Essays in Energy Economics and Entrepreneurial Finance

Essays in Energy Economics and Entrepreneurial Finance

A dissertation presented

by

Sabrina Tugendrajch Howell

to

The Department of Political Economy and Government (Economics Track)

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the Subject of

Political Economy and Government

Harvard University

Cambridge, Massachusetts

May 2015

Dissertation Advisors:
Professors Josh Lerner & David S. Scharfstein

Author:
Sabrina Tugendrajch Howell

Essays in Energy Economics and Entrepreneurial Finance

Abstract

When does government intervention successfully correct perceived market failures? What effects do such interventions have on firm decisions? These questions are especially vital to the energy sector, which features large negative externalities, volatile commodity prices, and intensive regulation. My dissertation examines energy policies in three otherwise disparate contexts: a U.S. national research and development (R&D) subsidy intended to expedite clean energy technology deployment; a U.S. state-level oil price risk management policy targeting highway paving firms; and a Chinese fuel economy standard aimed at reducing oil consumption and hastening technology adoption among Chinese automakers. Each analysis evaluates the public policy and uses it to glean insight into firm financial constraints and innovation investment. Together, the three chapters contribute to the literatures on entrepreneurial finance, corporate risk management, innovation, and industrial policy.

Motivating the first paper is the observation that governments regularly subsidize new ventures to spur innovation, often in the form of R&D grants. I examine the effects of such grants in the first large-sample, quasi-experimental evaluation of R&D subsidies. I implement a regression discontinuity design using data on ranked applicants to the Small Business Innovation Research grant program at the U.S. Department of Energy. An award approximately doubles the probability that a firm receives subsequent venture capital and has large, positive impacts on patenting and the likelihood of achieving revenue. The effects are stronger for more financially constrained firms. In the second part of the paper, I use a signal extraction model to identify why grants lead to future funding. The evidence is inconsistent with a certification effect, where the award contains information about firm quality. Instead, the grant money itself is valuable, possibly because it funds proof-of-concept work that reduces investor uncertainty about the technology.

The second chapter examines how firms manage oil price risk when oil is an important input cost. Despite a rich theoretical literature, there is little empirical evidence about risk management heterogeneity across firm types. I evaluate a policy that shifts oil price risk in highway procurement

from the private sector to the government, reducing the cost of hedging to zero. In a triple-differences design using data from Kansas and Iowa, I show that firms value hedging oil price risk between the auction and commencement of work. Consistent with the prediction that hedging is more valuable for financially constrained firms, I find higher risk premiums in private vis-à-vis public firms and in smaller vis-à-vis larger firms. I also find that family ownership and a lack of diversification are associated with higher risk premiums. Competition is highly imperfect in this industry. Monopoly power in product markets, together with market frictions in derivative hedging, may limit the pass-through of risk to financial markets, and thus prevent efficient allocation of risk.

I turn to China - a very different economic setting - in the third chapter. Technology absorption is critical to emerging market growth. To study this process I exploit fuel economy standards, which compel automakers to either acquire fuel efficiency technology or reduce vehicle quality. With novel, unique data on the Chinese auto market between 1999 and 2012, I evaluate the effect of China's 2009 fuel economy standards on firms' vehicle characteristic choices. Through differences-in-differences and triple differences designs, I show that Chinese firms responded to the new policy by manufacturing less powerful, cheaper, and lighter vehicles. Foreign firms manufacturing for the Chinese market, conversely, continued on their prior path. For example, domestic firms reduced model torque and price by 12% and 13% of their respective means relative to foreign firms. Private Chinese firms outperformed state-owned firms and were less affected by the standards, but Chinese firms in joint ventures with foreign firms suffered the largest negative effect regardless of ownership. My evidence suggests that fuel economy standards and joint venture mandates - both *intended* to increase technology transfer - have instead retarded Chinese firms' advancement up the automotive manufacturing quality ladder.

Contents

Acknowledgments

I thank my dissertation committee - Joseph Aldy, Raj Chetty, Josh Lerner, and David Scharfstein - for their support and unerring advice. David: for showing me that finance is important, policy-relevant, and exciting; for endlessly editing my introduction; for reminding me that good writing is central to the discipline; and for guiding me through the job market. Josh: for pushing me to pursue my JMP idea - directly inspired by his work - and for providing extremely useful kernels of professional wisdom. Joe: for sharing his passion for energy and the environment with me from my first year at Harvard; for guiding my research; and for being a supportive friend and mentor. Raj: for making me believe in myself through his interest in my work; for providing critical, granular input into my estimation and interpretation; and for welcoming me belatedly despite a full slate.

I also thank my "fifth advisor," Ramana Nanda, for his steadfast mentorship and guidance. I am grateful to the National Science Foundation for a Graduate Research Fellowship, and to Brenda Piquet and Nicole Tateosian for their patient help.

For help with Chapter 1, I thank Gary Chamberlain, Lee Fleming, Jeff Furman, Ed Glaeser, Shane Greenstein, Sam Hanson, Adam Jaffe, Larry Katz, Ramana Nanda, Ariel Pakes, Jeremy Stein, and Adi Sunderam, as well as the Harvard Finance, Labor, and IO lunch groups. I am indebted to Ken Alston, Jeff Dowd, Matthew Dunne, Carla Frisch, Carl Hebron, Tina Kaarsberg, Teryn Norris, and Jamie Vernon, all currently or formerly at the Department of Energy. Partial funding for this project is from the Harvard Lab for Economic Applications and Policy.

For help with Chapter 2, I am grateful to Gary Gorton, Ryan Kellogg, Jonathan Levin, Greg Lewis, Ariel Pakes, Robert Pindyck, Adriano Rampini, Tom Wollman, anonymous referees, and the Harvard Finance and IO lunch communities. I thank the Iowa and Kansas Departments of Transportation, in particular Roger Bierbaum, Steven Belzung, Kevin Martin, Abe Rezayazdi and LouAnn Hughes. For help with Chapter 3, I thank Henry Lee, Lu Mai, Ariel Pakes, Martin Rotemberg, Anthony Saich, Wang Qing, and Lilei Xu. I thank the China State Council Development Research Center for its support and data access. Funding for this project is from the Belfer Center at the Harvard Kennedy School.

Finally, I thank my parents for encouraging me to pursue a PhD and providing crucial emotional and financial aid. I am enormously grateful to my husband, Rob Berschinski, for editing sentences he didn't understand, for trusting that long distance was worth the effort, and for supporting me when the outcome was far from clear.

To my Dad, who taught me to read graphs on the Subway.

1 Financing Constraints as Barriers to Innovation: Evidence from R&D Grants to Energy Startups

1.1 Introduction

Governments regularly subsidize research and development (R&D) in new ventures.[1] One rationale for such subsidies is that the private sector does not internalize the social benefits of innovation.[2] Another is that financial frictions lead to underinvestment in early-stage R&D.[3] Yet critics contend that government R&D subsidies are ineffective because they crowd out private investment or allocate funds inefficiently (Wallsten 2000, Lerner 2009). Despite opposing theoretical arguments, we have little empirical evidence about the effectiveness of R&D subsidies. There is also little work on whether financing constraints are first-order barriers to innovative startups.

In the first quasi-experimental, large-sample evaluation of R&D grants to private firms, I show that the grants have statistically significant and economically large effects on measures of financial, innovative, and commercial success. I then provide evidence that the grants benefit firms because they ease financing constraints. Finally, I explore the specific mechanism through which grants alleviate financial frictions.

The study is based on a new, proprietary dataset of applications to the U.S. Department of Energy's (DOE) Small Business Innovation Research (SBIR) program. The data include 7,436 small high-tech firms and over $884 million in awards from 1983 to 2013. Awards typically fund testing or proof-of-concept of a new energy technology. DOE officials rank firms within competitions, and I exploit these ranks in a sharp regression discontinuity design that compares firms immediately around the award cutoff.

I show that a Phase 1 grant of $150,000 approximately doubles a firm's chance of subsequently receiving venture capital (VC) investment, increasing the long term probability by 9 percentage

[1] In addition to the federal SBIR, many U.S. states have similar programs. Parallels overseas include the UK's Innovation Investment Fund, China's Innofund, Israel's Chief Scientist incubator program, Germany's Mikromezzaninfonds and ZIM, Finland's Tekes, Russia's Skolkovo Foundation, and Chile's InnovaChile.

[2] For evidence that startups contribute disproportionately to economic growth, see Akcigit and Kerr (2011), Haltiwanger et al. (2013), and Audretsch, Keilbach and Lehmann (2006).

[3] Grants might increase investment if given to startups that face excessively costly external finance. Frictions that can lead to such costly finance and thwart privately profitable investment opportunities include information asymmetry, asset intangibility, and incomplete contracting (Akerlof 1970, Holmstrom 1989).

points from 10% to 19%. Within two years of the grant, the effect is 7 percentage points. These results imply that on average the grants do not crowd out private capital, and instead transform some awardees into privately profitable investment opportunities. I provide evidence that the effect does not reflect reallocation of capital from losers to winners within competitions.

Firms that tend to be more financially constrained receive the most benefit. First, the effect is strongest for the youngest firms, and I show that it declines with firm age. Second, the effect is larger and more robust for immature technologies, like geothermal and wave energy, which are likely the riskiest investments. Third, the effect is stronger in times when external finance is harder to access. Employing clean energy industry Tobin's Q as a proxy for investment opportunities, I find that when Q is lower, the grant effect is larger. The effect is also negatively correlated with total U.S. venture deal flow, a proxy for VC availability.

Beyond the consequences for future private financing, I also show that the Phase 1 grants influence real outcomes. A grant leads a firm to produce about 1.5 extra patents within three years, increasing the average from one patent to 2.5 patents. It is associated with greater technology commercialization; increasing the probability a firm achieves revenue from 52% to 63%. While grants do not affect firm survival, they do increase exit probability via IPO or acquisition. Like the results on future financing, these results are stronger for more constrained firms. Together, the VC, patent, and revenue results show that the early stage grants enable new technologies to go forward.

While Phase 1 grants have large, positive effects on financing and real outcomes, I find that later stage grants are ineffective. Phase 1 winners can apply for Phase 2 grants of $1 million, disbursed about two years after the Phase 1 award. Entrepreneurs' revealed preference indicates that they perceive relatively low benefits to the much larger grant. For example, among firms that get VC within two years of Phase 1, 55% opt *not to apply* to Phase 2. Regression discontinuity estimates using Phase 2 applicants yield tiny or negative effects on VC finance, and small positive effects on patents and patent citations. These findings suggest that - perhaps due to very high discount rates - Phase 2 is often not worthwhile for high-quality firms, and has little benefit among firms that do apply. However, the right to apply to Phase 2 may have option value in case of a bad post-Phase 1 state.

What mechanism might explain the early stage grants' impact on future financing? In a simple signal extraction model, I capture how the grant might influence investor beliefs to ease financing constraints. One mechanism is *certification*; the government's decision conveys positive information to venture capitalists about the firm's technology. Alternatively, the money itself may switch the net present value (NPV) of investing in the startup from negative to positive (a *funding* effect). The NPV may initially be negative because of financing frictions like information asymmetry and agency problems, or because the technology risk at such an early stage is too high. The funding effect has two possible channels. First, the grant could allow the entrepreneur to retain more equity; in the counterfactual, an investor might require such a large stake that entrepreneurial incentives could not be maintained. Second, the startup might use the grant to prove the viability of its technology. This *prototyping* channel could reduce investor uncertainty.

I use my empirical evidence to identify which mechanism most likely drives the grants' effect on VC. The certification test reveals an important fact about the grant program: officials seem unable to identify high-quality firms. The test asks whether applicant ranks are correlated with outcomes, conditional on award status. Rational investors should view the grant as a positive signal only if ranks are relevant to market outcomes. This is because a firm's rank within a competition, which the investor *does not* observe, maps directly to whether the firm wins, which the investor *does* observe.[4] Empirically, the ranks are uninformative about all outcomes that I observe. For example, conditional on winning, more highly ranked firms are not more likely to receive VC; the same is true conditional on losing. To the rational investor, the grant signal is pure noise. Thus certification is unlikely to explain the large jump at the discontinuity.

Instead, the evidence best supports the funding effect and is most consistent with the prototyping channel, where the grant enables proof-of-concept work that the firm cannot otherwise finance. Startups with a successful prototype can demonstrate to investors that their technology works as advertised. After Phase 1 prototyping, there is enough information for the private market to take over. At this later stage firms either prefer VC to government funds, or apply to Phase 2, in which case the larger grant crowds out private investment. In Section 1.5, I discuss the mechanisms and describe in detail how I tell them apart.

Seattle-based Oscilla Power, a wave energy startup, illustrates the prototyping hypothesis. Founded in 2009, Oscilla won its first DOE SBIR Phase 1 grant in May 2011 to conduct "testing activities to ensure the reliability of both the core power generation module as well as the mooring lines."[5] In an interview, CEO Rahul Shendure said that this proof-of-concept work helped Oscilla

[4] The decision about a competition's award cutoff is exogenous to the ranking process. Officials producing the ranks do not determine the cutoff and are uncertain about the number of awards.

[5] From the application abstract.

raise a \$1.6 million Series A round from venture investors in November 2011. "Phase 1 is not providing a material amount of money in terms of the investor's dollar," he said, "instead it's about running experiments, demonstrating that the idea you have works, or doesn't work." In his opinion, the grants "have no certification effect," a view shared by nearly all thirty of the venture capitalists I interviewed.

For startups like Oscilla, early stage grants appear to relieve a critical liquidity constraint on R&D investment. Such startups are an important middle ground between universities and national labs, which must undertake basic R&D, and large firms, which have the market-oriented discipline to efficiently conduct later stage, applied R&D (Griliches 1998; Aghion, Dewatripont and Stein 2008).[6] My results suggest that for early stage applied R&D in capital-intensive sectors, there may be space for a hybrid model that involves both government funding and startups.

Severe financing constraints at the "seed" stage, however, contrast with evidence from Phase 2 that later stage ("Series A") projects may not suffer from the same frictions. This study's main policy implications, therefore, are that the SBIR program - and potentially similar programs - could achieve better outcomes through reallocating money (1) from larger, later stage grants (Phase 2) to more numerous small, early-stage grants (Phase 1); and (2) from older firms and regular winners to younger firms and first-time applicants. I do not address the complex questions of optimal program size or whether government should be subsidizing private R&D.

This paper builds on the costly external finance literature, which finds evidence of financing constraints but has focused on large public companies and rarely studied R&D. I provide a novel and plausibly exogenous cash flow shock that identifies a causal relationship between financing constraints and investment responses.[7] In addition, this study relates to the literature on barriers

[6]Aghion, Dewatripont and Stein (2008) present a model describing the challenge of locating basic R&D in private firms. They use scientists' demand for research control rights to demonstrate why much early-stage research must be located in academia.

[7]Financing constraints are a central issue in corporate finance. A debate beginning with Fazzari, Hubbard and Petersen (1988) and Kaplan and Zingales (1997) has for the most part found investment to be sensitive to cash flow shocks (e.g. Lamont 1997, Rauh 2006, Whited and Wu 2006). However, it is difficult to establish that financial constraints *cause* this sensitivity, and there is little evidence on small or private firms (see Hall 2010). Zwick and Mahon (2014) use a tax policy change to find evidence of financing constraints that are more severe for smaller firms. Barrot (2014) shows that financial constraints can impeded entry and competition in the context of trade credit supply. Studies of intangible asset investment under imperfect capital markets include Himmelberg and Petersen (1994), Aghion et al. (2012), Bond, Harhoff and Van Reenen (2005), Brown and Petersen (2009), Hall (1992), Carpenter and Petersen (2002), and Czarnitzki and Hottenrott (2011). See Hall (2010) for discussion of the gaps in the literature on startups and R&D.

to entrepreneur entry (Chatterji and Seamans 2012, Hochberg, Ljungqvist, and Lu 2007, Black and Strahan 2002). Finally, in establishing a causal effect of grants on outcomes, I contribute to the literature evaluating R&D subsidy programs. This literature has not reached consensus. For example, while Lerner (2000) finds that SBIR awardees in the first few years of the program grew more than a matched sample, Wallsten (2000) finds that the program crowded out private funding, also using mid-1980s data. Most studies examine non-U.S. R&D programs and come to disparate conclusions, such as Lach (2002), Takalo, Tanayama and Toivanen (2013), and Almus and Czarnitzki (2003).[8]

Much of this literature focuses not on financing constraints as a rationale for R&D subsidies but rather on the extent to which the private sector fails to internalize knowledge spillovers and other positive externalities (Arrow 1962). I do not quantify this public good aspect to the grants, but two findings provide indirect evidence. I find no Phase 1 effect on patent citations, suggesting that proof-of-concept work may not lead to large knowledge spillovers. The grants do, however, seem to help internalize positive externalities from clean energy (Nordhaus 2013). I find the strongest effect in the cleanest sub-sectors, such as solar and wind, and the weakest effects in conventional sub-sectors like natural gas and coal.

The paper is organized as follows. In Section 1.2, I explain the DOE SBIR setting and the applicant data. Section 1.3 describes the regression discontinuity design and establishes its validity in my context. Section 1.4 contains the empirical results on financing and real outcomes. Section 1.5 uses a signal extraction model to frame how grants might affect investor decisions, and evaluates the model's hypotheses in light of the empirical evidence. I test the robustness of the empirical results in Section 1.6. Section 1.7 conducts a return calculation. Section 1.8 concludes. All Appendices here and in subsequent chapters are available as online supplemental material.[9]

---

[8]Evaluations of R&D subsidies mainly address European programs, with quite disparate findings, including Czarnitzki and Lopes-Bento (2012), Serrano-Velarde (2008), Busom (2000), Duguet (2003), González et al. (2005), González and Pazó (2008), Blasio, Fantino and Pellegrini (2014), and Henningsen et al. (2014). In the U.S., Nemet and Kammen (2007) find little evidence of crowding out in federal energy R&D, but Popp and Newell (2009) do. Link and Scott (2010) use SBIR Phase 2 awardee survey data to analyze the likelihood of commercialization. To my knowledge, only the working papers by Zhao and Ziedonis (2013) and Bronzini and Iachini (2011) use data on applicants to R&D incentive programs. The former evaluates a Michigan loan program (N=104), and the latter grants to large firms in Northern Italy (N=171). Both programs have private cost sharing, which SBIR does not. Other researchers have used RD to evaluate grants to university researchers, such as Jacob and Lefgren (2011) and Benavente et al. (2012).

[9]See http://dash.harvard.edu

1.2 The Setting: Context & Data Sources

In this section, I first discuss DOE's SBIR program and my applicant dataset. Section 1.2.2 summarizes the private finance data and matching. Section 1.2.3 describes data on patenting, revenue, and survival.

1.2.1 The SBIR Program at the Department of Energy

In the U.S., grants are a significant funding source for high-tech entrepreneurs.[10] The largest single program is the SBIR grant program, which disburses around $2.2 billion each year. Congress first authorized the SBIR program in 1982 to strengthen the U.S. high technology sector and support small firms. Today, 11 federal agencies must allocate 2.7% of their extramural R&D budgets to the SBIR program; the required set-aside will increase to 3.2% in 2017 and beyond. Though important in its own right, the SBIR program is also representative of the many targeted subsidy programs for high-tech new ventures at the state level and around the world.

Akin to staged VC funding, the SBIR program has two "Phases." Phase 1 grants fund proof-of-concept work intended to last nine months. Awardees are given the $150,000 in a lump sum (the amount has increased stepwise from $50,000 in 1983). DOE does not monitor how they use the money, but firms must demonstrate progress on their Phase 1 projects to apply for $1 million Phase 2 grants. Phase 2 funds more extensive or later stage demonstrations, and the money is awarded in two lump sums over two years.[11]

There is no required private cost sharing in the SBIR program. Also, the government neither takes equity in the firm nor assumes IP rights. Eligible applicants are for-profit, U.S.-based, and at least 51% American-owned firms with fewer than 500 employees. Although the SBIR grant is non-dilutive, it is not costless. In interviews, 30 VC investors and employees at ten startups described the application and reporting process as onerous. Applying for an SBIR grant can require two months of 1-2 employees working full time.

[10] A rough estimate suggests that federal and state R&D grants to high-tech new ventures were about $3 billion in 2013, compared total VC investments in the U.S. that year of $29.6 billion (NVCA 2014).

[11] Phase 2 grants are analyzed in Appendix 1E. Please find all appendices here: http://scholar.harvard.edu/showell/home. Phase 3 is commercialization of the technology. It is ineligible for SBIR funds except when agencies are purchasing the technology, which does not occur at DOE but is common at the Department of Defense.

Each year, DOE officials in technology-specific program offices (e.g. "Solar") develop a series of competitions. A firm applies to a relevant competition, proposing a project that fits within its scope. Examples of competitions include "Solar Powered Water Desalination," and "Improved Recovery Effectiveness In Tar Sands Reservoirs." My empirical strategy compares firms within competitions.

Three external experts from National Labs and universities review applications according to three criteria: 1) Strength of the scientific/technical approach; 2) Ability to carry out the project in a cost effective manner; and 3) Commercialization impact (Oliver 2012). Program officials rank applicants within each competition based on the written expert reviews and their own discretion. These ranks and losing applicant identities are strictly and indefinitely non-public information.[12] Program officials submit ordered lists to an independent, separate DOE SBIR office. The cutoff within each competition is unknown to the program officer when she produces the rankings. The SBIR office determines the competition's number of awards. This cutoff varies across competitions, so one competition may have one awardee while another has four; the average is 1.7. To the best of my knowledge the cutoff is arbitrary.[13] Figure 1.1 shows that there are no obvious differences among program offices in the average number of awards.[14]

In this study, I use complete data from the two largest applied offices at DOE, Fossil Energy (FE) and Energy Efficiency & Renewable Energy (EERE), which has eight technology-based program offices.[15] Together, EERE and FE awarded $884 million (2012 dollars) in SBIR grants over the course of my data from 1983 to 2013. Appendix 1D Figure 4 shows all applicants by office and

[12]It is only in my capacity as an unpaid DOE employee that I am able to use this data. Throughout the paper, specific references to companies will only include winners.

[13]The number of awards is determined by topic and program budget constraints, recent funding history, office commitments to projects such as large National Laboratory grants, and the overall number of ranked applicants the central SBIR office receives (the number of applicants deemed "fundable"). My understanding of the exogeneity of the cutoff to the ranking comes from conversations with stakeholders in the DOE SBIR program, and from historical email records containing rank submissions. I cannot predict the number of awards in a competition using any observable covariates, and fluctuation in the number of awards does not differ systematically by program office, technology topic, or time.

[14]The average number of applicants per competition by program office is in Appendix 1D Figure 1. Appendix 1D Figures 2 and 3 show the number of awards per office and per competition over time.

[15]Besides EERE and FE, the other offices are: Basic Energy Science; Nuclear Energy; Environmental Management; and Electricity Delivery & Energy Reliability. Within EERE, the eight program offices are: Solar Energy Technology, Biomass Program; Fuel Cell Technologies; Geothermal Technology; Wind & Hydropower Technology; Vehicle Technology; Building Technology and Advanced Manufacturing.

award status. The data include, for each applicant, the company name and address, funded status, grant amount, and award notice date. I have ranking information only since 1995, so my estimation starts in that year.

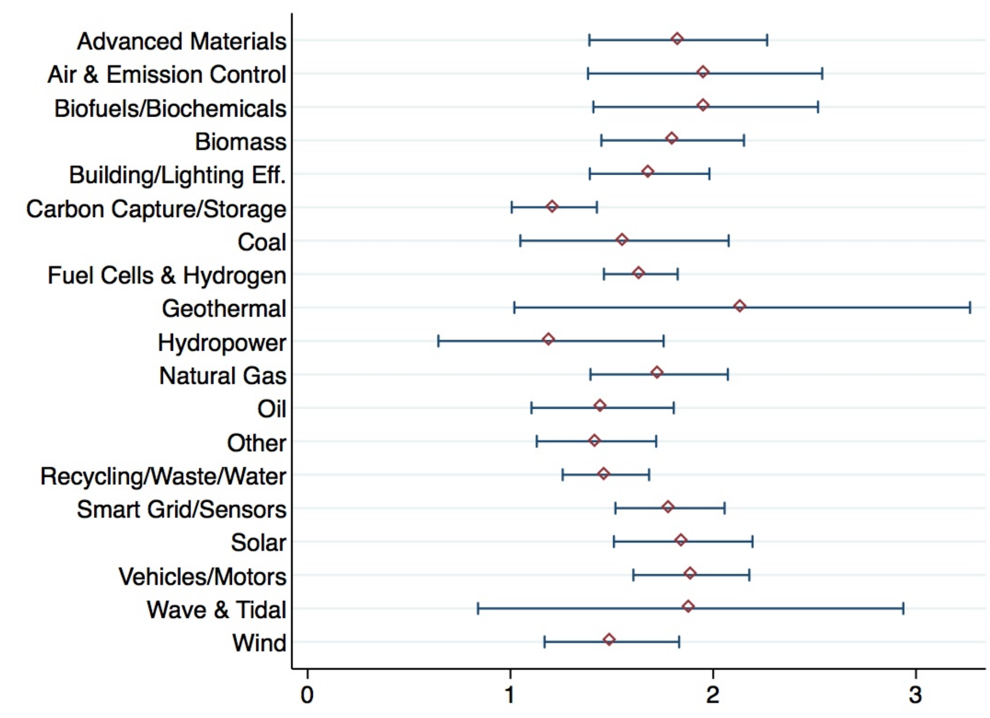

Figure 1.1: Average Number of Awards per Competition by Program Office
Note: This figure shows that within competitions, the average number of Phase 1 awards does not vary systematically across program offices (topics). It includes all DOE EERE & FE competitions from 1995 are included. Capped lines indicate 95% confidence intervals. For the number of awards per office and per competition over time, see Appendix 1D Figures 1-3. N=863.

Table 1.1 contains summary statistics about the applications and competitions, and Table 1.2 shows all variables used in estimation. Each competition has on average 9.8 applicants, with a standard deviation of eight. Of the 7,436 applicant firms, 71% applied only once, and a further 14% applied twice. Within my data, seven companies each submitted more than 50 applications. For discussion of "SBIR mills" and the grant effect by the number of awards, see Appendix 1F. Despite the presence of "SBIR mills," startups dominate the applicant pool; the firm median age is six years, and many firms are less than a year old.[16] Consistent with this fact, scholars have

[16]Among the 23 solar firms that have ever had an IPO, nine appear in my data; SBIR winners include Sunpower, First Solar, and Evergreen Solar (Cleantech Group i3). Although there is no strict definition of "startup," they must be young, small, and have location-unconstrained growth potential. This is why restaurants, plumbers, and other local small businesses are not startups.

used SBIR winners as representative samples of high-tech entrepreneurial firms. For example, Hsu (2006) uses a sample of SBIR awardees as a counterfactual for VC-funded startups. Gans and Stern (2003) use survey data on 71 SBIR grantees to test whether capital constraints or appropriability problems explain different performance across sectors.

Table 1.1: Summary Statistics of DOE SBIR Applicants

1983-2013	
# Phase 1 Applications	14,522
# Unique Phase 1 Applicant Firms	7,419
# Competitions	1,633
1995-2013	
# Phase 1 Applications	9,659
# Unique Phase 1 Applicant Firms	4,545
# Phase 1 Applications with ranking data used in RD	5,021
# Phase 1 Competitions used in RD[1]	428
Average # Phase 1 Applicants per Competition	10.6 (8.3)
Average # Phase 1 Awards per Competition	1.73 (1.13)
# Phase 2 Applications used in RD	919

[1]Competitions w/ $\geq$ 1 award
Note: This table summarizes the DOE Energy Efficiency & Renewable Energy (EERE) and Fossil Energy (FE) SBIR programs.

1.2.2 Private Finance Data

To match as many private financing deals to applicant companies as possible, I combined the ThompsonOne, Preqin, Cleantech Group i3, CrunchBase, and CapitalIQ databases. After matching by name and state, and hand-checking for accuracy, there are 838 firms with at least one private financing deal, of which 683 had at least one VC deal. Summary statistics about the matches are in Appendix 1D Table 2. Note that my private finance variables include IPOs and post-IPO transactions. I use “private” in the sense of non-government, as opposed to private equity. The matched VC deals by round type over time are in Appendix 1D Figure 6, and all private finance deals are in Appendix 1D Figure 7.[17]

[17]The paucity of matched deals before 2000 likely reflects the poorer quality of private transaction databases in earlier years and the lower volume of clean energy deals.

In Table 1.2, VC_i^{Post} is one if the firm ever received VC investment after its first grant award date.[18] This variable includes angel financing, which is qualitatively different from VC, but both target high-growth startups. I use binary indicators (or number of deals in robustness tests) and not dollar amounts for two reasons. First, VCs often report an investment but not the amount to survey firms, so the amount is available for a selected fraction of the deals. Second, there is rarely information about the pre-money valuation or how much the company sought to raise. A VC round of $1 million has a different value for a capital intensive battery company than for a smart phone energy efficiency app.

Table 1.2: Summary Statistics of Baseline Covariates and Dependent Variables

Covariate	Variable Type	Mean	Std. Dev.	Min	Max	N
MSA_i	0-1	0.304	0.46	0	1	5693
Age_i	Cont.	9.6	11.6	0	106	3808
$Minority_i$	0-1	0.081	0.27	0	1	1915
$Woman_i$	0-1	0.086	0.28	0	1	1915
$Exit_i^{\text{Post}}$	0-1	0.032	0.18	0	1	5693
$Exit_i^{\text{Prev}}$	0-1	0.033	0.18	0	1	5693
$\#SBIR_i^{\text{Prev}}$	Semi-Cont.	10.7	36.6	0	555	5693
VC_i^{Post}	0-1	0.11	0.31	0	1	5693
VC_i^{Prev}	0-1	0.077	0.27	0	1	5693
$Revenue_i$	0-1	0.55	0.50	0	1	5693
$Survival_i$	0-1	0.77	0.42	0	1	5365
$\#Patent_i^{\text{3 yrs Post}}$	Count	0.80	4.17	0	112	5693
$\#Patent_i^{\text{Prev}}$	Count	1.82	7.48	0	157	5693
$Citation_i^{\text{3 yrs Post}}$	Semi-Cont.	1.20	13.34	0	769.61	5693
$Citation_i^{\text{Prev}}$	Semi-Cont.	2.45	16.97	0	766.15	5693

Note: This table summarizes the variables used in the RD estimation. "Prev" indicates the variable prior to the firm's DOE SBIR application, and "Post" indicates afterward. See Appendix 1D Table 1 for additional statistics. First-time winners only. Year$\geq$ 1995

The variable $Exit_i$ takes a value of 1 if a firm has experienced an IPO or acquisition in the relevant time period. As in much of the literature, I am unable to distinguish acquisitions with high rates of return for investors from acquisitions that are an escape hatch, yielding modest or no

[18] For summary statistics on all private finance events and the number of deals, see Appendix 1D Table 1.

returns.[19] The majority of startups fail altogether, so a "selling for parts" exit at least indicates that the human capital or IP were valuable.

1.2.3 Real Outcome Data

I employ firm patents and a normalized citation metric as proxies for innovation quantity and quality, respectively. The data, from Berkeley's Fung Institute for Engineering Leadership, include all patents filed between 1976 and 2014. I matched non-reissue utility patents to applicant firms, and checked most by hand. Appendix 1D Table 4 contains summary statistics about the 2,109 firms with at least one patent. The pre- and post- treatment variables use the patent application date rather than the issue date, as is standard in the literature.

I do not normalize the patent count by USPTO classification or year because competition fixed effects control for sub-sector and date. For citations, however, I use the normalization from Lerner, Sorensen, and Strömberg (2011). It starts with a patent's forward citation count, which is the number of citations it receives from later patents within a three-year window after it was granted. I divide this count by the patent's class-year intensity.[20]

Data on firm survival and achieving revenue (commercialization) were collected by searching the internet for each firm to identify its current or historical status, website, and brief product description. Appendix 1D Table 3 summarizes the relevant information from this process. Roughly half of the companies in the estimation sample commercialized their technology, which I define as having ever sold their product or service. Less than a quarter are out of business as of May 2014. The revenue variable is not date-specific relative to the award. Section 1.4.3 discusses how this limits the interpretation of the RD estimates. Although the real outcome metrics are crude, an advantage is that I have data for each firm in my sample.

---

[19] Other papers that use all M&A events as positive exit outcomes include Gompers (1995), Hochberg, Ljungqvist, and Lu (2007), Puri and Zarutskie (2012), and Brander, Egan and Hellman (2008).

[20] This intensity is: $\gamma = \frac{\text{Total 3 Year Citations for a Class-Year}}{\text{Total Patents in a Class-Year}}$, where "Total 3 Year Citations for a Class-Year" are the number of citations made within 3 years to all patents in a given class-year.

1.3 Empirical Strategy

Regression discontinuity (RD) is a design that estimates a local average treatment effect around the cutoff in a rating variable - in my case the applicant's rank. The critical assumption in RD is that applicants cannot precisely manipulate their rank immediately around the cutoff. My institutional context, where firms are funded in rank order and the cutoff is exogenous to rank, permits a sharp RD comparing firms around the cutoff. As public agencies resist randomizing treatment to evaluate R&D subsidies (unlike new medicines), RD is the most plausibly exogenous variation possible (Jaffe 2002).

More specifically, a valid RD design must satisfy four conditions to be considered a local randomized experiment.[21] First, treatment cannot cause rank. This holds for the DOE SBIR program, as the award happens after ranking. To avoid contamination, I exclude applicants who previously won a grant within EERE/FE. Second, the cutoff must be exogenous to rank, which is true in my setting (Section 1.2.1). Third, the functional form must be correctly specified, else the estimator will be biased. I perform a goodness-of-fit test and show that rank is uninformative (Sections 1.4.1 and 1.7). Finally, to meet the key assumption that applicants cannot precisely manipulate their rank in the region around the cutoff, all observable factors must be shown to be locally continuous. To establish the necessary weak smoothness (see Hahn et al. 2001), I show continuity of covariates below.

Since the number of applicants and awards varies across competitions, I center the applicant ranks in each competition around zero at the cutoff. The lowest-ranked winner has centered rank $R_i = 1$, and the highest-ranked loser has $R_i = -1$. Each competition that I consider has at least this pair. As I expand the bandwidth, $[-r,\ r]$, I include higher ranked winners and lower ranked losers.[22]

I estimate variants of Equation 1, where Y_i^{Post} is the outcome and dependent variable. The coefficient of interest is τ on an indicator for treatment, and $f(R_{ic})$ is a polynomial controlling for the firm's rank within competition c.[23] The pre-assignment outcome variable is Y_i^{Prev}. I include a

[21] For more on RD, see Lee and Lemieux (2010).

[22] To assess composition issues, I also use percentile ranks and conduct a variety of tests, such as interacting raw rank with the number of awards in a competition.

[23] The standard RD implementation pools the data but allows the function to differ on either side of the cutoff by interacting the rank with treatment and non-treatment (Imbens and Lemieux 2008). However, I

full set of dummies for each competition δ_c, which are date-specific. X_i indicates other controls.[24] My estimations use OLS for binary dependent variables, negative binomial for count data, and two-part models for semi-continuous data.[25] Standard errors are robust and clustered by topic-year, to account for correlation in time and sector.

$$\begin{aligned} Y_{ic}^{\text{Post}} &= \alpha + \tau\left[\mathbf{1} \mid R_{ic} > 0\right] + f\left(R_{ic}\right) + \gamma_1 Y_{ic}^{\text{Prev}} + \gamma_2 X_{ic} + \delta_c + \varepsilon_{ic} \\ & \quad \text{where } -r \leq R_{ic} \leq r \end{aligned} \tag{1}$$

An important data limitation is the discreteness of my rating variable - competitions average ten applicants. Lee and Card (2008) note that discrete rating variables can require greater extrapolation of the outcome's conditional expectation at the cutoff. The fundamental econometrics are no different than with a continuous rating variable, however, as extrapolation is required in both cases. Section 1.7 demonstrates the robustness of my findings to this discreteness by, for example, separately considering competitions with certain numbers of awards.

To determine the appropriate polynomial, I employ Lee and Card's (2008) goodness-of-fit test for RD with discrete covariates, which compares unrestricted and restricted regressions. The former is a projection of the outcome on a full set of dummies for each of K ranks. The latter is a polynomial similar to Equation 1.[26] The null hypothesis is that the unrestricted model does not

---

potentially have too few points to the right of the cutoff to estimate a control function separately on both sides, so I rely on global polynomials for my primary specification. I show that my results are robust to allowing the slope coefficients to differ.

[24]The RD design does not require conditioning on baseline covariates, but doing so can reduce sampling variability. Lee and Lemieux (2010) advise including the pre-assignment dependent variable as they are usually correlated. Appendix 1G Table 1 projects rank on observable covariates. Previous non-DOE SBIR awards are the strongest predictor of rank. A one standard deviation increase in previous SBIR wins (the mean is 11.4 and the standard deviation is 38) increases the rank by nearly one unit. Previous VC deals also have a small positive impact. I include these two variables in my primary specifications.

[25]I use OLS for binary outcomes because many of the groups defined by fixed effects (competitions) have no successes (e.g. no subsequent VC). Logit drops the groups without successes. In such situations, Beck (2011) finds that OLS is superior despite his conclusion that logit is usually preferable with binary variables. Also, OLS with a binary variable is common in applied economics, following the arguments in Angrist (2001) that regression does as well as logit in estimating marginal effects and often better with binary treatment variables. My main results are intact with a logit specification (see Section 1.7).

[26]The goodness-of-fit statistic is: $G \equiv \frac{(ESS_{Restr.} - ESS_{Unrestr.})/(K-P)}{ESS_{Unrestr.}(N-K)}$, where ESS is the error sum of squares from regression, N is the number of observations, and P is the number of parameters in the restricted regression. G takes an F-distribution $F(K-P,\ N-K)$.

provide a better fit. If the goodness-of-fit statistic G exceeds its critical value for a certain level of confidence, then we can reject the null and turn to a higher order polynomial. The test results for each outcome metric are in Section 1.4.

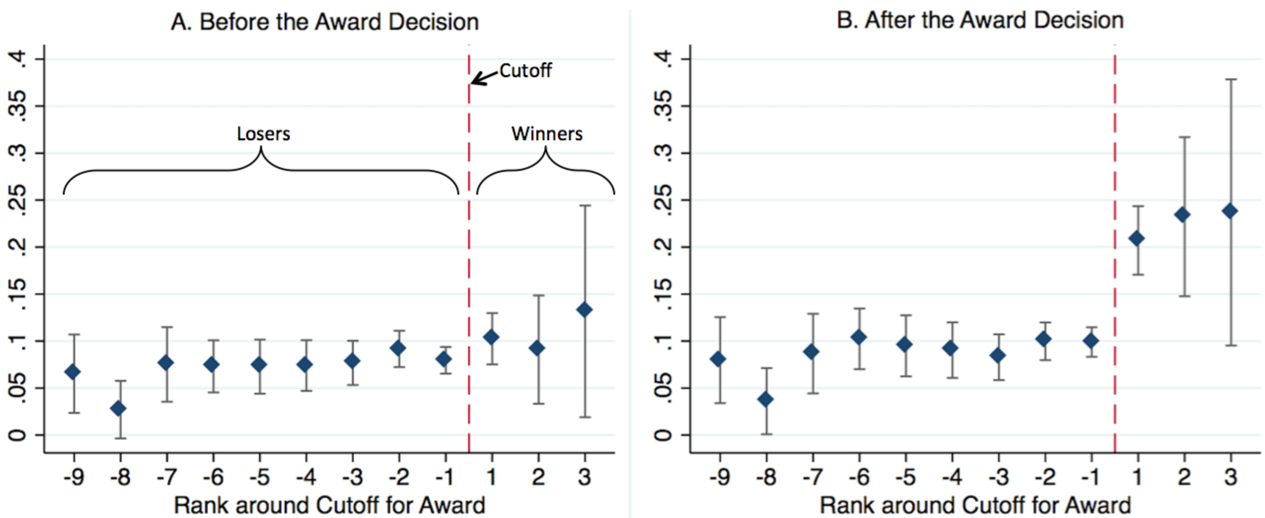

Figure 1.2: Probability of Venture Capital Financing Before and After Grant Decision by Rank
Note: This figure shows the fraction of applicants who ever received VC investment ever prior to (1A) and ever after (1B) the Phase 1 grant award decision. The applicants are binned by their DOE assigned rank, which I have centered so that Rank > 0 indicates a firm won an award. Capped lines indicate 95% confidence intervals. N=4,812.

I demonstrate smoothness in observable baseline covariates in three ways: visually, through an RD on baseline covariates, and through differences in means. First, I show at each rank the means of baseline covariates, most importantly the pre-assignment outcome variables VC investment (Figure 1.2 A), patenting (Figure 1.3 A), exit (Figure 1.4 A), and all private finance (Appendix 1D Figure 8). For ease of comparison, these are shown adjacent to the post-treatment variables. Four additional covariates are in Appendix 1H Figure 1; average age as well as the probability a firm is located in a major metro area, is woman owned, and is minority owned. In none of the eight figures is there any discontinuity around the cutoff visible, nor is there any trend in rank. A ninth covariate is the exception: previous non-DOE SBIR wins (Appendix 1H Figure 2). Rank is clearly increasing in previous wins, but again there is no discontinuity around the cutoff.[27]

[27]See Appendix 1F for analysis of multiple SBIR wins.

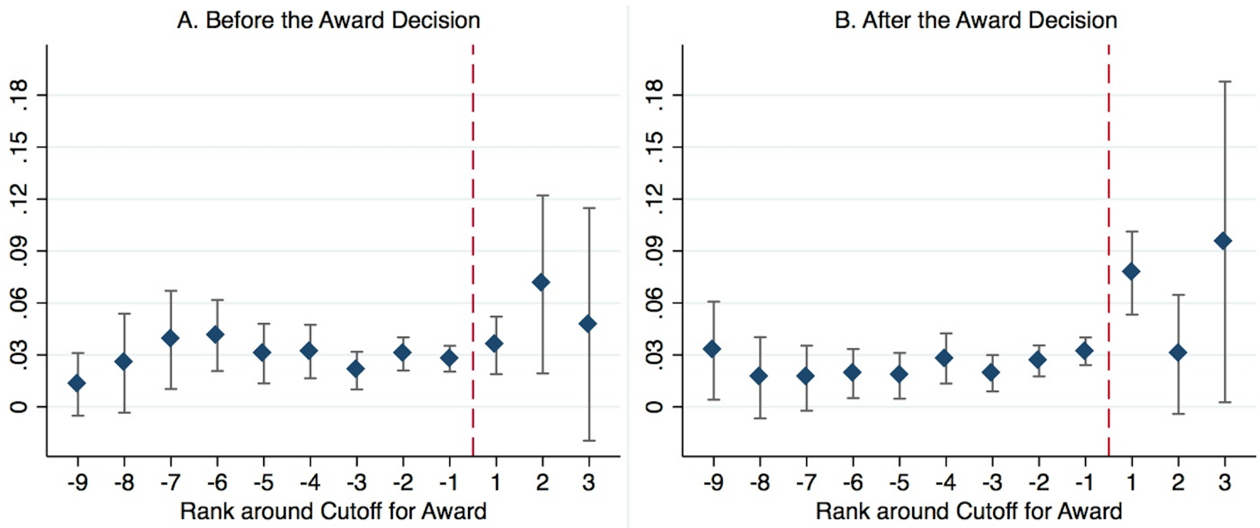

Figure 1.4: Probability of Exit (IPO or Acquisition) Before and After Grant Decision by Rank
Note: This figure shows the fraction of applicants who ever experienced an exit (IPO or acquisition) ever prior to (5A) and ever after (5B) the Phase 1 grant award decision. The applicants are binned by their DOE assigned rank, which I have centered so that Rank > 0 indicates a firm won an award. Capped lines indicate 95% confidence intervals. N=4,816.

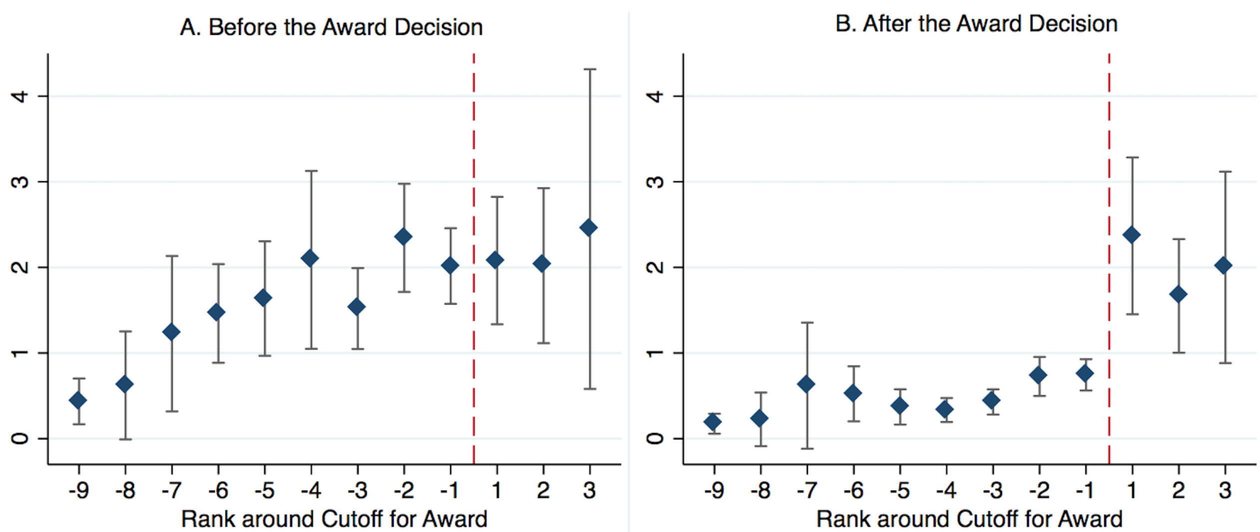

Figure 1.3: Number of Patents Before and After Grant Decision by Rank
Note: This figure shows firm patents ever prior to (2A) and within three years after (2B) the Phase 1 grant award decision. The applicants are binned by their DOE assigned rank, which I have centered so that Rank > 0 indicates a firm won an award. The date associated with a successful patent is the patent application date. Capped lines indicate 95% confidence intervals. N=4,816.

Second, I try to detect a discontinuity in the outcome predicted by the baseline covariates, following Card, Chetty and Weber (2007) and Imbens and Lemieux (2008). I use an OLS regression of the outcome of interest, Y_{ic}^{Post}, on baseline covariates and competition dummies to obtain a weighted average of the covariates by relevance to the outcome:

$$Y_{ic}^{\text{Post}} = \alpha + X_i\phi + \delta_c + \varepsilon_{ic} \tag{2}$$

For each applicant I then use the estimated coefficient vector to predict the probability of subsequent VC financing: $\hat{Y}_{ic}^{\text{Post}} = \hat{\alpha} + X_i\hat{\phi} + \hat{\delta}_c$. I average the probabilities for each rank and plot them in Appendix 1H Figure 3. There is no obvious discontinuity around the cutoff, in striking contrast to the actual outcome in Figure 1.2 B.

Third, I conduct a t-test for matched pair differences of means in baseline covariates immediately around the cutoff, as in Kerr et al. (2014). The null hypothesis is that the mean of the covariate for $R_i = -1$ applicants is the same as for $R_i = 1$ applicants: $H_o = \bar{X}_1 - \bar{X}_{-1} = 0$. The first alternative hypothesis is a two-tailed test: $H_1 = \bar{X}_1 - \bar{X}_{-1} \neq 0$. The second is a one-tailed test: $H_2 = \bar{X}_1 - \bar{X}_{-1} > 0$ (this is most relevant for the pre-application covariates). The results are in Appendix 1G Table 22. The two-tailed tests cannot reject the null at the 10% level for any covariate. The one-tailed tests find a significant difference only for previous citations (at the 10% level). However, adding or removing these covariates from the regression has essentially no effect on my results. I also estimate whether treatment can predict each covariate individually. In Appendix 1G Table 21, I regress each the 10 baseline covariates on treatment. None of the treatment effects have any significance.

Program officials observe more data than the econometrician, so it is impossible to fully test the assumption of no sorting on observables in the neighborhood of the cutoff. Nonetheless, this preponderance of evidence suggests the RD design is valid.

1.4 The Grant Impact on Firm Outcomes

I find strong effects of the grant on financial and real outcomes, summarized in Table 1.3. A Phase 1 award nearly doubles a firm's probability of venture capital finance and leads to almost three times as many patents. It also increases a firm's likelihood of reaching revenue and of achieving a liquidation event. The effects are consistently stronger for younger, more inexperienced firms. In contrast with the large Phase 1 impact, Phase 2 has no effect on any outcome other than patents, where it has a much weaker effect than Phase 1.

Table 1.3: Summary of Results

Outcome Metric	*A Phase 1 award:*	*A Phase 2 award:*
Venture Capital Finance	increases firm's probability of VC investment by 9 percentage points (average 12%) effect stronger for firms that: - are young - are in immature sectors - are in lean times - have no previous SBIR awards - are in VC-intensive regions	has no effect
Number of Patents	leads firm to produce 3 times more patents within three years (average 0.92 patents); has no long term effect effect stronger for firms that: - are young - have no previous patents - are in high propensity to patent sectors - have no previous SBIR awards	leads to 1.5 times more patents (average 2.2 patents)
Number of Normalized Patent Citations	has no effect	leads awardees to be 85% more likely to have positive citations[1]
Reaching Revenue	increases firm's probability of achieving revenue by 11 percentage points (average 56%)[2] effect stronger for firms that: - have no previous SBIR awards	has no effect
Survival	has no effect	has no effect
Exit (IPO or Acquisition)	increases firm's probability of exit by 3.5 percentage points (average 4%)[3]	has no effect

Note: This table summarizes the principal robust and precisely estimated results from the RD estimation. A firm first applies for a Phase 1 award of $150,000, and may then apply a year later for a Phase 2 award of $1,000,000. For the detailed results and variable descriptions, see Section 1.4 for VC, Section 1.5.1 for Revenue, Survival, and Exit, and 1.5.2 for Patents.

[1]This is a strong effect along the extensive margin. However, I find no effect along the intensive margin (conditional on firms having positive citations, there is no effect of the award).

[2]This variable is not date-specific, so while the estimated effect tells us that a positive impact exists, the magnitude cannot be interpreted as causal.

[3]This result is less visually and statistically significant than the others.

I begin with the long-term effect of the Phase 1 grant on VC. Subsequent sections use variation in firm characteristics and over time to reinforce the case that the grant eases financial constraints. I also test for reallocation of capital within competitions. Last, I evaluate the Phase 2 grant. Section 1.4.2 assesses the effect on patents and patent citations, considering heterogeneity across firms, the Phase 2 effect, and the relationship of VC finance to patenting. Finally, Section 1.4.3 examines commercialization, exit, and survival.

1.4.1 The Grant Impact on Venture Capital Investment

Startups' typically have little or no tangible collateral, so they often cannot initially access debt finance. VC is their main source of external capital outside of partnering with a larger corporation (Hall and Woodward 2007). VC accomplishes two important goals as an outcome metric. First, it tests whether the grants mobilize or crowd out private investment. Second, observing subsequent VC investment indicates that the company presents a privately profitable investment opportunity.

VC investment is not only a financial outcome, but is also as a good early-stage proxy for market success in a context where outcome data are difficult to collect. The literature has established that venture capitalists are important intermediaries in the U.S. innovation system.[28] They select innovative firms and bring new technologies to market quickly (Hellmann and Puri 2000, Sorenson 2007, Engel and Keilbach 2007). The VC commitment also makes debt finance easier to obtain (Hochberg, Serrano and Ziedonis 2014). VCs further provide non-monetary resources, such as intensive monitoring, improved governance, legal services, and networking. Chemmanur et al. (2011) find that VC-backed manufacturing firms have higher productivity prior to receiving VC finance, but that after controlling for this screening, VC-backed firms also subsequently experience faster growth. Kortum and Lerner (2000) exploit the 1979 pension fund policy shift and find that $1 of VC money produces 3-4 times more patents than $1 of corporate R&D. Further, DOE officials consider mobilizing private investment to be an important goal.

Visual evidence for a grant treatment effect on VC is in Figure 1.2 B. The probability of subsequent VC jumps from about 10% to 20% around the grant cutoff. Table 1.4 contains this

[28]The U.S. VC industry has grown dramatically since its origins in the 1960s. Over the past decade it has invested $20-$30 billion annually in portfolio companies, up from about $8 billion in 1995 (NVCA 2014). VC firms invested between $4 and $7 billion annually in U.S. clean energy in recent years (see Appendix 1D Figure 5).

Table 1.4: Impact of Phase 1 Grant on VC with Linear and Quadratic Control Functions

Dependent Variable: VC_i^{Post}							
Bandwidth:	1	2		3		All	
	I.	II.	III.	IV.	V.	VI.	VII.
$\mathbf{1} \mid R_i > 0$	**.098*****	**.09*****	**.14****	**.1*****	**.12****	**.11*****	**.072****
	(.032)	(.025)	(.058)	(.023)	(.058)	(.021)	(.033)
VC_i^{Prev}	.27***	.32***	.32***	.31***	.31***	.32***	.32***
	(.057)	(.038)	(.038)	(.036)	(.036)	(.029)	(.029)
$\#SBIR_i^{\text{Prev}}$	.0012***	.001***	.001***	.001***	.001***	.00087***	.00084***
	(.00034)	(.00029)	(.00029)	(.00027)	(.00027)	(.00024)	(.00024)
R_i			-.02		-.029		.0086
			(.021)		(.033)		(.0071)
R_i^2					.012		-.000074
					(.0088)		(.00043)
Competition f.e.	Y	Y	Y	Y	Y	Y	Y
N	1872	2836	2836	3368	3368	5021	5021
R^2	0.47	.39	.39	.34	.35	.27	.27

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on VC. The likelihood of receiving VC after the grant is 10.9%; among losers it is 9.4%, and among winners it is 21.3% (bandwidth=all specification). The specifications are variants of the model in Equation 1. The dependent variable ${VC^{\text{Post}}}_i$ is 1 if the company ever received VC after the award decision, and 0 if not. Specifications vary the bandwidth around the cutoff and control for rank linearly and quadratically. Standard errors are robust and clustered at the topic-year level. *** $p < .01$. Year$\geq$ 1995

difference in regression form. The dependent variable (VC_i^{Post}) is one if a firm ever subsequently received VC investment, and zero if it did not. Column I finds that an award increases the probability of subsequent venture funding by 9.8 percentage points (hereafter pp), significant at the 1% level, with the narrowest bandwidth possible of one rank on either side of the cutoff. Subsequent columns find effects between 7.2 and 14 pp using larger bandwidths of two, three, and all my data.[29] Note that the overall likelihood of receiving VC after the grant is 10.9%; among losers it is 9.4%, and among winners it is 21.3% (with the bandwidth=All specification). I control for centered rank linearly with a bandwidth of two ($f(R_{ic}) = \beta_1 R_{ic}$), and quadratically with wider bandwidths

[29]Appendix 1G Figure 1 depicts the predictive margins. It shows the conditional expectation of VC^{Post} by rank, calculated at the mean of all the other independent variables. I use a linear rank specification around the cutoff with BW=all.

