



Essays in Industrial Organization and Finance

Citation

Covert, Thomas Rutford. 2014. Essays in Industrial Organization and Finance. Doctoral dissertation, Harvard University.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:12274570>

Terms of Use

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

Share Your Story

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

[Accessibility](#)

Essays in Industrial Organization and Finance

A dissertation presented

by

Thomas Rutherford Covert

to

The Committee on Degrees in Business Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Business Economics

Harvard University

Cambridge, Massachusetts

April 2014

© 2014 Thomas Rutherford Covert

All rights reserved.

Dissertation Advisors:
Professor Ariel Pakes
Professor Bharat Anand

Author:
Thomas Rutherford Covert

Essays in Industrial Organization and Finance

Abstract

This dissertation consists of two essays on the behavior of traders in opaque financial markets and one on the behavior of firms while they are learning to use a new technology.

The first essay describes the market for borrowing corporate bonds using a comprehensive dataset from a major lender. The cost of borrowing corporate bonds is comparable to the cost of borrowing stock, between 10 and 20 basis points, and both have fallen over time. Factors that influence borrowing costs are loan size, percentage of inventory lent, rating, and borrower identity. There is no evidence that bond short sellers have private information. Bonds with CDS contracts are more actively lent than those without. Finally, the 2007 Credit Crunch does not affect average borrowing costs or loan volume, but does increase borrowing cost variance.

The second essay studies how mandatory transparency affects trading in the corporate bond market. In July 2002, TRACE began requiring the public dissemination of post-trade price and volume information for corporate bonds. Dissemination took place in Phases, with actively traded, investment grade bonds becoming transparent before thinly traded, high-yield bonds. Using new data and a differences-in-differences research design, this essay shows that transparency causes a significant decrease in price dispersion for all bonds and a significant decrease in trading activity for some categories of bonds. The largest decrease in daily price standard deviation, 24.7%, and the largest decrease in trading activity, 41.3%, occurs for bonds in the final Phase, which consisted primarily of high-yield bonds. These results indicate that mandated transparency may help some investors and dealers through a decline in price dispersion, while harming others through a reduction in trading activity.

The third essay examines firms' learning behavior using data on their operational choices, profits, and information sets. I study companies using hydraulic fracturing in North Dakota's Bakken Shale formation, where firms must learn the relationship between fracking input use and oil production. Using a new dataset that covers every well since the introduction of fracking to this formation, I find that firms made more profitable input choices over time, but did so slowly and incompletely, only capturing 68% of possible profits from fracking at the end of 2011. To understand what factors may have limited learning, I estimate a model of fracking input use in the presence of technology uncertainty. Firms are more likely to make fracking input choices with higher expected profits and lower standard deviation of profits, consistent with passive learning but not active experimentation. Most firms over-weight their own information relative to observable information generated by others. These results suggest the existence of economically important frictions in the learning process.

Contents

Abstract	iii
Acknowledgments	xi
1 The Market for Borrowing Corporate Bonds	1
1.1 Introduction	1
1.2 Related Literature	4
1.3 Shorting a Corporate Bond: Rationales, Mechanics, and Market Size	6
1.3.1 Rationales for Shorting Corporate Bonds	6
1.3.2 Mechanics of Shorting Bonds	7
1.3.3 Size of the Bond Loan Market	8
1.4 Description of Sample	10
1.4.1 Sample Characteristics	13
1.4.2 Properties of Short Positions	18
1.5 Costs of Borrowing Corporate Bonds	20
1.5.1 Time Series and Cross-Section of Borrowing Costs	22
1.5.2 Determinants of Borrowing Costs	25
1.5.3 Regression Analysis of Borrowing Costs	31
1.5.4 Borrowing Costs Around Credit Events	38
1.6 Relationship Between Bond and Stock Shorting	42
1.6.1 Matching Bonds and Stocks	42
1.6.2 Comparison of Bond and Stock Borrowing Costs	44
1.7 Returns to Shorting Bonds	50
1.7.1 Returns to Portfolios of Shorted Bonds	52
1.7.2 Profitability to Short Sellers of Corporate Bonds	55
1.8 Relationship Between the Market for Shorting Bonds and the CDS Market	57
1.9 The 2007 Credit Crunch	59
1.10 Conclusion and Implications	62
2 The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market	65
2.1 Introduction	65

2.2	TRACE and the Corporate Bond Market	71
2.2.1	History and Implementation of TRACE	71
2.2.2	Related Literature	76
2.3	Data and Descriptive Statistics	79
2.3.1	Historical TRACE Data and Phase Identification	79
2.3.2	Bond Characteristics	81
2.3.3	Measuring Trading Activity and Price Dispersion	83
2.4	Research Design and Main Results	91
2.4.1	Difference in Differences Framework	91
2.4.2	Estimates	96
2.5	Timing, Robustness, and Other Measures of Trading Activity and Price Dispersion	98
2.5.1	Event Study and Time Windows	98
2.5.2	Time Trends	100
2.5.3	Control Groups	103
2.5.4	Alternative Measures of Trading Activity and Price Dispersion	106
2.6	Heterogeneity in Credit Rating and Issue Size	110
2.7	NAIC	115
2.7.1	Trading Activity and Price Dispersion Using NAIC data	118
2.7.2	Dealer Trading Activity	120
2.8	Conclusions and Implications	121

3 Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale **125**

3.1	Introduction	125
3.1.1	Related literature	130
3.2	Institutional Background and Data	132
3.2.1	Fracking and US Oil Production	132
3.2.2	The Bakken Shale and Hydraulic Fracturing	133
3.2.3	The Information Environment in North Dakota	136
3.2.4	Data	137
3.2.5	Drilling and Fracking Costs	141
3.3	The Fracking Production Function	148
3.3.1	The Time Series of Oil Production	149
3.3.2	The Cross Section of Oil Production	150
3.3.3	Likelihood	155
3.3.4	Production Function Estimates	155
3.4	Evidence for Learning	161
3.4.1	<i>ex post</i> Comparisons	162

3.4.2	Profitability vs. Productivity	167
3.4.3	<i>ex ante</i> Comparisons	169
3.5	Fracking input choice model	172
3.5.1	Preferences Over Uncertainty	172
3.5.2	Own-data bias	179
3.6	Conclusion	180
References		182
Appendix A Appendix to Chapter 2		190
A.1	TRACE Data Appendix	190
A.2	NAIC Data Appendix	199
A.2.1	Market Share Analysis	205
Appendix B Appendix to Chapter 3		209
B.1	Likelihood Calculation	209
B.1.1	Step 1	209
B.1.2	Step 2	209
B.2	Expected Present Discounted Value of Oil Production	211
B.3	Weighted Gaussian Process Estimates	213
B.4	Geology Covariates	214
B.4.1	Available Data	215
B.4.2	Explanatory Power	218
B.5	Stability of the Production Function Relationship	219

List of Tables

1.1	Number and Par Value of Bonds in Corporate Bond Databases	12
1.2	Non-Rating Characteristics of Bonds in the Corporate Bond Databases	16
1.3	Rating Characteristics of Bonds in the Corporate Bond Databases	17
1.4	Loan Size, Loan Duration, and Changes in Loan Size	19
1.5	Distribution of New Loan Borrowing Costs	24
1.6	Corporate Bonds with the Highest Borrowing Costs	29
1.7	Regression Analysis of Determinants of Borrowing Costs	34
1.8	Competitive Races Between Brokers	37
1.9	Bond and Stock Borrowing Relationship	48
1.10	Monthly Returns to Long Bond Portfolio Positions	53
2.1	Timeline of Major TRACE Regulatory Changes	74
2.2	Bond Characteristics by Phase	82
2.3	Trading Activity and Price Dispersion for the 90-Day Window Around Phase Start	86
2.4	Difference-in-Difference Estimates of Transparency on Trading Activity and Price Dispersion	97
2.5	Alternative Specifications for Difference-in-Difference Regressions of Trading Activity and Price Dispersion	104
2.6	Difference-in-Difference Estimates for Alternative Measures of Trading Activity	107
2.7	Difference-in-Difference Estimates for Alternative Measure of Price Dispersion .	111
2.8	Difference-in-Difference Estimates by Credit Rating and Issue Size	113
2.9	NAIC Trading Activity and Price Dispersion for the 90-Day Window Around Phase Start	119
2.10	Estimates of Transparency Effects on NAIC Dealer Market Share for 12 Week Window Around Phase Start	122
3.1	Summary Statistics	139
3.2	Summary Statistics by Year	140
3.3	Average First Year’s Oil Production per Foot of Horizontal Length by Quintiles of Sand and Water Use	141

3.4	Spatial Heterogeneity in the Relationship Between Sand, Water and Oil Production	142
3.5	Wells Completed by the 8 Most Active Firms, by Location, Time and Well Characteristics	148
3.6	Production Function Model Estimates	157
3.7	Production Function Model Cross-Validation Statistics	160
3.8	Uncertainty Preference Model Estimates	176
3.9	Uncertainty Preference Model Estimates, With Interaction	178
3.10	Weighted Uncertainty Preference Model Estimates	180
A.1	Steps from Historical TRACE File to Cleaned Historical TRACE Sample	192
A.2	Steps from FINRA's Phase Listings to Cleaned Phase Sample	198
A.3	Steps from Historical NAIC File to Cleaned Historical Sample	201
A.4	Comparison of NAIC and TRACE Trading Activity 90 Days After Phase Start	204
B.1	Geology Covariates Summary Statistics	216
B.2	Explanatory Power of Geology Covariates	220
B.3	Stability of Production Function Estimates Over Time	222

List of Figures

1.1	Number and Par Value of Outstanding Loans	14
1.2	Equally-Weighted Monthly Distribution of Loan Borrowing Costs	23
1.3	Relationship Between Borrowing Cost and Percent of Inventory On Loan	27
1.4	Borrowing Costs Around Bankruptcies	40
1.5	Borrowing Costs Around Credit Events	41
1.6	Equally-Weighted Monthly Distribution of Stock Loan Borrowing Costs	45
1.7	Bond and Stock Borrowing Cost Differences	47
2.1	Weekly Trading Volume by Phase	85
2.2	Weekly Trading Volume/Issue Size by Phase	88
2.3	Weekly Price Standard Deviation by Phase	92
2.4	Event Study for Weekly Volume/Issue Size	101
2.5	Event Study for Weekly Price Standard Deviation	102
3.1	Diagram of a Hydraulically Fractured Bakken Shale well	134
3.2	Diagram of a hypothetical spacing unit	137
3.3	Quarterly Average Cushing Oil Prices	142
3.4	Fracking Cost Index	144
3.5	Fixed and Variable Costs of Drilling and Fracking	145
3.6	Variable Costs of Using Sand and Water in Fracking	146
3.7	Contour Plots of Production Function Estimates	159
3.8	Fraction of Positive Profits Captured and Maximal Profits by Year, ex post	163
3.9	Average Profit Maximizing Sand Use and Actual Sand Use Per Well, ex post	164
3.10	Average Profit Maximizing Water Use and Actual Water Use Per Well, ex post	165
3.11	Gaussian Process Year Effects	166
3.12	Fraction of Positive Profits Captured and Maximal Profits by Year, ex ante	168
3.13	Average Profit Maximizing Sand Use and Actual Sand Use Per Well, ex ante	170
3.14	Average Profit Maximizing Water Use and Actual Water Use Per Well, ex ante	171

Acknowledgments

I am deeply grateful for the teaching, guidance and encouragement I received from my committee members: Ariel Pakes, Bharat Anand, Parag Pathak and Greg Lewis. Though he was not on my committee, Paul Asquith has been a terrific mentor and he is the reason why I am an economist today. I received additional valuable feedback and advice from Michael Luca, Ben Edelman, Julie Mortimer, and Dennis Yao.

My graduate school colleagues have contributed enormously to my thought process, presentation skills and well being. Thanks especially to Thomas G. Wollman, Richard Sweeney, Bryce Millet-Steinberg, Duncan Gilchrist, Ben Lockwood, Stephanie Hurder, Alex Peysakhovich, Evan Herrnstat, and Joseph Shapiro. The third chapter of this dissertation would not exist without the numerous conversations about the oil and gas industry I had with Hugh Daigle and Heath Flowers. Funding from the Harvard Business School Doctoral Programs, a Harvard University Dissertation Completion Fellowship and the Sandra Ohrn Family Foundation is gratefully acknowledged.

Finally, I thank my family for their love and support. This dissertation is dedicated to my wife Riva, my parents Steven and Diane, and my sister Rachel.

Chapter 1

The Market for Borrowing Corporate Bonds¹

1.1 Introduction

This paper analyzes the market for borrowing and shorting corporate bonds. The corporate bond market is one of the largest over-the-counter (OTC) financial markets in the world. Between 2004 and 2007, the time period of our study, the value of outstanding corporate debt averaged \$6.6 trillion and, according to the Securities Industry and Financial Market Association (SIFMA), trading activity averaged \$17.3 billion per day. We estimate that shorting represents 19.1% of all corporate bond trades.

There is a large theoretical literature on short sales constraints and their impact on asset prices. Constraints on short selling may lead to mis-valuation because they limit the ability of some market participants to influence prices. The empirical literature on short sales, while also large, has focused almost exclusively on stocks. Our analysis of shorting corporate bonds allows us to determine if the empirical findings on shorting stocks are present in other markets. In addition, unlike stocks, where borrowing takes place in an OTC market and short selling takes place on an exchange, both borrowing and shorting activities take place OTC in the

¹with Paul Asquith, Andrea S. Au and Parag A. Pathak

corporate bond market. Thus, any effects of short sale constraints may be amplified in the bond market.

A major issue in the study of any OTC market is the availability of data. Unlike stock short positions, which are reported bimonthly by the stock exchanges, bond shorting is not regularly reported. In addition, while a number of studies have access to proprietary databases of stock lending for short periods (e.g., D'avolio (2002); Geczy *et al.* (2002)), comparable analyses of bond lending do not exist, with the exception of Nashikkar and Pedersen (2007).

This paper uses a large proprietary database of corporate bond loan transactions from a major depository institution for the four year period, January 1, 2004 through December 31, 2007. Although our data is only from one lender, the size and coverage of our database allows us to study the functioning of a relatively opaque, yet large market. Our lender's par value of loanable bond inventory averages \$193.3 billion daily and accounts for 2.9% of the overall par value of outstanding corporate bonds listed by the Fixed Income Securities Database (FISD). From this inventory, our lender loans an average daily par value of \$14.3 billion and 66.4% of bonds which appear in inventory are lent out at some point during our time period 2004-2007.

This paper describes the market for shorting corporate bonds along several dimensions. In Section 1.3 we discuss why and how corporate bonds are shorted and estimate the market's size. After describing our data sample in Section 1.4, we examine cross-sectional and time-series determinants of borrowing activity and costs in Section 1.5. In Section 1.6, we investigate the relationship between bond and stock shorting. In Section 1.7, we check if bond short sellers have private information. The next two sections consider how corporate bond shorting relates to the CDS market and whether it was impacted by the Credit Crunch of 2007. Finally, Section 1.10 outlines some implications of our results and concludes.

Since this paper is largely descriptive, after reading Sections 1.1-1.5, readers interested only in one of these later topics can skip directly to the relevant section. A brief summary of our results follows.

We find that the market for borrowing bonds is large and most lent bonds have small borrowing costs. In our database, the mean and median annual borrowing cost, equally-

weighted by loan, are 33 and 18 basis points (bps), respectively, for the entire sample period. In mid-2006, there is a dramatic narrowing in the distribution of bond borrowing costs. This compression causes a reduction in mean and median borrowing costs during the latter part of our sample period. By 2007, these rates fall to 19 and 13 bps, respectively. Borrowing costs are related to several factors. Four significant factors are loan size, the bond's credit rating, on-loan percentage, which is the fraction of the lender's inventory already lent, and the identity of the borrowing broker. Smaller loans (less than 100 bonds) and lower rated bonds have higher borrowing costs. In addition, borrowing costs increase after ratings downgrades and bankruptcy filings. Borrowing costs remain flat until on-loan percentage reaches approximately 70% and then rise sharply for high yield bonds. Finally, while our lender lends to 65 brokers, a select few borrow at significantly lower rates.

Borrowing costs for corporate bonds and stocks are linked. Since our lender has a significant market share of stock shorting, we construct a matched sample of corporate bond and stock loans for the same firms. The costs of borrowing the two securities are usually quite close and 63.7% of matched loan borrowing costs are within 10 bps of each other. When the borrowing costs of matched loans are not close, the stock is usually more expensive to borrow than the bond.

Bond shorting does not appear to be motivated by investors possessing private information since bond short sellers do not earn excess returns on average. Portfolios of bonds with a high on loan percentage or with high borrowing costs do not underperform the market portfolio of corporate bonds. In addition, mimicking the actual positions of bond short sellers (using the beginning and ending dates of bond loans) does not generate excess returns.

We examine two other aspects of the market for borrowing corporate bonds. The first is whether credit default swap (CDS) contracts impact bond borrowing activity. Almost half of our borrowed bonds also have CDS contracts available. These bonds are more actively borrowed and have higher average higher borrowing costs (1 bp) than those where CDS contracts are not available.

The second aspect is the Credit Crunch of 2007. We examine the second half of 2007, the

beginning of the Credit Crunch, separately to see if borrowing activity changes. In this period, borrowing costs became more volatile. However, the volume of bond shorting remained stable, as did the average level of borrowing costs. In addition, the average returns to shorting bonds did not change.

1.2 Related Literature

The theoretical literature on the effects of short sale constraints on asset prices is extensive. One modeling approach examines the implications of heterogeneous investor beliefs in the presence of short sale constraints and whether this causes mis-valuation. Miller (1977) argues that short sale constraints keep more pessimistic investors from participating in the market, so market prices reflect only optimists' valuations (see also Lintner 1971). Harrison and Kreps (1978) consider a dynamic environment and provide conditions where short sale constraints can drive the price above the valuation of even the most optimistic investor. More recent contributions include Chen *et al.* (2002) who relate differences of opinion between optimists and pessimists to measures of stock ownership, and Fostel and Geanakoplos (2008), who consider the additional effects of collateral constraints.

Another approach to studying the effects of short sale constraints focuses on search and bargaining frictions, which arise because investors must first locate securities to short (Duffie 1996; Duffie *et al.* 2002). Finally, there is theoretical literature in the rational expectations tradition, which examines how short sale constraints can impede the informativeness of prices (see Diamond and Verrecchia 1987; Bai *et al.* 2006).

The empirical literature on short sale constraints focuses almost entirely on stocks. An early strand of this literature examines the information content of short interest (see Asquith and Meulbroek 1995) where short interest is the number of shares shorted divided by the number of shares outstanding. This literature advanced in two directions as richer data sets became available. The first direction examines daily quantities of short sales by observing transactions either from proprietary order data (Boehmer *et al.* 2008) or from Regulation SHO data (Diether *et al.* 2009). Both papers find that short sellers possess private information and

that trading strategies based on observing their trades generate abnormal returns. The second direction in this literature examines the direct cost (or price) of borrowing stocks. These papers either use data from a unique time period when the market for borrowing stocks was public (Jones and Lamont 2002) or proprietary data from stock lenders (D’avolio 2002; Geczy *et al.* 2002; and Ofek *et al.* 2004). Jones and Lamont (2002) and Ofek *et al.* (2004) find that stocks with abnormally high rebate rates have lower subsequent returns, while Geczy *et al.* (2002) find that higher borrowing costs do not eliminate abnormal returns from various short selling strategies. D’avolio (2002) and the other three papers find that only a small number of stocks are expensive to borrow. Using data from 12 lenders, Kolasinski *et al.* (2013) find that the equity loan market is opaque, and this, in combination with search costs, results in borrowing costs varying across lenders.

A challenge identified in this literature is that short interest is a quantity and borrowing costs are a price, both of which are simultaneously determined by shorting demand and the supply of shares available to short. A high borrowing cost may indicate either a high shorting demand or a limited supply of shares available to short. As a result, some researchers have constructed proxies for demand and supply and have tried to isolate shifts in either demand or supply. Asquith *et al.* (2005) use institutional ownership as a proxy for the supply of shares available for shorting and find that stocks that have high short interest and low levels of institutional ownership significantly underperform the market on an equally-weighted basis, but not on a value-weighted basis. Using richer, proprietary loan-level data, Cohen *et al.* (2007) examine shifts in the demand for shorting, and find that an increase in shorting demand indicates negative abnormal returns for the stocks being shorted. Both papers highlight that their results only apply to a small fraction of outstanding stocks.

The only paper on corporate bond market shorting is Nashikkar and Pedersen (2007), who describe a proprietary dataset from a corporate bond lender between September 2005 and June 2006. Their examination of the cross-sectional determinants of borrowing costs complements ours, but they do not examine as extensively the differences between investment grade and high yield bonds, the relationship between bond and stock shorting, and the profitability of

short selling corporate bonds. Furthermore, our longer time period allows us to document several time-series patterns, such as the reduction in borrowing costs and the increase in volatility of borrowing costs during the 2007 Credit Crunch.

1.3 Shorting a Corporate Bond: Rationales, Mechanics, and Market Size

1.3.1 Rationales for Shorting Corporate Bonds

The primary purpose of borrowing a corporate bond is to facilitate a short sale of that bond. Aside from market making activities, investors short bonds for the same reason they short stocks: to bet that the security will decline in price. If short sellers focus on overvalued firms and can either short the stock or the bond, it would seem that they would target the stock due to the priority of claims. That is, since bond holders have a higher priority in bankruptcy, stock prices should decline before bond prices when there is a threat of financial distress. Thus, on first pass, short sellers short bonds only if they cannot find the stock to short or it is too expensive to short.

One potential reason why a firm's stock cannot be borrowed is that the firm is private, yet has publicly traded debt. That is, there is public debt but no public stock. In this case, taking a position that the firm is overvalued requires an investor to short bonds. We show below that we are unable to match our corporate bonds to publicly traded stock for 18.4% of our sample.²

If a stock is publicly traded and the stock and bond markets are linked, bond shorting is attractive if the net return for shorting bonds is greater than the net return for shorting stocks, adjusting for risk. We expect this to occur more frequently for lower credit quality bonds. This is because bonds without default risk trade at par (absent interest rate movements) while lower rated bonds will experience greater price fluctuations. Thus, investment grade bonds should not decrease in price as often as high yield bonds, and therefore, the market for shorting

²This does not mean that 18.4% of our bonds were issued by private firms, however. We discuss this further below.

high yield bonds should be different than the market for shorting investment grade bonds.

If the stock and bond markets are not linked, bonds may be shorted due to segmentation. One possibility is that bond short sellers are separate from stock short sellers and evaluate the firm's prospects independently. For instance, practitioners have told us that within an investment firm, the bond and stock trading desks may not trade in each other's instruments. Hence, the bond desk may short the bonds, while the stock desk shorts the stocks.

There are also reasons for shorting bonds that are not related to the value of the firm's stock. If there is a capital structure arbitrage, investors may go long one tier of the firm's capital structure and short another. Arbitrage is also possible between a firm's bonds and their CDS (or other securities reflecting the firm's credit).

Arbitrage trades involving bond shorting are not necessarily specific to an individual bond issue. Two examples are credit spread arbitrage (between different yield curves) and market-wide interest rate arbitrage. In the first case, if investors believe that yield curves are mispriced in relation to one another, they will short one credit category of bonds, and go long another. In this instance, it is not important which firm issued the bond, only the bond's credit rating. In the second case, we expect that investors who believe interest rates will rise prefer to short government bonds rather than corporate bonds because of their low credit risk. However, bond traders have told us that AAA-rated debt is occasionally used for this purpose because it is sometimes cheaper to borrow than treasuries. Here, it is not important which firm issued the AAA-rated debt.

Finally, corporate bonds may be borrowed short term to facilitate clearing of long trades in the presence of temporary frictions in the delivery process.

1.3.2 Mechanics of Shorting Bonds

The mechanics of shorting corporate bonds parallel those of shorting stocks. Shorted bonds must first be located and then borrowed. The investor has three days to locate the bonds after placing a short order. Investors usually borrow bonds through an intermediary such as a depository bank. Such banks serve as custodians for financial securities and pay depositors a

fee in exchange for the right to lend out securities. The borrower must post collateral of 102% of the market value of the borrowed bond, which is re-valued each day. Loans are typically collateralized with cash although US Treasuries may also be used. In our sample, 99.6% of bond loans are collateralized by cash. Investors subject to Federal Reserve Regulation T must post an additional 50% in margin, a requirement that can be satisfied with any security. The loan is “on-demand” meaning that the lender of the security may recall it at any time. Hence, most loans are effectively rolled over each night, and there is very little term lending.

The rebate rate determines the fee that the borrower pays for the bond loan. The rebate rate is the interest rate that is returned by the lender of the security for the use of the collateral. For example, if the parties agree to a bond loan fee of 20 bps, and the current market rate for collateral is 100 bps, then the lender of the corporate bond returns, or “rebates”, 80 bps back to the borrower undertaking the short position. There can be variability in the rebate rate for the same bond even on the same day. It is even possible that the rebate rate is negative, which means the borrower receives no rebate on their collateral and has to pay the lender. Finally, if a bond makes coupon payments or has other distributions, the borrower is responsible for making these payments back to the owner of the security.

1.3.3 Size of the Bond Loan Market

There is limited information about the size of the markets for shorting any security. For stocks, all three major stock exchanges release short interest statistics bimonthly.³ Short interest is the number of shares shorted at a particular point in time divided by the total shares outstanding and is often represented as a percentage. In addition, daily stock shorting information is available from January 2005 through July 2007 when Regulation SHO was in effect. Regulation SHO required all exchanges to mark stock trades as long or short. This is no longer the case.

To estimate the size of the market for shorting stocks, most researchers first examine stock short interest statistics released by the exchanges. Asquith *et al.* (2005) report that in 2002

³Prior to September 2007, all three exchanges reported short interest once a month.

the equally-weighted average short interest for stocks is approximately 2.4% for the NYSE and AMEX combined, and 2.5% for the NASDAQ-NMS. That is, 2.4% or 2.5% of the total number of shares are lent out on average. To examine short sales as a percentage of trading volume, Diether *et al.* (2009) use Regulation SHO data and find that short sales represent 31% of share volume for NASDAQ-listed stocks and 24% of share volume for NYSE-listed stocks in 2005. Asquith *et al.* (2006) report that short sales represent 29.8% of all stock trades on the NYSE, AMEX, and NASDAQ-NMS exchanges during the entire SHO period.⁴ Since bonds primarily trade OTC, comparable information on short interest does not exist and Regulation SHO did not apply.

To estimate the size of the market for shorting corporate bonds, we assume that our proprietary lender's share of the bond shorting market is identical to their share of the stock shorting market. Asquith *et al.* (2006) report that our proprietary lender made stock loans totaling 16.7% of all stock shorting volume on the NYSE, AMEX, and NASDAQ-NMS markets during the SHO period. From Table 1.1, discussed below, the average daily par value of the bonds on loan by our proprietary lender is \$14.3 billion. This measure is comparable to short interest, i.e. it is the daily average par value of bonds shorted over our sample period. If we assume that our lender represents 16.7% of the bonds lent, then total bonds lent for the entire market on an average day is \$85.6 billion. This is 1.3% of the average par value of corporate bonds outstanding as reported from the FISD database discussed below. Thus, by this measure, bond shorting is approximately half as large as stock shorting.

The average daily new corporate bond loan volume of our proprietary lender is \$550.3 million. If we again assume our proprietary lender is responsible for the same proportion of loans to bond short sellers as they are to stock short sellers, this implies that the average daily par value of corporate bonds shorted is \$3.3 billion. SIFMA reports that the average daily corporate bond trading volume for the years 2004-2007 is \$17.3 billion. By this measure, bond short selling would represent 19.1% of all corporate bond trades.

⁴Asquith *et al.* (2010) find for a sample of NYSE and NASDAQ stocks, that short trades are 27.9% of trading volume in 2005.

Using these estimates implies that shorting corporate bonds is an important market activity. The percentage of corporate bonds shorted, 1.3%, is slightly greater than half the percentage of stocks shorted, 2.5%. Furthermore, the percentage of all daily corporate bond trades that represents short selling, 19.1%, is almost two-thirds the percentage of stock trades that entails short selling, 29.8%. Thus, at any point in time the amount of corporate bonds shorted is large, and trading in the corporate bond market includes significant short sale activity.

1.4 Description of Sample

We use four separate databases, two that are commercially available and two that are proprietary, to construct the sample of corporate bonds used in this paper. All four databases cover the period from January 1, 2004 through December 31, 2007. The commercially available databases are the Trade Reporting and Compliance Engine database (TRACE) and the Fixed Income Securities Database (FISD). The two proprietary databases are a bond inventory database and a bond loan database. These databases were provided to us by one of the world's largest custodians of corporate bonds. The bond inventory database contains all corporate bonds available for lending, and the companion bond loan database describes the loans made from that inventory. The bond CUSIP is used as the common variable to link these four databases.

TRACE is a database of all OTC corporate bond transactions and was first implemented on a limited basis on July 1, 2002. TRACE reports the time, price, and quantity of bond trades, where the quantity is top-coded if the par value of the trade is \$5 million or more for investment grade bonds and \$1 million or more for high yield bonds. Over time, bond coverage expanded in phases, and the compliance time for reporting and dissemination of bond prices shortened. Our sample begins between Phase II and III of TRACE. Phase II was implemented on March 3, 2003, while Phase III was implemented in two stages, on October 1, 2004 and on February 7, 2005. Phase III required reporting on almost all public corporate

bond transactions.⁵ Since the vast majority of corporate bonds are traded over-the-counter, TRACE provides the first reliable daily pricing data for corporate bonds.

The FISD database contains detailed information on all corporate bond issues including the offering amount, issue date, maturity date, coupon rate, bond rating, whether the bond is fixed or floating rate, and whether it is issued under SEC Rule 144a. We exclude any corporate bond in the inventory file that we cannot match to FISD. In addition we also exclude all convertibles, exchangeables, equity-linked bonds, and unit deals.

