchange, LLC v. eBay, Inc.*
- CUMMING, DOUGLAS, SIEGEL, DONALD S, & WRIGHT, MIKE. 2007. Private equity, leveraged buyouts and governance. *Journal of Corporate Finance*, **13**(4), 439–460.
- DASS, NISHANT, NANDA, VIKRAM K, & XIAO, STEVEN CHONG. 2015. Truncation Bias in Patent Data: Does it Explain Why Stock-Liquidity Seemingly Reduces Innovation? *Working Paper*.
- DAVIS, STEVEN J, HALTIWANGER, JOHN, HANDLEY, KYLE, JARMIN, RON, LERNER, JOSH, & MIRANDA, JAVIER. 2014. Private Equity, Jobs, and Productivity. *American Economic Review*, **104**(12), 3956–90.
- DE MARCO, FILIPPO. 2014. Bank lending and the sovereign debt crisis. *Boston College Working Paper*.
- DEGRYSE, HANS, & ONGENA, STEVEN. 2005. Distance, lending relationships, and competition. *The Journal of Finance*, **60**(1), 231–266.
- DEMIRGÜÇ-KUNT, ASLI, & MAKSIMOVIC, VOJISLAV. 1998. Law, finance, and firm growth. *Journal of Finance*, 2107–2137.
- DETRAGIACHE, ENRICA, GARELLA, PAOLO, & GUISO, LUIGI. 2000. Multiple versus single banking relationships: Theory and evidence. *The Journal of Finance*, **55**(3), 1133–1161.
- DJANKOV, SIMEON, LA PORTA, RAFAEL, LOPEZ-DE SILANES, FLORENCIO, & SHLEIFER, ANDREI. 2003. Courts. *The Quarterly Journal of Economics*, 453–517.
- DUCHIN, RAN, OZBAS, OGUZHAN, & SENSOY, BERK A. 2010. Costly external finance, corporate investment, and the subprime mortgage credit crisis. *Journal of Financial Economics*, **97**(3), 418–435.
- DUYGAN-BUMP, BURCU, LEVKOV, ALEXEY, & MONTORIOL-GARRIGA, JUDIT. 2015. Financing constraints and unemployment: evidence from the Great Recession. *Journal of Monetary Economics*.
- EPSTEIN, RICHARD A. 2008. The property rights movement and intellectual property. *Regulation*, **30**(4), 2007–2008.
- FARRE-MENSA, JOAN, & LJUNGQVIST, ALEXANDER. 2015. Do Measures of Financial Constraints Measure Financial Constraints? *Review of Financial Studies*.
- FAZZARI, STEVEN M, HUBBARD, R GLENN, PETERSEN, BRUCE C, BLINDER, ALAN S, & POTERBA, JAMES M. 1988. Financing Constraints and Corporate Investment. *Brookings Papers on Economic Activity*, 141–206.
- FELDMAN, ROBIN. 2014. Patent Demands & Startup Companies: The View from the Venture Capital Community. *Yale Journal of Law & Technology*, **16**(2), 236.
- FENG, JIASHUO, & JARAVEL, XAVIER. 2015. Patent Trolls and the Patent Examination Process.
- FISHER, RONALD A. 1922. On the interpretation of χ^2 from contingency tables, and the calculation of P. *Journal of the Royal Statistical Society*, 87–94.
- GAN, JIE. 2007a. Collateral, debt capacity, and corporate investment: Evidence from a natural experiment. *Journal of Financial Economics*, **85**(3), 709–734.

- GAN, JIE. 2007b. The real effects of asset market bubbles: Loan-and firm-level evidence of a lending channel. *Review of Financial Studies*, **20**(6), 1941–1973.
- GENNAIOLI, NICOLA, MARTIN, ALBERTO, & ROSSI, STEFANO. 2013. Banks, government bonds, and default: what do the data say? *IMF Working Paper*.
- GENNAIOLI, NICOLA, MARTIN, ALBERTO, & ROSSI, STEFANO. 2014. Sovereign default, domestic banks, and financial institutions. *The Journal of Finance*, **69**(2), 819–866.
- GERTLER, MARK, & GILCHRIST, SIMON. 1994. Monetary Policy, Business Cycle, and the Behavior of Small Manufacturing Firms. *Quarterly Journal of Economics*, **109**(2), 309–340.
- GIORDANO, RAFFAELA, PERICOLI, MARCELLO, & TOMMASINO, PIETRO. 2013. Pure or Wake-up-Call Contagion? Another Look at the EMU Sovereign Debt Crisis. *International Finance*, **16**(2), 131–160.
- GOLDSTEIN, MORRIS. 1998. *The Asian financial crisis: Causes, cures, and systemic implications*. Vol. 55. Peterson Institute.
- GORMLEY, TODD A, & MATSA, DAVID A. 2014. Common errors: How to (and not to) control for unobserved heterogeneity. *Review of Financial Studies*, **27**(2), 617–661.
- GREENSTONE, MICHAEL, MAS, ALEXANDRE, & NGUYEN, HOAI-LUU. 2014. Do credit market shocks affect the real economy? Quasi-experimental evidence from the Great Recession and 'normal' economic times. *NBER Working paper*.
- GUIO, LUIGI, SAPIENZA, PAOLA, & ZINGALES, LUIGI. 2004. Does Local Financial Development Matter? *The Quarterly Journal of Economics*, **119**(3), 929–969.
- GUIO, LUIGI, PISTAFERRI, LUIGI, & SCHIVARDI, FABIANO. 2013. Credit within the Firm. *The Review of Economic Studies*, **80**(1), 211–247.
- HALL, BRONWYN H, & LERNER, JOSH. 2010. The financing of R&D and innovation. *Handbook of the Economics of Innovation*, **1**, 609–639.
- HALL, BRONWYN H, JAFFE, ADAM B, & TRAJTENBERG, MANUEL. 2001. *The NBER patent citation data file: Lessons, insights and methodological tools*. Tech. rept. National Bureau of Economic Research.
- HALL, BRONWYN H, JAFFE, ADAM, & TRAJTENBERG, MANUEL. 2005. Market value and patent citations. *RAND Journal of economics*, 16–38.
- HANSON, SAMUEL G, KASHYAP, ANIL K, & STEIN, JEREMY C. 2011. A Macroprudential Approach to Financial Regulation. *The Journal of Economic Perspectives*, **25**(1), 3–28.
- HARRIS, RICHARD, SIEGEL, DONALD S, & WRIGHT, MIKE. 2005. Assessing the impact of management buyouts on economic efficiency: Plant-level evidence from the United Kingdom. *Review of Economics and Statistics*, **87**(1), 148–153.
- HASLEM, BRUCE. 2005. Managerial opportunism during corporate litigation. *The Journal of Finance*, **60**(4), 2013–2041.
- HELM, JEREMIAH 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, RYAN T. 2015. The Misinterpretation of eBay v. MercExchange and Why: An Analysis of the Case History, Precedent, and Parties. *Chapman Law Review*.
- HOSHI, TAKEO, KASHYAP, ANIL, & SCHARFSTEIN, DAVID. 1990a. Bank monitoring and investment: Evidence from the changing structure of Japanese corporate banking relationships. In "**Asymmetric information, corporate finance, and investment**", R. Glenn Hubbard, editor: University of Chicago Press.
- HOSHI, TAKEO, KASHYAP, ANIL, & SCHARFSTEIN, DAVID. 1990b. The role of banks in reducing the costs of financial distress in Japan. *Journal of Financial Economics*, **27**(1), 67–88.
- HOSHI, TAKEO, KASHYAP, ANIL, & SCHARFSTEIN, DAVID. 1991. Corporate structure, liquidity, and investment: Evidence from Japanese industrial groups. *The Quarterly Journal of Economics*, 33–60.
- HOTCHKISS, EDIE, SMITH, DAVID C, & STRÖMBERG, PER. Private Equity and the Resolution of Financial Distress.
- HOUSE, WHITE. 2013. Patent Assertion and U.S. Innovation.
- HSIEH, CHANG-TAI, & KLENOW, PETER J. 2009. *Misallocation and Manufacturing TFP in China and India*. Tech. rept. 4.
- IVASHINA, VICTORIA, & KOVNER, ANNA. 2011. The private equity advantage: Leveraged buyout firms and relationship banking. *Review of Financial Studies*, hhr024.
- IVERSON, BENJAMIN CHARLES. 2014. Get in line: Chapter 11 restructuring in crowded bankruptcy courts. *Working Paper*.
- JAFFE, ADAM B, & LERNER, JOSH. 2011. *Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it*. Princeton University Press.
- JIMÉNEZ, GABRIEL, & ONGENA, STEVEN. 2012. Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications. *The American Economic Review*, **102**(5), 2301–2326.
- JIMÉNEZ, GABRIEL, ONGENA, STEVEN, PEYDRÓ, JOSÉ-LUIS, & SAURINA, JESÚS. 2014. Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking? *Econometrica*, **82**(2), 463–505.
- KAPLAN, STEVEN. 1989. The effects of management buyouts on operating performance and value. *Journal of Financial Economics*, **24**(2), 217–254.
- KAPLAN, STEVEN N, & STEIN, JEREMY C. 1993. The evolution of buyout pricing and financial structure (or, what went wrong) in the 1980s. *Journal of Applied Corporate Finance*, **6**(1), 72–88.
- KAPLAN, STEVEN N, & STROMBERG, PER. 2009. Leveraged Buyouts and Private Equity. *Journal of Economic Perspectives*, **23**(1), 121–46.
- KAPLAN, STEVEN N, & ZINGALES, LUIGI. 1997. Do investment-cash flow sensitivities provide useful measures of financing constraints? *The Quarterly Journal of Economics*, 169–215.
- KAPLOW, LOUIS, & SHAVELL, STEVEN. 1996. Property rules versus liability rules: An economic analysis. *Harvard Law Review*, 713–790.

- KASHYAP, ANIL K, & LAMONT, OWEN A. 1994. Credit conditions and the cyclical behavior of inventories. *Quarterly Journal of Economics*, **109**(3).
- KASHYAP, ANIL K, STEIN, JEREMY C, & WILCOX, DAVID W. 1992. Monetary policy and credit conditions: Evidence from the composition of external finance.
- KERR, WILLIAM R. 2010. Breakthrough inventions and migrating clusters of innovation. *Journal of Urban Economics*, **67**(1), 46–60.
- KHWAJA, ASIM IJAZ, & MIAN, ATIF. 2008. Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *The American Economic Review*, 1413–1442.
- KIM, IRENE, & SKINNER, DOUGLAS J. 2012. Measuring securities litigation risk. *Journal of Accounting and Economics*, **53**(1), 290–310.
- KING, ROBERT G, & LEVINE, ROSS. 1993. Finance and growth: Schumpeter might be right. *The Quarterly Journal of Economics*, 717–737.
- KOGAN, LEONID, PAPANIKOLAOU, DIMITRIS, SERU, AMIT, & STOFFMAN, NOAH. 2012. *Technological innovation, resource allocation, and growth*. Tech. rept. National Bureau of Economic Research.
- KORTUM, SAMUEL, & LERNER, JOSH. 2000. Assessing the contribution of venture capital to innovation. *RAND Journal of Economics*, 674–692.
- LA PORTA, RAFAEL, LOPEZ-DE SILANES, FLORENCIO, SHLEIFER, ANDREI, & VISHNY, ROBERT W. 1997. Legal determinants of external finance. *Journal of Finance*, 1131–1150.
- LANE, PHILIP R. 2012. The European sovereign debt crisis. *The Journal of Economic Perspectives*, **26**(3), 49–67.
- LANJOUW, JEAN, & SCHANKERMAN, MARK. 2001. Characteristics of patent litigation: a window on competition. *RAND Journal of Economics*, 129–151.
- LANJOUW, JEAN O, & LERNER, JOSH. 1998. The Enforcement of Intellectual Property Rights: A Survey of the Empirical Literature. *Annales d'Économie et de Statistique*, 223–246.
- LEMLEY, MARK A, & SHAPIRO, CARL. 2005. Probabilistic patents. *Journal of Economic Perspectives*, 75–98.
- LEMLEY, MARK A, & SHAPIRO, CARL. 2006. Patent holdup and royalty stacking. *Tex. L. Rev.*, **85**, 2163.
- LERNER, JOSH. 1995. Patenting in the Shadow of Competitors. *Journal of Law and Economics*, **38**(2), 463–95.
- LERNER, JOSH. 2006. Trolls on State Street?: The litigation of financial patents, 1976-2005. *Harvard Business School manuscript*.
- LERNER, JOSH. 2009. *The empirical impact of intellectual property rights on innovation: Puzzles and clues*.
- LERNER, JOSH, & SCHOAR, ANTOINETTE. 2005. Does legal enforcement affect financial transactions? The contractual channel in private equity. *The Quarterly Journal of Economics*, 223–246.
- LERNER, JOSH, & SERU, AMIT. 2015. The use and misuse of patent data: Issues for corporate finance and beyond.