$\left(f\left(R_{ic}\right)=\beta_1 R_{ic}+\beta_2 R_{ic}^2\right)$. My preferred estimate is 9 pp (column II).[30]

Table 1.5: Impact of Phase 1 Grant on VC with Percentile Rank Control (Quintiles)

Dependent Variable: VC_i^{Post}				
Bandwidth:	I. 1	II. 2	III. 3	IV. all
$\mathbf{1} \mid R_i > 0$	**.098*** **	**.1*** **	**.094*** **	**.1*** **
	(.032)	(.035)	(.033)	(.028)
VC_i^{Prev}	.27***	.32***	.31***	.32***
	(.057)	(.038)	(.036)	(.029)
$\#SBIR_i^{\text{Prev}}$	.0012***	.001***	.001***	.00085***
	(.00034)	(.00029)	(.00027)	(.00024)
R_iQ2		.016	-.01	.011
		(.032)	(.028)	(.022)
R_iQ3		.019	.0043	-.022
		(.042)	(.033)	(.022)
R_iQ4		.014	-.026	-.039
		(.047)	(.036)	(.026)
R_iQ5		-.026	-.05	-.044
		(.062)	(.041)	(.029)
Competition f.e.	Y	Y	Y	Y
N	1872	2836	3368	5021
R^2	.47	.39	.35	.27

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on VC. The specifications are variants of the model in Equation 1. The dependent variable $VC^{\text{Post}}{}_i$ is 1 if the company ever received VC after the award decision, and 0 if not. Ranks are transformed into the applicant's percentile rank within his competition. The highest quantile is omitted. Standard errors are robust and clustered at the topic-year level. *** $p < .01$. Year$\geq$ 1995

The models with and without rank controls in Table 1.4 yield fairly similar coefficients. The ranks do not contain much information about an applicant's chances of VC financing. The Lee and Card (2008) goodness-of-fit test reveals that once I control for award, no function is too

[30]Note that in specifications with bandwidth "all," the data are not symmetric around the cutoff. In Appendix 1G Table 2 I use quadratic specifications that do not restrict the slope to be the same on either side. The coefficients jump to 16.7 and 23.2 pp with BW=2 and BW=3, but return to 11.5 pp with BW=all. Compared with Table 1.4, the standard error increases when rank is added, indicating that rank is correlated with treatment. It is difficult to distinguish the effect of winning from the rank because of the coarseness of my rating variable. The confidence interval implied by the standard errors from Appendix 1G Table 2 include my preferred estimate of 9 pp. Any bias from excluding rank is downward rather than upward, which is reassuring if the concern is overstating the result.

restrictive.[31] We might worry that information in the raw rank is lost when I center the ranks around the cutoff. A firm with a centered rank of two in a competition with two awards might be of different quality than in a competition with four awards. I create percentile ranks to address this possibility. Regressions controlling for quintiles in rank within a competition, instead of centered rank, are in Table 1.5. The coefficients on treatment range from 9.3 to 10.1 pp, all significant at the 1% level.[32]

The grant effect on VC happens quickly. This confirms that the long term effect above is indeed due to the grant, and also tells us that whatever mechanism explains the grant effect must act rapidly. Within one year of the award a grantee is 5.8 pp more likely than a loser to receive VC, significant at the 1% level, (Table 1.6 column 1). This is more than half the total effect. Subsequent columns show the cumulative effect over time; for example, within two years the effect is 7.5 pp and within four years it is 8.2 pp, both also significant at the 1% level.

When I include all private financing events, such as IPOs, acquisitions, and debt, I find a slightly larger effect of about 12 pp. The probability of funding jumps from 12% to 26% around the cutoff, shown visually in Appendix 1D Figures 8 and 9. Appendix 1G Tables 4-7 replicate the VC findings with all private financing (PF_i^{Post}) as the dependent variable, and find analogous results.

Variation in the Effect Across Firm Age and Sector

If the grants ease financing constraints, then the estimated effect ought to be larger for more constrained firms. In this section and the next, I examine variation in the effect across firm characteristics and over time. Since these variables are not randomly assigned, the analysis is necessarily more speculative than the affirmative conclusions in the main result above.

---

[31] G-values from the goodness-of-fit test are tiny. With no control for rank, $G = 0.000028$, while the critical value above which I could reject the null even with 15% confidence is 1.27. In F-tests for regressions with linear and quadratic rank, I find that the G-value remains miniscule.

[32] I find the same result using quartile ranks (Appendix 1G Table 3). See Section 1.7 for further robustness tests, including regressions estimated on subsamples with specific numbers of awards, and dummies for raw rank interacted with the number of awards.

Table 1.6: Temporal Impact of Phase 1 Grant on VC

Dependent Var.:	I. $VC_i^{\text{0-1 yr Post}}$	II. $VC_i^{\text{0-2 yr Post}}$	III. $VC_i^{\text{0-3 yr Post}}$	IV. $VC_i^{\text{0-4 yr Post}}$	V. $VC_i^{\text{0-5 yr Post}}$	VI. $VC_i^{\text{0-6 yr Post}}$
$\mathbf{1} \mid R_i > 0$	**.058*** **	**.075*** **	**.074*** **	**.082*** **	**.079*** **	**.083*** **
	(.017)	(.019)	(.019)	(.021)	(.021)	(.021)
VC_i^{Prev}	.24***	.32***	.32***	.32***	.33***	.33***
	(.029)	(.033)	(.034)	(.035)	(.035)	(.035)
$\#SBIR_i^{\text{Prev}}$	-.000027	-.00004	-.000065	.000039	.00011	.000092
	(.00016)	(.0002)	(.0002)	(.00024)	(.00024)	(.00024)
Competition f.e.	Y	Y	Y	Y	Y	Y
N	3368	3368	3368	3368	3368	3368
R^2	.36	.38	.39	.38	.37	.37
Dependent Variable: VC_i^{Post}						
	VII. 1995-1999	VIII. 2000-2004	IX. 2005-2009	X. 2009-2013	XI. 2009-2011	XII. 2009
$\mathbf{1} \mid R_i > 0$	**.076***	**.047**	**.07****	**.19*** **	**.13*** **	**.1***
	(.04)	(.036)	(.031)	(.047)	(.039)	(.055)
VC_i^{Prev}	.096	.3***	.41***	.34***	.42***	.43***
	(.062)	(.078)	(.045)	(.049)	(.04)	(.066)
$\#SBIR_i^{\text{Prev}}$	.0019***	.0017***	.00039	-.001***	-.001***	-.00092*
	(.00025)	(.00034)	(.00028)	(.00038)	(.00025)	(.0005)
Competition f.e.	Y	Y	Y	Y	Y	Y
N	1392	1052	1970	3160	2192	893
R^2	.23	.3	.26	.39	.31	.26

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on VC over time. The specifications are variants of the model in Equation 1. The dependent variables in the top panel are indicators for whether a firm received VC investment within a certain number of years from the award. For example, $VC_i^{\text{0-1 yr Post}} = 1$ if the company received VC within one year of the award. The top panel uses BW=3. The bottom panel limits the sample to certain time periods, where years are inclusive, and uses BW=all. The dependent variable $VC^{\text{Post}}{}_i$ is 1 if the company ever received VC after the award decision, and 0 if not. Standard errors are robust and clustered at the topic-year level. *** $p < .01$. Year $\geq$ 1995

First, young firms tend to be more financially constrained - there is less information available about them, and they generally have fewer assets (e.g. Brown, Fazzari and Petersen 2009, Whited and Wu 2006). Indeed, young firms experience much stronger grant treatment effects. Table 1.7 Column I includes only firms less than three years old and finds that a grant increases the likelihood of subsequent VC by 17 pp (significant at the 5% level), while for firms older than three the effect is 9.2 pp (column II). Similarly, the effect for firms less than ten years old is 14 pp, significant at the

1% level, but for firms ten years or older, it is only 4.7 pp (columns IV and V). I jointly estimate the young and old regressions by fully interacting the variables, including fixed effects, with dummies for age group. The coefficient on the difference between the treatment effect for firms younger and older than nine is 9.3 pp, significant at the 5% level (column VI).[33]

This result is in keeping with the model in Acemoglu et al. (2013), where R&D subsidies to entrants increase welfare, but subsidies to incumbents decrease welfare. Policymakers might consider targeting young firms for grants. Not only do they experience the largest grant effects, but also young companies generate greater innovation and growth than simply small companies (Evans 1987, Calvo 2006).

Immature technologies without well-developed markets or supply chains, such as solar and geothermal, are riskier investments than incumbent technologies, such as coal and natural gas. I create a binary variable, $Immature_i$, which is one if the sector is solar, wind, geothermal, fuel cells, carbon capture and storage, biomass, or hydro/wave/tidal; and zero if the sector is oil, gas, coal, biofuels, or vehicles/motors/engines.[34] More ambiguous sectors are excluded. The grant effect is 18 pp for immature sectors, but only 7.2 pp for mature sectors (Table 1.7 columns X-XI). Both coefficients and their difference (column XII) are significant at conventional levels.[35]

---

[33]This is equivalent to an F-test for equality of the coefficients in the separate regressions.

[34]Most electric vehicle and hydrogen car competitions are classified as batteries or fuel cells. The sector categorizations are based on the topic to which the firm applied.

[35]The degree to which some of these sectors are mature may have changed over time, so Appendix 1G Table 8 considers the sample from 2007, and finds roughly the same results.

Table 1.7: Impact of Phase 1 Grant on VC by Firm Age, Location, & Sector Maturity

Dependent Variable: VC_i^{Post}	I. $Age_i \leq 2$	II. $Age_i > 2$	III. I & II	IV. $Age_i \leq 9$	V. $Age_i > 9$	VI. IV & V
$\mathbf{1} \mid R_i > 0$	**.17****	**.092*****	.092***	**.14*****	**.047***	.047*
	(.069)	(.021)	(.016)	(.031)	(.024)	(.024)
$\mathbf{1} \mid R_i > 0 \cdot (\mathbf{1} \mid Age_i \leq X)$			**.076***			**.093****
			(.043)			(.039)
VC_i^{Prev}	.44***	.31***	.31***	.37***	.18***	.18***
	(.11)	(.032)	(.021)	(.041)	(.053)	(.053)
$\#SBIR_i^{\text{Prev}}$	.0043	.001***	.001***	.0012**	.0012***	.0012***
	(.0027)	(.00024)	(.00014)	(.00053)	(.00028)	(.00028)
Topic f.e.	Y	Y	Y	Y	Y	Y
Topic f.e.$\cdot$ $(\mathbf{1} \mid X)$	N	N	Y	N	N	Y
N	576	2792	3368	1574	1876	3368
R^2	.52	.22	.31	.33	.23	.34

	VII. Same MSA	VIII. Different MSAs	IX. VII & VIII	X. Mature	XI. Immature	XII. X & XI
$\mathbf{1} \mid R_i > 0$	**.12*****	**.099*****	.099***	**.072****	**.18*****	.072**
	(.04)	(.021)	(.021)	(.036)	(.04)	(.036)
$\mathbf{1} \mid R_i > 0 \cdot (\mathbf{1} \mid \text{Same MSA})$			**.02**			
			(.044)			
$\mathbf{1} \mid R_i > 0 \cdot (\mathbf{1} \mid Imm.)$						**.11****
						(.054)
VC_i^{Prev}	.3***	.33***	.33***	.23***	.39***	.23***
	(.056)	(.034)	(.034)	(.059)	(.045)	(.059)
$\#SBIR_i^{\text{Prev}}$	.001***	.00095***	.00095***	.001**	.00028	.001***
	(.00038)	(.00023)	(.00023)	(.00038)	(.00034)	(.00038)
Topic f.e.	N	N	N	Y	Y	Y
Topic f.e.$\cdot$ $(\mathbf{1} \mid X)$	N	N	N	N	N	Y
Competition f.e.	Y	Y	Y	N	N	N
Competition f.e.$\cdot$ $(\mathbf{1} \mid X)$	N	N	Y	N	N	N
N	1380	4312	5692	1330	1820	3150
R^2	.23	.26	.26	.18	.2	.2

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on VC (variants of Equation 1 w/ BW=3). I to V divide the sample by firm age (years) at application. III & VI jointly estimate the preceding regressions to provide a std error on the difference (bold). VII-IX assess the reallocation effect w/ BW=all. VII includes firms on each side of the cutoff within a topic from the same city (MSA). VIII estimates the effect when competing firms are from different MSAs. X-XII use an indicator for immature sectors. I use topic f.e. where needed for sufficient within-group observations. Control coefficients not reported. Std errors robust and clustered by topic-year. *** $p < .01$. Year$\geq$ 1995

Table 1.8: Impact of Phase 1 Grant on VC Investment by Technology Type

Dependent Variable: VC_i^{Post}		
Technology (sub-sector)	Coefficient on treatment ($\mathbf{1} \mid R_i > 0$)	N
Geothermal	**.56*** (.24)	51
Hydropower, Wave & Tidal	**.51**** (.19)	181
Solar	**.25**** (.11)	421
Carbon Capture & Storage	**.2**** (.091)	211
Building & Lighting Efficiency	**.14**** (.057)	370
Vehicles, Motors, Engines, Batteries	**.12**** (.06)	726
Wind	**.11**** (.039)	194
Advanced Materials	**.11** (.071)	435
Biomass Production/ Generation	**.085** (.067)	308
Fuel Cells & Hydrogen	**.077** (.0723)	400
Natural Gas	**.06** (.074)	255
Recycling, Waste to energy & Water	**.045** (.053)	549
Smart Grid, Sensors & Power Converters	**.045** (.053)	634
Air & Emission Control	**.025** (.035)	300
Coal	**.024** (.053)	108
Biofuels & Biochemicals	**.014** (.054)	176

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on VC by technology (sub-sector) using BW=all. Here I report only the coefficient on treatment. A full table is in Appendix 1G Table 10. The specifications are variants of the model in Equation 1, but each includes only competitions whose topics fall within the specific technology. Other and "Oil" are omitted due to few observations. Control coefficients are not reported for brevity. Standard errors are robust and clustered at the topic-year level. *** $p < .01$. Year$\geq$ 1995.

Separate regressions for each clean energy technology (Table 1.8) confirm that the grants are most beneficial for emerging energy generation technologies. For example, a grant makes a solar company 25 pp more likely to get subsequent VC investment, increasing the probability from roughly 11% for losers to 35% for awardees. For wind companies, the grant increases the probability of subsequent VC from about 5% to 16%. There is no correlation between the grant effect on VC in a sector and that sector's propensity to receive VC.[36] These emerging energy sub-sectors have positive externalities from reduced pollution and greenhouse gases. Mitigating climate change does

[36]Without controlling for treatment, I project subsequent VC on sector dummies in Appendix 1G Table 9. Vehicles/batteries and advanced materials are among the most likely to receive VC, but have weak treatment effects. Meanwhile, solar and efficiency are relatively likely likely to receive VC and also have strong treatment effects. Wind is unlikely to receive VC, but the grant has a dramatic impact.

not enter most private sector return calculations, but it is one of DOE's central objectives. My results indicate that subsidies have the greatest impact when awarded to clean energy generation technologies, rather than to projects that improve efficiency in a mature sector.

Variation in the Effect over Time

The results thus far pool all years between 1995 and 2013, but the effect has actually changed somewhat over time. The bottom panel of Table 1.6 divides the sample into four five-year periods. Between 1995 and 1999, the effect is 7.6 pp. It drops to 4.7 pp between 2000 and 2004, perhaps because VC firms were focused on internet startups at the beginning of the period, then dramatically reduced investing when the internet bubble collapsed. The effect returns to 7 pp in 2005-2009. The strongest effect is between 2009 and 2013 at 19 pp. I focus on the ARRA years of 2009-2011, when DOE funding was unusually high, in columns XI and XII. Some investors I interviewed believed that in this period there was "too much government money chasing too few good projects." But the estimated grant effect is 13 pp for the whole Stimulus period. Despite a large spike in applicants in 2009, limiting the sample to that year yields the same effect as the whole sample.

The economic environment may explain these across-time period differences. Unlike large firms, startups cannot use cash reserves to smooth R&D investment over time and have little control over when their invention requires an infusion of capital (Himmelberg and Petersen 1994). If the grants mitigate entrepreneurs' financing constraints, they should be more powerful in lean times when external financing is more difficult to attain.

Tobin's Q, the ratio of a firm's market value to its book value, is widely employed in the literature to measure investment opportunities (e.g. Stein 2003, Gompers, Lerner and Scharfstein 2005). Q can also be interpreted as an indicator of financing availability, as in Baker, Stein and Wurgler (2003). I hypothesize that in low-Q environments firms face greater difficulty accessing external finance, making the grant more useful. But the grant could act pro-cyclically if, say, there are always more worthy startups seeking funding than willing investors, but the supply of entrepreneurs is positively elastic to hot markets.

My simplified measure of Q follows Kaplan and Zingales (2007) and Gompers, Ishii and Metrick (2003).[37] I use NAICS codes to identify companies in the clean energy sector, and calculate

[37] Q is calculated using the equation below, where BV is book value, MV is market value (price times shares outstanding), and DT is balance sheet deferred taxes. Data is from Compustat via Wharton Research Data Services. The book value is in fiscal year t and the common stock value is at the end of calendar year

Q annually by company.[38] I interact the treatment variable with median sector Q_{t+1} (Q in the four quarters following the award), which I demean so that the coefficient on treatment alone reflects the impact of the grant at mean Q. The results, in Table 1.9, show that the grant effect decreases significantly as Q increases. A one standard deviation increase in Q is associated with a 4 pp decrease in the grant effect. I also divide the years into periods of low and high Q, and find that the difference in the effect between periods is 9.2 pp, significant at the 5% level (column III of Appendix 1G Table 11).

The private sector's disinterest in funding startups when industry Q is low makes sense under both Q interpretations: low Q implies poor investment opportunities or that the market undervalues the investment opportunities. Under the investment opportunities interpretation, VC firms - who are relatively unconstrained and thus Q-sensitive - should invest less in clean energy startups when industry Q is low. Market failure occurs because startups' financing constraints disrupt the linkage between Q and investment. Worthwhile startups with the bad luck (or poor choice) to commercialize their invention when industry Q is low cannot substitute other resources for venture funding. They find the grant more valuable.

A different angle on access to finance is VC investment in portfolio companies, which is quite volatile (Nanda and Rhodes-Kropf 2012, Jeng and Wells 2000). This volatility may reflect irrational herding, as in Scharfstein and Stein (2000), or it may reflect shocks to investment opportunities, as in Gompers et al. (2008). I expect that when VC availability is high, firms are less financially constrained, so the grant effect is diluted.

The right panel of Table 1.9 explores how the grant effect varies with the total number of U.S. VC deals over the eight quarters following the grant.[39]

---

t.

$$Q_t = \frac{MV_t^{Assets}}{BV_t^{Assets}} = \frac{BV_t^{Assets} + MV_t^{\text{CommonStock}} - (BV_t^{\text{CommonStock}} + DT)}{BV_t^{Assets}}$$

[38]The sector median is plotted in Appendix 1D Figure 10, and summary statistics are in Appendix 1D Table 6. See Appendix 1D Table 5 for NAICS codes that define the clean energy sector.

[39]I use data from ThompsonOne (Appendix 1D Figure 10, summarized in Appendix 1D Table 6).

Table 1.9: Impact of Phase 1 Grant on VC with Varying External Capital Availability

Dependent Variable: VC_i^{Post}						
Time Series Variable:	Clean Energy Industry Tobin's Q (Q_{t+1})			Total U.S. VC Deals ($\#VC_{t+2}$)		
	I. BW=2	II. BW=3	III. BW=all	IV. BW=2	V. BW=3	VI. BW=all
$(\mathbf{1} \mid R_i > 0) \cdot Q_{t+1}$	**-.2**	**-.26****	**-.22****			
	(.14)	(.13)	(.11)			
$(\mathbf{1} \mid R_i > 0) \cdot \#VC_{t+2}$				**-.02***	**-.03****	**-.025****
				(.011)	(.012)	(.01)
$\mathbf{1} \mid R_i > 0$	.12***	.14***	.15***	.122***	.15***	.16***
	(.03)	(.031)	(.025)	(.031)	(.032)	(.027)
Q_{t+1}	.20	.26**	-26.49***			
	(.14)	(.13)	(.95)			
$\#VC_{t+2}$				.02*	.03**	-.63***
				(.011)	(.012)	(.022)
Competition f.e.	Y	Y	Y	Y	Y	Y
N	2836	3368	5021	2836	3368	5021
R^2	.32	.28	.18	.32	.28	.18

Note: This table reports regression estimates of the Phase 1 grant effect $(\mathbf{1} \mid R_i > 0)$ on VC interacted with time series metrics for Q and VC flow. The dependent variable $VC^{\text{Post}}{}_i$ is 1 if the company ever received VC after the award decision, and 0 if not. The specifications are variants of the model in Equation 1. The left panel uses a measure of clean energy industry Tobin's Q over the 4 quarters following the award decision. The right panel uses the total number of VC investments in U.S. companies over the 8 quarters following the award decision. Both variables are demeaned, and VC deals also divided by 1,000. Standard errors are robust and clustered at the topic-year level. *** $p < .01$. Year$\geq$ 1995

The coefficient on the interaction between treatment and number of deals is negative and significant at the 5% level. It implies that a one standard deviation increase in deal flow is associated with a 5.3 pp decrease in the grant's effect. The alternative specification finds that the difference in the treatment effect between high and low deal flow periods is 6.6 pp, significant at the 10% level (column VI of Appendix 1G Table 12). When I perform this exercise within only one year of the grant ($\#VC_{t+1}$), I find a smaller and insignificant difference.

It seems that a grant is more valuable in times of low Tobin's Q and low VC availability. This counter-cyclicality reinforces the conclusion that energy startups face severe financing constraints, like the across-period findings in Fazzari, Hubbard and Petersen (1988). Yet this heterogeneity analysis is an exercise in theory-motivated correlations, so other economic conditions may drive the relationships.[40] However, my counter-cyclical finding accords with Tian and Wang's (2014) conclusion that being financed by a failure-tolerant VC is more important for innovation when ventures are founded in recessions. Related research finds that R&D investment is pro-cyclical, declining in recessions due to financing constraints. This body of work includes Aghion et al. (2012), Campello, Graham, and Harvey (2010), and Ouyang (2011).

Testing for Spillovers

Thus far I have assumed that awardees do not affect losing applicants. But a grant might increase an awardee's chance of VC by *decreasing* the losers' chance. In this section I test whether my RD estimates reflect negative spillovers. Unfortunately, I cannot test whether capital is reallocated from non-applicant firms to winning firms, or whether total VC investment in clean energy changes as a result of the grant program.

To test for reallocation of capital within the applicant pool, I conduct two tests. First, I ask whether the likelihood of a losing firm obtaining VC varies with the number of winners in the competition. Recall that within a competition firms are doing very similar activities - they are in the same narrowly defined sub-sector. Also recall that the number of awards in a competition is unrelated to the technology type, program office, time period, and ranking process. Therefore, if there are negative spillovers from winners to losers, these should be more intense when there are multiple winners in the competition. I regress the outcome on the subset of losers and include in separate models dummies for having either more than one, or more than two, awards in the competition. I find that these dummies have no predictive power, suggesting that spillovers do not explain the main effect (see Appendix 1G Table 31).

Second, I exploit the robust finding in the literature that VC firms typically invest in geographic proximity to their offices, and indeed in firms located in their city (Sorenson and Stuart 2001, Samila and Sorenson 2011). Chen et al. (2010) point out that distant monitoring is costly, which is one reason why portfolio companies typically have at least one investor in the same metro region. Cumming and Dai (2010) also find strong local bias in VC investments. They calculate the average distance between a company and its venture investor at less than 200 miles since 1998.

Geographically close firms competing for an SBIR grant are much more likely than firms far

[40] I also tested the correlation of the grant effect with the business cycle using NBER recessions, but found no significant effects.

away from one another to also be competing for investment from the same VC firms. Therefore, if the grant causes reallocation, I should observe a larger treatment effect in competitions where winners and losers are from the same area. My first test identifies firms within competitions from the same metropolitan statistical area (MSA) and from different MSAs. The Phase 1 grant effect is slightly higher when competing firms are from the same MSA, at 11.9 pp compared to 9.9 pp (Table 1.7 columns VII and VIII). Column IX shows that the difference between these coefficients is insignificant.

In the geographical analysis (Appendix 1B), a second test examines specific within-region effects. I find that the grants are consistently most useful to firms in the San Francisco (SF) region, regardless of whether they are competing with firms locally or far away. Hochberg, Ljungqvist, and Lu (2007) also find that the benefits of early-stage resources are amplified in SF. Otherwise, the effect when competing firms are from the same MSA and when they are from different MSAs is not systematically different. Therefore, reallocation does not seem drive the main findings, although I cannot rule it out.

The grant effect is, however, systematically larger not just for firms from SF, but more broadly when the winner is located in a city with greater VC investment per unit of city output. I demonstrate this in Appendix 1B, the geographical analysis. The literature has found that firms, particularly startups, are *less* financially constrained in areas with deeper capital markets (Rajan and Zingales 1998, Berkowitz and White 2004). My other results point to the grant having a larger effect for firms that are *more* financially constrained. This is a puzzle.

One possible solution comes from Lerner (2000), who finds that SBIR awards stimulate firm growth only in regions with high venture investing, a more extreme result than mine. Lerner suggests that perhaps congressional efforts distort award allocation across regions. In Appendix 1C I use delegation congressional power in the House and Senate to predict spending to a jurisdiction. The regressions reveal a statistically significant positive effect of seniority on committees with relevant authority in both chambers. However, the effect is very small, which is not surprising since these awards are small, dispersed, and bureaucratized. While its direction supports Lerner's hypothesis, it seems unable to explain the much larger grant effect in cities with greater VC intensity. Lerner also hypothesizes that long-lived research firms, which win many awards and do not seek VC finance, could be disproportionately located in areas without high venture activity. This is *not* the case in my data. The correlation of all-government SBIR awards (i.e. the degree to which a firm is an "SBIR mill") and local VC intensity is 0.01. Of the 59 firms with at least 50 all-government SBIR awards, 20% are in Boston, 10% are in LA, and 11% are in SF.

What, then, explains the regional variation? Larger knowledge spillovers may play a role. High-tech employees in Silicon Valley exhibit extreme inter-firm labor mobility (Saxenian 1994, Fallick, Fleischman and Rebitzer 2006). Rapid job-hopping can increase agglomeration economies, but it imposes costs on employers who must invest in - and expose trade secrets to - fleeting human capital. Greater spillovers from R&D investment in high-tech clusters could make the grant more valuable for startups in these areas. A second factor could be that regions with high VC per unit output have more intense competition for venture finance.

Phase 2 Grant Impact on VC

Roughly a year after receiving a $150,000 Phase 1 award, a firm may apply for a $1 million Phase 2 grant. Successful applicants typically receive their Phase 2 money nearly two years after the Phase 1 award. In Appendix 1E, I analyze the Phase 2 grant effect in depth. Here, I summarize my results and their policy relevance.

The Phase 2 grant has no consistently positive effect on subsequent VC. RD estimations using the DOE ranking of Phase 2 applicants (a subset of Phase 1 winners) produce small, positive, but imprecise coefficients. When I jointly estimate the Phase 1 and 2 effects, shown in Table 1.10, I find the same robust Phase 1 effects, but coefficients on Phase 2 range from -4.2 pp to -0.003 pp. These coefficients have only slightly smaller standard errors than when I estimate Phase 2 alone. While Phase 2 may be useful for some firms, it is not for others. The true average effect is almost certainly smaller than Phase 1, if not negative. I find no heterogeneity across firm age, sector or over time; the coefficients are always small or negative, and insignificant.

One reason for this Phase 2 finding is adverse selection among Phase 1 winners in the decision to apply to Phase 2. Among Phase 1 winners, 37% *did not apply for Phase 2.* Of these non-applyers, 19% received VC investment within two years of their initial award. This is only 9% for firms who applied and lost Phase 2, and 8% for firms who applied and won. From a different angle, 55% of firms who receive VC within two years of the Phase 1 grant do not apply for Phase 2. Apparently, firms do not apply for Phase 1 - and VC firms do not fund Phase 1 winners - because of the Phase 2 expected value.

In interviews, grantees told me that the grant application and reporting processes are so onerous that once they receive external private finance, it is often not worthwhile to apply for additional government funding. Similarly, Gans and Stern (2003) hypothesize that private funding is preferred to SBIR funding. Startup Oscilla Power, introduced above, did win a Phase 2 grant. CEO Shendure said that the $1 million was significant relative to what the firm sought to raise from private sources. Had Oscilla raised a $10 million VC round, he added, applying to Phase 2 may not have been worthwhile.

Table 1.10: Impact of Phase 1 and Phase 2 Grants on VC

Dependent Variable : VC_i^{Post}				
Bandwidth:	I. 1	II. 2	III. 3	IV. all
$\mathbf{1} \mid R_i^{Ph1} > 0$	**.099*** **	**.1*** **	**.11*** **	**.11*** **
	(.034)	(.027)	(.027)	(.025)
$\mathbf{1} \mid R_i^{Ph2} > 0$	**-.003**	**-.042**	**-.032**	**-.017**
	(.078)	(.054)	(.048)	(.043)
VC_i^{Prev}	.27***	.32***	.31***	.32***
	(.057)	(.038)	(.036)	(.029)
$\#SBIR_i^{\text{Prev}}$	.0012***	.001***	.0011***	.00087***
	(.00034)	(.00029)	(.00027)	(.00024)
Competition f.e.	Y	Y	Y	Y
N	1872	2835	3367	5021
R^2	.47	.39	.35	.27

Note: This table reports regression estimates of the Phase 1 ($\mathbf{1} \mid R_i^{Ph1} > 0$) and Phase 2 grant ($\mathbf{1} \mid R_i^{Ph2} > 0$) effects on subsequent VC. The dependent variable $VC^{\text{Post}}{}_i$ is 1 if the company ever received VC after the award decision, and 0 if not. The specifications are variants of the model in Equation 1, but with an additional indicator that is 1 if the firm won Phase 2, and 0 if it did not or did not apply. Standard errors are robust and clustered at the topic-year level. *** $p < .01$. Year$\geq$ 1995

Extremely high discount rates could help explain why firms do not find applying to Phase 2 worthwhile. It may be that the value of the time required to apply exceeds the expected value of the $1 million Phase 2 grant. Note that roughly 40% of Phase 2 applicants win, and the Phase 2 money is split into two equal disbursements, one in the following year, and one two years after applying. At the seed stage, a VC's required rate of return is typically at least 50%, and as high as 80% (Sahlman and Scherlis 2009). If the entrepreneur uses an 80% discount rate to value his time, then if the application cost exceeds $172,000, it would not be worthwhile to apply.[41] My interviews

[41]Discounted Present Value= $.4\left[\frac{500{,}000}{(1+\delta)^1} + \frac{500{,}000}{(1+\delta)^2}\right]$

suggest that the application cost solely in employee time is one to two full months, apart from any consulting or legal costs the firm may incur. High-tech, early stage startups often place a very high value on their time, so in conjunction with a high discount rate, it is plausible that among Phase 1 winners, high quality startups seeking venture finance tend not to apply for Phase 2.

The SBIR program spends vastly more on Phase 2 than Phase 1, so the absence of a strong Phase 2 effect is worrisome from a policy perspective. At the high end of the confidence intervals, the impact of Phase 2 is still much weaker per public dollar than Phase 1. For example, suppose that the true effect of Phase 2 on the likelihood of subsequent VC is 12 pp, which is the highest end of the estimates' 95% confidence intervals. Then the effect of Phase 1 per grant dollar is six times that of Phase 2. Consider the following thought experiment. In 2012 DOE spent $111.9 million on 111 Phase 2 grants and $38.3 million on 257 Phase 1 grants. If all the Phase 2 money were reallocated to Phase 1, DOE could have provided 750 additional firms with Phase 1 grants, increasing by a factor of at least 2.5 the program's impact on the probability of additional VC funding. To test the hypothesis that the program could achieve better outcomes by reallocated Phase 2 funds to additional Phase 1 grants, it would be necessary to experiment with removing Phase 2 from randomly selected competitions and observing differences in applicant types and subsequent investment. If the right to apply to Phase 2 has option value that affects the Phase 1 application decision, then the Phase 1 effect might well change for the worse.

1.4.2 The Grant Impact on Patents and Patent Citations

I now turn to the grant's impact on real outcomes, starting with the best available proxy for innovation: patenting. Patents are only one way that firms protect IP, and they have an ambiguous relationship with technological progress (e.g. Arora, Ceccagnoli and Cohen 2008, Cohen, Nelson and Walsh 2000). Nonetheless, they are positively associated with economic value creation and stock market returns (Hall, Jaffe, and Trajtenberg 2005, Eaton and Kortum 1999). As explained in Section 1.2.4, I use raw patent counts to measure the quantity of innovation and a normalized 3-year forward citation metric to measure the quality.

Table 1.11: Impact of Phase 1 Grant on 3-year Patenting (Negative Binomial)

Dependent Variable: $\#Patent_i^{\text{3 yrs Post}}$							
Bandwidth:	1	2		3		All	
	I.	II.	III.	IV.	V.	VI.	VII.
$\mathbf{1} \mid R_i > 0$	**1.03*** **	**1.18*** **	**1.07*** **	**1.4*** **	**1.0*** **	**2*** **	**1.1*** **
	(0.17)	(0.14)	(0.25)	(0.13)	(0.210)	(.16)	(.21)
$\#Patent_i^{\text{Prev}}$	0.16***	0.11***	0.11***	0.112***	0.11***	.14***	.13***
	(0.042)	(0.019)	(0.019)	(0.02)	(0.02)	(.018)	(.017)
VC_i^{Prev}	1.22***	1.38***	1.36***	1.34***	1.33***	1.3***	1.1***
	(0.25)	(0.17)	(0.18)	(0.17)	(0.17)	(.16)	(.15)
$\#SBIR_i^{\text{Prev}}$	0.0094***	0.011***	0.011***	0.011***	0.011***	.011***	.011***
	(0.0023)	(0.0015)	(0.0015)	(0.0016)	(0.0016)	(.0015)	(.0015)
R_i			0.044		0.018		.19***
			(0.083)		(0.0873)		(.054)
R_i^2					0.06*		-.0054
					(0.034)		(.0041)
Topic f.e.	Y	Y	Y	Y	Y	Y	Y
N	1872	2836	2836	3368	3368	5021	5021
Pseudo-R^2	0.21	0.183	0.18	0.16	0.16	.16	.16
Log likelihood	-1351.7	-2054.8	-2054.7	-2421.9	-2419.3	-3219	-3208

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on patents. The mean number of patents within three years after the grant is 0.79; among losers it is 0.57, and among winners it is 2.2 (bandwidth=all specification). The specifications are variants of the model in Equation 1. The dependent variable $\#Patent_i^{\text{3 yrs Post}}$ is the number of successful patents that the firm applied for within three years of the grant award. Specifications vary the bandwidth around the cutoff and control for rank linearly and quadratically. Topic fixed effects are a higher level than competition to achieve convergence of the maximum likelihood function, but still within-year. Standard errors are robust. *** $p < .01$. Year$\geq$ 1995

A Phase 1 grant leads to at least one additional patent within three years of the Phase 1 award, depicted in Figure 1.3 B.[42] The mean number of patents within three years of the grant is 0.79; among losers it is 0.57, and among winners it is 2.2 (with the bandwidth=All specification). Table 1.11 reports the results of negative binomial regressions with quadratic rank controls.[43] The

[42] I find no statistically significant effect of the grant on long-term patenting (all subsequent patents).

[43] For patenting, the Pearson goodness-of-fit χ^2 suggests that the data are excessively dispersed for the Poisson regression model, so I rely on the negative binomial distribution. I also tried log transformations of the patent and citation metrics, as well as a binary variable for positive patenting/citations. The former provided a similar effect to that shown here, and the latter did not yield effects with statistical significance.

table reports Poisson coefficients, but in the text I exponentiate to give incident rate ratios (IRR).[44] The award causes 2.7-2.9 times more patents at bandwidths of one, two and three firms around the cutoff, a large effect. The sample mean is 0.92 patents. My preferred specification is an IRR of 2.7 (columns I and V). There is no information in rank about subsequent patenting, but in contrast to the earlier results the coefficients on treatment decline somewhat when I remove rank controls (columns II, IV and VI).

Two issues with this result bear mention. The literature finds investment in R&D and patenting to occur simultaneously (Pakes 1985, Hall, Griliches and Hausman 1986; Gurmu and Pérez-Sebastián 2008). However, in my setting firms might plausibly conduct the key research prior to the award and file patent applications after winning. Second, the result becomes less consistent when the control function is estimated separately around the cutoff (Appendix 1G Table 20).

To evaluate the impact on patent citations I use a two-part model, because it would be incorrect to assume normality of the errors for semicontinuous data (Duan et al. 1983, Mullahy 1986).[45] I find no short or long term effect of the Phase 1 grant on the citation metric.

Heterogeneity in the Effect Across Firm Characteristics

Young firms have fewer internal resources and their R&D investment is likely more affected by capital market imperfections (Hall 2008). Columns I-IV of Table 1.12 show that the grant effect on short-term patenting falls dramatically and loses all significance for older firms. The IRR is a staggering 12 for firms no more than two years old, significant at the 1% level (column I), whereas the IRR is only 1.75 for firms more than two years old, and is highly imprecise. For firms less than 10 years old, the IRR is 4.5, whereas for firms older than 10, it is 0.62 - a negative effect - and insignificant (columns III and IV). To my knowledge this is the first direct empirical evidence that young privately held firms face greater R&D investment financing constraints than older private

[44]Poisson regression models the log of the expected count. Coefficients indicate, for a one unit change in the covariate, the difference in the logs of expected counts. If λ is the Poisson rate (the number of patents), the model is $\log(\lambda) = \alpha + \tau\left[\mathbf{1} \mid R_{ic} > 0\right]$, where covariates other than treatment are omitted. We can write $\tau = \log(\lambda_{R_{ic}>0}) - \log(\lambda_{R_{ic}<0}) = \log\left(\frac{\lambda_{R_{ic}>0}}{\lambda_{R_{ic}<0}}\right)$. Exponentiating the coefficient τ gives the incidence rate ratio (IRR). (This term comes from interpreting the patent count as a rate.) The IRR tells us how many times more patents awardees are expected to have compared to losers.

[45]The first stage models zero versus positive citations (I use logit), and a second stage models observations with positive citations linearly assuming a log-normal distribution for the citations. The two-part model is preferred to the Tobit model, in which the same stochastic process arbitrarily censored from below determines both zero and the positive outcomes. The Tobit model nonetheless gives similar qualitative results.

firms, supporting the findings on public firms in Brown, Fazzari and Petersen (2009).

Table 1.12: Impact of Phase 1 Grant on 3-year Patenting by Firm Age, Technology Propensity to Patent, and Number of Previous Patents (Negative Binomial)

Dependent Variable: $\#Patent_i^{3\text{ yrs Post}}$	Firm Age in Years				Firm # Previous Patents		Tech. Patent Propensity	
	I. ≤ 2	II. > 2	III. ≤ 9	IV. > 9	V. 0	VI. ≥ 1	VII. High	VIII. Low
$\mathbf{1} \mid R_i > 0$	**2.5*****	**.56**	**1.5*****	**-.48**	**1.2*****	**1*****	**2.1*****	**.99*****
	(.38)	(.42)	(.28)	(1.1)	(.39)	(.23)	(.46)	(.22)
$\#Patent_i^{\text{Prev}}$	.21	.13***	.16***	.12***			.32***	.3***
	(.16)	(.023)	(.049)	(.022)			(.077)	(.046)
VC_i^{Prev}	1.8***	1.1***	1.5***	.73***	2***	1.1***	.59	1.4***
	(.39)	(.2)	(.25)	(.24)	(.37)	(.18)	(.37)	(.19)
$\#SBIR_i^{\text{Prev}}$	.0073	.0097***	.012***	.011***	.017***	.0051***	.011***	.01***
	(.0083)	(.0017)	(.0031)	(.0019)	(.0063)	(.00088)	(.0043)	(.002)
R_i	-.15**	.43*	.0059	1	.14	-.022	.14	-.17*
	(.072)	(.22)	(.075)	(.68)	(.11)	(.082)	(.17)	(.094)
R_i^2	-.072	-.081	-.016	-.24	.047	-.014	-.046	.14***
	(.046)	(.064)	(.034)	(.18)	(.055)	(.033)	(.061)	(.04)
Topic f.e.	N	N	N	N	N	N	Y	Y
Year f.e.	Y	Y	Y	Y	Y	Y	N	N
N	576	2790	1410	1958	2308	1058	834	2532
Pseudo-R^2	.14	.092	.1	.1	.083	.067	.15	.2
Log likelihood	-383	-2221	-1220	-1367	-794	-1646	-719	-1640

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on patents using BW=3. The specifications are variants of the model in Equation 1. The dependent variable $\#Patent_i^{3\text{ yrs Post}}$ is the number of successful patents that the firm applied for within three years of the grant award. The left panel divides the sample by an indicator for high propensity to patent, which is 1 if the firm's technology sub-sector is Smart Grid, Sensors & Power Converters, Advanced Materials, Solar, or Batteries. The middle panel divides the sample by firm age, and the right panel by the firm's number of patents prior to applying for the grant. For all three, I could not estimate difference equations due to non-convergence of the Poisson maximum likelihood. Standard errors are robust. *** $p < .01$. Year$\geq$ 1995

As with age, we might think there is more information available about firms with patents. Hsu and Ziedonis (2008) and Conti, Thursby and Thursby (2013) show that patents improve entrepreneurs' access to finance by signaling potential investors about a firm's quality. Patents may also serve as collateral, as in Mann (2014) and Hochberg, Serrano and Ziedonis (2014). The latter paper finds that among VC-backed startups with available patents available, 36% used the patents to secure loans. Columns V and VI Table 1.12 shows that the treatment effect declines when firms have previous patents: with no patents, the grant leads a firm to produce 3.3 times more patents than it would otherwise, significant at the 1% level. With at least one patent, the IRR is 2.7. More experienced, later stage firms who may have better access to debt finance seem to benefit less from the grants.

Last, the noisiness in the patent data (note the large confidence intervals in Figures 1.3 A and B) may reflect the wide variation in propensity to patent across technologies (Scherer 1983, Brouwer and Kleinknecht 1999). I create an indicator for high propensity to patent from the USPTO (2012) patent intensity estimations.[46] In high propensity industries, a grantee produces 8.1 times as many patents as a loser, significant at the 1% level (Table 1.12 column VII). In contrast, the IRR is only 2.7, significant at the 10% level, in low propensity industries (Table 1.12 column VIII).[47]

Phase 2 Grant Impact on Patents

In contrast to the financing results, I do find a positive effect of the Phase 2 grant on patenting *and* patent citations. The IRR for the Phase 2 effect on the number of patents is 1.5, half the Phase 1 effect (and thus much smaller on a per grant dollar basis). The average patents for this sample is 2.2. The two-part model for citations finds that the odds of positive citations for Phase 2 grantees are 85% higher than the odds for non-grantees.[48] The sample mean probability of positive subsequent citations is 0.31, so the odds (probability of positive citations divided by probability of no citations) are 0.44. The second stage, a regression within observations with positive citations, finds small and insignificant coefficients. For tables, see Appendix 1E.

The Phase 2 grant acts on the extensive margin of innovation quality, but not the intensive margin. I also find that among firms with at least one previous DOE SBIR win, the Phase 2 grant has no measurable effect on either patents or citations. A policy implication is that if the government's objective is to generate R&D, measured by patents and more highly cited patents, then Phase 2 awards are beneficial when awarded to firms without previous patenting or citation

[46]These are based on patents per 1,000 jobs in an industry. The indicator takes a value of 1 if the firm is in one of the following sectors: Smart Grid, Sensors & Power Converters, Advanced Materials, Solar, or Batteries, and 0 otherwise.

[47]For all three heterogeneity analyses in Table 1.12, I am unable to estimate difference equations due to non-convergence of the maximum likelihood function. Similarly, I cannot separately estimate regressions for each technology (sub-sector) because the sample sizes are too small for the negative binomial model.

[48]Logit coefficients give the change in the log odds of the outcome for a one unit increase in the predictor variable. This odds ratio is calculates as $OR = e^{\beta}$, where β is the logit coefficient.

histories.

Relationship of VC Finance to Patents

In light of the literature on the benefits of VC finance, I am not surprised to see a large positive coefficient on previous VC finance in the regressions with patents as the dependent variable. I explore the relationship between VC and patenting further using subsets of the data unaffected by the grant: firms prior to application, and firms that lose Phase 1. I find that VC finance is associated in both groups with more patents and higher quality patents, shown in the top panel of Table 1.13. For example, prior to the grant application firms with VC finance have 2.6 times as many patents as firms without VC finance (column I). With citations, I find the inverse of the Phase 2 effect. Along the extensive margin, the odds of having positive citations is just slightly larger if a firm has VC finance (logit in column IIa). The regression part reveals that conditional on having patent citations, VC financing increases by 12 the number of citations (relative to a mean of 11.8). I observe essentially the same pattern when I consider only Phase 1 losers, in columns III and IV.

This positive impact of VC on patents raises the concern that the estimated grant effect on patents may indirectly capture VC investment after the award. Rather than the grant funding useful R&D work, the grant might simply enable VC finance, which in turn leads to patents. However, I find that among firms with no VC investment prior to their grant application and no VC investment within three years of applying, the grant effect on patents within three years remains large and robust (bottom panel of Table 1.13). So the grant and VC finance *both* induce patents.

Thus the patent estimates imply quick use of plausibly exogenous cash for R&D, offering an alternative to the corporate finance estimates of R&D sensitivity to cash flow shocks. The ideal experiment observes whether firms invest exogenous cash in R&D, in which case costly external finance must have prevented the firm from exploiting existing profitable investment opportunities. Empirical work typically uses investment demand equations with adjustment costs, and although studies have established that R&D is rarely financed with debt, it has been difficult to definitively identify that financial constraints cause R&D cash flow sensitivity (see Hall 2010). Here I find that profitable R&D investment would not occur in the absence of a subsidy, contributing to the body of work arguing that financial constraints inhibit investment, especially for smaller firms, such as Li (2011), Faulkender and Petersen (2012), and Zwick and Mahon (2014).

Table 1.13: Relationship between VC Finance and Patenting/Citation Outcomes

Panel A: Impact of VC on Patents & Citations Prior to Applying and Among Phase 1 Losers

	All Applicants			Losers only		
Dependent Variable:	I. $\#Patent_i^{\text{Prev}}$	II. $Citation_i^{\text{Prev}}$		III. $\#Patent_i^{\text{3 yrs Post}}$	IV. $Citation_i^{\text{3 yrs Post}}$	
		IIa. Logit	IIb. Regress		IVa. Logit	IVb. Regress
VC_i^{Prev}	**.96*** **	**1.005*** **	**12.04*** **	**1.31*** **	**.78*** **	**21.66*** **
	(.12)	(.11)	(4.52)	(.16)	(.18)	(6.37)
Year f.e.	Y	Y	Y	Y	Y	Y
Sector f.e.	Y	Y	Y	Y	Y	Y
N	6324	6322	6322	5042	4677	4677
R^2			.06			.14
Pseudo-R^2	.016	.055		.094	.19	
Log lik.	-8390.7	-10101.4	-10101.4	-5098.0	-4840.1	-4840.1

Panel B: Impact of Grant on Patents for Firms with no VC before or within 3 Yrs of Applying

Dependent Variable: $\#Patent_i^{\text{3 yrs Post}}$

	V. BW=1	VI. BW=2	VII. BW=3	VIII. BW=all
$\mathbf{1} \mid R_i > 0$	**.89*** **	**.57** **	**.84*** **	**1.12*** **
	(.18)	(.29)	(.25)	(.26)
Year f.e.	Y	Y	Y	Y
Sector f.e.	Y	Y	Y	Y
N	1644	2482	2952	4424
Pseudo-R^2	.063	.064	.059	.056
Log lik.	-1248.3	-1833.8	-2129.7	-2851.1

Note: This table reports regression estimates of the relationship between VC funding and patenting/citation outcomes for Phase 1 applicants. The top panel estimates the impact of having VC finance prior to applying for the grant (VC_i^{Prev}) on outcomes. Columns I and II consider only events prior to application. Columns III and IV limit the sample to firms who applied for an SBIR and lost. For patents, I use the negative binomial model as in previous regressions. For citations I use the two-part (logit plus regression). The logit portion of estimates zero vs. positive citations (extensive margin), and then the regress part estimates the impact of the grant on observations with positive citations (intensive margin). The bottom panel estimates the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on patents as in Table 1.11, but includes only firms that did not previously receive VC prior to application, nor received VC finance within three years of application. Covariates omitted for brevity. Standard errors are robust. *** $p < .01$. Year$\geq$ 1995

1.4.3 The Grant Impact on Revenue, Survival & Exit

My final outcome metrics are binary variables for achieving revenue, survival, and exit (IPO or acquisition). As with financing, I find that rank has no predictive power over revenue, survival, or exit, so my preferred specifications, reported in Table 1.14, omit rank controls.[49]

Visual evidence for an increase in commercialization probability around the cutoff is in Figure 1.5, and the top panel of Table 1.14 shows the regression results. A Phase 1 grant increases a firm's probability of commercialization by roughly 11 pp, from around 52% to 63%. Unfortunately, I cannot center the commercialization variable around the application date, so a firm may have reached revenue before it applied. However, if the assumptions underlying the RD are sound, this probability should be the same for firms on either side of the cutoff. The magnitude of the estimated effect is not interpretable as a direct grant effect, but offers insight into whether there is an impact.

The majority of firms survive through 2014, depicted in Figure 1.6. Only about 23% were discovered to be out of business, bankrupt, or acquired. This is, however, likely a very conservative measure given the limitations of manual web scraping. Visually, there is a decline in the survival probability for losers as the cutoff approaches, and then a jump from around 70% to 85% survival. The regression results (middle panel of Table 1.14) yield coefficients of about 4 pp, but they are imprecise. When I add rank controls (Appendix 1G Tables 15-16), the coefficients further lose significance. I conclude that I cannot measure an effect on survival.

VC investors typically liquidate successful investments through an IPO or acquisition. The regression results in the bottom panel of Table 1.14 find a strong statistical impact of 3.3-4 pp. This is a dramatic increase in the probability of acquisition or IPO from roughly 4% to 7.5%, but it should be interpreted with some caution in light of visual inconsistency. Figure 1.4 B suggests there may be an effect of the grant on exit probability, but it disappears for firms with $R_i = 2$. As with financing, I find no effect of the Phase 2 grant on revenue, survival or exit (see Appendix 1E for results).

[49]The G-value from the goodness-of-fit test with no control for rank is 0.0001, orders of magnitude less than the critical value of 1.47 with 5% confidence. Appendix 1H Table 3 suggests that there are no major discontinuities besides the award cutoff. Specifications with rank controls are in Appendix 1G Tables 13-18.

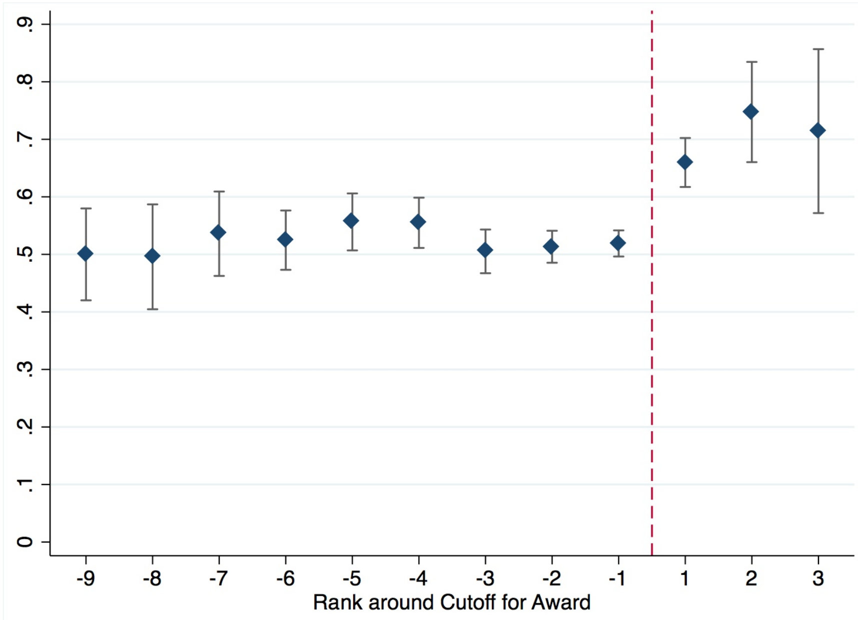

Figure 1.5: Probability of Achieving Revenue (Commercialization) by Rank
Note: This figure shows the fraction of applicants who achieved revenue. The applicants are binned by their DOE assigned rank, which I have centered so that Rank > 0 indicates a firm won an award. This variable is not dated, so I do not know if the firm achieved revenue before or after the grant. Capped lines indicate 95% confidence intervals. N=4,816.