The proprietary bond inventory database contains the number of bonds in inventory and number of bonds available to lend. From January 1, 2004 through March 30, 2005 we have end-of-the month inventory information for all bonds. The database reports daily inventory information from April 1, 2005 to December 31, 2007. In contrast to the inventory database, the loan database is updated daily for the entire period January 1, 2004 through December 31, 2007.⁶ For each day, the loan database includes which bonds are lent, the size of the loan, the rebate rate paid to the borrower, and an indicator of who borrows the bond. The proprietary loan database identifies 65 unique borrowers for corporate bonds. These borrowers are primarily brokerage firms and hedge funds.

Table 1.1 describes the match between the proprietary bond inventory and loan databases to the overall universe of FISD corporate bonds averaged by day. Panel A shows that from 2004 to 2007, the average number of bonds in the inventory database is 7,752. This represents 20.7% of all corporate bonds in FISD for an average day. The relationship between the number of bonds in FISD and the inventory is stable over each of the four years. Although not aggregated in Table 1.1, there are a total of 15,493 unique bonds in the bond inventory sample that match

⁵Phase I of TRACE covered transaction information on approximately 500 bonds. It required users to report transaction information on covered bonds to the NASD (since renamed FINRA) within 75 minutes. Phase II of TRACE expanded coverage of bonds to approximately 4,650 bonds. Coverage of additional 120 bonds was added on April 14, 2003. On October 1, 2003 the time to report was shortened to 45 minutes. A year later, on October 1, 2004, reporting time was shortened again to 30 minutes. Finally, on July 1, 2005 the reporting time was shortened to 15 minutes. Most reported trades are immediately disseminated by FINRA.

⁶There are several missing days in the loan database. On these days the file we obtained from the proprietary lender was either unreadable or a duplicate of an earlier daily file. These days are December 16-31, 2004, all of February 2005, June 7, 2006, and November 27, 2007.

Table 1.1: Number and Par Value of Bonds in Corporate Bond Databases

	Panel A: Daily Average Number of Bonds				
	2004 - 2007	2004	2005	2006	2007
Number of Corporate Bond CUSIPs in FISD	37,535	32,919	35,796	37,471	39,163
Number of Corporate Bond CUSIPs in Both Lender Database and FISD	7,752	7,592	7,669	7,750	7,827
Percent of FISD Represented in Lender Database	20.7%	23.1%	21.4%	20.7%	20.0%
Number of Corporate Bond CUSIPs in Lender Database and FISD That Go on Loan	2,901	2,612	2,797	2,841	3,054
Percent of Corporate Bond CUSIPs in Lender Database and FISD That Go on Loan	37.4%	34.4%	36.5%	36.7%	39.0%

	Panel B: Par Value of Bonds				
	2004 - 2007	2004	2005	2006	2007
Average Daily Par Value of Existing FISD Bonds (Billions of \$)	6,619	5,649	6,105	6,530	7,159
Average Daily Par Value of Existing FISD Bonds in Lender Inventory (Billions of \$)	193.3	183.4	186.7	195.5	196.8
Lender Inventory as a % of FISD Par Value	2.9%	3.2%	3.1%	3.0%	2.7%
Average Daily Par Value of Bonds On Loan in Lender Inventory (Billions of \$)	14.3	14.2	14.7	13.9	14.4
Lent as a % of Lender Inventory	7.4%	7.7%	7.9%	7.1%	7.3%

Table 1.1 reports the number and par value of bonds in the FISD Corporate Bond, Proprietary Bond Inventory, and Proprietary Bond Loan databases for the overall period and by year. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. All data is daily except for data from the proprietary inventory database which is only available monthly from January 1, 2004 to March 31, 2005.

to FISD at some point. In addition, 2,901 or 37.4% of bonds in the lender inventory are on loan on an average day. There is a slight upward trend in the fraction of bonds lent from inventory during 2004 to 2007. There are 10,293 unique bonds in the merged database that are lent at some point during the four-year period.

Table 1.1 Panel B reports similar comparisons using the par value of the bonds. The average daily par value of corporate bonds outstanding in the FISD database during the period 2004 to 2007 is \$6.6 trillion, while the average daily par value of corporate bond inventory in the database is \$193.3 billion. This represents 2.9% of the total par value of corporate bonds issued and listed in FISD. Of this inventory, an average \$14.3 billion, or 7.4% of the total par value of the inventory, is on loan each day.

In Figure 1.1, we plot our proprietary lender's number of loans outstanding, on the left hand axis, and the total par value of these loans, on the right hand axis, over time. On an average day, there are between 7,000 and 11,000 outstanding loans. The total par value of outstanding loans also fluctuates around the overall mean of \$14.3 billion, with a maximum of more than \$16.8 billion in October 2004, and a minimum of about \$10.5 billion in January 2004.

Table 1.1 and Figure 1.1 clearly demonstrate that the number and value of corporate bonds and corporate bond loans in the two proprietary databases are large. The bond inventory database covers 20.7% of the bonds in FISD. The par value of the inventory is \$193.3 billion on average, representing 2.9% of the \$6.6 trillion market. In total, the proprietary database consists of 367,751 loans, covering 10,293 bonds, and representing an average par value of \$14.3 billion per day. We believe this is of sufficient size to draw inferences about the overall market.

1.4.1 Sample Characteristics

Tables 1.2 and 1.3 compare various bond characteristics from FISD to the proprietary inventory and loan databases by year and for the entire period. It allows us to determine how representative the proprietary databases are of the entire corporate bond market. We

Figure 1.1: Number and Par Value of Outstanding Loans




Figure 1.1 plots the evolution of the corporate bond loans from the Proprietary Bond Inventory and Loan databases the over time. The left-hand axis reports the number of loans outstanding, while the right-hand axis shows the total par value of these loans. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007.

focus on characteristics that are likely to affect the demand and supply for corporate bond loans. The characteristics we examine are the size at issue, maturity, time since issuance, percent defaulted, percent floating rate, and percent subject to SEC Rule 144a. Rule 144a is a provision that allows for certain private resale of restricted securities to qualified institutional buyers.

Table 1.2 shows that the average bond in the inventory is much larger at issue (\$418.6 million) than the average FISD bond at issue (\$175.3 million). The average bond lent is even larger at issue with a size of \$487.4 million. The average maturity at issue of the bonds in the inventory database (11.3 years) is close to the average maturity at issue of the universe of all FISD corporate bonds (10.7 years). The average maturity at issue for lent bonds is 11.9 years. A comparison of time since issuance indicates that lent bonds are not outstanding as long as the average bond in the inventory or in FISD. There are no year-to-year trends in the values of these bond characteristics.⁷

Bonds in the FISD database are less likely to default (0.6%) than bonds in inventory (1.1%) and the default percentage for lent bonds is between the two (0.7%). Bonds on loan are much less likely to be floating rate bonds (10.6%) than bonds in either the FISD dataset (22.4%) or the inventory dataset (17.0%). The fraction of bonds that are subject to SEC Rule 144a is much lower for bonds on loan than for the FISD and inventory samples.

Table 1.3 reports Standard and Poor's (S&P) rating characteristics of corporate bonds. The coverage of the S&P ratings information in FISD is not as extensive as those characteristics reported in Panel A, however. For instance, there are 57,896 bonds in FISD where we observe the size at issue, while we observe S&P ratings for only 39,197 of these bonds. Fortunately, the limited coverage of ratings in FISD has a smaller impact on the inventory and loan samples. While we have issue size information for 10,293 lent bonds, we have an S&P rating for 9,822, or 95.4% of lent bonds.

⁷The values for some of the variables, e.g. maturity and time since issuance, over the entire period are outside the range of the per-year means. This is because each bond is only counted once for the entire period, but may be counted multiple times when counting the observations in the per-year columns. For example, the number of FISD, inventory, and lent bonds for the entire sample period is not the respective sums of the four separate years.

Table 1.2: Non-Rating Characteristics of Bonds in the Corporate Bond Databases

	2004 - 2007			2004			2005			2006			2007		
	FISD (57,896), Inventory (15,493), Lent (10,293)			FISD (38,075), Inventory (9,730), Lent (5,449)			FISD (40,835), Inventory (9,534), Lent (5,771)			FISD (43,189), Inventory (9,909), Lent (6,321)			FISD (44,807), Inventory (9,884), Lent (6,256)		
	Average	Standard Deviation	Deviaton	Average	Standard Deviation	Deviaton	Average	Standard Deviation	Deviaton	Average	Standard Deviation	Deviaton	Average	Standard Deviation	Deviaton
Number of Observations:															
Size At Issue (Millions of \$)															
<i>FISD</i>	\$175.3	\$324.8	\$168.1	\$288.7	\$168.5	\$296.1	\$175.4	\$314.0	\$183.8	\$339.1					
<i>Lender Inventory</i>	\$418.6	\$461.1	\$374.3	\$408.0	\$402.7	\$431.7	\$435.7	\$460.7	\$474.9	\$496.5					
<i>Lent</i>	\$487.4	\$481.7	\$466.2	\$461.2	\$484.1	\$471.4	\$505.3	\$478.1	\$555.9	\$518.6					
Maturity at Issuance (years)															
<i>FISD</i>	10.7	10.1	12.5	10.5	12.2	10.4	12.0	10.5	12.1	10.7					
<i>Lender Inventory</i>	11.3	10.1	12.0	10.1	12.1	10.3	12.0	10.7	12.4	11.1					
<i>Lent</i>	11.9	10.2	12.1	9.6	12.1	9.7	12.2	10.3	12.8	10.9					
Time Since Issuance (years)															
<i>FISD</i>	5.4	5.7	5.5	5.8	5.3	5.7	5.4	5.6	5.4	5.6					
<i>Lender Inventory</i>	4.4	4.0	4.3	3.8	4.3	3.9	4.4	4.0	4.4	4.1					
<i>Lent</i>	3.6	3.1	3.3	2.8	3.4	2.9	3.7	3.2	3.8	3.4					
% Defaulted															
<i>FISD</i>	0.6%		0.8%		0.7%		0.6%		0.5%						
<i>Lender Inventory</i>	1.1%		1.4%		1.1%		1.0%		1.0%						
<i>Lent</i>	0.7%		0.8%		0.6%		0.7%		0.5%						
% Floating Rate															
<i>FISD</i>	22.4%		15.9%		17.3%		19.5%		19.5%						
<i>Lender Inventory</i>	17.0%		10.3%		11.5%		15.2%		16.0%						
<i>Lent</i>	10.6%		5.5%		6.5%		9.1%		10.0%						
% Rule 144a															
<i>FISD</i>	20.6%		17.0%		17.8%		19.8%		19.6%						
<i>Lender Inventory</i>	23.0%		16.0%		16.1%		18.1%		18.8%						
<i>Lent</i>	14.2%		6.5%		8.6%		10.1%		10.4%						

Table 1.2 reports bond characteristics from the FISD Corporate Bond, Proprietary Bond Inventory, and Proprietary Bond Loan databases. All ratings are S&P Ratings. Ratings data is missing for some FISD bonds. Therefore, the FISD dataset in Panel B is a subset of the overall FISD dataset in Panel A. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. Time series variables are daily averages. "Rating at Issue" is defined as the first S&P rating. For rating and rating at issue, we report the median. The treasury spread variable is available over the entire sample period for 15,785, 8,601, and 6,236 bonds in FISD, Lender Inventory, and Lent, respectively. In 2004, it is available for 13,235, 5,960, and 3,527 bonds. In 2005, it is available for 12,917, 5,605, and 3,821 bonds. In 2006, it is available for 12,686, 5,523, and 3,821 bonds. In 2007, it is available for 12,576, 5,584, and 3,830 bonds.

Table 1.3: Rating Characteristics of Bonds in the Corporate Bond Databases

	2004 - 2007			2004			2005			2006			2007		
	FISD (39,197), Inventory (13,836), Lent (9,822)			FISD (27,513), Inventory (8,972), Lent (5,272)			FISD (30,338), Inventory (8,850), Lent (5,601)			FISD (31,841), Inventory (9,005), Lent (6,105)			FISD (32,729), Inventory (8,995), Lent (6,013)		
	Median / Average	Standard Deviation		Median / Average	Standard Deviation		Median / Average	Standard Deviation		Median / Average	Standard Deviation		Median / Average	Standard Deviation	
Median Rating at Issue															
<i>FISD</i>	A			A-			A-			A			A		
<i>Lender Inventory</i>	BBB			BBB+			BBB+			BBB+			BBB+		
<i>Lent</i>	BBB			BBB+			BBB+			BBB+			BBB+		
Median Rating over Period															
<i>FISD</i>	A-			BBB+			A-			A-			A-		
<i>Lender Inventory</i>	BBB+			BBB			BBB+			BBB+			BBB+		
<i>Lent</i>	BBB			BBB			BBB			BBB			BBB		
% Investment Grade at Issue															
<i>FISD</i>	79.2%			78.1%			78.7%			79.0%			79.1%		
<i>Lender Inventory</i>	69.0%			72.4%			73.8%			74.5%			74.3%		
<i>Lent</i>	68.9%			72.6%			71.3%			71.6%			72.0%		
% Investment Grade when Lent															
<i>FISD</i>	70.7%			70.9%			69.5%			70.2%			72.0%		
<i>Lender Inventory</i>	70.7%			69.8%			69.6%			71.1%			71.3%		
<i>Lent</i>	64.3%			64.0%			61.3%			64.8%			66.4%		
Treasury Spread (bps)*															
<i>FISD</i>	177.7		181.5	170.0		181.4	179.1		184.1	185.5		185.4	191.8		185.7
<i>Lender Inventory</i>	178.4		155.2	163.4		145.0	156.3		133.3	155.6		130.2	161.3		130.0
<i>Lent</i>	164.7		137.8	146.8		118.9	147.7		117.1	152.3		123.2	157.6		123.9

Table 1.3 reports bond characteristics from the FISD Corporate Bond, Proprietary Bond Inventory, and Proprietary Bond Loan databases. All ratings are S&P Ratings. Ratings data is missing for some FISD bonds. Therefore, the FISD dataset in Panel B is a subset of the overall FISD dataset in Panel A. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. Time series variables are daily averages. "Rating at Issue" is defined as the first S&P rating. For rating and rating at issue, we report the median. The treasury spread variable is available over the entire sample period for 15,785, 8,601, and 6,236 bonds in FISD, Lender Inventory, and Lent, respectively. In 2004, it is available for 13,235, 5,960, and 3,527 bonds. In 2005, it is available for 12,917, 5,605, and 3,821 bonds. In 2006, it is available for 12,686, 5,523, and 3,821 bonds. In 2007, it is available for 12,576, 5,584, and 3,830 bonds.

The bond inventory has a lower median rating at time of issue and over our time period than the universe of FISD corporate bonds. The sample of lent bonds has the same median rating at time of issue as inventory, but a lower rating over the entire period. The other rows of Panel B, which show percentage investment grade at issue and percentage investment grade as of the date of the loan, show a pattern consistent with the lower ratings for lent bonds than for FISD bonds.⁸

In summary, Tables 1.2 and 1.3 show that shorted bonds are much larger at issue, have a slightly longer maturity at issue, and have a lower median rating at issue than the average FISD bond. 68.9% of the lent bonds are investment grade, while 79.2% of all FISD bonds are. Lent bonds are also more likely to be fixed rate and less likely to be defaulted.

1.4.2 Properties of Short Positions

Each loan in the loan database has a unique loan number, which allows us to describe the time series properties of lent positions. Using the loan number, we are able to determine when the loan is initiated, the duration of the loan, and the number of bonds lent over the duration of the loan. Table 1.4 provides descriptive statistics for the new bond loans in the database in total and split by whether the bonds are investment grade or high yield and unrated.⁹ There are 10,293 unique bonds lent in the database, and 367,751 unique loans for an average of 35.7 loans per bond. Some bonds change ratings between investment grade and high yield over the sample period. There are 293,649 loans on investment grade bonds and 128,102 loans on high yield bonds.

The data in Table 1.4 indicates that the size and duration of loans are skewed and this skewness differs between investment grade and high yield loans. The mean loan size for

⁸The data on treasury spreads has a different pattern. The lent bonds have a smaller spread to treasuries than do our inventory or the FISD database. It is important to note, however, that the available information on treasury spreads is much smaller than that of bond ratings, and therefore these two descriptives are not directly comparable since the samples are different. The notes in Tables 1.2 and 1.3 give more information on this issue.

⁹There are only 13,884 loans to unrated bonds in our database and we have grouped them with high yield bonds. Holding the unrated bonds out as a separate sample does not change the analysis. We will refer to high yield and unrated bonds as high yield in the text going forward.

Table 1.4: Loan Size, Loan Duration, and Changes in Loan Size

Year	Investment Grade		High Yield and	
	All Bonds	Bonds	Unrated Bonds	Unrated Bonds
Number of New Loans	367,751	239,649	128,102	128,102
Size of New Loans (Bonds)				
Mean	1,444.1	1,267.3	1,774.8	1,774.8
Median	350	200	980	980
Mode	100	100	1000	1000
10th percentile	73	58	100	100
25th percentile	100	100	210	210
75th percentile	1,435	1,000	2,000	2,000
90th percentile	4,000	3,500	4,350	4,350
Duration of New Loans (Days)				
Mean	32.4	28.3	40.1	40.1
Median	11	10	13	13
Mode	1	1	1	1
10th percentile	1	1	1	1
25th percentile	3	3	3	3
75th percentile	34	30	43	43
90th percentile	83	71	107	107
Changes in Loan Size				
Percentage of loans that decrease in size	31.2%	29.3%	34.8%	34.8%
Average total decrease in loan size (for loans that decrease)	56.9%	55.7%	58.6%	58.6%
Average number of decreases (for loans that decrease)	1.9	1.9	1.8	1.8

Table 1.4 provides descriptive statistics for the new bond loans in the Proprietary Bond Inventory and Loan databases for all bonds and by credit rating. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. Size of New Loans is reported as the number of bonds lent. Duration of New Loans is reported as the number of days that bonds are lent. New loans are only defined when we have loan data for the previous day. That is, for the first day of data or the first day after missing data, no loans classified as new. Similarly, duration and changes in loan size are only defined when the last day of a loan is not the day before a missing day. Thus, in the duration and loan size calculations, there are only 359,754 loans in total, comprised of 234,450 in investment grade bonds and 125,306 in high yield and unrated bonds.

investment grade bonds is 1,267.3 bonds or approximately \$1.3 million at a par value of \$1,000. The median loan size is only 200 bonds or \$200,000. The mode loan size is \$100,000. High yield bonds have a higher mean, \$1.8 million, a higher median, \$980,000, and a mode of \$1.0 million. The mean new investment grade loan is outstanding for 28 calendar days, with a median time outstanding of 10 days and a mode of one day. For high yield bonds the mean is 40 days, the median is 13 days, and the mode is also one day. Thus, loans for high yield bonds are larger and longer than those for investment grade bonds.

The last three rows of Table 1.4 show how often loan size changes during the life of the loan. Changes to loan size may occur if borrowers partially repay the loan or if portions of their loan are recalled by the lender. In the sample, loan size decreases for 29.3% of investment grade loans and 34.8% of high yield loans before the loan is closed. Of the loans which change size, the average decreases of the initial loan size are 55.7% and 58.6% respectively, and the average number of decreases are 1.9 and 1.8. We do not observe increases in loan size, presumably because a borrower who wishes to borrow more bonds initiates a new loan.

Tables 1.1, 1.2, 1.3, 1.4 and Figure 1.1 show that the proprietary inventory and loan databases are extensive. The inventory database covers over 20% of all corporate bonds issued and the loan database contains over 367,000 loans on over 10,000 bonds. The average amount in inventory per day is \$193.3 billion, and the average amount on loan per day is \$14.3 billion. The lent bonds are larger, have a longer duration, and have a lower rating than the average bond in the FISD database. Loan activity is extensive throughout the entire period. New bond loans average over \$1.4 million and have an average duration of 32 days. Loans on high yield bonds are larger and longer than those on investment grade bonds. Finally, approximately one-third of loans are partially repaid before being closed out.

1.5 Costs of Borrowing Corporate Bonds

The borrowing cost for corporate bonds has two major components: the rebate rate paid by the lender and the market interest rate. The rebate rate is the interest rate the lender pays on the collateral posted by the borrower and is typically lower than the market rate that

the borrower could receive on the same funds invested at similar risk and duration elsewhere. Thus, we calculate the cost of borrowing as the difference between the market rate and the rebate rate. The loan database gives the rebate rate paid by the lender, but not the market rate. We use the one-month commercial paper rate as a proxy for the market rate.¹⁰

Even though most corporate bond loans are short term, as shown in Table 1.4, borrowing costs vary frequently over the life of the loan. Although not shown in a table, overall, 49.3% of the bond loans in the sample experience a change of at least 5 bps in their borrowing cost before repayment. These changes are due both to changes in the rebate rate and changes in the commercial paper rate. 42.3% of bond loans experience a rebate rate change of at least 5 bps, while 21.2% experience a change in the commercial paper rate of at least 5 bps.¹¹

It is possible for the lender to change the rebate rate frequently because all of the loans are demand loans. In addition, if supply and demand conditions for the bond improve, and if the lender does not raise the rebate rate, the borrower has the option of closing out the loan and borrowing from a different lender. For the loan sample, there is an average of 3.5 rebate rate changes of at least 5 bps per loan, or approximately 8 rebate rate changes for those loans with changes. Furthermore, rebate rate changes of at least 5 bps go in both directions. 38.4% of all loans have a rebate rate increase, 29.7% of all loans have a rebate rate decrease, and 25.8% of all loans have both. Hence, a considerable factor driving changes in the cost of borrowing is changes in the rebate rate on existing loans by the lender.

The frequent changes in borrowing costs suggest that existing loans should track current market conditions, although perhaps with a lag. Comparing new and existing loans, the average absolute difference in the borrowing costs for the same bonds on the same day is 4.3 bps, with a standard deviation of 27.6 bps. Moreover, for those bonds that have new and existing loans on the same day, 46.5% of new loans have an average borrowing cost that is

¹⁰An alternative to the commercial paper rate is the Fed Funds rate. We use the commercial paper rate because we think it more accurately represents the rate the borrowers could get on their collateral. For most of the period, January 1, 2004 through December 31, 2007, the commercial paper and Fed Funds rates correlate highly (the average difference across days is 4.9 bps and the coefficient of correlation is 0.998).

¹¹High yield loans are more likely than investment grade loans to experience a rebate rate change of at least 5 bps (45.9% versus 40.4%), though this may be due to their longer average duration.

more expensive than existing loans and 35.4% of new loans are cheaper than existing loans. Given these differences, the analyses below only use the borrowing cost for new loans unless otherwise stated. All loans start as new loans, and new loans must reflect current market conditions.

1.5.1 Time Series and Cross-Section of Borrowing Costs

Figure 1.2 plots the distribution of equally-weighted borrowing costs by quintile for each month of our sample period. The plot shows that the distribution of borrowing costs changes abruptly between March and July 2006. Before March 2006, the 60th and 80th percentiles of borrowing costs are usually at or above 50 bps for each month. After March 2006, the 60th percentile is at or below 20 bps for each month. The 80th percentile drops below 20 bps in August 2006 and is near or below 20 bps until the start of the Credit Crunch in August 2007. The plot of value-weighted loan borrowing costs, although not shown, shows a similar but less dramatic pattern during the same time period.

The reasons why borrowing costs are reduced in 2006 are not immediately clear. Table 1.1, Table 1.2, Table 1.3 and Figure 1.1 show that the lender's inventory of bonds and the amount lent do not change significantly after 2005. Furthermore, although not reported, the duration of bond loans also does not change significantly over time. To further investigate the decline in borrowing costs, Table 1.5 presents borrowing costs over time partitioned by loan size and credit quality. Over the sample period 2004-2007, shown in the first column, the equally-weighted mean and median borrowing costs are 33 bps and 18 bps, respectively.¹² The composition of loans by size and credit quality does not change dramatically in 2006. Table 1.5 shows that the percentage of large loans remains fairly constant (it decreases slightly by 2007) and the percentage of investment grade loans remains flat.

Panel A divides loans into those of 100 bonds or less (i.e., \$100,000 par value, the overall mode loan size) and those of more than 100 bonds. It shows that large loans have lower

¹²The borrowing costs in Table 1.5 are equally-weighted by loan. When value-weighting borrowing costs by loan size, the value-weighted mean borrowing cost is 22 bps and the median is 14 bps.

Figure 1.2: *Equally-Weighted Monthly Distribution of Loan Borrowing Costs*




Figure 1.2 plots the equally-weighted borrowing cost quintiles monthly from the Proprietary Bond Inventory and Loan databases over time. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007.

Table 1.5: *Distribution of New Loan Borrowing Costs*

Panel A: Borrowing Costs (bps) by Loan Size															
Year	2004-2007			2004			2005			2006			2007		
	All Loans	≤100	>100	≤100	>100	≤100	≤100	>100	≤100	>100	≤100	>100	≤100	>100	>100
Number of New Loans	367,751	109,124	258,627	23,127	58,994	24,067	64,854	27,126	67,194	34,804	67,585				
Mean	32.6	39.0	29.9	50.9	31.1	62.6	38.5	33.0	31.1	19.4	19.2				
Median	18.0	48.0	16.0	52.0	22.0	56.0	27.0	20.0	15.0	13.0	13.0				
Mode	13.0	13.0	13.0	51.0	51.0	49.0	14.0	13.0	13.0	13.0	13.0				
10th percentile	7.0	9.0	6.0	12.0	11.0	24.5	11.0	9.0	8.0	3.0	-1.0				
25th percentile	12.0	13.0	12.0	50.0	15.0	50.0	16.0	13.0	12.0	9.0	5.5				
75th percentile	51.0	54.0	49.0	65.0	51.0	67.0	54.0	50.0	35.0	21.8	15.0				
90th percentile	64.0	69.0	59.0	76.0	58.0	74.0	69.0	61.0	56.0	50.0	49.0				

Panel B: Borrowing Costs (bps) by Credit Rating															
Year	2004-2007			2004			2005			2006			2007		
	All Loans	Investment	Grade	Investment	Grade	High Yield	Investment	Grade	High Yield	Investment	Grade	High Yield	Investment	Grade	High Yield
Number of Loans	367,751	239,649	128,102	54,457	27,664	56,119	32,802	60,752	33,568	68,321	34,068				
Mean	32.6	30.0	37.4	38.4	33.4	44.7	45.6	26.0	41.8	14.8	28.2				
Median	18.0	18.0	18.0	50.0	23.9	50.0	27.0	15.0	17.0	13.0	13.0				
Mode	13.0	13.0	13.0	51.0	51.0	49.0	14.0	13.0	13.0	13.0	13.0				
10th percentile	7.0	7.0	7.0	11.0	11.0	13.0	11.8	8.0	9.0	-0.5	-0.5				
25th percentile	12.0	12.0	13.0	16.0	15.0	20.0	16.0	12.0	12.0	8.0	7.0				
75th percentile	51.0	51.0	50.0	53.0	51.0	60.0	55.0	48.0	49.0	15.0	28.0				
90th percentile	64.0	63.0	66.0	68.0	66.0	71.0	73.0	55.0	68.0	48.0	50.0				

Table 1.5 Panel A presents borrowing costs over time partitioned by loan size. Panel B presents borrowing costs over time partitioned by credit rating - investment grade vs. high yield. Unrated bonds are included with high yield bonds. All borrowing costs are calculated on an equally weighted basis. Data is from the Proprietary Bond Loan database for the overall period and by year. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. Loan Borrowing Costs are defined as the One-month Commercial Paper Rate minus the Rebate Rate. Loans are allocated to the year in which they are initiated, even if they extend into subsequent years. New loans are only defined when we have loan data for the previous day. That is, for the first day of data or the first day after missing data, no loans classified as new.

borrowing costs than small loans, but this difference diminishes over time. For example, in 2004 the mean borrowing cost for loans of 100 bonds or less is 51 bps. For loans of more than 100 bonds, the mean borrowing cost is 31 bps. By 2007, it appears that size is no longer priced as the mean borrowing cost for small loans is 19 bps, which is identical to that of large loans. The median borrowing costs behave similarly.

Panel B presents borrowing costs over time by credit rating. For the entire period high yield bonds have a higher average borrowing cost than investment grade bonds, 37.4 bps versus 30.0 bps, but identical medians of 18.0 bps. Borrowing costs for both investment grade and high yield bonds decline by 2007. The decline in both mean and median borrowing costs is greater for investment grade bonds than for high yield bonds.

Thus, Table 1.5 shows that average borrowing costs are usually lower for large loans and for investment grade loans. Borrowing costs generally decline over our sample period regardless of loan size and credit quality. In addition, the decline in borrowing costs is not explained by changes in the composition of large vs. small or investment grade vs. high yield loans.

Another factor why borrowing costs change over time may be greater transparency in bond market pricing related to the growth of TRACE during our sample period. The sample begins between Phases II and III of TRACE. As stated above, Phase II was implemented on April 14, 2003, while implementation of Phase III was completed by February 7, 2005. Phase III required reporting on almost all public corporate bond transactions. It seems unreasonable, however, that it would take more than a year, until April 2006, for the effects of this increased coverage to have an impact. Finally, the growth of the CDS market may have driven improvements in the liquidity of corporate bonds, and the narrowing of borrowing cost spreads may reflect this trend. We investigate the impact of the CDS market for the market for borrowing corporate bonds in Section 1.8.

1.5.2 Determinants of Borrowing Costs

We next investigate how the cost of borrowing is related to the available supply of bonds in the lender's inventory. As previously mentioned, we do not have daily inventory data from

January 2004 to March 2005, and thus cannot compute the daily available supply of bond inventory during this period. Figure 1.3 plots the relationship between the average borrowing cost and the amount of inventory on loan divided by investment grade and high yield loans for the periods April 2005 to March 2006 and April 2006 to December 2007. It also plots the Credit Crunch 2007 period from July 2007 to December 2007 which we will discuss later in Section 1.9. The vertical axis displays average borrowing cost and the horizontal axis displays amount of inventory lent.