- LI, GUAN-CHENG, LAI, RONALD, D'AMOUR, ALEXANDER, DOOLIN, DAVID M, SUN, YE, TORVIK, VETLE I, AMY, Z YU, & FLEMING, LEE. 2014. Disambiguation and co-authorship networks of the US patent inventor database (1975–2010). *Research Policy*, **43**(6), 941–955.
- MANN, WILLIAM. 2013. Creditor rights and innovation: Evidence from patent collateral. *Working Paper*.
- MERLER, SILVIA, & PISANI-FERRY, JEAN. 2012. Who's afraid of sovereign bonds?
- MICHAELY, RONI, & ROBERTS, MICHAEL R. 2012. Corporate dividend policies: Lessons from private firms. *Review of Financial Studies*, **25**(3), 711–746.
- MOSER, PETRA. 2013. Patents and innovation: evidence from economic history. *The Journal of Economic Perspectives*, 23–44.
- NERI, STEFANO, & ROPELE, TIZIANO. 2013. The macroeconomic effects of the sovereign debt crisis in the euro area.
- ONGENA, STEVEN, & SMITH, DAVID C. 2000. What determines the number of bank relationships? Cross-country evidence. *Journal of Financial Intermediation*, **9**(1), 26–56.
- PAKES, ARIEL. 1986. Patents as options: some estimates of the value of holding European patent stocks. *Econometrica*, **54**(4), 755–784.
- PANETTA, FABIO, CORREA, RICARDO, DAVIES, MICHAEL, DI CESARE, ANTONIO, MARQUES, JOSÉ-MANUEL, NADAL DE SIMONE, FRANCISCO, SIGNORETTI, FEDERICO, VESPRO, CRISTINA, VILDO, SIRET, & WIELAND, MARTIN. 2011. The impact of sovereign credit risk on bank funding conditions. *CGFS Papers, No. 5*.
- PERICOLI, MARCELLO, & SBRACIA, MASSIMO. 2003. A primer on financial contagion. *Journal of Economic Surveys*, **17**(4), 571–608.
- PETERSEN, MITCHELL A. 2009. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies*, **22**(1), 435–480.
- PETERSEN, MITCHELL A, & RAJAN, RAGHURAM G. 1994. The benefits of lending relationships: Evidence from small business data. *The Journal of Finance*, **49**(1), 3–37.
- PETERSEN, MITCHELL A, & RAJAN, RAGHURAM G. 2002. Does distance still matter? The information revolution in small business lending. *The Journal of Finance*, **57**(6), 2533–2570.
- PONTICELLI, JACOPO. 2013. Court enforcement and firm productivity: evidence from a bankruptcy reform in Brazil. *Chicago Booth Research Paper*.
- POPOV, ALEXANDER, & VAN HOREN, NEELTJE. 2013. The impact of sovereign debt exposure on bank lending: Evidence from the European debt crisis. *Review of Finance*, forthcoming.
- RAJAN, RAGHURAM G, & ZINGALES, LUIGI. 1998. Financial Dependence and Growth. *The American Economic Review*, **88**(3), 559–586.
- ROGERS, JONATHAN L, & VAN BUSKIRK, ANDREW. 2009. Shareholder litigation and changes in disclosure behavior. *Journal of Accounting and Economics*, **47**(1), 136–156.
- SAKAKIBARA, MARIKO, & BRANSTETTER, LEE. 2001. *Do stronger patents induce more innovation? Evidence from the 1988 Japanese patent law reforms*.

- SCHIVARDI, FABIANO, & TORRINI, ROBERTO. 2008. Identifying the effects of firing restrictions through size-contingent differences in regulation. *Labour Economics*, **15**(3), 482–511.
- SCHUMPETER, JOSEPH ALOIS. 1934. *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*. Vol. 55. Transaction publishers.
- SCHWARTZSTEIN, JOSHUA, & SHLEIFER, ANDREI. 2013. An Activity-Generating Theory of Regulation. *Journal of Law and Economics*, **56**, 1–38.
- SHAPIRO, CARL. 2010. Injunctions, Hold-Up, and Patent Royalties. *American Law and Economics Review*.
- SMEETS, ROGER. 2014. Does patent litigation reduce corporate R&D? an analysis of US public firms. *Working Paper*.
- STEIN, JEREMY C. 1998. An Adverse-Selection Model of Bank Asset and Liability Management with Implications for the Transmission of Monetary Policy. *RAND Journal of Economics*, **29**(3), 466–486.
- STIGLITZ, JOSEPH E, & WEISS, ANDREW. 1981. Credit rationing in markets with imperfect information. *The American Economic Review*, 393–410.
- SUFI, AMIR. 2009. Bank lines of credit in corporate finance: An empirical analysis. *Review of Financial Studies*, **22**(3), 1057–1088.
- TANG, YIXIN H. 2006. Future of Patent Enforcement after Ebay v. MercExchange, The. *Harv. JL & Tech.*, **20**, 235.
- TUCKER, CATHERINE. 2014. Patent Trolls and Technology Diffusion: The Case of Medical Imaging.
- TUCKER, CATHERINE. 2015. The Effect of Patent Litigation and Patent Assertion Entities on Entrepreneurial Activity. *Research Policy*.
- VENKATESAN, JAIDEEP. 2009. Compulsory Licensing on Nonpracticing Patentees after eBay v. MercExchange. *Va. JL & Tech.*, **14**, 26.
- WESENBERG, ERIC, & O’ROURKE, PETER. 2006. *The Toll on the Troll*.
- WILLIAMS, HEIDI L. 2015. Intellectual Property Rights and Innovation: Evidence from Health Care Markets. *In: Innovation Policy and the Economy, Volume 16*. University of Chicago Press.
- WILLIAMSON, OLIVER E. 1975. Markets and hierarchies. *New York*, 26–30.
- WILLIAMSON, OLIVER E. 1983. Credible commitments: Using hostages to support exchange. *The American Economic Review*, 519–540.

A. Appendix to Chapter 1

A.1. 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.¹ 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

¹The patent is USPTO number 5,845,265.

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.2. The timing of the decision

The ruling should affect the trajectory of innovation only if it was not anticipated by agents in this market.² 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 & 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.³ 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,

²The 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.

³MarketWatch 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

pundits see a potential split decision in the eBay v. MercExchange injunction case.”⁴ 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.⁵ 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.⁶

A.3. Data

A.3.1. Samples

This is a more detailed clarification on Section (1.3) 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 (Uni-

⁴See 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

⁵See 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>

⁶Helm (2006) argues that these difference stems from the different use of injunction for large firms across these industries.

versity of California at Berkeley),⁷ 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 statistics that USPTO provides online.⁸ 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%.⁹

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

⁷Data can be found: <http://funginstitute.berkeley.edu/tools-and-data> (downloaded in August 2014)

⁸Table can be found here: http://www.uspto.gov/web/offices/ac/ido/oeip/taf/h_at.htm#PartA2_1

⁹This 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.

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 headquarter in USA. I exclude firms with missing data in the quarterly reports -total asset and revenue, which should always be populated when there is a report- around the two years window and firms with systematic negative equity over the 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 (1.3), 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 & Schankerman 2001).

Using the online tool LitAlert, I searched for all the litigation involving patents between 1980 and 2006.¹⁰ 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 (1.3), 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 (1.2) to construct the final score at technology class level.

A.3.2. 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

¹⁰http://intranetsolutions.westlaw.com/practicepages/template/ip_litalert.asp?rs=IPP2.0&vr=1.0

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 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.¹¹ 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

¹¹The outcome variable is constructed looking only at the two years before and after the shock, as in the intensive margin measure.

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 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 which are consistently adjusted at quarterly level assuming equal R&D across quarters within the fiscal year.

A.3.3. 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.¹²

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.

¹²I use the stock price and the number of shares provided by CRSP to compute the market value of equity.


A.4. Additional Figures

Figure A.1.: 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 A.2.: 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.3.: Persistence of patent litigation over time




This figure provides a scatter plot of the size of litigation technology class level, as measured by equation (1.2), 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 (1.3). 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.4.: Effect of litigation over time



This Figure plots the β_t from equation 1.4, using the usual sample. With respect to the other Figure (A.4), 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.3). 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.5.: 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 (1.2.2). Abnormal returns are constructed with respect to the S&P500, as discussed in the Appendix (A.3.3). The straight red line correspond to the trading day right before the decision.

A.5. Additional Tables

Table A.1.: Distribution of Litigation

	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 (1.1), and in particular by equation (1.2).

Table A.2.: Effect of the policy change on patenting: Poisson model

	(1)	(2)	(3)
Poisson			
	$\#Patents_{jt}$		
$Post \cdot Exposure_j$	0.045***	0.049**	0.040**
	(0.016)	(0.022)	(0.020)
<i>Firm F.E.</i>	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y
<i>Indu. \times Time F.E.</i>		Y	Y
<i>Other Controls_{jt}</i>			Y
Observations	256,944	256,944	256,944

This Table reports the estimate of the standard difference-in-difference specification (equation 2.1) 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.3). 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.3.: Stock Market returns and Litigation Exposure

	<i>Event Day</i>		<i>Event [-1; +1]</i>		<i>Event [-5; -1]</i>		<i>Event [-20; -5]</i>		<i>Event [-40; -5]</i>	
	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>	<i>Ret.</i>	<i>A.Ret.</i>
<i>Mean</i>	-0.034*** (0.008)	-0.038*** (0.008)	-0.036** (0.012)	-0.033** (0.013)	-0.076*** (0.015)	-0.064*** (0.014)	-0.026 (0.045)	-0.071 (0.046)	0.129* (0.057)	0.063 (0.054)
<i>Obs.</i>	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 (1.2), 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.3.3). *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table A.4.: Robustness: Leave-out-Industry

Leave out:	$\ln(Patents_{jt})$							
	Chemical	Computer	Drug	Electronic	Mechanical	Others		
<i>Post · Exposure_j</i>	0.043*** (0.009)	0.038*** (0.009)	0.061*** (0.016)	0.037*** (0.011)	0.035*** (0.009)	0.033*** (0.011)	0.039*** (0.009)	0.036*** (0.011)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y
<i>Indu. × Time F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y
<i>Other Contrals_{jt}</i>	Y	Y	Y	Y	Y	Y	Y	Y
Observations	28,408	25,938	26,448	25,982	26,204	26,204	27,610	27,610

In this table we report the estimation of the equation 2.1. 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.5.: Effect of the policy change on patenting: alternative *Exposure* measure

	(1)	(2)	(3)	(4)	(5)	(6)
OLS		$\ln(Patents_{jt})$			$1\{Patent_{jt} > 0\}$	
$Post \cdot Exposure_{j}^{LARGE}$	0.050*** (0.012)	0.049*** (0.016)	0.047*** (0.015)	0.017*** (0.004)	0.038*** (0.004)	0.043*** (0.004)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
$Indu. \times Time F.E.$		Y	Y	Y	Y	Y
<i>Other Controls</i> _{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 (2.3), using an alternative measure to firm exposure to litigation. In particular, I estimate equation (2.1), which is $y_{jt} = \alpha_j + \alpha_t + \beta(Exposure_{j, Post}) + \gamma X_{jt} + \epsilon_{jt}$, 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.3). 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)	(2)	(3)	(4)	(5)	(6)
	$\ln(\text{Patents}_{jt})$	$\ln(\text{Patents}_{jt} + 1)$	$\text{Patent}_{jt}^{\text{scaled}}$	$1\{\text{Patent}_{jt} = \text{Top}^{10\%}\}$	$1\{\text{Patent}_{jt} = \text{Top}^{25\%}\}$	$\ln(\text{Citations}_{jt} + 1)$
<i>Post · Exposure_j</i>	0.0165* (0.009)	0.014*** (0.004)	0.010*** (0.003)	0.004* (0.002)	0.010*** (0.003)	0.015*** (0.007)
<i>Pre · Exposure_j</i>	0.003 (0.010)	0.004 (0.004)	0.002 (0.003)	0.001 (0.002)	0.001 (0.003)	0.006 (0.007)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	Y	Y	Y	Y	Y	Y
<i>R</i> ²	0.001	0.009	0.016	0.001	0.002	0.004
Observations	105,922	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 (2.1), 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 by Column (1). 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 of granted patent that firm *j* applied during period *t*. Column (2) has the (natural) logarithm plus one of granted patent that firm *j* applied during period *t*. Column (3) 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 (4) and (5) 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 (6) 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.3). 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 across technologies