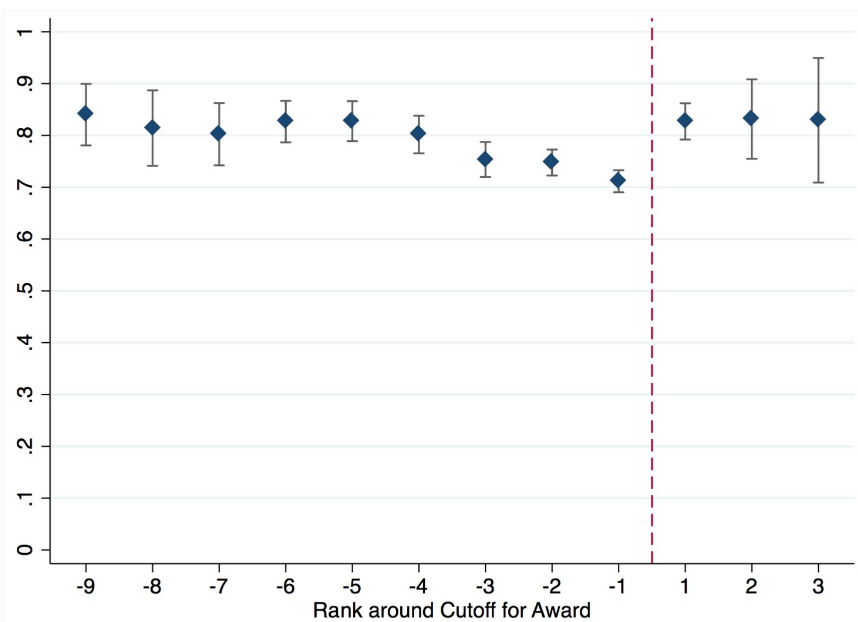

Figure 1.6: Probability of Survival After Grant Decision by Rank
Note: This figure shows the fraction of applicants who survived (as of May 2014) after the Phase 1 grant award decision. The applicants are binned by their DOE assigned rank, which I have centered so that Rank > 0 indicates a firm won an award. Capped lines indicate 95% confidence intervals. N=4,816.

Table 1.14: Impact of Phase 1 Grant on Firm Revenue, Survival and Exit

Dependent Variable: $Revenue_i$				
	I. BW=1	II. BW=2	III. BW=3	IV. BW=all
$\mathbf{1} \mid R_i > 0$	**.11*** **	**.09*** **	**.1*** **	**.12*** **
	(.038)	(.03)	(.028)	(.025)
VC_i^{Prev}	.17***	.17***	.18***	.23***
	(.05)	(.038)	(.033)	(.024)
$\#SBIR_i^{\text{Prev}}$	.0017***	.0017***	.0018***	.002***
	(.00028)	(.00022)	(.00022)	(.00019)
Competition f.e.	Y	Y	Y	Y
N	1872	2836	3368	4812
R^2	.41	.33	.3	.23
Dependent Variable: $Survival_i$				
	I. BW=1	II. BW=2	III. BW=3	IV. BW=all
$\mathbf{1} \mid R_i > 0$	**.072** **	**.046* **	**.039**	**.046** **
	(.036)	(.026)	(.024)	(.021)
VC_i^{Prev}	.086*	.11***	.096***	.1***
	(.047)	(.03)	(.028)	(.02)
$\#SBIR_i^{\text{Prev}}$	.00071***	.00072***	.00078***	.00079***
	(.00025)	(.00019)	(.00016)	(.00014)
Competition f.e.	Y	Y	Y	Y
N	1750	2660	3160	4533
R^2	.39	.32	.28	.23
Dependent Variable: $Exit_i^{\text{Post}}$				
	I. BW=1	II. BW=2	III. BW=3	IV. BW=all
$\mathbf{1} \mid R_i > 0$	**.044* **	**.033* **	**.041*** **	**.034*** **
	(.025)	(.017)	(.015)	(.012)
$Exit_i^{\text{Prev}}$	-.1***	-.099***	-.094***	-.084***
	(.039)	(.023)	(.018)	(.012)
VC_i^{Prev}	.14***	.12***	.13***	.13***
	(.043)	(.029)	(.025)	(.019)
$\#SBIR_i^{\text{Prev}}$	.00074**	.0007***	.00056***	.0003*
	(.0003)	(.00022)	(.00021)	(.00016)
Competition f.e..	Y	Y	Y	Y
N	1872	2836	3368	5021
R^2	.41	.31	.26	.18

Note: This table reports regression estimates of the effect of the Phase 1 grant ($\mathbf{1} \mid R_i > 0$) on revenue, survival, and exit with no rank controls (variants of Equation 1). Top: dep. var. is 1 if the firm ever reached revenue (not centered around award time). Middle: the dep. var. is 1 if the firm active as of May, 2014. Bottom: dep. var. is 1 if the firm experienced IPO/acq after award. Std errors robust and clustered by topic-year. *** $p < .01$. Year$\geq$ 1995.

1.5 How Does the Grant Affect Investor Decisions?

DOE SBIR grants positively impact a range of relevant outcomes. This fact, established in Section 1.4, is relevant to policy regardless of the mechanism. Yet understanding the source of the large effect is interesting and important. In this section I explore how the grants affect investor decisions, which also helps explain the real impacts.

The most obvious explanation for the Phase 1 grant's effect on VC investment is *certification*: the government's willingness to invest conveys positive information to venture capitalists that the firm has a promising technology. Thirty interviews I conducted with venture investors, mostly in 2013, consistently rebutted this hypothesis. The investors included experienced angels, partners at conventional VC firms, and leaders of corporate ("strategic") VC groups. Nearly all believe that while an SBIR grant can help a firm advance to an investment-grade stage, the grant itself has little informational value. "SBIRs have no signal value," Matthew Nordan, then a Vice President at Venrock, said. "We don't care - they're completely immaterial. The only time we would care is when it gives the company time to do proof-of-concept." Investors like Rachel Sheinbein, then a CMEA Capital partner, and Andrew Garman, Managing Partner at New Venture Partners, conveyed similar opinions.[50] The startups I spoke with also did not think the grants signaled the value of their technology.

With this field evidence in mind, I present a simple model in Section 1.5.1 containing the mechanisms that might explain the grants' impact on external investment. In Section 1.5.2 and 1.5.3 I discuss which channel is most likely in light of my empirical evidence.

1.5.1 A Signal Extraction Model

I consider the grant's effect on investor decision-making through the lens of a signal extraction problem, drawing from Phelps (1972) and Aigner and Cain (1977). Here I summarize the model and describe the hypotheses; the full model is in Appendix 1A.

Whether a technology proposal will work in practice is often inherently uncertain. Layered on the entrepreneur's own uncertainty are information asymmetries between the entrepreneur and

[50] For example, Sheinbein said: "Nothing about government due diligence is informative...They're more in business of fear." A few angel and strategic investors, notably Mitch Tyson, Partner at Clean Energy Venture Group, and Steve Taub, then Senior Investment Director for Energy at GE Ventures, said that there is a small positive signal in the grant about the technology.

potential investors (Gompers and Lerner 1999). Venture investors rely on noisy signals and heuristics to choose a few firms quickly out of hundreds of proposals (Metrick 2007, Kirsch, Goldfarb and Gira 2009). I do not portray this complicated process here, but seek to distill the key elements that are relevant to my reduced form evidence.

A grant might alleviate financial constraints for recipient firms through either (1) *certification*; or (2) *funding.* Certification is when informational content in the grant decision alleviates information asymmetries, and it requires DOE to identify or be perceived to identify better firms. The second channel is the money itself, which has two subcategories: (2a) *equity* and (2b) *prototyping*. In the former, the grant allows the entrepreneur to retain more equity, which reduces financial frictions. Without the grant, an investor might have to take such a large stake in the firm that maintaining entrepreneurial incentives would be impossible. The latter channel is prototyping, where grantees demonstrate their technology's viability by investing in proof-of-concept work. Prototyping reduces uncertainty about the technology, which can alleviate information asymmetry (a financial friction), or simply decrease the project's risk. Certification is proposed as a possible mechanism in Lerner (2000), as well as in other studies. I have devised the funding effect and its two channels to suit the present setting.

I begin with in the no-grant case. Let each startup have a uni-dimensional technology quality signal $T_i = \bar{t} + \tau_i$, where T is normally distributed with mean $\bar{t}$ and variance σ_T^2. Suppose there is a single venture capitalist. He forms rational expectations and is more likely to invest in firms with high expected technology qualities. The investor knows the T distribution but receives only a noisy signal from each startup $\tilde{T}_i = \bar{t} + \tau_i + \varepsilon_i$, where ε is normally distributed with mean 0 and variance σ_ε^2. The investor calculates the expected technology quality given the signal, $\mathbf{E}\left(T_i \mid \tilde{T}_i\right)$, putting more weight on the signal $\tilde{T}_i$ if it is reliable - σ_ε^2 is small - and more weight on the mean $\bar{t}$ if σ_ε^2 is large. The optimal weight on the signal is $\frac{\sigma_T^2}{\sigma_\varepsilon^2+\sigma_T^2} = \alpha$, so the expected technology quality is:

$$\mathbf{E}\left(T_i \mid \tilde{T}_i\right) = (1-\alpha)\bar{t} + \alpha\tilde{T}_i \qquad (3)$$

The first term is a group effect and the second term is an individual effect. The line in Equation 3 is depicted in Figure 1.7 A. Note that α is the slope coefficient of a linear regression of T on $\tilde{T}$ and a constant.

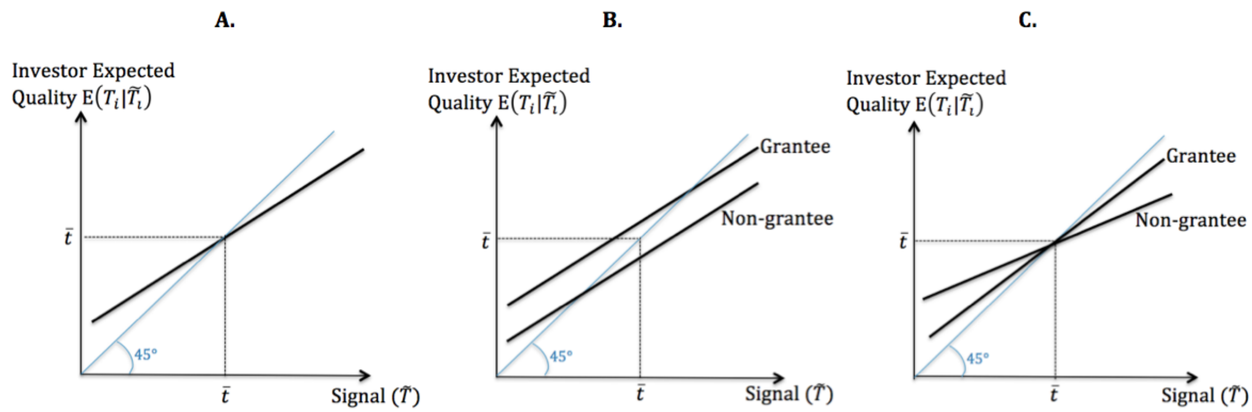

Figure 1.7: Possible grant effects on investor expected quality given firms' signal to investors
Note: Figure 1.7 A shows the investor's expected quality of the entrepreneur (y-axis) as a function of the noisy signal that the investor observes (x-axis). Figure 1.7 B shows that a certification or valuation effect increases the mean expected quality of grantees relative to non-grantees ($\bar{t}_g > \bar{t}_n$). Figure 1.7 C shows that a prototyping effect increases the slope of the grantee line relative to the non-grantee line. This occurs because the grant causes the grantee's signal to be more reliable, which for example may occur if prototyping decreases the variance of the noisy signal $\left(\sigma^2_{\varepsilon,g} < \sigma^2_{\varepsilon,n}\right)$.

The government also receives a signal about the firm, $\tilde{T}_i^G$, which neither the investor nor entrepreneurs observe.[51] The government awards grants to a subset of firms whose $\tilde{T}_i^G$ are located above a cutoff. Whether a firm has a grant (g) or does not (n) is a truncated dichotomous version of $\tilde{T}_i^G$. The investor observes this binary signal $x \in \{g, n)$.[52] The grant might affect the mean technology quality ($\bar{t}$), the quality variance (σ^2_T), and the signal variance (σ^2_ε). Any value of the grant money that is unrelated to its technology quality is μ_x, where $\mu_n = 0$ and $\mu_g \geq 0$. After the competition entrepreneurs have technology quality $T_{i,x} = \bar{t}_x + \mu_x + \tau_{i,x}$. Now $T_x \sim N\left(\bar{t}_x + \mu_x, \sigma^2_{T,x}\right)$, and the signal error becomes $\varepsilon_x \sim N\left(0, \sigma^2_{\varepsilon,x}\right)$.

Suppose two firms have the same noisy signal $\tilde{T}_i = \tilde{T}_j = k$, but one has a grant ($x = g$) and the other does not ($x = n$). The difference between their expected qualities, Equation 4, should reflect the grant.

$$\mathcal{D} = \mathbf{E}\left(T_i \mid \tilde{T}_i = k, x = g\right) - \mathbf{E}\left(T_j \mid \tilde{T}_j = k, x = n\right) \tag{4}$$

There are two broad mechanisms that might drive this difference away from zero:

[51] I need not make any functional form assumptions about $\tilde{T}_i^G$.

[52] The investor does not observe whether a non-grantee firm applied and lost or did not apply at all. The model is agnostic about whether the grant has a negative effect on losers (though this seems unlikely because the applicant firms form a small subset of the space of energy startups). In Section 1.4.1 I argue that negative spillovers seem absent.

1. *Certification Effect:* Suppose that the award process separates applicant firms into higher and lower technology quality types, but has no other effect. Now $\bar{t}_g > \bar{t}_n$, while $\mu_g = 0$, $\sigma^2_{T,g} = \sigma^2_{T,n} = \sigma^2_T$, and $\sigma^2_{\varepsilon,g} = \sigma^2_{\varepsilon,n} = \sigma^2_\varepsilon$. The difference in expected quality, shown in Figure 1.7 B, is:
$$\mathcal{D} = (\bar{t}_g - \bar{t}_n)\left(1 - \frac{\sigma^2_T}{\sigma^2_\varepsilon + \sigma^2_T}\right) \tag{5}$$

2. *Funding Effect*

   (a) *Equity Channel:* The grant increases the entrepreneur's internal resources, potentially making a VC deal tractable by allowing the entrepreneur to retain a larger share of the firm. This also manifests as a mean shifting effect for grantees, as the only difference between grantees and non-grantees is μ (Figure 1.7 B).
$$\mathcal{D} = \mu_g \left(1 - \frac{\sigma^2_T}{\sigma^2_\varepsilon + \sigma^2_T}\right) \tag{6}$$

   (b) *Prototyping Channel:* The award is invested in proof-of-concept work. This improves the signal's reliability (increasing α), which translates to a steeper line, shown in Figure 1.7 C. Grantees with above-average signals benefit from the slope change, and I assume that these high-type signal firms constitute the investor's consideration set. Prototyping occurs through increased signal precision, such that $\sigma^2_{\varepsilon,g} < \sigma^2_{\varepsilon,n}$.[53] With all else held the same, the difference is:
$$\mathcal{D} = \left(\bar{t} - \tilde{T}_k\right)\left(\frac{\sigma^2_T}{\sigma^2_{\varepsilon,n} + \sigma^2_T} - \frac{\sigma^2_T}{\sigma^2_{\varepsilon,g} + \sigma^2_T}\right) \tag{7}$$

Given these three possible mechanisms, I shift to the government perspective, and connect the model to the empirical design. Entrepreneurs have an ultimate observable quality T_i^O, which is a function of latent quality T_i and resources provided to the entrepreneur. Figure 1.8 shows the correlation of this outcome with the private government signal $\tilde{T}_i^G$. Applicant firms with $\tilde{T}_i^G$ to the right of the red cutoff line are awardees, while applicants to the left are losers. My regression

[53]See Appendix 1A for discussion of the alternative possibility for a higher α, which is when $\sigma^2_{T,g} > \sigma^2_{T,n}$.

discontinuity design approximates the difference in outcomes between two firms that present the same signal to the government, but where one has a grant and one does not. This is shown in Equation 8.

$$\mathcal{D} = \mathbf{E}\left(T_i^O \mid \tilde{T}_i^G = k, x = g\right) - \mathbf{E}\left(T_i^O \mid \tilde{T}_j^G = k, x = n\right) \tag{8}$$

First, when the grant has no effect, the observed outcome projected on the government signal is a horizontal line, and $\mathcal{D} = 0$. This is depicted in Figure 1.8 A. Second, if the signal is informative about outcomes, the regression line is upward sloping (Figure 1.8 B).[54] Here, the grant acts as a binary signal about firm quality, which the market learns is informative, so we observe a jump at the discontinuity due to certification ($\mathcal{D} > 0$). Investors are more likely to finance grantees because they have higher mean expected quality ($\bar{t}_g > \bar{t}_n$), even if the money itself has no effect. Figure 1.8 B, which describes actual investment outcomes as a function of the government signal, maps to Figure 1.7 B, which shows how the government signal affects investor beliefs.

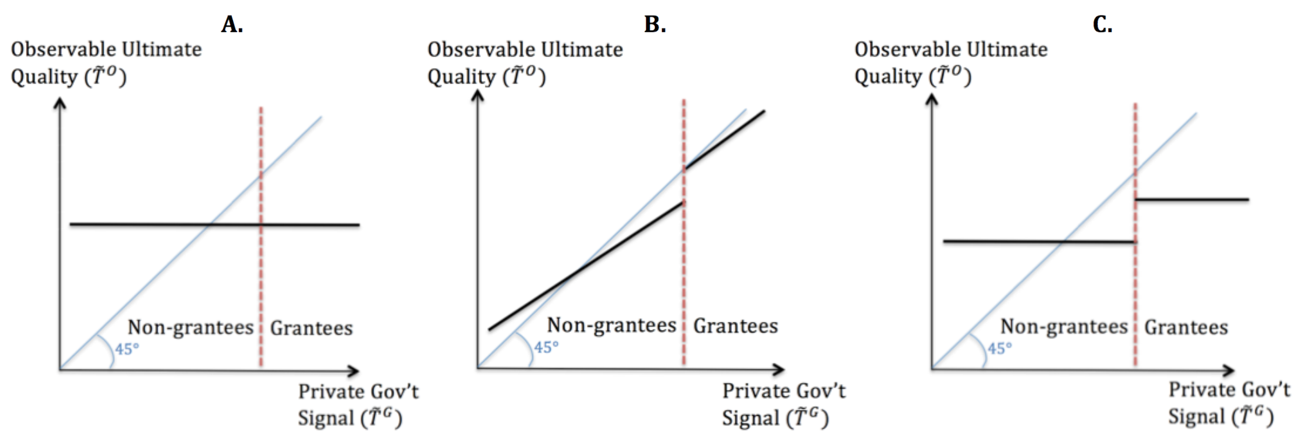

Figure 1.8: Possible grant effects on firm outcome given firms' private signal to government
Note: Figure 1.8 A shows this observable outcome (y-axis) as a function of the signal that the government receives from the firm, which is private to the government (x-axis). In this case, the government signal $\tilde{T}^G$ is wholly uninformative about outcomes, so the line is flat, and there can be no certification effect with rational investors. In Figure 1.8 A, there is both no certification effect and no effect of the grant money itself, so there is no jump at the discontinuity between non-grantees and grantees. Figure 1.8 B shows a prototyping or valuation effect increasing outcomes for grantees relative to non-grantees in the absence of certification ($\tilde{T}^G$ uninformative). Figure 1.8 C. shows the certification case, in which $\tilde{T}^G$ is informative and thus correlated with outcomes. In the absence of a valuation or prototyping effect, we nonetheless observe a jump at the discontinuity as the market accounts for information in the private government signal $\tilde{T}^G$.

[54]It is possible that the government signal is informative in the other direction; that is, it orders poor quality firms above higher quality firms on average. In this case the line will slope down, and we would expect a downward jump at the discontinuity.

Finally, if $\tilde{T}_i^G$ is uninformative but the grant money itself benefits recipients through either funding or prototyping, we observe a horizontal line with a jump at the discontinuity, shown in Figure 1.8 C. Because the funding channel is a mean-shifting effect ($\mu_g > 0$), it maps to Figure 1.7 B from the investor perspective. With only a prototyping channel, the government signal is uninformative (Figure 1.8 C), but prototyping changes the variance of the signal to investors and so maps to Figure 1.7 C.

1.5.2 Evaluating the Certification Hypothesis

Applicant ranks permit a test for the certification effect. Although the ranks are secret, the fact that rank maps directly to award means that investors should incorporate the grant as a positive signal *only* if DOE accurately ranks firms according to technological quality. This assertion requires rational investors. Irrational investors might consider DOE awards a valuable signal even if DOE has no ability to identify high quality firms.[55] Also, to contend that uninformative ranks reflect an absence of valuable information in the award, I need to establish that the centered ranks do not conceal information in raw rank, and that DOE program officials cannot predict the number of awards in a competition. This test, while requiring strong assumptions in this context, is novel and may prove useful in other settings as well.

As far as I can discern, neither of these issues are present. In Section 1.6, I show a variety of tests establishing that, controlling for winning, there is no information in raw rank regardless of the number of awards in the competition. My assertion that program officials are unsure of precisely the number of awards in any given competition is based on email correspondence included in the ranking data, and interviews at DOE with program officials who generate the ranks and SBIR office administrators. To the best of my knowledge, winning generates an effect, not rank nor the number of awards in the competition.

The actual probability of subsequent VC finance by rank depicted in Figure 1.2 B is most similar to Figure 1.8 C from the toy model - the slope of quality outcome (T_i^O) projected on the government signal ($\tilde{T}_i^G$) appears to be zero. The share of firms getting VC is flat in the DOE assigned rank, except immediately around the award cutoff. Ranks are also uninformative about the other outcome metrics. The ranks may reflect social benefits that I do not capture in my outcome metrics,

[55]See Baker and Wurgler (2011) on behavioral finance.

but I believe the ranks are essentially randomly assigned, particularly for the higher ranked firms immediately around the cutoff. In interviews, program officials told me that SBIR is a tax on their time; they view the grants as excessively small and a burdensome administrative duty imposed from outside. Their primary task is to provide the much larger university, national lab, and large firm grants, where each grant decision involves vastly more money than SBIR competitions. In any event, identifying high quality startups is no easy task (Kerr, Nanda and Rhodes-Kropf 2013). Regardless of the ranking process, the ranks appear to be pure noise from the investor perspective.

Phase 2 provides an additional argument against certification. DOE does a second round of selection to determine the Phase 2 winners, so with certification a Phase 2 grant should reveal further quality distinction. I observe no measurable Phase 2 effect on financing, suggesting that Phase 2 does not have a certification effect and therefore making it less likely that Phase 1 does. Although Phase 1 is more competitive, for the certification hypothesis to explain the Phase 1 effect would require us to assume that all Phase 1 winners are "good firms," or that the private sector believes there is something special about the Phase 1 decision.

If my understanding of the institutional setting is correct, and if we are willing to accept the rational expectations hypothesis for investors, then the grant - the public signal x - is likely pure noise. Although we cannot rule it out, certification alone seems incapable of explaining the discontinuity in the grant's effect on VC. This presents a puzzle, and we must turn to more subtle mechanisms.

1.5.3 Evaluating the Funding Hypothesis

If certification is not the main channel, then the money itself must be useful, either because it permits entrepreneurs to make deals with VC investors, or because entrepreneurs invest it in valuable R&D.

Equity

We might imagine a simple incentive constraint requiring the entrepreneur to retain a certain share of the firm, else agency problems become excessively severe. The Phase 1 grant is a positive wealth shock for the entrepreneur, and may render a deal tractable. The rapidity of the Phase 1 effect argues in favor of the equity channel; recall that two-thirds of the effect occurs within two years.[56]

[56]The valuation channel does not imply that the grant is a subsidy to VC firms. For example, if the

Yet compared to the average VC round size in my data of $9 million, the Phase 1 grant is quite small. It is hard to imagine that on average $150,000 can shift the startups in my sample from negative NPV to positive NPV simply by decreasing the required investor stake. Indeed, my Phase 2 evidence makes the equity channel less credible. If the cash acts by resolving the underinvestment problem, I should observe a much stronger effect of the $1 million Phase 2 grant. Yet the RD revealed that the Phase 2 effect is either negative or much smaller per grant dollar than the Phase 1 effect.

Also, we would expect that if Phase 1 enables access to VC finance because of the expected value of the Phase 2 effect, all Phase 1 awardees should apply for Phase 2. The revealed preference of awardees suggests that the cash as such is not critical. Surprisingly, 37% of Phase 1 winners opt not to apply for Phase 2. Also, the Phase 1 grant effect is much *stronger* for Phase 1 winners who choose not to apply or who lose Phase 2 than for the whole sample (see Appendix 1E). It seems unlikely that the large Phase 1 impact is purely an equity effect.

Prototyping

We are left with prototyping as the dominant channel for the Phase 1 effect on VC. The Phase 1 grant is supposed to fund the applicant's proposed small-scale testing or demonstration project. This may be how the money is actually used, on average, even though the government does not monitor expenditure. High-quality firms whose prototyping reveals positive information find it easier to secure an investor. That most of the Phase 1 impact on VC occurs within two years makes sense if the Phase 1 research is completed within the nine-month time frame set by the SBIR program.[57] By the Phase 2 stage there is sufficient information about the firms that Phase 2-funded work does not provide an incremental benefit.

Consistent with prototyping, the patent analysis finds that the grants fund valuable R&D in the short term. While both Phase 1 and and Phase 2 positively impact patents, only Phase 2

VC sector is competitive, the investor gets a break-even number of shares in the portfolio company. In equilibrium the grant causes the VC to get fewer shares, not a higher rate of return. The grant could also increase the entrepreneur's bargaining power.

[57]The VC deal flow analysis also suggests that VC availability is most relevant is at least six or eight months after the award. The grant effect is countercyclical with respect to deal flow in the two years after the grant (Section 1.4.1). Within only one year of the grant, I find a smaller and less significant effect. The grantee, under the prototyping hypothesis, must conduct its proof-of-concept work before it can effectively pitch to VCs.

impacts patent citations, which measure innovation quality. The proof-of-concept Phase 1 work does not seem to cause a change in the entrepreneur's technology quality (τ_i), while the larger Phase 2 project may do so. This story substantiates prototyping through signal precision, where $\sigma^2_{\varepsilon,g} < \sigma^2_{\varepsilon,n}$.[58] The Phase 1 grant funds valuable demonstration and testing on existing technology, alleviating uncertainty and potentially information asymmetry. Through this prototyping channel, the grant may reduce the cost of external finance.

Although I conclude that the evidence best supports protoyping, my data does not allow me to affirmatively identify one hypothesis and reject others. Furthermore, other stories are plausible. For example, the right to apply for Phase 2 could have option value in the event a firm does not get privately financed soon after Phase 1. That is, a firm whose prototyping does not yield a positive outcome, or who for some other reason fails to move forward, could then apply to Phase 2. A second example is that certification is plausible even with rational investors under, for example, a sunspot coordinating equilibrium.

1.6 Robustness Tests

This section addresses validity of the empirical results for the VC outcome. The appendices contain similar analyses for revenue, survival, exit, and patenting.

Five tests explore the issue of changing rank composition as I move away from the cutoff. First, Table 1.15 presents a regression in which variation in the number of awards across compositions identifies the grant effect on VC. I interact dummies for a firm's raw rank with dummies for the competition's number of awards. This estimates the treatment effect as, for example, the difference between a raw rank of two when there are two awards compared to one award in the competition (in the former case the firm is a winner, and in the latter a loser). It shows that the impact of raw rank does not change with the number of awards in the competition, and provides a stringent test of the conclusion that the treatment effect is explained by being above the cutoff, not rank.

Second, Appendix 1H Figures 4-6 show visual evidence that the discontinuity, and absence of information in rank, does not differ when I consider competitions with only one, two, and three awards. Third, Appendix 1H Table 1 separately considers competitions with only one, more than

[58] I expect only high-type signals enter the VC's consideration set, so $\sigma^2_{\varepsilon,g} < \sigma^2_{\varepsilon,n}$ leads the grant to have a positive impact on investment. If the full space is under consideration (perhaps some low-technology types have excellent business plans) then the grant may have no impact. The other avenue to a steeper regression line for grantees is to move their technology quality (τ_i) away from the mean, so that $\sigma^2_{T,g} > \sigma^2_{T,n}$.

one, two, three, and more than three awards. The results are consistent with the main specifications. Fourth, the last three columns of Appendix 1H Table 1 also show that the control functions do not differ by the cutoff; the coefficients on R_i and R_i^2 are quite consistent across specifications, and usually insignificant. Fifth, Appendix 1H Table 4 estimates regressions including dummies for raw rank rather than centered or percentile ranks. It again shows how little information is contained in rank compared to treatment. Thus pooling across competitions and centering of ranks does not conceal differences across cutoff points. The only variation that matters is winning versus losing.

I estimate the grant effect on the number of deals, rather than on indicators for VC or all private finance (Appendix 1G Tables 24-25). I use a negative binomial specification to best fit the over-dispersed count data. The results imply, using a conservative estimate, that the grant generates about 2.4 additional VC deals. I also test the grant's impact on early-stage venture capital (VCE^{Post}), which is a subset of VC^{Post} including only seed, angel, and Series A deals. This gave roughly the same results as for VC, albeit slightly smaller, shown in Appendix 1G Table 26.

A logit specification equivalent of Table 1.4 in the main text is in Appendix 1G Table 10. The results are strongly positive, but logit drops competitions without instances of financing. When I use the standard full set of competition dummies, more than half the observations are dropped and the coefficients are quite large. The odds ratio corresponding to the logit coefficient with BW=all implies that a winner is 3.2 times more likely to get VC finance than a loser (column VII), in contrast to the doubling I find with OLS. With topic dummies, fewer observations are dropped but the odds ratio is still 2.9 (column VIII). Clearly, logit grossly overestimates the effect.

Placebo tests check whether any difference between ranks 1 and 2 could be measured as a second discontinuity. Appendix 1H Table 5 runs the basic specification with ranks re-centered so that 0 lies between true ranks 1 and 2. The coefficients are mostly negative, all small, and all insignificant. I test the impact of fixed effects in Appendix 1H Table 6. The treatment effect is unchanged, so the within-competition comparison is apparently unimportant.

Table 1.15: Impact of Phase 1 Grant on VC where Identifying Variation is Number of Awards in Competition

Dependent Variable: VC_i^{Post}	
	I.
$\left(R_i^{Raw}=1\right)\cdot(\mathbf{1} \mid \#Awards=1)$	**0.0930****
	(0.0417)
$\left(R_i^{Raw}=2\right)\cdot(\mathbf{1} \mid \#Awards=1)$	**-0.0549**
	(0.0933)
$\left(R_i^{Raw}=3\right)\cdot(\mathbf{1} \mid \#Awards=1)$	**0.0603***
	(0.0362)
$\left(R_i^{Raw}=4\right)\cdot(\mathbf{1} \mid \#Awards=1)$	**0.0146**
	(0.0344)
$\left(R_i^{Raw}=1\right)\cdot(\mathbf{1} \mid \#Awards=2)$	**0.128****
	(0.0629)
$\left(R_i^{Raw}=2\right)\cdot(\mathbf{1} \mid \#Awards=2)$	**0.148***
	(0.0805)
$\left(R_i^{Raw}=3\right)\cdot(\mathbf{1} \mid \#Awards=2)$	**0.0636**
	(0.0651)
$\left(R_i^{Raw}=4\right)\cdot(\mathbf{1} \mid \#Awards=2)$	**0.0266**
	(0.0487)
$\left(R_i^{Raw}=1\right)\cdot(\mathbf{1} \mid \#Awards=3)$	**0.0946**
	(0.0869)
$\left(R_i^{Raw}=2\right)\cdot(\mathbf{1} \mid \#Awards=3)$	**0.164**
	(0.117)
$\left(R_i^{Raw}=3\right)\cdot(\mathbf{1} \mid \#Awards=3)$	**0.0545**
	(0.0903)
$\left(R_i^{Raw}=4\right)\cdot(\mathbf{1} \mid \#Awards=3)$	**-0.116**
	(0.104)
$\left(R_i^{Raw}=1\right)\cdot(\mathbf{1} \mid \#Awards=4)$	**0.0874**
	(0.164)
$\left(R_i^{Raw}=2\right)\cdot(\mathbf{1} \mid \#Awards=4)$	**0.182**
	(0.220)
$\left(R_i^{Raw}=3\right)\cdot(\mathbf{1} \mid \#Awards=4)$	**-0.0171**
	(0.170)
$\left(R_i^{Raw}=4\right)\cdot(\mathbf{1} \mid \#Awards=4)$	**0.135**
	(0.176)
N	3206
R^2	0.288

Note: This regression interacts raw (non-centered) rank dummies with dummies for the number of awards in the competition. Winning firms' coefficients in blue; losing firms' coefficients in red. The omitted dummy for each number of award group is $\left(R_i^{Raw}=5\right)\cdot(\mathbf{1} \mid \#Awards=x)$. Includes subsample of competitions with 1-4 awards and firms with raw ranks of 1-5. Standard errors robust and clustered at topic-year level. *** $p<.01$. Year$\geq$ 1995

Lee and Card (2008) suggest clustering standard errors by rank with discrete assignment variables. Appendix 1H Table 7 shows that with this method estimated effects are slightly higher than in the primary specification, but they remain significant at the 1% level. The primary results are also essentially unchanged when covariates are excluded and with additional covariates, such as location in a major metro area (Appendix 1H Tables 10-11).[59] Finally, Appendix 1H Table 12 provides permutations of rank for the VC outcome. The basic result is consistent and robust across specifications. Second and third degree polynomials in rank have tiny, insignificant coefficients.

1.7 Back-of-the-Envelope Return Calculation

The RD analysis relies on the *probability* of financing events as a measure of success. My data on ultimate firm valuation, albeit incomplete, provides some insight into the private return to the grants. First, I ask what stake a VC firm would require in order to be willing to invest the total grant amount. This amount - Phase 1 and Phase 2 grants - totals \$616 million (2012 dollars) between 1995 and 2013. The return consists of liquidation, or exit, events after the award: 10 IPOs and 43 acquisitions. Unfortunately, I have dollar amounts for only 14 of the acquisitions. After extrapolating the average acquisition amount to missing deals, the total deal amounts are \$3.01 billion in IPOs and \$2.18 billion in acquisitions (both in 2012 dollars). The average time between the award and the liquidation event is 8.6 years. If a VC firm requires a 30% IRR, it would need to take a 114% equity stake in order to be willing to invest \$0.62 billion in these firms and earn \$5.19 billion 8.6 years later.

Many awardees have not had time to exit because the investment data is censored in mid-2014. Using a Cox proportional hazards model, I estimate the probability of exit at each year from the firm's first award date. Appendix 1G Figure 4 shows the predicted probability of an IPO or acquisition as a function of years from award. I calculate from the estimates that a total of 152 IPOs and acquisitions are expected from the awardees, rather than 52. The gross deal amount is \$12.9 billion (based on the average deal), and the VC required investment stake with a 30% IRR is 46% - still quite high.

In order to maintain entrepreneurial incentives, it is untenable for a VC investor to take 46%

[59] However, when I add woman-ownership and minority-ownership, the sample size decreases precipitously and I lose significance.

of the firm for $150,000.[60] This back-of-the-envelope calculation helps explain why the subsidies might be necessary for firms to access finance, supporting the funds mechanism from Section 1.5. Private investment in the portfolio of grantees at the stage at which they got the grant is apparently unviable, either because the required stake is too large (equity channel) or because the company has not yet proven that its technology works (prototyping channel).

A second exercise considers the "return" from the government's perspective. The RD analysis in Section 1.4.1 found that the grant doubles a firm's probability of receiving any type of private finance. Therefore, I assume that DOE is responsible - as though it took a notional equity stake - for 50% of grantees' subsequent IPOs, acquisitions, and VC deals. I use VC deals only where a firm did not exit.[61] Note that while IPO and acquisition amounts are interpretable as company valuations, VC investments provide a lower bound on the valuation. I allocate an equal share of the total grant "investment" ($616 million) to each unique awardee firm ($630,000), and calculate each deal's IRR (also the CAGR in this case).[62] Summary statistics about the process and the results are in Table 1.16. The average IRR across all awardee firms is 8.5%. This is broken down by firm type as follows: For the 777 firms who never receive any type of private finance, the average return is -100%; for firms with only VC deals, it is 375%; for firms that are acquired it is 512%; and finally for firms that IPO it is 970%. The returns are highly dispersed, and medians are much lower than the means.

Finally, I calculate the Kaplan-Schoar (2005) PME to compare the grant investment to a similar investment in public equity markets, with the same assumption as above that the government takes a 50% stake.[63] I use the S&P 500 index value in the week in which the grant was awarded, and the week of the private financing event (I average daily values to get a weekly index).[64] As

---

[60] Usually in syndicates, VC investors typically own 40-75% of portfolio companies (Gompers and Lerner 2004, Mehta 2011).

[61] As with acquisitions, I extrapolate from the 268 VC deals where I have amounts to the 101 where I do not. I use only observed deals rather than the hazard model prediction.

[62] The compound annual growth rate (CAGR) is the discount rate that makes the NPV of investment cash flow zero, and its formula is: $CAGR = (\text{Deal Amount/Grant "Investment"})^{(1/\#\text{ years})} - 1$.

[63] See Kaplan and Schoar (2005) for an introduction and discussion.

[64] The formula can be simplified for my setting. I calculate a PME for each "investment" DOE makes in a firm and average over firms. The formula for each firm's PME is: $KS-PME = \frac{\text{Deal Amount}}{\text{Award Amount}} \cdot \frac{S\&P_{AwardDate}}{S\&P_{DealDate}}$.

shown in Table 1.16, the average PME across the whole sample of 978 firms is 2.68, which is quite high (a value of 1 indicates that the fund gives the same return as the market.) Again, the standard deviation is large, and the medians for all groups are much lower than the means.

Table 1.16: Back-of-the-Envelope Return Calculation 1995-2013 by Deal Type

	I. IPO	II. Acquisition	III. VC only	IV. No Finance	V. All Firms
# Awardee Firms	10	43	148	777	978
# Deals	10	43	353	0	406
# Deals missing amt	0	29	90	0	119
Mean deal amt (mill)	$301	$50.6	$8.99	0	$20.60
Total deal amt w/extrapolation (mill)	$3,013	$2,175	$3,897	0	$9,084
Grant "investment" per deal (mill)[1]	$.63	$.63	$.63	$.63	$.63
Mean years award to deal	10.46	6.87	3.10	-	3.68
Mean IRR w/ 50% gov't stake	**970%**	**512%**	**375%**	**-100%**	**8.5%**
Median IRR w/ 50% gov't stake	101%	80%	337%	-100%	-100%
Std Dev IRR w/ 50% gov't stake	2,423%	1,078%	355%	0%	414%
Mean KS-PME w/ 50% gov't stake	**14.6**	**23.2**	**10.3**	**0**	**2.68**
Median KS-PME w/ 50% gov't stake	3.48	3.46	4.57	0	0
Std Dev KS-PME w/ 50% gov't stake	25.8	119	21.2	0	26.4
Mean IRR w/ 10% gov't stake	**214%**	**106%**	**39%**	**0**	**-66%**
Mean KS-PME w/ 10% gov't stake	**2.9**	**4.63**	**2.06**	**0**	**0.54**

[1] $616 million/978

Note: This table documents a back-of-the-envelope calculation of the grant "investment" return based on ultimate company valuation, using standard return (IRR) and public market equivalent (PME) formulas. The IRR is the same as the compound annual growth rate (CAGR) here. I assign each deal an equal share of the total DOE SBIR grants given to all firms between 1995-2013. Based on this "investment" of $.82 million, I calculate an IRR and PME for each deal, assuming that the government takes a 50% stake or a 10% stake in the firm. The reported mean return is the average of these deal-specific IRRs and PMEs. Column I shows the return for awardees that experienced IPOs, and column II awardees that were acquired. Where a firm does not have an IPO or acquisition, I use VC deal amounts as a lower bound on firm valuation (column III). Column IV shows the -100% return for all firms with no subsequent private finance. For deals with missing amounts, I extrapolate using the average deal amount for that category. For firms with multiple VC deals, I use the total deal amount and average the time between award and deals. I assign deals that occurred less than 365 days after the award a time period of one year. The Kaplan-Schoar PME is calculated using the S&P 500 index average value during the week of the award and the week of the deal. Mechanically, awardee firms with no deal have a KS-PME of 0. All amounts in millions of 2012 dollars.

The average overall government IRR of 8.5% is slightly lower than calculations of VC fund returns, but the PME of 2.68 is much higher. Kaplan and Schoar (2005), using data from 1981 to 2001, calculate average VC fund IRR net of fees and carried interest at 17-18%, and a PME of

1.2. Preqin's database, using VC fund vintage years from 1981 to 2013, estimates an IRR of 13.5%. Harris, Jenkinson, Kaplan and Stucke (2014) use data from 1984-2008 and find an average IRR of 12.5% and a PME of 1.28.[65] My calculation is sensitive to the government stake assumption, and does not include the administrative cost that would be equivalent to a VC firm's carried interest and fees. However, the results suggest that the government portfolio is unlikely to provide an acceptable return to conventional investors in VC firms, who typically demand higher returns in exchange for the illiquidity and risky nature of the assets. At the same time, this portfolio seems to have substantially outperformed the S&P, which may reflect the somewhat countercyclical timing of the grants and deals.

1.8 Conclusion

Taking the government's objectives as given, this paper establishes that on average DOE SBIR money is not wasteful - it helps propel firms to the private market. For the early-stage projects in my sample, asset intangibility and uncertainty are at their most extreme. Further, energy technology startups are more capital intensive, have longer lead times, and carry higher project finance and market risk than the startups VCs typically finance in IT and biotech (Nanda, Younge and Fleming 2013). Finally, positive externalities motivate basic R&D and entrepreneurship in clean energy, but the absence of a carbon price makes commercialization challenging (Nordhaus 2013). My setting, therefore, is fertile ground for severe financing constraints and grants that provide additionality.

My results indicate that in this context, early-stage grants can alleviate financing constraints. Phase I grants lead recipients to generate more patents and be more likely to commercialize their technologies. Grantees are also nearly twice as likely to access VC finance. The mechanism, surprisingly, does not seem to be certification. Instead, the grants are useful because they increase firms' internal resources. Specifically, my evidence best supports a prototyping effect. Armed with a prototype that reduces uncertainty about its technology, the startup presents venture capitalists with a more viable investment opportunity. The problem, as Shane and Stuart (2002) explain, is that the information funders need to assess quality emerges only after the venture has enough funds to prove its potential. I find that the grants help overcome this Catch-22.

[65] Cochrane (2005) estimates the mean return to VC investments that result in an IPO or an acquisition, correcting for selection bias, at 59% between 1987 and 2000. His estimate includes returns both to the VC firm itself (fees and carried interest) and returns to the investor.

This insight into the grant mechanism contributes to the literature. Wallsten (2000) argues that grants must crowd out private capital because the SBIR program is explicitly designed to select high quality, or inframarginal, firms. Lerner (2000) considers this selection channel, but argues that instead certification explains the positive effects of SBIR grants.[66] My ranked application data indicate that officials do not or cannot choose firms based on their likelihood of success. This supports Lerner's argument against selection, and agrees with his broader argument that officials are unable to choose the "best" firms. I find support, however, for an alternative mechanism to explain the grant effect - the cash itself.

This paper also relates to the corporate finance literature on innovation. Seru (2014) and Bernstein (2012) find that target firms prior to acquisition and private firms prior to IPO, respectively, are more innovative than after the ownership change. Diversified conglomerates have been shown to underinvest (Ozbas and Scharfstein 2009). These and other studies provide grounds for locating R&D in more entrepreneurial, focused institutions. But for the economy to benefit from high-impact entrepreneurship, many startups must be given the chance to test their ideas with the expectation that most will fail (Hsu 2008). While the market effectively disciplines outcomes, initial experimentation may suffer from severe financial frictions. Gruber, MacMillan and Thompson (2008), and Hao and Jaffe (1993), among others, suggest that inadequate external financing hinders new technology development. There is limited direct empirical evidence, however. I extend the literature and provide strong evidence that high-tech startups face financing constraints, which impede innovation.

Governments, both in the U.S. and abroad, fund a large share of applied research. Since 2000, the federal government has spent between $130 and $150 billion per year on R&D, about 30% of total annual U.S. R&D (NSF 2012). To the extent public funds are used to subsidize applied private sector R&D, the findings in this paper suggest that one-time grants to small firms seeking to prototype their product may be more effective in stimulating innovation than large grants that seek to identify and support the "best" firms.

[66] Lerner (2000) reaches this conclusion primarily because the award impact in his sample is larger for more high-tech firms, and also because he finds decreasing returns to additional awards.

2 Risk Management: Evidence from Oil Price Hedging in Highway Procurement

2.1 Introduction

There are rich theoretical predictions of how nonlinearities in firm costs or other characteristics may lead to corporate risk management (e.g. Smith and Stulz 1985). Empirical tests have yielded some evidence that hedging against cash flow volatility is more valuable to more constrained firms (e.g. Vickery 2008, Haushalter 2000) and is associated with increased firm value (e.g. Mackay and Moeller 2007, Nance et al. 1993).[67] Other relationships are more contested or less well studied, in part because the literature has often relied on cross-sectional data, survey data, and data exclusively from publicly traded firms. Further, the risk in question can be correlated with determinants of firm value, especially demand. Last, risk management may be conflated with speculative activity when financial derivatives measure hedging. This paper exploits a natural experiment in a panel setting to assess the value to firms of relaxing constraints on risk management. It then examines how risk management varies by firm ownership, size, and diversification.

Highway procurement is a useful context to study risk management. Paving firms take on oil price risk in the period between a government auction of the project and commencement of work, and government highway "demand" is plausibly exogenous to oil price and other market risks. I use detailed procurement auction data from Iowa and Kansas between 1998 and 2012 to assess the impact of oil price volatility on firm bids to pave asphalt ("blacktop") roads, whose primary component is bitumen, an oil product. This industry is economically important; the U.S. spends around $150 billion annually on public highway construction and maintenance, of which about 85% goes to asphalt roads (CBO 2011).

Many U.S. state governments have recently shifted oil price risk from contractors to Departments of Transportation. These policies emerged from a belief - to my knowledge untested - that firms charge large risk premiums in oil-intensive projects, despite oil's near-zero CAPM-implied

[67] On the other hand, Rampini et al. (2014) show that more financial constraints may lead to less hedging in practice, and Brown et al. (2003) do not find that hedging improves performance. Additional work on benefits of hedging includes Carter et al. (2006) and Cornaggia (2013). On tax convexity, Graham and Smith (1999) and Graham and Rogers (2002) find opposing results. Guay and Kothari (2003), Stulz (1996) and Haushalter (2000) find greater hedging by larger firms (using derivatives). Panousi and Papanikolaou (2012) and Tufano (1996) find a positive relationship between manager ownership of the firm and effective risk aversion.

beta. In 2006 Kansas implemented such a risk removal for bitumen, but Iowa has never done so. The difference between the two otherwise observably similar states is apparently due to bureaucratic preference. After the policy, firms in Kansas were automatically fully hedged.

In a triple-differences design, I test the effect of additional oil price volatility on bids for bitumen in Kansas compared to Iowa after relative to before the policy. I show that a 100% increase in historical volatility leads bitumen bids in Kansas to be 16% lower than in Iowa after relative to before the policy. This effect is robust to a variety of tests, including placebo, falsification with non-oil bid items, alternative volatility metrics, and measures to address potential serial correlation in the variables. Fully hedged firm bids are much less sensitive to oil price volatility, suggesting that firms place a high value on hedging.

Hedging a diversifiable risk is worthwhile only if market imperfections cause cash flow variability to be costly to the firm (Stulz 1996). Froot, Scharfstein and Stein (1993, henceforth FSS) show that financially constrained firms are most likely to hedge when cash flows are negatively correlated with investment opportunities. My setting conforms to this situation; highway contractors are not paid until work is well underway, but must fund labor, materials and equipment costs at the beginning of the construction season. They tend to be cash flow constrained precisely when they are most exposed to oil price risk. In Rampini and Viswanathan's (2010, 2013, henceforth RV) model, constrained firms value hedging more but must weigh its benefits against those of current investment. I show that highway paving firms place a large value on being fully hedged, consistent with risk aversion among constrained firms in both FSS and RV.

Idiosyncratic volatility should not have a measurable effect on bids if firms hedge in financial markets. Interviews with executives of the firms in my data indicate that only the largest, publicly traded firms regularly hedge using oil futures or options. A lack of sophistication, basis risk, economies of scale in hedging, information costs, and the opportunity cost of capital dedicated to hedging are reasons not to hedge in financial markets. Imperfect competition in the industry prevents firms with higher hedging costs from being priced out of the market.[68] Monopoly power in product markets may impede efficient allocation of risk to financial markets, instead leading firms to pass higher risk premiums to the consumer. Relatedly, Scharfstein and Sunderam (2013) find that imperfect competition in mortgage lending decreases the pass-through of lower mortgage-backed security yields to mortgage rates, vitiating government policies aimed at home buyers.

Highway paving firms are heterogeneous, and many are privately owned or family owned

[68] See Section 2.3, as well as Bajari and Ye (2003) and Porter (2005).

firms. Private firms make up 99.9% of the 5.7 million firms in the U.S. (Asker et al. 2014). Yet to my knowledge, there has been no comparison of public and private firm risk management decisions, and there is only a small literature on private firm financial constraints (e.g. Saunders and Steffen 2011). Family firms raise interesting corporate governance questions and are also economically important (Schulze et al. 2001). The theoretical literature provides diverse but largely untested theories about how family firms should manage risk (Schmid et al. 2008, Shleifer and Vishny 1986).

To assess heterogeneity, I build a simple model of the firm's bidding decision that includes a risk premium in the markup on the bitumen item bid. I present a novel reduced form strategy for estimating the impact of risk - measured as the interaction between oil price volatility and time to work start - on firm bids. Using Iowa data matched to firm characteristic data, I show that firms that are publicly listed, diversified across industries, not family-owned, and larger are less responsive to oil price risk than their respective counterpart firms. The strongest results are for the public-private and diverse-concentrated relationships. For example, with 11 months between auction and work start, at the 90th percentile of oil price volatility, a firm whose only business is asphalt paving charges a markup that is roughly \$25 per ton higher than a diversified firm, equivalent to half the average markup. While the increased risk premium of private firms vis-à-vis public firms is driven by the time to work start, the increased premium of non-diversified firms and family firms is driven by oil price volatility.

The firms that are most constrained and have the most concentrated ownership appear to place the highest value on hedging, consistent with the theoretical literature and with a number of empirical studies, including Tufano (1996), Panousi and Papanikolaou (2012), and Geczy et al. (1997). A limitation is that I do not observe hedging directly (though I do observe 100 physical forward hedging contracts for one large firm). I cannot determine whether underlying risk aversion or hedging efficiency drives my results. For example, if public and private firms have the same risk preferences but only public firms hedge in financial markets, then private firms might appear more risk averse when they are simply less sophisticated or do not have adequate scale. While I provide evidence of heterogeneous risk management, further research is needed to identify underlying risk preferences and to establish external validity beyond this context.

I explore the real outcomes of Kansas' risk removal policy in two ways. First, I examine whether different types of firms in Kansas experienced different auction results after the policy. The policy benefited private firms at the expense of public firms, but had no measurable effect on family owned vis-a-vis non-family owned firms. Second, I show that the risk removal policy reduced the price that Kansas paid per ton of bitumen by \$37, or 12% of the average. This is relevant for procurement policy. Unfortunately, I do not have access to variables like profitability and employment to fully assess how the risk removal policy affected firms' real outcomes.

My findings contribute not only to the risk management literature, but also to work on the relationship of oil prices to real outcomes, a literature that includes Bond and Cummins (2004), Henrique and Sadorsky (2011), and Bulan (2005). Finally, this paper is related to the industrial organization literature on auctions (Hendricks and Porter 1988, Athey and Levin 2001) and the highway procurement application (Bajari and Ye 2003, Krasnokutskaya 2011, Jofre-Bonet and Pesendorfer 2003).

Section 2.2 explains the setting and firm risk management practices, including the implications of the conventional CAPM model. It also discusses monopoly power in the industry and describes the risk removal policy. Section 2.3 presents the triple-difference estimation strategy and data on Iowa and Kansas highway auctions, and Section 2.4 presents the risk removal policy triple-difference results. In Section 2.5, I assess the real effects of the policy. Section 2.6 evaluates risk premium heterogeneity across firms, and includes a bidding model to motivate the empirical approach.

2.2 Context: Risk Management and the Risk Removal Policy

In Iowa and Kansas, as in most states, the state Department of Transportation (DOT) procures highway construction projects via simultaneous sealed-bid first-price auctions. First, DOT prepares a public proposal for the project detailing the location and type of work, which includes estimated quantities of materials needed and the expected date of work start. DOT also estimates the cost of each item, but these estimates are not public either before or after the auction.[69] Bidders submit itemized bids with a price specified for each item, including a per ton bid for bitumen (also called asphalt binder). The bidder with the lowest vector sum of unit item bids wins the auction. By submitting a lower bid, a paver ensures a higher chance of winning, but takes the risk that high cost realizations could make the project unprofitable.[70] Once a winner is announced and the

[69] The "engineer's estimate" is calculated in both states by computer programs that use recent bid history for similar projects. Materials with volatile prices, like oil products, are updated manually. There is no reserve price; the secret estimate serves as a guide for what is reasonable.

[70] The unit item bids are analytically meaningful. Contracts are contingent rather than fixed price, so the paver is paid for the miles of guardrail actually installed, or the cubic meters of earth actually excavated, rather than the estimated quantity in the proposal. It has been widely noted in the auction literature

contract signed, time passes (5 months on average) before work begins.

2.2.1 Firm Risk Management

Highway pavers face cost uncertainty at the time they place their bid. Primary risks are weather and oil prices, and secondary risks are labor shortages and equipment failure. For asphalt paving, the largest risk is that an oil price spike between the time of the auction and the time of work will lead to unexpectedly high bitumen costs. An oil product, bitumen is the black, sticky material mixed with rock pieces to make asphalt. My analysis includes only paving (also called resurfacing) projects that are very asphalt-intensive.[71] Although its price is highly correlated with crude oil, there is no liquid market or futures contracts for bitumen in the US. Instead, firms purchase bitumen in one-off, non-public transactions with local suppliers, who store bitumen purchased from refineries.