For both periods, April 2005 to March 2006 and April 2006 to December 2007, high yield bonds are more expensive to borrow than investment grade bonds at higher on loan percentages. When the loan percentage is below 40-45%, there is no noticeable difference in borrowing costs between high yield and investment grade bonds. However, when the on-loan percentage is greater than 45% high yield bonds become more expensive while the cost of borrowing investment grade bonds remain flat. Finally, at approximately 70% on loan there is a steep increase in the average borrowing cost for high yield bonds: each 10% increase in the amount on loan is associated with a greater than 10 bps increase in the average borrowing cost.¹³ In contrast, the borrowing costs for investment grade bonds continue to be insensitive to on loan percentage.

Figure 1.3 also shows that borrowing costs are significantly lower in the latter period, April 2006 through December 2007, compared to the earlier period, April 2005 to March 2006. This is true for both high yield and investment grade bonds. This result is consistent with Figure 1.3 and Table 1.5, which show a decrease in borrowing costs after April 2006. Note, the kink at 70% of available inventory still exists, and although borrowing costs are lower in the latter period, the slope of that segment is similar. This suggests that the reduction in borrowing costs in the latter half of our sample period is not due to changes in how inventory impacts borrowing costs. Finally, the pattern for high yield and investment grade bonds during the 2007 Credit Crunch is similar.

¹³The pattern for high yield bonds is consistent with the results of D'avolio (2002) and Kolasinski *et al.* (2013) for the equity loan market. Neither paper divides the equity loan market by credit quality.

Figure 1.3: Relationship Between Borrowing Cost and Percent of Inventory On Loan

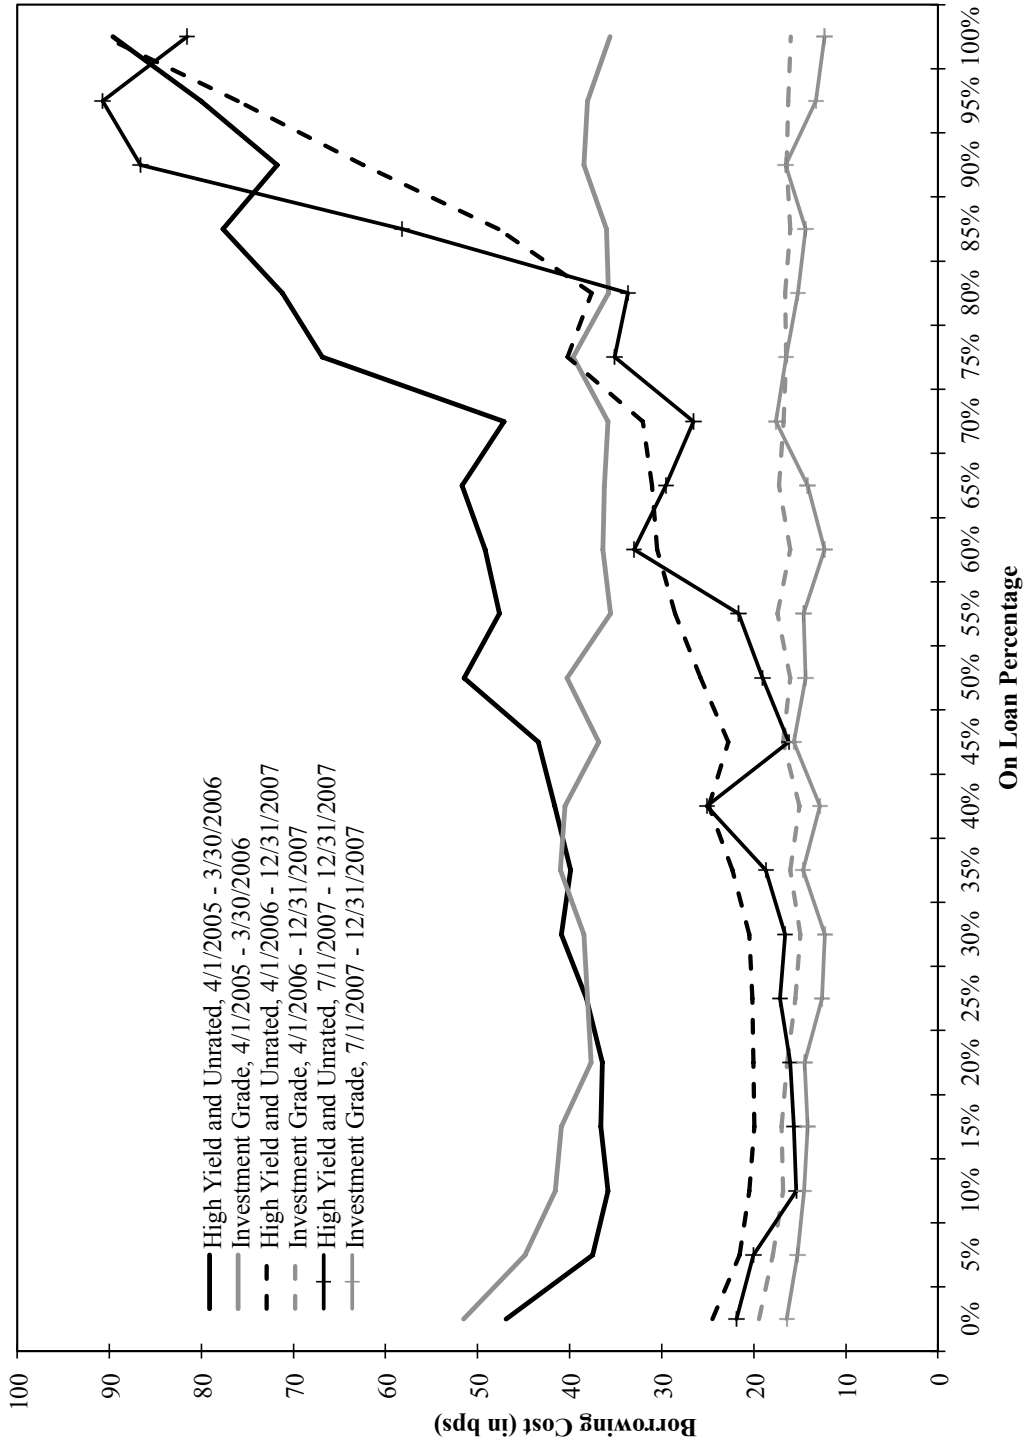


Figure 1.3 plots the relationship between the average borrowing cost and the amount of inventory on loan for the period April 2005 to December 2007 and for several sub-periods, by credit status. Data is from the Proprietary Bond Inventory and Loan databases. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded.

Table 1.6 presents the 35 corporate bonds with the highest borrowing costs in the sample. Each bond is listed once, together with its maximum loan borrowing cost and the date and borrowing cost corresponding to that maximum. Since there is a great deal of clustering by firm of the most expensive bonds to borrow, the last column of Table 1.6 indicates the number of bonds from that issuer where the borrowing cost is greater than the 250th most expensive to borrow bond in the sample. For example, the borrowing cost of the most expensive loan on the Calpine Corp bond with CUSIP 131347AW6 is 14.50%, but there are 10 other Calpine Corp bonds which have borrowing costs above the 250th most expensive to borrow bond in the sample.

There are three features of the bonds in Table 1.6 that are worth noting. First, these bonds are highly lent out. The average percentage on loan is 79.7%, well above the 70% “kink” observed in Figure 1.3. Second, most of the firms in Table 1.6 experience credit problems around the date they appeared on our list and as seen in the next to the last column, all of the bonds are high yield. Of the 35 firms on the list, 9 are bankrupt as of the date of the loan, while another 6, while not filing for bankruptcy, were downgraded in the prior year. In addition, 7 of the firms, while not bankrupt or downgraded, were frequently mentioned in the press in the previous year as “financially struggling.” Interestingly, 8 of the remaining firms undertook an LBO during this period. Although we did not check explicitly, we infer the increased leverage from the LBO impacted the bond’s borrowing cost.

A third feature of Table 1.6 is that a large fraction of the most expensive bond loans take place during the latter half of 2007. Thirteen out of 35 bond loans in our list are after July 1, 2007, and 8 of these are on one day, October 31, 2007. Importantly, all 8 have negative rebate rates on that date. This means their inclusion cannot be explained solely by that day’s reported commercial paper rate.

Calculated borrowing costs are not always positive. A negative borrowing cost is the result of the lender paying a rebate rate above the commercial paper rate, and it implies that the lender loses money on the loan. In total, we have 11,971 loans (or 3.3% of the total) with negative borrowing costs in the sample. Most of the loans with negative borrowing costs

Table 1.6: Corporate Bonds with the Highest Borrowing Costs

CUSIP	Issuing Company Name	Date	Rebate Rate (in bps)	Borrowing Cost (in bps)	On Loan %	Credit Rating	Number of Bonds
13134VAA1	CALPINE CDA ENERGY FIN ULC	5/10/06	-1,000	1,501	100.00%	D	1
131347AW6	CALPINE CORP	2/15/06	-1,000	1,450	75.87%	D	10
26632QAK9	DURA OPER CORP	2/28/07	-700	1,223	21.62%	D	2
247126AC9	DELPHI AUTOMOTIVE SYS CORP	2/2/06	-700	1,150	51.59%	D	3
07556QAN5	BEAZER HOMES USA INC	10/31/07	-479	932	100.00%	B+	2
45661YAA8	INEOS GROUP HLDGS PLC	10/31/07	-479	932	65.26%	B-	1
729136AF8	PLIANT CORP	10/31/07	-479	932	100.00%	CCC	3
909279AW1	UNITED AIR LINES INC	12/13/05	-500	927	90.18%	D	1
256605AD8	DOLE FOOD INC	10/31/07	-413	866	38.82%	B-	1
15101QAC2	CELESTICA INC	10/31/07	-400	853	100.00%	B-	1
800907AK3	SANMINA - SCI CORP	10/31/07	-400	853	76.75%	B-	1
194832AD3	COLLINS & AIKMAN PRODS CO	6/23/06	-300	824	99.46%	D	2
001765AU0	AMR CORP DEL	3/14/07	-250	775	76.26%	CCC+	1
370442BT1	GENERAL MTRS CORP	10/31/07	-288	741	88.29%	B-	6
35687MAP2	FREESCALE SEMICONDUCTOR INC	9/6/07	-200	728	84.32%	B	2
984756AD8	YANKEE ACQUISITION CORP	8/7/07	-200	728	100.00%	CCC+	2
85375CAK7	STANDARD PAC CORP NEW	10/31/07	-200	653	100.00%	B+	3
978093AE2	WOLVERINE TUBE INC	2/1/06	-200	648	64.29%	CCC	1
624581AB0	MOVIE GALLERY INC	10/24/06	-100	625	34.69%	CCC-	1
256669AD4	DOLLAR GEN CORP	10/16/07	-100	583	99.89%	CCC+	1
179584AG2	CLAIRES STORES INC	12/26/07	-75	539	99.00%	CCC+	2
767754AD6	RUTE AID CORP	8/2/06	0	532	10.73%	B-	2
156503AH7	CENTURY COMMUNICATIONS CORP	7/31/06	0	531	73.22%	D	3
373200AT1	GEORGIA GULF CORP	9/11/07	0	531	100.00%	B-	2
667281AM1	NORTHWEST AIRLIS INC	8/1/06	0	531	80.00%	D	4
640204AH6	NEIMAN MARCUS GROUP INC	7/18/06	0	530	97.22%	B-	1
651715AD6	NEWPAGE CORP	7/27/06	0	530	84.70%	CCC+	1
75040KAC3	RADIOLOGIX INC	7/18/06	0	530	98.81%	CCC+	1
872962AD7	TECHNICAL OLYMPIC USA INC	6/26/07	0	530	100.00%	CCC+	1
247361XY9	DELTA AIR LINES INC DEL	7/17/06	0	529	99.66%	D	4
420029AD2	HAWAIIAN TELCOM COMMUNICATIONS INC	7/26/06	0	529	82.81%	CCC+	5
721467AF5	PILGRIMS PRIDE CORP	8/7/07	0	528	99.76%	B	2
79546VAF3	SALLY HLDGS LLC / SALLY CAP INC	9/6/07	0	528	85.02%	CCC+	2
87971KAA5	TEMBEC INDS INC	12/12/06	0	528	14.27%	CCC-	3
682391AC1	155 EAST TROPICANA LLC / 155 EAST TROPICANA FIN CORP	6/29/06	0	527	99.45%	B-	1

Table 1.6 presents the 35 corporate bonds with the highest borrowing costs in our sample. Data is from the Proprietary Bond Loan database for the overall period and by year. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. Each bond is listed once with corresponding S&P credit rating, date, rebate rate, maximum loan borrowing cost, and on loan percentage. Number of Bonds is the number of bonds issued by a given firm that ever had borrowing costs greater than the 250th most expensive to borrow bond in our sample.

coincide with the 2007 Credit Crunch from August 2007 until December 2007. This can be seen in Figure 1.2, which shows that the borrowing cost of the bottom quintile becomes negative after July 2007. Of the 11,971 loans with negative borrowing costs, 8,832 of them occur between August and December 2007, of which 7,960 are on only 26 different days.

There is more than one possible reason why the cost of borrowing is negative for some bond loans. It is possible that the reported one-month commercial paper rate, which we take from the Federal Reserve Board's website, is not representative of the true market conditions for all days. This is particularly true for those days with very large intra-day interest rate movements. During the 2007 Credit Crunch, the Fed eased credit and dropped the Fed Funds rate several times, causing the commercial paper rate to fall as well. It is also possible that the proprietary lender is slow to respond to changes in credit conditions.

Finally, it should be noted that during the Credit Crunch in the last half of 2007, the Fed's intervention caused short-term rates to fall substantially below medium-term rates. If the reinvestment rate on collateral received by the lending institution is above short-term rates, the lender can still make a profit on their bond loans even with negative borrowing costs.¹⁴ Alternatively, the Credit Crunch of 2007 may have caused borrowers of the bond to want to close out their short position and have their collateral returned. If the lender has invested the collateral in illiquid securities which have lost value, they may have difficulty in returning collateral on demand. In this instance, they may subsidize borrowers to avoid reducing their collateral pool. This scenario was reported in the financial press and a number of lenders reported losses on their collateral during this period.¹⁵ To determine if the market for lending bonds in the period July to December 2007 is different, we examine this time period separately

¹⁴Our loan database provides a reinvestment rate which the lender estimates they will receive on the collateral. This rate is not constant across all loans or even across all loans on one particular bond at a point in time. The reason for this is that the lender invests the collateral in a number of different funds at the direction of the bond's owner. These funds can have a different duration and risk than that represented by investing short term at the commercial paper rate. We ignore these reinvestment rates when calculating borrowing costs since they do not represent the opportunity cost of the borrower's collateral.

¹⁵See Weiss, *AIG to Absorb \$5 Billion Loss on Securities Lending*, Bloomberg News, June 27, 2008 and Karmin and Scism, *Securities-Lending Sector Feels Credit Squeeze*, Wall Street Journal, October 30, 2008. Also, see State Street Press Release on July 7, 2010, *State Street Records Second-Quarter After-Tax Charge of \$251 Million, or \$0.50 Per Share*.

in Section 1.9.

1.5.3 Regression Analysis of Borrowing Costs

Although we know that borrowing costs are lower in 2006 and 2007 than they are in 2004 and 2005 and that borrowing costs are dependent on the size of the loan, the credit rating of the bond, and the available inventory to borrow, it is hard to determine the relative importance of these factors from the univariate comparisons we have made so far. We next conduct a multivariate analysis, which allows us to simultaneously control for the factors we have examined.

Even though bond loans are fully collateralized, bond characteristics may affect borrowing costs because they reflect supply and demand conditions. A bond's time since issuance may be important if it affects how widely the bond is held, and thus how difficult it is to locate, or if investor beliefs become more heterogeneous the longer the bond is outstanding. The availability to borrow may also be proxied by whether the debt is public or private (Rule 144a), as private debt may be harder to sell short. Smaller issue size may also make the bonds harder to find, increasing borrowing costs. Another factor is whether the bond is fixed or floating rate. Floating rate bonds re-price with interest rate movements and are thus less likely to deviate from par.

In addition, a bond's rating may be an important determinant of borrowing costs. As stated earlier, high yield bonds might attract more shorting activity because they are more likely to deviate from par than investment grade bonds. In our sample, 5.0% of the inventory for investment grade bonds is lent out, while for high yield it is 13.7% of inventory. Moreover, ratings will impact borrowing costs if lower rated bonds are in short supply. For our lender, investment grade bonds represent 70.8% of inventory, while high yield bonds are 29.2%. This 29.2% is under-represented (relative to FISD), where high yield bonds constitute 43.5% of the FISD universe.

Borrowing costs may also differ for a given bond because of loan characteristics. As Figure 1.3 shows for high-yield bonds, a larger percentage of bonds already on loan may lead to higher

borrowing costs. However, holding inventory constant, larger loans may have lower borrowing costs if there is a size discount. Further, borrowing costs may differ by borrower if the lender either gives a discount to large volume borrowers or if some borrowers are more knowledgeable about the lending market than others.

Our regression model incorporates the data on bond characteristics from Tables 1.2 and 1.3 as well as on loan percentage, loan size, and loan initiation day dummies. In some specifications, we also include dummy variables for each bond's CUSIP and the identity of the borrowing broker. The CUSIP controls allow us to examine how pricing varies across loan market variables, while fixing bond characteristics. Since daily inventory data is only available after March 2005, the regression analysis covers the period April 2005 through December 2007. The models we estimate are variations of the following model for the borrowing cost of loan i on bond b on day t :

$$\begin{aligned} \text{Borrowing Cost}_{ibt} = \text{CPrate}_t - \text{RR}_{ibt} = & \beta_1 \text{on loan } \%_{bt} + \beta_2 \text{loan size}_i + \beta_3 \text{rating}_{bt} \\ & + \beta_4 \text{issue size}_b + \beta_5 \text{time since issue}_{bt} + \beta_6 \text{floating rate}_b + \beta_7 \text{rule144a}_b \\ & + \delta_t + \kappa_b + \lambda_{\text{broker}} + \epsilon_{ibt} \end{aligned}$$

where CPrate is the one month financial commercial paper rate (in our model 100 basis points = 1.00) and RR is the rebate rate (with the same scale as the CPrate). The on loan % is the percentage of daily inventory already lent, and loan size is the total number of bonds lent in thousands of bonds (that is, the loan value in millions of dollars). Rating is the bond's S&P rating at the time of the loan (where AAA is given a value of 1, D is given a value of 22, and all intermediate ratings are given consecutive values between 1 and 22). Issue size is the size of the initial bond offering (in \$100 millions). The time since issue variable is the time since the bond was issued (in years). The floating rate variable is a dummy variable equal to 1 if the bond pays a floating rate coupon and 0 if the bond has a fixed rate coupon. The Rule 144a variable is a dummy variable equal to 1 if the bond was issued under SEC Rule 144a and 0 otherwise. δ_t represents a set of dummies for each trading day in the sample. κ_b represents a set of dummies for each bond CUSIP in the sample, and λ_{broker} are a set of dummies for each

unique borrower in the sample who borrows 100 or more times during our sample period.¹⁶ We report heteroscedasticity-robust standard errors.

Table 1.7 reports estimates from four specifications of the regression for all bonds: one without broker or bond CUSIP dummies, one with broker dummies, one with bond CUSIP dummies, and one with both broker and bond CUSIP dummies. The specifications with bond CUSIP dummies do not include issue size, floating rate, and Rule 144a since these characteristics are completely captured by the bond-specific controls. We also exclude time since issuance when we have bond and date controls since these controls together capture nearly all of the variation in this variable. In addition, Table 1.7 also reports estimates from two specifications for only high-yield bonds using both broker and bond CUSIP dummies.

In the first four specifications (i.e. for all bonds), the on loan % coefficient is positive and significant. In the two specifications without CUSIP dummies, the coefficient is 26.30 without broker dummies and 26.23 with broker dummies. When we add the bond-specific controls, the estimates fall to 3.19 and 4.38. The coefficients decrease because the bond-specific controls pick up much of the variation in bond inventory. Still, consistent with the pattern we observed in Figure 1.3, the larger the percentage of the inventory lent, the higher the borrowing cost. Increasing the percentage lent by 10% is associated with an increase in borrowing costs by 2.6 bps across the sample of all bonds. For a specific bond, a 10% increase in on loan percentage is associated with an increase of 0.3 to 0.4 bps on average.

For the sample of high-yield bonds in column (5), the coefficient for on loan percentage is 10.36, which is more than double the estimate for the same regression model on the entire sample in column (4). This is consistent with the plots in Figure 1.3 for high yield versus investment grade bonds. We also estimate a version of our regression model where we include two additional measures of on loan percentage. These allow the slope to differ when on loan percentage is greater than 50% and when it is greater than 70%. We choose these percentages

¹⁶Our lender identifies 65 borrowers. 40 make 100 or more loans and 25 make less than 100 loans during our sample period. The average number of loans made by the largest 40 is 9,178 and the average made by the smallest 25 is 25. Restricting our sample to the period covered by the regression, there are a total of 62 borrowers, 38 of whom make 100 or more loans.

Table 1.7: Regression Analysis of Determinants of Borrowing Costs

	All Bonds						High Yield Bonds Only	
	[1]	[2]	[3]	[4]	[5]	[6]		
On Loan %	26.30 *** (47.42)	26.23 *** (48.33)	3.19 *** (6.10)	4.38 *** (8.58)	10.36 *** (10.78)	4.44 *** (2.82)		
On Loan % * I (if On Loan % > 50)	0.05 (0.06)		
On Loan % * I (if On Loan % > 70)	5.09 *** (6.05)		
Loan Size (thousands)	-2.16 *** (-49.10)	-1.74 *** (-39.30)	-1.67 *** (-48.91)	-1.36 *** (-40.20)	-1.41 *** (-19.59)	-1.38 *** (-19.25)		
Bond Rating (where AAA=1, ..., D=22)	1.12 *** (35.12)	1.41 *** (40.64)	3.33 *** (16.30)	3.23 *** (16.05)	5.52 *** (15.84)	5.54 *** (15.93)		
Bond Issue Size (\$100M)	0.31 *** (22.15)	0.31 *** (22.46)						
Bond Time Since Issuance (years)	0.74 *** (17.90)	0.72 *** (17.87)						
Bond Floating	-5.86 *** (-13.32)	-5.72 *** (-13.10)						
Bond Rule 144a	3.34 *** (4.31)	3.05 *** (3.94)						
Broker Dummies	N	Y	N	Y	Y	Y		
CUSIP Dummies	N	N	Y	Y	Y	Y		
Broker Effects								
F-test	...	F = 969.21 ***	...	F = 1172.05 ***	198.02 ***	197.04 ***		
p-value	...	p < 0.0001	...	p < 0.0001	p < 0.0001	p < 0.0001		
max - min	...	59.28 ***	...	59.14 ***	60.07 ***	59.85 ***		
p-value	...	p < 0.0010	...	p < 0.0010	p < 0.0001	p < 0.0001		
p_75-p_25	...	23.72 ***	...	19.76 ***	15.54 ***	15.53 ***		
p-value	...	p < 0.0010	...	p < 0.0010	p < 0.0001	p < 0.0001		
R ²	0.3924	0.4328	0.5486	0.5888	0.5786	0.5788		
N	258,060	258,060	258,060	258,060	85,537	85,537		

Table 1.7 reports estimates of the following equation:

$$\text{Borrowing Cost}_{ibt} = \text{CPrate}_t - \text{RR}_{ibt} = \beta_1 \text{on loan } \%_{bt} + \beta_2 \text{loan size}_t + \beta_3 \text{rating}_{bt} + \beta_4 \text{issue size}_b + \beta_5 \text{time since issue}_{bt} + \beta_6 \text{floating rate}_b + \beta_7 \text{rule144a}_b + \delta_t + \kappa_b + \lambda_{\text{broker}} + \epsilon_{ibt}$$

where the on loan % is the percentage of daily inventory already lent, and loan size is the total number of bonds lent in thousands of bonds (that is, the loan value in \$ millions). Rating is the bond's S&P rating at the time of the loan (where AAA is given a value of 1, D is given a value of 22, and all intermediate ratings are given consecutive values between 1 and 22). Issue size is the size of the initial bond offering (in \$100 millions). The time since issue variable is the time since the bond was issued (in years). The floating rate variable is a dummy variable equal to 1 if the bond pays a floating rate coupon and 0 if the bond has a fixed rate coupon. The Rule 144a variable is a dummy variable equal to 1 if the bond was issued under SEC Rule 144a and 0 otherwise. d represents a set of dummies for each trading day in the sample. k represents a set of dummies for each bond CUSIP in the sample, and l are a set of dummies for each unique borrower in the sample who borrows 100 or more times during our sample period. Subscripts i, b, and t correspond to loan i, bond b, and day t. There are 62 brokers that borrow from the lender during the period covered by the regression, 38 make 100 or more loans and 24 make less than 100 loans. Standard errors are heteroscedasticity-robust and t-statistics are reported in parenthesis. The data is from the FISD Corporate Bond, Proprietary Bond Inventory, and Proprietary Bond Loan databases. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is April 1, 2005 through December 31, 2007. * significant at 0.10; ** significant at 0.05; *** significant at 0.01.

given the discussion of Figure 1.3 above. The difference in coefficients between columns (5) and (6) is due to loans where the on loan percentage is greater than 70%. The coefficient for on loan percentage alone in column (6) is 4.44, comparable to 4.38 in column (4). The coefficient for on loan % greater than 70% is an incremental 5.09. There is no statistically significant differential effect for on loan percentage greater than 50

Loan size is negative and significant for all six specifications showing that the larger the loan, the lower the borrowing cost. The magnitude of the coefficient is economically large and similar across all regression models, ranging from -1.36 to -2.16. This means that adding 1,000 bonds to loan size decreases borrowing costs by 1.36 to 2.16 bps.

The coefficients on bond ratings are positive and significant in all six specifications. The lower rated the bond, the higher the borrowing costs. The magnitude of the estimate is larger when we include bond-specific controls. For the specification in column (4), with broker and CUSIP dummies, the estimates imply that a full letter downgrade raises borrowing costs by 9.69 bps (three times the regression coefficient estimate of 3.23). The magnitude is even larger when we restrict the sample to high-yield bonds in column (5) where a full letter downgrade raises borrowing costs by 16.56 bps.

The estimated coefficient for issue size is small, but positive and significant for the first two specifications. Issue size must increase by \$300 million for borrowing costs to increase by 1 bp. The coefficient on time since issuance is positive and significant in the two specifications, implying that the longer a bond is outstanding, the higher the borrowing cost. For every year a bond is outstanding, the borrowing cost increases by 0.7 bps.

The last two bond characteristics from Table 1.2 are indicators for floating rate bonds and for whether a bond is Rule 144a. The estimates imply that fixed rate bonds are almost 6 bps more expensive to borrow than floating rate bonds and that the borrowing costs for Rule 144a bonds are about 3 bps more expensive.

The identity of the borrower who initiates a loan is also important in determining borrowing costs. The proprietary database only allows us to observe the borrowing broker (or hedge fund); it does not allow us to determine the final party undertaking the short sale transaction.

In the database each bond is lent to one of 65 unique brokers who then either delivers the bonds to their own institutional and retail clients for short selling or keeps them for its own account.

The specifications in Table 1.7 columns (2), (4), (5) and (6) include 38 broker dummies, each of which borrowed 100 or more bonds from April 2005 to December 2007. For all specifications, we can reject the hypothesis that all broker coefficients are zero. The difference between maximum and minimum broker coefficients and the 75th and 25th percentile broker coefficients are also reported. In column (4), the “best” broker receives borrowing costs 59 bps less than the “worst” broker. This means that on the same day for the same CUSIP and loan size, the lowest cost broker is able to borrow at a rate 59 bps lower than the rate for the highest cost broker. This difference is considerably larger than the average borrowing cost of 33 bps as reported in Table 1.5. The difference between the 75th and 25th percentiles in column (4) is 20 bps. All differences are statistically significant.¹⁷

Table 1.8 further explores whether some brokers obtain lower borrowing costs. We examine all days where two or more brokers borrow the same bond. Requiring that a broker “compete” with another broker on the same day at least 100 times restricts us to consider 26 brokers. For this group, we rank each broker’s “performance” on that day for that bond by evaluating whether they received a lower, higher, or the same borrowing cost as another competing broker.¹⁸ Those results are summarized in Table 1.8 and show that some brokers receive consistently lower borrowing costs. We ran two sets of “competitive” races per borrower. One set was between two brokers only; the second set was between three or more brokers. The top-rated broker received the lowest borrowing cost for any given day and bond 92.5% of the time when there were two brokers and 78.9% of the time when there were three or more brokers for the same bond on the same day.

¹⁷While we find differences in costs across borrowers, Kolasinski *et al.* (2013) find significant differences in costs across lenders for the equity lending market.

¹⁸The last line of Table 1.8 with Broker ID “Remainder” is a summary line that consolidates the other 39 brokers as one competitor. The competitive race results in columns 5-8 represent contests between the combined 39 brokers and any of the 26 brokers above. It does not include contests that the 39 remaining brokers have with each other.