	$\ln(Patents_{jt})$					
	1 Year After		2 Years After		3 Years After	
$Post \cdot Exposure_j$	0.002		0.037***		0.054***	
	(0.001)		(0.010)		(0.010)	
$Post \cdot Exposure_j \cdot$ $\cdot Computer$	0.092***		0.023		-0.006	
	(0.020)		(0.021)		(0.022)	
$Post \cdot Exposure_j^{LARGE}$		0.001		0.035***		0.052***
		(0.009)		(0.010)		(0.010)
$Post \cdot Exposure_j^{LARGE} \cdot$ $\cdot Computer$		0.096***		0.029		0.001
		(0.019)		(0.020)		(0.022)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y
<i>Time F.E.</i>	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 2.1, where I interact the shock measure with a dummy for firms that are in the Computer industry, as defined in Appendix A.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 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** (0.002)	0.006** (0.003)		
$Pre \cdot Exposure_j^{LARGE}$	0.001 (0.002)	0.001 (0.002)		
$Post \cdot Exposure_j$			0.002* (0.001)	0.004** (0.001)
$Pre \cdot Exposure_j$			0.001 (0.001)	-0.001 (0.002)
$Firm\&Time\ F.E.$	Y	Y	Y	Y
$Indu. \times Time\ F.E.$		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 (2.1), 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 two 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.3). 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 (0.017)	-0.009 (0.015)	0.040*** (0.002)	0.027*** (0.002)
$Pre \cdot 1\{Risk_r\}$	-0.016 (0.018)	0.0015 (0.015)	-0.002 (0.003)	0.001 (0.002)
<i>Split</i>	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 (I.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 - 2006:1- 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: robustness

	(1)	(2)	(3)	(4)
	<i>R&D_{jt}/Asset_{jt}</i>			
	<i>Standard</i>	<i>Add Growth</i>	<i>Standard</i>	<i>Add Growth</i>
<i>Post · Exposure_j · Small_j^{REV}</i>	0.004** (0.002)	0.004** (0.002)	0.003* (0.002)	0.003** (0.002)
<i>Post · Exposure_j · Small_j^{EMP}</i>	0.004** (0.002)	0.004** (0.002)	0.003 (0.002)	0.003* (0.002)
<i>Post · Exposure_j · 1{Div_j = 0}</i>	0.005** (0.002)	0.007*** (0.002)	0.005* (0.003)	0.007*** (0.002)
<i>Post · Exposure_j · 1{Rating_j = NO}</i>	0.004** (0.002)	0.003** (0.002)	0.002 (0.002)	0.002 (0.002)
<i>Firm F.E. & Time F.E.</i>	Y	Y	Y	Y
<i>Indu. × Time F.E.</i>		Y		Y
<i>Other Controls_{jt}</i>		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.3). 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%.

B. Appendix to Chapter 2

B.1. Data and variable construction

All the data in the paper comes from Amadeus/Orbis produced by Bureau Van Dijk. To minimize the chances that data errors could drive our results, we winsorize at 1% every ratio used in the analysis. The winsorization is done at the level of the full company sample, in order to benchmark the ratios to the whole population distribution.


The main variables we used in the analysis are the following: (a) Capex/Asset, where Capex is constructed as the change in asset in the last year, plus the depreciation; (b) (Net) Equity contribution/Asset, where the equity contribution is measured as the difference in total equity (shareholder value) over the last year, minus the profit; (c) Leverage, which is simply total debt (short and long-term) scaled by asset; (d) $\ln(\text{Debt})$, which is simply the natural log of total debt; (e) Working Capital/Asset, as defined in the data set; (f) ROA: profit before taxes over total assets.

In the paper, we also use different methodologies to determine companies that are more or less likely to be financial constrained as the financial crisis unfolds. First, we use size by looking at the top quartile of revenue or employment at 2007. This is constructed within the analyzed firms. Second, we identify firms that are at the top quartile of leverage at 2007, as previously defined. Third, we identify companies operating in industries more dependent

on external finance. The index is constructed using Compustat in the following way. Using the whole set of consolidated reports in Compustat between 1980 and 2007, we construct a score for every two digit SIC code which is median of CAPEX minus operating cash flow, scaled by CAPEX.

B.2. Additional Figures

Figure B.1.: Private Equity Activity



These Figures report some data on the time series of private equity activity in USA. The first Figure reports the amount of private equity fundraising across three major categories: VC, Buyouts and Other (residual group). The series go from 1969 and 2013. The scale is in real 2013 US Billion \$. The second Figure reports the value multiple in private equity investment over time, in particular between 1980 and 2012. As multiple, we report the enterprise value scaled by EBITDA. Lastly, the third Figure reports the equity share in Buyout transaction between 1980 and 2013. By construction, the number is between 0 and 100.

B.3. Additional Tables

Table B.1.: Summary Statistics alternative matching: Level

Level	Matching on ROA-Assets-Leverage-2 Digit SIC						
	N.	Mean	SD	N.	Mean	SD	Diff.
	PE-Backed Company			Non PE Company			
Assets	630	118.52	661.75	2740	78.77	261.54	39.75*
Revenue	613	89.57	219.87	2650	78.07	256.06	11.50
Employment	568	443.67	1085.18	2324	272.14	702.84	171.5***
ROA	630	0.08	0.24	2740	0.06	0.2	0.013
CAPEX % As.	523	0.17	0.29	2296	0.19	0.24	-0.02
Net Contribution % As.	534	-0.01	0.16	2510	0.01	0.13	-0.02*
Tot. Liabilities % As.	630	0.73	0.44	2740	0.68	0.38	0.05**
Long Debt % As.	629	0.19	0.28	2737	0.15	0.25	0.04**
Bank Debt % As.	571	0.25	0.4	2434	0.23	0.42	0.03
Debt/EBITDA	600	5.62	30	2470	5.11	28.91	0.51
Long Debt/EBITDA	599	1.45	7.6	2467	1.4	6.68	0.06
Bank Debt/EBITDA	591	2.61	17.6	2432	2.78	17.45	-0.18

This Table reports the summary statistics of the companies at 2007 across treated (PE-backed companies) and untreated firms (non PE companies), as well as the mean difference across the two groups when we perform the matching based on leverage, ROA, size and industry. In particular, it reports the mean characteristics in level at 2007.

Table B.2.: Summary Statistics alternative matching: Changes

Panel A	Matching on ROA-Assets-Leverage-2 Digit SIC						
	N.	Mean	SD	N.	Mean	SD	Diff.
Growth rates 1 yr.	PE-Backed Company			Non PE Company			
Growth Asset	478	0.27	0.68	2064	0.32	0.74	-0.05
Growth Revenue	457	0.54	2.39	1909	0.52	2.34	0.02
Growth Employment	425	0.07	0.31	1682	0.08	0.28	-0.01
Growth CAPEX	364	1.68	7.27	1659	1.84	7.22	-0.16
Growth Contribution	378	-0.04	5.33	1778	0.14	5.59	-0.18
Growth Leverage	426	0.01	0.25	1957	0.01	0.24	-0.01
Change rates 1 yr.							
Change Asset	481	17.48	94.73	2078	12.26	96.32	5.23
Change Revenue	457	8.31	72.99	1909	7.52	71.31	0.79
Change Employment	425	27.65	189.19	1682	10.23	147.52	17.43*
Change Ln(Asset)	478	0.2	0.75	2064	0.21	0.46	-0.02
Change Ln(Revenue)	457	0.11	0.89	1909	0.14	0.77	-0.03
Change Ln(Employment)	425	0.05	0.43	1682	0.04	0.34	0.01
Change CAPEX	364	0.02	0.39	1659	0.01	0.33	0.01
Change Contribution	379	0.02	0.19	1780	0.01	0.18	0.01
Change Leverage	427	0.01	0.31	1959	-0.01	0.24	0.01
Panel B							
Growth rates 2 yr.	PE-Backed Company			Non PE Company			
Growth Asset	346	0.61	1.6	1432	0.78	1.98	-0.17
Growth Revenue	319	0.95	5.16	1296	1.06	4.93	-0.11
Growth Employment	486	0.14	0.57	1995	0.17	0.59	-0.03
Growth CAPEX	246	2.47	7.68	1121	2.65	7.38	-0.18
Growth Contribution	256	2.37	14.58	1177	2.12	12.47	0.25
Growth Leverage	300	0.04	0.4	1331	0.02	0.34	0.02
Change rates 2 yr.							
Change Asset	348	32.2	187.17	1440	23.21	185.16	8.99
Change Revenue	319	15.7	89.49	1296	16.12	79.03	-0.42
Change Employment	298	36.87	220.52	1149	20.07	202.2	16.80
Change Ln(Asset)	346	0.31	0.53	1432	0.41	0.8	-0.10*
Change Ln(Revenue)	319	0.17	0.96	1296	0.28	0.88	-0.12*
Change Ln(Employment)	298	0.03	0.52	1149	0.08	0.46	-0.05
Change CAPEX	246	0.09	0.33	1121	0.11	0.33	-0.02
Change Contribution	256	0.05	0.22	1177	0.04	0.18	0.01
Change Leverage	300	0.01	0.34	1335	-0.02	0.32	0.03

This Table reports the summary statistics of the companies at 2007 across treated (PE-backed companies) and untreated firms (non PE companies), as well as the mean difference across the two groups when we perform the matching based on leverage, ROA, size and industry. Panel A reports the one-year trend in the characteristics. Panel B reports the two-year trend of the characteristics. More information in the variable definition is available in the Appendix.

Table B.3.: Robustness: alternative matching model

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS										
$PEFirm_i \cdot Crisis$	0.057*** (0.012)	0.054*** (0.012)	-0.004 (0.014)	-0.002 (0.014)	0.051 (0.033)	0.056* (0.032)	0.025*** (0.007)	0.021*** (0.007)	-0.011* (0.006)	-0.011* (0.006)
$Firm F.E.$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Year F.E.$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Firm Controls_{it}$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Observations$	18426	17217	21427	18463	21367	18412	19804	17490	21027	18134
$Clusters$	3117	2796	3370	2834	3370	2834	3332	2810	3364	2832
$Adj.R^2$	0.151	0.155	0.002	0.025	0.042	0.049	0.048	0.061	0.002	0.007

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS										
$PEFirm_i \cdot Crisis$	2.848 (3.168)	-0.035 (2.733)	0.025 (0.025)	0.006 (0.026)	0.007 (0.050)	-0.019 (0.049)	-0.015 (0.026)	-0.006 (0.027)	0.000 (0.008)	-0.001 (0.008)
$Firm F.E.$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Year F.E.$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Firm Controls_{it}$	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
$Observations$	17892	16592	19111	17814	20633	18211	18538	16854	21097	18297
$Clusters$	3006	2706	3168	2832	3354	2834	3104	2709	3370	2834
$Adj.R^2$	0.006	0.012	0.128	0.140	0.350	0.436	0.003	0.012	0.006	0.028

This Table reports a robustness to the main results of the paper, where we estimate the standard difference-in-difference fixed effect model on investment and funding variables using an alternative matching sample (Panel A). Similarly, we also do the same with performance measures (Panel B). In particular, we match companies using also leverage on top of size, ROA and industry. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In panel (A), we have the following outcomes. In columns (1) and (2) the outcome is CAPEX scaled by asset; in columns (3) and (4) is total leverage; in columns (5) and (6) is the natural log of total debt; in columns (7) and (8) is the net equity contribution over asset; in columns (9) and (10) is working capital over asset. In panel (B) instead we have the following outcomes. In columns (1) and (2) we have firm EBITDA scaled by total employment; in columns (3) and (4) we have EBITDA scaled by revenue; in columns (5) and (6) we have the natural logarithm of revenue; in columns (7) and (8) we have the same logarithm for employment; in columns (9) and (10) we have ROA, which is computed as profit over asset. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table B.4.: Robustness: no MBO

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	<i>Capex/Asset</i>		<i>Debt/Asset</i>		<i>Ln(Debt)</i>		<i>NetContr./Asset</i>		<i>Wkcap/Asset</i>	
<i>PEFirm_i · Crisis</i>	0.048*** (0.015)	0.046*** (0.014)	0.010 (0.017)	0.016 (0.018)	0.071** (0.036)	0.099*** (0.034)	0.024*** (0.009)	0.021** (0.009)	-0.007 (0.006)	-0.008 (0.007)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	14130	13131	17000	14178	16824	14086	15711	13406	16731	14032
<i>Clusters</i>	2488	2205	2773	2253	2764	2249	2724	2220	2780	2262
<i>Adj.R²</i>	0.140	0.146	0.002	0.012	0.030	0.039	0.034	0.053	0.002	0.010