Firms can manage risk via hedges, insurance, or diversification (Merton 1995). I conducted a phone survey of the twelve largest bidders in Iowa and three large bidders in Kansas, and spoke either with a President, a Vice President, or an Estimator (who writes up the bids for DOT auctions). The firms - which range from family-owned firms with a dozen full-time employees to some of the world's largest construction conglomerates - describe themselves as very risk-averse toward input costs. They usually hedge bitumen risk with physical forward contracts signed at the time of auction with local third party suppliers. Firms sometimes wait to sign later, or buy spot at the time of work and either don't hedge at all or occasionally hedge in financial markets.[72] Storing bitumen

that scoring rules, or unit-price contracts, generate incentives to skew; to overbid on items that DOT has underestimated and to underbid on items that DOT has overestimated. DOT pays the winning contractor based on quantities actually used, so it is in the contractor's interest to put his profit margin on items that are likely to overrun. Skewing is sufficiently pervasive that IDOT explicitly forbids it, reserving the right to reject bids it deems "unbalanced." In practice, this is achieved through rules of thumb; when bid items appear to be weighted in a manner that causes them to differ appreciably from the engineer's estimate, the bid is rejected (about 3% of bids are rejected for this reason). Skewing incentives do not bias my risk management findings.

[71] I do not study diesel, another oil product used in highway paving, because it is much smaller as a percentage of the total bid, and is not a bid line item but rather goes into a line item for general overhead.

[72] The physical forward contracts are based on quotes that pavers request from bitumen suppliers before the auction. The paver often signs a contract with one supplier committing to purchase the bitumen at the quoted price at the time of work start should he win the project. The price is good only for the DOT project specified in the contract, and the bitumen can be taken typically any time during the construction season (roughly mid-April to the end of October, because paving requires a road temperature no less than 55° F). The supplier must have sufficient bitumen stored to cover all contracted supply. Although end-use demand for bitumen in Iowa only exists for 1/2 the year, oil refineries produce bitumen year-round as a byproduct.

is costly, so most paving firms do not own storage facilities. Since the suppliers do store bitumen, at the time of the auction they are partially physically hedged against the short positions they are taking in their contracts with paving firms. Thus in the supplier-paver relationship, the supplier has downside risk while the paver hedges against upside risk. If the supplier has total bargaining power, the price could include *both* firms' risk premiums.

I have 100 forward physical contracts from one firm (who I call Firm Z to protect its identity) with all three local suppliers. Firm Z is among the top three firms in number of total bids submitted, and has about the mean percentage of contracts won among regular bidders. An example (fictional) contract might be for 1,200 tons of bitumen at $510 per ton, dated January 23, 2009 for IDOT project STP-038-3(46)–2C-53, effective from April 15, 2009 to November 15, 2009. Figure 2.1 shows the actual per ton price specified in the contract alongside the 1-month futures contract price. Figure 2.2 shows the markup (the bid in the auction less the forward contract price). The markup is fairly stable at about $22 per ton regardless of oil price levels and volatility. Any risk premium is apparently included in the forward contract. In the pre-2005 period, when prices were relatively low and stable, the markup averaged 13.7 percent of the bitumen bid. Post-2005, when prices were higher and more volatile, it was 5.4 percent.

A rich literature, beginning with is Mayers and Smith (1982) and Smith and Stulz (1985), is devoted to identifying market imperfections that might lead firms to hedge. Froot, Scharfstein and Stein (1993) show that hedging can mitigate an underinvestment problem that emerges when firm cash flows are not strongly correlated with investment opportunities and firms face financial distress costs. Hedging allows firms to take advantage of profitable investment opportunities even in bad cash flow states. Rampini and Viswanathan's (2013) dynamic model of collateralized financing has a more complex relationship between hedging and financial constraints; constrained firms must weigh the benefits of current investment against those of hedging.

The refineries typically don't store bitumen, so they sell it to third parties who own terminals (storage capacity). These third parties, my "suppliers," start storing binder in early winter, but sign contracts with pavers during the winter that go far beyond their storage capacity.

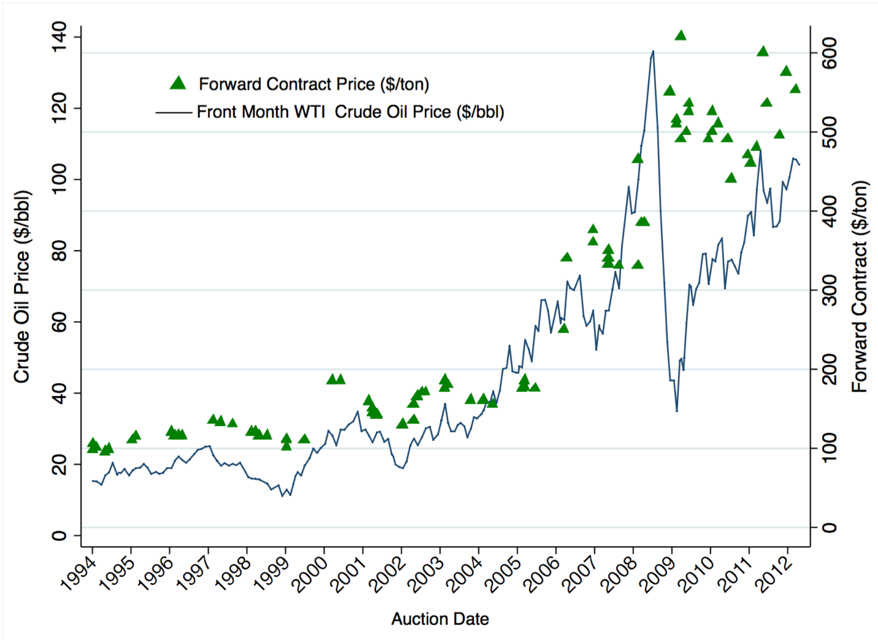

Figure 2.1: Firm Z Forward Physical Bitumen Contract Prices
Note: This figure shows the bitumen prices in 100 forward physical contracts between one large paving firm and bitumen suppliers, as well as the spot oil price.

Figure 2.2: Firm Z Bitumen Bid Markup Over Forward Physical Contract with Supplier
Note: This figure shows the paving firm's bitumen markup, defined as the project-specific unit bid less the project-specific forward contract price. It uses 100 forward physical contracts between one large paving firm and bitumen suppliers, as well as the bid observed in auctions.

Firms might not hedge in financial markets because of a lack of sophistication, basis risk, economies of scale in hedging, information costs, and the opportunity cost of capital dedicated to hedging. First, the paving firms are mostly small, local, and typically do not have in-house financial sophistication. In interviews, firm executives consistently expressed distrust in Wall Street, and viewed hedging in financial markets as gambling. Second, pavers face two kinds of basis risk that separate the spot price change from the futures contract: different assets (oil vs. bitumen); and different contract time horizons. Although in theory pavers can choose a hedge ratio to minimize the variance of the hedge, basis risk nonetheless lowers the value of hedging.[73] Haushalter (2000) finds a strong negative correlation between basis risk in hedging instruments and the fraction of production hedged, as well as a strong positive correlation between firm assets and the likelihood of hedging. Third, small firms may not hedge because of economies of scale in hedging in financial markets (Mian 1996, Geczy et al. 1997).

Hedging ties up firm capital, a fourth reason that firms might not hedge in financial markets. While the transaction costs are low, hedging requires the firm to maintain a margin if via futures contracts or to buy calls if via options on futures. In a hypothetical calculation based on information from OptionsXpress, a brokerage firm, fully hedging against oil price increases for a typical project using oil futures might require a margin in the account of $40,000.[74] If prices fall, the broker will likely issue a margin call, requiring the immediate wiring of funds - and thus a dedicated employee - into the account to prevent the contracts from being closed. Alternatively, purchasing call options on the same number of contracts might cost around $10,000.[75] However, the firm must purchase more options than the underlying oil quantity being hedged to achieve a 1-to-1 hedge, navigating the declining delta of the option as it moves out of the money. Firms are cash flow constrained during the winter, when they participate in most auctions and establish their hedging positions.

---

[73] The ideal hedge is to have any change in the spot price equal to the change in the futures contract with which you are hedging.

[74] The typical project would need uses the bitumen equivalent of roughly 4,400 barrels of oil (see http://www.eia.gov/dnav/pet/tbldefs/pet_pnp_pct_tbldef2.asp). If a firm bought five six-month futures contracts at $80 per barrel with a 10% margin requirement, it would need $40,000 in the account. More importantly, if prices fall, the broker will likely issue a margin call, requiring the immediate wiring of funds - and thus a dedicated employee - into the account to prevent the contracts from being closed. To bring the amount in the account up to the initial margin requirement, the firm may have to add roughly an extra $500 for each dollar the price of oil drops.

[75] The firm could purchase a call expiring in six months. If the call costs $2, the cost of options on 5,000 barrels is $10,000.

The cost of hedging all their expected bitumen consumption for the following construction season may be prohibitively high, preventing investment in equipment or labor. This is the argument in Rampini, Sufi and Viswanathan (2014), who show that constrained firms may not hedge because of inadequate resources.

Last, there is evidence in the literature that financial intermediaries may reduce the benefits to hedging. For example, Etula (2013) demonstrates a link between broker-dealer effective risk aversion (broker-dealers take the other side in OTC hedging contracts) and commodity price risk premiums, a link empirically particularly strong for energy returns. Investing in a fund may not be ideal either. Bhardwaj, Gorton and Rouwenhorst (2014) show that commodity trading advisors on average provide excess returns (after fees) to investors of roughly zero, while gross excess returns (before fees) are 6.1%. They conclude that the best rationale for investors' continued use of these vehicles is information asymmetry. In practice, studies of oil producers and airlines have found that hedging occurs in a minority of time periods (Haushalter 2000, Carter et al. 2006, Jin and Jorion 2006).

In the conventional CAPM, the expected future oil price's relationship to the futures contract price depends on the risk-adjusted discount rate, the risk-free rate, and the spot price of oil. Abstracting from storage and transport costs, if oil covaries with the market return, the risk premium should in general be positive. That is, the discount rate should exceed the risk-free rate, equivalent to a positive CAPM beta. Investors are compensated for systematic risk by a futures price that is lower than the expected future price. Thus with perfectly functioning capital markets, firms bidding in highway auctions should charge the government the CAPM-implied beta.

The crude oil beta appears to be near zero.[76] Figure 2.3 shows rolling betas for a conventional strict CAPM regression.[77] The sign of the oil beta is sensitive to the period chosen, because oil's correlation with the market depends on whether the oil price movement is driven by demand or

[76] A beta of zero indicates the asset has no relationship with the market portfolio, a beta of 1 indicates that the asset has the same risk (moves with) the market. Negative beta indicates that the asset tends to move in the opposite direction of the market. In this latter case, the negative risk premium indicates that the return on the asset should be less than the risk-free rate.

[77] The expected return on the asset is $E(R_i) = R_f + \beta_i(E(R_m) - R_f)$, where returns for asset i are calculated monthly using the price time series as $R_{i,t} = \frac{p_{i,t} - p_{i,t-1}}{p_{i,t-1}}$ and $\beta_i = \frac{\text{cov}(R_i, R_m)}{\text{var}(R_m)}$. Denoting the market premium $\rho_{m,t} = R_{m,t} - R_{f,t}$, and the oil premium $\rho_{o,t} = R_{o,t} - R_{f,t}$, the regression equation is $\rho_{o,t} = \beta_o \rho_{m,t} + \varepsilon_{o,t}$. I use the front-month WTI oil futures, the S&P 500 index price for the market portfolio, and 3 month Treasury Bill interest rates for the risk free rate.

supply. With a beta of 0.5 (roughly the level post-2009), and assuming annual stock market returns of 6%, the return on bearing the oil risk (in theory the difference between the 6 month futures price and the expected spot price in 6 months) should be at most 1.5%. Similarly, Ahn and Kogan (2011) find oil equity beta from a standard CAPM procedure between 1971 and 2010 to be 0.01. Thus risk-neutral firms should charge little if any premium for holding oil price risk.

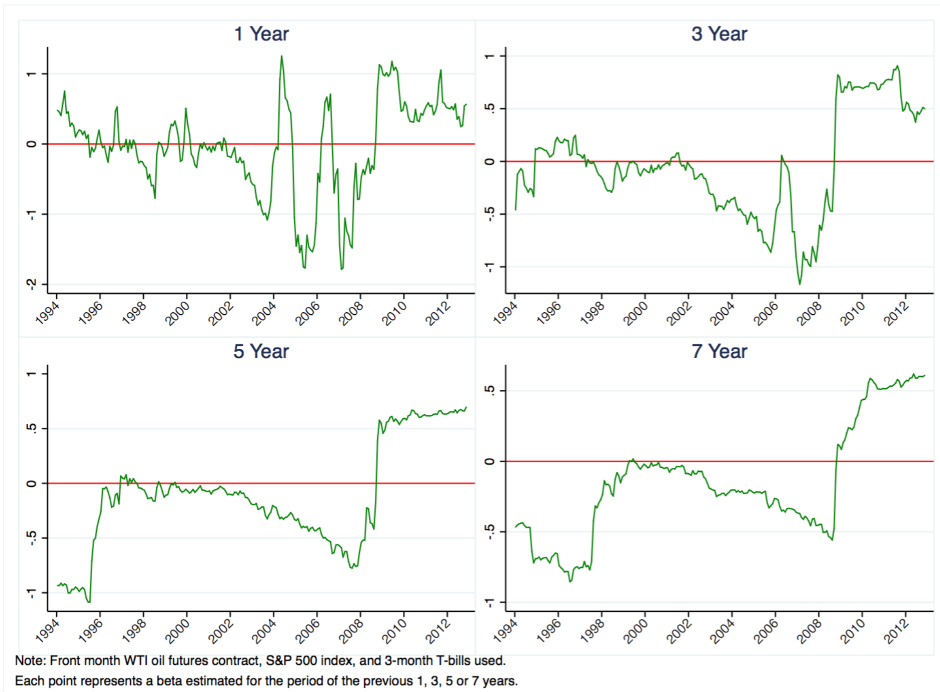

Figure 2.3: Rolling Betas for Crude Oil

Note: This figure shows rolling crude oil betas for a conventional strict CAPM regression. I use front-month WTI oil futures, the S&P 500 index price for the market portfolio, and 3 month Treasury Bill interest rates for the risk free rate. The sign of the beta is sensitive to the period over which it is calculated, but on average is close to zero.

Much of the theoretical literature relies in non-linearities or asymmetries in firm cost structure or information sets to explain the value of hedging (e.g. FSS and Demarzo and Duffie 1991 and 1995). Most closely related to the input cost risk studied here, the intuition in Mackay and Moeller (2007) is that convex costs can make hedging valuable (as can concave revenues). In their model, the second derivative of the cost function with respect to the input price exists because the quantity used is a function of the input price. In contrast, in my setting the quantity of the input is exogenously fixed at the required project amount. Therefore, their line of reasoning does not rationalize hedging among paving contractors. Instead, my results suggest that financial constraints best explain

I do not address the risk of losing the auction. Anecdotal evidence from interviews suggests that paving firms are risk-averse towards input costs but risk-neutral towards an individual auction for a particular project. Firms participate in many auctions and seem to treat them as a portfolio. While the risk of losing any given auction is idiosyncratic, oil price risk for the coming construction season is highly correlated across projects.

2.2.2 An Imperfectly Competitive Environment

In a competitive environment firms would bid down the price of bitumen risk to the cost for the least risk averse agent. Instead, paving firms and bitumen suppliers are in oligopolistic, territorial equilibria. Highway procurement is characterized by inelastic demand, high barriers to entry, information asymmetry, easy defection detection, auction setups where phony bids are possible, and a static market environment, which are all conducive to collusion (Porter 2005). Porter and Zona (1993), Ishii (2008) and Pesendorfer (2000) demonstrate collusive bidding in highway procurement contracts, and Bajari and Ye (2003) note the widespread incidence of cartels in procurement auctions. Gupta (2002) finds collusion in Florida highway procurement and estimates that this type of auction is not competitive until there are 8 bidders participating. In my data, the average number of bidders is 3.4.

Many of the regular bidders in my dataset have well-defined territories; Figure 2.4 shows the location of auction wins and losses for a large bidder, Norris. Wins are obviously concentrated around its headquarters. Appendix 2C contains similar maps for all the top bidders, suggesting that each has a distinct territory. This may be due to tacit collusion or transportation costs (with perfect competition the rents are zero on territory boundaries and positive within). In my phone survey, one CEO suggested of his own accord that the imperfectly competitive nature of the business permitted even very risk averse pavers to stay in business.

Only a handful of bitumen suppliers serve a given local market; in Iowa there are just three. The high cost of transporting bitumen permits a spatial oligopolistic equilibrium. Like the paving firms, they enjoy markups within their territories at least as large as the differential transportation cost for the next-closest supplier. In Firm Z's 100 forward physical contracts, I find distinct territories for the suppliers, shown in Appendix 2C Figure 1. The three suppliers provide quotes to paving firms for on average 153 Iowa DOT auctions per year. Bids (including bitumen items) are published immediately after the auction. In interviews, the suppliers suggested that recent auctions may provide a signaling mechanism, which aligns with Friedman's (1971) seminal discussion.[78] The suppliers charge the pavers if not their full cost of risk, at least a significant portion. Imperfect competition in two layers of product markets permits firms with higher effective risk aversion to remain in the market.

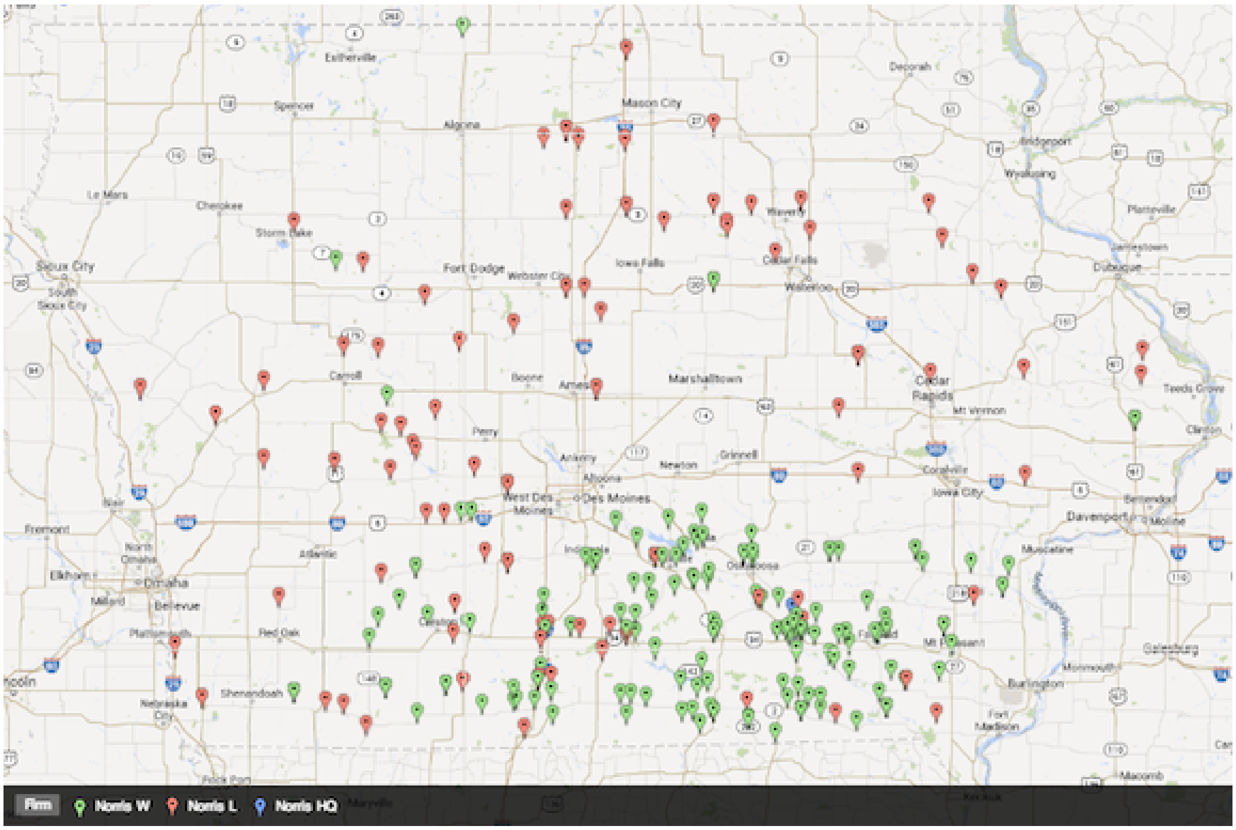

Figure 2.4: Norris Bids in Iowa Highway Procurement Auctions
Note: This figure shows the location of auction wins (green) and losses (red) for a large bidder, Norris. The firm headquarters is blue. Wins are obviously concentrated around its headquarters. Appendix C contains similar maps for all the top bidders, suggesting that each has a distinct territory.

2.2.3 A Natural Experiment: Bitumen Price Adjustment Policies

In the mid-2000s, states began to adopt oil price "adjustment" policies to transfer the risk of volatile oil prices from the private sector to the government. The Kansas DOT implemented its

[78] "...It seems unsatisfactory for firms to achieve only the profits of the Cournot point when each firm must realize more can be simultaneously obtained by each. This line of argument often leads to something called 'tacit collusion' under which firms are presumed to act as if they colluded. How they do this is not entirely clear, though one explanation is that their market moves are interpretable as messages" (Friedman 1971).

bitumen risk removal policy in August, 2006, at which time its director of operations announced that "The volatile price of the asphalt oil has led contractors to make bids that are more costly than necessary" (Shaad 2006). Kansas adjusts its payment to the paving firm by the amount that an oil-based price index has increased or decreased since the auction. When prices go up, the firm is paid his bid plus the index's increase, and when prices go down, the firm receives less than his bid.[79] The introduction of the policy means that Kansas firms are automatically fully hedged by the government. Thus private hedging contracts with local suppliers are unnecessary; put another way, the supply contracts that formerly may have included a risk premium no longer do.

Iowa has not yet removed oil price risk from its private contractors, apparently at least partly due to official inertia. It seems that certain members of the Kansas DOT leadership became interested in oil price volatility, which ultimately led to the policy, while the Iowa leadership did not develop such an interest. I have no reason to believe that a systematic difference between the states led to the different policy outcomes. Both DOTs report that some firms were in favor, and some opposed, to the policy. Kansas, located immediately to the southwest of Iowa, has similar weather patterns, road systems, and auction characteristics. There are also similar proportions of public, private and family owned firms in the two states (see Section 2.3.2).

The consensus in the policy community is that these price "adjustment" policies, now used by most states, reduce firms' input cost uncertainty, and the cost to the government of bearing oil price risk is offset by lower bids (Skolnik 2011). However, to my knowledge there is no public evaluation of the policies' impact on procurement cost. In the only analysis thus far, Kosmopoulou and Zhou (2014) examine only at one state, Oklahoma, so they cannot control for economic and other factors. They attribute their finding that firms bid more aggressively after the policy to the winner's curse effect, and assume firms are risk-neutral.

[79]Specifically, each month KDOT publishes an Asphalt Material Index (AMI), which they purchase from Poten & Partners. Bidders incorporate the current month's AMI into their bid for asphalt. The AMI for the month of the letting becomes the Starting Asphalt Index (SAI) for the duration of the contract. KDOT technicians take samples from the mix being placed. This serves both to monitor quality and to obtain a percent bitumen content to adjust payment based on the change in the AMI. The difference between the SAI and the AMI to the nearest dollar becomes the adjustment factor, applied to work completed during that month. The adjustment only occurs when the AMI differs from the SAI by $10 or more. The Kansas price index is almost identical to the Argus Media spot price index I use elsewhere in the paper. Both are created from surveys of recent bilateral transactions. The KDOT index is for PG 64-22 but KDOT applies it to all grades. For the index, see: http://www.ksdot.org/burconsmain/ppreq/asphaltpriceindex.asp. For the specifications, see: http://www.ksdot.org/burconsmain/specprov/pdf/90m-0295-r01.pdf.

2.3 Empirical Strategy and Data

2.3.1 Triple-Differences Design

Iowa and Kansas' bifurcated policies towards oil price risk offer a unique opportunity to assess firm risk preferences. I use a triple-difference design to assess whether oil price volatility affects bitumen bids in Iowa relative to Kansas since the implementation of the bitumen price adjustment policy in 2006. The triple-difference design is more robust than any differences-in-differences approach (Imbens and Wooldridge 2007). The underlying intuition, however, is the same. If two groups are ex-ante similar and one is subject to treatment in the second of two time periods, then with controls for treatment and state the estimated coefficient on the treated state should be the average difference between the treatment group and the control group, without bias from trends over time and from permanent differences between the groups. Following the suggestion in Bertrand et al. (2004), I cluster standard errors by firm in my primary specification, and demonstrate that the results are robust to alternative groupings.

I estimate the effect of oil price volatility on bids in Kansas ($\mathbf{I}_{KS_j} = 1$) after the policy was implemented ($\mathbf{I}_{Policy_t} = 1$). The regression, where i indexes bidders, j indexes auctions, and t indexes letting day, is:

$$\begin{aligned}\ln b_{ijt} = \beta_0 &+ \beta_1 \mathbf{I}_{KS_j} \cdot \mathbf{I}_{Policy_t} \cdot \ln V_t^{\text{oil}} + \beta_2 \ln V_t + \beta_3 \ln p_t^{\text{oil}} + \gamma' \cdot \text{Auction/Bidder Chars} \\ &+ \beta_7 \mathbf{I}_{\text{state}_j} + \beta_8 \mathbf{I}_{\text{policy}_j} + \beta_9 \mathbf{I}_{\text{policy}_t} \cdot \ln V_t^{\text{oil}} + \beta_{10} \mathbf{I}_{\text{policy}_t} \cdot \mathbf{I}_{KS_j} + \beta_{11} \mathbf{I}_{KS_j} \cdot \ln V_t^{\text{oil}} \\ &+ \delta \mathbf{I}_{county_j} + \delta \mathbf{I}_{month_j} + \delta \mathbf{I}_{year_j} + \epsilon_{ijt} \qquad (9)\end{aligned}$$

$$\begin{aligned}\gamma' \cdot \text{Auction/Bidder Chars} &= \gamma_1 \# Bidders_j + \gamma_2 \ln AveTotalBid_j + \gamma_3 \ln ReqTons_j \\ &\quad + \gamma_4 \ln OtherBidItems_{ij} + \gamma_5 \ln Distance_{ij}\end{aligned}$$

where $\ln V_t^{\text{oil}}$ is oil volatility, p_t^{oil} is the 6-month futures price, $\# Bidders_j$ is the number of bidders in the auction, $AveTotalBid_j$ is the average total bid in the auction (a measure of the project scale), $ReqTons_j$ is the estimated required tons of asphalt, $\ln OtherBidItems_{ij}$ is the sum total of all other bid items that the bidder submits, and $Distance_{ij}$ is the distance in miles between the bidder and the project site. Equation 9 includes controls for the state, the time period (whether the policy was in effect), the three sets of interactions (policy-oil volatility, policy-state, and state-oil volatility), and fixed effects for the county, the month of the year, and the year. Kansas' competitive equilibrium both among pavers and between pavers and suppliers may have changed after the policy. However, changes that are unrelated to oil price risk should be controlled for by the state and state-time fixed effects.

2.3.2 Data on Highway Auctions and Oil Prices

My estimation of Equation 9 employs comprehensive, detailed data on auctions and payments from the Iowa and Kansas DOTs between 1998 and 2012. I focus on straightforward paving projects projects that are bitumen-intensive. In these contracts the bitumen cost comprises 11.3% of the total bid on average, but can be up to 40%.[80] Figure 2.5 shows Iowa and Kansas bitumen bids (bid unit items within the larger total project bid) over time, as well as the crude oil price and historical oil price volatility. Appendix 2A Figure 1 shows the bids and oil volatility in the years immediately around the 2006 policy.

Summary statistics are in Table 2.1, with bid data in the top panel and contract (auction) data in the bottom panel. There are more projects in Iowa than in Kansas; in total there are 4,618 bids in Iowa and 1,438 bids in Kansas. Iowa projects are more bitumen-intensive, but the item bids for bitumen are very similar across the two states, at \$303 for Iowa and \$362 for Kansas over the whole period. Kansas firms are also slightly further away from the project site than Iowa firms, at 113 miles compared to 88 miles. The auction characteristics are similar across the two states, with an average of 3.3 bidders in Iowa auctions and 3.6 bidders in Kansas auctions. Money on the table (the percent difference between the second lowest and the winning bid) is 6.3% for Iowa and 5.3% for Kansas. These figures are close to Krasnokutskaya (2011), who used data from Michigan and found that money on the table averaged 7%. The time to work start is also similar, at 4.6 months for Iowa and 5.2 months for Kansas. None of the differences are statistically significant.

I use two dependent variables. One is the unit item bid on bitumen, b^A, which is depicted in Figure 2.5 and is the most direct measure. However, the same percentage markup is likely not

[80]These projects do not include bridge work or extensive earthwork. For Kansas, the work types I include are called overlay and surfacing, codes 20,53,55,64,65,66, and 67. For Iowa, they are generally called paving and resurfacing, codes 1521, 1522, 1523, 1524, 1525, 1021 and 1022.

applied universally to all items, and strategizing over where to place profit across items might distort the true effect of volatility on the metric that matters to DOT, which the overall bid for the project. Therefore, I also use the total bid for the project per ton of required bitumen, b^T.[81]

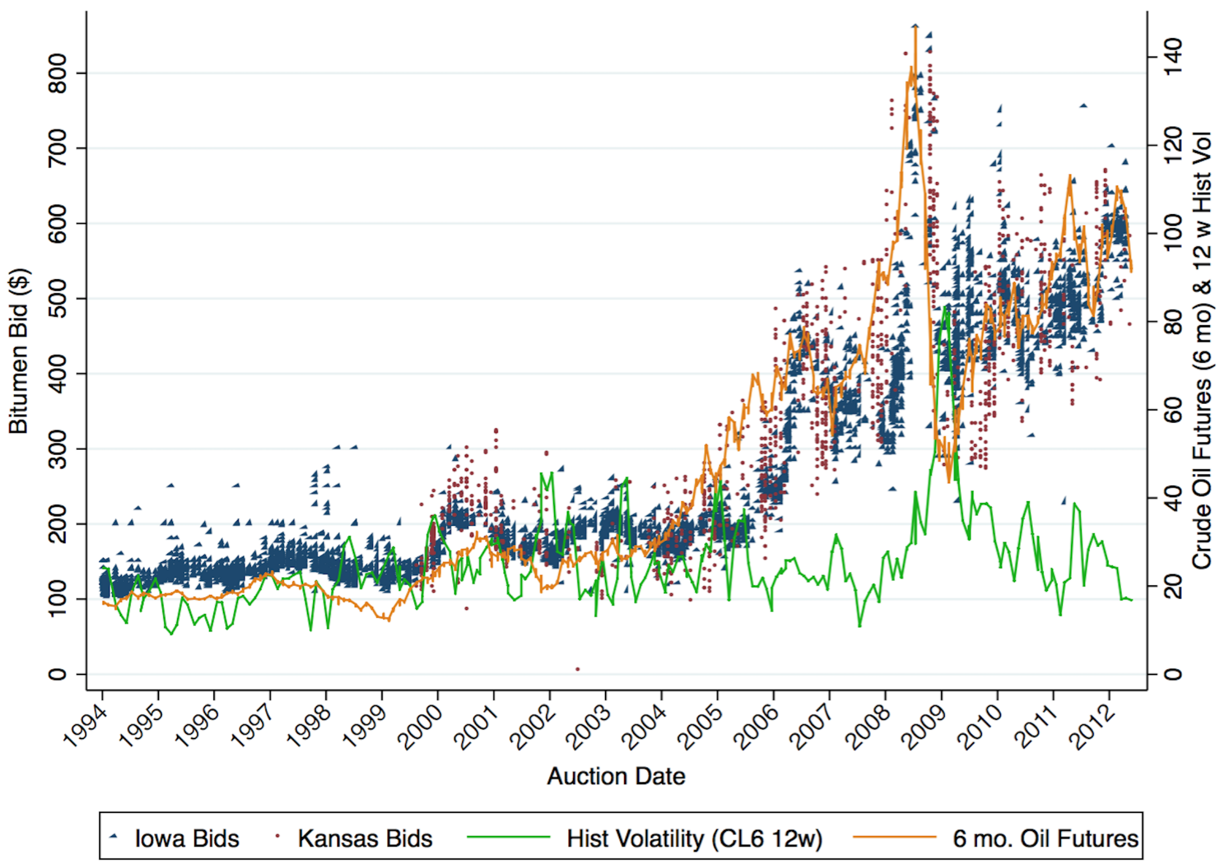

Figure 2.5: Bitumen Bids in Iowa and Kansas, Oil Price, and Oil Volatility
Note: This figure shows all of the highway asphalt paving procurement auction bids in my data from Iowa (blue) and Kansas (red). The Iowa data begins in 1994, and Kansas in 1998. It also shows the 6-month WTI oil futures price and the 12-week historical oil price volatility using the 6-month WTI contract.

In my primary specifications I use historical and implied volatility. The former is an annualized standard deviation of daily returns. Implied volatility is based on the observed prices of put and call options placed on the futures contract and derived by numerically inverting the Black-Scholes (1973) option pricing formula. The options implicitly contain information about market partici-

[81]In order to ensure that bitumen is a meaningful part of the project, I only use projects in which the portion of the total bid that is bitumen is at least $50,000. The effect of this term may also reflect the importance of diesel fuel in the total bid. Diesel price risk is highly correlated with bitumen price risk.

pants' expectations of price volatility (Chalamandaris and Tsekrekos 2009).[82] Historical volatility seems a more natural measure from the paving firms' perspectives. They are very cognizant of recent oil price trends, but do not report looking at options on oil futures, much less implied volatility. Specifically, I use historical volatility over the past 12 weeks, and Bloomberg data on at-the-money implied volatility for options expiring in 3 months. Since they are both over a 3 month period, the 12-week historical volatility and the implied volatility are considered directly comparable. I show that my results are robust to alternative measures, such as 26 week historical volatility. To control for the expected oil price, I use six-month oil futures.[83]

2.4 Triple-Differences Estimation

2.4.1 Results

I find that the Kansas risk removal policy caused bids in Kansas to be less responsive to oil price volatility than bids in Iowa. The results from the primary specification is shown in Table 2.2, using both historical and implied volatility and, for the dependent variable, either the bitumen bid (b^A) or the total bid for the project per ton of bitumen (b^T). The coefficient of interest is β_1, which represents the effect of a 1% increase in volatility on bids in Kansas relative to Iowa after oil price risk shifted to the public sector.

The coefficients on the triple interactions are all negative and significant; in column I, the -.16 means that a 100% increase in volatility, which occurs frequently in my data, results in a 16% decrease in b^A in Kansas relative to Iowa, significant at the 5% level. Using b^T as the dependent variable gives a similar coefficient of -.15, significant at the 10% level. Both use 12-week historical volatility. The effects with implied volatility (columns III and IV) are larger, with coefficients of -.35 and -.47, both significant at the 5% level.

[82]"Model-free" option-implied volatility metrics have been developed to deal with perceived issues with Black-Scholes, but these are beyond the scope of this paper (see Bollerslev et al. 2011).

[83]There is disagreement about whether the futures price or the current spot price is the best forecast of future oil prices (Alquist and Kilian 2010, Kellogg 2010). Here I use the six month futures price, following convention in the literature on volatility and the fact that the average time to work start is five months. Futures contracts not purchased for physical delivery close or roll over at the end of the month prior to the delivery month.

Table 2.1: Summary Statistics of Iowa and Kansas Auction Data, 1998-2012

Panel 1: Bids

	Iowa		Kansas		All	
	Mean (sd)	N	Mean (sd)	N	Mean (sd)	N
Total bid	\$2,254,777 (\$3,271,562)	4,618	\$2,772,247 (\$5,214,448)	1,438	\$2,377,651 (\$3,829,190)	6,056
Bitumen bid item (per ton)	\$303 (\$149)	4,618	\$362 (\$165)	1,438	\$317 (\$156)	6,056
Total bid per ton bitumen required	\$9,784 (\$28,872)	4,618	\$26,861 (\$104,318)	1,438	\$13,737 (\$56,658)	6,056
Miles to project	88 (91)	4,618	113 (227)	1,438	94 (136)	6,056
Total bid for non-bitumen items	\$1,983,852 (\$3,158,923)	4,618	\$2,513,103 (\$5,145,371)	1,438	\$2,109,523 (\$3,734,004)	6,056

Panel 2: Contracts (Auctions)

	Iowa		Kansas		All	
	Mean (sd)	N	Mean (sd)	N	Mean (sd)	N
# bidders	3.33 (2.06)	1,287	3.59 (1.73)	364	3.39 (1.99)	1,688
Money on the table	6.3% (7.6%)	1,287	5.3% (8.5%)	364	6.1% (7.8%)	1,688
Proposed tons bitumen	811 (1,099)	1,287	625 (657)	364	771 (1,023)	1,688
Months between auction and work start	4.6 (2.8)	1,287	5.2 (2.8)	364	4.7 (2.8)	1,688

Note: This table summarizes the bitumen-intensive projects (highway paving) used in the regression analysis. Money on the table is the the percent difference between the second lowest and the winning bid: $100 * \frac{(B^{Second} - B^{Win})}{B^{Win}}$. Auctions with only 1 bidder are excluded for this metric. Miles to project is Vicenty distance calculated using the latitude and longitude of the project site.

The CAPM exercise in Section 2.2.2 indicated that the crude oil beta, or volatility of oil relative to the market portfolio, has been on average zero and rarely greater than 0.5, implying a risk-adjusted return of at most 1.5% over six months. Although not directly comparable, my Iowa-Kansas analysis reveals that the magnitude of the risk premium in Iowa is at least an order of magnitude greater than CAPM would imply, keeping in mind that my statistic is relative to Kansas and relative to the pre-policy period.

The strong effect of volatility on bids suggests that at least some firms do not exploit financial markets. Imperfect competition may prevent higher cost firms from being priced out of the market. If the suppliers were competitive, the bitumen price to paving firms should reflect the supplier's cost of hedging (physically and in financial markets), not the value to the paver of reducing risk. With monopoly power, suppliers can charge the paver his full cost of risk. Both the bitumen supply market and the procurement auction market are imperfectly competitive. I cannot explore further how the policy affects the supplier-paver relationship, but evidence of a substantial risk premium - in the context of a near-zero CAPM-implied beta - suggests that with two layers of imperfect competition in product markets, the consumer does not get the benefit of the least risk averse agent.

2.4.2 Robustness Tests

Key robustness tests are in Tables 2.3 and 2.4. For both, I use bitumen bids (b^A) and historical volatility; Appendix 2A Tables 1 and 2 contains these specifications using implied volatility and total bid per ton bitumen $\left(b^T\right)$. First, Table 2.3 shows single differences and single interactions in columns I-IV. There is a strong positive effect of the $\mathbf{I}_{\text{policy}_t} \cdot \ln V_t^{\text{oil}}$ interaction when the triple interaction is omitted (column III). Otherwise, the single interactions are not significant. In column V, I omit the control covariates (such as the spot oil price and number of bidders). The coefficient increases somewhat in magnitude to -.21, still significant at the 5% level. Omitting month and county fixed effects, in column VI, and year effects in column VII, give similar results to the primary specification of -18 and -17, respectively, both significant at the 5% level.

Table 2.4 first conducts a series of placebo tests, in which the policy implementation year is artificially set to 2002, 2004, or 2008. In 2002 and 2008, the coefficient on the triple difference is near zero and insignificant. In 2004, it is -.22, significant at the 10% level. This placebo is close to the actual policy, and so a noisier policy effect. In column IV I demonstrate that prior to the policy, there was no measurable difference between Kansas and Iowa in their response to oil price volatility. This is shown in the interaction $\mathbf{I}_{KS_j} \cdot \ln V_t^{\text{oil}}$, which gives an insignificant coefficient of .068. Finally, column V is a falsification test with the total bid less bitumen as the dependent variable. The coefficient on the triple difference is now .058, with no significance.

Table 2.2: Triple Difference Results using Risk Removal Policy

Vol Metric Used:	Historical Volatility (12 w)		Implied Volatility	
Dependent variable:	I: Log bitumen bid $\left(b^A\right)$	II. Log bid total per ton bitumen $\left(b^T\right)$	III: Log bitumen bid $\left(b^A\right)$	IV. Log bid total per ton bitumen $\left(b^T\right)$
$\mathbf{I}_{KS_j} \cdot \mathbf{I}_{Policy_t} \cdot \ln V_t^{\text{oil}}$	**-.16*****	**-.15****	**-.35*****	**-.47*****
	(.036)	(.072)	(.069)	(.14)
$\mathbf{I}_{\text{policy}_t} \cdot \ln V_t^{\text{oil}}$	.79***	.33***	.67***	.62***
	(.04)	(.089)	(.05)	(.036)
$\mathbf{I}_{\text{policy}_t} \cdot \mathbf{I}_{KS_j}$	.5***	.44*	1.2***	1.6***
	(.12)	(.24)	(.24)	(.49)
$\mathbf{I}_{KS_j} \cdot \ln V_t^{\text{oil}}$	.041	.17**	.22***	.57***
	(.029)	(.068)	(.051)	(.13)
$\ln V_t$	-.0015	.0056	-.018	.078**
	(.0089)	(.01)	(.022)	(.03)
$\mathbf{I}_{\text{state}_j}$	-.021	2.1***	-.65***	.72
	(.096)	(.23)	(.18)	(.45)
$\mathbf{I}_{\text{policy}_j}$	-2.4***	-.93***	-2.1***	-.013
	(.13)	(.25)	(.18)	(.096)
$\ln price_t^{\text{oil}}$	.27***	.14***	.29***	.074**
	(.031)	(.042)	(.026)	(.029)
$\#Bidders_j$	-.0057***	.0099***	-.0053***	.011***
	(.0011)	(.0026)	(.0012)	(.0028)
$\ln AveTotalBid_j$	-.027	.95***	-.018	.95***
	(.023)	(.015)	(.022)	(.015)
$\ln ReqTons_j$	-.006**	-.97***	-.0055*	-.97***
	(.0029)	(.0099)	(.0029)	(.01)
$\ln OtherBidItems_{ij}$	.023		.015	
	(.021)		(.02)	
$\ln Distance_{ij}$	-.0071***	.007**	-.0072***	.0062*
	(.0026)	(.0034)	(.0026)	(.0033)
County f.e.	Y	Y	Y	Y
Year f.e.	Y	Y	Y	Y
Month of year f.e.	Y	Y	Y	Y
N	6107	4542	6107	4542
R^2	.92	.97	.92	.97

Note: This table reports regression estimates of the effect of the risk removal policy on an additional unit of oil price volatility on bids in Kansas relative to Iowa after vs before the policy (Equation 9). The sample size is smaller in regressions II and IV because only projects with bitumen bid totals $\geq$ \$50,000 are used. Standard errors clustered by firm. *** $p < .01$. $1998 \leq \text{Year} \leq 2012$.

Table 2.3: Robustness Tests of Triple Difference Estimation, Part I

Dependent variable: Log bitumen bid (b^A)							
	Interactions				Controls		
	I. None	II. Kansas-Policy	III. Policy-Vol	IV. Kansas-Vol	V. No covariates	VI. No month or county f.e.	VII. No year f.e.
$\mathbf{I}_{KS_j} \cdot \mathbf{I}_{Policy_t} \cdot \ln V_t^{\text{oil}}$					**-.21*****	**-.18*****	**-.17*****
					(.036)	(.038)	(.037)
$\mathbf{I}_{\text{policy}_t} \cdot \ln V_t^{\text{oil}}$			.74***		.61***	.85***	.81***
			(.042)		(.037)	(.04)	(.039)
$\mathbf{I}_{\text{policy}_t} \cdot \mathbf{I}_{KS_j}$		-.013			.64***	.59***	.55***
		(.016)			(.12)	(.13)	(.12)
$\mathbf{I}_{KS_j} \cdot \ln V_t^{\text{oil}}$				.0067	.072**	.047	.038
				(.02)	(.03)	(.031)	(.03)
$\ln V_t$	.053***		.0093	.04***	-.028***	.0074	.0045
	(.013)		(.0093)	(.015)	(.0087)	(.0088)	(.0095)
$\mathbf{I}_{\text{state}_j}$	.11***	.12***		.095	-.12	-.025	.0012
	(.011)	(.012)		(.065)	(.097)	(.1)	(.098)
$\mathbf{I}_{\text{policy}_j}$	.12***	.11***	-2.2***		-1.8***	-2.5***	-2.4***
	(.03)	(.032)	(.12)		(.12)	(.13)	(.13)
$\ln price_t^{\text{oil}}$	.053*	.0078	.32***	.025		.32***	.28***
	(.029)	(.029)	(.031)	(.029)		(.029)	(.033)
$\#Bidders_j$	-.006***	-.005***	-.006***	-.005***		-.001***	-.01***
	(.0012)	(.0012)	(.0012)	(.0012)		(.0012)	(.0013)
$\ln AveTotalBid_j$	.013	.017	-.1***	.0071		-.02	-.02
	(.024)	(.023)	(.021)	(.024)		(.024)	(.024)
$\ln ReqTons_j$	-.009***	-.009***	.008***	-.008***		-.011***	-.011***
	(.003)	(.0031)	(.0026)	(.003)		(.003)	(.003)
$\ln OtherBidItems_{ij}$	-.012	-.016	.079***	-.0073		.025	.025
	(.022)	(.021)	(.02)	(.022)		(.022)	(.022)
$\ln Distance_{ij}$	-.009***	-.01***	-.005	-.009***		-.012***	-.012***
	(.0026)	(.0027)	(.0029)	(.0027)		(.0032)	(.0032)
County f.e.	Y	Y	Y	Y	Y	N	Y
Year f.e.	Y	Y	Y	Y	Y	Y	N
Month of year f.e.	Y	Y	Y	Y	Y	N	Y
N	6107	6107	6107	6107	6107	6107	6107
R^2	.91	.91	.92	.91	.92	.92	.92

Note: This table reports estimates of the effect of the risk removal policy on an additional unit of historical oil price volatility on bids in Kansas relative to Iowa after vs before the policy (Equation 9). Standard errors clustered by firm. *** $p < .01$. $1998 \leq \text{Year} \leq 2012$.

Table 2.4: Robustness Tests of Triple Difference Estimation, Part II

Dependent variable:	Log bitumen bid $\left(b^A\right)$			Log non-bitumen bid items
	Placebo test at year			
	I: 2002	II. 2008	III. Parallel trends (before policy)	IV. Falsification
$\mathbf{I}_{KS_j} \cdot \mathbf{I}_{Policy_t} \cdot \ln V_t^{\text{oil}}$	-.096**	-.047		.058**
	(.04)	(.044)		(.024)
$\mathbf{I}_{\text{policy}_t} \cdot \ln V_t^{\text{oil}}$	-.035	.2***		-.14***
	(.026)	(.032)		(.028)
$\mathbf{I}_{\text{policy}_t} \cdot \mathbf{I}_{KS_j}$	.33**	.13		-.12
	(.13)	(.15)		(.078)
$\mathbf{I}_{KS_j} \cdot \ln V_t^{\text{oil}}$	.087**	.023	-.013	-.11***
	(.036)	(.026)	(.032)	(.02)
$\ln V_t$	.068***	-.019**	.024***	.0013
	(.018)	(.0079)	(.0089)	(.0084)
$\mathbf{I}_{\text{state}_j}$	-.18	.048	.16	.45***
	(.12)	(.083)	(.11)	(.07)
$\mathbf{I}_{\text{policy}_j}$	1.5***	.62***		.42***
	(.096)	(.12)		(.097)
$\ln price_t^{\text{oil}}$	.021	.087***	.36***	-.093***
	(.029)	(.028)	(.011)	(.021)
$\#Bidders_j$	-.0048***	-.004***	-.0015	.0039***
	(.0012)	(.0012)	(.0014)	(.0012)
$\ln AveTotalBid_j$	.0078	.0021	.061	1.1***
	(.024)	(.023)	(.047)	(.0046)
$\ln ReqTons_j$	-.0084**	-.0078**	-.0053	-.078***
	(.0033)	(.003)	(.0042)	(.0039)
$\ln OtherBidItems_{ij}$	-.0072	-.005	-.064	
	(.022)	(.021)	(.045)	
$\ln Distance_{ij}$	-.0092***	-.0086***	-.0078**	1.0e-07
	(.0027)	(.0027)	(.0034)	(.0024)
County f.e.	Y	Y	Y	N
Year f.e.	Y	Y	Y	Y
Month of year f.e.	Y	Y	Y	N
N	6107	6107	3528	6107
R^2	.91	.91	.7	.99

Note: This table reports estimates of the effect of the risk removal policy on an additional unit of historical oil price volatility on bids in Kansas relative to Iowa after vs before the policy (Equation 9). The sample size is smaller in regressions II and IV because only projects with bitumen bid totals ≥ \$50,000 are used. Standard errors clustered by firm. *** $p < .01$. 1998 ≤ Year ≤ 2012.

Table 2.5: Robustness Tests of Triple Difference Estimation, Part III

Dependent variable: Log bitumen bid $\left(b^A\right)$

Standard errors clustered by:	I. None (robust)	II. State-month	III. Firm-month	IV. Firm-month of year	V. Firm-state	VI. State-year
$\mathbf{I}_{KS_j} \cdot \mathbf{I}_{Policy_t} \cdot \ln V_t^{\text{oil}}$	**-.16*** **	**-.16** **	**-.16*** **	**-.16*** **	**-.16*** **	**-.16* **
	(.04)	(.072)	(.046)	(.044)	(.036)	(.09)
$\mathbf{I}_{\text{policy}_t} \cdot \ln V_t^{\text{oil}}$	.79***	.79***	.79***	.79***	.79***	.79***
	(.036)	(.12)	(.045)	(.05)	(.04)	(.13)
$\mathbf{I}_{\text{policy}_t} \cdot \mathbf{I}_{KS_j}$	.5***	.5**	.5***	.5***	.5***	.5
	(.13)	(.23)	(.15)	(.14)	(.12)	(.31)
$\mathbf{I}_{KS_j} \cdot \ln V_t^{\text{oil}}$	.041	.041	.041	.041	.041	.041
	(.034)	(.052)	(.039)	(.038)	(.029)	(.081)
$\ln V_t$	-.0015	-.0015	-.0015	-.0015	-.0015	-.0015
	(.0077)	(.023)	(.011)	(.01)	(.0089)	(.033)
$\mathbf{I}_{\text{state}_j}$	-.021	-.021	-.021	-.021	-.021	-.021
	(.11)	(.16)	(.13)	(.12)	(.096)	(.27)
$\mathbf{I}_{\text{policy}_j}$	-2.4***	-2.4***	-2.4***	-2.4***	-2.4***	-2.4***
	(.12)	(.38)	(.15)	(.16)	(.13)	(.41)
$\ln price_t^{\text{oil}}$	.27***	.27***	.27***	.27***	.27***	.27***
	(.021)	(.058)	(.028)	(.029)	(.031)	(.098)
$\#Bidders_j$	-.0057***	-.0057**	-.0057***	-.0057***	-.0057***	-.0057***
	(.00096)	(.0023)	(.0011)	(.0011)	(.0011)	(.002)
$\ln AveTotalBid_j$	-.027	-.027	-.027	-.027	-.027	-.027
	(.019)	(.028)	(.02)	(.02)	(.023)	(.028)
$\ln ReqTons_j$	-.006**	-.006	-.006**	-.006**	-.006**	-.006
	(.0024)	(.0047)	(.0025)	(.0024)	(.0029)	(.0054)
$\ln OtherBidItems_{ij}$	.023	.023	.023	.023	.023	.023
	(.018)	(.024)	(.019)	(.018)	(.021)	(.025)
$\ln Distance_{ij}$	-.0071***	-.0071***	-.0071***	-.0071***	-.0071***	-.0071***
	(.0021)	(.002)	(.0023)	(.0023)	(.0026)	(.0023)
County f.e.	Y	Y	Y	Y	Y	Y
Year f.e.	Y	Y	Y	Y	Y	Y
Month of year f.e.	Y	Y	Y	Y	Y	Y
N	6107	6107	6107	6107	6107	6107
R^2	.92	.92	.92	.92	.92	.92

Note: This table reports estimates of the effect of the risk removal policy on an additional unit of historical oil price volatility on bids in Kansas relative to Iowa after vs before the policy. Specifications are variants on Equation 9. Standard errors clustered as described. *** $p < .01$. $1998 \leq \text{Year} \leq 2012$.

Table 2.5 contains alternative assumptions about standard errors. Standard errors are not clustered at all in column I, and clustered by state-month, firm-month, firm-month of year, firm-state, and state-year in subsequent columns. The coefficient on the triple interaction maintains significance at the 1% level in all specifications except state-month (5% level) and state-year (10%) level. Appendix 2A Tables 3 and 4 show these alternative cluster assumptions using implied volatility and b^T. With implied volatility, significance is at the 1% level for all but state-month clusters (5%), and with b^T as the dependent variable significance declines to the 10% level with state-month clusters, and no significance for state-year clusters.

Many additional robustness checks are shown in Appendix 2A. For example, Table 2.5 columns I-II show that the coefficients do not change dramatically using 26 week historical volatility with the 6th month futures contract instead of 12 week volatility. Columns III and IV use the 5th month futures contract instead of the 6th month contract, and show slightly smaller effects with both b^A and b^T. Appendix 2A Table 6 shows specifications that exclude various covariates, and Table 7 shows that when 2008 - a year of unprecedented volatility - is excluded, the results are roughly the same, but somewhat larger with implied volatility.

2.5 Real Effects of the Policy in Kansas

In this section I explore the effects of the policy in Kansas in two ways. First, I examine whether auction outcomes for public and family owned firms changed after the policy, relative to private and non-family owned firms. Second, I calculate whether the Kansas government paid more for bitumen and for projects overall after the policy, given the realization of oil prices. Ideally I would also evaluate how the policy affected real variables like profitability and employment. Unfortunately I do not have access to such data, nor do I have the rich firm characteristics data that I have for Iowa.

I find that the policy seems to have benefited privately owned firms at the expense of publicly held firms, but that the policy had no apparent effect on family owned firms vis-a-vis non-family owned firms. Table 2.6 compares the win percentages (number of wins divided by number of bids) before and after the policy for all firms with at least 10 total bids.[84] The top panel shows that before the policy, publicly held firms won 36% of the auctions they bid in, while privately held firms won 33%. The t-test p-value for differences in means is 0.31. After the policy, the win percentages were

[84] I identified whether Kansas firms were publicly or privately owned by manually searching the internet.