Table 1.8: *Competitive Races Between Brokers*

Broker ID	# of Loans	# of Bonds Borrowed	Total Lending Fees Paid	2 Broker Races		3+ Broker Races	
				# Competitive Races / Wins	Percentage of Wins	# Competitive Races / Wins	Percentage of Wins
A	40,994	41,714,394	\$13,090,277	6,478 / 5,993	92.5% ***	1,561 / 1,231	78.9% ***
B	2,595	2,075,390	\$63,271	546 / 479	87.7% ***	164 / 127	77.4% ***
C	12,773	42,175,029	\$6,994,331	1,780 / 1,423	79.9% ***	719 / 476	66.2% ***
D	5,816	24,283,893	\$7,006,490	790 / 622	78.7% ***	361 / 239	66.2% ***
E	11,132	28,620,944	\$4,632,767	1,668 / 1,261	75.6% ***	574 / 328	57.1% ***
F	1,755	7,944,398	\$2,764,846	257 / 189	73.5% ***	118 / 59	50.0% ***
G	4,190	12,189,596	\$3,043,453	556 / 406	73.0% ***	252 / 151	59.9% ***
H	35,259	90,905,675	\$22,738,738	3,444 / 2,128	61.8% ***	1,244 / 532	42.8% ***
I	972	2,639,919	\$189,152	125 / 76	60.8%	55 / 25	45.5%
J	2,209	5,404,871	\$1,420,770	345 / 194	56.2%	129 / 46	35.7%
K	3,767	11,597,273	\$9,623,957	366 / 195	53.3%	183 / 68	37.2%
L	3,011	8,902,543	\$2,063,986	399 / 206	51.6%	184 / 77	41.8%
M	11,762	26,925,386	\$3,697,178	1,444 / 695	48.1%	584 / 226	38.7%
N	21,355	38,973,071	\$10,798,318	2,323 / 976	42.0%	861 / 332	38.6%
O	5,428	6,060,740	\$1,565,975	503 / 177	35.2%	195 / 50	25.6%
P	87,613	84,175,234	\$40,545,670	6,992 / 2,399	34.3%	2,056 / 518	25.2%
Q	6,633	18,783,575	\$7,711,792	645 / 217	33.6%	318 / 83	26.1%
R	14,339	23,432,851	\$15,138,170	1,404 / 403	28.7%	607 / 144	23.7%
S	43,344	22,503,842	\$4,825,499	2,951 / 839	28.4%	1,109 / 241	21.7%
T	2,662	1,787,228	\$260,718	287 / 41	14.3%	136 / 19	14.0%
U	2,244	535,303	\$88,309	237 / 29	12.2%	139 / 4	2.9%
V	10,638	3,875,297	\$858,456	996 / 113	11.3%	395 / 25	6.3%
W	14,407	5,641,386	\$1,195,550	1,464 / 94	6.4%	442 / 26	5.9%
X	2,646	1,213,004	\$253,763	309 / 19	6.1%	136 / 4	2.9%
Y	3,460	1,323,795	\$272,237	518 / 24	4.6%	175 / 11	6.3%
Z	11,813	5,726,357	\$1,701,179	1,577 / 54	3.4%	458 / 29	6.3%
Remainder	4,934	11,640,326	4,764,471	682 / 291	42.7%	256 / 90	35.2%

Table 1.8 uses data from the Proprietary Bond Loan database and compares broker borrowing costs by examining all days where two or more brokers borrow the same bond. 26 identified brokers have at least 100 competitive races. Success in 2 Broker and 3+ Broker Competitive Races is defined as having the lowest borrowing cost for a new loan in the same bond on the same day. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007. *** indicates Percentage of Wins that are significantly different than 50% and 33.33% at 0.01 one tailed probability for 2 and 3+ brokers, respectively.

The two winning percentages of the top-rated broker are both significant using the sign test. In fact, the top eight brokers all have winning percentages which are significantly greater than 50% at the 1% level when “competing” with one other broker and significantly greater than 33% when competing with two or more brokers. This success in the competitive races is not dependent on the number of loans or the amount borrowed by the borrower. Rank order correlations between placement in the competitive races and either the number of loans or the dollar amount of the bonds borrowed are not significant. Furthermore, the differences are not due to differences in the credit quality of brokers. Using broker bond ratings from S&P and 5-year CDS spreads from Markit, there is no significant relationship between either and broker rank in the competitive race.¹⁹ Thus it appears that differences in borrowing costs between borrowers reflect differences in market knowledge and abilities to negotiate borrowing costs.²⁰

To summarize, the borrowing cost regression results in Table 1.7 show that a smaller loan size, a higher percentage of inventory lent, and a lower bond rating lead to higher borrowing costs. These results hold for all specifications of the model, although the coefficients for on loan percentage are weaker when CUSIP dummies are included. Finally, the identity of the borrowing broker significantly influences borrowing costs, both in aggregate and when comparing loans for the same bond, regardless of the broker’s volume.

1.5.4 Borrowing Costs Around Credit Events

We next look at borrowing costs in the 30 days before and after significant credit events. The events we examine are bankruptcy filings and large credit rating changes. We define a large credit rating change as a movement of three or more S&P ratings, or one full letter or more, e.g. going from an A+ to a B+ or from a BB- to an AA-. There are 241 bonds in the inventory database of corporate bonds involved in a bankruptcy, representing 93 unique bankruptcies. However, only 88 bonds have lending activity during the period from 30 trading days before

¹⁹We only found ratings and CDS spreads for 15 and 17 of the 31 brokers, respectively.

²⁰Each unique broker’s identity is available to us from the proprietary database, although we are not allowed, for confidentiality reasons, to disclose it. The differences in borrowing costs are consistent with our perceptions of reputation.

until 30 trading days after the bankruptcy, which corresponds to 42 unique bankruptcies.

The average borrowing cost of these bonds for each of the 61 days is plotted in Figure 1.4. Since there are new loans for only 2.9 bankrupt bonds per day in the period -30 to +30 days around bankruptcy, we expand the sample by including old loans (which, as we discussed above, are re-priced). This expands the number of bonds per day in Figure 1.4 to an average of 60. However, not all bonds have a loan outstanding for all 61 days. We have also done the analysis only on new loans and only on bonds that have loans for all 61 days. Although there are far fewer observations, the results are qualitatively similar.

Figure 1.4 shows that bond borrowing costs are high for the entire period from -30 days to +30 days, where Day 0 is the bankruptcy filing date. The average equally-weighted bond borrowing cost for firms that file bankruptcy is 173 bps during the 30 days before filing. This is substantially greater than the average 33 bps reported for all new loans in Table 1.5 and indicates that these bonds are difficult to borrow before bankruptcy. After bankruptcy, bond borrowing costs increase further to an average of 245 bps for the 30 days after the filing. Thus, the borrowing costs indicate that short sellers identify firms in financial distress prior to bankruptcy, but the bankruptcy filing is not completely anticipated since borrowing costs rise after that date.

In Figure 1.5, we report a similar analysis for large bond downgrades and upgrades. There are 292 full-letter upgrade events on bonds in the inventory, covering 281 unique bonds as some bonds have multiple upgrades. Our loan data covers 125 of these events, which correspond to 122 unique bonds. The plot for these upgrade events shows that the average upgraded bond borrowing cost is close to the average for all bonds before the upgrade and does not vary much after the rating change. The average borrowing cost for the 30 days before the upgrade is 29.9 bps, and the average borrowing cost for the 30 days after the upgrade is 32.1 bps.

The bond borrowing costs for downgrades are much lower than those for bankruptcies, but are above the average of all bonds and increase after the downgrade. There are 381 full-letter downgrade events during our time period on 356 unique bonds. The data covers 206 of these events on 193 bonds. The average borrowing cost for the bonds involved in a full-letter

Figure 1.4: *Borrowing Costs Around Bankruptcies*




Figure 1.4 plots borrowing costs around bankruptcy filings. Data is from the Proprietary Bond Inventory and Loan databases. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. There are 241 bonds in the inventory database involved in a bankruptcy, representing 93 unique bankruptcies. However, only 88 bonds have any lending activity (either new or existing loans) during the period from 30 trading days before until 30 trading days after the bankruptcy. These bonds correspond to 42 unique bankruptcies.

Figure 1.5: Borrowing Costs Around Credit Events




Figure 1.5 plots borrowing costs around credit rating changes. Data is from the Proprietary Bond Inventory and Loan databases. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. We define a large credit rating change as a movement of three or more S&P ratings, or one full letter or more, e.g. going from an A+ to a B+ or from a BB- to an AA-. There are 292 full-letter upgrades on bonds in the inventory database, which correspond to 281 unique bonds. Our data covers 125 of these upgrades, corresponding to 122 unique bonds. There are 381 full-letter downgrades during our time period on 356 unique bonds. Our data covers 206 of these downgrades, corresponding to 193 bonds.

downgrade is 38.4 bps in the 30 days before the downgrade and 52.3 bps in the 30 days after the downgrade. It is important to remember that all downgrades are included, including those between investment grades, i.e. from an A+ to a BBB+, and thus all downgrades do not signal financial distress.

Thus, Figures 1.4 and 1.5 show that bankruptcies and large credit downgrades increase a bond's borrowing cost, while large credit upgrades do not decrease a bond's borrowing cost. The reasons for these changes in borrowing costs around bankruptcies and downgrades are difficult to discern. Although not shown, for bankruptcies, the supply of bankrupt bonds in inventory falls by 12.5% in the 30 days after bankruptcy is announced compared to the 30 days before. That is, our lender has 12.5% less bonds to lend. The amount lent also falls after bankruptcy by 23.9% comparing 30 days after to 30 days before. Hence, while the cost of borrowing bankrupt bonds increases, we cannot definitively rule out that the reason is a decrease in supply versus changes in demand. For downgrades, however, the amount of inventory actually increases by 4%, and the quantity of bonds borrowed increases by 6%. Since the cost of borrowing also goes up after downgrades, we can infer that there is increased demand for borrowing these bonds.

1.6 Relationship Between Bond and Stock Shorting

1.6.1 Matching Bonds and Stocks

We next investigate how the market for shorting corporate bonds is related to the market for shorting stocks. If the purpose of borrowing securities is to short the firm, we expect the two markets to be integrated. As mentioned above, given the priority of claims, the stock of a firm should lose its value before the debt, suggesting that investors who wish to express a negative view about the firm may prefer to short stocks. This is consistent with loan activity by our proprietary lender who made 367,751 bond loans and made 7,241,173 stock loans during our sample time period.²¹

²¹The stock loan database has some borrowing costs that are suggestive of data errors. In particular, there are stock loans that occur at large negative borrowing costs, implying that borrowers were being paid

To understand how the market for shorting corporate bonds is related to the market for shorting stocks, we matched each firm's bonds to its corresponding common stock. We match the first 6 digits of the bond CUSIP to the first 6 digits of the common stock CUSIP. This match was not complete since many of the bonds in the dataset are subsidiaries or private firms and thus have 6 digit CUSIPs which do not directly correspond to a common stock CUSIP. To add the subsidiary bonds (which may have a different 6 digit CUSIP), we hand matched the remaining bonds using SEC filings and CUSIP.com. To avoid potential biases that hand matching may introduce, we analyze our results for both methods separately, i.e. those that were matched with 6 digit CUSIPs versus those which were hand matched. There are 15,493 bond CUSIPs in the inventory file. We were able to match 11,591 bond CUSIPs, 5,997 using the 6-digit CUSIP match, and an additional 5,594 were matched by hand. We found no significant differences in results between the two subsamples.

Another matching problem is that there are many firms with multiple bond issues. For instance, there are 124 different GM bonds in inventory, and we want to relate the borrowing costs of all of those bonds to the cost of borrowing GM's common stock. We group all issues of bonds together for this analysis. The reason we group in this way is that for any given day, within the same firm, bond rebate rates are close. When different bonds from the same firm have a new loan on the same day, the median absolute value of the difference in bond borrowing costs is zero bps. This means that for more than half the firm-day observations, the borrowing costs are the same for all bonds of a given firm. Furthermore, the 75th percentile of this distribution is only 4 bps.

As a result, for our bond and stock analysis, if a firm has more than one new bond loan on a given day, we aggregate the borrowing costs across all bonds and all new loans by computing a value-weighted median borrowing cost. Likewise, for stocks we take the median

a significant amount to borrow the stock. We eliminate the 53,481 stock loans with borrowing costs below -5%. Also, there are some stock loans at high borrowing costs which would require that a significant amount of the borrower's collateral would be consumed by lending fees. We eliminate 4,883 stock loans where the borrowing cost is greater than or equal to 6% if that is the most expensive loan for a stock on a given day and the borrowing cost for the next most expensive loan for that stock on that day is no more than 3 times the general collateral rate.

stock borrowing cost for new loans weighted by shares lent. Hence the unit of observation in this section is a matched firm-day, corresponding to a firm's median value-weighted borrowing cost across bonds and the firm's median value weighted stock borrowing cost.

There are 336,449 bond loans which are matched to a stock in our sample. This represents 91.5% of all bond loans. There are 2,304,127 stock loans which are matched to a bond in our sample, which is 31.8% of all stock loans. Thus, it is much more likely that a bond loan occurs in conjunction with a stock loan, than vice versa.

1.6.2 Comparison of Bond and Stock Borrowing Costs

Figure 1.6 plots the equally-weighted distribution of stock loan borrowing costs over time by quintile for matched stock loans. It is comparable to Figure 1.2, which plots a similar time series for bond borrowing costs. Comparing Figure 1.6 to Figure 1.2 shows that the 20th and 40th percentiles of bond and stock borrowing costs are similar. However, the 60th and 80th percentiles of stock borrowing cost are less expensive than bonds until mid 2006. At that point, there is a compression in the distribution of stock borrowing costs generated by the large drop at the top quintiles. This compression occurs at the same time as the compression in bond borrowing costs discussed extensively above and seen in Figure 1.2. After mid 2006, stock and bond borrowing costs are similar at all quintiles.²²

To compare borrowing costs for stocks and bonds within a firm, it is necessary to impose the restriction that stock and bond loans occur on the same day. This restriction reduces our sample to 238,940 bond loans and 316,216 stock loans, corresponding to 113,548 matched firm-days.

For most firms, there is a fixed link between bond and stock borrowing costs. In particular, the difference between stock and bond borrowing costs is one of six distinct values: -10 bps, -5 bps, -1 bp, 0 bps, +35 bps, and +40 bps for 75.5% of the firm-days in the matched sample. This is seen in Figure 1.7, which plots the percentage of loans in the matched sample in each

²²The extreme values of stock borrowing costs are significantly greater than those of bond borrowing costs throughout the period. For example, the borrowing cost for the 35th most expensive stock loan is still three times the most expensive bond loan shown in Table 1.6.

Figure 1.6: *Equally-Weighted Monthly Distribution of Stock Loan Borrowing Costs*




Figure 1.6 plots the equally-weighted borrowing cost quintiles monthly from the Proprietary Stock Inventory and Loan databases over time. The time period analyzed is January 1, 2004 through December 31, 2007. Only stocks that are matched to bonds are used.

of these six categories over time.

The largest category in Figure 1.7 is new bond loans with borrowing costs 1 bp below new stock loans. For the matched loans, this category accounts for an average of 39.4% of observations. This 1 bp difference is impossible to explain if bond and stock borrowing costs are not related. There are two other major fixed borrowing cost differences where bonds are cheaper to borrow than stocks. They are -5 bps and -10 bps, which average 14.0% together.

The second largest category of fixed borrowing cost differences is bond loans with borrowing costs 35 bps more expensive than stock loans. This relationship changes, however, during our sample period. For the period from December 2004 until March 2006, the mean number of observations in this category is 22.8%. For the period from April 2006 until December 2007, the mean number in this category is 6.7%. This drop is clearly shown in Figure 1.7 and April 2006 appears to be a fundamental shift in the pricing relationship between bond and stock loans. Moreover, the +40 bps category, where bond loans are 40 bps more expensive than stocks, disappears by June 2006. These changes coincide with the reduction in the premium charged for small bond loans in April 2006, as described in Section 1.5.

There is a category that expands dramatically after March 2006: bond and stock loans that have the same borrowing cost. Before March 2006, the average percentage of matched loans in this category is 0.2%, while after March 2006, it is 7.1%. The percentage of loans in this category expands exactly when the percentage of loans in the +35 bps category decreases, although not by equal amounts. The -1 bp category also increases after March 2006.

While Figure 1.7 graphs the differences in bond and stock borrowing costs, Table 1.9 considers these differences by credit quality and compares expensive bond and stock loans. The first part of Table 1.9 confirms Figure 1.7 and shows that 63.7% of loans in the matched sample have borrowing costs within 10 bps of each other. There is no significant difference between investment grade firms and high yield firms which have 63.6% and 64.0% of matched loans within 10 bps of each other respectively for the sample period.

For expensive matched loans the borrowing costs are not close to one another; the stock loan is more likely to be expensive. In particular, only 1.3% of all matched bond loans are over

Figure 1.7: Bond and Stock Borrowing Cost Differences




Figure 1.7 examines differences in borrowing costs between matched corporate bonds and stocks. Data is from the Proprietary Loan databases for the overall period and by year. Only bonds that can be matched to a unique stock for a given loan and date are included. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with “KNOCK”, “REVERSE”, or “EQUITY” in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007.

Table 1.9: Bond and Stock Borrowing Relationship

	Investment		High Yield and
	All Firms	Grade	Unrated
N	113,548	72,051	41,497
% bond > stock	28.8%	28.9%	28.5%
% bond = stock	3.3%	3.4%	3.1%
% bond < stock	68.0%	67.7%	68.4%
% bond and stocks within +/- 10 bps	63.7%	63.6%	64.0%
% bond > stock by more than 10 bps	25.6%	25.8%	25.3%
# bond > 75 bps	3,131	1,114	2,017
% of all matched loans	2.8%	1.5%	4.9%
# stock > 75 bps	7,626	4,720	2,906
% of all matched loans	6.7%	6.6%	7.0%
if bond > 75 bps, % stock > 75 bps	11.6%	5.7%	14.9%
if stock > 75 bps, % bond > 75 bps	4.8%	1.3%	10.4%
# bond > 100 bps	1,425	184	1,241
% of all matched loans	1.3%	0.3%	3.0%
# stocks > 100 bps	7,015	4,348	2,667
% of all matched loans	6.2%	6.0%	6.4%
if bond > 100 bps, % stock > 100 bps	15.6%	9.8%	16.4%
if stock > 100 bps, % bond > 100 bps	3.2%	0.4%	7.6%

Table 1.9 examines differences in borrowing costs between matched corporate bonds and stocks. The unit of observation are CUSIP-days aggregated by loans by firms within a day. Data is from the Proprietary Loan databases for all firms and by credit status. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007.

100 bps, while 6.2% of matched stock loans are over 100 bps. Furthermore, if a bond borrowing cost is more than 100 bps, 15.6% of matched stock borrowing costs also costs more than 100 bps. For the converse, if a stock borrowing cost is more than 100 bps, only 3.2% of the matched bond borrowing costs are over 100 bps. This means that it is more common for stocks to be hard to borrow (as measured by borrowing costs) than it is for bonds. Furthermore, when a bond is harder to borrow, the stock is more likely to be as well. While not definitive, this pattern is consistent with stock borrowing activity leading bond borrowing activity.

These aggregate differences in stock and bond loan percentages are largely driven by high yield bonds. Only 0.3% of all matched investment grade bond loans are over 100 bps, while 3.0% of all matched high yield bond loans are over 100 bps. For matching stock loans there is little difference between investment grade and high yield (6.0% vs. 6.4%, respectively). In addition, for investment grade bonds if a stock borrowing cost is more than 100 bps, only 0.4% of bonds are greater than 100bps, whereas for high yield bonds, the corresponding number is 7.6%. This indicates that loan costs are more likely to be linked for high yield securities, which is consistent with high yield bonds serving as substitutes for stocks when stocks are expensive to borrow. These patterns also hold for borrowing costs greater than 75 bps. Finally, since borrowing costs for stocks are insensitive to investment grade status, while borrowing costs for bonds are, supports credit quality as an important determinant of borrowing activity for bonds.

To summarize, there are three main results on the relationship between bond and stock market shorting. First, most bond and stock loans for the same firm differ by one of six fixed amounts, which do not depend on the day of the loan. For example, the most common differences in borrowing costs between bonds and stocks, which are -1 bps and +35 bps, constitute 55.1% of the matched sample. Second, bond borrowing costs are very close to stock borrowing costs for most matched loans. For matched bond and stock loans from the same firm on the same day, 63.7% of the borrowing costs are within +/- 10 bps of each other. Finally, if neither the bond nor the stock is hard to borrow, they are priced very similarly. However, on a day when a stock is expensive to borrow, bonds from the same firm are usually

not, and vice versa. This suggests that for low levels of borrowing costs bond and stock lending markets are similar, but when borrowing costs are high the bond and stock lending markets are fragmented.

1.7 Returns to Shorting Bonds

In the last two sections, we calculated bond borrowing costs, described their cross-sectional and time-series distribution, and examined some of their important determining factors. In this section, we perform similar analysis on the returns to shorting bonds. As mentioned above, we do not know if all borrowed bonds are necessarily shorted, but for the purposes of this section we assume they are. The literature on stock shorting that uses proprietary lending databases makes a similar, although usually unstated, assumption. The literature on shorting stocks infers that excess returns from highly shorted stocks imply the existence of private information among short sellers and/or borrowing constraints. We make the same inference for the market for shorting bonds.

To calculate bond returns over any holding period, it is necessary to have bond prices at the beginning and end of the period. Following the approach of Bao and Pan (2013) we match the proprietary databases of bond inventory and loans to the FISD TRACE database, which provides transaction bond prices. The number of bonds covered in TRACE increased during our sample period. This increase ostensibly extended TRACE’s coverage to all US corporate bonds by February 7, 2005. Even with universal TRACE coverage, there are difficulties in computing bond returns. (See Bessembinder *et al.* (2009) for the difficulty of working with bond returns in general and TRACE in particular.)

We calculate bond returns with the following formula²³:

$$\text{return} = \frac{\text{sale price} - \text{buy price} + \text{sale accrued interest} - \text{buy accrued interest} + \text{coupons paid}}{\text{buy price} + \text{buy accrued interest}}$$

In this formula, the return is computed from the point of view of a long holder of the bond.

²³This is a formula from Bessembinder *et al.* (2009) with a correction for a typographical error in that paper.

That is, the returns are positive if the bond prices increase. A short seller of the bond, therefore, benefits if the return is negative. In the formula, sale and buy prices are “clean”, meaning net of accrued interest, which is the way prices are reported in TRACE. In some databases bond prices are “dirty”, meaning they include accrued interest, and the above formula has to be modified appropriately.

Of the 10,293 bonds that are ever loaned in the bond loan database, 8,212 bonds have at least one TRACE price observation, and 8,033 have at least ten TRACE price observations. Since a bond must only be delivered to a buyer within three trading days after a short sale, a bond loan does not always occur on the same day as the linked trade. They can either be located first and then sold short, or sold short and then located within 3 days after the sale. Of the 367,751 bond loans during the sample period, 301,167 have TRACE prices both within three days before or after the initiation of the loan and three days before or after the loan’s termination.²⁴

The fact that bonds do not trade every day and that short sales may occur on different days than the bond loans complicates calculating holding-period returns. As a result, our approach to calculating monthly returns for a bond is not precisely over thirty days because the bond may not trade exactly one month apart. We compute a monthly bond return when a bond has a trade in two consecutive calendar months. If there is more than one bond trade in a calendar month, we use the price of the last trade in that month. If there are multiple bond trades on this day, we use the trade-size-weighted median price for the day. Following Bessembinder *et al.* (2009) we exclude bond trades that are cancelled, modified, or include commissions. An equally-weighted monthly portfolio return is then calculated by equally weighting the monthly returns of the individual bonds in the portfolio. We also calculate an issue-size value-weighted monthly portfolio return by weighting monthly returns by the bond’s issue amount. Weekly returns are calculated in a similar manner.

²⁴After February 7, 2005 when TRACE’s universal coverage became effective, 245,508 out of 277,220 bond loans have TRACE prices both three days before or after the initiation of the loan and three days before or after the loan’s termination.

1.7.1 Returns to Portfolios of Shorted Bonds

In Table 1.10, we form monthly portfolios of bonds sorted by either percent of inventory on loan or borrowing cost. Panel A reports the returns from taking long positions in portfolios of bonds based on the percentage of inventory lent as of the last day of the month. The first two rows of Panel A report the monthly returns for portfolios of bonds that are not lent as well as those that are. In addition for each month, we calculate on loan percentage quintiles and assign the lent bonds to one of five portfolios. We also construct portfolios of bonds in the 95th and 99th percentiles of the on loan percentage distribution. These portfolios are formed conditional on the bonds being lent; that is, e.g, the 95th percentile portfolio is only selected from the universe of lent bonds. We report four different one-month returns for these portfolios.

In column 1, we report the number of bonds in each portfolio. Quintile sizes are not exactly equal because some values of on loan percentage are identical. Column 3 reports the equally-weighted raw portfolio return, while column 7 reports equally-weighted excess portfolio returns. Columns 5 and 9 report issue-size value-weighted raw and excess portfolio returns.²⁵ We calculate excess returns by subtracting equally-weighted and issue-size value-weighted TRACE index returns from the corresponding portfolio's raw returns.²⁶

The results in Table 1.10 Panel A show that there is no significant difference in the raw or excess returns between portfolios of bonds that are not lent and those that are lent. In fact, the mean issue-size value-weighted excess return in column 9 for the portfolio of lent bonds is -0.02%. Moreover, Panel A does not support the hypothesis that bonds which have higher on loan percentage are more likely to have lower returns in the future. In fact, both

²⁵We calculated value-weighted returns several ways including using the bond price times issue size as the weight. This results in no significant differences relative to the discussion below.

²⁶It is customary to use the Lehman Brothers (now Barclays) Corporate Bond Index when calculating bond excess returns (see, e.g., Bessembinder *et al.* 2009 and Bao and Pan 2013). While we also used this benchmark, we calculated a separate TRACE bond index using corporate bond prices from TRACE that were also in our FISD sample. We do this for two reasons. First, the Lehman Index uses matrix pricing while our TRACE index uses transaction prices. Second, the Lehman Index is a single aggregate number and does not match as closely our sample, e.g., the Lehman Corporate Bond Index does not include high yield bonds, but we include them in our TRACE index, since they are in our sample.

Table 1.10: Monthly Returns to Long Bond Portfolio Positions

Panel A: Bond Portfolios Which Are Formed According To Percent of Inventory On Loan																					
Portfolio	# of Bonds with TRACE										Issue-size Value-weighted Excess Returns (Net of TRACE)										
	# of Bonds in Coverage in All Months		Equally-weighted Raw Returns		Issue-size Value-weighted Raw Returns		Equally-weighted Excess Returns (Net of TRACE)		Issue-size Value-weighted Excess Returns (Net of TRACE)		Mean	Std. Dev	Std. Dev								
	Mean	[1]	Mean	[2]	Mean	[3]	Mean	[4]	Mean	[5]				Mean	[6]	Mean	[7]	Mean	[8]	Mean	[9]
Not Lent	5,013.5	2,574.1	0.40%	0.70%	0.37%	0.69%	-0.04%	0.32%	-0.05%	0.17%	-0.05%	0.17%									
Lent	2,821.5	2,246.9	0.40%	0.98%	0.40%	0.96%	-0.04%	0.39%	-0.02%	0.17%	-0.02%	0.17%									
1st Quintile	564.8	478.4	0.37%	0.92%	0.37%	0.91%	-0.08%	0.48%	-0.05%	0.28%	-0.05%	0.28%									
2nd Quintile	564.3	466.8	0.39%	0.93%	0.38%	0.93%	-0.06%	0.43%	-0.03%	0.24%	-0.03%	0.24%									
3rd Quintile	564.3	454.1	0.40%	0.99%	0.39%	1.02%	-0.05%	0.43%	-0.02%	0.27%	-0.02%	0.27%									
4th Quintile	564.3	442.0	0.41%	1.03%	0.40%	1.02%	-0.03%	0.46%	-0.01%	0.26%	-0.01%	0.26%									
5th Quintile	563.9	405.6	0.47%	1.32%	0.48%	1.28%	0.03%	0.82%	0.07%	0.79%	0.07%	0.79%									
95th Percentile	141.5	93.9	0.44%	1.97%	0.47%	1.98%	0.00%	1.68%	0.06%	1.72%	0.06%	1.72%									
99th Percentile	57.7	35.4	0.37%	2.38%	0.43%	2.64%	-0.07%	2.21%	0.02%	2.51%	0.02%	2.51%									

Panel B: Bond Portfolios Which Are Formed According To Borrowing Cost																					
Portfolio	# of Bonds with TRACE										Issue-size Value-weighted Excess Returns (Net of TRACE)										
	# of Bonds in Coverage in All Months		Equally-weighted Raw Returns		Issue-size Value-weighted Raw Returns		Equally-weighted Excess Returns (Net of TRACE)		Issue-size Value-weighted Excess Returns (Net of TRACE)		Mean	Std. Dev	Std. Dev								
	Mean	[1]	Mean	[2]	Mean	[3]	Mean	[4]	Mean	[5]				Mean	[6]	Mean	[7]	Mean	[8]	Mean	[9]
All New Loans	2,360.9	1,937.3	0.43%	0.89%	0.43%	0.90%	-0.06%	0.40%	-0.03%	0.17%	-0.03%	0.17%									
1st Quintile	536.0	432.2	0.45%	0.93%	0.44%	0.92%	-0.05%	0.46%	-0.02%	0.23%	-0.02%	0.23%									
2nd Quintile	469.6	373.3	0.45%	0.90%	0.45%	0.91%	-0.04%	0.42%	-0.01%	0.22%	-0.01%	0.22%									
3rd Quintile	509.7	417.3	0.40%	0.92%	0.40%	0.92%	-0.09%	0.41%	-0.06%	0.20%	-0.06%	0.20%									
4th Quintile	451.7	380.5	0.38%	0.85%	0.38%	0.86%	-0.10%	0.44%	-0.06%	0.22%	-0.06%	0.22%									
5th Quintile	403.5	342.1	0.46%	0.96%	0.46%	0.94%	-0.03%	0.50%	0.00%	0.30%	0.00%	0.30%									
95th Percentile	163.3	134.0	0.56%	1.36%	0.56%	1.38%	0.07%	0.96%	0.10%	0.91%	0.10%	0.91%									
99th Percentile	26.4	19.5	0.81%	3.60%	1.12%	4.67%	0.33%	3.39%	0.67%	4.48%	0.67%	4.48%									

Table 1.10 uses the TRACE database and computes returns for portfolios of bonds that are borrowed. Equally-weighted and issue-size value weighted returns are computed for each month, both raw and excess (net of TRACE). Portfolio quintiles are calculated at the beginning of each period based on the set of bonds that go on loan in that period. Equally-weighted raw returns are the unweighted average of returns, defined by

$$\text{return} = \frac{\text{sale price} - \text{buy price} + \text{sale accrued interest} - \text{buy accrued interest} + \text{coupons paid}}{\text{buy price} + \text{buy accrued interest}}$$

Equally-weighted excess returns are the unweighted average of raw returns minus the TRACE portfolio return. The TRACE portfolio return is the return from holding a portfolio of all bonds in TRACE. The issue-size value-weighted raw returns are the average of raw returns, weighted by the bond's issue size. Issue-size value-weighted excess returns subtract the issue-size value-weighted TRACE portfolio return. Convertibles, exchangeables, unit deals, perpetual bonds, bonds with missing or nonsensical offering amount data, and all bonds with "KNOCK", "REVERSE", or "EQUITY" in their description are excluded. The time period analyzed is January 1, 2004 through December 31, 2007.

the equally-weighted and issue-size value-weighted returns for the 5th quintile, which has the highest amount lent, are larger than those for all of the other quintiles in columns 3, 5, 7, and 9. Across quintile portfolios, the equally-weighted portfolio excess returns in column 7, though mostly negative, are small, and the issue-size value-weighted portfolio excess returns in column 9 are all within 8 bps of zero. Finally, the standard deviations of all portfolios returns, both equally- and issue-size value-weighted, are much larger than the means. As a result, none of the excess returns are significantly different from zero or from each other.