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	<i>EBITDA/Emp.</i>		<i>EBITDA/Rev.</i>		<i>Ln(Rev.)</i>		<i>Ln(Emp.)</i>		<i>ROA</i>	
<i>PEFirm_i · Crisis</i>	0.043 (0.058)	0.001 (0.057)	0.029 (0.034)	0.012 (0.035)	1.182 (4.481)	-0.648 (3.628)	0.020 (0.033)	0.037 (0.033)	-0.005 (0.010)	-0.002 (0.010)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	16488	14219	15023	13863	14212	12970	14819	13186	17174	14334
<i>Clusters</i>	2767	2265	2569	2261	2449	2161	2547	2166	2792	2266
<i>Adj.R²</i>	0.328	0.406	0.128	0.136	0.005	0.016	0.001	0.008	0.006	0.045

This Table reports a robustness to the main result of the paper, where we estimate the standard difference-in-difference fixed effect model on investment and funding variables dropping the PE-backed companies whose deal is identified as a management buyout (MBO) and the corresponding matched companies. Similarly, we also do the same with performance measures (Panel B). Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In columns (1) and (2) the outcome is CAPEX scaled by asset; in columns (3) and (4) is total leverage; in columns (5) and (6) is the natural log of total debt; in columns (7) and (8) is the net equity contribution over asset; in columns (9) and (10) is working capital over asset. In panel (B) instead we have the following outcomes. In columns (1) and (2) we have firm EBITDA scaled by total employment; in columns (3) and (4) we have EBITDA scaled by revenue; in columns (5) and (6) we have the natural logarithm of revenue; in columns (7) and (8) we have the same logarithm for employment; in columns (9) and (10) we have ROA, which is computed as profit over asset. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table B.5.: Robustness: only 2007-2008

OLS	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	<i>Capex/Asset</i>	<i>Capex/Asset</i>	<i>NetContr./Asset</i>	<i>NetContr./Asset</i>	<i>Ln(Debt)</i>	<i>Ln(Debt)</i>	<i>Debt/Asset</i>	<i>Debt/Asset</i>	<i>Wkcap/Asset</i>	<i>Wkcap/Asset</i>	<i>Wkcap/Asset</i>	<i>ROA</i>
<i>PEFirm_i · Crisis</i>	0.063*** (0.020)	0.068*** (0.020)	0.037*** (0.012)	0.033*** (0.012)	0.055*** (0.022)	0.051*** (0.022)	0.003 (0.009)	0.006 (0.009)	-0.005 (0.005)	-0.006 (0.005)	-0.001 (0.012)	0.003 (0.011)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	6181	5812	6852	5868	7252	6037	7309	6059	7254	6035	7569	6225
<i>Clusters</i>	3220	2958	3616	3021	3767	3105	3789	3114	3855	3166	3941	3206
<i>Adj.R²</i>	0.333	0.337	0.116	0.191	0.041	0.059	0.008	0.008	0.001	0.005	0.010	0.104

This Table reports a robustness to the main result of the paper, where we estimate the standard difference-in-difference fixed effect model on various outcomes using only data from 2007 and 2008. This correspond to the last year before the crisis and the first one in the crisis. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In columns (1) and (2) the outcome is CAPEX scaled by asset; in columns (3) and (4) is net equity contribution over asset; in columns (5) and (6) is the natural log of total debt; in columns (7) and (8) is total leverage; in columns (9) and (10) is working capital over asset; (11) and (12) is ROA. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table B.6.: Robustness: only no exit companies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
OLS	<i>Capex/Asset</i>	<i>Capex/Asset</i>	<i>NetContr./Asset</i>	<i>NetContr./Asset</i>	<i>Ln(Debt)</i>	<i>Ln(Debt)</i>	<i>Debt/Asset</i>	<i>Debt/Asset</i>	<i>Wkcap/Asset</i>	<i>Wkcap/Asset</i>	<i>ROA</i>	<i>ROA</i>
<i>PEFirm_i · Crisis</i>	0.046*** (0.013)	0.043*** (0.013)	0.023*** (0.008)	0.021*** (0.008)	0.033 (0.037)	0.031 (0.034)	0.000 (0.015)	0.001 (0.015)	-0.012* (0.006)	-0.012** (0.006)	-0.002 (0.008)	-0.002 (0.008)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	14103	13158	15275	13390	16295	14035	16422	14103	16185	13920	16459	14140
<i>Clusters</i>	2396	2134	2557	2148	2559	2150	2564	2152	2558	2149	2564	2152
<i>Adj. R²</i>	0.145	0.149	0.046	0.060	0.029	0.038	0.008	0.024	0.002	0.004	0.006	0.030

This Table reports a robustness to the main result of the paper, where we estimate the standard difference-in-difference fixed effect model on various outcomes using only data only for groups of matched firms where no company is identified as leaving the data by 2011 (survivorship bias free). Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level, using all the data available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In columns (1) and (2) the outcome is CAPEX scaled by asset; in columns (3) and (4) is net equity contribution over asset; in columns (5) and (6) is the natural log of total debt; in columns (7) and (8) is total leverage; in columns (9) and (10) is working capital over asset; (11) and (12) is ROA. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table B.7.: Robustness: Industry time-varying effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	<i>Capecx/Asset</i>		<i>Debt/Asset</i>		<i>Ln(Debt)</i>		<i>NetContr./Asset</i>		<i>Wkcap/Asset</i>	
<i>PEFirm_i · Crisis</i>	0.051*** (0.012)	0.049*** (0.011)	0.005 (0.014)	0.015 (0.014)	0.055* (0.032)	0.062** (0.030)	0.023*** (0.007)	0.018*** (0.007)	-0.010* (0.005)	-0.009* (0.005)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Ind * Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	20315	18874	24403	20453	24192	20342	22449	19260	24172	20327
<i>Clusters</i>	3525	3121	3914	3189	3903	3184	3842	3142	3926	3201
<i>Adj.R²</i>	0.144	0.149	0.006	0.019	0.035	0.044	0.036	0.053	0.009	0.014

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
OLS	<i>EBITDA/Emp.</i>		<i>EBITDA/Rev.</i>		<i>Ln(Rev.)</i>		<i>Ln(Emp.)</i>		<i>ROA</i>	
<i>PEFirm_i · Crisis</i>	1.023 (3.235)	-0.340 (2.806)	0.006 (0.024)	-0.007 (0.025)	0.003 (0.048)	0.031 (0.046)	-0.015 (0.026)	-0.005 (0.026)	-0.004 (0.008)	-0.001 (0.008)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Ind * Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Observations</i>	20200	18476	21582	19944	23654	20463	21024	18773	24554	20601
<i>Clusters</i>	3439	3034	3639	3201	3902	3205	3570	3041	3941	3206
<i>Adj.R²</i>	0.017	0.023	0.142	0.150	0.343	0.433	0.009	0.017	0.008	0.035

This Table reports a robustness to the main results of the paper, where we estimate the standard difference-in-difference fixed effect model on various outcomes adding set of fixed effects generated as the product of industry (two digit SIC) and the post dummy. Similarly, we also do the same with performance measures (Panel B). Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007. Data are at firm-year level, where we use every data point available between 2004 and 2011. Odd columns contain the baseline regression where instead even columns augment the baseline model with a set of firm level controls measured before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). In columns (1) and (2) the outcome is CAPEX scaled by asset; in columns (3) and (4) is total leverage; in columns (5) and (6) is the natural log of total debt; in columns (7) and (8) is the net equity contribution over asset; in columns (9) and (10) is working capital over asset. In panel (B) instead we have the following outcomes. In columns (1) and (2) we have firm EBITDA scaled by total employment; in columns (3) and (4) we have EBITDA scaled by revenue; in columns (5) and (6) we have the natural logarithm of revenue; in columns (7) and (8) we have the same logarithm for employment; in columns (9) and (10) we have ROA, which is computed as profit over asset. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table B.8.: Exit across more vs. less equity injection

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Marginal Eff. Logit	$1\{M\&A\}$	$1\{M\&A\}$	$1\{M\&A\ No\ Distr.\}$	$1\{M\&A\ No\ Distr.\}$	$1\{Out\ Business\}$	$1\{Out\ Business\}$	$1\{Bad\ Exit\}$	$1\{Bad\ Exit\}$
$PEFirm_i$	0.387*** (0.014)	0.120** (0.048)	0.381*** (0.017)	0.083** (0.040)	0.181*** (0.052)	0.170** (0.083)	0.144** (0.066)	0.144* (0.081)
$PEFirm_i \cdot 1\{Top25\%Eq\}$	-0.141* (0.082)	-0.066 (0.049)	-0.103 (0.080)	-0.032 (0.031)	-0.399*** (0.113)	-0.202* (0.112)	-0.310** (0.155)	-0.206* (0.117)
$1\{Top25\%Eq\}$	0.083** (0.033)	0.033** (0.017)	0.077** (0.036)	0.023* (0.013)	0.107** (0.046)	0.111*** (0.035)	0.055 (0.064)	0.062 (0.060)
$Firm\ Controls_{it}$		Y		Y		Y		Y
$Observations$	2907	2907	2907	2907	2665	2665	2640	2640
$Clusters$	40	40	40	40	33	33	32	32

This Table reports the marginal value (at the mean) of a conditional logit model where we study the effect of being a PE-backed company on various outcomes across companies that receive more or less equity contributions in the post crisis period. The conditioning in the model is given by two digit industry. Odd columns do not have any additional controls while even columns have firm level controls at 2007. In Columns (1) and (2) the outcomes is a dummy if the company was the target of an M&A activity in the post-crisis period; in Columns (3) and (4) is instead a dummy if the company was still target of an M&A activity and the company does not exit from the data in the same time frame; in columns (5) and (6) is the dummy equal to one if the company exit the data set in the post period; lastly in columns (7) and (8) is a dummy if the company exit the data and it reported some financial difficulties before the exit. The dummy (1Top25% Eq) identifies those companies that are in the top 25% of overall contribution in the post period. See the Appendix and the paper for more info on the variables. Standard errors are clustered at industry level. The number of observations correspond to the number of observations that were effectively used by the logit in the estimation, which are observations that present some within-industry (condition) variation in the outcomes. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table B.9.: Activity across more vs. less equity injection

OLS	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
		<i>Capex/Asset</i>		<i>Ln(Empl.)</i>	<i>Ln(Rev.)</i>		<i>EBITDA/Rev.</i>			<i>ROA</i>
<i>PEFirm_i · Crisis</i>	0.0347*** (0.012)	0.0346*** (0.012)	-0.026 (0.030)	-0.023 (0.029)	-0.003 (0.050)	0.037 (0.049)	0.016 (0.030)	-0.001 (0.030)	-0.013* (0.008)	-0.003 (0.007)
<i>PEFirm_i · 1{Top25%Eq}</i>	-0.0133 (0.014)	-0.002 (0.015)	0.079** (0.032)	0.102*** (0.036)	0.204*** (0.059)	0.045 (0.059)	-0.068*** (0.026)	-0.066** (0.029)	0.025*** (0.009)	-0.026** (0.011)
<i>PEFirm_i · 1{Top25%Eq} · Crisis</i>	0.081** (0.037)	0.068* (0.035)	0.073 (0.068)	0.106 (0.071)	0.120 (0.139)	0.065 (0.131)	-0.037 (0.0500)	-0.038 (0.0540)	0.0486** (0.022)	0.025 (0.023)
<i>Firm F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Year F.E.</i>	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
<i>Firm Controls_{it}</i>		Y		Y	Y	Y		Y		Y
<i>Observations</i>	20,159	18,741	20,624	18,471	23,235	20,107	21,194	19,599	24,037	20,233
<i>Clusters</i>	3,468	3,072	3,442	2,951	3,754	3,096	3,500	3,092	3,775	3,096
<i>Adj. R²</i>	0.142	0.147	0.004	0.012	0.330	0.423	0.131	0.137	0.007	0.036

This Table reports an heterogeneity of the main results across companies that receive more or less equity contribution in the post-period. In particular, I estimate the standard difference-in-difference fixed effect model on various outcomes. Every specification contains a set of firm and year fixed effects. The main parameter of interest is the interaction between the post dummy and a dummy identifying firms that are PE-backed at 2007 and the financial constraint dummy. Data are at firm-year level, where we use every data point available before the crisis and interacted with the post dummy. These variables are firm size (log of revenue), growth revenue, cash flow over asset, ROA, leverage (long term debt over asset). The dummy (1Top25%Eq) identifies those companies that are in the top 25% of overall contribution in the post period. More information on the variables are available in the Appendix. Standard errors are always clustered at firm level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

C. Appendix Chapter 3

C.1. A Simple model of loan supply

Borrowing from Khwaja & Mian (2008) model of bank lending with costly external finance (Stein, 1998), we present a stylized discrete that helps gaining intuition on the identification challenges of our analysis.