18% for public firms and 27% for private firms, and the difference is significant with 99% confidence. The bottom panel conducts a similar analysis for family ownership, and shows that there was no difference in win percentages in either of the two periods.

Table 2.6: Kansas Firms' Win Percentages Before and After 2006 Risk Removal Policy

	Public Ownership		
Firm type	Win percentage before 2006 policy	Win percentage after 2006 policy	T-test p-value
Publicly listed	.36 (.00)	.18 (0.01)	**.05**
Privately owned	.33 (.03)	.27 (.04)	**.01**
T-test p-value	**.31**	**.01**	
N	37	29	
	Family Ownership		
Firm type	Win percentage before 2006 policy	Win percentage after 2006 policy	T-test p-value
Family owned	.34 (0.03)	.25 (.04)	**.16**
Not family owned	.35 (0.05)	.26 (.03)	**.21**
T-test p-value	**.78**	**.91**	
N	37	29	

Note: This table provides firm mean win percentages across all auctions in which the firm bids, as well as the p-values from t-tests for differences of means. These tests look across firm ownership types before and after the 2006 risk removal policy. I use only firms with at least 10 bids. The t-tests assume unequal variances.

After the policy there are fewer large bidders (firms with at least 10 bids) and they have a lower win percentage (Table 2.6). This is in part because the distribution of winning bids changed, shown in Figure 2.6. The bar heights indicate the number of firms in each category of auction win percentage. The distributions are strongly skewed left, but the skewness declines from 1.0 before the policy to 0.8 after the policy.[85] Kurtosis, or peakedness and fatness of tails, declines more dramatically from 4.9 to 3, where 3 is precisely the kurtosis of the normal distribution. This means that the "winningness" of firms became more evenly distributed across firms after the policy. The reverse occurred in the total number of bids by firm, shown in the bottom graph. Skewness and

[85] Skewness measures a distribution's symmetry, where a normal distribution has a skewness coefficient of 0. When the coefficient is positive, the median is less than the mean and the distribution is skewed right, and vice versa when it is negative. A skewness coefficient greater than 1 indicates that the distribution is highly skewed. Kurtosis measures the peakedness of the distribution, where the normal distribution has kurtosis of 3. Kurtosis greater than 3 has more observations closer to the mean and fatter tails than the normal distribution.

kurtosis, both much more extreme than percent wins, increase after the policy. Thus after the policy a few firms were submitting more of the bids, but nonetheless wins were more evenly spread across firms.

My second test answers a public finance question: Can a state lower its asphalt paving costs by removing oil price risk from the private sector? If firms are risk neutral or charge simply the CAPM-implied price of risk, then this policy should have been quite costly for Kansas over the course of my data because between 2006 and 2012 oil prices rose between auctions and work start overwhelmingly more than they declined. State governments, with plentiful access to finance, should be risk-neutral. Only if firms charge large risk premiums will the policy be beneficial.

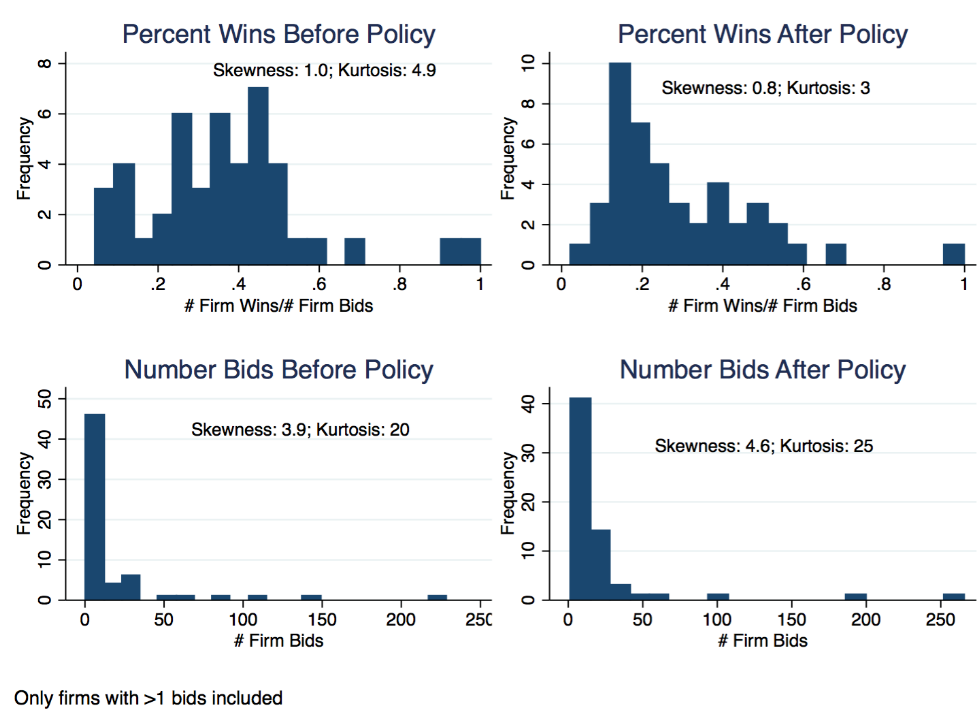

Figure 2.6: Kansas Firms Win Percentage Before and After the 2006 Risk Removal Policy
Note: The top two graphs in this figure show the frequency of of firms in various categories of win percentages. The bottom two graphs show the frequency of firms by the firms' total number of bids, with the firms sorted by their number of bids. The changing distributions indicate that after the policy a few firms were submitting more of the bids, but wins were more evenly spread across firms.

Using auction and payments data, I compare in Appendix 2B how much each state paid for bitumen after the 2006 risk-shifting policy. The mean bitumen bid was about the same in Iowa and Kansas prior to the policy intervention, at \$210 and \$205 per ton, respectively. Both states experienced cost escalation post-2006. On average Kansas paid \$489 per ton after the policy, whereas

Iowa paid \$513. A simple least squares differences-in-differences design finds that the policy reduced the cost of bitumen for Kansas by \$37 per ton, or 12% of the average bid over the period. This implies that over the 166 projects post-2006, Kansas saved \$5.1 million. Although the estimated benefit is small, it bears repeating that the policy should be *least* beneficial to Kansas when prices are increasing, which they did over most of the period.

2.6 Heterogeneity in Firm Risk Premium

Motivated by the evidence of risk aversion towards an idiosyncratic risk, in this section I present a simple model of how risk aversion could enter a firm's bid, which motivates a reduced form estimation. The diversity of paving firms provides a unique opportunity to examine how firm ownership and other characteristics impact risk management.

2.6.1 A Simple Model

Consider the following simple model describing a paving firm's profit maximization problem when bidding in a highway procurement auction. The firm's estimator (often a dedicated employee) submits unit price bids for bitumen, b_B, and for everything else, b_O. He knows the actual quantities that he will use, q_B^a and q_O^a, but his total bid B is calculated based on the DOT estimated quantities, q_B^e and q_O^e.[86] Conditional on his optimal total bid B, which determines his chances of winning the project (he wins if he submits the lowest total bid), the paver solves:

$$\begin{aligned} \max_{b_B, b_O} \pi &= b_B q_B^a + b_O q_O^a - \frac{1}{2}\eta \left[(b_B - \tilde{c}_B)^2 - (b_O - c_O)^2 \right] \\ \text{s.t.} \quad & b_B q_B^e + b_O q_O^e \leq B \end{aligned} \tag{10}$$

The firm's cost for each item is c. For bitumen, this is $\tilde{c}_B = c_B + \rho$, where ρ is a non-negative risk premium, or a value to the firm of hedging. I do not microfound this parameter; it may be due to financial constraints, agency problems, or owner preferences. The last term in Equation 10 reflects a penalty for excessive skewing. To the extent that $q^e \neq q^a$ on any item, the bidder has an incentive to skew his bid toward the quantity that has been underestimated. Then he stands a higher chance of winning (a lower B) but will in expectation be paid more based on q^a . Following Bajari, Houghton

[86] This simplifies the notion that the paving firm is more informed than the state about the quantities he will use.

and Tadelis (2010), I use a quadratic penalty η for for deviating from the engineering cost estimates.

The firm's FOC is:

$$\frac{\partial}{\partial b_B} : \; q_B^a - \lambda q_B^e = \eta \left(b_B - \tilde{c}_B \right) \tag{11}$$

where λ is a Lagrange multiplier, which I assume may be a function of everything that does not have to do with oil price risk, such as the marginal benefit of bidding a bid that scores higher, how others are skewing, etc.[87] Solving for the bid b_B gives:

$$b_B = \tilde{c}_B + \frac{1}{\eta} \left(q_B^a - \lambda q_B^e \right) = c_B + \rho + \frac{1}{\eta} \left(q_B^a - \lambda q_B^e \right) \tag{12}$$

We can think of this unit item bid as the paver's expected cost plus a markup that includes the paver's cost of risk. This markup is:

$$m_B = b_B - c_B = \rho + \frac{1}{\eta} \left(q_B^a - \lambda q_B^e \right). \tag{13}$$

Equation 13 leads to the reduced form estimation in Section 2.6.2.

2.6.2 Risk Premium Heterogeneity Estimation Strategy

I estimate the risk premium as the impact on bids of the forward market interacted with oil price volatility. The dependent variable is a large panel of proxy markups, $\hat{m}_B = b_B - c_B = \rho + \frac{1}{\eta} \left(q_B^a - \lambda q_B^e \right)$. I observe the bitumen bid b_B and proxy for c_B with the Argus "spot" index (B_t^{Argus}) reflecting the underlying cost of bitumen at the time of the auction.[88] The Argus "spot" index is closely correlated to the cost of oil (correlation of 0.79). The markup measure is thus: $\hat{m}_{B,j,i} = b_{B,j,i} - B_t^{Argus}$. Appendix 2A Figure 2 shows that as prices rose and became more volatile after, roughly, 2006, the difference between the bid unit item and the spot price transitioned from being tightly packed between \$10 and \$50 per ton to being more dispersed, with higher bids coincident

[87] Specifically, $\lambda = \frac{\eta}{(q_B^e)^2 + (q_O^e)^2} \left[q_B^e \tilde{c}_B + q_O^e c_O + \frac{1}{\eta} \left(q_B^a q_B^e + q_O^a q_O^e \right) - B \right]$

[88] Argus is a market research firm that kindly provided their price indices for the Eastern and Western regions of Iowa to me. Since there is no liquid bitumen market, Argus surveys transactions between suppliers and contractors by phone to get prices on bitumen for-delivery in the current week. These deliveries include purchases to pave commercial projects (e.g. a Wal-Mart parking lot), state projects (e.g. a highway), and by asphalt storage firms, who act as intermediaries between the refiners and the pavers. During the winter, when no actual paving is happening this index reflects only sales to the intermediaries. These intermediaries will be suppliers in the summer, alongside the refineries, to the paving firms.

with periods of high volatility.

I evaluate what portion of the markup can be attributed to oil price risk; specifically how a given level of oil price volatility affects bids at different distances in time from the work start date. A longer time between the auction and work start means that the contractor is taking on more oil price risk. Likewise, for a project that starts very soon after the auction there is no risk regardless of recent volatility. The forward market is the number of months between the auction and work start ($Wait_j$) interacted with oil price volatility $\left(V_t^{\text{oil}}\right)$. Unfortunately I do not have enough data to estimate ρ for each firm. Instead, I test whether ρ is larger for different kinds of firms, and am also able to assess whether a larger ρ for a given group of firms is driven by $Wait_j$ or V_t^{oil}. The estimating equation for the public-private analysis, where where I_{Public_i} indicates that the firm is publicly listed, is below. The unit of observation is project j auctioned at time t.

$$\begin{aligned}\hat{m}_{B,j,i} &= \alpha + \rho I_{\text{Public}_i} \cdot Wait_j \cdot \ln V_t^{\text{oil}} + \gamma' \cdot \text{Auction/Bidder Chars} \\ &+ \beta_1 I_{\text{Public}_i} + \beta_2 \ln V_t^{\text{oil}} + \beta_3 Wait_j + \beta_4 I_{\text{Public}_i} \cdot Wait_j + \beta_5 \mathbf{I}_{Public_i} \cdot \ln V_t^{\text{oil}} + \beta_6 Wait_j \cdot \ln V_t^{\text{oil}} \\ &+ \delta \mathbf{I}_{month_t} + \delta \mathbf{I}_{year_t} + \epsilon_{ijt} \end{aligned} \tag{14}$$

where:

$$\begin{aligned}\gamma' \cdot \text{Auction/Bidder Chars} &= \gamma_1 \#Bidders_j + \gamma_2 \ln AveTotalBid_j + \gamma_3 \ln ReqTons_j \\ &+ \gamma_4 \ln Distance_{ij} + \gamma_5 \text{Firm Size (Emp)}_i + \gamma_6 \text{Firm Size (Rev)}_i\end{aligned}$$

The specification includes double interactions and individual effects of $Wait_j$, V_t^{oil}, and I_{Public_i}. Other specifications replace I_{Public_i} with a different firm characteristic, and the regressions on firm size exclude those variables from the controls. I control for year and month of the year fixed effects (the latter is especially important because of capacity constraints that firms face as the construction season progresses). I control for firm size using both employment and revenue variables. I cluster standard errors by firm. Table 2.7 contains descriptive statistics for the variables used in this analysis.

Table 2.7: Iowa Risk Premium Heterogeneity Analysis Summary Statistics

Panel 1: Firm Characteristic Variables

Variable	Description	N	Mean	Std. Dev.	Min	Max	Source
$\mathbf{I}_{\text{Public}_i}$	Publicly Listed	8207	0.08	0.272	0	1	CIQ
$\mathbf{I}_{\text{Family}_i}$	Family-owned and managed	8125	0.71	0.46	0	1	HC
$\#\text{SIC}_i$	# 8-digit SIC codes	8207	2.1	1.5	1	8	DB
$\mathbf{I}_{\text{Not Div}_i}$	Only 1 8-digit SIC code	8207	0.55	0.50	0	1	DB
$\mathbf{I}_{\text{Paving Primary}_i}$	Paving asphalt is primary business	8207	0.74	0.44	0	1	DB, HC
$\mathbf{I}_{\text{Small (Emp)}_i}$	Small business ($\leq$ 100 employees)	8162	0.56	0.49	0	1	DB, CIQ
$\mathbf{I}_{\text{Small (Rev)}_i}$	Small business ($\leq$ \$15 mill in revenue)	8095	0.60	0.49	0	1	DB, CIQ
Firm Size (Emp)$_i$	# employees	8162	304	644	1	14700	DB, CIQ
Firm Size (Rev)$_i$	Annual Revenue (\$)	8095	66	370	0.02	13563	DB, CIQ
$\mathbf{I}_{\text{Related}_i}$	Owners or officers related (blood or marriage) with another firm	8207	0.09	0.29	0	1	IDOT
$\mathbf{I}_{\text{Subsid}_i}$	Subsidiary of another firm	8207	0.48	0.50	0	1	IDOT
$\mathbf{I}_{\text{JV}_i}$	Has JV or partnership with another firm	8207	0.19	0.40	0	1	IDOT

Panel 2: Bid Variables

Variable	Description	N	Mean	Std. Dev.	Min	Max	
$\hat{m}_{B,j,i}$	Markup of bid over argus "spot" price	8207	52	49	-259	575	IDOT
$Wait_j$	Months between auction and work start	8087	4.7	2.8	0	17	IDOT
$\ln V_t^{\text{oil}}$	Log of 12-week oil price volatility of CL6	8207	3.2	0.39	2.2	4.4	IDOT
N_j	# bidders in auction	8207	4.4	2.3	1	12	IDOT
$\ln T_j$	Log estimated tons bitumen for project	8207	5.5	2.0	-1.9	9.3	IDOT
$\ln M_{ij}$	Log miles between firm primary address and project site	8207	4.0	1.1	-1.4	8.8	IDOT
$\ln \bar{b}_j$	Log average total bid in auction	8207	14	1.1	9.2	18	IDOT

Note: This table provides summary statistics of data used in the heterogeneity analysis. Sources are CapitalIQ (CIQ), Dunn & Bradstreet Hoovers (DB), the Iowa DOT (IDOT) and hand collection via the web (HC) (primarily company websites, the Iowa Department of State, and business record websites like Manta).

The coefficient ρ on the impact of the forward market interacted with oil price volatility should reflect only differences in firm bids that relate to oil price risk. I interpret ρ as a measure of the oil price risk premium, but ρ does not provide information about underlying risk preferences, nor does it indicate how firms are hedging. If some firms hedge efficiently using oil futures, their estimated risk premium may be much lower than a firm who always hedges through a supplier. If firms manage this idiosyncratic risk by passing it to financial markets, there should be no (or a very small) premium. Differences in ρ across firms means that some firms either are more risk averse or hedge less efficiently. In either case, they pass a premium for bearing diversifiable risk to the consumer (the government).

2.6.3 Observable Firm Characteristics

The primary reason that most research on firm risk management has focused on publicly listed companies is because little data exists about private companies. My Iowa data provides a unique, albeit incomplete, window into the relationship of firm characteristics to risk management. I do not have firm-specific data for Kansas, so this heterogeneity analysis uses only Iowa data.

The top panel of Table 2.7 provides summary statistics of the characteristics. The first variable is a dummy categorical variable for a firm being publicly held ($\mathbf{I}_{\text{Public}_i}$), which changes from 0 to 1 over time for the firms that are acquired by a public company. The family firm variable, also a binary dummy ($\mathbf{I}_{\text{Famil}y_i}$), indicates that the company meets at least one of three conditions: a) the company states it is family-owned and managed on its website; b) records of executives indicate that a President or CEO has the same name as the company; or c) more than 2 top executives share the same last name.

I use the number of 8-digit SIC codes a firm does business in as a measure of diversification ($\#\text{SIC}_i$), where a higher number in the index indicates greater diversification. Firms that are less diversified may have more of the managers' wealth tied to oil prices and thus may exhibit greater risk aversion. This is obviously a crude measure of diversification, and additional SIC code industries could also be very oil-intensive. However, examination of the SIC codes suggests that additional codes are usually other types of construction, such as metalwork and sewer lines, that are less oil intensive. I also use two dummies for a firm not being diversified: $\mathbf{I}_{\text{Not Div}_i}$ takes a value of 1 if the firm operates in only one SIC code, and 0 if more; and $\mathbf{I}_{\text{Paving Primary}_i}$ takes a value of 1 if asphalt paving is the firm's primary activity, and 0 otherwise. If a firm is primarily a bridge builder, they are likely less exposed to oil price shocks than a firm that primarily paves.

My metrics for business size are the number of employees and annual revenue. Unfortunately, I only observe these for 2012 or the latest year the company was active. Thus these variables are a very rough indication of firm size. However, the stability of this industry makes these single-year measures more appropriate than they would be for other sectors (Porter and Zona 2003). I use two continuous and two binary measures. Firm Size $(\text{Emp})_i$ is the raw number of employees and Firm Size $(\text{Rev})_i$ is annual revenue. $\mathbf{I}_{\text{Small (Emp)}_i}$ is 1 for firms with less than 100 employees and 0 otherwise; $\mathbf{I}_{\text{Small (Rev)}_i}$ is 1 for firms with less than \$15 million in revenue and 0 otherwise (these are the 40th percentiles of the respective samples).

Finally, I exploit a database from the Iowa DOT describing affiliations among its contractors. The variables are a dummy variables for whether a firm's owners or officers are related by blood or marriage to another contractor $\left(\mathbf{I}_{\text{Related}_i}\right)$, whether a firm is a subsidiary company to another contractor $\left(\mathbf{I}_{\text{Subsid}_i}\right)$, and whether it has a joint venture or partnership with another contractor $\left(\mathbf{I}_{\text{JV}_i}\right)$. Unfortunately, I have no measure of leverage.

All pairwise correlations among characteristics are in Appendix 2A Table 9, where correlation coefficients significant at the 5% level or better are starred. The correlation between $\mathbf{I}_{\text{Family}_i}$ and the diversification variables is quite low, as is the correlation between $\mathbf{I}_{\text{Public}_i}$ and the diversification variables. For example, the correlation between $\mathbf{I}_{\text{Family}_i}$ and the $\mathbf{I}_{\text{Paving Primary}_i}$ (indicator for paving being the firm's primary activity) is only 0.02. Only two correlations are greater that 0.5 in either direction: Firm Size $(\text{Emp})_i$ with $\mathbf{I}_{\text{Public}_i}$ (0.52*), Firm Size $(\text{Emp})_i$ with $\#\text{SIC}_i$ (0.67*), $\mathbf{I}_{\text{Small (Rev)}_i}$ with $\mathbf{I}_{\text{Small (Emp)}_i}$ (0.88*), and Firm Size $(\text{Rev})_i$ with Firm Size $(\text{Emp})_i$(0.68*).

For a more granular sense of the data, Appendix 2A Table 8 shows the top 30 bidders' number of bids, win percentage, public ownership status, family ownership status, and first and last bid date. The publicly owned firms were all originally private firms acquired by either MDU Resources, an energy and infrastructure company, or Oldcastle Materials, the US arm of CRH Plc. Nearly all 30 firms submit bids spanning the full period (1994-2012), and most were founded decades before 1994. The remaining 185 firms not included in this table collectively have 1,924 bids.

2.6.4 Risk Premium Heterogeneity Estimation Results

Public vs Private Firm Ownership

Theory suggests that privately held firms are likely to be more risk averse than publicly held firms. The rationale from FSS implies that if external finance is more costly for private firms than for public firms, managing risk will be more valuable to private firms. On the other hand, Rampini and Viswanathan (2013) and Rampini, Sufi and Viswanathan (2014) show that constrained firms may face greater costs to hedging, so regardless of their risk aversion, they may hedge less. A different line of argument is that private firms might be more risk averse if their owners are less diversified than the shareholders of public companies (Panousi and Papanikolaou 2012, Jin and Jorion 2006). However, agency problems could push the other direction. Following Stulz (1984) and Asker et al. (2012), if publicly held firms are subject to managerial agency problems and their managers are risk averse, there may be no significant difference between public and private firms.

Estimates of Equation 14 are shown in Table 2.8. The coefficient of interest ρ is on the triple interaction in the first row, $\left(\mathbf{I}_{\text{Public}_i} = 1\right) \cdot Wait_j \cdot \ln V_t^{\text{oil}}$. The measure of risk is the interaction $Wait_j \cdot \ln V_t^{\text{oil}}$. Employing interactions (also referred to as moderating or conditioning effects) in this regression implies that coefficients should be interpreted as being relative to the base, or reference level. In the case of $\mathbf{I}_{\text{Public}_i}$, for example, the coefficient is relative to being privately held.[89]

I find a strong negative coefficient on the triple interaction, indicating that publicly held firms charge the state a lower risk premium than privately held firms. Each unit increase in the measure of risk reduces the bid markup for public firms relative to private firms by about $5. Although the negative and significant value for ρ is informative, a more meaningful interpretation arises from graphical representation of the marginal effects of the triple interaction. I fix either $Wait_j$ or $\ln V_t^{\text{oil}}$ and allow the other to move, and then calculate the conditional marginal effect of $\mathbf{I}_{\text{Public}_i}$ on the markup. Figure 2.7 shows in the left panel oil volatility held at its 50th percentile level. As time-to-start moves from 1 month to 14 months, the effect of being publicly listed becomes strongly negative. At 14 months a publicly listed firm bids on average $20 less than a privately held firm. The average markup is $53. Figure 2.8 shows this latter effect more clearly. The graphs fix oil price volatility at its 10th, 40th, 60th, and 90th percentiles. The "forward curve" of the impact

[89] Note that the coefficients on the controls, such as on $\left(\mathbf{I}_{\text{Public}_i} = 1\right) \cdot Wait_j$, are not general effects. For example, the coefficient on $Wait_j$ reflects the impact of $Wait_j$ when $\mathbf{I}_{\text{Public}_i}$ is zero and oil price volatility is zero. The coefficient on $\left(\mathbf{I}_{\text{Public}_i} = 1\right) \cdot Wait_j$ is the effect of the interaction between being public and months-to-start when oil price volatility is zero. That is, these are conditional effects that give the covariates' impact when other variables involved in the interaction are zero, and are different from unconditional effects.

of months-to-start rotates counterclockwise as oil price volatility increases.

Table 2.8: Markup Analysis - Impact of Public and Family Ownership

Dependent Variable: Estimated Markup ($\hat{m}_{B,j,i}$)

	I. Impact of public ownership		II. Impact of family ownership
$\mathbf{I}_{\text{Public}_i} = 1 \cdot Wait_j \cdot \ln V_t^{\text{oil}}$	**-5****	$\mathbf{I}_{\textbf{Family}_i} = 1 \cdot Wait_j \cdot \ln V_t^{\textbf{oil}}$	**2.7****
	(2)		(1.2)
$\mathbf{I}_{\text{Public}_i} = 1 \cdot Wait_j$	14**	$\mathbf{I}_{\text{Family}_i} = 1 \cdot Wait_j$	-8.3**
	(6.4)		(3.7)
$\mathbf{I}_{\text{Public}_i} = 1 \cdot \ln V_t^{\text{oil}}$	44***	$\mathbf{I}_{\text{Family}_i} = 1 \cdot \ln V_t^{\text{oil}}$	-7.5
	(14)		(8.6)
$Wait_j \cdot \ln V_t^{\text{oil}}$	-1.3***	$Wait_j \cdot \ln V_t^{\text{oil}}$	-3.6***
	(.44)		(1.1)
$\mathbf{I}_{\text{Public}_i} = 1$	-136***	$\mathbf{I}_{\text{Family}_i} = 1$	26
	(43)		(25)
$Wait_j$	4.6***	$Wait_j$	12***
	(1.4)		(3.5)
$\ln V_t^{\text{oil}}$	9.2***	$\ln V_t^{\text{oil}}$	18**
	(2.7)		(7.9)
$\#Bidders_j$	-.63*	$\#Bidders_j$	-.66**
	(.33)		(.33)
$\ln ReqTons_j$	-6.5***	$\ln ReqTons_j$	-6.5***
	(.58)		(.55)
$\ln Distance_{ij}$	-1.1	$\ln Distance_{ij}$	-1.3
	(1.1)		(1.2)
$\ln AveTotalBid_j$	3***	$\ln AveTotalBid_j$	3***
	(.8)		(.8)
Firm Size (Emp)$_i$	.00072	Firm Size (Emp)$_i$	.0013
	(.0029)		(.004)
Firm Size (Rev)$_i$	.0004	Firm Size (Rev)$_i$	-.000025
	(.0033)		(.0044)
Year f.e.	Y	Year f.e.	Y
Month-of-year f.e	Y	Month-of-year f.e	Y
N	7970	N	7927
R^2	.49	R^2	.48

Note: This table reports results from the markup estimation in Equation 14. The coefficient of interest on the triple interaction gives the impact of risk, measured as the interaction between the time-to-start and oil volatility, for public firms relative to private firms in column I, and family owned relative to non-family owned firms in column II. Standard errors clustered by firm. *** $p < .01$. $1994 \leq \text{Year} \leq 2012$.

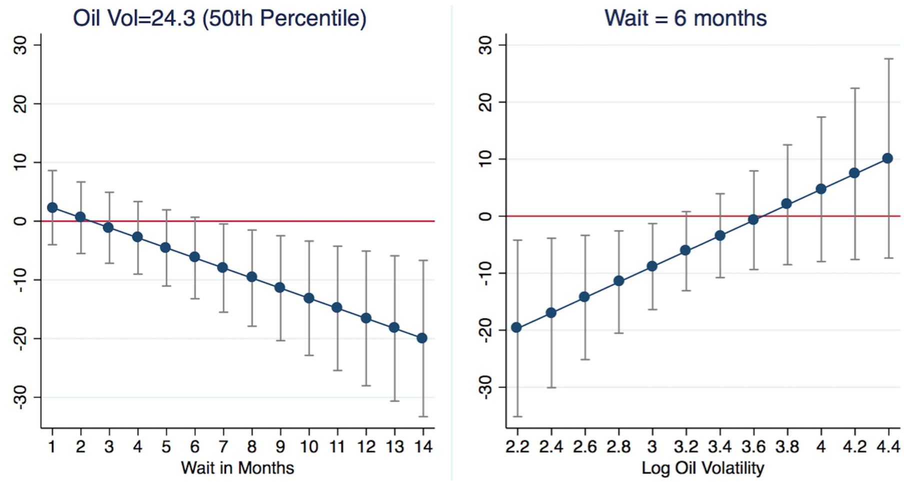

Figure 2.7: Conditional Marginal Effect on Bitumen Bid Markup of Public Ownership
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.8 Column I). In the left graph volatility is fixed at its 50th percentile, and the y-axis indicates the conditional marginal effect of a firm being publicly rather than privately owned as the time to start increases. In the right graph the time to start is fixed at its mean and the same conditional effect is calculated as volatility increases. 95% confidence intervals shown.

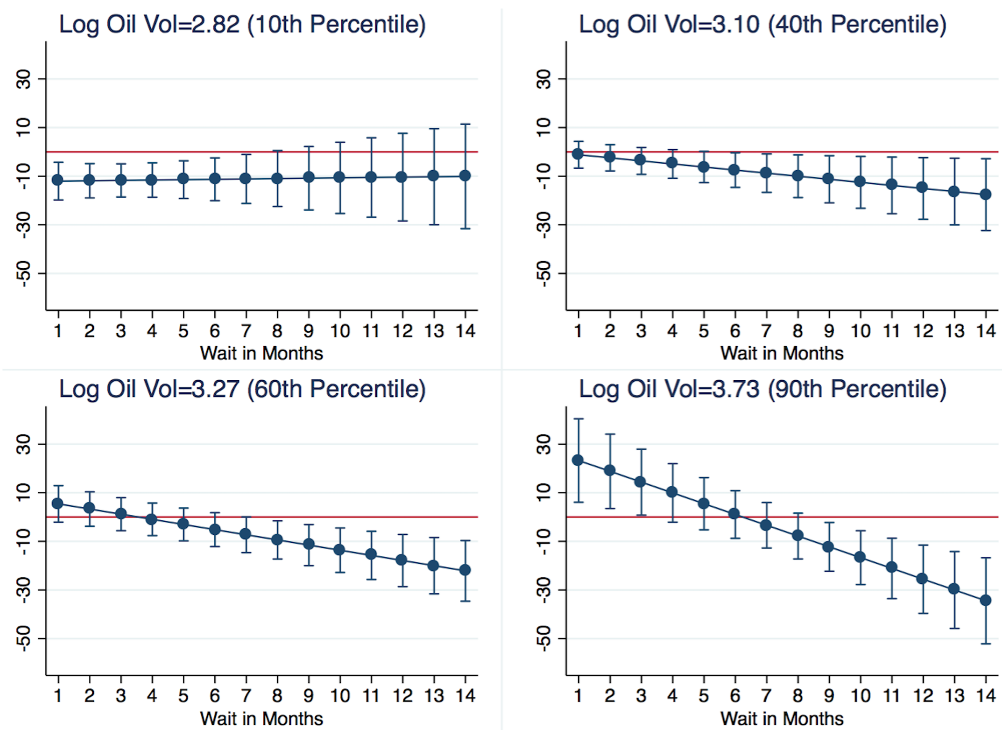

Figure 2.8: Conditional Marginal Effect on Bitumen Bid Markup of Public Ownership
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.8 Column I). The graphs fix volatility at its 10th, 40th, 60th, and 90th percentiles. They show the conditional marginal effect of a firm being publicly rather than privately owned as the time to start increases. 95% confidence intervals shown.

The right panel of Figure 2.7 fixes time-to-start at 6 months. As oil price volatility increases, the difference between public and private firms narrows, and at high levels of volatility public firms bid higher than private firms. This suggests that although the interaction coefficient is large and significant, the overall negative impact of being public is mostly driven by the time-to-start. Private firms' risk premium is more associated with the delay, suggesting they pay less attention to recent oil price volatility.

In part because I do not directly observe the paver-supplier contracts, I am unable to explain *why* different types of firms have different risk management behavior. It is possible that public firms are fully hedging on financial markets, but this is much cheaper than hedging with suppliers, and thus it appears that private firms are more risk averse, when in fact they either are less sophisticated or do not have the scale to hedge efficiently in financial markets. It is also possible that publicly held firms have greater bargaining power with suppliers than private firms, and this bargaining power varies with time-to-start. However, the total volume of bitumen used, as well as average project size, does not differ substantially across public and private firms (nor across the other dimensions, with the exception of firm size). Thus the bargaining power hypothesis seems less plausible.

Family vs Non-Family Firm Ownership

I conduct a similar exercise for family-owned and managed firms. On one hand, if the owners of family firms seek to maximize personal utility and smooth income, the family firms may be more risk averse (Bertrand and Schoar 2006, Shleifer and Vishny 1986, Schulze et al. 2001). On the other hand, if family firms use intense monitoring to overcome manager-shareholder agency problems, they may do less risk management (Anderson and Reeb 2003, Schmid et al. 2008).

The main specification, in Table 2.8 column II, reveals that family owned firms charge substantially larger risk premiums than non-family owned firms. Graphically, Figure 2.9 uses marginal effects to show that it is oil price volatility, not time-to-start, that drives this result. The left panel shows that the risk premium does not vary with time-to-start, while the right panel shows that at the average time-to-start the risk effect is sharply increasing in oil price volatility. This contrasts with the public vs. private analysis. However, this result is somewhat imprecise; note that the 95% confidence interval does not exclude zero in either panel. In Figure 2.10, I show that it does exclude zero in high oil price volatility environments. Table A5 has a waterfall series of regressions.

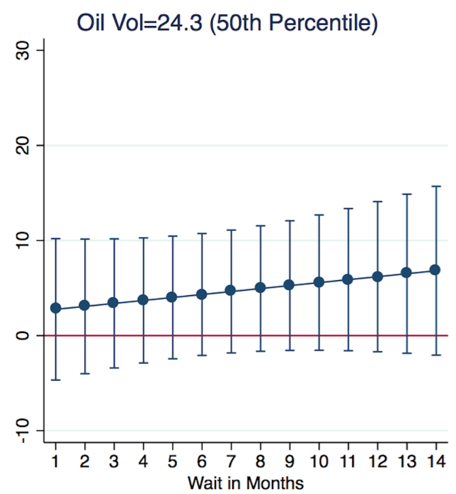

Figure 2.9: Conditional Marginal Effect on Bitumen Bid Markup of Family Ownership
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.8 Column II). In the left graph volatility is fixed at its 50th percentile, and the y-axis indicates the conditional marginal effect of a firm being family rather than non-family owned as the time to start increases. In the right graph the time to start is fixed at its mean and the same conditional effect is calculated as volatility increases. 95% confidence intervals shown.

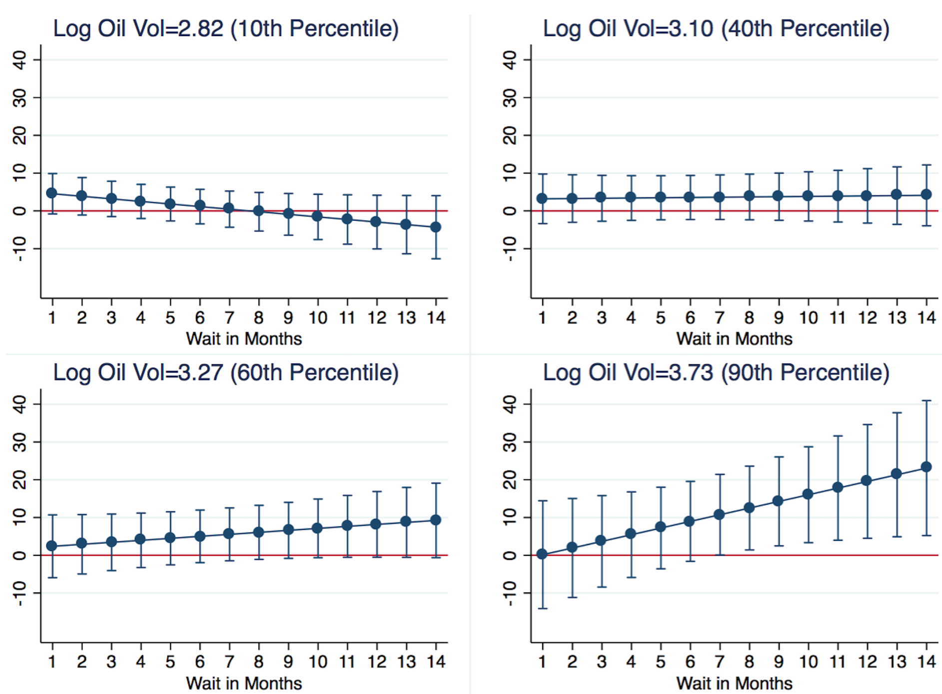

Figure 2.10: Conditional Marginal Effect on Bitumen Bid Markup of Family Ownership
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.8 Column II). The graphs fix volatility at its 10th, 40th, 60th, and 90th percentiles. They show the conditional marginal effect of a firm being family rather than non-family owned as the time to start increases. 95% confidence intervals shown.

Diversified Firms vs. Concentrated Firms

Firms may manage oil price risk by diversifying to non-oil intensive sectors. Well-diversified firms in general may be less risk averse (Mackay and Moeller 2007, Faccio et al. 2011). Panel 1 of Table 2.9 shows that diverse firms charge consistently and robustly lower risk premiums than not diverse firms. (Table 2.9 shows only the coefficient ρ on the triple interaction term and graphical marginal effects. For the full tables, see Appendix 2A tables 10, 11, 12, and 13.)The continuous variable for diversification has a strong negative slope, and the two categorial variables for being concentrated in one industry and for paving asphalt roads being the firm's primary activity have strong positive impacts. Figures 2.11 and 2.12 show that non-diverse firms (operating in only 1 8-digit SIC code) have higher risk premiums. This result has strong statistical significance. As with family ownership, the marginal effects graphs show that this effect primarily increases not with months-to-start, but with oil price volatility.

Small Firms vs. Large Firms

Small firms usually have less assets than large firms to use as collateral, and therefore are often assumed to face more severe financing constraints (FSS, Hennessy and Whited 2007, Vickery 2008, Nance et al. 1993). However, some empirical literature finds that small firms do less risk management, possibly due to economies of scale in hedging (e.g. Stulz 1996). Alternatively, Rampini and Viswanathan (2013) and Rampini, Sufi and Viswanathan (2014) theorize that collateral constraints may cause hedging to be suboptimal for small firms.

The tests of firm size in Panel 2 of Table 2.9 suggest that larger firms are less risk averse, but the effects are smaller and much weaker than some of the other characteristics I test. The strongest result is from the continuous variable for firm size by revenue. Figures 2.13 and 2.14 reveal that this effect is both a function of time-to-start and oil price volatility; that is, the marginal effects line slopes significantly downward in both time-to-start and oil price volatility when the other variable is fixed.

Table 2.9: Markup Analysis - Diversification, Firm Size, and Relationships

Dependent Variable: Estimated Markup ($\hat{m}_{B,j,i}$)

Panel 1: Industrial Sector Diversification

	I.	II.	III.	
Diversification Variable:	# SIC codes	**1** | Not Diversified (1 SIC code)	**1** | Paving Primary Activity	
Div Var$_i \cdot Wait_j \cdot \ln V_t^{\text{oil}}$	**-.77****	**2.1****	**2.7****	
	(.33)	(.97)	(1.3)	

Panel 2: Impact of Firm Size

	IV.	V.	VI.	VII.
Size Variable:	**1** | Small (Emp)	**1** | Small (Rev)	Firm Size (Emp)	Firm Size (Rev)
Size Var$_i \cdot Wait_j \cdot \ln V_t^{\text{oil}}$	**1.7***	**.43**	**-.0019***	**-.0081****
	(.99)	(.98)	(.00099)	(.0035)

Panel 3: Relationship to other Iowa Contractors

	VIII.	IX.	X.	
Relation Variable:	**1** | Related	**1** | Subsidiary	**1** | JV	
Rel Var$_i \cdot Wait_j \cdot \ln V_t^{\text{oil}}$	**1.8***	**2.7****	**.32**	
	(1)	(1.1)	(.98)	

Panel 4: Interaction of Family Ownership with Diversification, Size and Relationship

	XI.	XII,	XIII.	XIV.
Family Variable:	**1** | Family·**1** | Not Diverse	**1** | Family·**1** | Related	**1** | Family·**1** | Subsidiary	**1** | Family·**1** | Small (Emp)
Fam Var$_i \cdot Wait_j \cdot \ln V_t^{\text{oil}}$	**2.5*****	**2.2****	**2.7*****	**2.5*****
	(.93)	(.95)	(.92)	(.9)

Note: This table reports the coefficient of interest from the markup estimation in Equation 14. For example, in Column I the coefficient of interest on the triple interaction gives the impact of risk, measured as the interaction between the time-to-start and oil volatility, and the degree of firm diversification, measured by the number of SIC codes in which the firm is active. See Appendix 2A for the full specification. Standard errors clustered by firm. *** $p < .01$. $1994 \leq \text{Year} \leq 2012$.

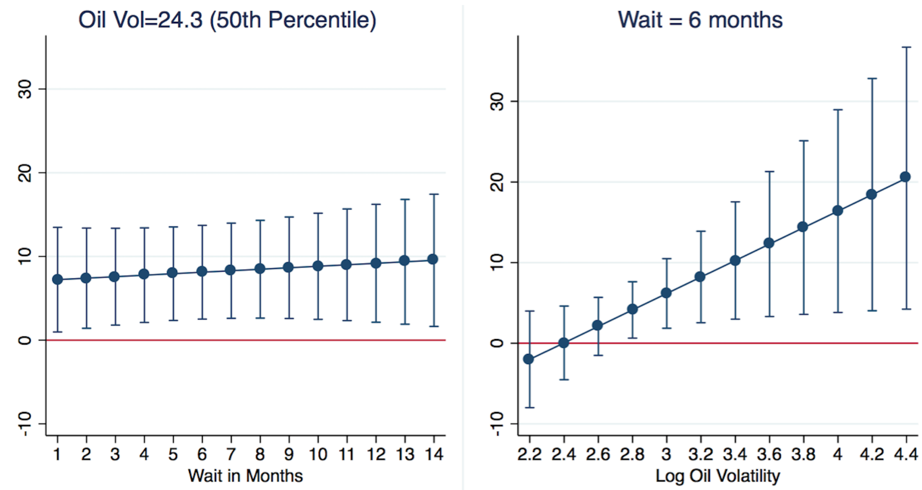

Figure 2.11: Conditional Marginal Effect on Bitumen Bid Markup of Firm Not Diversified
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.9 Column II). In the left graph volatility is fixed at its 50th percentile, and the y-axis indicates the conditional marginal effect of a firm being not diversified (defined as being active in only one SIC code). In the right graph the time to start is fixed at its mean and the same conditional effect is calculated as volatility increases. 95% confidence intervals shown.

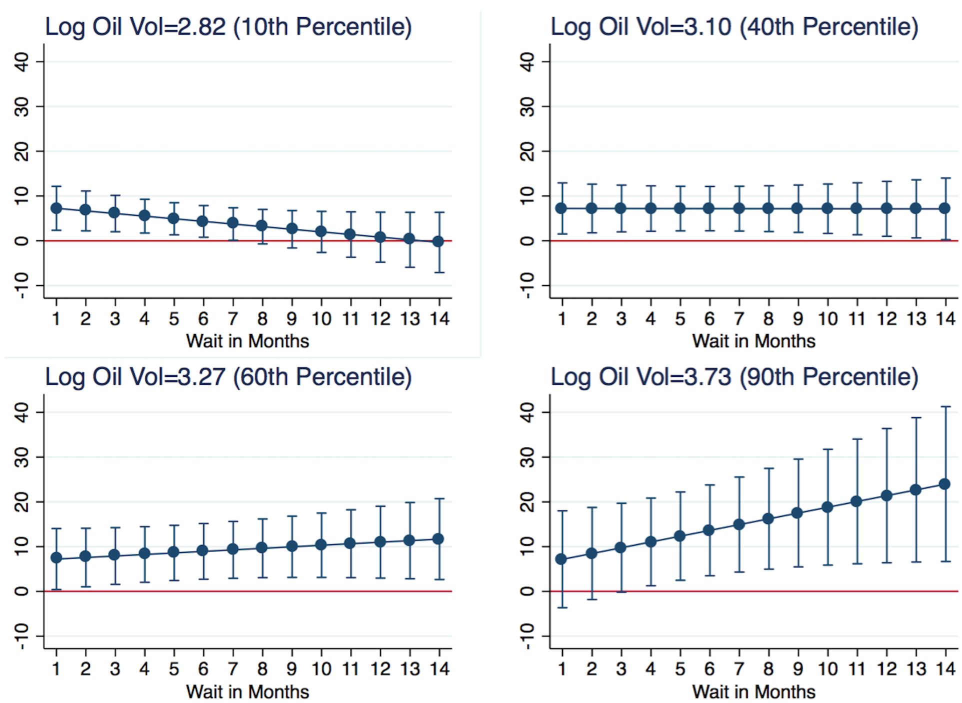

Figure 2.12: Conditional Marginal Effect on Bitumen Bid Markup of Firm Not Diversified
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.9 Column II). The graphs fix volatility at its 10th, 40th, 60th, and 90th percentiles. They show the conditional marginal effect of a firm being not diversified (defined as being active in only one SIC code) as the time to start increases. 95% confidence intervals shown.

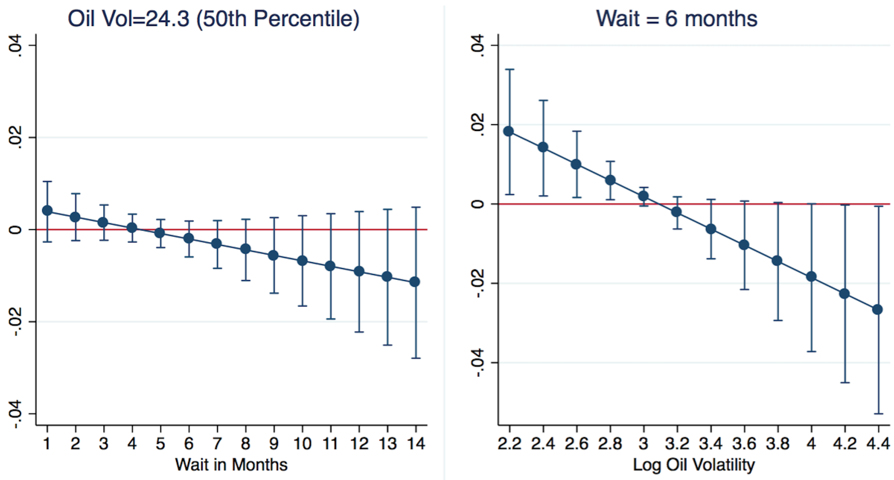

Figure 2.13: Conditional Marginal Effect on Bitumen Bid Markup of Firm Size (Revenue)
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.9 Column VII). In the left graph volatility is fixed at its 50th percentile, and the y-axis indicates the conditional marginal effect of log firm revenue. In the right graph the time to start is fixed at its mean and the same conditional effect is calculated as volatility increases. 95% confidence intervals shown.

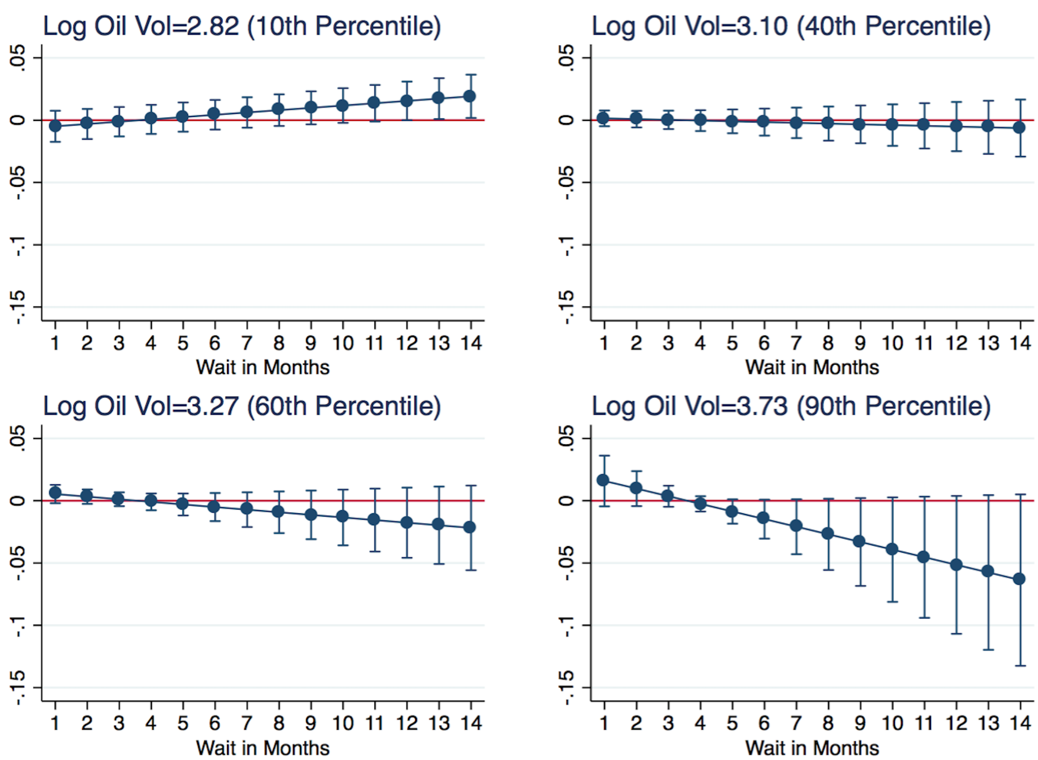

Figure 2.14: Conditional Marginal Effect on Bitumen Bid Markup of Firm Size (Revenue)
Note: This figure shows marginal effects of oil price risk on bitumen bid markup (Equation 14, results in Table 2.9 Column VII). The graphs fix volatility at its 10th, 40th, 60th, and 90th percentiles. They show the conditional marginal effect of log firm revenue as the time to start increases. 95% confidence intervals shown.

Inter-Firm Relations

Subsidiary firms and firms whose owners and officers are related to those at other Iowa contractors may have greater access to capital and more diversified owners, leading to less risk management. Panel 3 of Table 2.9 examines the relationship variables. I find that firms whose owners and officers are related to those at other Iowa contractors exhibit *greater* risk premiums. This is potentially because of the large overlap between family ownership and the relationship variables. The largest risk premiums are among firms that are both family owned and fall into one of the following categories: non-diversified, have relatives at a another firm, are a subsidiary, and are small. Coefficients on these interaction terms are shown in Panel 4 of Table 2.9. Figures 2.15 and 2.16 show the conditional marginal effects for being both family-owned and having owners/officers related to those at another contractor.[90]

Robustness

Robustness tests for the risk premium heterogeneity results are in Appendix 2A. For example, Tables 14 and 15 show a variety of alternative standard error assumptions for the public and family ownership estimations. Standard errors are not clustered at all in column I, and clustered by state-month, firm-month, firm-month of year, firm-state, and state-year in subsequent columns. For both public and family ownership, the coefficient on the triple interaction remains significant. For public ownership, significance drops to the 10% level with state-month, firm-month, and state-year clusters, but for family ownership significance drops to the 10% level only for state-month clusters. Similar results hold for the other firm characteristics, including the diversification variables.

[90] There are a number of additional firm characteristics that I tested, but found no significant effects. A binary variable for whether the firm's primary state is Iowa yielded a positive but not significant impact on the coefficient of interest (the triple interaction proxy for risk). Similarly, a continuous variable for the distance between the firm's primary address and the project produced a negative but insignificant result. A variable for the firm's age at bid yielded a positive but near-zero insignificant coefficient. The Dunn & Bradstreet database contains an interesting variable on whether the firm owns or rents its primary facility. In theory, a firm that owns its facility has more collateral and might be expected to be less risk averse. Unfortunately, the database only has this variable for about half of my firms. Using this subsample, the regression yielded a large positive coefficient on ownership, but without any significance.

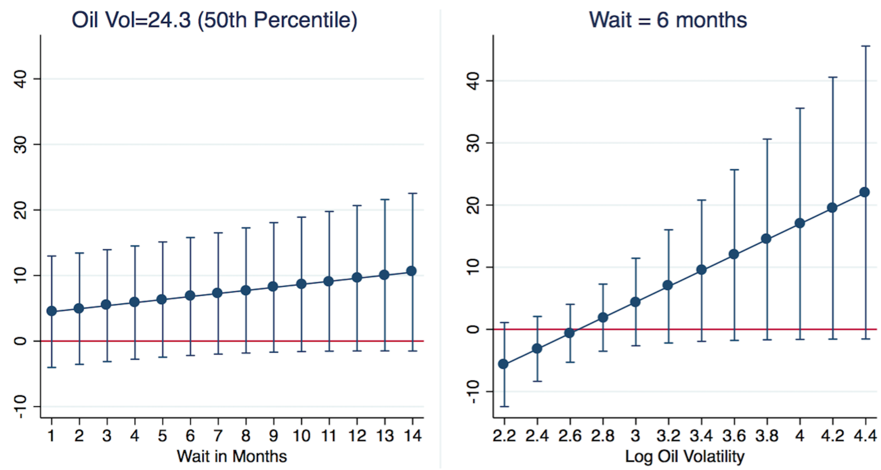

Figure 2.15: Conditional Marginal Effect on Bitumen Bid Markup of being Family Owned & Owners Related to Another Firm
Note: This figure shows marginal effects oil price risk on bitumen bid markup (Equation 14, results in Table 2.9 Column XII). In the left graph volatility is fixed at its 50th percentile, and the y-axis indicates the conditional marginal effect of a firm being family rather than non-family owned as the time to start increases. In the right graph the time to start is fixed at its mean and the same conditional effect is calculated as volatility increases. 95% confidence intervals shown.