In Panel B, we form monthly portfolios based on the borrowing cost of the bonds. The first row of the Panel reports returns for all new loans. Each bond is then assigned a borrowing cost equal to the borrowing cost of the last new loans in the month, median-weighted by loan size. Then, for each month we calculate borrowing cost quintiles and assign bonds to one of the five portfolios. As in Panel A, we report one-month returns for these portfolios as well as for portfolios that include only bonds in the 95th and 99th percentiles of borrowing costs. Panel B has fewer observations than Panel A because it includes only bonds with new loans, whereas Panel A includes bonds with existing loans.

The results in Panel B do not support the hypothesis that bonds which are more expensive to borrow are more likely to have lower returns in the future. The 95th and 99th percentile portfolios have the highest borrowing costs, but they also have the highest average returns across all measures. Furthermore, the returns for the quintiles are not monotonic. Overall, the results in Panel B parallel those in Panel A: there are no significant results for any of the portfolios or any of the differences between the portfolios.

Table 1.10 shows that none of the portfolio returns or differences in Panels A or B are statistically significant. That is, neither the bond's on loan percentage nor the borrowing cost predicts future returns. Although not shown, we also calculated one week, two week, and three-month returns for all of the portfolios in Table 1.10. In no instances were any of the excess returns significantly different from zero. In addition, we also did the analysis in Table 1.10 split into investment-grade or high-yield bonds. Neither of those results are significantly different from zero, nor are they statistically different from each other.

1.7.2 Profitability to Short Sellers of Corporate Bonds

Table 1.10 indicates that shorting portfolios of bonds with high on loan percentage or high borrowing costs are not strategies that yield abnormal returns to short sellers. These results are based on shorting portfolios of bonds that are already highly shorted. They may indicate, but do not accurately measure, whether short sellers made money on their short positions. To evaluate the profitability of actual short trades, we must know the period the short position was held, and we must net out the borrowing costs and the overall movements in the bond market. The bond loan database, which has the start and end date of bond loans and their borrowing costs, allows us to undertake this analysis.

To calculate short sellers' profitability, we compute a return on capital net of coupons paid, accrued interest, and borrowing costs. We assume that the beginning and ending dates of a short position are the same as the beginning and ending dates of a bond loan. Since corporate bonds do not necessarily trade every day, we take as the starting price the TRACE price closest to the loan's actual start date in the period three trading days before until three trading days after the loan's initiation. The ending price is computed analogously. If there are multiple trades in one day, we take the trade-size-weighted median price of all trades that day.

Loans where the nearest trades are more than three days removed from either the loan start or end date are eliminated. We also eliminate loans where the starting and ending dates are matched to the same TRACE trade. This can occur if the loan is short term and there is only one reported TRACE trade during the time period from three days before the initiation until three days after loan termination. The profit from each loan, net of borrowing costs, accrued interest, and coupon payments, is then summed to obtain aggregate short sellers' profits over some period. This amount is then divided by the average capital invested during that period. Average capital invested is the summed daily par value of new and old outstanding loans divided by the number of days in the time period. Thus, the net return on capital is calculated as total net profit divided by average capital invested over a time period.

As an example, for the entire four-year period, the total profit assuming all borrowed bonds were shorted is -\$2.4 billion, which is a loss for short sellers. The borrowing cost for all loans

over the same period totaled \$112 billion. The average amount of bond loans outstanding per day is \$12.4 billion.²⁷ Thus, the average monthly return over the four-year sample period is -48 bps. This is consistent with positive monthly returns to long portfolios of shorted bonds in Table 1.10. For example, in Panel A, the raw portfolio returns for equally-weighted and issue-size value-weighted for all lent bonds are both 40 bps, and in Panel B, the comparable returns for all new loans are 43 bps. These values do not account for the average 2.8 bps monthly borrowing cost.

We next evaluate short seller profits by several loan characteristics, including loan size, duration, and borrowing cost. Loan size and duration do not substantively change the result reported above, but borrowing costs appear to be responsible for some variation in short seller profits. The return on capital for loans where the borrowing cost is greater than 100 bps is substantially lower than the return on loans where the borrowing cost is less than 100 bps. The return on capital is -123 bps per month for the more expensive loans and -46 bps per month for the less expensive loans. Even though borrowing costs are higher for the more expensive loans, they only account for 31 bps of the difference. This finding of larger losses for high borrowing cost loans parallels the finding of high positive returns for the 95th and 99th borrowing cost portfolios in Table 1.10.

Table 1.10 shows that portfolios formed on the basis of bond shorting activity do not earn significant excess returns. Examining realized profits from the actual short trades indicates that short sellers do not have private information. In fact, the average monthly return for short sellers is negative and almost the opposite of the returns from holding the bond market. This result is consistent with short selling being used as a hedging activity with short sellers paying for the hedge.

²⁷This number differs from the average daily par value of bonds on loan in the lender inventory in Table 1.1 because we only compute profits when we have both beginning and ending TRACE prices, and the loan must begin and end during our four-year period.

1.8 Relationship Between the Market for Shorting Bonds and the CDS Market

Rather than shorting a bond, another way for an investor to profit from a bond price decline is to purchase a credit default swap. This is similar to a stock investor purchasing a put. Unlike the options market for equities, which is smaller in notional amount than the stock market, the notional amount of the CDS market has become larger than the market value of corporate bonds. In mid 2009, the par value of corporate bonds was \$6.8 trillion, while the notional principal amount of CDS on corporate debt was \$12.1 trillion.²⁸

There is a documented link between shorting stocks and the equity options market. Many dealers who write equity puts hedge their positions by shorting stocks. There is also a link between option put-call parity and shorting constraints in the stock market (see, for example, Figlewski and Webb (1993) and Ofek et al., (2004)).

We use Markit as the source for the CDS data. Markit collects data from various financial institutions, inter-dealer brokers, and electronic trading platforms. The data consist of daily CDS spreads for reference securities. Each CDS contract is assigned a REDCODE number by Markit, which we then map to individual bond CUSIPs. Because of cross-default provisions, CDS contracts can correspond to more than one bond for any given firm. As a result, we ultimately match individual CDS to multiple bonds based on the first six digits of the bond CUSIPs.

Of the 15,493 bonds in the lender's inventory, we are able to match 7,033 (45.4%) to a CDS. The percentage of bonds lent with a CDS is higher: of the 10,293 bonds ever lent, 5,540 (53.8%) had a corresponding CDS at some point during our sample period. Furthermore, of the 367,751 new loans in the sample, 77.8% are of bonds with CDS. Thus, inventory bonds matched with CDS are more likely to be lent and constitute a much larger fraction of new loans. This suggests that there are common factors that determine which bonds have CDS

²⁸Corporate bond value is from SIFMA (2009), and CDS market value is from Depository Trust and Clearing Corporation (DTCC). This data is from 2009 because we are unable to find the breakout of corporate debt CDS during our sample period. The par value of outstanding corporate bonds in 2007 is \$7.2 trillion.

contracts and which bonds are lent.

We next use the bond characteristics in Tables 1.2 and 1.3 to examine the differences between bonds with CDS and those without. Lent bonds with CDS tend to be larger and have much higher credit quality than lent bonds without CDS. For example, 70.7% of the lent bonds with CDS are investment grade at the time of the loan, while only 50.4% of the lent bonds without CDS are. Examining loan size and duration in a manner similar to Table 1.4, we find that loans on bonds with CDS have similar sizes and median duration to those without. Importantly, the distribution of borrowing costs is almost identical between bonds with CDS and those without. For example, the mean and median equally-weighted borrowing costs for bonds with CDS are 33 and 19 bps, while they are 32 and 18 bps for bonds without CDS.

When we include an indicator for CDS in the borrowing cost regression presented in Table 1.7, we find that the presence of a CDS results in a significant increase in borrowing costs of 1.5-2.0 bps and has no discernible impact on the relative importance of the other factors we previously examined. This cross-sectional comparison does not imply that the presence of CDS causes higher borrowing costs; rather it may reflect the fact that bonds that are most likely to be shorted are more expensive to short and, are also most likely to have a CDS contract.

To look at the impact of CDS on borrowing costs, we next examine the introduction of a CDS contract. We plot the borrowing cost on individual bonds for the 30 days before and after Markit first lists a CDS on those bonds. This time series comparison holds fixed all other bond attributes unlike the previous cross-sectional comparisons. There are 332 new CDS introductions during our sample period, representing 1,589 lent bonds. 820 of these bonds have borrowing cost data in the 61-day window. There is no noticeable change in borrowing costs over this period. The average borrowing cost for the 30 days prior to the introduction of a CDS contract is 27.2 bps, while the average for the 30 days after is 25.3 bps. There is also no noticeable increase or decrease in the amount lent. Since Markit does not collect information from all dealers, there is the possibility that CDS contracts exist for some bonds before they first appear in Markit.

In summary, bonds with CDS tend to have higher loan activity than bonds without. In

addition, borrowing costs for loans with CDS are slightly higher than those without. Finally, the introduction of a CDS contract does not materially affect borrowing costs in the short term. All of these facts suggest that CDS are correlated with bond shorting, but do not substantially replace it.

1.9 The 2007 Credit Crunch

The Credit Crunch of 2007-2008 started in late July or early August 2007. The 3-month LIBOR-OIS rate, the difference between LIBOR and the overnight indexed swap rate, increased from 12.3 bps on August 1st to 40.0 bps on August 8th. By September 10th, the rate was 94.9 bps. The LIBOR-OIS rate is considered by many to be a “barometer of fears of bank insolvency.”²⁹ This increase occurred shortly after Bear Stearns announced they were liquidating two hedge funds investing in mortgage-backed securities on July 31, 2007. The Federal Reserve Bank took immediate action, reducing interest rates starting in mid-August 2007.

We examine the impact of this credit market turmoil on the market for borrowing corporate bonds. Although we do not have data from the entire Credit Crunch of 2007-2008 in our sample period, we are able to investigate the first six months, from July – December 2007. In particular, we investigate the impact of the 2007 Credit Crunch on lending activity, borrowing costs, and their determinants.

Figure 1.1 indicates that there was no distinguishable change in the number or par value of outstanding loans during the period July 2007 to December 2007 compared to the first half of 2007. Moreover, in Table 1.1, the average daily par value of bonds on loans in 2007 is \$14.4 billion and the percentage of inventory lent is 7.3%. Although not shown, the average daily par value of bonds on loan for the first and second half of 2007 are both \$14.4 billion, and the percentage of inventory lent changes from 7.1% to 7.5%. Both measures of loan activity are greater than those in 2006, but below the activity in 2005. The average characteristics of bonds lent reported in Tables 1.2 and 1.3 also do not change between the first and the second

²⁹ Alan Greenspan quoted in Thornton (2009).

half of 2007. The size and duration of lent bonds reported in Table 1.4 also do not change in any meaningful way even when dividing the sample by investment grade and high yield.

While the number of bonds lent, their characteristics, and loan size do not change in the second half of 2007, borrowing costs do. Figure 1.2 shows that following the March 2006 period, the distribution of borrowing costs is compressed. During the first half of 2007, the spread between the 20th and the 80th percentile borrowing cost averages 6 bps per month. In the second half, the spread expands and the average difference between the 20th and the 80th percentile is 28 bps per month. This increase in spread is due to both an increase and decrease in borrowing costs. As seen in Figure 1.2, the borrowing costs for the 80th percentile climbs from an average of 14 bps to 28 bps. At the same time, the borrowing cost for the 20th percentile falls from an average of 8 bps to 0 bps with three months showing negative borrowing costs.

This increase in volatility of borrowing costs does not affect the mean or median borrowing costs substantially. The mean equally-weighted and value-weighted borrowing costs for the first half of 2007 are 19 and 13 bps, respectively. The comparable mean borrowing costs for the second half of 2007 are 20 and 13 bps. The median equally-weighted and value-weighted borrowing costs behave similarly: they are 13 and 8 bps in the first half of 2007 and 13 and 7 bps in the second.

Borrowing costs becomes more volatile in the second half of 2007 because both components of borrowing costs, the commercial paper rate and the rebate rate, are more volatile. Although not shown, in the first half of 2007, only 6.7% of loans experienced commercial paper rate changes of at least 5 bps, while in the second half, 59.0% of loans experienced commercial paper rate changes of at least 5 bps. There is also a large increase in the percentage of loans that have a change in their loan rebate rate during the second half of 2007. For the first half of 2007 the percentage with rebate rate changes is 29.4%, while for the second half it is 63.4%. Thus, during the Credit Crunch of 2007 borrowing costs are reset more frequently than previously.

There are also a large number of loans with negative borrowing costs during the 2007

Credit Crunch period. This differs from the earlier sample period. During the second half of 2007, 17.6% of the loans have negative borrowing costs as compared to 3.4% during the first half of 2007. Interestingly, 90% of the loans with negative borrowing costs in the second half of 2007 occur on only 26 days. As discussed in Section 1.5, these negative borrowing costs may occur for two reasons. First, during this period short-term rates fell substantially below medium-term rates and, as a result, reported commercial paper rates may not reflect market conditions.³⁰ Second, these negative borrowing costs may arise if the lender is subsidizing borrowers to maintain collateral pools.

This large number of loans with negative borrowing costs is the reason why in Figure 1.3, where we plot borrowing costs against inventory lent, the lines for the July 2007 to December 2007 period are below the other plotted lines for most of the range. This is true for both investment grade and high yield bonds. The slope of the high yield line from this period continues to have a kink at 70%, and is similar to that of high yield lines from earlier periods.

Since the distribution of borrowing costs widens during the second half of 2007, we re-estimate the borrowing cost regression presented in Table 1.7 using only data from the second half of 2007. For all four specifications of the model, the coefficients for the second half of 2007 have similar magnitudes as the entire period presented in Table 1.7. All coefficients also remain significant.

In summary, the Credit Crunch of 2007 affected the market for borrowing corporate bonds primarily by widening the distribution of borrowing costs. The number of loans, the types of bonds lent, the size of loans, and the average borrowing costs all remained relatively stable in the second half of 2007 compared to the prior period. Thus, the change we document in March 2006 appears to be more of a structural change than that occurring during the Credit Crunch of 2007.

³⁰Our use of commercial paper rates as the market rate is not responsible for these negative borrowing costs. If we use the Fed Funds rate, loans with negative borrowing costs are still prevalent.

1.10 Conclusion and Implications

This paper presents the first complete examination of short selling for securities traded in an OTC market. It does this by utilizing a detailed proprietary database of corporate bond loans from 2004 to 2007. Short selling activity in corporate bonds is large and substantial. We estimate that short selling constitutes 19.1% of trading activity in the corporate bond market. This is about two-thirds of the percentage of short selling in equity markets.

Borrowing costs for corporate bonds are comparable to stocks and have become cheaper over time. The average borrowing cost of loans in the sample is 33 bps per year on an equally-weighted basis. There is a structural change in the pricing of corporate bond loans starting in April 2006 when the entire distribution of borrowing costs is compressed. As a result, the average equally weighted borrowing costs by 2007 is 19 bps.

Our analysis shows that bond borrowing costs are related to loan size, the bond's credit rating, and the lender's inventory. The importance of loan size on borrowing costs diminishes over our sample period. At the beginning, the median borrowing cost of a small loan is three times that of a large loan, while by the end, loan size is no longer priced. Credit rating and inventory remain important throughout our sample period. High yield bonds are more expensive to borrow than investment grade. Furthermore, borrowing costs increase substantially following bankruptcy and bonds with credit downgrades, not involving bankruptcy, also experience increases in borrowing costs.

A bond's credit quality also impacts the relationship between inventory and borrowing costs. When the lender has greater than 70% of its available bonds lent out, borrowing costs for high yield bonds rise sharply. In contrast, borrowing costs for investment grade bonds are not positively related to the percent of inventory lent. This holds before and after the mid 2006 structural shift.

Another factor impacting borrowing costs is the identity of the borrower. Broker effects are significant both in our regression analysis and in our competitive broker races. Moreover, our results do not indicate that this pricing differential is due to loan volume or the credit quality of the borrowing broker.

The market for borrowing corporate bonds is linked closely to the market for borrowing stock. 63.7% of the matched borrowing costs are within +/- 10 bps of each other, and 42.6% are within 1 bp. In fact, borrowing costs for 75.5% of the matched bond and stock loans for the same firm on the same day differ by one of only six distinct amounts. The distribution of stock borrowing costs also becomes compressed starting in April 2006, like the bond borrowing costs.

After examining returns to short selling, there is no evidence that, on average, bond short sellers have private information. Portfolios formed on the basis of corporate bond borrowing costs or levels of borrowing activity do not generate excess returns. Moreover, in aggregate, bond short-sellers do not realize a profit from their trades. In addition, borrowing costs have a very small influence on overall trade performance. Finally, there is strong evidence that short sellers, on average, pay a small cost for shorting corporate bonds.

We also investigate the impact of the CDS market on the market for borrowing corporate bonds tangentially. We find that bonds that have higher lending activity are more likely to have CDS contracts. Furthermore, we find that these bonds have small, but significantly higher borrowing costs (one or two bps) than bonds without CDS contracts. These differences are after controlling for other factors such as percent on loan, loan size, and bond rating. We conclude that the CDS market is correlated with bond shorting and is not a perfect substitute.

Finally, we examine six months of the 2007 Credit Crunch and compare it to the remainder of our period. We find that the volume and average pricing of corporate bond loans do not change. We do find, however, that the distribution of borrowing costs widen substantially during this period. There may be effects of the 2007 Credit Crunch on this market that do not appear until 2008, which our analysis does not capture.

An important caveat to our work is that we only examine data from one proprietary lender. We do not know with certainty if the patterns we document are particular to our lender or are market-wide. However, given the number of bonds and the size of lending activity by our lender, our analysis applies to a large portion of the market for shorting corporate bonds.

Our results speak to the larger literature on short sale constraints and their effects on asset

prices. That literature has argued that short sale constraints may generate mis-valuation. We find, at least for the sample of bonds covered by our lender, that while short selling is a large and important market activity, constraints, as measured by borrowing costs, do not have a measurable impact on corporate bond pricing. In addition, we find that shorting securities that are traded in an over-the-counter market is very similar to shorting exchange-listed securities, in particular stocks. Moreover, the fact that portfolios of heavily shorted bonds do not generate excess returns suggests that private information is not driving shorting activity. Finally, our results indicate that short selling is not responsible for the growth of the CDS market, nor is it being replaced by it.

Chapter 2

The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market¹

“...capital markets ... [are] replete with problems in the economics of information:
[e.g.] What over-the-counter transactions should be required to be reported?”

Stigler (1964)

2.1 Introduction

Trading in many financial securities takes place in environments with a great deal of transparency. For instance, nearly all U.S. stocks trade on exchanges with real-time reporting of pre-trade bid and ask quotes and post-trade transaction prices and volume. On the other hand, some securities, such as credit default swaps and collateralized debt obligations, have

¹with Paul Asquith and Parag A. Pathak

historically traded over-the-counter without even post-trade information about previous transactions. This paper studies the effects of a dramatic increase in transparency in the corporate bond market. We find that transparency significantly reduces price dispersion for all bonds and significantly reduces trading activity for some categories of bonds.

Corporate bonds were largely exchange-traded in the 1930s, which meant that post-trade prices and volume were publicly available (Biais and Green (2007)). After World War II, however, trading in this market migrated to over-the-counter, with private bilateral negotiations and no public reporting of transaction details. If investors wanted information on a bond's market price, they had a limited set of options: they could contact corporate bond dealers and ask for quotes or they could consult a vendor that provides estimated prices (widely known as "matrix prices").

The corporate bond market underwent a significant change in July 2002 when information on the prices and volume of completed transactions were once again publicly disclosed. FINRA (then the NASD) mandated transparency in the corporate bond market through the Trade Reporting and Compliance Engine (TRACE) program. FINRA required that all transactions in U.S. corporate bonds by regulated market participants be reported on a timely basis to TRACE. Corporate bonds are one of the world's largest over-the-counter markets with average transactions of \$4.2 trillion a year over this period (SIFMA 2013). FINRA then made this information transparent by publicly releasing (in their words "disseminating") the prices and volume of completed bond trades. Bond trade dissemination was Phased-in on four separate dates over a three-and-a-half year period. The increase in information available to market participants was so significant that it has been compared to the early 20th century introduction of stock market tickers and electronic screens for Treasuries (Vames 2003).

Studies of changes in market design for opaque markets are usually limited because, although data sometimes exists after the new design is implemented, there is rarely comprehensive information on market behavior beforehand. Prior to 2010, FINRA did not release any information regarding a bond's trades until after the dissemination Phase for that bond began. In 2010, however, FINRA released transactions data on all bonds, disseminated and not

disseminated, since the start of TRACE. With this newly released dataset, it is now possible to observe changes in the trading behavior of corporate bonds using data from periods before and after their trades are disseminated. Moreover, this comprehensive record of transactions makes it possible to provide a definitive account of the effect of TRACE across all categories of bonds.²

Even before FINRA released this historical transaction-level data, TRACE had become a template for how financial market reform and regulation should proceed. Difficulties evaluating the trading and value of over-the-counter instruments during the 2008 financial crisis motivated some to propose reforms inspired by TRACE. See, for example, Acharya *et al.* (2009) or the recommendations of the Squam Lake Group (French *et al.* 2010) which state:

Regulators should promote greater transparency in the CDS market for the more liquid and standardized index and single-name contracts. Consideration should be given to the introduction of a trade reporting system for these contracts similar to the TRACE system.

Furthermore, TRACE was expanded in March 2010 to include Agency-Backed Securities and in May 2011 to include Asset-Backed Securities (Shenn and Scheer 2009). In April 2013, the FINRA board approved a proposal, currently awaiting SEC approval, to publicly disseminate 144A transactions. There are also on-going efforts to mimic TRACE for European corporate bonds (Lerner 2011). Finally, Title VII of the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank) requires that swaps (including credit default swaps, interest rate swaps, collateralized debt obligations, and other derivatives) be traded and cleared centrally on exchanges. Dodd-Frank follows TRACE's definition of transparency by requiring public dissemination of post-trade transaction information regarding price and volume.

Proponents of TRACE argue that transparency makes the corporate bond market accessible to retail clients, enhances market integrity and stability, and provides regulators greater ability to monitor the market. They reason that with the introduction of transparency, price discovery

²Because of data limitations, earlier studies of TRACE focused on part of TRACE's implementation and, therefore, on particular subsets of bonds. For instance, Bessembinder *et al.* (2006) primarily study the effect of Phase 1 of TRACE on using data from the National Association of Insurance Commissioners. Edwards *et al.* (2007) and Goldstein *et al.* (2007) study the effect of Phase 2 on different samples of bonds.

and the bargaining power of previously uninformed participants should improve (NASD 2005a). This in turn should be reflected in a decrease in bond price dispersion and, if more stable prices attract additional participants, an increase in trading activity (Levitt 1999).

Opponents of TRACE object to mandatory transparency, saying that is unnecessary and potentially harmful. They argue that “transparency would add little or no value” to highly liquid and investment grade bonds since these issues often trade based on widely known US Treasury benchmarks (NASD 2006). They further argue that if additional information about trades was indeed valuable, then third-party participants would already collect and provide it, a view that dates back to Stigler (1964). Opponents also forecast adverse consequences for investors since, if price transparency reduces dealer margins, dealers would be less willing to commit capital to hold certain securities in inventory making it more difficult to trade in these securities. The Bond Market Association argued that the adverse effects of transparency may be exacerbated for lower-rated and less frequently traded bonds (Mullen 2004). Lastly, opponents saw TRACE as imposing heavy compliance costs, particularly for small firms who do not self-clear (Jamieson (2006)). Thus, opponents argue that market transparency reduces overall trading activity and the depth of the market. Not surprisingly, similar arguments for and against transparency have resurfaced in response to the recent introduction of the Dodd-Frank’s post-trade transparency requirements for swaps (Economist (2011)).

The implementation of TRACE and the release of the new database provide a unique opportunity to study the impacts of mandated transparency on market behavior. TRACE’s dissemination of price and volume data was not implemented on all bonds simultaneously. In July 2002, FINRA began collecting price and volume information for all corporate bond trades. On the same day, FINRA began dissemination of this information for just a subset of bonds. There were three other major “Phase-ins,” Phase 2, 3A, and 3B, expanding the set of bonds covered. Bonds were assigned to Phases using bond issue size, credit quality, and previous levels of trading activity. By February 2005, the price and volume of every corporate bond trade was publicly disseminated shortly after the trade’s execution. Thus, between 2002 and 2005, corporate bond market participants went from having little knowledge of trading

activity to having post-trade knowledge similar to equity market participants.

Our empirical strategy exploits these Phases to construct a before-and-after comparison between bonds subject to a change in transparency and bonds that are not. This difference-in-difference research design gives us the chance to avoid confounding the effects of transparency with unobserved shocks to the corporate bond market. Although our approach does not cover the first Phase of TRACE (where there is no TRACE data beforehand), it covers the remaining Phases, which represent 98% of bonds in the Phases.

The new database and our research design allow us to ask questions previous researchers were unable to investigate. Previous work on TRACE focused on imputed transaction costs. In this paper, we focus on TRACE's impact on market behavior, in particular its effect on trading activity and price dispersion. Earlier work also focused only on Phase 1 and/or Phase 2. This paper covers the entire TRACE implementation period, which is important because the types of bonds covered by TRACE in later Phases differ from that of earlier Phases by design. In particular, bonds covered in earlier Phases had large issue sizes and investment grade ratings, while bonds covered in later Phases of TRACE were bonds with smaller issue sizes and lower credit quality. These latter bonds are exactly the ones that opponents of TRACE warned would have the most adverse consequences.

We find that post-trade transparency of price and volume leads to a significant reduction in trading activity and price dispersion. Using our main measure of trading activity, trading volume/issue size, and our preferred difference-in-differences specification, we find a significant 15.2% reduction in trading activity in the 90 days after TRACE's introduction for the pooled sample across Phases 2, 3A, and 3B, i.e., the Phases where we can observe trading before and after dissemination. This result is driven primarily by Phase 3B bonds, which experience a significant 41.3% reduction in volume/issue size. Phase 3B bonds are largely bonds with credit ratings below investment grade that trade infrequently. Event studies show that the reduction in trading activity for Phase 3B bonds occurs immediately upon dissemination. In addition, these results are robust to alternative difference-in-differences specifications that vary time trends and control groups. The reduction in trading activity caused by TRACE is also seen

using several other measures of trading activity such as volume, number of trades, and average trade size.

Transparency also causes a significant reduction in price dispersion. We find a significant 8.5% reduction in within-day price standard deviation in the 90 days after TRACE's introduction for the pooled sample, and significant reductions for Phases 2, 3A, and 3B when examined individually. The largest reduction is for Phase 3B bonds, which is a significant 24.7%. The reduction for Phases 2 and 3A are also both significant at 7.3% and 6.5%, respectively. Event studies show that price dispersion falls immediately upon dissemination for all three Phases. In addition, these results are robust to trends and alternative assumptions about control groups. The reduction is also evident using other measures of price dispersion such as the difference between the maximum and minimum price on a given day and price standard deviation measures computed over longer time windows.

FINRA implemented TRACE in Phases because of concerns about the possible negative impact of transparency on thinly traded, small issue and low-credit rated bonds. Examining issue size across all Phases, we find that trading activity decreases more for large issue size bonds, but that the reduction in price dispersion is uncorrelated with issue size. Credit ratings, however, matter for both trading activity and price dispersion. High-yield bonds experience a large and significant reduction in trading activity, while the results are mixed for investment grade bonds. High-yield bonds also experience the largest decrease in price dispersion, but price dispersion significantly falls across all credit qualities. Therefore, the introduction of transparency in the corporate bond market has heterogeneous effects across sizes and rating classes.