Assume that bank b can only raise funds using either short-term funding $S_{b,t}$ or bonds $B_{b,t}$.¹ Short-term funding bears no interest and can be considered risk-less as it is backed by a bank's holding of sovereign bonds which can always be liquidated for a fraction of their face value. That is $S_{b,t} \leq \gamma_t G_{b,t-1}$, where $G_{b,t-1}$ is the nominal value of bank b 's sovereign portfolio at the end of period $t - 1$, which bank b takes as given in period t . The parameter $\gamma_t \leq 1$ is the liquidation value of sovereign securities capturing current and future market value of government debt. Alternatively, banks can raise funds issuing bonds ($B_{b,t}$), which have a constant marginal cost of $\alpha_B B_{b,t}$ ($\alpha_B > 0$).

Two investment opportunities are available. Banks b can either lend to firm j ($L_{bj,t}$) or buy sovereign securities. We assume that the marginal return on a loan $\hat{\theta}_{j,t} - \alpha_L L_{ib,t}$ has two components as in Khwaja & Mian (2008). The first component is time-varying and firm-specific, $\hat{\theta}_{j,t} = \theta_{j,t} + \theta_t$. This captures both time-varying firm investment opportunities

¹We could generalize the model to include costless deposits without change its message.

$(\theta_{j,t})$ and more aggregate factors affecting the returns of all borrowers in the economy (θ_t) . The second component $(\alpha_L L_{bj,t})$ determines the diminishing returns. In particular, returns decreases at rate α_L in the size of the loan. Sovereign bonds $(G_{b,t})$ pay no interest and, abstracting from reasons that induce banks to invest in sovereign securities, we assume a constant and exogenous bank demand $G_{bt} = G_{bt-1} = G_b$.²

The equilibrium quantity of loans is determined by equating the marginal cost of loans to their marginal return, and by imposing the bank accounting equality $L_{bj,t} + \bar{G}_b = S_{b,t} + B_{b,t}$. In particular, at time t , we have that:

$$L_{bj,t} = \frac{\theta_{j,t} + \theta_t - \alpha_B G_b (1 - \gamma_t)}{\alpha_B + \alpha_L}$$

Without loss of generality, we set $\gamma_1 = \gamma$ and $\gamma_2 = \gamma - \tau$, where $0 \leq \tau \leq \gamma$ is a shock to the market value of sovereign held in bank's portfolio. Then, the change in loans between period 1 and period 2 is equal to:

$$\Delta L_{bj} = \frac{1}{\alpha_B + \alpha_L} (\Delta \theta_j + \Delta \theta + \alpha_B G_b \tau) \quad (\text{C.1})$$

In the end, this equation describes as, at the optimum, the change in loan for bank b to firm j will depend to firm demand, economy-wide shocks and bank's balance sheet factors.

C.2. Sovereign Holdings and Banks: A Cross-Country Comparison

Focusing on Italy, we are able to exploit the Italian Credit Register data and therefore we can provide a very detailed and precise account of the effects of sovereign holdings during period of macroeconomic distress. However, this choice may generate concerns regarding

²In the empirical application, we will relax this assumption by allowing the demand of sovereign in the pre-period to be a function of bank's characteristics.

the external validity of our results. To address these concerns, we show that both the characteristics of the banking sector and its exposure to sovereign risk are not substantially different other Western countries. Consider that also regulation is relatively homogeneous across other developed countries, in particular within Europe, this analysis can help us to generalize our results outside the specific case of Italy.

We collect data on annual balance sheet and income statement information of banks operating in a large number of countries from Bankscope.³ We focus on the sub-sample of banks that are active in Europe and United States, and compared them to Italian banks along a number of dimensions. The summary statistics per country are available in Table C.1. All data refer to fiscal year 2009, the last one available before the Greek bailout.

A cross-country comparison of capital structure across the financial institutions in our sample (Panel (b)) displays that Italian banks are comparable to other intermediaries in terms of capitalization and maturity of liabilities, especially intermediaries from other European countries. This is not surprising given the emphasis placed on capital requirements by the Basel regulation and the similarities in regulations across Western countries. With the exception of German and British banks, equity is typically around 12-14 percent of total liabilities. The average bank in Italy, France, Ireland and the Netherlands appears to have ample reliance on long-term funding, coherently with the business model prevailing in Europe. In particular, Italian banks have one of the largest share of assets funded by long-term debt, which should make the banking sector overall more resilience to a crisis episode. Furthermore, loans represent a slightly larger share of asset for Italian banks relative to other European countries, but this number is still similar to US banks. At the same time, we do not observe significant differences between Italian banks and the rest of the sample in terms of net income and impaired loans.

³Bankscope is a database managed by Bureau van Dijk Electronic Publishing (BvD). This data has been used in other works and its quality has been also recently scrutinized by Gennaioli *et al.* (2013).

Most importantly, Italian banks are not unique in terms of their exposure to sovereign securities. According to Bankscope, Italian banks hold around 14 percent of asset in government issued bonds. While this number is on the upper tail of the distribution of government bond holdings in our sample, in every country banks tend to be very exposed to sovereign securities. In fact, excluding Italy, the average portion of total assets invested in sovereign securities is about 8 percent for the banks in our sample, with Netherlands, Ireland and Greece having around 10 percent of their assets held in government debt.⁴


Bankscope does not provide the share of the total sovereign assets that are issued by the bank's own government. This information is instead provided by Merler & Pisani-Ferry (2012), who collected data on the share of total public debt that is held by domestic banks for a sample Western countries. In Figure C.1, we report the share of national debt held by resident financial institutions at the end of calendar years 2008 to 2011. According to these estimates, Italian banks hold 12 percent of outstanding national debt at the end of 2008, and gradually increased their holdings in the following three years, reaching 16.5 percent at the end of 2011. These percentages are not very different from other countries in the euro area - such as France, Ireland and Greece (before 2011). Italian banks hold more debt issued by their own government than intermediaries in the UK, the US and the Netherlands. On the contrary, banks from Belgium, Spain, Portugal and Germany are more exposed to national sovereign debt than Italian intermediaries. Indeed, if we rank these countries according to the percentage of national debt held by domestic banks, the Italian banking system positions itself in the middle of the distribution.

This descriptive evidence suggests that the Italian banking system shows no anomalies when compared to that of other developed countries, both in terms of profitability and capital structure. On average, Italian banks have a higher fraction of assets represented by sovereign

⁴For the USA, the variable total government assets is comparable to the variable Treasury and Agency Securities from the Flow of Funds data.

bonds, but this investment strategy is common to other Western countries as well.

Figure C.1.: Holding of Domestic Sovereign by Domestic Banks



This graph reports share of domestic public debt held by domestic financial institutions for a selected sample of countries, across different years. Source: data from Merler & Pisani-Ferry (2012).

C.3. Correlation Between Greek and Italian Sovereign Yields

We focus our analysis around the Greek bailout in April 2010. As discussed in the paper, this timing is justified by at least two considerations: (i) The Greek bailout led to an unexpected shock to sovereign holdings, which was orthogonal to the sovereign holding or lending behavior of Italian banks at that time; (ii) The Greek bailout led to an unprecedented change in sovereign markets. As we discuss in Section (3.2), many papers share this view and provide evidence in line with these statements.

Table C.1.: International Comparison of Banking Systems

Panel (a)											
	BE	DE	ES	FR	GR	IR	IT	NL	PT	UK	US
Tot. Sovereign Securities / Tot. Assets	0.18 (0.23)	0.02 (0.04)	0.05 (0.05)	0.06 (0.11)	0.09 (0.09)	0.1 (0.21)	0.14 (0.11)	0.09 (0.08)	0.04 (0.06)	0.07 (0.09)	0.1 (0.11)
Tot. Loans / Tot. Assets	0.42 (0.28)	0.54 (0.16)	0.63 (0.23)	0.58 (0.27)	0.67 (0.13)	0.42 (0.31)	0.66 (0.19)	0.57 (0.25)	0.55 (0.27)	0.45 (0.28)	0.63 (0.17)
Return on Avg. Asset	0.71 (2.65)	0.26 (0.39)	0.36 (0.97)	0.95 (3.54)	-0.13 (1.15)	-1.11 (4.07)	0.38 (0.79)	1.29 (4.09)	0.34 (1.56)	0.08 (7.38)	-0.09 (1.94)
Return on Avg. Equity	6.18 (13.53)	3.86 (4.8)	2.89 (16.9)	5.34 (15.48)	-5.52 (20.68)	-20.01 (61.86)	3.59 (6.78)	4.44 (14.5)	4.78 (10.96)	1.1 (32.9)	-2.33 (24.35)
Net Income / Tot. Assets	0.01 (0.03)	0 (0.01)	0 (0.02)	0.01 (0.05)	0 (0.01)	-0.01 (0.05)	0 (0.01)	0.01 (0.04)	0 (0.01)	-0.14 (2.58)	0 (0.03)
Impaired Loans / Tot. Assets	0.01 (0.01)	0.04 (0.03)	0.03 (0.02)	0.03 (0.04)	0.07 (0.06)	0.05 (0.07)	0.05 (0.04)	0.02 (0.01)	0.04 (0.05)	0.03 (0.06)	0.03 (0.03)

Panel (b)											
	BE	DE	ES	FR	GR	IR	IT	NL	PT	UK	US
Equity / Tot. Assets	0.18 (0.23)	0.08 (0.04)	0.12 (0.05)	0.14 (0.11)	0.12 (0.09)	0.08 (0.21)	0.12 (0.11)	0.16 (0.08)	0.14 (0.06)	0.22 (0.09)	0.11 (0.11)
Long-term Debt / Tot. Assets	0.11 (0.14)	0.04 (0.07)	0.2 (0.18)	0.17 (0.8)	0.06 (0.04)	0.21 (0.27)	0.3 (0.14)	.19 (0.21)	0.2 (0.17)	0.11 (0.19)	0.05 (0.1)
(Deposits + Short Term Debt) / Tot. Assets	0.69 (0.27)	0.86 (0.12)	0.78 (0.17)	0.72 (0.26)	0.8 (0.16)	0.61 (0.28)	0.56 (0.14)	0.76 (1.12)	0.64 (0.23)	0.69 (0.29)	0.82 (0.15)
Tier 1 Ratio	12.99 (3.21)	10.83 (3.57)	10.28 (5.67)	12.39 (5.69)	15.23 (10.27)	12.25 (6.47)	16.14 (8.68)	15.28 (11.24)	12.59 (6.94)	15.22 (9.53)	14.95 (10.27)

This table presents a sequence of descriptive statistics describing the financial system of 11 countries: Italy (IT), Germany (DE), France (FR), Belgium (BE), Netherlands (NL), United Kingdom (UK), United States (US), Greece (GR), Ireland (IR), Portugal (PT) and Spain (ES). **Panel (a)** refers to the asset structure, while **Panel (b)** focuses on the capital structure of the financial intermediaries operating in each country. All variables are measured at the end of fiscal year 2009. We present the average and standard deviation (in parenthesis) of the variables of interest across all banks operating in each country. Government bonds are reported at book value. Source: Bankscope.

We provide more evidence in line with these hypotheses by studying the correlation between the spread of Italy and Greece between 2006 and 2012 (Table C.2). In particular, we focus our analysis in four periods. First, until the beginning of 2009, the two spreads were very small, and they strongly co-moved together. This pattern changes when, in the fall of 2009, the news regarding the fiscal misconducts of Greece starts to become public. While the Greek spread started to increase, the Italian one remained stable. As a consequence, the correlation during this window is much lower and not significantly different from zero.⁵ However, in the second quarter of 2010, when the shock to sovereign risk was transmitted to other European countries, among which also Italy, the correlation increased again, remaining at this high level thereafter.

Table C.2.: Correlation between the Italian and Greek spread

Period	2006-2009Q1	2009Q2-2010Q1	2010Q2-2011Q1	2011Q2-2012
$\hat{\rho}$	0.984***	0.264	0.778***	0.641***
(p-value)	0.000	0.361	0.002	0.002

This table reports the the correlation between the Italian and Greek spread over German bonds, calculate for four different temporal windows. We use monthly yields on zero coupon bonds with 10 years maturity. All regressions include a constant. Data are publicly available on the ECB web page. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

C.4. Semi-parametric estimation of the bank lending channel

This Section provides more details on the semi-parametric tests presented in Figures C.2 and 3.2 of Section (3.4). We proceed in steps, beginning with the procedure to construct

⁵decreasing the window and keeping only the months since the summer, it actually becomes negative

Figure C.2.