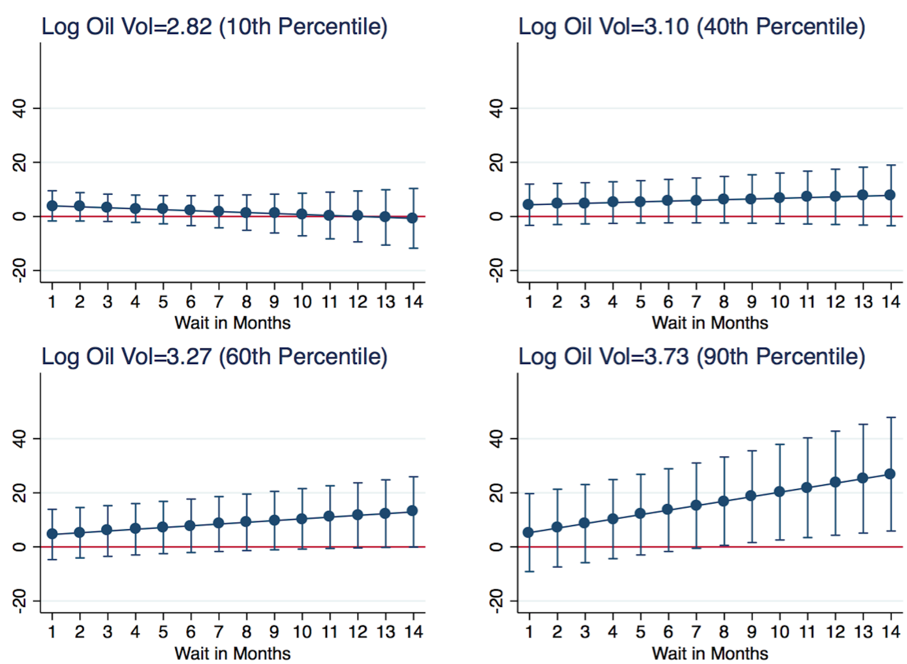

Figure 2.16: Conditional Marginal Effect on Bitumen Bid Markup of Family Owned & Owners Related to Another Firm
Note: This figure shows marginal effects oil price risk on bitumen bid markup (Equation 14, results in Table 2.9 Column XII). The graphs fix volatility at its 10th, 40th, 60th, and 90th percentiles. They show the conditional marginal effect of a firm both being family owned and having its owners or officers related to another firm in the data as the time to start increases. 95% confidence intervals shown.

Appendix 2A Table 16 shows alternative specifications for the impact of public ownership with varying covariates. Excluding month-of-year and year fixed effects increases the estimated coefficient on the triple interaction to -3.8, though it remains significant. Excluding the contract level covariates, such as the number of bidders and distance to the project, and the firm size controls, reduces the coefficient slightly to -5.3 (that is, making the estimated risk premium slightly larger for public firms relative to private firms). Omitting all the individual and single interactions increases the coefficient substantially to -0.22. Table 17 addresses alternative volatility measures for the public and family ownership regressions; with implied volatility, the difference between public and private firms widens, with a coefficient on the triple interaction of -7.7, significant at the 1% level. Similarly, implied volatility increases the difference between family and non-family owned firms, giving a coefficient of 3.7, significant at the 5% level. Using the 5th month futures contract for historical volatility instead of the 6th month contract increases the effect of public ownership slightly to -4.6, but the impact of family ownership is the same as with the 6th month contract at 2.7 (both significant at the 5% level).

2.7 Conclusion

In this paper I find strong evidence that firms in an imperfectly competitive setting exhibit risk averse behavior, and that different types of firms manage risk in ways broadly consistent with key theoretic hypotheses. I establish using a quasi-experiment that oil price volatility increases bitumen bids in Iowa relative to Kansas since Kansas removed oil price risk from the private sector in 2006. Given the imperfect nature of competition in this industry, firms appear to be able to be able to pass through inefficiently high costs of risk to the government without being priced out of the market.

The generality of these results is, of course, limited by my focus on public procurement and asphalt paving. However, the firms in my sample much more adequately represent the size and ownership distribution of U.S. firms than the majority of past studies on risk management, which focus on large publicly listed corporations. The industry, in which government is the consumer, also lends itself to relatively clean identification of plausibly exogenous policy changes, whereas most settings face confounding demand variables.

The risk premium heterogeneity analysis finds that firms that are publicly listed, diversified, not family-owned, and larger tend to charge economically and statistically significantly lower risk premiums relative to their counterpart firms. There is a very strong negative relationship between risk aversion and diversification away from oil-intensive construction. Greater effective risk aversion among private firms and family owned firms supports the theory that concentrated ownership generates greater demand for risk management. It is also evidence against the hypotheses that managerial agency problems may be so dire that public firms manage risk as if private, or that family firms are less risk averse because their owner-managers are smoothing personal income. I find support, albeit weaker, that small firms are more financially constrained and therefore behave in a more risk averse manner. In sum, my results support the theoretical prediction that firms more likely to be financially constrained have a higher value of hedging.

3 Incentives to Invest in New Technology: The Effect of Fuel Economy Standards on China's Automakers

"We have been trying to exchange market access for technology, but we have barely gotten hold of any key technologies in the past 30 years"

- Liao Xionghui, Vice President of Lifan, a Chinese automaker (Ying 2012)

"New proposed fuel-economy standards for passenger cars...[leave] foreign makers well positioned to inject new technology...That leaves locals such as Great Wall and Geely with the most work to do."

- Bhattacharya (Wall Street Journal, 2014)

3.1 Introduction

Despite 35 years of industrial policy targeting a globally competitive, high quality indigenous auto sector, domestic Chinese automakers have not acquired the technology necessary to compete with foreign automakers. Foreign brands dominate China's passenger vehicle market - the world's largest since 2010 - by quality, price, and market share.[91] Chinese exports are minimal, and near zero to the developed world. Not only are quality and exports associated with higher profit margins in automotive manufacturing, but domestic firms' decision to produce low quality, low price vehicles runs counter to explicit government directives. Further, domestic firms' lagging technological capacity contrasts with other high-tech sectors, where Chinese firms, some majority state-owned, produce and export at globally competitive quality levels. It also contrasts with the rapid progress achieved by Japanese and Korean automakers in the 1970s and 1980s.[92]

Why did Chinese automakers choose to make small, cheap, low quality cars instead of investing in the power-train, safety, and design technology needed to compete in higher value segments? In this paper I propose two hypotheses: 1) the sudden imposition of fuel economy standards increased the barrier to entry in high quality segments; and 2) the joint venture mandate for foreign entry disincentivized technology upgrading. Both were enabled by high tariffs that precluded imports.

High quality vehicles have greater torque (acceleration) and horsepower. They have more

[91] For cross-country market size comparisons, see Wang et al. (2013).

[92] See Appendix B for a comparison of the auto sector with other sectors in China and with auto sectors in other countries.

accessories like air conditioning, and are usually bigger and heavier, and thus safer.[93] There is a basic tradeoff between these characteristics and fuel economy. An automaker faced with fuel economy standards can either reduce quality to meet the standard or invest in developing or acquiring fuel efficiency technologies. China's sudden and stringent fuel economy standards aimed to hasten technology transfer. Yet they imposed fixed costs of technology acquisition on domestic firms. Foreign firms, who already faced high standards in other markets, incurred only the variable cost of including their technologies in Chinese production. A fixed cost disadvantage may have pushed domestic firms to produce lighter, smaller, and cheaper cars to meet the standard without investing in new fuel efficiency technologies.

I use detailed, novel data on model-level sales and characteristics between 1999 and 2012.[94] I assess the fuel economy standards' effect on model characteristic decisions of domestic Chinese automakers. My primary approach compares foreign to domestic branded vehicles before and after the policy, in a differences-in-differences design. I find that domestic Chinese firms responded to the 2009 fuel economy standards by reducing their models' torque, horsepower, weight, and price.[95] I find weak evidence that they also reduced vehicle height and length. Foreign model characteristics continued on their pre-policy trajectory.

Specifically, I find that the standards reduced vehicle torque in domestic models relative to foreign models by 17 nm, or about 11.5% of mean torque among domestic firms, and horsepower by 6.3 kw, or 8% of the mean. The standards reduced domestic model price by $2,784 (13% of the mean), weight by 55 kg (4.3% of the mean), and length by 91 m (2.1% of the mean), all relative to foreign models. The policy's effect on all characteristics grows when I restrict the sample to models with larger sales volume. In my primary specification, I pool data on both sides of the policy and cluster standard errors by group (firm) to reduce bias from serial correlation of the variables. I also include firm fixed effects. I conduct a rich array of robustness tests, including placebo with other years, different bandwidths around the policy, different types of fixed effects, and alternative assumptions about standard errors.

---

[93]The relationship between weight and safety is well-documented. See, for example, Anderson and Auffhammer (2014) and Consumer Guide Automotive (2014).

[94]The sales data are from the State Council Development Research Center in Beijing, and are linked to model-year characteristics collected from the internet. I also have data from 2013, which will be incorporated in a future version.

[95]China imposed fuel economy standards in phases from 2005-2009, but the more stringent binding standards only came into force in 2009 (See Section 3.2.2).

A second empirical approach exploits the standards' staged implementation for new models in 2008, and continuing models in 2009. For example, the 2008 Great Wall Peri was a new model as it was not produced in 2007, while the 2008 Volkswagen Jetta was a continuing model. A triple differences design reveals that in 2008, domestic firms' continuing models (not yet subject to the policy) were more powerful, more expensive, larger, and heavier than new models. As with the differences-in-differences design, the primary triple differences specification examines within-firm variation, netting out the foreign-domestic firm and 2008-2007 differences.

My results establish that the fuel economy policy failed to achieve its original motivation of forcing increased technology transfer. From a social welfare perspective, although China's fleet became more fuel efficient, it also became more unsafe. An increasing share of vehicles are either very heavy or very light, making crashes more likely fatal, and poor quality in Chinese vehicles is accompanied by reduced safety (see Sections 3.2 and 3.6). From a private welfare perspective, Chinese firms may maximize profits by producing at the bottom end of the quality-price distribution. Yet the absence of Chinese exports despite explicit government export targets, evidence from the global market that exports are positively associated with profits, and the failure of Chinese firms to gain market share together suggest that thus far the down-market strategy has not been successful.[96]

However, China's automotive industry is changing rapidly. New organizational structures, including independent engineering and design firms that allow domestic automakers to outsource R&D, may allow Chinese firms in the future to undercut foreign competition for small, cheap cars in China and elsewhere (Shirouzu 2012). The results in this paper apply only to the industry through 2012.

I explore the mechanism driving Chinese firms' poor quality outcomes by comparing the performance of firms along two dimensions: whether the firm has a JV with a foreign firm, and whether it is privately owned or is a state-owned enterprise (SOE) at the central or local level. All foreign firms that manufacture vehicles in China do so through JVs, enterprises that produce foreign branded vehicles (such as the Mazda 6) but that pass about 50% of profits to a Chinese partner (in Mazda's case, FAW Auto), which produces domestic brand vehicles in separate plants. In theory, the domestic partner has greater access to the foreign firm's R&D and manufacturing capabilities than it would without a JV, and thus a lower cost of technology absorption. The popular press and some political science literature has argued, however, that the JV policy failed to spur technology transfer (e.g. Thun 2004).

[96]See Clerides, Lach and Tybout (1998), Melitz and Redding (2014), and State Council (2009).

Is there an "innovation cost" to FDI through JVs? This is an important policy question in many developing economies whose industrial policy has required FDI through JVs (e.g. Mathews 2002). China leveraged its bargaining power - access to the domestic market - to mandate JVs. On one hand, Chinese firms extracted large rents and the JVs created many jobs, some high skill. Yet dynamically the industry structure may have reduced domestic firm innovation incentives. In a stylized model, I show how domestic firms with JVs could be disincentivized from producing substitutes to their foreign partners' models to avoid cannibalizing their share of foreign brand profits. That is, the negative effect of increasing own quality on the share of JV profits might outweigh the JV's technology acquisition cost advantage.

I evaluate the effect of the fuel economy policy on subsets of firms, and show that SOEs with JVs were primarily responsible for the negative effects of the policy on domestic firm quality and price. Private firms without JVs responded least to the policy, and private firms also generally outperformed SOEs over the whole period. This is consistent with previous literature documenting greater productivity of private firms in China (e.g. Khandelwal et al. 2012, Lin et al. 1998). However, the negative effect of having a JV appears stronger than the negative effect of being state-owned. This suggests that requiring JVs in order to accelerate technology transfer may be misguided.

Technology diffusion is central to economic development (Lucas 1993, Young 1991, Nelson and Phelps 1966). In particular, increasing the quality of manufactures is often assumed necessary for export success and growth (Kremer 1993, Grossman and Helpman 1991a). Guided by the empirical fact that successful emerging markets in the post-WWII period developed innovation capacity by first obtaining foreign technology, the literature typically posits that growth depends on the rate of technology adoption (Parente and Prescott 1994, Grossman and Helpman 1991b). When and at what rate firms learn helps explain income disparities across countries, and is pivotal to the effectiveness of infant industry protection. However, the evidence about these policies is mixed; in particular, FDI's role in technology diffusion and growth is contested in both research and policy (Blalock and Gertler 2007, Hale and Long 2012).

Acquiring technology is costly, whether by own development, licensing, JVs, M&A, imitation, or theft. This is especially true in the modern automotive industry, where technology absorption involves considerable tacit knowledge in engineering, manufacturing, and other types of human capital (Ahrens 2013). I present evidence that a set of distortionary policies designed to protect (high tariffs), nurture (JVs), and prod (fuel economy standards) an infant industry backfired.

My analysis departs from much of the past literature by focusing on the technical quality of firm products, rather than accounting-based measures of productivity like labor cost. Though my findings are limited to a specific sector, country, and time period, the question of how government policy affects incentives to invest in technology upgrading is broadly applicable. This paper contributes to the literatures on industrial policy, technology transfer, the Chinese economy, and the impacts of energy efficiency regulation. I show that standards based on weight and vehicle type perversely incentivize automakers to produce more SUVs, which relates to the literature on the counterproductive effects of attribute-based regulations, such as the U.S. Energy Star program for household appliances (Aldy and Houde 2015).

In Section 3.2, I provide historical context about the Chinese auto sector and explain the fuel economy standards. I present the data and provide descriptive statistics in Section 3.3. I propose the estimation strategy in Section 3.4. Section 3.5 contains the main results and robustness tests. In Section 3.6, I analyze the role of JVs and SOEs in the auto sector, and assess their relative response to the fuel economy standards. 3.7 concludes.

3.2 Context: Industry Structure and Fuel Economy Standards

3.2.1 China's Auto Sector in Historical Context

Chinese policymakers considered light-duty passenger vehicles to be inessential luxury goods until the "Opening and Reform" of 1978. Indeed, before 1984 personal vehicle ownership was technically illegal (Anderson 2012). But in 1986, the central government designated the automotive sector a "Pillar Industry," and it has subsequently described automobile production as key to China's development.[97] Even the most recent automotive sector plan states that "Development of the automobile industry, including transformational upgrading, is an urgent task and is important for new economic growth and international competitive advantage" (State Council 2012).

From the early 1980's, the central government's auto policy focused on inducing technology transfer from foreign to domestic firms, primarily through encouraging foreign direct investment (FDI) (e.g. State Council 2006). Initially widely perceived as an avenue to knowledge spillovers, the role of FDI in technology diffusion is now contested in the empirical literature.[98] In practice,

[97]The 7th Five-Year Plan issued in 1986 instructed policymakers to consider the "automotive industry as an important pillar industry, and it should follow the principles of 'high starting point, mass production, and specialization' to establish backbone enterprises as leaders." See Chu (2011)

[98]Borensztein et al. (1998) find in a large sample of countries that FDI has larger positive effect on growth than domestic investment. Similarly, Xu (2008) finds that FDI positively impacts innovation (patenting) in

many countries subsidize FDI while others restrict it, and some switch between the two (UNCTAD 2012). China allowed foreign firms to manufacture light duty vehicles in China *only* in partnerships with domestic firms; the idea was to exchange market access for foreign technology. The JV is a stand-alone enterprise no more than 50% owned by the foreign automaker. Initially the domestic partner was hand-picked by the government, but in the past ten years JVs have merely requred government approval (Richet and Ruet 2008). During the period I study, JV enterprise plants produced essentially only foreign-brand vehicles, and the foreign partner was responsible for designing, controlling, and operating the plant. However, there is usually 50-50 profit share agreement (for more detailed discussion of how the JVs operate, see Section 6). Beijing explicitly intended the domestic partners to evolve into multinationals competing in foreign markets.

The JV policy was not systematically applied to other Chinese sectors, but other countries have taken similar approaches, including Malaysia, India, Russia, and a number of Latin American countries. Some research has found a positive effect of JVs on the innovative capabilities of local firms (e.g. Lyles and Salk 1996 and Mathews 2002). However, other work has found JVs to have negative effects on the partner firm, despite local managers and engineer learning (e.g. Inkpen and Crossan 1995, Doner 1991, Grieco 1984). In China, Gao (2004) finds negative impacts of JVs on firm innovation, and Jing and Zhou (2011) suggest that many JVs in a sector can lead to dependence on foreign technology.

China's protectionism likely exacerbated these incentive problems. Import tariffs of 180-220% through 1994, 70-150% through 2001, 30% through 2005, and 25% thereafter restricted the vast majority of Chinese consumers to vehicles produced in China. Appendix 3A Figure 1 shows that less than 0.5 million vehicles were imported until 2010, and since imports have risen - driven by SUVs - to a little over 1 million. Initially, the absence of competition enabled the few foreign firms manufacturing in China through JVs to use outdated technology, thus limiting the potential cost of any technology transfer (Moran 1998). In the early 2000s, 60% of domestic brand models were outdated foreign designs purchased or stolen from foreign automakers (Oliver et al. 2009). Subsequent policy required JVs to have "the capacity for manufacturing products which attain the international technological levels of the 1990s" as well as an R&D center (Walsh 1999).

China, and Haskel et al. (2007) find a positive effect of FDI on TFP in the UK. Blalock and Gertler (2007) find strong evidence that foreign investment generates Pareto improving technology transfer, increasing productivity, profits and output in the local market. Other work, such as Haddad and Harrison (1993), Konings (2001), and Aitken and Harrison (1999) find negative effects of FDI on productivity in Morocco, Eastern Europe, and Venezuela, respectively. See Hale and Long (2011 and 2012) for a review

During the 1990s, state-owned automakers were corporatized, largely separated from direct government control, and many were partially listed on stock exchanges (Andrews-Speed 2012). Deng and Ma (2010) estimate markups in Chinese auto industry between 1995 and 2001, and found that large automakers set high markups; Volkswagen, for example, had estimated markups of 42%, with a 41% market share in the late 1990s. Following WTO accession in 2001, the government gradually removed barriers to entry for both independent private firms and foreign firms establishing new JVs. In this period demand grew dramatically, and new foreign firm entry led to more competition and updated models (Oliver et al 2009). Although WTO terms forbid market access-technology transfer *quid pro quo*, the government continued to enforce the technology transfer requirements of its 1994 auto sector policy.

Beijing has called for "self-reliant Chinese car manufacturers who ranked among the 500 largest global firms" (NDRC 2004). More recent policies emphasize the auto industry as a key locus of economic upgrading, and focus on independent R&D ("indigenous innovation," 自主创新) and "new energy" vehicles (State Council 2006, 2012). Throughout the reform period, central industrial planners have sought to consolidate the auto industry, aiming to mimic the scale of the Big Three American companies[99]. Despite achieving a few large SOE mergers, these consolidation efforts were broadly unsuccessful. Privately owned firms entered the market and provincial governments established new local state-owned automakers, ignoring the central targets (Oliver et al 2009).

In 2009, an industry analyst concluded:

> "Two-and-a-half decades have passed and dozens of such joint ventures have been built in China. But no domestic automaker has achieved what the government wanted. While some own-brand cars are built on platforms transferred from global automakers, almost all of the rest are products of the reverse engineering of international models. Some domestic firms continue to resort to outright copying" (Yang 2009).

Similarly, a study of patents found that local Chinese automakers lagged far behind in conventional power-train technologies (Medhi 2006).

Growth literature typically posits that income disparities across countries depend on varying rates of technology adoption. For example, Parente and Prescott (1994) theorize that barriers to technology adoption - including regulatory constraints, corruption, or threat of violence - increase

[99]The Automotive Industry Policy of 1994 was quite specific, designating 8 companies that were permitted to manufacture passenger cars "The Big Three, Small Three and Mini Two" permitted to produce passenger cars were, in order: FAW, SAIC, Dongfeng, BAIC, TAIC, GAC, Changan and Ghizou Aviation. See State Council 1994.

the cost of adoption, accounting for much of the income disparity. But how this technology adoption occurs remains unclear. For Lucas (1993), the engine of growth is climbing the quality ladder through local industry exports. Melitz (2003) shows that new export opportunities and intense competition create aggregate productivity gains by reallocating resources from less to more productive firms. This is consistent with the finding in Clerides, Lach and Tybout (1998) that more productive firms select into exporting, but exporting itself does not increase a firm's technical efficiency. Limited local markets, competition, and export-oriented industrial policy apparently allowed firms in the "East Asian Tigers" to learn quickly in order to compete in foreign markets.

Unlike Japan and Korea, China's automotive industrial policy was not successful. Foreign brands dominate the Chinese market (see Section 3.3), and the little exports thus far are concentrated in privately owned firms without JVs (see Section 3.6). China's high import tariffs and JV requirements are forms of infant industry protection. In general, there is no consensus in the literature about the effectiveness of infant industry policies (Grossman and Helpman 1994, Nunn and Trefler 2010). Models of infant industry protection and the effects of trade on growth depend on how firms learn (e.g. Melitz 2005, Young 1991, and Clemout and Wan 1970). In this paper, I present evidence of when firms *do not* learn, hopefully shedding some light on this debate.

3.2.2 The Fuel Economy Standards

In 2004, China's National Development and Reform Commission announced that China would, for the first time, adopt fuel economy standards. The policy had two aims: 1) to decrease oil consumption for energy security purposes; and 2) to increase technology transfer by forcing foreign firms to bring more up-to-date technology to China (Wagner et al. 2009, UNEP 2010, Oliver et al. 2009).

There is a basic tradeoff between vehicle fuel economy and, primarily, weight, torque and horsepower. An automaker faced with fuel economy standards can build lighter, less powerful cars that will meet the standards without new technology. Alternatively, the automaker can maintain or improve quality by acquiring fuel efficiency technologies. These include discrete engine parts like catalytic converters and whole-vehicle design improvements in the power-train, aerodynamics and rolling resistance.[100] Importantly, high quality vehicles - particularly heavy and powerful ones -

[100]Other specific technologies include reducing transmission losses, direct fuel injection, variable valve timing, turbochargers, superchargers.

have higher profit margins than lower segments (IMF 2006).

Foreign automakers have faced stringent fuel economy standards in Japan and Europe for decades, and have developed technologies permitting heavy, powerful cars to meet those standards. Knittel (2011) examines the trade-offs in the U.S. auto industry between 1980 and 2006, establishing that decreasing weight in passenger cars by 10% is associated with a 4.2% increase in fuel economy, and decreasing horsepower by 10% is associated with a 2.6% increase in fuel economy. He documents that U.S. automakers improved fuel efficiency technology dramatically but used those improvements to increase engine power and weight but improve fuel economy only slightly.

Some of the technologies - particularly in the engine - are often outsourced to suppliers, but to integrate the technology and effectively model its trade-offs the branded automaker must invest in engineering and design competency, as well as the relationship with the supplier (Morris et al. 2004, Chanaron 2001). Industry analysis typically assumes that the locus of innovation is the branded automaker, especially for fuel efficiency technologies (Oliver Wyman 2013). Unfortunately, I do not have data on the fixed and variable costs of fuel efficiency technologies. However, the variable costs are not insignificant; Mckinsey (2012) estimates that new U.S. fuel economy standards will increase component costs in American vehicles by 20% between 2012 and 2020.

In general, fuel economy standards generate an incentive to down-weight certain classes of vehicles. Jacobsen (2013) and Anderson and Auffhammer (2014) show that down-weighting in response to fuel economy standards produced large negative welfare effects in the U.S., because when the fleet has widely varying weight, crashes are more likely fatal for passengers in small cars. While the standards in the U.S. and Europe are based on targets for an automaker's overall fleet, China and Japan use a weight-based step system that applies to each individual vehicle.[101] This generates the perverse incentive to meet standards by either increasing fuel economy within a class (potentially by decreasing weight) or jumping to a higher weight class with a more lenient standard. In Japan, weight-based standards are estimated to impose large safety costs (Sallee and Ito 2013). China is currently increasing the stringency of its standards, and is shifting to a fleet-based system. The policy agenda is now much more oriented towards using fuel economy and emissions standards to reduce urban pollution, rather than generate technology transfer (Shen and Takada 2014).

[101]Wagner et al. (2009) suggest that because China had so many small manufacturers producing only one or two models, a fleet average approach was not meaningful. Oliver et al (2009) point out that "vehicle sales figures in China have been historically secret, unknown, and/or difficult to obtain, making a sales-weighted average approach unpractical."

China's Phase 1 fuel economy standards were implemented in July 2005 for new models and January 2006 for continuing models. Phase 2 came into effect in January 2008 for new models, and January 2009 for continued models.[102] The Phase 2 standards are graphed in Figure 3.1, and Appendix 3A Table 1 lists the standards by weight class.[103] Phase 2 is more stringent than current U.S. standards, but much less stringent than Japanese and European standards (Appendix 3A Figure 2 compares standards across countries). The Chinese standards are designed to be "bottom heavy," meaning that they are stricter for heavier vehicle classes (An et al. 2011).

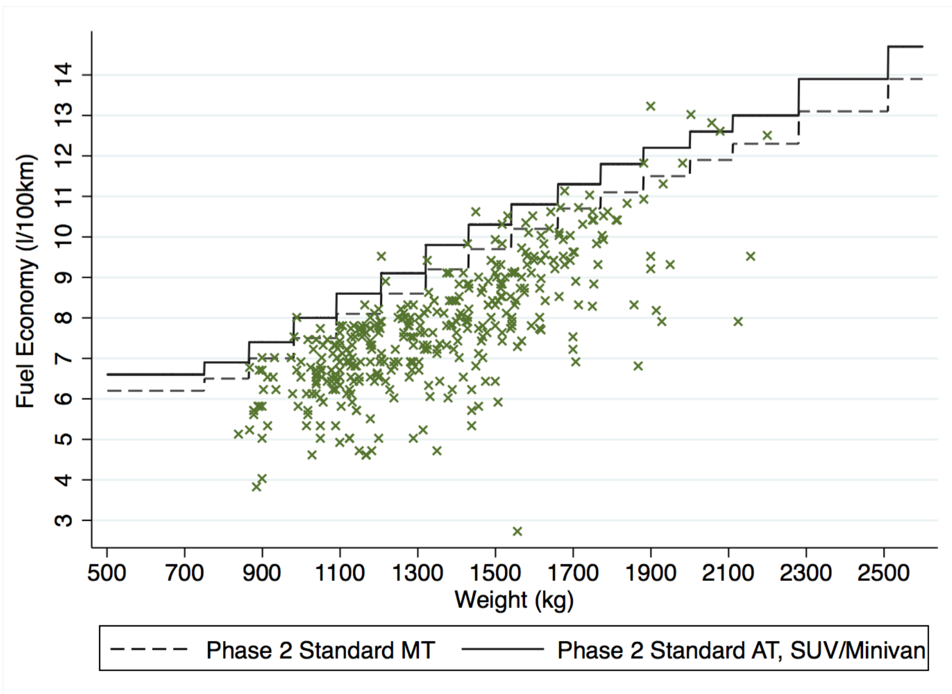

Figure 3.1: Model Fuel Economy and Weight, with Phase 2 Standards, 2010

Note: This figure shows China's weight-based Phase 2 fuel economy standards, which were imposed in 2008 and 2009. The dotted line indicates the standard for manual transmission vehicles, while the line indicates the standard for automatic transmission vehicles and all SUVs and minivans.

Before the standards, automakers selling vehicles in China did not have to report fuel economy. However, assessments of the standards have concluded that the initial 2006-07 standards were

[102]Phase III was phased in between 2012 and 2015. Phase III alters the previous program by adding corporate average fuel economy targets to the weight-based system. According to the 2012 Energy-Saving and New Energy Vehicle Industrialization Plan, the goal is to achieve a fleet average of 6.9 L/100km by 2015, and 5.0L/100km by 2020.

[103]The Phase 2 standards are roughly equivalent to Euro IV. China uses the New European Driving Cycle (NEDC) testing method, rather than the CAFE method used in the U.S.

not binding (Wagner et al. 2009, Oliver et al. 2009, An et al. 2007). But Wagner et al. (2009) estimate that 32% of models in 2007-07 would not meet the Phase 2 standards. More generally, prior to the Phase 2 implementation, government inspection and enforcement was lax, particularly for domestic automakers. If automakers made a model's fuel economy public, they often provided no indication about the driving cycle (city vs highway driving). It is thus difficult to compare fuel economy before and after the standards. My interviews in 2013 at the the government-affiliated China Automotive Research and Technology Center (CATARC) in Tianjin, which has been partially responsible for developing fuel economy standards and testing vehicles, confirmed that meaningful enforcement of the standards and consistent fuel economy testing began in 2008-2009. My primary estimation therefore takes 2009 as the policy implementation year. Figure 3.1 shows the fuel economy reported for new vehicles in 2010 alongside the Phase 2 standards. Assuming accurate reporting, it seems that the vast majority of models meet the standards.

3.3 Data and Descriptive Statistics

This paper is based on a unique, non-public dataset of all passenger vehicle sales in China between 1999 and 2012. Each observation is a model-year, and includes the ultimate Original Equipment Manufacturer (OEM), brand, model name, vehicle class, engine displacement, and power-train (all in Chinese).[104] The data is from the State Council Development Research Center (DRC), which is the policy analysis organization for China's top-level state (i.e. not Party) governing apparatus. The sales data is quite reliable, as it originates in police registration data that is provided to the DRC.[105] In this section I describe the data, present summary statistics, and demonstrate parallel trends for foreign and domestic firms.

I acquired model-year characteristics through web scraping. The model characteristics are: price (MSRP), maximum torque (nm), peak power (kw), curb weight (kg), length (mm), height (mm), and fuel economy (l/100 km).[106] I convert price into dollars using the average monthly

[104] OEM refers to the firms that design, assemble and brand vehicles such as Ford and Hyundai. Class is either city car, sedan, SUV, minivan, or van. Engine displacement is in liters, and is not used. Power-train is either internal combustion engine, natural gas, electric, or hybrid electric.

[105] I acquired this data in my capacity as a visiting scholar at the DRC (中国发展研究基金会), which was possible because of an invitation secured by Harvard Kennedy School Professor Anthony Saich from Lu Mai, the Secretary General of the DRC. The data itself was provided through the head researcher at DRC's Institute of Market Economy. I now have 2013 data, and will incorporate it in a future draft.

[106] The webscraping did not find characteristics for some model-years. There is coverage for 82% of models

exchange rate that year, and all price figures are nominal. As discussed in Section 3.2.2, fuel economy is rarely reported and unreliable in the pre-policy period.

Vehicle torque, responsible for acceleration and power, is a useful measure of vehicle quality.[107] Torque depends not only on the engine but also transmission ratios, weight, and many other aspects of overall vehicle integration. A car with more torque will have a better driving feel, and usually better engineering and design. In my data, the correlation between torque and price for all model-years is 0.83. When torque is multiplied by a given speed (usually in rpm), it gives horsepower (usually in kilowatts). Power is the amount of energy the engine can produce and determines the top speed of the vehicle. Its correlation with price is 0.84, and with torque 0.9. I treat torque, power and price as measures of vehicle quality, but also show the effects of the policy on vehicle weight, height and length. In general, larger, heavier cars have more amenities and are safer. The correlation between weight and price in my data is 0.67.

I use brands as the unit of analysis in descriptive statistics and primary estimations. Examples of brands are Ford, Audi, BYD, and Roewe. To avoid confusion, I term brands "firm," but the reader should be aware that in many cases the firms I refer to are in fact subsidiaries of an OEM. While Ford and BYD are both their respective OEM's only brand, Audi is a Volkswagen subsidiary, and Roewe is a brand of Shanghai-government owned SAIC. I use brands because they are the unit of observation most relevant to understanding quality; design, engineering and final assembly generally take place at the brand (firm) level, rather than the OEM level. This is especially true in China, where some OEMs are JVs producing domestic and foreign brand vehicles, albeit at different plants. I show that my empirical results are robust to grouping at the OEM level, but focus descriptive statistics at the firm level.

Figure 3.2 compares foreign and domestic market share, providing visual evidence for parallel trends between foreign and domestic firms in volume and market share in the years preceding the 2009 policy. The number of vehicles produced in China rose from 0.6 million in 1999 to nearly 16 million in 2012. Variety increased as well; the number of models rose roughly linearly from 23 in

---

(slightly more for foreign models (88%) than domestic (73%), and slightly better in later years). Models without characteristics have much lower sales; the mean sales volume is 13,629 for models lacking characteristics data compared with 25,824 for models with characteristics data.

[107] Torque is the amount of force the engine can apply in a rotational manner, measured in nanometers.

1999 to 426 in 2012.[108] In the 2004-2006 period, domestic firms gained market share. The strong relationship between torque and price, as well as the marked difference between foreign and domestic firms, are depicted in Figure 3.3. Figure 3.4 shows how foreign and domestic firm sales-weighted firm characteristics have diverged over time, specifically after 2009.[109] Domestic firms show slight decreases while foreign firms improve vehicle quality, additional evidence for parallel trends in sales-weighted characteristics prior to the policy. Below I discuss more rigorous regression tests. Figure 3.5 shows 2010 sales and price figures for the the largest foreign (top graph) and domestic (bottom graph) firms.[110] The same graphs at the OEM level are in Appendix 3A Figure 7.

---

[108]Here versions of the same model with different engine sizes are not treated as different models

[109]A firm j's sales-weighted torque (SWT) is calculated as follows, where i denotes model and t denotes year: $\text{SWT}_{i,t} = \sum_{i \in j} \left(\frac{s_{i,j,t}}{\sum_j s_{i,j,t}} \cdot \text{torque}_{i,j,t} \right)$. The figures show SWT averaged across firms within a firm type (foreign or domestic).

[110]The highest volume domestic firms are Wuling, with about 1.2 million vehicles sold in 2012, and then Chang'an, Chery, Great Wall, and BYD (in that order). Chang'an sold 0.8 million vehicles in 2012. However, Wuling is a JV between GM, SAIC and Liuzhou Wuling Motors, and Wuling is counted by GM as part of its China production. Volkswagen is the dominant foreign firm, at about 2.1 million vehicles in 2012. Hyundai follows, at around 0.8, and the next largest are Nissan, Toyota, Buick, Chevrolet, and Honda, in that order.

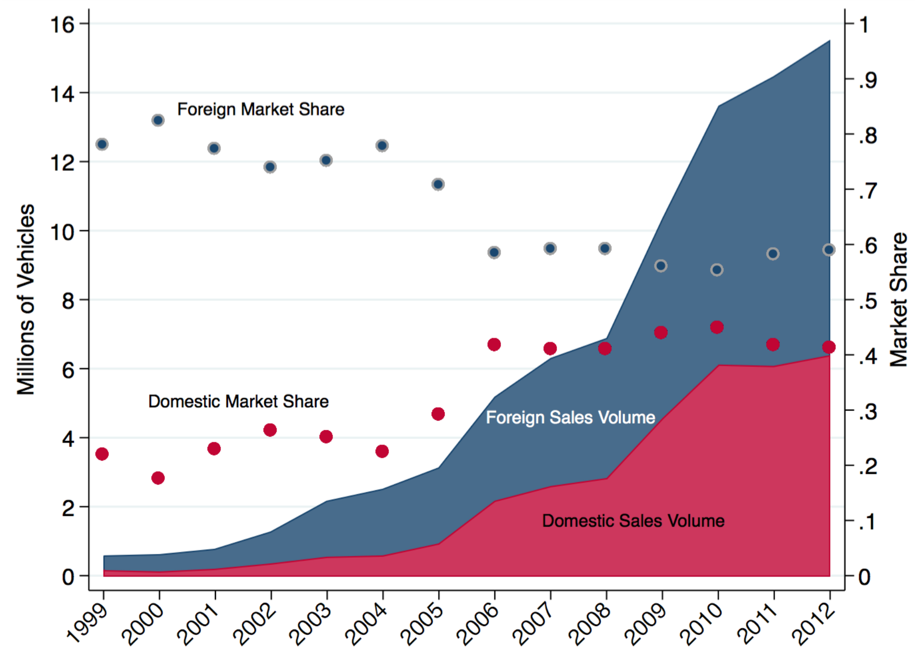

Figure 3.2: Sales Volume and Market Share by Firm Type

Note: This figure shows foreign and domestic brand Chinese sales volume (number of new vehicles sold in a given year) on the left axis, where the blue area is foreign and the red area is domestic. Market share of sales volume is on the right axis and in the foreign (blue) and domestic (red) scatterplot.

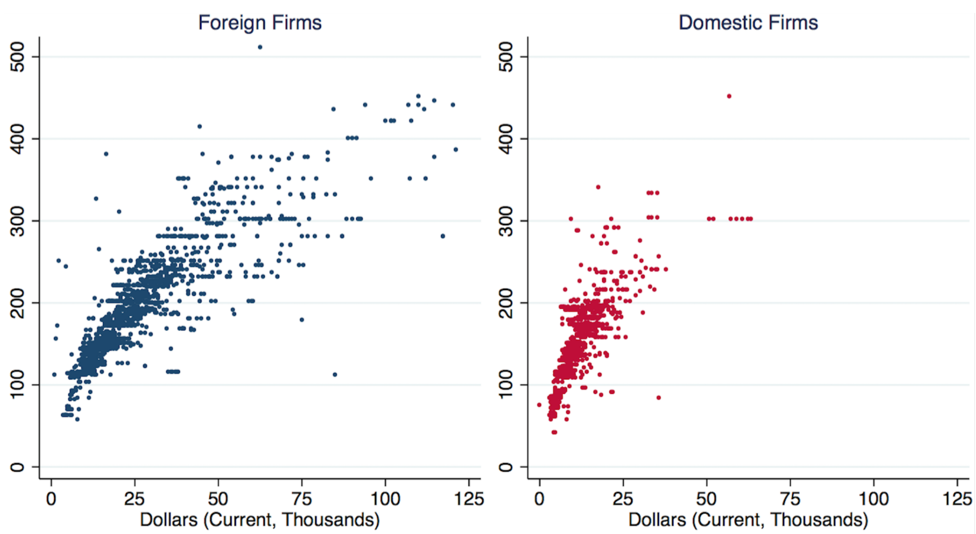

Figure 3.3: Model Torque (nm) and Price by Firm Type

Note: This figure shows model torque (y-axis) and price (x-axis) for foreign firms and domestic firms. Each observation is a model-year, and all models sold between 1999 and 2012 are included.

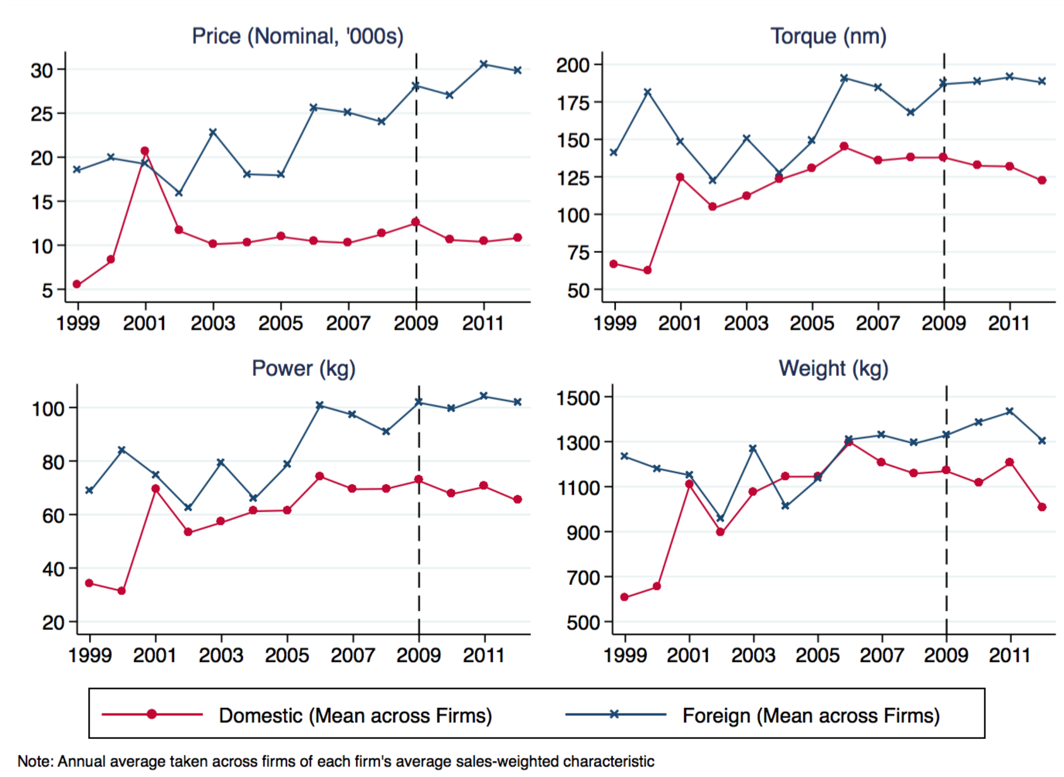

Figure 3.4: Sales-Weighted Characteristics by Firm Type

Note: This figure shows foreign (blue) and domestic (red) sales-weighted characteristics. The annual means are calculated by averaging across firms each firm's average sales weighted characteristic.

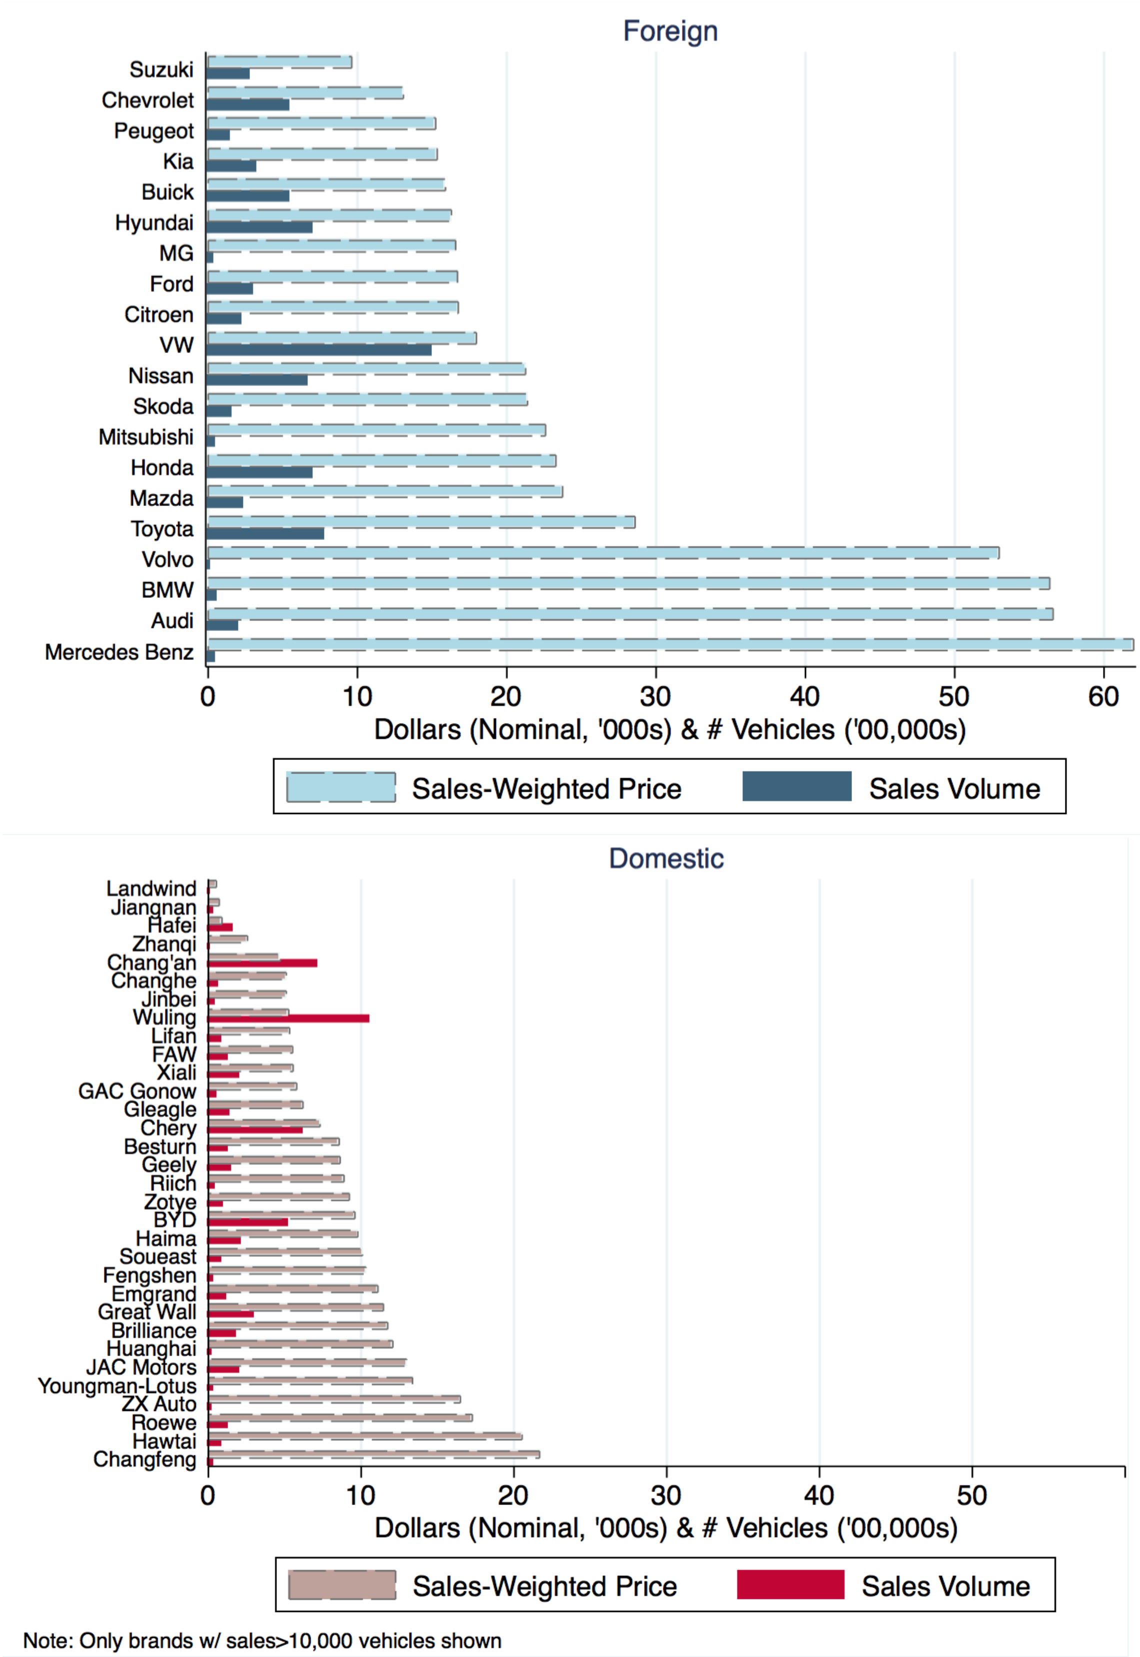

Figure 3.5: Firm-Specific Sales Volume and Sales-Weighted Price, 2010

Note: This figure shows firm sales volume (number of vehicles) and sales-weighted average price across models sold for foreign firms (top graph) and domestic firms (bottom graph). Only data from 2010 is included.

Summary statistics of the firms in my data are in Table 3.1, and of model characteristics in Table 3.2. In both tables, the first three columns divide the sample into three periods, the first two prior to the fuel economy standards (1999-2004 and 2005-2008) and the latter after the policy (2009-2012). This reveals average changes in the data around the policy and also offers a sense of how the market has changed over time. Column IV contains all years. Table 3.1 shows that the number of domestic firms nearly doubled over the course of the data, with most entry occurring prior to the 2009 policy. Average domestic firm sales volume doubled between each period, increasing from 52,000 vehicles per year in the 2005-08 period to 116,000 vehicles in the 2009-12 period. Similarly, average foreign firm sales volume increased from 146,000 per year to 320,000 per year. Amid this massive growth, domestic prices in nominal dollars have stayed essentially constant - thus decreasing in real terms - while foreign prices have increased significantly. For the full span of the data, the average foreign firm sales-weighted mean price is $24,200, while for domestic firms it is $10,800.

Table 3.2, where the model-year is the unit of observation, and characteristics are not sales-weighted, gives an average foreign model price of $26,500, compared with $12,200 for domestic models. Domestic Chinese model torque, power and weight increased between the first and second periods, but decreased or remained stable between the second and third periods (which bracket the fuel economy standards implementation). This contrasts with the foreign firms, for whom the means of all six characteristics increase between each period. However, none of the differences in means are significant, as characteristics within groups vary widely.

I estimate the relationship of characteristics to vehicle price using Equation 15, where j denotes firm and i denotes model:

$$Y_{it} = \alpha + \beta_1 \text{Domestic}_j + \beta_2 \text{Torque}_i + \beta_3 \text{Weight}_i + \beta_2 \text{Height}_i + \beta_2 \text{Length}_i + \delta' \mathbf{1} \mid \text{Class}_i + \varepsilon_{ijt}. \quad (15)$$

The results, shown in Table 3.3, indicate that there is a large and robust premium associated with torque; a one standard deviation increase in torque increases price by $11,500, or 50% of the average price. This relationship is consistent across the three periods (columns II-IV). The domestic firm discount increased over time; domestic firms were associated with a $3,300 discount between 1999 and 2004, and a $5,700 discount between 2009 and 2012. There is a large discount for SUVs relative to compact cars (the omitted class dummy), but there is no measurable relationship between the other classes and price. This may be because domestic firms have disproportionately increased their SUV sales relative to foreign firms, which is shown in Appendix 3A Figure 3.

Table 3.1: Sales Volume and Firm Sales-Weighted Price ('000s) by Firm Type

	I. 1999-2004		II. 2005-2008		III. 2009-2012		IV. All Years	
	Mean (sd)	N	Mean (sd)	N	Mean (sd)	N	Mean (sd)	N
			A. All Firms					
# **Active Firms**	-	45	-	78	-	86	-	94
Sales Volume	22.7 (29.0)	81	52.1 (89.5)	162	116 (192)	201	75.3 (148)	444
Sales-Wgtd Price	10.6 (8.2)	56	10.7 (7.0)	124	11.2 (8.1)	167	10.8 (8.0)	347
			B. Domestic (Chinese)					
# **Active Firms**	-	27	-	52	-	61	-	68
Sales Volume	22.7 (29.0)	81	52.1 (89.5)	162	116 (192)	201	75.3 (148)	444
Sales-Wgtd Price	10.6 (8.2)	56	10.7 (7.0)	124	11.2 (8.1)	167	10.8 (8.0)	347
OEM has JV	-	15	-	27	-	35	-	40
Sales Volume	24.4 (32.9)	42	53.8 (99.0)	87	118 (221)	115	79 (168)	244
Sales-Wgtd Price	9.9 (7.1)	29	10.2 (7.2)	64	11.3 (9.9)	100	10.7 (8.6)	193
Privately Owned	-	11	-	22	-	23	-	27
Sales Volume	13.5 (18.9)	37	32.7 (39.6)	67	97.1 (124)	77	56.2 (91.9)	181
Sales-Wgtd Price	10.7 (7.1)	22	12.0 (6.8)	52	10.8 (6.3)	63	11.2 (7.1)	137
Central SOE	-	10	-	17	-	21	-	38
Sales Volume	34.5 (36.2)	28	60.4 (75.7)	54	102 (149)	71	74.8 (115)	153
Sales-Wgtd Price	9.3 (7.0)	23	10.1 (7.8)	41	11.7 (12.0)	58	10.7 (9.9)	122
Local SOE	-	6	-	13	-	17	-	29
Sales Volume	23.3 (28.9)	16	72.9 (144)	41	160 (305)	53	108 (234)	110
Sales-Wgtd Price	12.9 (6.7)	11	9.2 (5.8)	31	11.0 ($5.4)	46	10.6 (5.7)	88
			C. Foreign (Non-Chinese; 100% have JVs)					
# **Firms**	-	18	-	26	-	25	-	26
Sales Volume	82.7 (123)	73	146 (173)	89	320 (383)	96	193 (281)	258
Sales-Wgtd Price	19.2 (13.2)	71	23.4 (15.4)	85	29.0 (17.3)	91	24.2 (160)	247

Note: This table shows means of firm sales volume (number of vehicles) and sales-weighted price ('000s of nominal US dollars at contemporary exchange rates). Sales volume is the average across firms of each firm's average annual vehicle sales over the specified time period, where each observation is a firm-year. The sales-weighted price is the mean annual sales weighted price of a firm's models, which is then averaged across firm-years. Prices are in nominal US dollars, at the average annual contemporary exchange rate. In columns I-III, the mean is taken across firm-years for all firms active in the period specified. JV= joint venture between foreign and domestic firm. SOE=state owned enterprise. I define firm at the brand level; a parallel table at the OEM level can be found in Appendix 3A.

The estimation strategy in Section 3.4 will compare domestic firms' response to the fuel economy policy with that of foreign firms. The results are the difference in the two types of firms' reaction to the policy. If foreign and domestic firms' model characteristics were on similar growth paths, the effects that I observe are more readily interpretable as reactions to the increase in fixed costs that domestic firms experienced but foreign firms did not. That is, the higher fixed cost to build high quality vehicles is like a "treatment."