Lastly, we report on a complementary analysis using transactions data from the National Association of Insurance Commissioners (NAIC) in an attempt to investigate the effect of TRACE on Phase 1 bonds. This analysis is inconclusive. However, since NAIC data reports the identity of the security dealer doing each trade, we analyze that data and show that TRACE causes a reduction in dealer volume and number of trades for the largest dealers for all Phases.

The rest of this paper is organized as follows. Section 2.2 presents additional background on TRACE and reviews the related literature. Section 2.3 describes the historical TRACE database and presents descriptive statistics. Section 2.4 describes our research design and the main results. Section 2.5 examines the robustness of our findings and reports on TRACE's effect on alternative measures of trading activity and price dispersion. In Section 2.6, we further explore heterogeneity in our findings based on ratings and issue size. Section 2.7 reports on an investigation of corporate bond trading using the NAIC database. The last section states our conclusions and discusses the implications of our findings.

2.2 TRACE and the Corporate Bond Market

2.2.1 History and Implementation of TRACE

The Trade Reporting and Compliance Engine (TRACE) was launched in July 2002, but it has its origins in the late 1990s when the Securities and Exchange Commission (SEC) reviewed issues related to price transparency in U.S. debt markets. After this review, the SEC asked the National Association of Security Dealers (NASD) to take three steps to enhance the transparency and the integrity of the corporate debt market: 1) adopt rules to report all transactions in U.S. corporate bonds to NASD and develop systems to receive and distribute transaction prices on an immediate basis; 2) create a database of transactions in corporate bonds to enable NASD and other regulators to take a proactive role in supervising the corporate debt market; and 3) create a surveillance program to better detect misconduct and foster investor confidence in the corporate debt market. The NASD changed its name to the Financial Industry Regulatory Agency (FINRA) in 2007.³

By January 2001, the SEC approved rules requiring NASD members to report all over-the-counter (OTC) market transactions in eligible fixed income securities to the NASD and mandating that certain market transactions be disseminated. NASD developed a platform, TRACE, to facilitate this mandatory reporting. The rules, referred to as the "TRACE Rules,"

³<http://www.finra.org/Industry/Compliance/MarketTransparency/TRACE/FAQ/P085430>, Last accessed: July 14, 2012.

are contained in the new Rule 6200 Series that replaced the old Rule 6200 Series, which governed the Fixed Income Pricing System (FIPS). FIPS started in April 1994 with reported transactions information on approximately 50 high-yield bonds at any point in time.

NASD's stated rationale for the introduction of TRACE was to bring transparency to the corporate bond market. Advocates of transparency anticipated that almost everyone would benefit because of increased market participation. For instance, SEC commissioner Arthur Levitt remarked that "this participation means more trading, more market liquidity, and perhaps even new business for bond dealers" (Levitt 1999). Doug Shulman, NASD's President of Markets, Services and Information stated as much: "By disseminating accurate and timely trading information, TRACE enhances the integrity of the corporate bond market and creates a level playing field for all investors" (NASD 2005a). The 2005 TRACE Fact Book adds: "From a regulatory standpoint, such levels of transparency better enable regulators to monitor the market, pricing and execution quality" (NASD 2005b).

Critics were concerned about how disclosure would impact the incentives of dealers and traders (see e.g., Bravo 2003 and Decker 2007) and in turn the operation of the corporate bond market. The Bond Market Association warned of "serious concerns about the potential harm to liquidity resulting from rapid transaction data on lower rated, less frequently traded issues" (Mullen 2004). In particular, there was a concern that dealers may be less likely to commit capital to hold inventory in illiquid securities when information about their transactions was made public. If bid-ask spreads subsidize dealers inventory holding costs and if TRACE reduces these spreads, it may become too costly for dealers to hold some less actively traded securities.

Another concern was that making trades public, particularly large trades, would disadvantage dealers. If large dealers buy in quantity and then provide liquidity to the market, having the price and quantity they buy at known may cap the resale price they can charge. Thus, as Duffie (2011) states, censoring trade information allows dealers to "have the chance to reduce inventory imbalances stemming from large trades with less concern that the size of a trade or their reservation price will be used to the bargaining advantage of their next counterparties."

These concerns ultimately motivated the NASD to censor trade size reports at \$1,000,000 for high-yield bonds and \$5,000,000 for investment grade bonds (Vames 2003).

On July 1, 2002, FINRA implemented TRACE, requiring dealers to report all bond transactions on TRACE-eligible securities within 75 minutes. As described in Table 2.1, FINRA began disseminating price and volume data for trades in selected investment-grade bonds with initial issue of \$1 billion or greater (i.e., Phase 1 bonds). FINRA's dissemination occurred immediately upon reporting for these bonds. A "TRACE-eligible security" is any US dollar-denominated debt security that is depository-eligible and registered by the SEC, or issued pursuant to Section 4(2) of the Securities Act of 1933 and purchased or sold pursuant to Rule 144a.⁴ Additionally, the 50 high-yield securities disseminated under FIPS were transferred to TRACE, which now disseminated their trades.⁵ We denote these bonds the FINRA50. About 520 securities had their information disseminated by the end of 2002.

At the start of Phase 1, it was not certain when and to what extent TRACE would be expanded. After all, the FIPS program had existed without expansion for eight years. Initially, a bond transactions reporting committee comprised of NASD and the Bond Market Association members was established to study TRACE's impact. Their mandate was to focus not on the largest, highest quality credit and actively traded issues, but rather on the rest of the market (Vames 2003). Their recommendation was to expand TRACE's coverage. The NASD approved the expansion of TRACE on November 21, 2002 and by the SEC on February 28, 2003.

Phase 2 of TRACE was implemented on March 3, 2003, and it expanded dissemination to include smaller investment grade issues. The new dissemination requirements included securities with at least \$100 million par value or greater and ratings of A- or higher. In addition, dissemination began on April 14, 2003 for a group of 120 Investment-Grade securities

⁴The list of eligible security types is: (1) Investment-grade debt, including Rule 144A/DTCC eligible securities, (2) High-yield and unrated debt of U.S. companies and foreign private companies, (3) Medium-term notes, (4) Convertible debt and other equity-linked corporate debt not listed on a national securities exchange, (5) Capital trust securities, (6) Equipment trust securities, (7) Floating rate notes, (8) Global bonds issued by U.S. companies and foreign private companies, and (9) Risk-linked debt securities (e.g., "catastrophe bonds"). TRACE-eligible securities exclude debt that is not depository-eligible, sovereign debt, development bank debt, mortgage- and asset-backed securities, collateralized mortgage obligations, and money market instruments.

⁵Alexander *et al.* (2000) examine the liquidity of the bonds in the FIPS dataset.

Table 2.1: Timeline of Major TRACE Regulatory Changes

Sample	Date	Bonds Affected	Time to Report
Phase 1	July 1, 2002	Investment Grade TRACE-eligible bonds having an initial issue of \$1 billion or greater	75 Minutes
FINRA50	July 1, 2002	50 Non-Investment Grade (High-Yield) bonds disseminated under Fixed Income Pricing System (FIPS). First day is 7/1/02, last day is 7/14/04	75 Minutes
Phase 2	March 3, 2003	All Investment Grade TRACE-eligible bonds of at least \$100 million par value (original issue size) or greater rated A- or higher; and 50 Non-Investment Grade bonds	75 Minutes
FINRA120	April 14, 2003	120 Investment Grade TRACE-eligible bonds rated BBB	75 Minutes
n/a	October 1, 2003	All currently disseminated bonds	45 Minutes
Phase 3A	October 1, 2004	All bonds that are not eligible for delayed dissemination (bonds with rating of BBB- or higher)	30 Minutes
Phase 3B	February 7, 2005	All bonds potentially subject to delayed dissemination (bonds with ratings BB+ or lower)	30 Minutes
n/a	July 1, 2005	All currently disseminated bonds	15 Minutes
n/a	January 9, 2006	All currently disseminated bonds	Immediate

Information from FINRA press releases available at finra.org. Time to report is the amount of time the dealer has to report the transaction to FINRA. FINRA collected information on all TRACE-eligible securities on July 1, 2002. A TRACE-eligible security means all US dollar-denominated debt securities that are depository-eligible and registered by the SEC, or issued pursuant to Section 4(2) of the Securities Act of 1933 and purchased or sold pursuant to Rule 144a. FINRA disseminates the transaction for Bonds Affected immediately after the report, except for bonds subject to delayed dissemination. Bonds subject to delayed dissemination must meet certain trading, size, and rating criteria described by Rule 6250(b).

rated BBB. We denote these BBB bonds as the FINRA120.⁶ After Phase 2 was implemented, the number of disseminated bonds increased to approximately 4,650 bonds. Meanwhile, the FINRA50 subset did not remain constant over our time period. On July 13, 2003, the FINRA50 list was updated, and the list was then updated quarterly for the next 5 quarters.⁷

Finally, on April 22, 2004, after TRACE had been in effect for some bonds for almost two years, the NASD approved the expansion of TRACE to almost all bonds. The last Phase came in two parts, which FINRA designates as Phase 3A and Phase 3B. The distinction between Phase 3A and 3B is that Phase 3B bonds are eligible for delayed dissemination. Dissemination is delayed if a transaction is over \$1 million and occurs in a bond that trades infrequently and is rated BB or below. In addition, dissemination is delayed for trades immediately following the offering of TRACE-eligible securities rated BBB or below. In Phase 3A, effective on October 1, 2004, 9,558 new bonds started having their information about trades disseminated. In Phase 3B, effective on February 7, 2005, an additional 3,016 bonds started dissemination, though sometimes with delay.⁸ According to the NASD at that point, there was “real-time dissemination of transaction and price data for 99 percent of corporate bond trades” (NASD 2005).

In an effort parallel to increasing the number of bonds with disseminated trade information, FINRA reduced the time delay for reporting a transaction from 75 minutes on July 1, 2002, to 45 minutes on October 1, 2003, to 30 minutes on October 1, 2004, and to 15 minutes on July 1, 2005. On January 9, 2006, the time delay for dissemination was eliminated. Since most bond trades infrequently, our trading analysis uses one day as the basic unit of time. In our sample the average number of trades per day for a bond is 0.68. Therefore, we do not focus

⁶The FINRA120 sample was selected by FINRA to study the impact of dissemination on market behavior and has been studied by Goldstein *et al.* (2007).

⁷The FINRA50 list was updated on July 13, 2003, October 15, 2003, January 15, 2004, April 14, 2004, and July 14, 2004.

⁸Rule 6250(b)(2)(A) states: “Transactions that are greater than \$1 million (par value) in BB-rated TRACE-eligible securities that trade an average of less than one time per day will be disseminated two business days from the time of execution.” Rule 6250(b)(2)(B) states: “Transactions that are greater than \$1 million (par value) in TRACE-eligible securities rated B or lower that trade an average of less than one time per day will be disseminated four business days from the time of execution.” On January 9, 2006, this exception changed and there was immediate dissemination of all trades.

on changes in time to dissemination, but instead on new dissemination.

2.2.2 Related Literature

There are three main studies of TRACE, each of which focuses on either Phase 1 or Phase 2. Bessembinder *et al.* (2006) study 439 bonds in Phase 1 using transaction data from the National Association of Insurance Commissioners. They formulate and estimate a structural model of transaction costs and report a 4.9-7.9 basis point reduction in transaction costs for Phase 1 bonds in a before-and-after comparison. They also find that after Phase 1, there is a decline in the concentration ratio for the 12 largest dealers.

Two other studies examine transaction costs for Phase 2 bonds. Using a then proprietary database of all bond trades (which is now publicly available), Edwards *et al.* (2007) also examine imputed transaction costs. They find that transparent bonds have lower transaction costs. Since this result may be due to bond characteristics rather than the effect of transparency, they also report on a difference-in-difference analysis, which compares the transactions costs of bonds which are newly disseminated to three distinct control groups of bonds that do not change dissemination status. The transactions costs of newly disseminated bonds decrease relative to each control group across the entire range of trade sizes.

Goldstein *et al.* (2007) report on a controlled experiment, commissioned by the NASD, of 120 BBB Phase 2 bonds, 90 of which are actively traded and 30 of which are relatively inactive. Through cooperation with the NASD, the authors construct a matched sample of the 90 actively traded bonds based on industry, average trades per day, bond age, and time to maturity. When the 90 actively traded bonds were disseminated on April 14, 2003, the matched bond was not. To increase power, they also compare the disseminated sample to a larger portfolio of non-disseminated bonds. For the 90 actively traded bonds, they find declines in transaction costs for all but the group with the smallest trade size. There is no evidence of a reduction in transaction costs for inactively traded bonds. In subsequent work, Hotchkiss and Goldstein and Hotchkiss (2007) study new issues of corporate bonds, and find a secular decline in price dispersion from July 2002 through February 2007 for newly issued bonds.

While these studies provide evidence that TRACE reduces transaction costs for Phase 1 and Phase 2 bonds, there is little evidence about TRACE's effect on trading activity. For their sample of 120 BBB bonds, Goldstein *et al.* (2007) report that TRACE did not cause an increase in daily trading volume and the number of transactions per day. Despite this small sample size and time period, Duffie (2011) concludes "the empirical evidence does not generally support prior concerns by dealers that the introduction of TRACE would reduce market liquidity." Others, including the SEC, saw the evidence as inconclusive, stating that concerns about liquidity were also not rejected.⁹

The absence of any trading activity results is surprising in light of the negative reaction to TRACE from many market participants. For instance, Bessembinder and Maxwell (2008) report that the near universal perception among bond dealers is that trading became more difficult after TRACE. (See also Jamieson 2006 and Decker 2007). Bessembinder and Maxwell (2008) are skeptical of these claims given that there was an upward trend in aggregate corporate bond trading from 2002-2007. This increase in aggregate bond trading does not imply TRACE increased trading activity, however, since there was also an upward trend in the amount of corporate debt outstanding due to new issues. When we hold the number of bonds constant by examining bonds covered in TRACE's four Phases, there is a strong downward trend in average daily volume (see Figure 2.1). In addition, we believe another the reason that previous work did not detect significant adverse effects on trading activity is that it did not examine the later Phases of TRACE, where the decline in trading activity is strongest.

Also relevant is a set of studies on municipal bonds. Green *et al.* (2007b) find significant price dispersion in new issues of municipal bonds, which they attribute to the decentralized and opaque market design. Green *et al.* (2007a) analyze broker-dealer and customer trades, and report that dealers exercise substantial market power. On January 31, 2005 the Municipal Securities Rulemaking Board started requiring that information about trades in municipal

⁹The SEC's Director of Market Regulation Nazareth stated "the NASD commissioned two studies to address this issue [the impact of TRACE on liquidity]. Neither study provided significant evidence that transparency harms liquidity. However, neither study was extensive enough to address all concerns raised by dealers and other market participants" (Nazareth 2004). The industry group, the Bond Market Association, described these studies as largely inconclusive (Mullen 2004)

bonds be reported within 15 minutes, similar to TRACE. Schultz (2012) compares price dispersion at offering date for municipal bonds before and after this change and finds that it falls sharply. He does not, however, study post-offer trading activity.

There is also empirical research on the effects of transparency in settings other than the bond market. Greenstone *et al.* (2006) study the mandatory disclosure requirements of the 1964 Securities Act Amendment. These requirements required OTC firms to register with the SEC, provide regular updates on financial positions, issue proxy statements, and report on insider holdings and trades. They find that these newly registered OTC firms experience positive abnormal returns post-disclosure. Further afield, Jensen (2007) investigates the impact of increased information on price dispersion among fishermen in southern India. After mobile phones became available, he finds a sharp reduction in price dispersion and a reduction of waste due to excess fish.

Finally, the theoretical work on the impact of dissemination highlights various mechanisms through which dissemination can impact trading behavior. (See Biais *et al.* 2005 for a review of the literature on the impact of transparency on financial markets). Madhavan (1995) demonstrates that dealers may prefer not to disclose trades because they benefit from the reduction in price competition. Pagano and Röell (1996) argue that well-informed dealers may be able to extract rents from less well-informed customers in an opaque market, and that transparency may result in more uninformed traders entering the market. Bloomfield and O'Hara (1999) show that transparency can reduce market makers incentives to supply liquidity, if the market maker has more difficulty unwinding inventory following large trades. On the other hand, Naik *et al.* (1999) show how transparency can improve dealers' ability to share risks, which decreases their inventory costs and therefore customers' costs of trading.

2.3 Data and Descriptive Statistics

2.3.1 Historical TRACE Data and Phase Identification

Beginning in July 2002, TRACE publicly provided price and volume data for disseminated trades for Phase 1 bonds.¹⁰ This and later publicly disseminated trade data constitutes the “Public” TRACE database available to market participants at the time. Simultaneously, FINRA also collected non-disseminated trade data. This non-disseminated data represents all trades on corporate bonds in the period before public dissemination. In March 2010, FINRA released a “Historical” TRACE dataset, which includes both disseminated and non-disseminated transaction records, starting from TRACE’s initiation in July 2002. We use the Historical TRACE database to examine the period from July 1, 2002 through December 31, 2006. Since Phase 3B, the last major Phase of TRACE, concluded in February 2005, our time period covers all four TRACE Phases.

The information in the FINRA databases (both Public and Historical) is self-reported by bond dealers who are FINRA members. Dealers are required to report the bond’s CUSIP, the trade’s execution time and date, the transaction price (\$100 = par), and the volume traded (in dollars of par). In addition, dealers are required to indicate whether they were the buyer or the seller, and whether the counterparty to the trade was a dealer or a customer. Unlike the Public TRACE database, the Historical TRACE database does not censor volume at \$1 million or \$5 million. Finally, dealers are required to correct errors in previously reported trades with flags corresponding to trade cancels, modifies, or reversals.

There are a number of steps required to process this raw data into the analysis dataset that we use. These steps and their rationale are described in detail in the Data Appendix and outlined in Table A.1. Two of the major steps are to eliminate all bonds not contained in the Mergent Fixed Income Securities Database (FISD), and to drop all bonds with an equity-like component since partial price information may be available from the stock market. Next we

¹⁰FINRA censored reported trading volume at \$1 million for high-yield bonds and \$5 million for investment-grade bonds. That is, for trades greater than this amount, the actual trading volume was not reported and TRACE only reported that the trade size exceeded the cap.

eliminate some of the trading records for the remaining bonds. There are three main reasons. First, there are records for trades that do not actually take place since they are cancelled, modified, or reversed. Second, there are records corresponding to trades that are reported more than once. Third, there are records with issues concerning their price, size, or timing. Table A.1 enumerates the number of bonds and trade records affected by each step.¹¹ After applying the filters described in Table A.1, there are 21,149,525 trades, corresponding to 30,643 CUSIPs, remaining in the Cleaned Historical TRACE database.

Phase Identification

FINRA's criterion for a bond's dissemination Phase is presented in Table 2.1. The main criteria are the initial issue size and the credit rating. FINRA does not indicate a bond's Phase designation in either the Historical or Public FINRA dataset. As a result, we contacted FINRA and obtained their listings of the bonds included at the start of Phases 2, 3A, and 3B. We obtained the list of bonds that are in the FINRA50 or FINRA120 directly from the FINRA website.¹²

FINRA did not provide us a list of bonds in Phase 1. To construct the Phase 1 list, we require a bond to have an initial issue size of \$1 billion or more, be investment grade (following the criteria FINRA used as outlined in Table 2.1), and have a publicly disseminated trade before the start of Phase 2. Bonds which are simultaneously classified in a Phase and in either the FINRA50 or FINRA120 are excluded from our Phase lists. The Data Appendix and Table A.2 further describe the steps involved in matching the Phase lists to the Cleaned Historical TRACE database.

Table A.2 shows that after cleaning, there are 343 Phase 1 bonds, 2,538 Phase 2 bonds, 11,087 Phase 3A bonds, and 2,853 Phase 3B bonds. We designate these 16,825 bonds and 14,210,328 trades as the Cleaned Phase Sample. The remaining bonds in the Cleaned Historical

¹¹We do not exclude bonds trades that occurred on the NYSE's Automated Bond System. Even though they take place on an exchange and therefore are transparent, they constitute a tiny fraction of the market. For instance, Goldstein *et al.* (2007) state that 99.9% of corporate bond trading in 2004 takes place over-the-counter.

¹²The list is available at <http://www.finra.org/Industry/Compliance/MarketTransparency/TRACE/Announcements/P117685>, last accessed January 28, 2013.

TRACE database are not associated with any Phase. 7,669 bonds are always disseminated (they were issued after the beginning of their Phases and always disseminated) and 1,708 bonds are never disseminated (they matured before the start of what would have been their Phase). Finally, 671 bonds are not disseminated consistent with FINRA's guidelines. They either have some non-disseminated trades after a bond's Phase began or some disseminated trades before the Phase's start date.

Although the number of bonds disseminated in Phase 1 and Phase 2 is lower than the number in Phases 3A and 3B, bonds in the earlier Phases account for a larger number of trades per bond. For instance, bonds in Phase 1 are heavily traded with a total of 10,208 trades per bond over our sample period. In contrast, bonds in Phase 3B have only 351 trades per bond.

2.3.2 Bond Characteristics

Table 2.2 shows the distribution of issue size, credit rating, coupon rate, and maturity for our sample of bonds by Phases. As mentioned above, when assigning bonds to Phases, FINRA uses issue size and rating as criteria. Table 2.2 shows the mean bond issue size decreases from Phase 1 to Phase 3A, consistent with the rules set by FINRA outlined in Table 2.1. Phase 1 bonds have by far the largest issue size with a mean of \$1.466 billion and Phase 3A bonds are the smallest with mean issue sizes of \$82 million. Phase 3B bonds have a larger mean issue size of \$181 million.

We also report the quartiles of the issue size distribution as well as the 10th and 90th percentiles. These quantiles show that there is overlap in issue size between Phases 2, 3A, and 3B. For example, the median of Phase 3B bonds equals the 25th percentile of Phase 2 bonds and the 75th percentile of Phase 3A bonds is close to the 25th percentile of Phase 3B bonds. These overlapping intervals allow us to compare bonds with similar issue sizes across Phases 2, 3A, and 3B.

Data on credit ratings comes from two sources. We first use ratings information from S&P RatingsXpress if it is available. This covers 74.5% of bonds for the four Phases. If ratings are

Table 2.2: *Bond Characteristics by Phase*

	Phase 1 (1)	Phase 2 (2)	Phase 3A (3)	Phase 3B (4)
Number of Bonds	343	2,538	11,087	2,853
Size at Issue (\$M)				
mean	1,466	263	82	181
p10	1,000	100	1	8
p25	1,000	150	3	85
median	1,250	200	12	150
p75	1,750	300	75	232
p90	2,500	500	288	350
Rating at Phase Start				
mean	A	A+	A-	B
p10	AA-	AA	AA	BB+
p25	A+	A+	A	BB-
median	A	A+	A-	B
p75	BBB+	A	BBB	CCC
p90	BBB	A-	BBB-	D
# where rating is from S&P RatingsXpress	331	2,191	7,733	2,274
# where rating is from FISD	12	345	3,319	489
Fixed Coupon Rate				
mean	6.7	6.9	5.8	9.0
median	6.8	6.9	5.9	8.8
number fixed coupon	309	2,155	10,149	2,632
Maturity at Issue (years)				
mean	9.0	15.0	11.8	12.4
median	5.1	10.0	10.0	9.7
Years since Issue (at Phase Start)				
mean	1.9	5.5	3.4	5.9
median	1.5	5.1	1.9	5.7

Bond issue size, coupon, maturity, and issue date characteristics are from FISD. Bond rating are the most recent rating before the Phase starts. Bond rating characteristics are from S&P RatingsXpress database. If ratings are not available in S&P RatingsXpress, we use ratings from FISD. To assign a FISD rating, we first use the S&P value if it exists, otherwise the Moody's value, otherwise the Fitch value, and otherwise the Duff and Phelps value. Mean ratings are computed by first converting each rating to a number (AAA=1, AA+=2, AA=3, ..., and D=22) and then converting the number back to a letter rating.

not available in S&P RatingsXpress, we use ratings from FISD.^{13,14} FISD includes ratings from S&P, Moody's, Fitch and Duff and Phelps. To assign a FISD rating, we first use the S&P value if it exists, otherwise the Moody's value, otherwise the Fitch value, and otherwise the Duff and Phelps value. If FISD does not have a rating from any of the four, we classify the bond as unrated. Using both sources, there are ratings for 99.2% of bonds, and 127 bonds are classified as unrated.

Table 2.2 shows the distribution of credit ratings at the start of each Phase. The average rating at the beginning of the Phase is similar between Phases 1, 2, and 3A. Bonds in Phase 3B have significantly lower credit ratings. While there is overlap between the ratings in Phases 1, 2, and 3A, there is little or no overlap in ratings between Phase 3B and the other Phases. The 10th percentile rating in Phase 3B is a BB+, while the 90th percentile rating in Phase 1, 2, and 3A are BBB, A-, and BBB-, respectively.

Table 2.2 also describes bond characteristics not used by FINRA when assigning Phases. For example, most bonds have fixed coupon rates. The only Phase with less than 90% fixed coupons is Phase 2 and even these bonds have fixed coupons 84.9% of the time. Consistent with ratings, the highest coupon rates are for Phase 3B. In addition, Phase 1 bonds have the lowest maturity at issue with a mean of 8.98 years and a median of 5.10 years. All three of the other Phases have a mean maturity greater than 11.8 years and a median maturity greater than 9.7 years.

2.3.3 Measuring Trading Activity and Price Dispersion

We measure trading activity in several ways. Our first measure is trading volume, which we define as the number of bonds traded times their par value. Figure 2.1 plots the daily trading volume averaged by week for the bonds in Phases 2, 3A, and 3B from July 2002 through

¹³Akins (2012) states that the S&P RatingsXpress database is more complete than FISD's S&P ratings database.

¹⁴FINRA does not rely exclusively on S&P ratings. It also uses ratings from other nationally recognized statistical rating organizations. If a bond is unrated or split rated, FINRA has specific rules determining the bond's rating for the purposes of Phase classification.

December 2006.¹⁵ The three vertical lines correspond to the starting date for each of the three Phases.¹⁶ For all three Phases, the average daily trading volume fell by about a half over the entire period July 2002 to December 2006. While this volume drop may be due to TRACE, we cannot, at this point, exclude the possibility that there is a pre-existing downward trend in volume independent of TRACE.

To focus on changes in the immediate time period surrounding dissemination, the first section of Table 2.3 reports the mean and quartiles of daily volume for the period 90-days before and 90-days after the beginning of each Phase.¹⁷ Table 2.3 shows the mean trading volume is lower in the 90-day period after the start of each Phase than in the 90-day period before each Phase. The declines in Phases 2 and 3A are not as large as that for Phase 3B, where the average 90-day trading volume falls 41.9%. For Phase 2 and 3A, the percentage declines are 4.9% and 5.5%, respectively.

Table 2.3 also shows how skewed the distribution of trading volume is across our sample. The mean trading volume in Phases 3A and 3B is roughly 100 times greater than the medians in the period before dissemination. In addition, more than half of the Phase 3B bonds do not trade in the 90 days after dissemination. Moreover, the average trading volume for Phase 1 bonds is more than 50 times greater than the average trading volume for Phase 3B bonds for the post 90-day period. Taken together, these facts suggest substantial heterogeneity in trading volume within and across our bond samples.

These differences in trading volume across Phases may be due to difference in bond issue sizes. A larger bond issue may generate more after-market trading simply because there are more bonds to trade. As shown in Table 2.2, the mean issue size of Phase 1 bonds is almost six times greater than those in Phase 2. Phase 2 bonds' mean issue size is three times those of

¹⁵Figure 2.1 does not include trading days that SIFMA recommends that bond dealers take off or operate for less than a full day. Additionally, Figure 2.1 does not include the two weeks spanning Christmas and New Year's Day due to significantly reduced volume.

¹⁶Bonds in Phase 1 are not plotted in Figure 2.1 because of scaling. Phase 1 bonds have an average daily volume of 7,513,772 for the sample period.

¹⁷Since bonds trade infrequently, we use a 90-day window to capture changes in trading behavior. In Table 2.4, we also look at 30- and 60-day windows.

Figure 2.1.1: Weekly Trading Volume by Phase




Figure does not include trading days that SIFMA recommends that bond dealers take off or operate for less than a full day. Figure also does not include the two weeks spanning Christmas and New Year's Day.