First, we sort banks into “*High Sovereign*” group (our “treatment group”) and a “*Low Sovereign*” group (our “control group”) based their (conditional) holding of Italian sovereigns in the last quarter before the shock. To do so, we run a cross-sectional of $\text{Sovereigns}_{b,2010Q1}$ on a battery of bank-level characteristics and balance sheet variables,

$$\text{Sovereigns}_{b,2010Q1} = \phi_0 + \Gamma \cdot X_{b,2010Q1} + \epsilon_{b,2010Q1}$$

where $X_{b,2010Q1}$ are bank-specific controls measured at the end of the first quarter of 2010. We extract the estimated residuals of this regression and we classify a bank as “*High sovereign*” whenever $\hat{\epsilon}_{b,2010Q1}$

is above the median of the cross-sectional distribution of residuals, and “*Low sovereign*” otherwise. Sorting banks according to their *residual* sovereign holdings helps us to focus only on the cross-sectional variation of their exposure that is not imputable to different bank-specific characteristics. Second, we aggregate by quarter all corporate loans ($\ln(\text{Loans}_{bj,t})$) belonging to our sample that have been granted by banks classified as “*Low sovereign*”. We repeat this aggregation for loans granted by banks belonging to the “*High sovereign*” group. For graphical clarity, in Figure (C.2) we normalize each one of the two time series to zero in 2010:Q1 (the last quarter before the sovereign shock).

Before the shock, the (unconditional) credit supply of banks with higher holdings of sovereigns - if anything - was growing faster than that of banks with lower holdings. Immediately after the shock, we observe a sharp reversal of the lending trend for the group with high sovereign holdings, while banks with lower holdings kept a similar dynamic for the first three quarters of the post-shock period. This evidence excludes that our results could be driven by a failure of parallel trend, where more exposed banks were cutting credit more than less exposed banks also before the actual shock to sovereign markets.

Figure C.2.: The bank lending channel




Figure C.2 illustrates the bank lending channel semi-parametrically by comparing lending to firms from banks with high holdings of Italian sovereign bonds, the most exposed to the sovereign shock, and banks with lower holdings. See appendix C.4 for a detailed description of the procedure use to conduct the this figure.

Figure 3.2 offers a refinement of the non-parametric estimation presented in Figure C.2, by restricting our attention to the variation in bank credit (on the y-axis in Figure C.2) which is not explained by bank’s balance sheet characteristics. To do so, we run a loan-level regression of loans $\ln(\text{Loans}_{bj,t})$ on a set of bank-level characteristics, and relationship-specific characteristics:

$$\ln(\text{Loans}_{bj,t}) = \psi_0 + \Omega \cdot \bar{X}_{b,t} + \Lambda \cdot \bar{Z}_{bj,t} + \varepsilon_{bj,t}$$

where $\ln(\text{Loans}_{bj,t})$ is the natural logarithm of the value of outstanding loan issued bank b in favor of firm j in quarter t , and $\bar{X}_{b,t}$ and $\bar{Z}_{bj,t}$ are four-quarters moving average of our bank-specific and relationship-specific controls. We extract the residuals $(\hat{\varepsilon}_{jb,t})$ of this regression. Sorting banks into “High” and “Low Sovereign” as described above, we aggregate $\hat{\varepsilon}_{jb,t}$ into

two time series. As above, we normalize each time series such that they take value zero in 2010:Q1, and plot them over time in Figure 3.2. As discussed in the paper, in this case we find that the two groups present very similar patterns in lending before the decision but more exposed banks start experience a larger decline in lending in the quarters after the Greek bailout. More discussion on the interpretation of these results and on other tests is available in the Section (3.4) of the paper.

C.5. Estimates of effect of the sovereign crisis on aggregates bank credit via the lending channel

In the paper, we use our micro estimates to quantify the aggregate impact of the sovereign shock on lending. In particular, we want to assess what percentage of the drop in aggregate credit can be imputed to the transmission of the sovereign shock via the bank lending channel. Using the estimates of the linear model (3.5), for every firm j we compute the percentage change in bank credit which can be imputed to its lenders' exposure to sovereigns ($\hat{s}_j = \hat{\alpha}_1 \text{Sovereigns}_{j,2010Q1}^{AVE}$). This percentage change is then transformed in a (euro) amount by multiplying it for the overall lending of firm j in the pre-period. Lastly, we can then aggregate the overall drop and compare it to the overall level of lending in the pre-period.⁶ This simple analysis provides us with a proxy of the overall decline in lending explained by the shock to sovereign markets. Our estimates suggest that about 2% of the decline in aggregate lending to non-financial corporations can be imputed to the transmission of the sovereign crisis through the balance sheet of financial institutions.

⁶Overall, this is $\frac{\sum_j \hat{s}_j \cdot \text{Tot Loans}_{j,PRE}}{\sum_j \text{Tot Loans}_{j,PRE}}$.

C.6. Data selection and other information on data construction

In this section, we discuss the data construction process. Starting from the universe of all business credit relationships appearing in the Italian Credit Register, we classify firms into two groups. The first sub-sample includes a random sample of seventy percent of all borrowers which established, over the sample period we consider, credit relationships with only one lender. The second group includes every firms which established multiple, simultaneous lending relationships with several banks.

For each of these two sub-samples, we exclude a number of observations. We drop defaulted loans as well as new credit granted to borrowers who already have some other relations in default, as these positions may no longer reflect genuine demand and supply dynamics, but rather capture debt restructuring operations or some other agreement due to the default procedures. We drop observations for which we have no information about the lender. We exclude credit provided by special purpose vehicles, non-bank financial intermediaries, and branches of foreign banks.⁷ We drop observations referring to borrowers which operate in the financial and insurance sector, utilities and government-related industries.⁸ We eliminate firms with more than seven contemporaneous credit relationship, i.e. firms belonging to the top 5% of the distribution of lending relationships.⁹

Finally, for the intensive margin estimates, our sample includes only firms which appear in the Credit Register at least one quarter in our pre-period (2009:Q2-2010:Q1) and at least one quarter in the post-period (2010:Q2-2011:Q1). Instead, the sample used to estimate the

⁷We exclude all lending activity of Italian firms with branches of foreign banks for which no detailed balance sheet information is available. As estimated by Cingano *et al.* (2014), these lenders grant only a small share of total loans to Italian firms (about 6 percent).

⁸We exclude firms operating in the education sector and utilities because, for a majority of the cases, the government either runs them directly or indirectly subsidizes their activity.

⁹Our inspection of the data suggest that some of the credit relationships of firms with a high number of lending relationships do not reflect genuine credit relationships.

extensive margin effects of the bank lending channel includes also firms that disappear from the Credit Register after the Greek bailout. We end up with 538,348 unique firm-bank credit relationships (424,191 referring to multi-lender firms) and more than 4.5 million quarterly observations for the two-years period going from 2009:Q2 to 2011:Q1. Our database includes 539 different banks and 302,538 firms (188,381 with multiple lending relationships). This sample is fully representative of the universe of firms and banks operating in Italy.

Table C.3.: Distribution of Firms Across Industries

Industry	Whole Sample (%)	Low Sovereigns (%)	High Sovereigns (%)
Agriculture and Fishing	1	0.8	1.2
Mining and Quarrying	0.7	0.6	0.8
Manufacturing	50.6	47.7	54.5
Construction	8.7	9.2	8.2
Trade (retail & wholesale)	28.0	29.5	26.5
Transportation	3.5	3.8	3.1
Real Estate	0.1	7.3	4.2
Services	5.7	4.3	3.2
Communications	1.8	1.9	1.6

This table reports the distribution of firms across different macro industries. Macro industries are defined as broad aggregates of SIC codes. The sample includes firms which established multiple lending relationships, after the application of the filters described in Appendix C.6. The first column reports the industry composition of the whole sample. The second and third column report the industry composition within the subsamples of firms borrowing from banks with sovereigns exposure ($\text{Sovereigns}_{j,2010Q1}^{AVE}$) above the median and below the median, respectively. Source: Italian Credit Register, Bank of Italy.

C.7. Variables description

Relationship-specific variables

All relationship-specific variables come from the Italian Credit Register. The variables $\Delta\ln(\text{Loans})$, $\Delta\ln(\text{Credit Lines})$, and $\Delta\ln(\text{Tot Credit})$ measure the change in log average

loans, credit lines, and total credit granted by bank b to firm j between the pre-shock (2009:Q2-2010:Q1) and post-shock period (2010:Q2-2011:Q1); Cut Credit is a dummy variable equal one whenever a term loan granted by bank b to firm j in the pre-period was not renewed in the one-year period following the Greek bailout; Lenght Relationship_{2010Q1} indicates the length of the lending relationship (in quarters) between borrower j and bank b , which we measured as the number of quarters the relationship has been in place between 2006:Q1 and 2010:Q1; Share Relationship_{2010Q1} is the fraction of borrower j total bank credit provided by the lender b .

Bank-specific variables

All bank-specific variables come from the Bank of Italy Supervisory Records, and are measured at the end of 2010:Q1 (the quarter before the sovereign shock). These variables include the stock of Italian sovereigns over risk-weighted assets (Sovereigns_{2010Q1}), the stock of Italian sovereigns over Tier1 (Sovereigns_{2010Q1}^{TIER1}), the fraction of total sovereign portfolio invested in Italian government bonds (Sovereigns_{2010Q1} over Total Sovereigns_{2010Q1}), profitability (ROA_{2010Q1}), bank size (Size_{2010Q1}, the log of Risk Weighted Assets (RWA)), Tier1 ratio (Tier1_{2010Q1} over RWA), deposits ratio (Deposits_{2010Q1} over RWA), liquidity ratio (Liquidity_{2010Q1} over RWA), interbank market participation (Net Interbank Debt_{2010Q1} over RWA), quality of lending portfolio (Bad Loans_{2010Q1} over RWA), an indicator variable for cooperative banks (BCC), total stock sovereign securities over RWA (Total Sovereigns_{2010Q1}), total stock sovereign securities issued by GIPSI (GR,IR,PR,SP, and IT) over RWA (Total Sovereigns GIPSI_{2010Q1}), total stock of sovereign securities issued by GIPS (GIPSI less IT) over RWA (Total Sovereigns GIPS_{2010Q1}), and total stock of German sovereigns over RWA (Total Sovereigns DE_{2010Q1}); a dummy indicating Tier1 ratio < 10% (Low Capital Ratio_{2010Q1} over RWA).

Firm-specific variables

All firm-specific variables come from the CADS database. They are available for the subsample of firms present in both the Italian Credit Registry and in the CADS database. Fixed assets₂₀₀₉ (thousand euros), Revenues₂₀₀₉ (thousand euros), and Employees₂₀₀₉ (units) refer to fiscal year 2009. $1\{\Delta\text{Invest} < 0\}$ ($1\{\Delta\text{Empl} < 0\}$) is a dummy variable equal to

one if the firm decreased its fixed assets (employment) between the end of 2009 and the end of 2011; $\% \Delta \text{Invest}$ ($\% \Delta \text{Empl}$) is the log-difference of fixed asset (employees) of the firm between 2011 and 2009. Rating_{2010Q1} is the credit rating (Altman's z-score) of the firm at the end of 2010:Q1.

RZ Index is the Rajan & Zingales (1998) index of dependence on external finance. Following Rajan & Zingales (1998), we construct the RZ index for each industry (SIC 2-digits) as the median of $(\text{CapEx} - \text{Cash from Operations}) / \text{CapEx}$, using data from the firms in Compustat North America between 1980 and 2008. We manually matched the Italian Counterpart (ATECO code) to the corresponding SIC code. The matching algorithm is available under request.

$\text{Sovereigns Province}_{2010Q1}$ is the average exposure to Italian sovereigns at the end of 2010:Q1 of banks operating in the same province of firm j . At the time of our analysis, there were 110 provinces in Italy, which can be roughly compared to US counties. As pointed out by Guiso *et al.* (2013), provinces represent a proper boundaries of the local market for bank credit. Indeed, provinces have been historically used by Bank of Italy to decide opening of new branches, and by the antitrust authority to assess and regulate deposit market concentration.

C.8. Additional Figures

Figure C.3.: 10-year spread of European bonds over German bonds






Figure C.3 shows the dynamic of the spread between the yield of 10-year zero-coupon bonds of a selected sample of European countries and that of bonds issued by Germany. On the left-hand side y-axis: Italy (IT), France (FR), Netherlands (NL), Spain (ES). On the right-hand side y-axis: Greece (GR), Portugal (PT), Ireland (IR).