Table 3.2: Model Characteristics by Firm Type

	I. 1999-2004		II. 2005-2008		III. 2009-2012		IV. All Years	
	Mean (sd)	N	Mean (sd)	N	Mean (sd)	N	Mean (sd)	N
A. All Firms								
Max Torque (nm)	160 (59.3)	280	173 (65.1)	916	177 (63.2)	1646	174 (63.6)	2842
Max Power (kw)	82.0 (31.8)	282	90.4 (34.8)	922	96.1 (33.9)	1651	92.9 (24.2)	2855
Weight (kg, '000s)	1.27 (0.32)	276	1.34 (0.32)	899	1.37 (0.30)	1587	1.35 (0.31)	2762
Height (m)	1.51 (0.14)	285	1.55 (0.15)	916	1.55 (0.16)	1640	1.54 (0.16)	2841
Length (m)	4.37 (0.47)	285	4.41 (0.44)	916	4.45 (0.40)	1641	4.43 (0.42)	2842
Price ('000s)	20.0 (13.5)	300	19.9 (16.2)	931	22.0 (17.1)	1654	21.1 (16.5)	2885
B. Domestic (Chinese) Firms								
Max Torque (nm)	129 (50.0)	78	151 (57.0)	350	147 (46.0)	653	147 (50.3)	1081
Max Power (kw)	65.5 (26.5)	80	76.2 (27.8)	354	79.0 (22.4)	658	77.1 (24.8)	1092
Weight (kg, '000s)	1.16 (0.35)	70	1.30 (0.35)	335	1.29 (0.28)	617	1.29 (0.31)	1022
Height (m)	1.54 (0.19)	77	1.61 (0.19)	344	1.59 (0.20)	643	1.59 (0.20)	1064
Length (m)	4.19 (0.58)	77	4.33 (0.50)	344	4.35 (0.43)	644	4.33 (0.47)	1065
Price ('000s)	12.2 (7.64)	87	12.1 (8.45)	354	12.3 (6.70)	651	12.2 (7.38)	1092
C. Foreign (Non-Chinese) Firms								
Max Torque (nm)	172 (58.4)	202	186 (66.3)	566	197 (64.9)	993	191 (65)	1761
Max Power (kw)	88.5 (31.5)	202	99.3 (35.8)	568	107 (35.4)	993	103 (36)	1763
Weight (kg, '000s)	1.30 (0.30)	206	1.37 (0.29)	564	1.41 (0.31)	970	1.38 (0.31)	1740
Height (m)	1.50 (0.12)	208	1.51 (0.11)	571	1.52 (0.12)	997	1.52 (0.12)	1777
Length (m)	4.44 (0.40)	208	4.46 (0.40)	572	4.52 (0.36)	997	4.50 (0.38)	1777
Price ('000s)	23.1 (14.1)	213	24.7 (17.9)	577	28.3 (18.8)	1003	26.5 (18.1)	1793

Note: This table shows means of firm model characteristics. The reported mean is the average across firms of each firm's average annual characteristic over the specified time period, where each observation is a firm-year. Prices are in nominal US dollars, at the average annual contemporary exchange rate. The unit of observation is the model-year. In the regressions, height is in millimeters. JV= joint venture between foreign and domestic firm. SOE=state owned enterprise. I define firm at the brand level. A parallel table at the OEM level, as well as a table where statistics are broken down by domestic firm ownership, can be found in Appendix 3A.

Table 3.3: Determinants of Vehicle Price by Time Period

Dependent Variable: Price (current dollars)				
Time Period:	I. All years	II. 1999-2004	III. 2005-2008	IV. 2009-2012
Domestic_j	**-5103*** **	**-3331*** **	**-4176*** **	**-5657*** **
	(476)	(1207)	(670)	(623)
Torque_i (nm)	**182*** **	**137** **	**201*** **	**178*** **
	(12)	(54)	(23)	(16)
Weight_i (kg)	**16*** **	**14**	**11*** **	**20*** **
	(3.1)	(12)	(3.4)	(4.9)
Height_i (mm)	**-10*** **	**-1.5**	**-5.9* **	**-9.8*** **
	(2.8)	(9.8)	(3.5)	(3.7)
Length_i (mm)	**-7.6*** **	**-1.1**	**-8.1*** **	**-7.6*** **
	(1.3)	(3.5)	(1.8)	(1.7)
$\mathbf{1}\,\|\text{Minivan}_i$	**-370**	**-2411**	**572**	**-2220**
	(761)	(2074)	(946)	(1465)
$\mathbf{1}\,\|\text{SUV}_i$	**-2802*** **	**-7532*** **	**-3258* **	**-3265** **
	(950)	(2616)	(1818)	(1297)
$\mathbf{1}\,\|\text{Sedan}_i$	**720**	**1866**	**2709*** **	**-1334** **
	(523)	(1172)	(944)	(649)
N	2720	267	883	1570
R^2	.74	.69	.73	.76

Note: This table reports regression estimates of the relationship between price and vehicle characteristics (Equation 15). The Domestic variable is 1 if the brand is domestic (Chinese), and 0 if foreign. There are fixed effects for 4 vehicle classes: Compact, Minivan, SUV and Sedan (Compact is omitted). The unit of observation is the model-year. There are no brand or year fixed effects. Standard errors are robust and clustered by brand-year. $1999 \leq \text{Year} \leq 2012$; *** indicates $p < .01$.

In Table 3.4, I present regressions that test for statistically different trends over time in model characteristics prior to the policy. The regressions, in which i indexes models, j indexes firms, and t indexes years, are of the form

$$Y_{it} = \alpha + \beta\left(\text{Year}_t \cdot \text{Domestic}_j\right) + \gamma_1 \text{Year}_t + \gamma_2 \text{Domestic}_j + \varepsilon_{ijt}, \tag{16}$$

where Year_t is a continuous variable ranging, in Panel A, from 2003 to 2008, and Domestic_j is an indicator for the firm being domestic (Chinese) rather than foreign. Y_{it} is a model-year characteristic, such as horsepower or price. Table 3.4 shows that there is no statistically significant difference in trends between foreign and domestic firms prior to the policy, except for length (column VI), which has a difference of 31 mm (relative to a sample average of 4,430 mm), significant at the 10%

level. The large standard errors mean that I cannot rule out a difference in trends. However, the coefficients are an order of magnitude smaller than the treatment effects I demonstrate in Section 3.5. For example, the treatment effect on torque in my primary specification is 17 nm, compared to an estimated difference in growth path of -1.3 to 1.6 nm shown in Table 3.4.

Table 3.4: Parallel trends among foreign and domestic firms prior to the policy

A. All Characteristics, $2003 \leq \textbf{Year}_t \leq 2008$

Dependent Variable:	I. Price (nom. $)	II. Torque (nm)	III. Power (kw)	IV. Weight (kg)	V. Height (mm)	VI. Length (mm)
$\text{Year}_t \cdot \text{Domestic}_j$	**-332**	**1.6**	**.56**	**22**	**4.4**	**31***
	(594)	(2.6)	(1.4)	(14)	(6.1)	(18)
Year_t	541	2.1	2.2***	7.6	1.4	13
	(347)	(1.5)	(.79)	(7.6)	(3.5)	(11)
Domestic_j	653111	-3160	-1143	-44947	-8678	-62417*
	(1190890)	(5199)	(2714)	(27392)	(12279)	(36948)
N	1113	1086	1092	1067	1090	1090
R^2	.15	.077	.12	.021	.081	.036

B. Alternative Specifications using Torque as Dependent Variable

Test:	VII. 2005-09	VIII. Firm f.e.	IX. Cluster s.e. by firm	X. Cluster s.e. by firm-yr	XI. Firm f.e., cluster s.e by firm-yr
$\text{Year}_t \cdot \text{Domestic}_j$	**-3.1**	**-1.3**	**1.6**	**1.6**	**-1.3**
	(4.2)	(2.1)	(2.4)	(3.4)	(2.3)
Year_t	.26	2.4**	2.1	2.1	2.4
	(2.5)	(1.1)	(1.6)	(2.5)	(1.7)
Domestic_j	6242	2630	-3160	-3160	2630
	(8469)	(4142)	(4723)	(6890)	(4621)
Firm f.e.	N	Y	N	N	Y
N	916	1086	1086	1086	1086
R^2	.067	.53	.077	.077	.53

Note: This table reports regression estimates testing whether the model characteristics of foreign and domestic firms were on different growth paths prior to the 2009 fuel economy policy (Equation 16). The Domestic indicator variable is 1 if the brand is domestic (Chinese), and 0 if foreign. The variable Year_t is continuous. The unit of observation is the model-year. Standard errors are OLS unless otherwise specified. *** indicates $p < .01$.

This paper does not address auto parts suppliers. In recent years, automakers sometimes purchase as much as 70% of the vehicle value added from parts suppliers (Canis and Morrison 2013). However, key vehicle design and technological challenges, particularly from a fuel economy perspective, are accomplished at the automaker level. A passenger car includes at least 15,000 parts, which must fit perfectly and function consistently in order to meet Western consumer expectations. Although component suppliers are an important part of the overall automotive industry, they are a separate sector from branded automakers and are beyond the scope of this paper.

3.4 Empirical Strategy

My analysis of the impact of the fuel economy standards relies primarily on a differences-in-differences (DD) design. I compare foreign and domestic firms' vehicle characteristics before and after the 2009 fuel economy policy. I also exploit the staged policy for new and continuing models in a triple-difference specification.

The standard DD design involves two groups, one of which is subject to a treatment in the second of two time periods. If the two groups are ex-ante similar and have similar time trends, then inclusion of controls for treatment and state should yield an estimated coefficient on the treated state that is the average difference between the treatment group and the control group. However, in practice DD estimators pose two potential problems. First, DD design will fail if the policy is endogenous to the group studied. The fuel economy standard affected both foreign and domestic firms, and I have been unable to identify other policies or market structure changes that would have affected domestic firms within the bandwidth of time in which I find an effect. Also, one of the policy's stated goals was to increase domestic firm technology quality. Therefore, endogeneity should work in the opposite direction than my results point.

The second issue is that serial correlation in variables may cause downward bias in the standard errors. This is especially problematic with relatively long time series and DD implementation via time fixed effects. As in most DD designs, the dependent variables here (e.g. model torque) are serially correlated. Pooling the data on either side of the treatment and clustering standard errors by group rather than time solves the problem, particularly when the number of groups is large (see Bertrand, Duflo and Mullainathan 2004, and Donald and Lang 2007). In my primary specification, I pool the data on either side of the cutoff with a bandwidth of three years around the policy, and cluster standard errors in 78 groups. In robustness tests I show that, among other things, using all years and including year fixed effects (as in most DD implementations) yields roughly the same results. I also demonstrate that conventional firm-year clusters do cause downward bias in standard errors.

My primary DD specification, where i is the vehicle model (e.g. the BYD F6 or the Chevrolet

Spark), j the firm (e.g. Chery or Honda), and t the year, is as follows:

$$Y_{it} = \alpha + \beta \left(\text{Policy}_t \cdot \text{Domestic}_j\right) + \gamma_1 \text{Policy}_t + \gamma_2 \text{Domestic}_j + \lambda_j + \varepsilon_{ijt}. \tag{17}$$

The outcome of interest is Y_{it}, such as model torque or price. The indicator Policy_t is 1 if the year is 2009 or later, and 0 otherwise. The indicator Domestic_j is 1 if the firm is Chinese (e.g. BYD or Chery), and 0 if it is foreign (e.g. Nissan or GM). I include firm fixed effects λ_j, which should control for unobserved firm-specific variables related to characteristic choice. The primary specification requires the model sales volume to be at least 1,000 vehicles.

The coefficient of interest, β, gives $\left(\bar{Y}_{i=\text{Domestic},1} - \bar{Y}_{i=\text{Domestic},0}\right) - \left(\bar{Y}_{i=\text{Foreign},1} - \bar{Y}_{i=\text{Foreign},0}\right)$, or the effect of the policy on domestic firms relative to foreign firms. β indicates how domestic firms responded differently to the standards than foreign firms. The parallel trend assumption - that the error term is uncorrelated with the other variables - is not directly testable, but evidence in Section 3.3 on model characteristics supports it. Although the Chinese auto industry grew and changed dramatically between 2006 and 2012, the specification is valid if shocks affected both foreign and domestic firms. Placebo tests show that similar treatment effects do not appear until 2009.

The second specification is a triple-difference design, which is more robust than any DD approach (Imbens and Wooldridge 2007). I exploit the staged policy implementation; only new models were required to meet the standard in 2008, and then in 2009 the standard applied to both new and continuing models. Automakers sensitive to the policy may have changed new model but not continuing model characteristics in 2008. The primary specification is:

$$\begin{aligned} Y_{it} \quad = \quad & \alpha + \beta \left(\text{Policy}_t^{2008} \cdot \text{Domestic}_j \cdot \text{Continuing}_{it}\right) + \gamma_1 \left(\text{Policy}_t^{2008} \cdot \text{Domestic}_j\right) \\ & + \gamma_2 \left(\text{Policy}_t^{2008} \cdot \text{Continuing}_{it}\right) + \gamma_3 \left(\text{Continuing}_{it} \cdot \text{Domestic}_j\right) \\ & + \gamma_4 \text{Policy}_t + \gamma_5 \text{Domestic}_j + \gamma_6 \text{Continuing}_{it} + \lambda_j + \varepsilon_{ijt}. \end{aligned} \tag{18}$$

The Policy_t^{2008} variable is 1 if the year is 2008, and 0 if 2007 or 2006. In the primary specification, I use two years before the policy in order to have a larger sample size. Here, the interpretation of the coefficient of interest β is slightly more complicated; it is the effect of being a continuing model relative to a new model, netting out the change in means in firm type (domestic vs. foreign) and in time period (after vs. before the 2008 policy).

3.5 Effect of Fuel Economy Standards on Vehicle Characteristics

3.5.1 Differences-in-Differences Results

The differences-in-differences estimation finds that domestic firms responded to the 2009 fuel economy standard by manufacturing less powerful, cheaper, smaller, and lighter vehicles. Throughout my specifications and robustness tests, the effects on torque and price are the strongest and most significant, with the effects on weight and size smaller and less robust.

The results of my primary specification (Equation 17) on all six vehicle characteristics are in Panel A of Table 3.5. I find that the standards reduced vehicle torque in domestic models relative to foreign models by 17 nm, or about 12% of mean torque among domestic firms (significant at the 1% level). The partial effect of the policy on torque is 11 nm for foreign firms, and -6 nm for domestic firms. Note that because I use firm fixed effects, the partial effect of being a domestic firm on torque requires omitting these effects. This is shown in Table 3.9 column III; the partial effect of being domestic is -32 nm of torque before the policy and -52 nm after.

Domestic automakers reduced price by \$2,784 relative to foreign automakers, which is 23% of average domestic firm price and 13% of average price across all models (significant at the 1% level). The standards reduced power by 6.3 kw, or 8% of average power among domestic firms (significant at the 5% level). They reduced weight is reduced by 55 kg, and length by 91 mm, which are 4.3% and 2.1%, respectively, of the domestic firm averages (both significant at the 10% level). Panel B of Table 3.5 uses a bandwidth of two years around the policy, and finds quite similar results, albeit slightly smaller. As the effects on height and length are not robust, and are also not strong measures of quality, I omit them in subsequent tables.

The effect of the policy on all characteristics grows as the sample is restricted to models with increasing required sales volume. In Table 3.6, I show the increasing effect on price; with no sales volume requirement (column I) the effect is -\$1,616, significant only at the 10% level, but at a required sales volume of 5,000 vehicles (column IV), the effect is -\$3,453, significant at the 1% level. Appendix 3A Table 2 shows a similar pattern for all characteristics. For example, when sales volume is required to be more than 5,000 vehicles, the effect on weight is -92 kg, almost twice the magnitude of the coefficient in the primary specification.

Table 3.5: Differences-in-Differences Estimation of the Fuel Economy Standard's Impact on Domestic vs. Foreign Model Characteristics

A. Bandwidth of 3 years around 2009 policy (primary specification)

Dependent Variable:	I. Torque (nm)	II. Power (kw)	III. Price (nom. $)	IV. Weight (kg)	V. Height (mm)	VI. Length (mm)
$\text{Policy}_t \cdot \text{Domestic}_j$	**-17*** **	**-6.3** **	**-2784*** **	**-55***	**-18**	**-91***
	(5.3)	(2.8)	(763)	(32)	(21)	(52)
Policy_t	11***	5.9***	2821***	29**	14***	30
	(3.5)	(1.9)	(627)	(14)	(3.8)	(23)
Domestic_j	59***	70***	4479***	248***	-39***	437***
	(2.7)	(1.5)	(488)	(11)	(2.9)	(18)
Firm f.e.	Y	Y	Y	Y	Y	Y
N	1646	1651	1653	1599	1630	1631
R^2	.5	.48	.63	.47	.39	.44

B. Bandwidth of 2 years around 2009 policy

Dependent Variable:	VII. Torque (nm)	VIII. Power (kw)	IX. Price (nom. $)	X. Weight (kg)	XI. Height (mm)	XII. Length (mm)
$\text{Policy}_t \cdot \text{Domestic}_j$	**-16*** **	**-5.4** **	**-2121** **	**-47**	**-9.9**	**-74**
	(4.6)	(2.7)	(801)	(29)	(21)	(50)
Policy_t	7.8***	3.7*	1708**	15	12***	7.5
	(2.9)	(2)	(688)	(15)	(4.2)	(25)
Domestic_j	53***	67***	2979***	221***	-35***	377***
	(2.2)	(1.5)	(516)	(11)	(3.2)	(19)
Firm f.e.	Y	Y	Y	Y	Y	Y
N	1088	1088	1079	1047	1069	1070
R^2	.49	.49	.63	.46	.34	.44

Note: This table reports regression estimates of the effect of the 2009 fuel economy standards on model characteristics, with a bandwidth of three years around 2009 policy (Equation 17). Sales volume is the number of units sold of that model-year vehicle. Domestic_j is 1 if the brand is domestic (Chinese), and 0 if foreign. Policy_t is 1 if the year is 2009 or later, and 0 if 2008 or before.. The unit of observation is the model-year. Only models with at least 1,000 units sold are included. Standard errors are robust and clustered by firm. *** indicates $p < .01$.

Table 3.7 contains the triple-difference estimation using the 2008 implementation of the policy for new models.[111] The coefficient of interest gives the effect of being a continuing vs. a new

[111] In 2009 the standards applied to both new and continuing models, so it is impossible to do a similar exercise with the 2009 rule.

model after vs. before the policy for domestic firms vs. foreign firms.[112] This coefficient is positive and significant for all four characteristics, showing that continuing models not subject to the policy were more powerful, more expensive, and heavier than new models for domestic firms relative to foreign firms. Note that the coefficients on the individual indicators and interactions are not direct effects.[113] The predictive margins of the policy's effect are graphed in Figure 3.6, where each line holds fixed whether the firm is domestic or foreign, and whether the model is continuing or new. Domestic firm new model torque decreased, while continuing models increased slightly. Foreign firm new model torque increased, while continuing model torque stayed roughly constant. Again, the firm fixed effects mean the that the level of torque is not meaningful, only the relative changes.

Table 3.6: Differences-in-Differences Estimation of the Fuel Economy Standard's Impact on New Vehicle Price by Model Sales Volume

Dependent Variable: Price (nominal dollars)				
Min. Model Sales Volume:	I. 0	II. 500	III. 1,000	IV. 5,000
$\text{Policy}_t \cdot \text{Domestic}_j$	**-1616***	**-2459*****	**-2784*****	**-3453*****
	(902)	(675)	(763)	(1232)
Policy_t	2560***	2730***	2821***	3589***
	(654)	(485)	(627)	(1075)
Domestic_j	7740***	4258***	4479***	-11010***
	(509)	(397)	(488)	(2255)
Firm f.e.	Y	Y	Y	N
N	2078	1775	1653	1177
R^2	.64	.64	.63	.21

Note: This table reports regression estimates of the effect of the 2009 fuel economy standards on model price, with a bandwidth of three years around 2009 policy (Equation 17). Sales volume is the number of units sold of that model-year vehicle. Domestic_j is 1 if the brand is domestic (Chinese), and 0 if foreign. Policy_t is 1 if the year is 2009 or later, and 0 if 2008 or before. Regressions IV and V omit brand dummies because they generate collinearity with the variables of interest. The unit of observation is the model-year. Standard errors are robust and clustered by firm. *** indicates $p < .01$.

[112] The proportion of new models was slightly higher than average in the policy implementation year. The average number of new models among all firms per year between 2006 and 2012 is 13%, and 15% in 2008. For domestic firms, the average is 26%, and is also 31% in 2008.

[113] For example, the -17 nm effect of $\text{Policy}_t^{2008} \cdot \text{Domestic}_j$ on torque is the interaction of the policy and being domestic within new models (when Continuing_{it} is zero). The coefficient of 39 on Domestic_j is the effect of being domestic, when the other two indicators and firm fixed effects are zero.

Table 3.7: Triple Difference Estimation of the Fuel Economy Standard's Impact on New Vehicle Characteristics (exploiting the 2008 implementation for new models, not continuing models)

Dependent Variable:	I. Torque (nm)	II. Power (kw)	III. Price (nom. $)	IV. Weight (kg)
$\text{Policy}_t^{2008} \cdot \text{Domestic}_j \cdot \text{Continuing}_i$	**20****	**8.6****	**4020*****	**114***
	(4.5)	(1.8)	(312)	(30)
$\text{Policy}_t^{2008} \cdot \text{Domestic}_j$	-17**	-4.8***	-3295***	-110*
	(2.9)	(.39)	(188)	(29)
$\text{Domestic}_j \cdot \text{Continuing}_i$	.068	-.28	-855	-46
	(3.7)	(1.5)	(430)	(26)
$\text{Policy}_t^{2008} \cdot \text{Continuing}_i$	-9.3*	-8.4	-1623***	-21**
	(3)	(4.7)	(159)	(3.9)
Policy_t^{2008}	8.9**	8*	2910**	23*
	(1.8)	(2.4)	(381)	(7.1)
Domestic_j	39***	61***	3799**	258***
	(1.2)	(1.9)	(532)	(3.2)
Continuing_i	-6.9	-1.8	23	27**
	(3.2)	(4.4)	(40)	(4)
Firm f.e.	Y	Y	Y	Y
N	638	641	646	626
R^2	.53	.52	.63	.55

Note: This table reports regression estimates of the effect of the 2008 fuel economy standards on model characteristics (Equation 18). The 2008 policy applied only to new models, not continuing models. Policy_t^{2008} is 1 if the year is 2008, and 0 if 2007 or 2006. Continuing_i variable is 1 if the model is a continuing model in 2008 (i.e. one that was already sold in 2007, like the VW Jetta, and 0 if the model is new, like the Great Wall Peri. Domestic_j is 1 if the brand is domestic (Chinese), and 0 if foreign.\ The unit of observation is the model-year. Only models with at least 1,000 units sold are included. Standard errors are robust and clustered by firm. *** indicates $p < .01$.

I find across specifications that the policy's effect on weight and height is less statistically significant than its effect on the other characteristics. This reflects the weight-based standards, which create perverse incentives to either jump up a weight class or reduce weight within a class. The standards are also more at lenient at each weight class for SUVs and minivans (see Appendix 3A Table 1). Domestic automakers may have responded to the standards *both* by producing more SUVs, which are heavier, and by downsizing sedans and compact vehicles, which made them smaller. Domestic firms produce disproportionately more SUVs and minivans relative to foreign firms.[114] Appendix

[114] Since 2006, when the first and non-binding standards were implemented, domestic firm SUV sales have roughly equalled sedan sales, whereas foreign firm sedan sales have been 6-10 times SUV sales since the 2009

3A Table 3 shows that with controls for vehicle class, the negative effect on weight is larger, but just barely insignificant. Domestic firms may have produced more SUVs because of the policy, or because foreign firms were relatively absent from the segment (e.g. AP 2013). The data do not permit establishing causality of the standards at the class level. The weight distribution of new sales has in general gotten less peaked; a higher proportion of vehicles are either very light or very heavy. This is shown in Appendix A Figure 8. As explained in Section 3.2.2, this in general reduces safety and has been shown to have large negative social welfare effects.

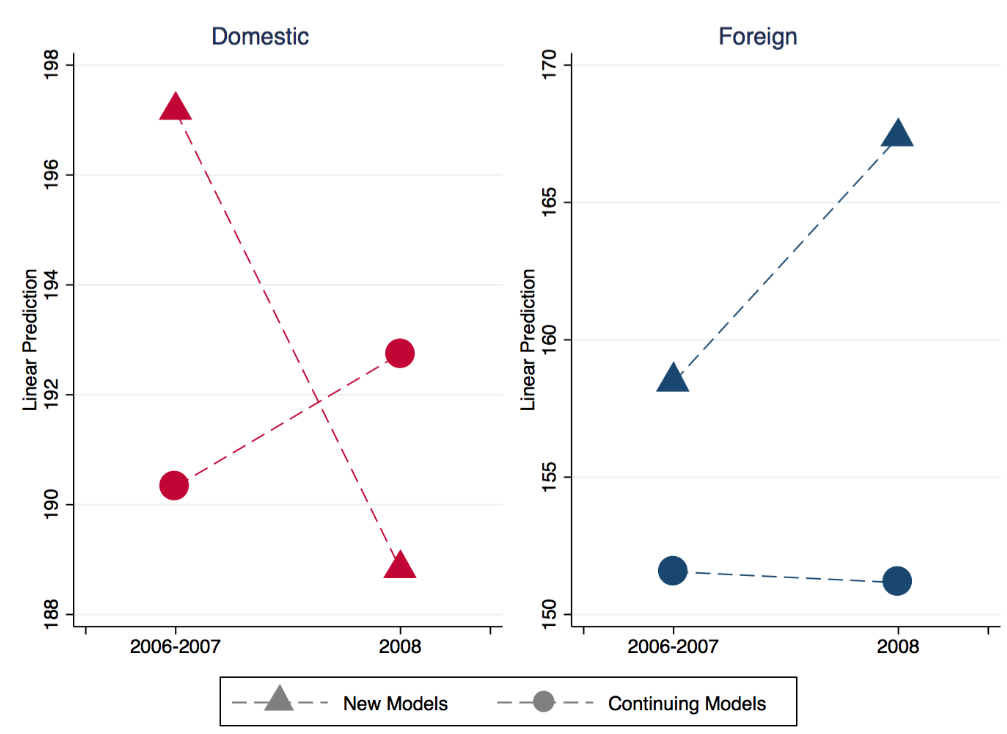

Figure 3.6: Triple Difference Results for Torque: Predictive Margins by Firm and Model Type
Note: This figure shows the predictive margins of the 2008 policy for new models, where torque is the dependent variable. In the left graph, the triangles hold fixed Domestic_j at 1 and Continuing_i at 0, and show how torque declined for new domestic models in 2008 relative to before the policy, in 2006-07. The policy did not apply to continuing models, and the circles show that domestic firms increased torque of continuing models slightly. Similarly, the blue triangles in the right graph show that foreign firms increased torque of new models while keeping that of continuing models roughly constant.

A potential alternative explanation for the effects I observe is that Chinese firms reduced price and vehicle quality to gain market share at the low end of the market, simultaneously but unrelated to the policy. However, they did not gain market share in any segment after the policy,

policy, and more than 20 times SUV sales in earlier years.

making this possibility less likely. In Appendix 3A Figure 4, I show that for models priced below the 25th percentile, domestic firm market share was increasing rapidly until 2009, when it flattened out at a bit more than 80%. Foreign firm model characteristics did not measurably respond to the policy at all, instead continuing along their prior growth path. This is clear in Table 3.2 and Figure 3.4, and also from regressions comparing post-standard firm-type specific trends to prior trends.

3.5.2 Robustness

This section focuses on selected important robustness tests shown for torque in Tables 3.8 and 3.9. It also discusses the same tests for the other characteristics, and further robustness tests, which the reader can find in Appendix 3A.

Standard errors are clustered at the firm level in my primary specification to reduce their potential downward bias from serial correlation of the variables (Bertrand et al. 2004). Columns IV and V of Table 3.8 show that the effect on torque remains significant at the 1% level with no clusters and with firm-year clusters. Appendix 3A Table 4 shows four alternative assumptions on the errors for all four characteristics: homoscedasticity, robust (sandwich estimator), robust with year clusters, and robust with firm-year clusters. The coefficient of interest is highly significant in all four alternatives, for all four characteristics. The bias problem appears most severe with firm-year clusters.

Using all the data instead of a bandwidth around the policy yields a result similar to the primary specification (Table 3.8 column I). Appendix 3A Table 5 shows this larger bandwidth specification for all characteristics, and also illustrates the downward bias in standard errors that occurs in the commonly used DD design of including all years, year fixed effects, and clustering standard errors and at the firm-year level. This specification, in Panel B, yields significant coefficients for all characteristics, a misleading finding, and greater significance for power and weight, which are only moderately significant in the primary specification.

Columns II and III of Table 3.8 vary the required sales volume of models included in the regression. My primary specification, which requires at least 1,000 units sold, yields a coefficient on torque of -17 nm. When the required sales volume is only 100 vehicles, the coefficient declines to -12 nm, and when it is more than 5,000 vehicles, the coefficient is -19 nm. Appendix 3A Table 2 shows how for all four characteristics, the effect becomes stronger and more robust as the sales volume requirement increases.

Table 3.8: Key Robustness Tests of Difference-in-Difference Estimation of the Fuel Economy Standard's Impact on New Vehicle Torque Part 1

Dependent Variable: Torque (nm)									
		Model sales volume		Standard error clustering		Placebo test with artificial policy at year			
Test:	I. All yrs (1999-2012)	II. >100	III. >5000	IV. None (robust)	V. Firm-year	VI. 2006	VII. 2007	VIII. 2008	IX. 2010
$\text{Policy}_t \cdot \text{Domestic}_j$	**-17****	**-12****	**-19*****	**-17*****	**-17*****	**-.44**	**-6.6**	**-11***	**-9***
	(6.9)	(5.5)	(6.1)	(4.4)	(3.7)	(7.3)	(5.6)	(5.9)	(5.2)
Policy_t	15***	11***	13***	11***	11***	3.2	5.1	6.3*	8.1**
	(5.2)	(3.2)	(3.8)	(3.1)	(2.7)	(5.1)	(3.4)	(3.3)	(4)
Domestic_j	65***	62***	-34***	59***	59***	46***	54***	53***	-43***
	(4.1)	(2.5)	(10)	(14)	(5.2)	(7.3)	(5.6)	(3.3)	(9.3)
Firm f.e.	Y	Y	Y	Y	Y	Y	Y	Y	N
N	2339	1927	1180	1646	1646	825	1055	1283	1250
R^2	.48	.5	.16	.5	.5	.52	.5	.49	.16

Note: This table reports regression estimates of the effect of the fuel economy standards on model torque (Equation 17). Sales volume is the number of units sold of that model-year vehicle. Except for columns II and III, only models with at least 1,000 units sold are included. Except for columns VI-VIII, a bandwidth of three years around 2009 policy is used. In columns VI-VIII, a bandwidth of three years around the specified year is used. Column IX does not include firm fixed effects as they are collinear with the Domestic indicator. Domestic_j is 1 if the brand is domestic (Chinese), and 0 if foreign. Policy_t is 1 if the year is 2009 or later, and 0 if 2008 or before. The unit of observation is the model-year. Except for regressions IV and V, standard errors are robust and clustered by firm. See Appendix A for a wide variety of additional tests, and all these tests using other dependent variables. *** indicates $p < .01$.

I conduct placebo tests for every year possible with varying bandwidths. That is, I estimate the DD design as though the policy had been implemented in an alternative year. I did not find a significant effect in any, except moderate significance in placebo tests where the bandwidth includes the actual policy year of 2009. Columns VI-IX of Table 3.8 show the placebo test results for 2006, 2007, 2008, and 2010. In 2006, the coefficient is -.44 nm, and in 2007 it is -6.6 nm, neither with any significance. The 2008 and 2010 placebo tests yield impacts of -11 and -9 nm, both significant at the 10% level. Note that these bandwidths include the policy. Appendix 3A Table 6 shows the placebo tests for all four characteristics with the artificial year set to either 2005, 2006, 2007 or 2008. The reader may be concerned that the global recession coincided with the policy. However, China recovered quickly relative to other countries in the second half of 2009, returning to its pre-crisis growth path by 2010 (Diao et al. 2012).

Table 3.9: Key Robustness Tests of Difference-in-Difference Estimation of the Fuel Economy Standard's Impact on New Vehicle Torque Part 2

Dependent Variable: Torque (nm)							
Test:	I. No individual effects	II. No interaction	III. No f.e.	IV. Year and firm f.e.	V. Year f.e.	VI. Class f.e.	VII. OEM f.e.
$\text{Policy}_t \cdot \text{Domestic}_j$	**-40*** **		**-20*** **	**-17*** **	**-20*** **	**-15** **	**-16*** **
	(7.2)		(6.2)	(5.3)	(6.2)	(5.8)	(5.6)
Policy_t		4	11***	15***	15**	9.4**	11**
		(3.5)	(4.1)	(5.5)	(6)	(3.8)	(4)
Domestic_j		-44***	-32***	58***	-32***	-38***	-25**
		(9.1)	(9.9)	(2.9)	(9.9)	(8.7)	(9.6)
$\mathbf{1}\ \|\text{Minivan}_i$						51***	
						(12)	
$\mathbf{1}\ \|\text{SUV}_i$						82***	
						(14)	
$\mathbf{1}\ \|\text{Sedan}_i$						47***	
						(8.1)	
Firm f.e.	N	N	N	Y	N	Y	N
Year f.e.	N	N	N	Y	Y	N	N
OEM f.e.	N	N	N	N	N	N	Y
N	1646	1646	1646	1646	1646	1646	1646
R^2	.084	.13	.14	.5	.14	.23	.42

Note: This table reports regression estimates of the effect of the 2009 fuel economy standards on model torque, with a bandwidth of three years around 2009 policy (Equation 17). In column VI, there are fixed effects whether the vehicle class is Compact, Minivan, SUV and Sedan (Compact is omitted). In column VII, OEM refers to Original Equipment Manufacturer. Domestic_j is 1 if the brand is domestic (Chinese), and 0 if foreign. Policy_t is 1 if the year is 2009 or later, and 0 if 2008 or before. The unit of observation is the model-year. Standard errors are robust and clustered by firm, except in column VII, where they are clustered at the OEM level. See Appendix A for a wide variety of additional tests, and all these tests using other dependent variables. *** indicates $p < .01$.

Alternative individual and fixed effects are considered in Table 3.9, also using torque as the dependent variable. Column I shows that when I exclude the individual effects ($\text{Policy}_t, \text{Domestic}_j$) the coefficient on the interaction term increases in magnitude to -40 nm. With no interaction term, in column II, the effect of being a domestic firm is -44 nm, and policy has an insignificant effect of 4. I omit firm fixed effects here so that they do not soak up the negative effect of being domestic. Appendix 3A Table 7 shows these specifications for all the characteristics. The positive effect of being domestic on vehicle height is because a larger share of domestic firm production is SUVs. The primary specification without firm fixed effects produces a slightly stronger effect of the policy on domestic firm torque of -20 nm (column III). Both year and firm fixed effects gives a coefficient of -17 nm (column IV). Appendix 3A Table 8 shows that for all four characteristics, there is a significant negative effect of the policy regardless of whether I use year, firm and year, or no fixed effects. The coefficients are all of similar magnitude to my primary specification, but somewhat smaller with both firm and year fixed effects.

Column VI of Table 3.9 considers vehicle class fixed effects, in addition to firm fixed effects. There are four vehicle classes: compact, sedan, minivan, and SUV.[115] In the regression, the omitted class is compact; as expected, the other three classes have large positive effects on torque relative to compact cars. The effect of sedans or minivans relative to compacts is about 50 nm, and 82 nm for SUVs. The coefficients for all the characteristics decline slightly (shown in Appendix 3A Table 3); the coefficient on torque is now -15 nm. Finally, column VII uses OEM fixed effects, and also clusters standard errors at the OEM level. As discussed in Section 3.3, some OEMs have multiple brands, which are treated separately (and called firms) in the main specification. The number of groups is smaller at the OEM level; there are 69 groups (and so clusters) in OEM compared to 78 groups in firm. The coefficients and their significance are essentially unchanged with OEM fixed effects. The coefficient on torque is -16 nm, and the other characteristics are shown in Appendix 3A Table 3.

Further test in Appendix 3A include a bandwidth of only one year (data only from 2009 and 2008). New models already faced the standard in 2008, so even though the majority of models are continuing, I would expect that the result would be more diluted with this specification. Appendix 3A Table 9 shows that it generates very similar effects on torque and power as the primary specification. However, the other characteristics lose their significance. Appendix 3A Table 10 adds covariates to the specifications, including vehicle class, weight, height and length. As these are correlated with power and to a lesser degree torque, the effects decline and lose some of their precision.

Appendix 3A also contains a variety of robustness tests for the triple difference estimation. Table 11 shows that using alternative required sales volumes yields similar coefficients as the primary specification, with equal or more precision. Table 12 shows that the primary specification is sensitive to fixed effects; with no fixed effects at all, power and price lose their significance, but torque and

[115]The DRC data included three additional classes. I include mini vehicles in the compact category, sedan hatchbacks in the sedan category, and pickup trucks in the SUV category.

weight remain positive and significant, albeit of smaller magnitude. However, with OEM fixed effects neither torque, power, nor price are significant; though weight actually increases slightly in magnitude *and* improves its significance from the 10% level to the 5% level. Thus there seems to be strong evidence that domestic firms down-weighted new models relative to continuing models in order to meet the 2008 standards. It may have been cheaper or faster to initially reduce weight in certain new models being prepared for production rather than alter the engine, transmission, and other components.

Finally, Appendix 3A Tables 13 and 14 show placebo tests for the triple difference design, using the years 2005, 2007, 2009, and 2011. In order to have adequate data, I require the sales volume to be more than 100 units. The placebo tests generate negative coefficients for 2005, and mostly positive coefficients for the later years, but all of these are insignificant except for price in 2007, which has a coefficient of $899 (less than 1/4 the estimated policy effect).

3.6 The Role of Joint Ventures and State Ownership

The domestic Chinese auto manufacturers can be divided along two dimensions: state ownership and joint venture (JV) status. This section addresses firms' performance and response to the standards along these two dimensions.

3.6.1 Background on JVs and Hypotheses about Firm Incentives

Much of the literature on Chinese economics demonstrates the inefficiency of state owned enterprises (SOEs) compared to private firms (e.g. Khandelwal et al. 2012, Bajona and Chu 2010, Jefferson et al. 2003, Lin et al. 1998). However, in some high-tech sectors, such as shipping, SOEs have become globally competitive, dominating the domestic market and making significant inroads into the the export market (see Appendix 3B). Recent work suggests Chinese SOEs are gaining in size and profits relative to the private sector. Hsieh and Song (2015) show that in the 2000s SOEs had faster total factor productivity growth than private firms and higher labor productivity, but lower capital productivity. My data includes firms that are majority owned by provincial governments (local SOEs), the central government (central SOEs), and privately owned firms. Many of the SOEs are partially listed on stock exchanges.

Both SOEs and privately owned firms have joint ventures (JVs) with foreign firms. A JV enterprise is the foreign firm's China operation; for example, all of BMW's China production occurs in a JV with Brilliance Auto, a domestic firm. But the JV manufacturing plant produces only BMW models. Brilliance receives 50% of the profits from each BMW sold, and provides a non-disclosed portion of the fixed and operating costs. The industry press suggests that the JVs failed to achieve technology transfer to domestic firms and that foreign partners remain the JVs' source of technology (Holmes et al 2013, Economist 2013, Sanford C. Bernstein 2013). Foreign firms operating in China source 25-75% of their parts in China, but the most advanced parts are still imported (Takada 2013, Yang 2008). Through case studies, Gallagher (2006) concludes that domestic Chinese companies have remained dependent on their foreign partners for technology, not learning how to innovate or design vehicles.

A bit of history illuminates this stylized fact. The first JV was announced in 1984 between SAIC, owned by the Shanghai government, and Volkswagen (VW). VW had substantial bargaining power and benefited from information asymmetry about auto manufacturing. Though the balance of power shifted as China's market grew, incomplete contracting and moral hazard continued to bedevil implementation of the JV arrangements (Thun 2004). To start fresh – and put pressure on VW – SAIC signed a second JV agreement with GM in 1997. GM aggressively marketed itself as a purveyor of useful technology, establishing a joint research center with SAIC called the Pan Asia Technical Automotive Center (PATAC). But PATAC was and continues to tweak existing GM-branded models for the Chinese market, not design new models. Further, most GM-branded models initially chosen for China were Daewoo or Opel designs, further distancing GM's Chinese operation from Detroit's state-of-the-art (Tang 2012).

A Wall Street Journal article concluded in 2012 that

> "Chinese auto regulators find themselves in a tight spot: their 30-year quest to build an industry dominated by Chinese car brands has backfired. The problem: joint ventures with foreign carmakers that have proven just a tad too comfortable."

According to He Guangyan, a former machinery industry minister, the JVs are "like opium" for the domestic firms (Dunne 2012).

China sought to exchange foreign access to the domestic market for technology transfer. Yet dynamically this industry structure may have reduced domestic firm innovation incentives. A domestic firm in a JV is disincentivized from designing substitutes to its foreign partner's vehicles because doing so would cannibalize the rents it earns from JV profits. To illustrates the intuition behind this tradeoff, consider the following stylized profit functions of domestic firm j with and without joint ventures, where ϕ is the technology quality of a model-year, and s is the share (typically

50%) of the foreign firm's profits that a firm with a JV earns:

$$\text{Firm without JV: } \pi_j = \sum_{i \in j} q_i(\mathbf{p}, \phi_i) \left(p_i - C_{i,\text{No } JV} \right) \tag{19}$$

$$\text{Firm with JV: } \pi_j = s\pi_{JV}^{foreign} + \sum_{i \in j} q_i(\mathbf{p}, \phi_i) \left(p_i - C_{i,JV} \right) \tag{20}$$

Suppose that ϕ is an increasing function of torque and horsepower: $\frac{\partial \phi_i}{\partial \text{Torque}_i} > 0$; $\frac{\partial \phi_i}{\partial \text{Power}_i} > 0$. The equilibrium vehicle price, p_i, increases with quality $\left(\frac{\partial p_i}{\partial \phi_i} > 0 \right)$, and also depends on all models in the market. The firm's cost function $(C_i = \mathcal{F}(\cdot, \phi_\text{i}))$ is also increasing in quality $\left(\frac{\partial \mathcal{F}}{\partial \phi} > 0 \right)$.

Suppose that fuel economy standard implementation requires that for a model with a given weight, the firm must invest some fixed additive amount to acquire fuel efficiency technology in order to meet the standards and maintain torque and power at their previous levels. Now $C_i = \mathcal{F}(\cdot, \phi_\text{i} + \text{F}_\text{j}(\phi_\text{i}) \mid \text{Weight}_\text{i})$.[116] Firms with JVs may have greater access to foreign firm technology than firms without JVs, so $F_{j \in JV} \leq F_{j \in No\ JV}$. The foreign firm already possesses the technology, so within its cost function $F_{foreign} = 0$. Thus holding other aspects of the cost function fixed, it may be cheaper for firms with JVs to provide a unit of quality (high torque and power) than firms without JVs under the standards; $\frac{\partial C_{i,JV}}{\partial \phi_i} \leq \frac{\partial C_{i,No\ JV}}{\partial \phi_i}$.

If domestic firms are price-takers, which is likely given their low concentration, then conditional on a given price vector for all models in the market, quantity sold should increase with quality and decrease with model price: $\frac{\partial q_i(\mathbf{p},\phi_i)}{\partial \phi_i} > 0$ and $\frac{\partial q_i(\mathbf{p},\phi_i)}{\partial p_i} < 0$. The first order conditions in quality are:

$$\begin{aligned} \text{Firm without JV: } \frac{\partial \pi_j}{\partial \phi_i} &= q_i(\mathbf{p}, \phi_i) \left[\frac{\partial p_i}{\partial \phi_i} - \frac{\partial C_{i,No\ JV}}{\partial \phi_i} \right] \\ &\quad + \frac{\partial q_i(\mathbf{p}, \phi_i)}{\partial \phi_i} \left(p_i - C_{i,No\ JV} \right) + \sum_{k \neq i \in j} \left[\frac{\partial q_k(\mathbf{p}, \phi_k)}{\partial \phi_i} \left(p_k - C_{k,No\ JV} \right) \right] \end{aligned} \tag{21}$$

[116] For simplicity, suppose the fixed cost is spread equally across models and firms have equal numbers of models.

$$\text{Firm with JV: } \frac{\partial \pi_j}{\partial \phi_i} = s\frac{\partial \pi_{foreign,\ JV}}{\partial \phi_i} + q_i(\mathbf{p}, \phi_i)\left[\frac{\partial p_i}{\partial \phi_i} - \frac{\partial C_{i,JV}}{\partial \phi_i}\right] + \frac{\partial q_i(\mathbf{p}, \phi_i)}{\partial \phi_i}(p_i - C_{i,JV}) + \sum_{k \neq i \in j}\left[\frac{\partial q_k(\mathbf{p}, \phi_k)}{\partial \phi_i}(p_k - C_{k,JV})\right] \tag{22}$$

Using the assumptions outlined above, the foreign firm's profit decreases in a competitor's quality $\left(\frac{\partial \pi_{foreign,\ JV}}{\partial \phi_i} < 0\right)$. The domestic firm's investment in own quality reduces its marginal profit from the JV.[117]

Is the equilibrium ϕ for a firm with a JV greater than that for a firm without a JV? Holding all else equal between the two types of firms, this depends on whether the negative effect on ϕ of access to the foreign firm's profits outweighs the positive effect of a lower technology acquisition cost:

$$\phi_{i,JV} > \phi_{i,No\ JV} \text{ if } s\left[\frac{\partial \pi_{foreign,\ JV}}{\partial \phi_i}\right] - \frac{\partial C_{i,JV}}{\partial \phi_i} > -\frac{\partial C_{i,,No\ JV}}{\partial \phi_i} \tag{23}$$

where the the first term has a negative sign, and the second two have positive signs (recall that $\frac{\partial C_{i,JV}}{\partial \phi_i} < \frac{\partial C_{i,,No\ JV}}{\partial \phi_i}$). In the following section, I try to test whether $\phi_{i,JV} > \phi_{i,No\ JV}$ or vice versa.

3.6.2 Descriptive Statistics

This section shows that private firms have achieved substantially better sales and revenue growth than state-owned firms, and have also dominated China's small volume of passenger vehicle exports. Domestic firms with JVs have higher sales and revenue than firms without JVs, but almost all of the exporting is done by firms without JVs. However, a challenge to separately assessing the association of JVs and state ownership with performance is the overlap between the two categories. Table 3.10 is a matrix of the number of firms (brands) and OEMs in each category. At the firm level, of the 43 SOEs, 37 have a JV. Of the 25 privately owned firms, only 3 have a JV. The correlation between state ownership and having a JV is 0.7. In the estimations below, I isolate firms that fall into the non-SOE, JV cell and the private, non-JV cell, but the reader should keep in mind that there are few firms in these categories.

[117] All firms have the same variable cost of producing more fuel efficient vehicles.

Table 3.10: Domestic Firm Ownership Matrix

	Firm (brand) level		
	SOE	Privately owned	**Total**
Firms with JV	37	3	**40**
Firms without JV	6	22	**28**
Total	**43**	**25**	

	OEM level		
	SOE	Privately owned	**Total**
Firms with JV	20	2	**22**
Firms without JV	5	20	**25**
Total	**25**	**22**	

Note: This table shows the number of unique firms and OEMs that fall into various categories: being a locally or centrally state owned enterprise (SOE), being privately owned, having a joint venture (JV) with a foreign firm, and not having a JV. Note that many firms classified as SOEs are only majority held by the state and are partially privately owned and even publicly listed.

The top left graph of Figure 3.7 shows that central SOEs, local SOEs, and private domestic firms have experienced quite similar sales volume trajectories since 2002. However, the bottom left graph shows that since 2005, private firms have experienced much higher revenue growth and today have annual revenue that is twice the level of either local SOEs or central SOEs. There does not seem to be any appreciable difference between local and central SOEs.[118] The right-hand graphs show that firms with JVs have experienced higher sales volume and much higher annual revenue than firms without JVs. This is partly because firms without JVs export a large fraction of their production. All four graphs contrast with Figure 3.8, which shows sales volume and revenue for foreign firms (note the difference in scale).

[118] I estimate revenue by multiplying each model's sales volume by its price, and then summing annually over brands within categories.

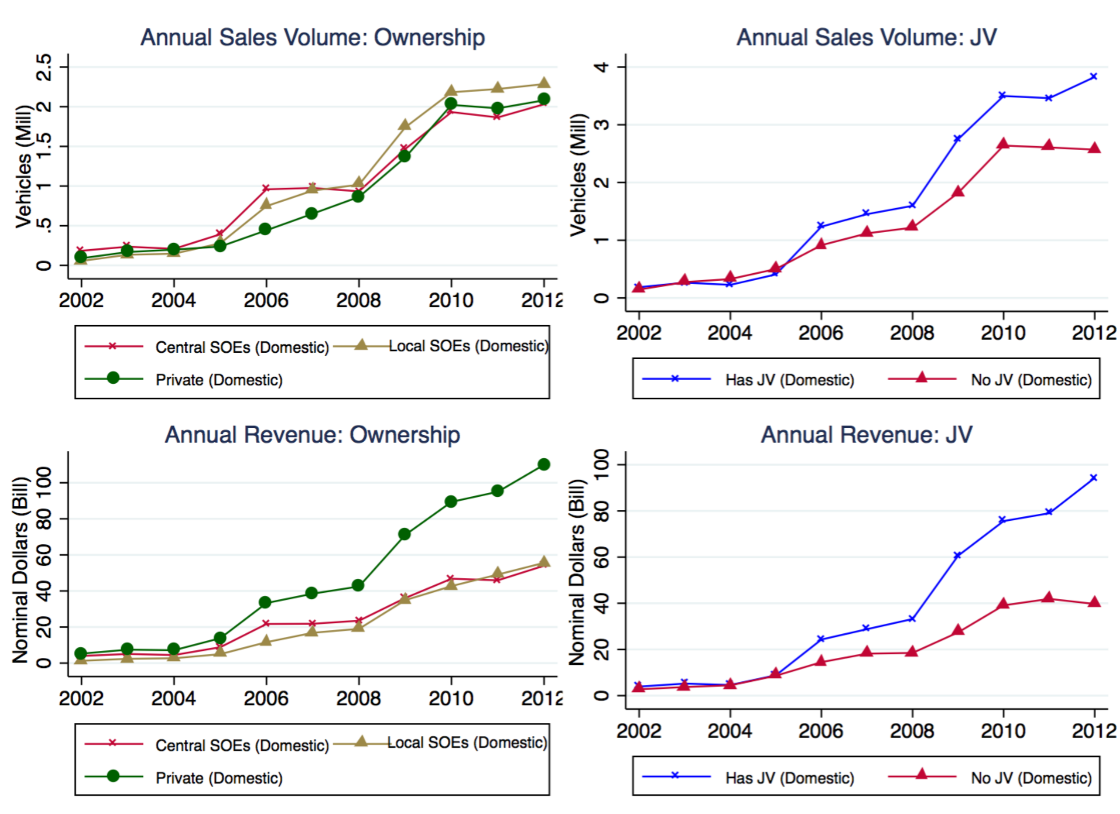

Figure 3.7: Total Sales Volume & Revenue among Domestic Firms by Ownership and JV Status
Note: This figure shows sales volume in the top two graphs, and revenue in the bottom two graphs. The left two graphs divide the firms by ownership type, where Central State Owned Enterprises (SOEs) are owned by China's central government and local SOEs are owned by provincial governments. The right two graphs divide the firms by whether they have a joint venture (JV) with a foreign firm or not.

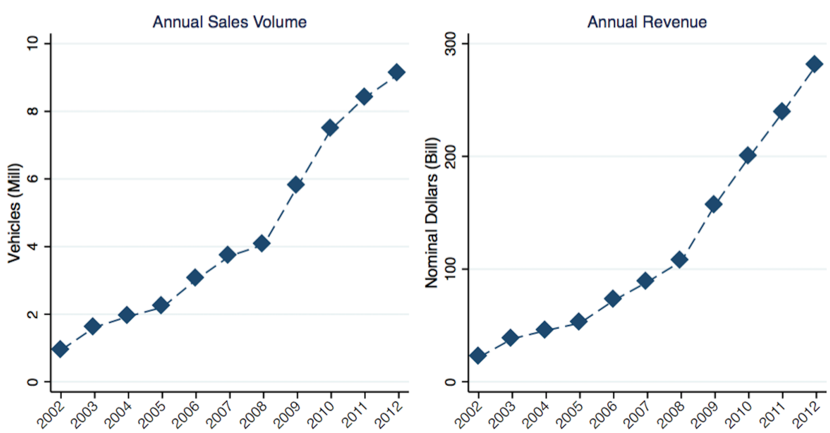

Figure 3.8: Total Sales Volume and Estimated Revenue among Foreign Firms
Note: This figure shows annual sales volume and revenue for foreign firms' sales in China. All foreign firm production occurs in JVs with Chinese firms, but there is no overlap with Figure 3.7 because the Chinese firms produce own-branded models, while the JVs produce, typically, only foreign brands.

Exporting is strongly associated with firm productivity and competitiveness, (Melitz and Redding 2014, Giles and Williams 2000, Wakelin 1998). It is also an explicit Chinese auto industrial policy goal (State Council 2009). Yet total domestic firm exports in 2012 were 0.6 million vehicles compared to production for domestic consumption of about 6 million vehicles. Although exports are increasing rapidly, it is clear that the industry is far from meeting government export targets (Roland Berger 2013). One reason for the failure to export is a number of high profile Western crash test failures. In 2007, Germany and Russia tested Chinese sedans made by Brilliance Jinbei and Chery, respectively. The former crash test was described by the German officials as "catastrophic," while the Russian testers described the performance as among the worst it had ever encountered (Osborn 2007).