Table 2.3: *Trading Activity and Price Dispersion for the 90-Day Window Around Phase Start*

	Percentiles										Within-Bond Comparison						
	Mean		25th				50th				75th		Before >		Before = After		
	Before (1)	After (2)	Before (3)	After (4)	Before (5)	After (6)	Before (7)	After (8)	Before (9)	After (10)	zero (11)	non-zero (12)					
A. Trading Activity																	
Volume																	
Phase 1 (N=343)	...	11,445,643	...	4,251,581	...	7,828,097	...	13,986,145
Phase 2 (N=2,538)	888,352	844,458	57,250	52,459	319,339	299,180	950,000	844,361	51.4%	43.7%	4.9%	0.0%					
Phase 3A (N=11,087)	335,026	316,497	0	105	3,361	4,035	30,738	31,228	39.3%	43.9%	16.8%	0.0%					
Phase 3B (N=2,853)	366,526	213,035	0	0	3,818	0	357,000	91,475	45.1%	15.2%	39.7%	0.0%					
Volume/Issue Size																	
Phase 1 (N=343)	...	0.76%	...	0.36%	...	0.59%	...	0.97%
Phase 2 (N=2,538)	0.28%	0.27%	0.03%	0.03%	0.15%	0.13%	0.35%	0.32%	51.4%	43.7%	4.9%	0.0%					
Phase 3A (N=11,087)	0.16%	0.16%	0.00%	0.00%	0.04%	0.05%	0.13%	0.15%	39.3%	43.9%	16.8%	0.0%					
Phase 3B (N=2,853)	0.18%	0.09%	0.00%	0.00%	0.01%	0.00%	0.23%	0.06%	45.1%	15.2%	39.7%	0.0%					
B. Price Dispersion																	
Price Standard Deviation																	
Phase 1 (N=340)	...	0.88	...	0.54	...	0.78	...	1.17
Phase 2 (N=2,023)	0.83	0.76	0.37	0.33	0.67	0.65	1.12	1.04	56.6%	43.2%	0.1%	0.0%					
Phase 3A (N=6,319)	0.78	0.68	0.35	0.31	0.65	0.57	1.09	0.95	59.6%	40.2%	0.2%	0.0%					
Phase 3B (N=1,129)	0.65	0.45	0.24	0.16	0.40	0.30	0.73	0.59	63.5%	36.1%	0.1%	0.3%					

Average daily volume is averaged over all bond-days in either the 90 calendar days before or after the Phase start. If a bond does not trade, it contributes zero daily volume for that day. Average daily volume/issue size is the average of daily volume/issue size calculated in the same manner as average daily volume. For average daily price standard deviation, the sample of bonds is restricted to bonds where there is at least one day in the 90 days before the Phase start with at least two trades and there is at least one day in the 90 days after the Phase start with at least two trades. After computing the within-day price standard deviation for each bond for all days with at least two trades, we average across these days during either the 90 days before or after the Phase start. Reported average daily price standard deviation is the average across these bonds. There is no transaction information for Phase 1 bonds in the 90 days before Phase 1 starts. N refers to the number of bonds that change dissemination status in the Phase. In column (9), Before > After reports the fraction of bonds where the measured outcome is greater in the 90 days before the Phase start than in the 90 days after. In column (10), After > Before represents the reverse of column (9). In column (11), Before=After (zero) reports the fraction of bonds where the measured outcome is equal to zero both before and after. In column (12), Before=After (non-zero) reports the fraction where the measured outcome is non-zero and is equal both before and after.

bonds in Phase 3A. Comparing median issue sizes in Table 2.2 across Phases sometimes leads to even larger differences. For example, the median issue size in Phase 2 is \$200 million, while the median issue size in Phase 3A is \$12 million.

To address the issue of whether the difference in volume across Phases is driven by differences in issue size, we next examine volume divided by issue size. Figure 2.2 plots volume divided by issue size for each of the four Phases. While the time-series of volume/issue size in Figure 2.2 follows the time-series for volume in Figure 2.1, dividing volume by issue size makes the plots of trading activity for Phases 2, 3A, and 3B closer to one another than volume alone. In addition, the second section of Table 2.3, which reports statistics on volume/issue size by Phase, reinforces this conclusion.¹⁸ Normalizing by issue size reduces the skewness in comparisons both within and across Phases. Comparing within Phases, the mean of volume/issue size in Phases 3A and 3B is four and 18 times the median respectively. This compares to a ratio of about 100 for volume as discussed above. Comparing across Phases, the mean of volume/issue size in Phase 1 is eight times that in Phase 3B in the 90 days after dissemination. This compares to 50 times when using volume. Consequently, the remainder of the paper reports volume/issue size as our primary measure of trading activity. We also conduct the entire analysis using volume alone, but to save space we only report those results when discussing alternative measures of trading in Table 2.6.

Table 2.3 also reports a within-bond metric, by computing the fraction of bonds for which trading volume increases, decreases or remains the same in the 90 days before and after the Phase initiation date. Since the comparison is before vs. after for a given bond, the numbers are identical whether using volume or volume/issue size. Phase 3B bonds show a pronounced decline in trading activity in the within-bond comparisons. 45.1% of Phase 3B bonds have more trading volume before dissemination, while 15.2% of Phase 3B bonds have more trading volume afterwards. A large percentage of Phase 3B bonds, 39.7%, do not trade in the 90-days before or after the beginning of the Phase. The within-bond results for Phase 2 bonds also

¹⁸An alternative normalization would be log volume. As seen in Table 2.3, this is infeasible since volume is equal to zero for many bonds in the 90 days surrounding the Phase starts.

Figure 2.2: Weekly Trading Volume/Issue Size by Phase




Figure does not include trading days that SIFMA recommends that bond dealers take off or operate for less than a full day. Figure also does not include the two weeks spanning Christmas and New Year's Day.

show a decline but not as much, from 51.4% to 43.7%. The results for Phase 3A are mixed. The fraction of bonds with higher volume post dissemination is slightly greater than for before dissemination, however, the mean volume declines from the period before to after.

Price Dispersion

We also examine the impact of transparency on price dispersion. We begin by focusing on the daily price standard deviation, defined for bond i on day t as

$$\sigma_{it} = \left(\frac{1}{N_{it}} \sum_j (p_{itj} - \bar{p}_{it})^2 \right)^{\frac{1}{2}}$$

where N_{it} is the number of trades for bond i on day t , p_{itj} is the price of bond i for trade j on day t and \bar{p}_{it} is the average price of bond i on day t . We focus on price standard deviation because it does not depend on assumptions about the relationship between transaction prices and order flow. We examine other measures of price dispersion in Section 2.5. All measures of daily price standard deviation are in units of dollars.

To compute a daily price standard deviation, it is necessary to observe at least two bond trades in a day. Given the lack of trading in many bonds, we often do not observe two trades.¹⁹ Further, to measure the effects of dissemination on price dispersion, we require that there is at least one daily price standard deviation observation both in the 90 days before and in the 90 days after the bond's change in dissemination. As a result, the number of bonds used in our price standard deviation analysis is substantially smaller than the number used in the volume analysis. This can be seen in Table 2.3's sample counts for each Phase. For example, only 57.0% of Phase 3A and 40.0% of Phase 3B bonds in the volume sample are also in the price standard deviation sample. Although not shown, the bonds for which we can compute price standard deviation tend to have a larger size at issue and higher rating than the volume sample.

¹⁹Measures of transaction costs such as direct round trip or imputed transaction costs also present difficulties for less actively traded bonds since they require observing multiple trades within a short time horizon. For instance, Edwards *et al.* (2007)'s method requires that a bond trades at least nine times.

There is a potential bias in our price standard deviation measure since the sample is defined based on trading behavior both before and after changes in dissemination. If dissemination causes an increase or decrease in bond trading, this may change the number of bonds for which we can compute price standard deviation.²⁰ Thus, if the bonds that would have traded without dissemination substantially differ from the bonds that do trade with dissemination, then it may be difficult to interpret changes in price standard deviation.²¹ This appears to not to be an issue for our sample. To further investigate the robustness of our price standard deviation findings, in Section 2.5 we construct a matched sample of bonds holding constant the observable characteristics of bonds before and after dissemination.

Figure 2.3 plots the daily price standard deviation averaged by week from July 2002 through December 2006.²² Just as with trading volume, there is a reduction in price standard deviation over the entire time period. In fact, the price standard deviation falls by over a half from July 2002 to December 2006. However, unlike trading volume, the decline in price standard deviation seems to initiate at TRACE's launch, and continues through 2005. Another pattern in Figure 2.3 is that price standard deviation, over the entire period, is usually highest for Phase 3A bonds, and is lowest for Phase 1. Furthermore, standard deviation for Phase

²⁰This problem does not affect our volume calculations because when a bond does not trade, it counts as having zero trading volume.

²¹The probability that any of the Phase 2, 3A, or 3B bonds trade at least twice on a day in the 90 days before dissemination is 12.5%. To test whether this probability changes after TRACE, we estimate the effect of TRACE on the probability that a bond trades twice or more on a given day. The estimates come from a difference-in-difference regression similar to those estimated in Table 2.6, where the dependent variable is an indicator for whether a bond trades two or more times in a day. (The next section introduces our difference-in-difference methodology.) There is a statistically significant 0.37% reduction in the probability of trading for treated bonds across all three Phases. Assuming that the likelihood of trading is independent across days, this implies that TRACE causes a negligible reduction in the probability that a bond's price standard deviation can be measured across 90 calendar days. The estimated probability that a bond is no longer in the price standard deviation sample due to TRACE is less than 0.01%. This is calculated as follows: the probability that in any day among the 90 calendar days before there are at least two trades on the same day and that in any day among the 90 calendar days after dissemination there at least two trades on the same day is equal $(1 - (1 - Pr(\text{at least two trades on day}))^{64}) * (1 - (1 - Pr(\text{at least two trades on day}))^{64})$, where 64 is the average number of trading days among 90 calendar days. The 0.37% reduction in the probability of at least two trades on a day from estimated probability of at least two trades before TRACE of 12.5% yields a 0.01% reduction in the probability that a bond will be in price standard deviation sample due to TRACE.

²²Following Figure 2.1, Figure 2.2 does not include trading days that SIFMA recommends that bond dealers take off or operate for less than a full day and does not include the two weeks spanning Christmas and New Year's Day.

1 bonds is lower than for Phase 2 and Phase 3A in the early part of the sample period, but converges by the end of our sample period.

Table 2.3 also reports on price standard deviation in the 90-day window around when a bond changes its dissemination status. There is a reduction in price standard deviation, measured in dollars, for bonds in all three Phases. The average Phase 2 bond's price standard deviation falls from \$0.83 to \$0.76, a 8.4% reduction, while the median Phase 2 bond's price standard deviation falls from \$0.67 to \$0.65. The percentage of bonds with higher standard deviation before the start of Phase 2 is 56.6%. The drop in price standard deviation is even greater for Phase 3A and 3B bonds. The average Phase 3A bond's price standard deviation falls by \$0.10, which is a 13.1% decrease, while the average Phase 3B falls by \$0.20, which is a 30.8% decrease. The median bond's price standard deviation drops by \$0.08 and \$0.10, respectively. Column (5) of Table 2.3 shows that the number of bonds for whom the price standard deviation is greater beforehand is 59.6% and 63.5% for Phases 3A and 3B, respectively.

Thus, Figures 2.1, 2.2, and 2.3 and Table 2.3 show that TRACE coincides with a decrease in trading volume for Phase 3B bonds. Moreover, there are sharp reductions in price standard deviation in each of the three Phases within a short 90-day window surrounding dissemination. However, changes in either volume or price standard deviation are contemporaneous with an overall downward trend in trading activity and in price standard deviation during our sample period. As a result, we cannot immediately conclude that any changes or lack of changes are the result of TRACE alone. Our next task is to adjust for market trends.

2.4 Research Design and Main Results

2.4.1 Difference in Differences Framework

Although the before-and-after comparisons in Table 2.3 show that price standard deviation falls for bonds in all Phases and trading volume declines for Phase 3B bonds, a before-and-after comparison is not sufficient to attribute the changes to dissemination alone. We adjust for market trends by comparing the changes in the treated sample to those in a control group by

Figure 2.3: *Weekly Price Standard Deviation by Phase*




Figure does not include trading days that SIFMA recommends that bond dealers take off or operate for less than a full day. Figure also does not include the two weeks spanning Christmas and New Year's Day.

estimating difference-in-differences models of the form:

$$y_{it} = \alpha + \gamma_0 \text{Disseminate}_i + \gamma_1 \text{Post}_t + \lambda \text{Disseminate}_i \times \text{Post}_t + \epsilon_{it} \quad (2.1)$$

where y_{it} is bond i 's outcome (i.e., measures of trading activity or price dispersion) on day t , Disseminate_i is an indicator for whether the bond changes dissemination status (i.e., is in the treated group) and Post_t is an indicator for the trade outcomes on days after the dissemination period. Since there are repeated observations per bond, in all estimates, the standard errors are clustered by bond.

In equation 2.1, any pre-existing difference between bonds that change dissemination status and those that do not are captured by γ_0 . Any effects of dissemination that accrue to all bonds - that is, effects that are not limited to only bonds that change their dissemination status in the Phase - are absorbed by time effects γ_1 . The coefficient of interest is λ , which estimates the direct effect of transparency on a bond's trading outcome. The coefficient λ reflects the change in trading outcomes for bonds that change dissemination status compared to the change in trading outcomes for bonds that do not change dissemination status. Estimates of λ , therefore, net out aggregate changes in bond trading outcomes.

It is possible that changes in dissemination will also affect bonds that do not change dissemination if the market impounds that information into all trading activity. Indeed, the overall downward trend in trading activity and price standard deviation in Figures 2.1 and 2.2 may be the consequence of TRACE's introduction in July 2002. However, we cannot assert that TRACE caused this decrease because we do not observe trading activity before Phase 1. The overall downward trend could instead be due to macroeconomic factors affecting the corporate bond market. For example, the Federal Reserve Bank raised interest rates 17 times from June 2004 through June 2006 (NASD (2006)). In our regression equation, the time effects incorporate all of these potential factors, and therefore we cannot interpret the estimates of γ_1 as a causal effect of dissemination.

For λ to provide unbiased estimates of the causal effect of transparency there are several important necessary assumptions. First, transparency and its consequences must not have

been fully anticipated by market participants; to the degree that impacts were foreseen by traders and dealers, the impacts on trading activity and price dispersion would appear before the actual change in dissemination status. If all trade outcomes responded immediately at Phase 1, our TRACE results for Phases 2, 3A, and 3B would only measure the incremental impact of later Phases of TRACE. Bessembinder *et al.* (2006) first emphasized this point when they argued that TRACE's initiation affected all bonds, not only those in Phase 1. In this case, our estimates understate the true impact of TRACE. (In Section 2.7, we investigate Phase 1 using a separate data set from the National Association of Insurance Commissioners.)

It seems unlikely that the effects of TRACE occurred in their entirety at the beginning of Phase 1. Even though TRACE started collecting information on trade activity for all bonds from July 1, 2002, the schedule of when transaction data would be disseminated remained uncertain. The timing of the expansions was not initially known and took place incrementally, depending on both FINRA and SEC approval. For example, FINRA, then NASD, approved Phase 2 on November 21, 2002, but the SEC did not approve it until February 28, 2003. Phase 2 was implemented on March 3, 2003. Thus, participants knew in advance that dissemination would expand, but they did not exact timing until shortly before it occurred.

The second assumption for λ to be a causal estimate is that there are no other changes simultaneous with the Phase start date that affects the trading activity for those bonds changing dissemination status. That is, in equation (2), the interaction between Disseminate and Post is uncorrelated with other unmeasured factors that affect trade activity that change at the same time as the change in dissemination status (but are not caused by the change in dissemination status). There are trends in the bond market trading during our time period, but we are unaware of any changes to bond market trading that coincide with the Phase start dates.

Finally, a third assumption is that we can measure the counterfactual difference in bond trading with the bonds that do not change dissemination status. That is, we assume that the change over time in control bonds' behavior reveals what would have occurred to treated bonds if there had been no change in their dissemination status. Note this assumption does not mean

that control bonds must have the same characteristics as treated bonds, but rather that the change in their behavior captures the counterfactual time path. This is important because our treated bonds have different attributes than our control bonds by definition. FINRA selected bonds for Phases based on characteristics such as ratings and issue size. For instance, Phase 2 bonds are investment grade and have an original issue size of at least \$100 million. Hence, our third assumption will be violated if the bond trading activity varies substantially over time due to different bond characteristics. We examine the sensitivity of our results to these three assumptions in the next section.

To estimate equation 2.1, there are two implementation decisions. First, it is necessary to specify the estimation window. Since bonds trade infrequently, a longer time window may be needed to observe changes in trading activity. A longer time window, however, may lead us to misattribute the effect of a change in dissemination to underlying market trends. For these reasons, we report estimates of equation 2.1 for three different estimation windows covering 30, 60, and 90 days surrounding the Phase start dates.

The second implementation decision is how to define the control bonds for any Phase for these regressions. Because of the four distinct TRACE Phases, there are several possibilities for defining control bonds. Control bonds can be defined as bonds that were already disseminated before the Phase begins. For example, to measure the impact of transparency on Phase 2 bonds, we can compare the trading behavior of Phase 2 bonds with the trading behavior of Phase 1 bonds. Alternatively, a control group can be defined as bonds that are disseminated in a later Phase. For example, for Phase 2 bonds, the control group could be Phase 3A and Phase 3B bonds.

We defined our control group several ways, both including Phase 1 bonds that were already disseminated and also excluding Phase 1 and only including bonds from later Phases that were not disseminated. We find that including or excluding Phase 1 bonds does not change our results in any meaningful way. With the exception of our robustness tests in Table 2.6, our Tables all use Phase 1 bonds in the control groups.

Another issue with control groups that we must confront is that Phase 3A and Phase 3B

occur just over four months apart, on October 1, 2004 and February 7, 2005, respectively. If we use a 90-day window before and after a Phase to capture the effects of dissemination, the post-dissemination trading of Phase 3A overlaps with the pre-dissemination trading of Phase 3B. Therefore, we do not use Phase 3B bonds as controls for Phase 3A bonds, and vice versa. When we present the analysis below, we use the bonds in Phases 1, 3A, and 3B as control bonds for Phase 2, and we use the bonds in Phases 1 and 2 as control bonds for Phases 3A and 3B.

2.4.2 Estimates

Table 2.4 reports estimates of equation 2.1 for 30, 60, and 90-day windows for bonds in Phases 2, 3A, and 3B, separately. It also reports pooled estimates, based on equation 2.1, with data stacked across the three Phases. That is, there are separate intercepts α for each Phase and γ_0 and γ_1 is also allowed to differ by Phase, while λ does not differ by Phase.

The estimate of the effect of TRACE on trading volume/issue size, pooled across all three Phases, is negative and significant for all three estimation windows. Across all Phases, volume/issue size (in percent, i.e., multiplied by 100) drops by 0.027 in the 90-day window around dissemination, which is significant at the 1% level. This is a 15.2% reduction from 0.178, the average level before dissemination. Across Phases, the only statistically significant reduction in volume/issue size for all estimation windows is for Phase 3B, which is significant at the 1% level.

In the 90-day window, TRACE reduces the average volume/issue size (in percent) for Phase 3B bonds by 0.074. This represents a 41.3% drop from the average level before dissemination. These findings reinforce the within-bond comparisons reported in column (9) of Table 2.3, which shows that three times as many bonds in Phase 3B have lower volume after dissemination than before.

Price standard deviation, reported in columns (6), (7) and (8), drops significantly (at the 1% level) after dissemination for all estimation windows, and for both the pooled sample and each Phase separately. In the 90-day window the pooled estimate of the reduction in price

Table 2.4: *Difference-in-Difference Estimates of Transparency on Trading Activity and Price Dispersion*

	Volume / Issue Size (in percent)				Price Standard Deviation			
	Mean for Disseminated (1)	30 days (2)	60 days (3)	90 days (4)	Mean for Disseminated (5)	30 days (6)	60 days (7)	90 days (8)
All Three Phases	0.178	-0.017** (0.008) -9.6%	-0.020*** (0.006) -11.2%	-0.027*** (0.005) -15.2%	0.901	-0.090*** (0.013) -9.7%	-0.095*** (0.010) -10.4%	-0.077*** (0.009) -8.5%
		A. Pooled						
Phase 2	0.279	0.017 (0.015) 6.1%	-0.001 (0.013) -0.4%	-0.026** (0.011) -9.3%	0.953	-0.099*** (0.024) -9.9%	-0.090*** (0.017) -9.4%	-0.070*** (0.015) -7.3%
		B. By Phase						
Phase 3A	0.156	-0.011 (0.013) -7.1%	-0.005 (0.009) -3.2%	0.003 (0.007) 1.9%	0.903	-0.057*** (0.014) -5.9%	-0.071*** (0.010) -7.8%	-0.058*** (0.009) -6.5%
Phase 3B	0.179	-0.082*** (0.011) -45.8%	-0.074*** (0.008) -41.3%	-0.074*** (0.008) -41.3%	0.679	-0.167*** (0.037) -24.4%	-0.198*** (0.035) -28.7%	-0.168*** (0.030) -24.7%
H ₀ : Phase effects equal		0.000	0.000	0.000		0.011	0.002	0.003
# of Phase 2 bonds		2,538	2,538	2,538		1,671	1,921	2,023
# of Phase 3A bonds		11,087	11,087	11,087		4,008	5,463	6,319
# of Phase 3B bonds		2,853	2,853	2,853		797	1,028	1,129
# of bond-days		1,155,677	2,313,806	3,410,347		183,214	376,174	557,057

This table reports estimates of Disseminate x Post for volume/issue size and price standard deviation. Panel A reports estimates from Phases 2, 3A, and 3B pooled together, while panel B reports estimates for each Phase separately. Robust standard errors clustered by bond and Phase are in parenthesis immediately below the estimates. Mean for Disseminated is the 90-day average for newly disseminated bonds immediately before the Phase start. The third entry in each row is the percentage effect as computed by dividing the estimate by the Mean for Disseminated. Phase effects equal reports p-values of tests that the three Phase estimates are equal. The 30-day regressions use trades from 30 calendar days before and after the Phase change. The 60- and 90-day regressions are defined analogously. For volume/issue size, there are 8,299 Phase 2, 2,020 Phase 3A, and 2,098 Phase 3B control bonds in columns (2)-(4). For price standard deviation, there are 4,057 Phase 2, 1,452 Phase 3A, and 1,430 Phase 3B control bonds in column (6), 5,057 Phase 2, 1,681 Phase 3A, and 1,587 Phase 3B control bonds in column (7), and 5,545 Phase 2, 1,769 Phase 3A, and 1,670 Phase 3B control bonds in column (8). * significant at 10%; ** significant at 5%; *** significant at 1%

standard deviation is 7.7 cents and is highly significant. Across the Phases, the smallest 90-day drop is for Phase 3A bonds. These bonds experience a significant reduction of 5.9 cents in their daily price standard deviation, which represents a 6.5% decrease relative to before the start of the Phase 3A. The largest drop is for bonds in Phase 3B. These bonds experience a significant reduction by 16.8 cents, which corresponds to a 24.7% reduction from the previous level. This pattern mirrors those the price standard deviation results in the within-bond comparisons reported in Table 2.3.²³

In summary, the estimates in Table 2.4 suggest that transparency causes a significant reduction in volume/issue size for Phase 3B bonds. In addition, daily price standard deviation falls significantly across all Phases. Since for each Phase our results are more precisely estimated at the 90-day window than at the 30 or 60-day window in subsequent tables, we report estimates from the 90-day estimation windows.

2.5 Timing, Robustness, and Other Measures of Trading Activity and Price Dispersion

In this section, we revisit the assumptions underlying the difference-in-differences estimates above and report estimates for other measures of trading activity and price dispersion.

2.5.1 Event Study and Time Windows

Table 2.4 does not tell us how long it takes for the market to react to a change in dissemination. Changes may be immediate if market participants anticipate the effects of dissemination in advance of Phase start dates. On the other hand, changes due to dissemination may occur with delay because of adjustment costs, such as rebalancing inventories, faced by market

²³The mean daily price standard deviation in column (5) of Table 2.4 is not identical to the mean daily price standard deviation in column (1) of Table 2.3. In Table 2.3, we compute the average daily price standard deviation, equally weighted by bond. In Table 2.4, we compute the average daily price standard deviation without weighting by bond, and cluster by bond in the regression. Since we require at least two trades on a day to calculate daily price standard deviation, unlike volume, we do not observe price standard deviation for each day and, hence, the calculated daily price standard deviation differs between Table 2.3 and 4 due to weighting. The measured daily price standard deviation in Tables 2.3 and 2.4 are close, and the relative sizes by Phase are similar.

participants. Delays may also occur if participants require time to utilize the newly available data. Moreover, the relative infrequency of bond trading may make it difficult to detect the effects of dissemination in short estimation periods.

To examine when the effects of dissemination begin, we estimate an “event-study” version of the regression model that allows the effects to differ by one-week intervals:

$$y_{it} = \alpha + \gamma_0 \text{Disseminate}_i + \gamma_w \text{One-Week Interval}_t + \lambda_w \text{Disseminate}_i \times \text{One-Week Interval}_t + \epsilon_{it} \quad (2.2)$$

where the $\text{One-Week Interval}_t$ is an indicator of whether day t is in week w . Equation 2.2 is estimated for each Phase separately. γ_0 captures any pre-existing difference between disseminated and non-disseminated bonds, while γ_w captures the overall trend in trading outcome in week w .

The estimate of λ_w is the amount by which the average newly disseminated bond deviates in trading outcome (either volume/issue size or price standard deviation) from control bonds during the one-week interval w . If there is a trend in the market that only affects bonds that change dissemination status, it should be reflected in the relative levels of λ_w . For example, if volume in newly disseminated bonds is trending down in the time period before a change in dissemination, the λ_w 's will be higher before than after. Since the estimates of λ_w are based on one-week contrasts, they will be estimated less precisely than models which impose a common effect for the period before and a separate common effect for the period after as in equation 2.1.

Figure 2.4 plots values of λ_w for trading volume/issue size for each week by Phase. We adopt the convention that week 0 includes the dissemination date and the six calendar days following it. We normalize λ_w to be zero in the week before the change in dissemination (i.e., week -1) and we add a vertical line to the plot for that week.²⁴ The patterns in Figure 2.4 for Phase 2 and 3A are consistent with the results in Tables 2.3 and 2.4. Volume/issue size is not affected by transparency for bonds since there is no shift in the level of coefficient estimates

²⁴Since the event study includes the period from 90 days before and 90 days after day 0, there is one fewer calendar day in week -13.

after dissemination in the Figure.

The Phase 3B plot in Figure 2.4 shows a sharp and significant drop in volume/issue size from the week immediately preceding dissemination to the first week after it. This suggests that the negative volume/issue size results for Phase 3B in Tables 2.3 and 2.4 are caused by dissemination and occur shortly after Phase 3B starts. In addition, for Phase 3B, the level of trading activity remains lower for the 12 weeks after dissemination begins. This persistent reduction is consistent with the Table 2.4 Phase 3B difference-in-differences estimates for 30, 60, and 90-days being similar.

For price standard deviation, the event study plots in Figure 2.5 show a clear drop at dissemination for all three Phases. The coefficients for each Phase are at or above zero before dissemination, and are clearly below zero after dissemination. Importantly, there is a pronounced drop in price standard deviation between week -1 and the first week of dissemination in each of the three Phases. The absence of visual evidence of trends provides support for a causal interpretation of TRACE's effect on price standard deviation.

In summary, the event-study plots in Figure 2.4 show a volume effect only for Phase 3B bonds, while Figure 2.5 shows a decline in price standard deviation for all three Phases. Furthermore, there is no pre-trend in price standard deviation for newly disseminated bonds. This fact provides support for our identification assumptions of incomplete anticipation and no simultaneous non-dissemination related changes in the bond market. Moreover, a large percentage of the overall effect for price standard deviation occurs immediately after dissemination.

2.5.2 Time Trends

Another assumption underlying the difference-in-differences estimates is common parallel trends. That is, we assume that if treated bonds had not changed their dissemination status, their trading behavior would follow the same trajectory as the control group bonds. However, it is possible that trading outcomes for treated bonds follow different trajectories than control bonds. As discussed above in Section 2.4, one reason for this possibility is that the control

Figure 2.4: *Event Study for Weekly Volume/Issue Size*




Figure plots coefficients on Disseminate x Week variables from event study regressions where the dependent variable is weekly Volume/Issue Size. Disseminate is an indicator for a bond that becomes disseminated in the Phase. The x-axes represent weeks, where, 0 is week of March 3, 2003 for Phase 2, 0 is week of October 1, 2004 for Phase 3A, and 0 is week of February 7, 2005 for Phase 3B.

Figure 2.5: *Event Study for Weekly Price Standard Deviation*

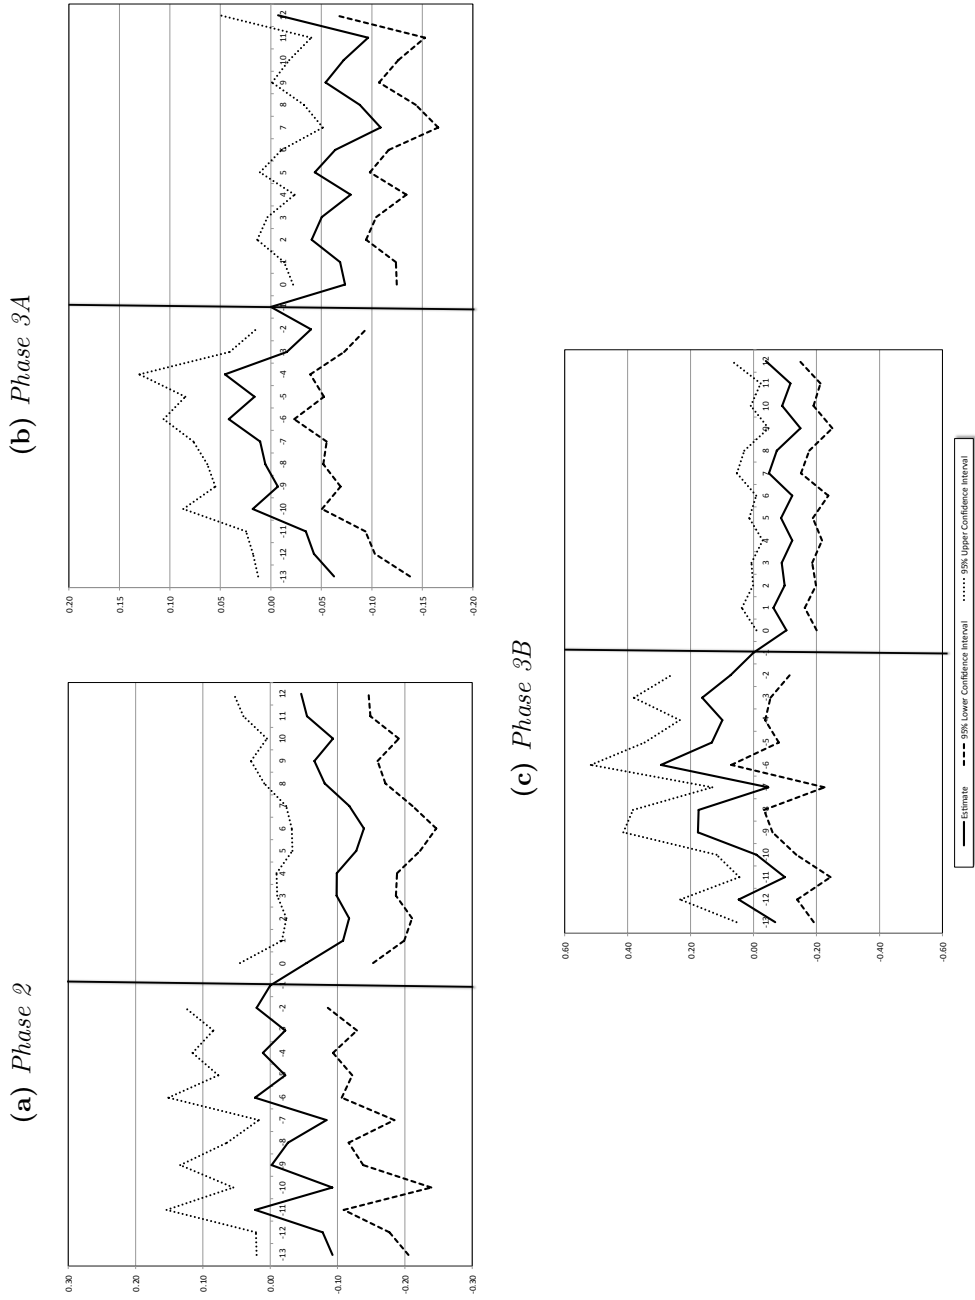


Figure plots coefficients on Disseminate x Week variables from event study regressions where the dependent variable is weekly price standard deviation. Disseminate is an indicator for a bond that becomes disseminated in the Phase. The x-axes represent weeks, where, 0 is week of March 3, 2003 for Phase 2, 0 is week of October 1, 2004 for Phase 3A, and 0 is week of February 7, 2005 for Phase 3B.

bonds have different characteristics than treated bonds, particularly since FINRA uses size and credit ratings to determine Phase classifications.