Figure C.4.: Investors pool: European Bank and Insurance CEO conference



(a) Threat to bank's share prices



(b) Expected performance of financial institutions

Figure C.5a and Figure C.5b display the survey results of two questions from the investor poll conducted during the European Bank and Insurance CEO conference organized in September 2010 in London by Bank of America Merrill Lynch. Source: investor poll conducted during the European Bank and Insurance CEO. Charts available at <http://ftalphaville.ft.com/2010/10/04/359726/european-bank-watch-past-present-and-future/>.

C.9. Additional Tables

Table C.4.: Banks Characteristics and Sovereign Holdings

	Below Median of Sovereigns _{2010Q1}	Above Median of Sovereigns _{2010Q1}	Difference Below-Above	Correlation with Sovereigns _{2010Q1}
ROA _{2010Q1}	0.00 (0.00)	0.00 (0.00)	-0.00 (0.00)	0.04 (0.36)
Size _{2010Q1}	6.19 (1.77)	5.03 (1.24)	1.16*** (0.13)	-0.26*** (0.00)
Tier1 _{2010Q1}	0.14 (0.12)	0.20 (0.16)	-0.06*** (0.01)	0.45*** (0.00)
Deposits _{2010Q1}	0.61 (0.32)	0.99 (0.51)	-0.37*** (0.04)	0.75*** (0.00)
Liquidity _{2010Q1}	0.00 (0.00)	0.01 (0.01)	-0.00*** (0.00)	0.26*** (0.00)
Net Interb. _{2010Q1}	-0.06 (0.33)	-0.11 (0.15)	0.05** (0.02)	-0.11*** (0.01)
Bad Loans _{2010Q1}	0.03 (0.03)	0.05 (0.04)	-0.010*** (0.00)	0.1052** (0.01)
BCC	0.61 (0.48)	0.88 (0.32)	-0.27*** (0.03)	0.09*** (0.04)

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table shows the relation between intermediaries' exposure to the sovereign crisis (fraction of bank's RWA invested in Italian sovereign securities, Sovereigns_{2010Q1}) and a host of bank-specific characteristics: profitability, size, capitalization, deposits ratio, liquidity ratio, interbank market participation, quality of lending portfolio, and status of the bank as a cooperative bank. All variables are measured at the end 2010:Q1. The first and second column report, respectively, the mean and standard deviation (in parenthesis) of bank's characteristics sorting bank into two groups: above and below the median exposure. The third column shows the difference between the first and the second column and the standard errors of a two-sample t-test of the equality of the means (in parenthesis). The fourth column shows the pairwise correlation between Sovereigns_{2010Q1} and banks characteristics and the p-value of this correlation (in parenthesis). *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.5.: Banks Characteristics and Credit Supply

	Below Median of $\Delta\ln(\text{Corp. Loans})$	Above Median of $\Delta\ln(\text{Corp. Loans})$	Difference Below-Above	Correlation with $\Delta\ln(\text{Corp. Loans})$
ROA _{2010Q1}	0.00 (0.00)	0.00 (0.00)	-0.00 (0.00)	-0.26*** (0.00)
Size _{2010Q1}	5.75 (1.69)	5.45 (1.56)	0.30** (0.14)	-0.09** (0.03)
Tier1 _{2010Q1}	0.17 (0.13)	0.18 (0.16)	-0.01 (0.01)	0.08* (0.06)
Deposits _{2010Q1}	0.75 (0.51)	0.85 (0.40)	-0.10*** (0.04)	0.00*** (0.88)
Liquidity _{2010Q1}	0.01 (0.00)	0.01 (0.01)	-0.00*** (0.00)	0.04 (0.32)
Net Interb. _{2010Q1}	-0.08 (0.30)	-0.09 (0.17)	0.01 (0.02)	-0.08* (0.06)
Bad Loans _{2010Q1}	0.04 (0.04)	0.04 (0.03)	0.00** (0.00)	-0.08* (0.06)
BCC	0.77 (0.42)	0.74 (0.44)	0.03** (0.04)	-0.11*** (0.01)

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table shows the relation between the percentage change in bank b corporate credit between the post- and pre-period ($\Delta\ln(\text{Corp. Loans})$), and a host of bank-specific characteristics: profitability, size, capitalization, deposits, liquidity, interbank market participation, bad loans, and status of the bank as a cooperative bank. All variables are measured at the end of first quarter of 2010. The first and second column report, respectively, the mean and standard deviation (in parenthesis) of bank's characteristics sorting banks into two groups: above and below the median exposure. The third column shows the difference between the first and the second column and the standard errors of a two-sample t-test of the equality of the means (in parenthesis). The fourth column shows the pairwise correlation between $\Delta\ln(\text{Corp. Loans})$ and banks characteristics and the p-value of this correlation (in parenthesis). *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.6.: Pre-Trending and Placebo Regressions

<i>Dep. Var :</i>	$\Delta \ln(\text{Loans})_{2010Q1-2009Q2}$ (1)
Sovereigns _{2010Q1}	0.028 (0.079)
Bank Controls _{2010Q1}	Y
Relationship Controls _{2010Q1}	Y
Firm FE	Y
Sample	Multiple
Cluster	Bank
Adj. R ²	0.417
Observations	249966

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table reports the results of the pre-trend test and of the placebo test described in Section (3.4). We run the empirical model of equation (3.2) to test for pre-trending. We investigate the correlation between holdings of sovereigns of bank b and the percentage change in loans to firm j in the pre-shock period. The change is measured as delta log loans from b to j between 2010:Q1 and 2009:Q2. The main independent variable is the exposure of lender j to Italian sovereigns in 2010:Q1 scaled by RWA in 2010:Q1 (Sovereigns_{2010Q1}). The regression includes a set of bank-specific and relationship-specific controls measured at the end of 2010:Q1 and firm fixed effects. The sample includes only firms with multiple lending relationships appearing in the Credit Register. Standard Errors are clustered at bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.7.: The Bank Lending Channel: Firms' Heterogeneity

<i>Dep. Var :</i>	$\Delta\ln(\text{Loans})$ (1)
Sovereigns _{2010Q1}	-0.321** (0.140)
Sovereigns _{2010Q1} x Small Firm ₂₀₀₉	0.008 (0.052)
Bank Controls _{2010Q1}	Y
Relationship Controls _{2010Q1}	Y
Firm FE	Y
Sample	CADS
Cluster	Bank
Adj. R ²	0.413
Observations	147,970

Sample: Firms with multiple lending relationships appearing in the CADS database

This table examines the heterogeneity of bank lending channel across different types of firms. It reports the estimates obtained from model (3.2) on the sample of firms with multiple lending relationships which also appear in the CADS database. We include in the baseline regression of Model (3.2) a set of interactions between bank b 's exposure to the shock and a host of bank-specific characteristics. The outcome variable in Columns (1) and (2) ($\Delta\ln(\text{Loans})$) is the log-difference in average loans granted by bank b to firm j between (2010:Q2-2011:Q1) and (2009:Q2-2010:Q1). The main independent variables is the stock of Italian sovereigns held by the lender at the end of 2010:Q1 scaled by RWA (Sovereigns_{2010Q1}). The interaction variables is a dummy equal to one if firm j 's revenues in 2009 are below the median across firms at the time (Small Firm₂₀₀₉). All regressions include a set of bank-specific and relationship-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at firm or bank level depending on the specification. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.8.: The Bank Lending Channel: Different Sovereign Holdings

<i>Dep. Var :</i>	$\Delta \ln(\text{Loans})$					
	(1)	(2)	(3)	(4)	(5)	(6)
Sovereigns _{2010Q1}	-0.285*** (0.105)					-0.273*** (0.098)
GIPS Sovereigns _{2010Q1}		-0.218 (0.261)				-0.129 (0.258)
GIPSI Sovereigns _{2010Q1}			-0.285*** (0.106)			
German Sovereigns _{2010Q1}				0.287 (0.501)		0.550 (0.563)
Total Sovereigns _{2010Q1}					-0.281*** (0.106)	
Bank Controls _{2010Q1}	Y	Y	Y	Y	Y	Y
Relationship Controls _{2010Q1}	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
Sample	Multiple	Multiple	Multiple	Multiple	Multiple	Multiple
Cluster	Bank	Bank	Bank	Bank	Bank	Bank
Adj. R ²	0.599	0.599	0.599	0.599	0.599	0.599
Observations	424191	424191	424191	424191	424191	424191

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table examines the bank lending channel using alternative measures of sovereign exposure. It reports the estimates obtained from model (3.2) on the sample of firms with multiple lending relationships appearing in the Credit Register. The outcome variable is the log-difference in average loans granted by bank b to firm j between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis ($\Delta \ln(\text{Loans})$). The main independent variables are different measures of bank's exposure to the sovereign crisis. In Column (1) we use the stock of Italian sovereigns over RWA, in Column (2) the stock of total GIPS sovereigns (Greece, Ireland, Portugal, and Spain) over RWA, in Column (3) the stock of GIPSI sovereigns (GIPS plus Italy) over RWA, in Column (4) the stock of German sovereigns over RWA, and in Column (5) we use total the stock of sovereign securities over RWA. Column (6) presents an horse-race among the alternative measures. All proxies of exposure are measured at the end of 2010:Q1. All regressions include a set of bank-specific and relationship-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at firm or bank level depending on the specification. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.9.: Transmission Mechanism of The Bank Lending Channel: Foreign and BCC

<i>Dep. Var :</i>	$\Delta \ln(\text{Loans})$	
	(1)	(2)
Sovereigns _{2010Q1}	-0.284** (0.114)	-0.316** (0.141)
<i>Foreign Bank</i>	-0.025 (0.045)	
Sovereigns _{2010Q1} x <i>Foreign Bank</i>	0.152 (0.309)	
BCC		0.098*** (0.032)
Sovereigns _{2010Q1} x <i>BCC</i>		0.070 (0.127)
Sold Sovereigns _{Post}		
Bank Controls _{2010Q1}	Y	Y
Relationship Controls _{2010Q1}	Y	Y
Firm FE	Y	Y
Sample	Multiple	Multiple
Cluster	Bank	Bank
Adj. R ²	0.599	0.599
Observations	424191	424191

Sample: Firms with multiple lending relationships appearing in the Credit Register

This table investigates the channels of transmission of the sovereign shock through banks' balance sheet. It reports the estimates obtained from model (3.2) on the sample of firms with multiple lending relationships appearing in the Credit Register. We interact exposure to the sovereign shock with a set of bank characteristics which are proxies for alternative balance sheet channels of transmission. The outcome variable is the log-difference in average loans granted by bank b to firm j between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis ($\Delta \ln(\text{Loans})$). The main independent variable is the exposure of the lender to Italian sovereigns (Sovereigns_{2010Q1}), and its interactions with different proxies of the transmission channels. The interaction variables include: Foreign Bank (a dummy equal one if the bank is a subsidiary of a foreign bank), BCC (a dummy equal one if the bank is a Cooperative bank). All regressions include a set of bank-specific and relationship-specific controls measured at the end 2010:Q1. Every specification contains firm fixed effects. Standard Errors are clustered at bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.10.: Fixed Effects and Demand-Side Shocks

<i>Dep. Var.</i>	Estimate Firm FE ($\hat{\rho}_j$)				
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln(\text{Revenues}_{2009-2011})$	0.075*** (0.005)			0.064*** (0.005)	0.071*** (0.005)
$\Delta \ln(\text{Assets}_{2009-2011})$		0.064*** (0.003)		0.056*** (0.003)	0.054*** (0.003)
Rating ₂₀₁₁			-0.052*** (0.007)	-0.059*** (0.007)	-0.075*** (0.007)
$\ln(\text{Revenues}_{2009})$					0.006*** (0.001)
Constant	0.045*** (0.002)	0.050*** (0.002)	0.057*** (0.002)	0.048*** (0.002)	0.049*** (0.002)
Sample	CADS	CADS	CADS	CADS	CADS
Standard Errors	Robust	Robust	Robust	Robust	Robust
Adj. R ²	0.015	0.017	0.006	0.030	0.031
Observations	35165	35165	35165	35165	35165

Sample: Firms with multiple lending relationships appearing in the CADS database

This table investigates the correlation between the fixed effects estimated by model (3.2) on the sample of firms with multiple lending relationships appearing in the CADS database with proxies of firms demand, investment opportunities, and creditworthiness. The right-hand side variables in include revenues' growth between the fiscal years 2009 and 2011 ($\Delta \ln(\text{Revenues}_{2009-2011})$), and assets' growth between 2009 and 2011 ($\Delta \ln(\text{Assets}_{2009-2011})$), credit rating at the end of fiscal year 2009 (Rating₂₀₀₉, low values indicate higher credit rating), the natural logarithm of revenues in 2009 ($\ln(\text{Revenues}_{2009})$). All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.11.: Real Effect: Sovereign Exposure on Investment and Employment - Pre-Trending