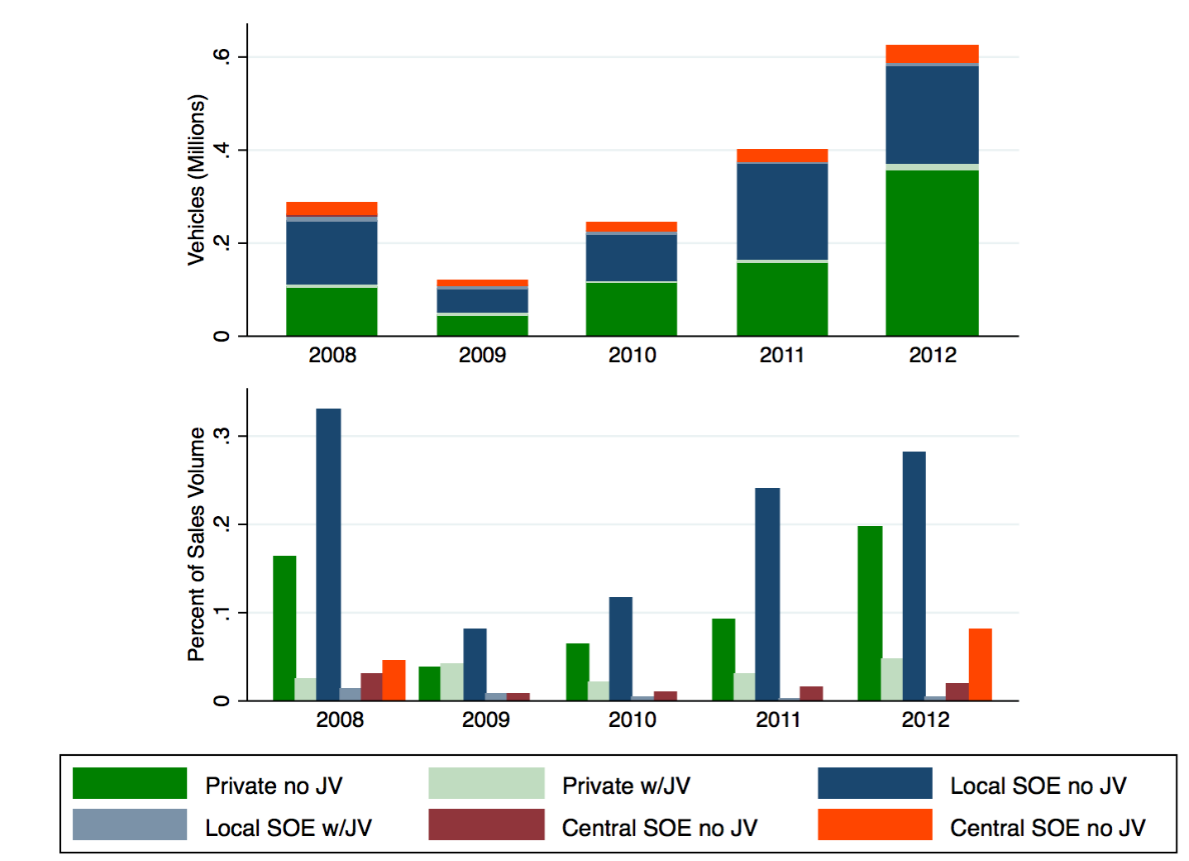

Figure 3.9: Domestic Firm Export Volume and Percent of Total Sales Volume 2008-2012
Note: This figure shows Chinese domestic firm vehicle exports (there are essentially no foreign firm exports from China). The top graph shows the annual number of vehicles exported by ownership type, and the bottom graph shows this number as a share of the firms' total sales volume. For example, the first green bar in the bottom graph shows the total number of vehicle exports divided by the total number of vehicles sold among all firms who are privately owned and do not have a joint venture (JV).

My data reveals that private firms and local SOEs without JVs have been responsible for almost all passenger vehicle exports, depicted in Figure 3.9.[119] The private and local SOEs without JVs are very small, so their exports as a percentage of sales is high. Between 2008 and 2012, private firms without JVs exported 10-20% of their total sales, and local SOEs without JVs exported 10-30%.

Descriptive statistics offer no obvious differentiation in model characteristics across firm types. Appendix 3A Table 15 shows average model characteristics by firm ownership type and JV status. Appendix 3A Figure 5 depicts domestic brand sales-weighted characteristics over time. There are no significant differences across private and state-owned firms or firms with and without JVs. Central SOE models on average are somewhat less powerful and command lower prices, but high variation within categories means that the differences in means are not significant. For example, the mean model price between 2009 and 2012 for private firms was \$12,500, for central SOEs \$11,600 and for local SOEs \$12,700. Mean torque was 151 nm for private firms, 137 nm for central SOEs, and 152 nm for local SOEs.

3.6.3 Response to the Fuel Economy Standards by Firm Type

I evaluate whether incentives to innovate and acquire technology vary by firm type using the DD design proposed in Section 3.4, estimated on subsamples of domestic firms. I find that the strong negative effect of the policy on measures of quality appears concentrated in firms with JVs, as well as in SOEs. First, I perform regressions on separate subsamples (Table 3.11), and then combine them into a single specification (Table 3.12).

The dependent variables in Table 3.11 are torque in the left panels, and price in the right panels.[120] Columns I(a)-I(c) include only domestic firms with JVs and foreign firms; I(a) includes all such firms, I(b) only state-owned firms with JVs, and I(c) only privately owned firms with JVs. The coefficients on the interaction term are all negative, significant, and of at least as large a magnitude

[119]The biggest exporters are Great Wall (privately owned, Hebei province-based, listed on the Shanghai stock exchange with no JV), Chery (SOE of the Anhui provincial government with no JV), Geely (privately owned, listed on the Hong Kong stock exchange with no JV), JAC Motors (SOE majority owned by the Anhui provincial government and partially listed on the Shanghai stock exchange with no JV), and Lifan (privately owned, listed on the Shanghai stock exchange with no JV). My classification of JV status is by year of sales and ends in 2012. Some companies have since established JVs. Chery, after many previous abortive attempts to form JVs, began producing vehicles through a JV with an Israeli company under the brand name Qoros in 2013, and in 2012 allied with Tata to produce Jaguar Land Rover models in China from 2014.

[120]These regressions do not include firm fixed effects because the number of firms in certain categories is quite small.

as the primary specification with all firms. Specification III considers SOEs, and specification V considers private firms. Looking down the first column, the strongest negative impact of the policy is for the subset of firms with JVs. Columns III(b) and V(b) show that for the subset of SOEs without JVs and privately owned firms without JVs, the coefficient is smaller and less precise. Further, it seems that private firms reduced model maximum torque more than SOEs in response to the policy; that is, the coefficients in specification V are more negative than those in specification III.

The right-hand panels of Table 3.11 conduct the same exercise with price as the dependent variable. In column (a), the strongest result is for the subset of firms with JVs. This is also clear in comparing SOEs with and without JVs (IV(b-c)) and private firms with and without JVs (VI(b-c)). Firms with JVs reduced model price by \$3,458 relative to foreign firms after the policy, compared to \$2,791 for SOEs without JVs and \$2,586 for privately owned firms without JVs. SOEs and privately owned firms with JVs, in contrast, decreased price by \$3,378 and \$3,750, respectively. Here there is no appreciable difference betweenSOE and private firm price reduction.

I combine these effects into a single regression in Table 3.12.[121] I interact the policy with indicator variables for firm type and firm fixed effects. Column I shows that firms with a JV reacted more strongly than firms without a JV to the policy, relative to foreign firms. For firms with a JV, the interaction coefficient is -24 nm, significant at the 1% level. For firms without a JV it is -12 nm, significant at the 10% level. A similar result for price is in column IV. Columns II and V show that SOEs decreased torque and price much more than private firms; the coefficient on the interaction for price, for example, is -\$3,473 for SOEs and -\$1,951 for private firms. Columns III and VI subdivide firms without JVs into SOEs and privately owned firms. The policy's effect on SOEs without JVs is much more negative and more precise than that on private firms.

Wald tests on the Table 3.12 regressions reject the hypothesis that the coefficients on the interactions with price as the dependent variable are equal at the 10% level. However, I am not able to reject the hypothesis that they are equal for the torque specifications. These regressions use a stringent standard error assumption, clustering at the firm level. When I cluster at the firm-year level, as is often done in DD designs, I can reject the null that firms with and without JVs responded similarly to the policy with 95% confidence (p-value of 0.02) for torque as well as price. This is also true when I cluster at the year level or do not cluster at all.

[121] Appendix 3A Table 16 conducts the same regressions as in Table 3.12, for weight and power, with analogous results.

Table 3.11: Difference-in-Difference Estimation of Policy's Impact in Firm Type Subsamples

Dependent Variable:	I. Torque (nm)			II. Price (nominal dollars)		
Domestic sample:	a. All Firms w/ JV	b. SOEs w/JVs	c. Privately owned w/JVs	a. All Firms w/ JV	b. SOEs w/JVs	c. Privately owned w/JVs
$\text{Policy}_t \cdot \text{Domestic}_j^{JV}$	**-22****	**-20***	**-31*****	**-3458*****	**-3378****	**-3750*****
	(8.9)	(10)	(11)	(1290)	(1430)	(941)
Policy_t	11**	11**	11**	3109***	3109***	3109***
	(4.2)	(4.2)	(4.2)	(905)	(905)	(913)
Domestic_j^{JV}	-34***	-39***	-14	-10996***	-11577***	-8551***
	(12)	(13)	(9.2)	(2492)	(2582)	(2179)
N	1295	1242	1068	1303	1251	1081
R^2	.11	.11	.023	.12	.11	.028
Dependent Var:	III. Torque (nm)			IV. Price (nominal dollars)		
Domestic sample:	a. All SOEs	b. SOEs w/o JVs	c. SOEs w/ JVs	a. All SOEs	b. SOEs w/o JVs	c. SOEs w/ JVs
$\text{Policy}_t \cdot \text{Domestic}_j^{SOE}$	**-16****	**-10**	**-20***	**-3144*****	**-2791****	**-3378****
	(6.8)	(9.4)	(10)	(1131)	(1193)	(1430)
Policy_t	11**	11**	11**	3109***	3109***	3109***
	(4.2)	(4.2)	(4.2)	(904)	(911)	(905)
Domestic_j^{SOE}	-41***	-44***	-39***	-11968***	-12580***	-11577***
	(10)	(11)	(13)	(2359)	(2367)	(2582)
N	1381	1154	1242	1388	1166	1251
R^2	.14	.08	.11	.15	.086	.11
Dependent Var:	V. Torque (nm)			VI. Price (nominal dollars)		
Domestic sample:	a. All privately owned	b. Privately owned w/o JV	c. Privately owned w/ JV	a. All privately owned	b. Privately owned w/o JV	c. Privately owned w/ JV
$\text{Policy}_t \cdot \text{Domestic}_j^{Priv.}$	**-21****	**-18***	**-31*****	**-2773****	**-2586***	**-3750*****
	(8.7)	(10)	(11)	(1351)	(1492)	(941)
Policy_t	11**	11**	11**	3109***	3109***	3109***
	(4.2)	(4.2)	(4.2)	(907)	(908)	(913)
$\text{Domestic}_j^{Priv.}$	-23*	-25	-14	-11442***	-12112***	-8551***
	(13)	(15)	(9.2)	(2520)	(2595)	(2179)
N	1280	1227	1068	1294	1242	1081
R^2	.066	.058	.023	.12	.11	.028

Note: This table reports regression estimates of the effect of the fuel economy standards on model torque and price, using a bandwidth of 3 years around the policy (Equation 17). Only certain subsets of domestic firms are used, as described in each specification. For example, I.a. compares domestic firms with joint ventures (JVs) with foreign firms, before and after the policy (domestic firms without JVs excluded). Only models with at least 1,000 units sold are included. Domestic_j is 1 if the brand is domestic (Chinese), and 0 if foreign. Policy_t is 1 if the year is 2009 or later, and 0 if 2008 or before. The unit of observation is the model-year. Standard errors are robust and clustered by firm. *** indicates $p < .01$.

Table 3.12: Difference-in-Difference Estimation of the Fuel Economy Standard's Impact on Firm Type Subsamples in Single Regression

Dependent Variable:	Torque (nm)			Price (nominal dollars)		
	I.	II.	III.	IV.	V.	VI.
$\text{Policy}_t \cdot \text{Domestic}_j^{JV}$	**-24*** **		**-24*** **	**-3557*** **		**-3554*** **
	(7.9)		(8)	(848)		(848)
$\text{Policy}_t \cdot \text{Domestic}_j^{\text{no JV}}$	**-12* **			**-2223** **		
	(6.4)			(903)		
$\text{Policy}_t \cdot \text{Domestic}_j^{SOE}$		**-20*** **			**-3473*** **	
		(6)			(802)	
$\text{Policy}_t \cdot \text{Domestic}_j^{Priv.}$		**-13* **			**-1951** **	
		(7.8)			(964)	
$\text{Policy}_t \cdot \text{Domestic}_j^{\text{SOE no JV}}$			**-18*** **			**-3364*** **
			(4.4)			(1011)
$\text{Policy}_t \cdot \text{Domestic}_j^{\text{Priv. no JV}}$			**-8.6**			**-1568**
			(8.6)			(1044)
Policy_t	11***	11***	11***	2858***	2858***	2858***
	(3.5)	(3.5)	(3.5)	(630)	(630)	(631)
Domestic_j^{JV}	59***		59***	4508***		4508***
	(2.7)		(2.7)	(490)		(490)
$\text{Domestic}_j^{\text{no JV}}$	53***			4579***		
	(7.3)			(862)		
Domestic_j^{SOE}		59***			4508***	
		(2.7)			(490)	
$\text{Domestic}_j^{Priv.}$		54***			6258***	
		(5.6)			(697)	
$\text{Domestic}_j^{\text{SOE no JV}}$			54***			4204***
			(7.2)			(722)
$\text{Domestic}_j^{\text{Priv. no JV}}$			50***			6177***
			(7.6)			(750)
Firm f.e.	Y	Y	Y	Y	Y	Y
N	1646	1646	1646	1653	1653	1653
R^2	.5	.5	.5	.63	.63	.63

Note: This table reports regression estimates of the effect of the fuel economy standards on model torque and price, using a bandwidth of 3 years around the policy (Equation 17). Only models with at least 1,000 units sold are included. Domestic_j^X is 1 if the brand is domestic (Chinese), and 0 if foreign, and fits into the category X (e.g. not being a SOE). Policy_t is 1 if the year is 2009 or later, and 0 if 2008 or before. The unit of observation is the model-year. Standard errors are robust and clustered by firm. *** indicates $p < .01$.

In sum, SOEs with JVs appear primarily responsible for the domestic firm response to the policy of reducing model quality and price. Private firms without JVs responded the least to the policy, though the results for private firms are in general less precise. The equilibrium quality choice for firms without JVs after the policy is higher than for firms with JVs, or in the terminology of the toy model above, $\phi_{i,JV} < \phi_{i,No\ JV}$. This is consistent with a story in which the negative effect of own ϕ_i on the foreign partner's profits $\left(\frac{\partial \pi_{foreign,\ JV}}{\partial \phi_i}\right)$ outweighs any technology acquisition cost advantage that the domestic firm with a JV may have $\left(\frac{\partial C_{i,JV}}{\partial \phi_i} < \frac{\partial C_{i,,No\ JV}}{\partial \phi_i}\right)$.

3.7 Conclusion

Using a novel and reliable data set, I assess how the fuel economy standards affected the model characteristic decisions of domestic Chinese automakers. Through a differences-in-differences design, I show that domestic Chinese firms responded to the 2009 implementation of fuel economy standards by reducing the torque, horsepower, weight, and price of their models. A triple differences design exploiting the staged policy for new and continuing models in 2008-09 finds that when domestic firms' continuing models were not yet subject to the policy, they were more powerful, more expensive, larger, and heavier than new models.

I then turn to the relative performance of Chinese firms by ownership and JV status. I show in a simple model that domestic firms might be disincentivized from producing substitutes for their foreign partners' models, even when they have a lower cost of technology acquisition than domestic firms without JVs. Competing with the foreign partner would cannibalize the domestic firm's share of foreign brand profits. I show that SOEs with JVs were primarily responsible for the negative effects of the policy on domestic firm quality and price. The negative effect of having a JV appears stronger than the negative effect of being state-owned, consistent with a story in which the negative effect of increasing own quality on the firm's share of JV profits outweighs any technology acquisition cost advantage that the firm reaps from its JV.

Conventional trade models, such as McGrattan and Prescott (2009, 2010) grossly overestimate both FDI inflows to and outflows from China, due to an assumption that foreign firms bring their technological capital to China, which Chinese firms accumulate. When Holmes, McGrattan and Prescott (2013) add China's requirement that foreign firms transfer technology in order to invest, they are much better able to match their multicountry dynamic general equilibrium model to moments in the data. FDI decreases when foreign firms must transfer technologies. They also find that JV-owned patents tend not to extend beyond China's borders; their primary case study is GM's patents with SAIC. They conclude that less foreign capital enters due to the "technology capital tax," and Chinese firms prefer to appropriate the foreign capital rather than innovate themselves. This is precisely what I observe in China's auto sector: foreign firms bring minimum technology to China because they cannot protect their intellectual property. Chinese firms, especially those with joint ventures, do little innovation on their own because a distorted market structure disincentivizes them from acquiring technology.

References - Chapter 1

Acemoglu, D., Akcigit, U., Bloom, N. & Kerr, W. R. 2013. Innovation, reallocation and growth. NBER Working Paper 18993.

Aghion, P. & Bolton, P. 1992. An incomplete contracts approach to financial contracting. Review of Economic Studies 59.

Aghion, P., Dewatripont, M. & Stein, J. C. 2008. Academic freedom, private sector focus, and the process of innovation. The RAND Journal of Economics, 39(3).

Aghion, P., Askenazy, P., Berman, N., Cette, G., & Eymard, L. 2012. Credit constraints and the cyclicality of R&D investment: Evidence from France. Journal of the European Economic Association, 10(5).

Aigner, D. J. & Cain, G. G. 1977. Statistical theories of discrimination in labor markets. Industrial and Labor relations review.

Akcigit, U. & Kerr, W. R. 2011. Growth through Heterogeneous Innovations. NBER Working Paper No. 16443.

Almus, M., & Czarnitzki, D. 2003. The effects of public R&D subsidies on firms' innovation activities: the case of Eastern Germany. Journal of Business & Economic Statistics, 21(2).

Angrist, J. D. 2001. Estimation of Limited Dependent Variable Models with Dummy Endogenous Regressors: Simple Strategies for Empirical Practice. Journal of Business & Economic Statistics 19.

Arora, A., Ceccagnoli, M., & Cohen, W. M. 2008. R&D and the Patent Premium. International journal of industrial organization, 26(5).

Audretsch, D. B., Keilbach, M. C., & Lehmann, E. E. 2006. Entrepreneurship and economic growth. Oxford University Press.

Baker, M., Stein, J. C., & Wurgler, J. 2003. When Does the Market Matter? Stock Prices and the Investment of Equity-Dependent Firms. The Quarterly Journal of Economics.

Baker, M., & Wurgler, J. 2011. Behavioral corporate finance: An updated survey (No. w17333). National Bureau of Economic Research.

Barrot, J. N. 2014. Trade Credit and Industry Dynamics: Evidence from Trucking Firms. Working Paper.

Beck, N. 2011. Is OLS with a binary dependent variable really OK?: Estimating (mostly) TSCS models with binary dependent variables and fixed effects. Unpublished working paper, NYU.

Benavente, J. M., Crespi, G., Figal Garone, L., & Maffioli, A. 2012. The impact of national research funds: A regression discontinuity approach to the Chilean FONDECYT. Research Policy, 41(8).

Berkowitz, J. and M. White. 2004. Bankruptcy and small firms' access to credit. RAND Journal of Economics, 35 (1).

Black, S. E., & Strahan, P. E. 2002. Entrepreneurship and bank credit availability. The Journal of Finance, 57(6).

Blasio, G., Fantino, D., & Pellegrini, G. 2011. Evaluating the impact of innovation incentives: evidence from an unexpected shortage of funds. Industrial and Corporate Change, 23(5).

Bond, S., Harhoff, D., & Van Reenen, J. 2005. Investment, R&D and Financial Constraints in Britain and Germany. Annales d'Économie et de Statistique.

Busom, I. 2000. An empirical evaluation of the effects of R&D subsidies, Economics of Innovation and New Technology 9(2).

Brander, J. A., Egan, E., & Hellmann, T. F. 2008. Government sponsored versus private venture capital: Canadian evidence (No. w14029). National Bureau of Economic Research.

Bronzini, R. & Iachini, E. 2011. Are Incentives for R&D Effective? Evidence from a Regression Discontinuity Approach. Working Paper, Banca d'Italia.

Brouwer, E. & Kleinknecht, A. 1999. Innovative output, and a firm's propensity to patent.: An exploration of CIS micro data. Research Policy, 28(6), 615-624.

Brown, J. R. & Petersen, B. C. 2009. Why has the investment-cash flow sensitivity declined so sharply? Rising R&D and equity market developments. Journal of Banking & Finance, 33(5).

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

Bureau of Economic Analaysis (BEA). 2013. GDP by Metropolitan Area. Available at http://www.bea.gov.

Calvo, J. L. 2006. Testing Gibrat's law for small, young and innovating firms. Small Business Economics, 26(2).

Campello, M., Graham, J. R., & Harvey, C. R. 2010. The real effects of financial constraints: Evidence from a financial crisis. Journal of Financial Economics, 97(3).

Card, D., Chetty, R. & Weber, A. 2007. Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market. Quarterly Journal of Economics, November.

Chatterji, A. K., & Seamans, R. C. 2012. Entrepreneurial finance, credit cards, and race. Journal of Financial Economics, 106(1).

Chemmanur, T. J., Krishnan, K., & Nandy, D. K. 2011. How does venture capital financing improve efficiency in private firms? A look beneath the surface. Review of Financial Studies, hhr096.

Chen, H., Gompers, P., Kovner, A., & Lerner, J. 2010. Buy local? The geography of venture capital. Journal of Urban Economics, 67(1).

Cochrane, J. H. 2005. The risk and return of venture capital. Journal of financial economics, 75(1).

Cohen, W., & Klepper, S. 1996. Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D. Review of Economics and Statistics.

Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2000). Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not) (No. w7552). National Bureau of Economic Research.

Conti, A., Thursby, J., & Thursby, M. 2013. Patents as Signals for Startup Financing. The Journal of Industrial Economics, 61(3).

Cumming, D., & Dai, N. 2010. Local bias in venture capital investments. Journal of Empirical Finance, 17(3).

Czarnitzki, D., & Bento, C. L. 2012. Evaluation of public R&D policies: a cross–country comparison. World Review of Science, Technology and Sustainable Development, 9(2).

Czarnitzki, D., & Hottenrott, H. 2011. R&D investment and financing constraints of small and medium-sized firms. Small Business Economics, 36(1).

Duan, N., Manning, W. G. Jr., Morris, C. N., and Newhouse, J. P. 1983. A comparison of alternative models for the demand for medical care (Corr: V2 P413). Journal of Business and Economic Statistics, 1.

Duguet E. 2004. Are R&D subsidies a substitute or a complement to privately funded R&D? Evidence from France using propensity score methods for non experimental data, Revue d'EconomiePolitique 114(2).

Eaton, J., Kortum, S., 1999. International technology diffusion, theory and measurement. International Economic Review 40.

Engel, D, & Keilbach, M. 2007. Firm-Level Implications of Early Stage Venture Capital Investment—An Empirical Investigation. Journal of Empirical Finance 14.

Evans, D. S. 1987. The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries. The Journal of Industrial Economics.

Faulkender, M., & Petersen, M. 2012. Investment and capital constraints: repatriations under the American Jobs Creation Act. Review of Financial Studies, 25(11).

Fazzari, S. M., Hubbard, R. G., Petersen. 1988. Financing Constraints and Corporate Investment. Brookings Papers on Economic Activity.

Feyrer, J., & Sacerdote, B. 2011. Did the stimulus stimulate? Real time estimates of the effects of the American Recovery and Reinvestment Act (No. w16759). National Bureau of Economic Research.

Florida, R. 2014. Startup City: The Urban Shift in Venture Capital and High Technology. Martin Prosperity Institute Report.

Gompers, P. A. 1995. Optimal investment, monitoring, and the staging of venture capital. The journal of finance, 50(5).

Gompers, P. A., & Lerner, J. 2004. *The venture capital cycle.* MIT press.

Gompers, P., Kovner, A., Lerner, J., & Scharfstein, D. 2008. Venture capital investment cycles: The impact of public markets. Journal of Financial Economics, 87(1).

Gompers, P., Ishii, J., & Metrick, A. 2003. Corporate governance and equity prices. The Quarterly Journal of Economics, 118(1), 107-156.

Gompers, P., Lerner, J., & Scharfstein, D. 2005. Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999. The Journal of Finance, 60(2).

Gompers, P. A., & Lerner, J. 1999. Capital formation and investment in venture markets: implications for the Advanced Technology Program. National Institute of Standards and Technology and U.S. Department of Commerce, GCR 99-784.

Gompers, P., A. Kovner, and J. Lerner. 2009. Specialization and Success: Evidence from Venture Capital. Journal of Economics and Management Strategy 18.

González, X., & Pazó, C. 2008. Do public subsidies stimulate private R&D spending?. Research Policy, 37(3).

Griliches, Z. 1998. R&D and Productivity: The Econometric Evidence. University of Chicago Press, Chicago.

Gruber, M., MacMillan, I. C., & Thompson, J. D. 2008. Look before you leap: Market opportunity identification in emerging technology firms. Management Science, 54(9).

Hahn, J., P. Todd, W. van de Klaauw. 2001. Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. Econometrica 69 (1).

Hall, B. H., Griliches, Z., & Hausman, J. A. 1986. Patents and R&D: Is There a Lag?. International Economic Review.

Hall, B. H. 2002. The financing of research and development. Oxford review of economic policy, 18(1),.

Hall, B. H., Jaffe, A., & Trajtenberg, M. 2005. Market value and patent citations. RAND Journal of economics.

Hall, B. H. 2008. The financing of innovation. In S. Shane (ed.), *Handbook of Technology and Innovation Management*. Blackwell Publishers, Ltd: Oxford, pp. 409–430.

Hall, B. H. 2010. The financing of innovative firms. Review of Economics and Institutions, 1(1).

Hall, B. H., & Lerner, J. 2009. The financing of R&D and innovation (No. w15325). National Bureau of Economic Research.

Hall, R. E., & Woodward, S. E. 2007. The Quantitative Economics of Venture Capital. NBER Working Paper 13056.

Haltiwanger, J., Jarmin, R.S., & Miranda, J. 2013. Who creates jobs? small versus large versus young. Review of Economics and Statistics 95.

Hao, K. Y., & Jaffe, A. B. 1993. Effect of liquidity on firms' R&D spending. Economics of Innovation and New technology, 2(4).

Hellman, T., & Puri, M. 2000. The interaction between product market and financing strategy: The role of venture capital. Review of Financial studies,13(4).

Henningsen, M. S., Hægeland, T., & Møen, J. 2014. Estimating the additionality of R&D subsidies using proposal evaluation data to control for research intentions. The Journal of Technology Transfer, 1-25.

Himmelberg, C. P., & Petersen, B. C. 1994. R&D and internal finance: A panel study of small firms in high-tech industries. The Review of Economics and Statistics.

Hochberg, Y. V., Ljungqvist, A., & Lu, Y. 2007. Whom you know matters: Venture capital networks and investment performance. The Journal of Finance, 62(1).

Hochberg, Y. V., Serrano, C. J., & Ziedonis, R. H. 2014. Patent Collateral, Investor Commitment, and the Market for Venture Lending. Investor Commitment, and the Market for Venture Lending.

Holmstrom, B. 1989. Agency costs and innovation. Journal of Economic Behavior & Organization, 12(2).

Hsu, D. H. 2006. Venture capitalists and cooperative start-up commercialization strategy. Management Science, 52(2).

Hsu, D. H. 2008. Technology-based Entrepreneurship. In S. Shane (ed.), *Handbook of Technology and Innovation Management*. Blackwell Publishers, Ltd: Oxford, pp. 367–387.

Hsu, D. H., & Ziedonis, R. H. 2008. Patents as quality signals for entrepreneurial ventures. In Academy of Management Proceedings (Vol. 2008, No. 1, pp. 1-6). Academy of Management.

Imbens, G.W. & Kalyanaraman, K. 2009. Optimal Bandwidth Choice for the Regression Discontinuity Estimator. January.

Imbens, G. W. & Lemieux, T. 2008. Regression Discontinuity Designs: A Guide to Practice. Journal of Econometrics 142 (2).

Jacob, B. A., & Lefgren, L. 2011. The impact of research grant funding on scientific productivity. Journal of Public Economics, 95(9).

Jaffe, A. B. 2002. Building programme evaluation into the design of public research support programmes. Oxford Review of Economic Policy, 18(1).

Jeng, L. A., & Wells, P. C. 2000. The determinants of venture capital funding: evidence across countries. Journal of Corporate Finance, 6(3).

Kaplan, S. N., & Zingales, L. 1997. Do investment-cash flow sensitivities provide useful measures of financing constraints?. The Quarterly Journal of Economics.

Kerr, W. R., Nanda, R., & Rhodes-Kropf, M. 2013. Entrepreneurship as experimentation. Journal of Economic Perspectives, Forthcoming.

Kerr, W. R., & Nanda, R. 2009. Democratizing entry: Banking deregulations, financing constraints, and entrepreneurship. Journal of Financial Economics, 94(1).

Kirsch, D., Goldfarb, B. & Gera, A. 2009. Form or substance: The role of business plans in venture capital decision making. Strategic Mangement Journal 30.

Knight, B. 2002. Endogenous federal grants and crowd-out of state government spending: Theory and evidence from the federal highway aid program. American Economic Review.

Kortum, S., & Lerner, J. 2000. Assessing the contribution of venture capital to innovation. RAND Journal of Economics.

Lach, S. 2002. Do R&D subsidies stimulate or displace private R&D? Evidence from Israel, Journal of Industrial Economics 50(4),0.

Lamont, O. 1997. Cash flow and investment: Evidence from internal capital markets. The Journal of Finance, 52(1).

Lee, D. S. & David Card, D. Regression Discontinuity Inference with Specification Error," Journal of Econometrics, February 2008, 142 (2).

Lee, D.S., & T. Lemieux. 2010. Regression Discontinuity Designs in Economics. Journal of Economic Literature 48.

Leleux, B. & Surlemont, B. 2003. Public versus private venture capital: seeding or crowding out? A pan-European analysis. Journal of Business Venturing, 18.

Lerner, J. 2000. The government as venture capitalist: the long-run impact of the SBIR program. The Journal of Private Equity, 3(2).

Lerner, J. 2002. When bureaucrats meet entrepreneurs: the design of effective public venture capital programmes. The Economic Journal, 112(477).

Lerner, J. 2009. *Boulevard of broken dreams: why public efforts to boost entrepreneurship and venture capital have failed–and what to do about it.* Princeton University Press.

Lerner, J., Sorensen, M., & Strömberg, P. 2011. Private Equity and Long-Run Investment: The Case of Innovation. The Journal of Finance, 66(2).

Li, D. 2011. Financial constraints, R&D investment, and stock returns. Review of Financial Studies.

Link, Albert N., and John T. Scott. "Government as entrepreneur: Evaluating the commercialization success of SBIR projects." Research Policy 39.5 (2010).

Mann, W. 2013. Creditor rights and innovation: Evidence from patent collateral. Working Paper, available at SSRN 2356015.

Mehta, Monica. 2011. Don't undercut your equity stake. Bloomberg Businessweek, August 12.

Metrick, A. 2007. Venture Capital and the Finance of Innovation. John Wiley & Sons: Hoboken, NJ.

Mullahy, J. 1986. Specification and testing of some modified count data models. Journal of econometrics, 33(3).

Nanda, R., & Rhodes-Kropf, M. 2012. Investment Cycles and Startup Innovation (Vol. 105, No. 12-032). Harvard Business School Working Paper.

Nanda, R., Younge, K., & Fleming, L. 2013. Innovation and Entrepreneurship in Renewable Energy. NBER Chapters.

National Venture Capital Association (NVCA). 2014. Yearbook 2014. Prepared by Thomson Reuters.

National Science Foundation (NSF). 2012. Science and Engineering Indicators 2012. Washington, DC.

Nemet, G. F., & Kammen, D. M. 2007. US energy research and development: Declining investment, increasing need, and the feasibility of expansion. Energy Policy, 35(1).

Oliver, M. 2012. Overview of the DOE's Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Programs. DOE Webinar, November 30.

Ouyang, M. 2011. On the Cyclicality of R&D. Review of Economics and Statistics, 93(2).

Ozbas, O., & Scharfstein, D. S. 2009. Evidence on the dark side of internal capital markets. Review of Financial Studies, hhp071.

Pakes, A. 1985. On Patents, R & D, and the Stock Market Rate of Return. The Journal of Political Economy, 390-409.

Phelps, E. S. 1972. The statistical theory of racism and sexism. The American Economic Review.

Popp, D. & Newell, R. G. 2009. Where does energy R&D come from? Examining crowding out from environmentally-friendly R&D (No. w15423). National Bureau of Economic Research.

Puri, M. & Zarutskie, R. 2012. On the Life Cycle Dynamics of Venture Capital and Non-Venture Capital-Financed Firms. The Journal of Finance, 67(6).

Rajan, R.G. and L. Zingales. 1998. Financial dependence and growth. The American Economic Review, 88 (3).

Rauh, J. D. 2006. Investment and financing constraints: Evidence from the funding of corporate pension plans. The Journal of Finance, 61(1).

Sahlman, W. A., & Scherlis, D. R. (1989). A Method for Valuing High-risk, Long-term Investments: The "Venture Capital Method." Revised 2009.

Saxenian, A. 1994. *Regional advantage: culture and competition in Silicon Valley and Route 128.* Harvard University Press.

Scharfstein, D. & Stein, J. C. 2000. Herd behavior and investment: Reply. American Economic Review, 90(3).

Scherer, F. M. 1983. The propensity to patent. International Journal of Industrial Organization, 1(1).

Serrano-Velarde, N. 2008. The Financing Structure of Corporate R&D-Evidence from Regression Discontinuity Design. Working Paper.

Seru, A. 2014. Firm boundaries matter: Evidence from conglomerates and R&D activity. Journal of Financial Economics, 111(2).

Shane, S., & Stuart, T. 2002. Organizational endowments and the performance of university start-ups. Management science, 48(1).

Sørensen, M. 2007. How smart is smart money? A two-sided matching model of venture capital. The Journal of Finance, 62(6).

Sorenson, O. & T. Stuart. 2001. "Syndication Networks and the Spatial Distribution of Venture Capital Investments," American Journal of Sociology 106.

Stein, J.C. 2003. Agency, Information and Corporate Investment" Chapter 2 in Constantinides, G.M et al, eds. Handbook of the Economics of Finance. Elsevier Science.

Takalo, T., Tanayama, T., & Toivanen, O. 2013. Estimating the benefits of targeted R&D subsidies. Review of Economics and Statistics, 95(1).

U.S. Department of Commerce. 2012. Intellectual Property and the U.S. Economy: Industries in Focus. Available at http://www.uspto.gov/news/publications/IP_Report_March_2012.pdf

U.S. General Accountability Office. 1992. Small Business Innovation Research Shows Success but can be Strengthened. Report to Congressional Committees, GAO-92-37.

Wallsten, S. J. 2000. The effects of government-industry R&D programs on private R&D: the case of the Small Business Innovation Research program. RAND Journal of Economics, 31(1),.

Whited, T. M., & Wu, G. 2006. Financial constraints risk. Review of Financial Studies, 19(2).

Zwick, E., & Mahon, J. 2014. Do Financial Frictions Amplify Fiscal Policy? Evidence from Business Investment Stimulus. Working Paper.

References - Chapter 2

Ahn, D., & Kogan, L. 2010. Crude or Refined: Identifying Oil Price Dynamics through the Crack Spread. Working paper.

Alquist, R., & Kilian, L. 2010. What do we learn from the price of crude oil futures? Journal of Applied Econometrics 25.

Anderson, R. & D. Reeb. 2003. Founding family ownership and firm performance: Evidence from the S&P 500. Journal of Finance 58.

Asker, J. et al. 2014. Corporate Investment and Stock Market Listing: A Puzzle? European Corporate Governance Institute (ECGI)-Finance Research Paper Series.

Athey, S. & J. Levin. 2001. Information and Competition in U.S. Forest Service Timber Auctions. Journal of Political Economy 109.

Bajari, P. et al. 2010. Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs. Working Paper.

Bajari, P. & L. Ye. 2003. "Deciding between Competition and Collusion," The Review of Economics and Statistics 85.

Bhardwaj, G., G. Gorton & K. Rouwenhorst. 2014. Fooling Some of the People All of the Time: The Inefficient Performance and Persistence of Commodity Trading Advisors. Review of Financial Studies 27(11).

Bertrand, M., & Schoar, A. 2006. The role of family in family firms. The Journal of Economic Perspectives 73-96.

Black, F. & M. Scholes. 1973. The Pricing of Options and Corporate Liabilities. Journal of Political Economy 81.

Bollerslev, T. et al. 2011. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. Journal of Econometrics 160.

Bond, S. & J. Cummins. 2004. Uncertainty and company investment, an empirical model using data on analysts' profits forecasts, Mimeo, Institute for Fiscal Studies.

Brown, G. W., Crabb, P. R., & Haushalter, D. 2006. Are Firms Successful at Selective Hedging?*. The Journal of Business, 79(6).

Carter, D. et al. 2006. Does hedging affect firm value? Evidence from the US airline industry. Financial Management, 35.

Chalamandaris, G. & A. Tsekrekos. 2009. Predictable dynamics in implied volatility surfaces from OTC currency options. Journal of Banking and Finance 34.

Congressional Budget Office. 2011. Spending and Funding for Highways. Economic and Budget Issue Brief.

Cornaggia, J. 2013. Does risk management matter? Evidence from the US agricultural industry. Journal of Financial Economics, 109(2).

Etula, Erkko. 2013. Broker-Dealer Risk Appetite and Commodity Returns. Journal of Financial Econometrics 0(0).

Faccio, M. et al. 2011. Large shareholder diversification and corporate risk-taking. Review of Financial Studies 24.

Friedman, J. W. 1971. A non-cooperative equilibrium for supergames. The Review of Economic Studies 38.

Froot, K. et al. 1993. Risk managements coordinating corporate investment and financing policies. the Journal of Finance 48.

Géczy, C. et al. 1997. Why firms use currency derivatives. Journal of Finance 52.

Graham, J. R., & Smith, C. W. 1999. Tax incentives to hedge. The Journal of Finance, 54(6).

Graham, J. R., & Rogers, D. A. 2002. Do firms hedge in response to tax incentives?. The Journal of Finance, 57(2).

Guay, W., & Kothari, S. 2003. How much do firms hedge with derivatives?. Journal of Financial Economics 70.

Gupta, S. 2002. Competition and collusion in a government procurement auction market. Atlantic Economic Journal 30.

Haushalter, G. 2000. Financing policy, basis risk, and corporate hedging: Evidence from oil and gas producers. The Journal of Finance, 55.

Hendricks, K. & R. Porter. 1988. An Empirical Study of an Auction with Asymmetric Information. The American Economic Review 78.

Hennessy, C., & Whited, T. 2007. How costly is external financing? Evidence from a structural estimation. The Journal of Finance, 62(4).

Henriques, I. & Sadorsky, P. 2011. The effect of oil price volatility on strategic investment. Energy Economics, 33(1).

Imbens, G. & J. Wooldridge. 2007. Difference-in-differences estimation. NBER Summer 2007, What's New in Econometrics? Lecture Notes 10.

Ishii, R. 2008. Collusion in repeated procurement auction: A study of a paving market in Japan. Osaka University Institute of Social and Economic Research Discussion Paper 710.

Jin, Y., & Jorion, P. 2006. Firm value and hedging: Evidence from US oil and gas producers. The Journal of Finance 61.

Jofre-Bonet, M., & Pesendorfer, M. 2003. Estimation of a dynamic auction game. Econometrica 71.

Kellogg, R. 2010. The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling. NBER Working Paper 1641.

Kosmopoulou, G., & X. Zhou. 2014. Price Adjustment Policies in Procurement Contracting: An Analysis of Bidding Behavior. The Journal of Industrial Economics, 62(1).

Krasnokutskaya, E. 2011. Identification and estimation of auction models with unobserved heterogeneity. The Review of Economic Studies 78.

MacKay, P., & Moeller, S. B. 2007. The value of corporate risk management. The Journal of Finance, 62(3).

Mayers, D. & C. Smith. 1982. On the corporate demand for insurance. Journal of Business 55.

Merton, R. 1995. A functional perspective of financial intermediation. Financial Management, 23-41.

Mian, S. 1996. Evidence on corporate hedging policy. Journal of Financial and Quantitative Analysis 31.

Nance, D. et al. 1993. On the determinants of corporate hedging. The Journal of Finance 48.

Panousi, V. & D. Papanikolaou, 2012. Investment, idiosyncratic risk, and ownership. The Journal of Finance 67.

Pesendorfer, M. 2000. A Study of Collusion in First-Price Auctions. Review of Economic Studies, 67.

Porter, R. & J. Zona. 1993. Detection of Bid Rigging in Procurement Auctions. Journal of Political Economy 101.

Porter, R. 2005. Detecting collusion. Review of Industrial Organization 26.

Rampini, A. et al. 2014. Dynamic risk management. Journal of Financial Economics 111.

Rampini, A. A., & Viswanathan, S. 2010. Collateral, risk management, and the distribution of debt capacity. The Journal of Finance, 65(6).

Rampini, A. A., & Viswanathan, S. 2013. Collateral and capital structure. Journal of Financial Economics, 109(2),.

Saunders, A., & Steffen, S. 2011. The costs of being private: Evidence from the loan market. Review of Financial Studies, 24(12), 4091-4122.

Scharfstein, D, & Sunderam, A. 2013. Concentration in Mortgage Lending, Refinancing Activity and Mortgage Rates. NBER Working Paper 19156.

Schmid, T. et al. 2008. Family firms, agency costs and risk aversion-Empirical evidence from diversification and hedging decisions. Technische Universität München Center for Entrepreneurial and Financial Studies Working Paper No. 2008-13.

Schulze, W. et al. 2001. Agency relationships in family firms: Theory and evidence. Organization science 12.

Shaad, J. 2006. Asphalt Price Indexes Smooth Process for Road Project Bids. Kansas City Business Journal, July 30.

Shleifer, A., & R. Vishny. 1986. Large shareholders and corporate control. The Journal of Political Economy 94.

Skolnik, J. 2011. Price Indexing in Transportation Construction Contracts. Jack Faucett Associates, in Association with Oman Systems.

Stulz, R. 1984. Optimal hedging policies, Journal of Financial and Quantitative Analysis 19.

Stulz, R.. 1996. Rethinking risk management. Journal of applied corporate finance 9.

Tufano, P. 1996. Who manages risk? An empirical examination of risk management practices in the gold mining industry. The Journal of Finance 51.

Vickery, J. 2008. How and why do small firms manage interest rate risk? Journal of Financial Economics 87.

References - Chapter 3

Aitken, Brian J., and Ann E. Harrison, "Do Domestic Firms Benefit from Foreign Direct Investment? Evidence from Venezuela," American Economic Review (June 1999).

Ahrens, Nathaniel. 2013. China's Competitiveness: Myth, Reality, and Lessons for the United States and Japan; Case Study: SAIC Motor Corporation. Center for Strategic and International Studies Report.

An, Feng et al. 2007. Passenger Vehicle Greenhouse Gas and Fuel Economy Standards: A Global Update. International Council on Clean Transportation.

An, F., R. Earley R, and L. Green-Weiskel. 2011. Global Overview on Fuel Efficiency and Motor Vehicle Emission Standards: Policy Options and Perspectives for International Cooperation. Commission on Sustainable Development, The Innovation Center for Energy and Transportation (iCET).

Anderson, G. E. 2012. *Designated Drivers: How China Plans to Dominate the Global Auto Industry.* Singapore: John Wiley & Sons.

Anderson, M. L., & Auffhammer, M. 2014. Pounds that kill: the external costs of vehicle weight. The Review of Economic Studies, 81(2).

Andrews-Speed, P. 2012. *The governance of energy in China: transition to a low-carbon economy.* Palgrave Macmillan.

Associated Press. 2013. China's struggling automakers jump on SUV boom. May 7.

Bajona, C., & Chu, T. 2010. Reforming state owned enterprises in China: Effects of WTO accession. Review of Economic Dynamics, 13(4).

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How Much Should We Trust Differences-In-Differences Estimates? The Quarterly Journal of Economics, 119(1).

Bhattacharya, Abheek. 2014. "Purifying Air for China's Car Makers." Wall Street Journal, March 25.

Canis, B, and M. Wayne. 2013. US-Chinese Motor Vehicle Trade: Overview and Issues. Congressional Research Service Report for Congress.

Chanaron, J. J. 2001. Implementing technological and organisational innovations and management of core competencies: lessons from the automotive industry. International Journal of Automotive Technology and Management, Vol 1.

Chu, W. 2011. How the Chinese government promoted a global automobile industry. Industrial and Corporate Change, Vol 20.

Deng, H. and A. Ma. 2010. Market Structure and Pricing Strategy of China's Automobile Industry. The Journal of Industrial Economics, Vol 58.

Diao, X., Zhang, Y., & Chen, K. Z. 2012. The global recession and China's stimulus package: A general equilibrium assessment of country level impacts. China Economic Review, 23(1).

Donald, S. G., & Lang, K. 2007. Inference with difference-in-differences and other panel data. The Review of Economics and Statistics, 89(2).

Doner, R. F. 1991. *Driving Bargains: Automobile Industrialization and Japanese Firms in Southeast Asia.* Berkeley: University of California Press.

Dunne, Michael. 2012. Chinese Auto Makers: Joint-Venture Junkies. The Wall Street Journal, September 11.

DRC Intranet 国务院发展研究中心信息网. 2013. "中国起草新政为新能源汽车创造条件." Aug 6.

Economist. 2013. Voting with their Wallets: Chinese car buyers overhwlmeingly prefer foreign brands. Special Report: Cars. April 20.

Gallagher, K. S. 2006. China Shifts Gears: Automakers, oil, pollution and development. Cambridge: MIT Press.

Gao, P. 2004. Shaping the future of China's auto industry. McKinsey Quarterly (3).

Giles, J., & Williams, C. L. 2000. Export-led growth: a survey of the empirical literature and some non-causality results. Journal of International Trade & Economic Development, 9(3).

Grieco, J. M. 1984. *Between Dependency and Autonomy: India's Experience with the International Computer Industry.* Berkeley: University of California Press.

Grossman, G. M. and E. Helpman. 1991a. Quality Ladders and Product Cycles. Quarterly Journal of Economics, 106.

Grossman, G. M. and E. Helpman. 1991b. Endogenous Product Cycles. The Economic Journal, Vol 101.

Grossman, G. M. and E. Helpman. 1994. Protection for sale. American Economic Review, Vol 84.

Haddad, M., and Ann E. Harrison, "Are there Positive Spillovers from Direct Foreign Investment?" Journal of Development Economics 42 (1993).

Hale, G. and C. Long. 2011. "Are There Productivity Spillovers from Foreign Direct Investment in China?" Pacific Economic Review, Vol 16.

Hale, G. and C. Long. 2012. *Foreign Direct Investment in China: Winners and Losers.* Singapore: World Scientific.

Haskel, J. E., Pereira, S. C., & Slaughter, M. J. 2007. Does inward foreign direct investment boost the productivity of domestic firms?. The Review of Economics and Statistics, 89(3).

Holmes, T., E. McGrattan, and E. C. Prescott. 2013. "Quid pro quo: Technology capital transfers for market access in China." Federal Reserve Bank of Minneapolis Research Department Staff Report 486.

Holweg, M. et al. 2005. "The Past, Present and Future of China's Automotive Industry: A Value Chain Perspective." UNIDO Global Value Chain Project Working Paper.

Hsieh, C. T., & Song, Z. M. 2015. Grasp the large, let go of the small: the transformation of the state sector in China. National Bureau of Economic Research Working Paper 21006.

Imbens, G. & J. Wooldridge. 2007. Difference-in-differences estimation. NBER Summer 2007, What's New in Econometrics? Lecture Notes 10.

IMF. 2006. The Automobile Industry in Central Europe. Research Note.

Inkpen, A. C. and Crossan, M. M. 1995. Believing Is Seeing: Joint Ventures and Organization Learning*. Journal of Management Studies, Vol 32.

Ito, K., & Sallee, J. M. 2013. The Economics of Attribute-Based Regulation: Theory and Evidence from Fuel-Economy Standards. draft paper, Boston University, Boston, November.

Jacobsen, M. R. 2013. Fuel Economy and Safety: The Influences of Vehicle Class and Driver Behavior. American Economic Journal: Applied Economics, 5(3).

Jefferson, G., Albert, G. Z., Xiaojing, G., & Xiaoyun, Y. (2003). Ownership, performance, and innovation in China's large-and medium-size industrial enterprise sector. China economic review, 14(1).

Khandelwal, A. K., Schott, P. K., and Wei, S. J. 2011. Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters (No. w17524). National Bureau of Economic Research.

Knittel, C. R. 2011. Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector. American Economic Review, Vol 101.

Lin, J. Y., Cai, F., & Li, Z. 1998. Competition, policy burdens, and state-owned enterprise reform. American Economic Review.

Lyles, M. A., & Salk, J. E. 1996. Knowledge acquisition from foreign parents in international joint ventures: An empirical examination in the Hungarian context. Journal of International Business Studies, Vol 29.

Lucas, Robert E. 1993. Making a Miracle. Econometrica, Vol 61, No 2.

Malerba, F. 1992. Learning by firms and incremental technical change. The Economic Journal, Vol 102.

Mathews, J. A. 2002. Competitive advantages of the latecomer firm: A resource-based 32 account of industrial catch-up strategies. Asia Pacific Journal of Management, Vol 19.

McGrattan, E. and E.C. Prescott. 2009. Openness, Technology Capital, and Development. Journal of Economic Theory, Vol 144.

McGrattan, E. and E.C. Prescott. 2010. Technology Capital and the U.S. Current Account. American Economic Review, Vol 100.

McKinsey & Company. 2012. The Future of the North American Automotive Supplier Industry: Evolution of Component Costs, Penetration, and Value Creation Potential Through 2020. Report.

Medhi, N. 2006. Patent Tales: trailing emission control technologies in the developing world. Department of Social Sciences, Center for Policy Research, The Maxwell School, Syracuse University.

Melitz, M. J. 2003. The impact of trade on intraindustry reallocations and aggregate industry productivity. Econometrica, 71(6).

Melitz, M. 2005. When and how should infant industries be protected? Journal of International Economics, Vol 66.

Melitz, Mark J. and Stephen J. Redding. 2014. Heterogeneous Firms and Trade. Chapter 1 in Gopinath, G., Helpman, E., & Rogoff, K. (Eds.). *Handbook of international economics (Vol. 4).* Elsevier.

Moran, T. H. 1998. *Foreign Direct Investment and Development: The New Policy Agenda for Developing Countries and Economies-in-transition.* Washington DC: Peterson Institute for International Economics.

Morris, D., Donnelly, T., and Donnelly, T. 2004. Supplier parks in the automotive industry. Supply Chain Management: An International Journal, Vol 9.

National Development and Reform Commission 国家发展和改革委员会. 2004. 汽车产业发展政策. May 21. Available at: http://www.sdpc.gov.cn/zcfb/zcfbl/zcfbl2004/t20050614_7501.htm

Nelson, R., Phelps, E., 1966. Investment in humans, technological diffusion, and economic growth. American Economic Review: Papers and Proceedings 61.

Nunn, N. and Trefler, D. 2010. The structure of tariffs and long-term growth. American Economic Journal: Macroeconomics, Vol 2.

Oliver, Hongyan H. et al. 2009. "China's Fuel Economy Standards for Passenger Vehicles: Rationale, Policy Process, and Impacts." Energy Policy, Vol 37.

Oliver Wyman. 2013. "Automotive manager: Trends, opportunities and solutions along the entire value chan." January.

Osborn, Andrew. 2007. Crash course in quality for Chinese car. The Wall Street Journal, August 8.

Parente, S. L., & Prescott, E. C. 1994. Barriers to technology adoption and development. Journal of political Economy.

Richet, X., & Ruet, J. 2008. The Chinese and Indian Automobile Industry in Perspective: Technology Appropriation, Catching-up and Development. Transition Studies Review, 15(3).

Roland Berger. 2013. Chinese vehicles in Europe: Myth or reality? Issue Paper, Brussels, June.

Sanford C. Bernstein. 2013. Chinese Autos, Part 1: The Quest for Global Competitiveness—Technology, Competence, Ambition and Politics"; and "Part 2: Can China Build a Competitive Car? A Unique Teardown Analysis. February.

Shen, Samuel and Kazunori Takada. 2014. Global auto component makers gear up for China's tougher emission rules. Reuters, June 8.

Shirouzu, Norihiko. 2012. China's car makers cut corners to success. Reuters, September 18.

State Council. 1994. 中华人民共和国国务院. "汽车工业产业政策(1994年). Available at http://news.xinhuanet.com/auto/2004-06/02/content_1503431.htm.

State Council 中华人民共和国国务院. 2006. "国家中长期科学和技术发展规划纲要." February 9.

State Council 中华人民共和国国务院. 2009. "汽车产业调整和振兴规划." March 20. Available at http://www.gov.cn/zwgk/2009-03/20/content_1264324.htm#.

State Council 中华人民共和国国务院. 2012. "节能与新能源汽车产业发展规划（2012—2020年)". ("New Energy Vehicle Plan for 2012-2020"). June 28.

Takada, Kazunori. 2013. Going local: Japanese carmakers turn to Chinese parts for China market. Reuters, April 18.

Tang, Rachel. 2012. "China's Auto Sector Development and Policies: Issues and Implications." Congressional Research Service Report for Congress.

Thun, Eric. 2004. "Industrial Policy, Chinese-Style: FDI, Regulation and Dreams of National Champions in the Auto Sector," Journal of East Asian Studies, 4.

United Nations Environment Program. 2010. "The Chinese Automotive Fuel Economy Policy." UNEP Autotool Fuel Economy Case Studies. Available at http://www.unep.org/transport/gfei/autotool/case

United Nations Conference on Trade and Development. 2012. World Investment Report.

Wagner, David Vance et al. 2009. Structure and impacts of fuel economy standards for passenger cars in China. Energy Policy 37.

Wakelin, K. 1998. Innovation and export behaviour at the firm level. Research policy, 26(7).

Walsh, K. A. 1999. U.S. Commercial Technology Transfers to the People's Republic of China. Report to the Office of Strategic Industries and Economic Security, Bureau of Export Administration, www.bis.doc.gov.

Wang, Arthur et al. 2013. Bigger, better, broader: A perspective on China's auto market in 2020. McKinsey & Co Automotive and Assembly Practice, November.

Yang Jian, 2009. Chinese Car Companies Resort to Buying Brands Rather Than Creating Them. Automotive News. July 15.

Yang Jian, 2008. As suppliers source in China, impact of trade ruling declines. Automotive News. February 15.

Ying, Tian. 2012. China's Auto Joint Ventures Failing to Build Local Brands. Bloomberg News, August 22.

Young, A. 1991. Learning by doing and the dynamic effects of international trade. The Quarterly Journal of Economics, 106(2).