Dep. Var :	1(Δ Invest > 0)		% Δ Invest		1{ Δ Empl > 0}		% Δ Empl	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sovereigns ^{AVE} _{2010Q1}	-0.045 (0.090)	0.050 (0.097)	-0.257* (0.131)	0.191 (0.137)	-0.054 (0.136)	-0.002 (0.136)	0.058 (0.069)	-0.164 (0.125)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉		-0.136 (0.123)		-0.312 (0.197)		-0.032 (0.178)		0.080 (0.129)
Small Firm ₂₀₀₉		-0.024** (0.011)		0.014 (0.022)		-0.179*** (0.019)		-0.093*** (0.013)
Estimated Firm FE ($\hat{\rho}_j$)	0.073*** (0.007)	0.071*** (0.006)	0.253*** (0.010)	0.120*** (0.012)	0.081*** (0.008)	0.069*** (0.009)	0.071*** (0.007)	0.065*** (0.007)
Bank & Relationship Contr. ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Sample	CADS	CADS	CADS	CADS	CADS	CADS	CADS	CADS
Cluster	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank
Adj. R ²	0.030	0.031	0.030	0.030	0.034	0.059	0.040	0.054
Observations	32961	32961	32961	32961	17175	17175	17175	17175

This table examines the effects of the sovereign crisis on employment transmitted via the lending channel. It reports the estimates obtained from model (3.5) where the dependent variables are two alternative proxies of firm's investments and employment. In Columns (1) and (2) we use a dummy which is equal to 1 if the firm's invested or increased its labor force between 2007 and 2009. In Columns (3) and (4) we look at the growth rate in investments and employment over the same period. The main independent variable is the weighted average of the exposure to Italian sovereigns of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). Columns (1) and (3) show the baseline effect on investments and employment. Columns (2) and (4) investigate the heterogeneity of the effect on investments across firms of different size. All regressions include a set of weighted averaged bank-specific and relationship-specific controls measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1, and we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated in the baseline regression of the bank lending channel (equation (3.2)). Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.12.: The Firm Borrowing Channel: Adding controls

<i>Dep. Var.</i>	(1)	(2)	(3)	(4)	(5)
	$\Delta \ln(\text{Tot Loans})$				
Sovereigns ^{AVE} _{2010Q1}	-0.135** (0.066)	-0.121*** (0.033)	0.064 (0.094)	-0.061 (0.047)	-0.134 (0.129)
Sovereigns ^{AVE} _{2010Q1} x Sovereigns Prov _{2010Q1}			-1.766* (0.930)		0.636 (1.308)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉				-0.106** (0.051)	0.296 (0.235)
Sovereigns ^{AVE} _{2010Q1} x Sovereigns Prov _{2010Q1}					-4.140** (1.709)
Small Firm ₂₀₀₉				0.011** (0.005)	-0.021 (0.022)
Small Firm ₂₀₀₉ x Sovereigns Prov _{2010Q1}					0.341* (0.177)
Estimated Firm FE ($\hat{\rho}_j$)		0.779*** (0.005)	0.779*** (0.005)	0.779*** (0.005)	0.779*** (0.005)
Bank & Relationship Contr ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y
Firm Controls	Y	Y	Y	Y	Y
Sample	CADS	CADS	CADS	CADS	CADS
Cluster	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank
Adj. R ²	0.058	0.644	0.644	0.644	0.644
Observations	34729	34729	34729	34729	34729

This table examines presents a robustness to Table (3.8), where we augment the baseline model with firm controls. It reports the estimates obtained from model (3.5) on the sample of firms with multiple lending relationships appearing in the Credit Register and on the subs-sample appearing in the CADS database. The outcome variable is the log-difference in average total bank loans granted to firm j between after (2010:Q2-2011:Q1) and before (2009:Q2-2010:Q1) the onset of the sovereign crisis ($\Delta \ln(\text{Tot Loans})$). The main independent variable is the weighted average of the exposure to Italian sovereign scaled by RWA of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). We augment all the columns with a fixed set of controls which are firm size, as measured by log-revenue, profitability and leverage, as total bank loan to asset. All these variables are measured at the end of 2009. In Column (1) we control for unobserved demand-side shocks (β_j) estimated from the baseline regression of the bank lending channel (equation (3.2)). Column (2) presents the results of the same econometric model estimated in Column (1), but without controlling for demand-side shocks (β_j). Columns (3) and (4) replicate the regressions in Column (1) and (2) on the CADS sample. Columns (5)-(9) investigate the heterogeneity of the effect. The interaction variables include: a dummy equal to one if firm j 's revenue in 2009 are below the median across firms at the time (Small Firm₂₀₀₉) using in sample observations only, and the average exposure of banks operating in the same province of the firm (Sovereigns Province_{2010Q1}). All regressions include a set of weighted averaged bank-specific and relationship-specific controls measured at the end of 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1. Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.13.: Real Effects: Adding Firm-Controls

<i>Dep. Var :</i>	1(Δ Invest < 0) (1)	% Δ Invest (2)	(3)	% Δ Invest (4)	1{ Δ Empl < 0} (5)	(6)	% Δ Empl (7)	(8)
Sovereigns ^{AVE} _{2010Q1}	0.104 (0.079)	-0.093 (0.083)	-0.148 (0.094)	0.032 (0.103)	-0.066 (0.149)	0.029 (0.159)	0.016 (0.069)	-0.019 (0.082)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉		0.350*** (0.100)		-0.322** (0.134)		-0.198 (0.202)		0.072 (0.098)
Small Firm ₂₀₀₉		-0.022 (0.013)		0.039*** (0.014)		0.036 (0.022)		-0.012 (0.011)
Estimated Firm FE ($\hat{\rho}_j$)	-0.202*** (0.008)	-0.202*** (0.008)	0.251*** (0.009)	0.252*** (0.009)	-0.152*** (0.010)	-0.152*** (0.010)	0.093*** (0.007)	0.092*** (0.007)
Bank & Relationship Contr. ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm Controls	Y	Y	Y	Y	Y	Y	Y	Y
Sample	CADS	CADS	CADS	CADS	CADS	CADS	CADS	CADS
Cluster	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank
Adj. R ²	0.032	0.032	0.033	0.033	0.037	0.043	0.044	0.044
Observations	34771	34771	34771	34771	16671	16671	16671	16671

This Table presents a robustness to Table (3.9), where we augment the baseline model with firm level covariates measured at 2009. It reports the estimates obtained from model (3.5) where the dependent variables are two alternative proxies of firm's investments and employment. In Columns (1) and (2) we use a dummy which is equal to 1 if the firm's invested or increased its labor force between 2009 and 2011. All these variables are measured at the end of 2009. In Columns (3) and (4) we look at the growth rate in investments and employment over the same period. The main independent variable is the weighted average of the exposure to Italian sovereigns of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). We augment all the columns with a fixed set of controls which are firm size, as measured by log-revenue, profitability and leverage, as total bank loan to asset. Columns (1) and (3) show the baseline effect on investments and employment. Columns (2) and (4) investigate the heterogeneity of the effect on investments across firms of different size. All regressions include a set of weighted averaged bank-specific and relationship-specific controls measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1, and we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated in the baseline regression of the bank lending channel (equation (3.2)). Standard Errors are clustered at lead bank level. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.14.: Real Effects across External Dependence: Adding Controls

<i>Dep. Var.</i>	1(Δ Invest < 0)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1(Δ Invest < 0)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sovereigns ^{AVE} _{2010Q1}	0.102 (0.078)	-0.097 (0.088)	-0.175* (0.091)	0.038 (0.103)	-0.015 (0.156)	0.026 (0.171)	-0.030 (0.071)	-0.050 (0.084)
Sovereigns ^{AVE} _{2010Q1} x RZ Index ₂₀₀₉	-0.004 (0.019)	-0.006 (0.028)	-0.023 (0.019)	-0.001 (0.023)	0.054** (0.024)	-0.007 (0.043)	-0.041* (0.021)	-0.025 (0.022)
Sovereigns ^{AVE} _{2010Q1} x Small Firm ₂₀₀₉		0.354*** (0.113)		-0.403*** (0.130)		-0.084 (0.218)		0.041 (0.097)
Sovereigns ^{AVE} _{2010Q1} x RZ Index ₂₀₀₉ x Small Firm ₂₀₀₉		0.001 (0.040)		-0.057 (0.050)		0.102* (0.052)		-0.028 (0.041)
RZ Index ₂₀₀₉ x Small Firm ₂₀₀₉		-0.001 (0.003)		0.007 (0.005)		-0.006 (0.006)		0.001 (0.003)
Small Firm ₂₀₀₉		-0.025* (0.014)		0.049*** (0.013)		0.029 (0.023)		-0.010 (0.013)
Estimated Firm FE ($\hat{\rho}_j$)	-0.202*** (0.008)	-0.203*** (0.008)	0.255*** (0.011)	0.255*** (0.011)	0.131*** (0.010)	-0.150*** (0.010)	0.092*** (0.007)	0.092*** (0.007)
Bank & Relationship Contr. ^{AVE} _{2010Q1}	Y	Y	Y	Y	Y	Y	Y	Y
Province & Industry FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm Controls	Y	Y	Y	Y	Y	Y	Y	Y
Sample	CADS	CADS	CADS	CADS	CADS	CADS	CADS	CADS
Cluster	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank	L. Bank
Adj. R ²	0.031	0.031	0.032	0.033	0.042	0.043	0.044	0.045
Observations	34354	34354	34354	34354	16444	16444	16444	16444

This Table presents a robustness to Table (3.10), where we augment the baseline model with firm level covariates measured at 2009. It reports the estimates obtained from model (3.5) on the sample of firms with multiple lending relationships appearing in the CADS database. The dependent variables are two proxies of firm investments and employment. The main independent variable is the weighted average of the exposure to Italian sovereigns of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). We interact with the firm level shock with a proxy of firm's size (Small Firm₂₀₀₉) and dependence on external finance (RZ Index). We augment all the columns with a fixed set of controls which are firm size, as measured by log-revenue, profitability and leverage, as total bank loan to asset. All regressions include a set of weighted averaged bank-specific and relationship-specific controls are measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1, and we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated in the baseline regression of the bank lending channel (equation (3.2)). Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.

Table C.15.: Real Effects: Labor Efficiency

<i>Dep. Var :</i>	$1\{\Delta\text{Empl} < 0\}$ (1)	$\%\Delta\text{Empl}$ (2)
Sovereigns ^{AVE} _{2010Q1}	0.462 (1.002)	-1.040* (0.600)
Sovereigns ^{AVE} _{2010Q1} x APL ₂₀₀₉	-0.051 (0.084)	0.084* (0.048)
APL ₂₀₀₉	0.065*** (0.009)	0.054*** (0.005)
Estimated Firm FE ($\hat{\rho}_j$)	-0.125*** (0.010)	0.086*** (0.007)
Bank Controls ^{AVE} _{2010Q1}	Y	Y
Relationship Controls ^{AVE} _{2010Q1}	Y	Y
Province & Industry FE	Y	Y
Sample	CADS	CADS
Cluster	Lead	Lead
	Bank	Bank
Adj. R ²	0.049	0.079
Observations	16670	16670

This table examines the effects of the sovereign crisis on employment transmitted via the lending channel across firms characterized by different level of labor productivity. It reports the estimates obtained from model (3.5) on the sample of firms with multiple lending relationships appearing in the CADS database. Our dependent variables are two alternative proxies of firm employment. In Columns (1) and (2), we use a dummy which is equal to 1 if the firm's employees in 2011 were higher than firm's employees in 2009. In Columns (3) and (4) we look at the growth rate in employment between 2009 and 2007. The main independent variable is the weighted average of the exposure to Italian sovereigns of firm j 's lenders (Sovereigns^{AVE}_{2010Q1}). We interact with the firm level shock with the firm's labor efficiency in before the shock (Average Product of Labor 2009Q4) measured as the log ratio of value added to employment. All regressions include a set of weighted averaged bank-specific and relationship-specific controls are measured at the end 2010:Q1. All regressions include province fixed effects and industry fixed effects measured at the end of 2010:Q1, and we control for unobserved demand-side shocks ($\hat{\rho}_j$) estimated in the baseline regression of the bank lending channel (equation (3.2)). Standard Errors are clustered at lead bank level. All regressions include a constant. *** denotes significance at the 1% level, ** at the 5%, and * at the 10%.