To relax the common trends assumption, in Table 2.5, we estimate specifications allowing the trade outcomes for bonds to evolve over time depending on whether they are investment-grade or not. Specifically, we estimate models with linear and quadratic time trends by including Phase-specific quadratic functions of time in equation 2.1 as follows:

$$y_{it} = \alpha + \gamma_0 \text{Disseminate}_i + \gamma_{01} \text{Investment Grade}_i \times t + \gamma_{02} \text{Investment Grade}_i \times t^2 + \gamma_1 \text{Post}_t + \lambda \text{Disseminate}_i \times \text{Post}_t + \epsilon_{it} \quad (2.3)$$

where $\text{Investment Grade}_i$ is an indicator for bond ratings of BBB- and above. For each Phase, the variable t starts at zero at the beginning of the time window. For the pooled estimate, we estimate separate Phase-specific trends.

Since equation 2.3 adds more flexible time trends to our difference-in-differences regression, we anticipate a reduction in the precision of the estimates in Table 2.5 compared to Table 2.4. The precision of each significant estimate in Table 2.5 column (2) is lower than that in Table 2.4 (which is repeated for convenience as column (1)). The pooled estimate of volume/issue size although smaller is still significant at the 5% level. The estimate for Phase 2 volume/issue size becomes insignificant with trends. Importantly, the estimate for Phase 3B remains significant at the 1% level.

When estimating equation 2.3 for price standard deviation, Table 2.5 column (6) shows that the results are robust to the addition of trends. For each Phase separately, as well as pooled, the estimates remain negative and significant at the 1% level.

2.5.3 Control Groups

We also address the robustness of the Table 2.4 results by considering two variations on the control group. First, we eliminate Phase 1 bonds from the control group. As discussed above, Phase 1 bonds are larger and more actively traded than bonds in any other Phase. It is therefore possible that the change in trading behavior of Phase 1 bonds does not provide an adequate counterfactual for bonds that change their dissemination status. In columns (3)

Table 2.5: *Alternative Specifications for Difference-in-Difference Regressions of Trading Activity and Price Dispersion*

	Volume / Issue Size (in percent)			Price Standard Deviation				
	Estimate from column (4) of Table 4 (1)	With linear and quadratic trends specific to investment grade for each Phase (2)	Without Phase 1 bonds as controls (3)	Sample matched on size, credit rating, time to maturity, and years since issuance (4)	Estimate from column (8) of Table 4 (5)	With linear and quadratic trends specific to each Phase investment grade for bonds as controls (6)	Without Phase 1 bonds as controls (7)	Sample matched on size, credit rating, time to maturity, and years since issuance (8)
All Three Phases	-0.027*** (0.005)	-0.013** (0.006)	-0.026*** (0.005)	-0.018** (0.007)	-0.077*** (0.009)	-0.075*** (0.009)	-0.080*** (0.010)	-0.077*** (0.012)
Phase 2	-0.026** (0.011)	-0.017 (0.011)	-0.028*** (0.011)	-0.014 (0.013)	-0.070*** (0.015)	-0.086*** (0.015)	-0.073*** (0.017)	-0.079*** (0.017)
Phase 3A	0.003 (0.007)	0.007 (0.007)	0.007 (0.007)	-0.014* (0.008)	-0.058*** (0.009)	-0.053*** (0.009)	-0.062*** (0.010)	-0.077*** (0.016)
Phase 3B	-0.074*** (0.008)	-0.077*** (0.012)	-0.073*** (0.008)	-0.066*** (0.023)	-0.168*** (0.030)	-0.169*** (0.041)	-0.156*** (0.031)	-0.045 (0.051)
H ₀ : Phase effects equal	0.000	0.000	0.000	0.096	0.003	0.006	0.015	0.825
# of Phase 2 bonds	2,538	2,536	2,538	2,536	2,023	2,021	2,023	2,021
# of Phase 3A bonds	11,087	11,052	11,087	4,552	6,319	6,307	6,319	3,014
# of Phase 3B bonds	2,853	2,763	2,853	808	1,129	1,111	1,129	325
# of bond-days	3,410,347	3,379,514	3,314,283	1,867,658	557,057	556,024	478,817	430,962

This table reports estimates of Disseminate x Post from alternative regression specifications. Panel A reports estimates from Phases 2, 3A, and 3B pooled together, while panel B reports estimates for each Phase separately. Robust standard errors clustered by bond and Phase are in parenthesis immediately below the estimates. The sample in columns (2), (4), (6), and (8) excludes unrated bonds. Models with trends in columns (2) and (6) include linear and quadratic functions of time for investment grade and high-yield bonds specific to each Phase. The characteristics used to construct the matched sample in columns (4) and (8) are issue size, credit rating at Phase start, time to maturity at Phase start, and years since issue at Phase start. We divide the sample into four issue size quartiles, and two groups for the other three characteristics: investment grade and high-yield, and above/below the median time to maturity and years since issue. We exclude bonds in cells with 5 or fewer treated bonds or 5 or fewer control bonds. Phase effects equal reports p-values of tests that the three Phase estimates are equal. * significant at 10%; ** significant at 5%; *** significant at 1%

and (7) of Table 2.5, we report estimates for volume/issue size and price standard deviation where Phase 1 bonds are not used as controls. This means that for Phase 2, the control bonds are from Phase 3A and 3B. For Phase 3A and 3B, the control bonds are from Phase 2. The estimates reported in columns (3) and (7) are nearly identical to our base results in columns (1) and (5), respectively.

Second, we construct a matched sample, restricting the treated sample to bonds for which there is a suitable control bond with similar pre-treatment characteristics. The pre-treatment bond characteristics we use to construct the matched sample are issue size, credit rating at Phase start, time to maturity at Phase start, and years since issue at Phase start.²⁵ To construct the matched sample, we divide the sample (which includes Phase 1 bonds) by issue size into four quartiles. For the other three characteristics, we divide in two groups: investment grade and high-yield, above and below the median time to maturity, and above and below the median years since issue. This results in 32 potential cells for each Phase. We exclude a cell if there are either fewer than 5 treated bonds or fewer than 5 control bonds. When we define the matched sample using our four bond characteristics, we cover 99.6% of Phase 2 bonds in our volume sample, but for Phases 3A and 3B the treated sample is only 41.1% and 28.3%, respectively for volume/issue size. For price standard deviation, we cover 99.9% of Phase 2 bonds in our price standard deviation sample, 47.7% of Phase 3A bonds, and 28.8% of Phase 3B bonds.

The estimates for the matched-sample difference-in-differences regression are in columns (4) and (8) of Table 2.5. To control for bond attributes, we add a dummy variable for each cell to equation 2.1, and interact the cell dummy with Post and treated. Their inclusion means that our estimates are a weighted average of the within-cell difference-and-differences estimates. For the matched sample, the volume/issue size estimates in column (4) for the pooled sample and Phase 3B remain negative and significant. Thus, the negative and significant effect of dissemination on volume/issue size documented in Table 2.4 for Phase 3B and the pooled sample is robust to the alternative specifications in columns (2)-(4).

²⁵We eliminate bonds that are unrated from the matched sample.

The price standard deviation results for the matched sample in column (8) are similar to those in columns (5)-(7) for the both the pooled and Phase samples. The only difference worth highlighting is that for Phase 3B, the effect on price standard deviation is no longer significant. This reduction in significance may be due to the small sample size of only 325 treated and 1,582 control bonds. Thus, examining columns (5)-(8) of Table 2.5 shows that the negative and significant effect of dissemination documented in Table 2.4 is robust across all alternative specifications for the pooled sample, and Phases 2 and 3A. The results are also robust for two of the three alternative specifications for Phase 3B.

2.5.4 Alternative Measures of Trading Activity and Price Dispersion

Trading Activity

So far, we've focused our investigation on volume/issue size and price standard deviation as the measures of TRACE's impact on bond trading. Next, we consider some alternative measures of trading activity and price dispersion in Table 2.6 and 2.7, respectively. Both Tables report estimates from the difference-in-differences regressions with 90-day windows used in Table 2.4, but with different outcomes.

As described above, TRACE proponents expected that transparency would increase trading activity, expand market participation, and attract greater retail interest.²⁶ In Table 2.6, we consider volume (not normalized by issue size), the probability of trade, the probability of a large trade, the number of trades, and the average trade size. The probability of trade is the percentage of days a bond trades during our sample period. The odd-numbered columns of Table 2.6 report the average value of the dependent variable for treated bonds in the 90 days before dissemination. The even numbered columns report the difference-in-differences estimate for each of the outcomes.

Before turning to the effects of dissemination on alternative measures of trading activity, we note some important differences in trading activity across Phases in means reported in the odd-numbered columns in Table 2.6. Trade sizes for Phase 3B bonds are quite large, but

²⁶FINRA defines retail trades as \$100,000 or less (Ketchum (2012)).

Table 2.6: Difference-in-Difference Estimates for Alternative Measures of Trading Activity

	Volume		Probability of Trade		Probability of Trade \geq \$1M		# of Trades		Average Trade Size	
	Mean for Disseminated (1)	Estimate (2)	Mean for Disseminated (3)	Estimate (4)	Mean for Disseminated (5)	Estimate (6)	Mean for Disseminated (7)	Estimate (8)	Mean for Disseminated (9)	Estimate (10)
All Three Phases	420,725	-72,439** (24,753) -17.2%	0.197	-0.002* (0.001) -1.0%	A. Pooled 0.046	-0.005*** (0.001) -10.9%	0.677	-0.131*** (0.039) -19.1%	669,198	-63,367*** (15,002) -9.5%
Phase 2	888,352	-33,792 (36,570) -3.8%	0.383	-0.018*** (0.002) -4.7%	B. By Phase 0.088	-0.007*** (0.002) -8.0%	1.344	0.016 (0.047) 1.5%	782,646	4,854 (26,537) 0.6%
Phase 3A	335,026	-96,218** (46,086) -28.7%	0.177	0.015*** (0.002) 9.0%	0.036	0.001 (0.002) 2.8%	0.625	-0.056 (0.037) -9.5%	535,926	-38,038* (19,645) -7.1%
Phase 3B	366,526	-98,344** (45,222) -26.8%	0.116	-0.005* (0.003) -4.3%	0.052	-0.012*** (0.002) -23.1%	0.298	-0.487*** (0.130) -163.3%	1,205,940	-344,005*** (33,039) -28.5%
H ₀ : Phase effects equal		0.428		0.000		0.000		0.001		0.000
# of Phase 2 bonds		2,538		2,538		2,538		2,538		2,194
# of Phase 3A bonds		11,087		11,087		11,087		11,087		7,478
# of Phase 3B bonds		2,853		2,853		2,853		2,853		1,320
# of bond-days		3,410,347		3,410,347		3,410,347		3,410,347		831,000

This table reports estimates of Disseminate x Post for alternative measures of trading activity following the 90-day estimates in Table 2.4. Panel A reports estimates from Phases 2, 3A, and 3B pooled together, while panel B reports estimates for each Phase separately. Robust standard errors clustered by bond and Phase are in parenthesis immediately below the estimates. Volume is the total daily par value of volume, Probability of Trade is 1 if the bond trades at all on the day and 0 otherwise, Probability of Trade \geq \$1M is 1 if there is a bond trade greater than or equal to \$1M and 0 otherwise, # of Trades is the number of trades per day, and Average Trade Size is the average size of the trades in a day, conditional on trading. Mean for disseminated is the 90-day average for newly disseminated bonds immediately before the Phase start. Percentage effects are computed by dividing the estimate by the prior mean. Phase effects equal reports p-values of tests that the three Phase estimates are equal. For the first four outcomes, there are 8,299 control bonds in Phase 2, 2,202 control bonds in Phase 3A, 2,098 control bonds in Phase 3B. For Average Trade Size, there are 6,206 control bonds in Phase 2, 1,876 control bonds in Phase 3A, and 1,785 control bonds in Phase 3B. * significant at 10%; ** significant at 5%; *** significant at 1%

Phase 3B bonds trade infrequently. For instance, the average trade size for Phase 3B bonds is 1,205,940, which is much larger than Phase 2 bonds and more than twice the average size of Phase 3A bonds. Despite this larger trade size, volume for Phase 3B bonds is much smaller than Phase 2 bonds, and approximately the same size as Phase 3A bonds. This is explained by the much lower probability of trading for Phase 3B bonds.

Dissemination causes a significant reduction in volume for the pooled sample and for Phases 3A and 3B separately as seen in column (2). For the pooled sample, there is 17.2% percent reduction in volume / issue after dissemination, significant at the 1% level. For Phase 3A bonds, the reduction is 28.7%, while for Phase 3B bonds, the reduction is 26.8%, both significant at the 5% level. The percentage reduction volume for in Phase 3B is not as large as the percentage reduction in volume/issue size in Table 2.4 and the significance level is lower. This difference may be due to greater skewness for trading volume, caused by idiosyncratic large trades, compared to volume/issue size when Phase 1 are included in the controls. Although not shown in the Table, when we eliminate Phase 1 bonds from our difference-in-differences regression on volume, only the Phase 3B and pooled estimates are negative and both are significant at the 1% level.²⁷

In the next two columns of Table 2.6, we fit models of the probability of any trade and the probability of a trade over \$1 million in size. In the Public TRACE dataset, TRACE censored the reporting of trades greater than \$1 million (for high-yield) and \$5 million (for investment grade). This was due to objections from dealers and certain institutional market participants who claimed that it would be possible to infer their trading positions from the release of large trade sizes and therefore place them at a competitive disadvantage.

Our estimates for the probability of any trade indicate that in the pooled sample, TRACE reduces trading. However, there are significant opposite patterns by Phase. The probability of trade for Phase 2 bonds decreases significantly at the 1% level, the probability of trade for Phase 3A bonds increases significantly at the 1% level, and the probability of trade in Phase

²⁷The estimate for Phase 3B volume without Phase 1 as a control is -96,507.7, similar to our estimate of -98,343.6 in Table 2.6, but the standard error is 19,386.6, much below the standard error in Table 2.6 of 45,222.4.

3B decreases significantly at the 10% level. When we measure of probability of trades over \$1 million in size, the effect for Phase 3A is no longer significantly positive, but the effect for Phase 2 and 3B remain significantly negative. For Phase 3B, the reduction in the probability of a large trade is -0.012, which is a 23.1% reduction from the mean level of 0.052. Thus, these two findings suggest that TRACE's influence on participation, as measured by probability of trade, is not positive as proponents anticipated.

he results for the number of trades are also similar to that for volume/issue size. In column (8), the change in the number of trades for the pooled sample and Phase 3B is negative and significant at the 1% level. Interestingly, the 0.49 reduction in the number of trades in Phase 3B of is greater than the mean number of trades, 0.30, prior to dissemination. The reason for this is that the number of trades for Phase 3B bonds which trade most frequently experience a greater reduction than the number of trades for Phase 3B bonds which trade infrequently.²⁸

We also examine average trade size in columns (9) and (10). Those results repeat the pattern of a significant decline for the pooled result and for Phase 3B. It's worth noting that trade sizes are larger for Phase 3B than in any other Phase. The reduction in trade sizes occurs even though certain infrequently traded Phase 3B bonds were subject to delayed dissemination if their transaction size was \$1 million or greater.²⁹ These results imply that the decline of large trades in Phase 3B play a large role in our overall volume findings.

Finally, in unreported tabulations, we also find that TRACE does not increase the likelihood of retail size trades. For instance, the pooled estimate for the probability of a trade less than \$100,000 is 0.000 with standard error 0.00128. In Phase 2, the estimate is significantly negative, -0.0127 with standard error 0.00218. Hence, TRACE did not increase the likelihood of retail size trades.

²⁸In an unreported analysis, we further investigated the reduction in the number of trades. There is a gradient in the percentage reduction in the probability that a bond trades multiple times a day. The percentage reduction in the likelihood of trading at least 20 times a day is greater than the percentage reduction at least 10 times a day, which in turn is greater than the percentage reduction in the probability of trading at least 5 times a day.

²⁹An infrequently traded bond is one that does not average one or more trades per day over last 20 business days of a 90-day period determined each quarter by NASD.

In summary, the results in Table 2.6 show that volume, probability of a large trade, number of trades, and trade size follow the same pattern as volume/issue size. Thus, TRACE does not appear to have increased market participation even from retail investors.

Price Dispersion

A weakness of our daily price dispersion measure is that since we require at least two trades in a day, it cannot be computed for all bonds. It is possible that TRACE also affects price dispersion for bonds that do not trade at least twice a day. To examine this possibility, in Table 2.7, we consider three additional measures of price dispersion: the intra-day absolute spread (max price minus min price)³⁰, the price standard deviation of all trades in 10-day windows, and the price standard deviation of all trades in 30-day windows. Using the 10-day and 30-day price standard deviation increases our sample sizes slightly. For instance, with the 30-day measure our coverage of Phase 3A bonds is 63.0% and Phase 3B bonds is 42.1% compared to 57.0% and 40.0% respectively with the intraday measure.

The results on other measures of price dispersion in Table 2.7 confirm the price standard deviation results in Tables 2.4 and 2.5. Every measure for the pooled sample and for each Phase is negative and significant. As with daily price standard deviation, the largest effect of dissemination occurs in Phase 3B for all three measures of dispersion. For the absolute spread, a reduction of 39.7 cents represents 28.6% of the average spread pre-transparency. This percentage reduction is similar to the 24.7% reduction for daily price standard deviation. Thus, transparency reduces price dispersion for four different metrics for all Phases.

2.6 Heterogeneity in Credit Rating and Issue Size

While the price dispersion results are consistent across all Phases, the results on trading activity differ for Phase 3B. What is different about the bonds in Phase 3B? FINRA selects the bonds in each Phase using credit rating, issue size, and trading activity. Examining credit

³⁰Using equity data from TAQ, Corwin and Schultz (2012) demonstrate that intraday absolute spread is highly correlated with bid ask spreads and show that it also outperforms other low-frequency spread measures.

Table 2.7: Difference-in-Difference Estimates for Alternative Measure of Price Dispersion

	Absolute Spread		10-Day Price Standard Deviation		30-Day Price Standard Deviation	
	Mean for Disseminated (1)	Estimate (2)	Mean for Disseminated (3)	Estimate (4)	Mean for Disseminated (5)	Estimate (6)
All Three Phases	1.914	-0.182*** (0.022) -9.5%	1.154	-0.100*** (0.013) -8.7%	1.354	-0.140*** (0.019) -10.3%
			A. Pooled			
			B. By Phase			
Phase 2	2.031	-0.172*** (0.036) -8.5%	1.300	-0.082*** (0.025) -6.3%	1.564	-0.104*** (0.037) -6.6%
Phase 3A	1.921	-0.132*** (0.023) -6.9%	1.133	-0.112*** (0.011) -9.9%	1.295	-0.145*** (0.014) -11.2%
Phase 3B	1.386	-0.397*** (0.095) -28.6%	0.943	-0.116*** (0.032) -12.3%	1.283	-0.214*** (0.045) -16.7%
H ₀ : Phase effects equal	0.020			0.529		0.172
# of Phase 2 bonds	2,023			2,083		2,103
# of Phase 3A bonds	6,319			6,788		6,951
# of Phase 3B bonds	1,129			1,172		1,201
# of observations	557,057			220,538		112,084

This table reports estimates of Disseminate x Post for alternative measures of price dispersion following the 90-day estimates in Table 2.4. Panel A reports estimates from Phases 2, 3A, and 3B pooled together, while panel B reports estimates for each Phase separately. Robust standard errors clustered by bond and Phase are in parenthesis immediately below the estimates. Absolute Spread is the maximum price minus minimum price traded for a bond in a day, Price Standard Deviation over 10 Days is the standard deviation of prices for all trades occurring in 10-day bins, and Price Standard Deviation over 30 Days is the standard deviation of prices for all trades occurring in 30-day bins. Mean for disseminated is the 90-day average for newly disseminated bonds immediately before the Phase start. Percentage effects are computed by dividing the estimate by the prior mean. * significant at 10%; ** significant at 5%; *** significant at 1%

rating and issue size in Table 2.2 shows that Phase 3B differs from the other Phases because it is the only Phase with a majority of high-yield bonds. However, there is some overlap of credit rating and issue size between Phases, making it possible to identify whether credit rating or size is the main determinants of the Phase 3B results.

In Table 2.8, we pool the Phases, and split the treated sample by credit rating and issue size. We split credit ratings into investment grade, BBB- or above, and high-yield, BB+ or below. We split issue size into bonds with issue size less than or greater than or equal to \$100 million. These criteria follow FINRA's breakpoints for Phase 2 classification. The control bonds remain the same across columns. The overlap between Phases on credit quality and issue size is shown in Table 2.8. For the 3,164 high-yield bonds in our sample, 634 are from Phase 3A, while the remainder is in Phase 3B. For 9,087 bonds with issue size less than \$100 million, 677 are from Phase 3B, while 8,410 are from Phase 3A and 10 are from Phase 2. Thus, pooling the high-yield sample amounts to combining most of Phase 3B with a portion of Phase 3A, while pooling the small issue size sample amounts to combining most of Phase 3A with a portion of Phase 3B.

The effect of dissemination on volume/issue size on high-yield bonds is a highly significant -0.057, while it is a smaller and less significant -0.013 for investment grade bonds, as shown in columns (1) and (2) of Table 2.8. This 4.4 ratio of effects represents a statistically significant difference as shown by the p-value from the Chi-square test reported below the estimates. Turning to issue size, the effect of dissemination on volume/issue size is primarily driven by bonds with issue size \geq \$100 million. The estimate for bonds with issue size $<$ \$100 million is not statistically significant and close to zero. This is consistent with the results in Table 2.4, which show that Phase 3A bonds do not experience a reduction in trading activity. These bonds by definition primarily have issue size less than \$100 million. Thus, the volume/issue size findings appear to be driven by low credit bonds or bonds with issue size \geq \$100 million.

To examine which feature is more responsible for driving the volume/issue size results, we next report a two-way split of the sample. In column (5) and (6), we split the investment grade sample into small and large issue size bonds. In column (7) and (8), we split the high-yield

Table 2.8: Difference-in-Difference Estimates by Credit Rating and Issue Size

	Investment Grade (1)	High Yield (2)	Issue Size < \$100M (3)	Issue Size ≥ \$100M (4)	Investment grade		High yield	
					Issue Size < \$100M (5)	Issue Size ≥ \$100M (6)	Issue Size < \$100M (7)	Issue Size ≥ \$100M (8)
All Three Phases	-0.013** (0.006)	-0.057*** (0.011)	-0.005 (0.006)	-0.038*** (0.006)	A. Volume / Issue Size 0.010 (0.007)	-0.028*** (0.007)	-0.071*** (0.017)	-0.056*** (0.012)
H ₀ : Effects equal between p-value	(1)-(2) 0.001			(3)-(4) 0.001		(5)-(6) 0.002	(7)-(8) 0.488	
# newly disseminated Phase 2	2,536	0	10	2,526	10	2,526	0	0
# newly disseminated Phase 3A	10,418	634	8,410	2,642	8,382	2,036	28	606
# newly disseminated Phase 3B	233	2,530	677	2,086	81	152	596	1,934
# of bond-days	3,011,573	1,826,801	2,530,591	2,307,783	2,458,151	2,011,931	1,530,949	1,754,361
All Three Phases	-0.058*** (0.009)	-0.124*** (0.017)	-0.053*** (0.010)	-0.084*** (0.010)	B. Daily Price Standard Deviation -0.048*** (0.010)	-0.068*** (0.011)	-0.170*** (0.059)	-0.122*** (0.017)
H ₀ : Effects equal between p-value	(1)-(2) 0.006			(3)-(4) 0.023		(5)-(6) 0.169	(7)-(8) 0.442	
# of Phase 2 bonds	2,021	0	6	2,015	6	2,015	0	0
# of Phase 3A bonds	5,695	612	4,175	2,132	4,154	1,541	21	591
# of Phase 3B bonds	90	1,021	104	1,007	6	84	98	923
# of bond-days	503,844	373,258	385,712	491,390	383,887	440,940	322,808	371,433

This table reports estimates of Disseminate x Post for 90-days by credit and issue size categories pooling together Phases 2, 3A, and 3B. Bonds that are unrated are excluded. Standard errors clustered by bond are in parentheses. In Panel A, there are 8,165 control bonds in Phase 2, 2,202 control bonds in Phase 3A, 2,098 control bonds in Phase 3B. In Panel B, there are 5,512 control bonds in Phase 2, 1,769 control bonds in Phase 3A, and 1,670 control bonds in Phase 3B. P-values reported from Chi-Square tests for equality of estimates between specifications. * significant at 10%; ** significant at 5%; *** significant at 1%

sample into small and large issue size bonds. The estimate for small investment grade bonds is not significant, but the estimate for large investment grade bonds is negative and significant. This estimate, however, is smaller than either estimate for high-yield bonds, which are both negative and similar in size for both small and large issue size bonds. Therefore, it appears that the results for volume/issue are affected more by credit ratings than issue size.

The second panel of Table 2.8 reports on price standard deviation split by ratings and issue size. Each of the estimates is negative and highly significant for all subgroups, but the reduction in price standard deviation is significantly larger for high-yield bonds than for investment grade bonds throughout. When examining issue size, the reduction in price dispersion is only slightly larger for bonds with issue size \geq \$100 million.

Thus, the reason the results on Phase 3B are different than the other Phases is largely because of the high proportion of high-yield bonds in that Phase. Although not shown, the other measures of trading activity in Table 2.6 and the measures of price dispersion in Table 2.7 decrease more for high-yield bonds than for investment grade bonds. Therefore, our initial question in this subsection of why the bonds in Phase 3B behave differently needs to be recast to ask why do high-yield bonds behave differently?

The fact that investment grade and high-yield bonds behave differently is not a surprise. Investment grade bonds trade near par except for price fluctuations due to market interest rate movements. This means that they can be treated as substitutes with one another within credit rating categories. High-yield bonds, even within the same rating category, are not as close as substitutes since they are subject to idiosyncratic, firm-specific risks.³¹ Moreover, some market participants such as pension and mutual funds have rules restricting ownership of high-yield bonds. Furthermore, since investment grade bonds trade more frequently than high-yield, they are also less opaque. For instance, the probability of a trade on any given day (pre-TRACE) is more than three times higher for the investment-grade sample in Phase 2 compared to the mostly high-yield sample in Phase 3B. Given these differences, TRACE

³¹Asquith *et al.* (2013) document significant differences between investment grade and high-yield bonds in the market for borrowing bonds.

probably provided more incremental information on trading activity for high-yield bonds than for investment grade bonds.

In addition, the bond market is a dealer market, so dealer inventory will affect trading levels and the potential impacts of TRACE. Dealers only hold inventory in those bonds with sufficient trading activity to cover their carry cost. Thinly traded bonds may require dealers to have higher spreads to cover their holding costs. Since TRACE reduces price dispersion significantly, the benefit of holding bonds in inventory decreases. TRACE reduces price dispersion the most for high-yield bonds, so the incentive to reduce inventory is strongest for those bonds. Thus, lower trading activity in high-yield bonds post-TRACE may be the result of a supply-side response of dealers.

2.7 NAIC

The evidence so far leaves open the question of TRACE's impact on Phase 1 bonds. TRACE data does not exist before July 2, 2002 when Phase 1 begins; therefore, our analysis of the effects of transparency using trades both before and after dissemination in TRACE is limited to Phases 2, 3A, and 3B. Phase 1 is important because, as discussed above in Section IV, dissemination of Phase 1 bonds may affect the corporate bond market behavior more broadly if transparency in part of the market influences trading in the rest of the market. As described in Section II, Bessembinder *et al.* (2006) examine trading costs in Phase 1 using data from the National Association of Insurance Companies (NAIC). While the NAIC database is not as complete as TRACE because it only contains transaction data for insurance companies, the NAIC data begins in 1994.

In this section, we describe the NAIC data and use that database from January 1, 2000 through December 31, 2006 to examine the effects of Phase 1 of TRACE as well as to verify our results for Phases 2, 3A, and 3B. The NAIC database also contains information about dealer activity not available in TRACE, which we use to examine how TRACE affected dealer market share.

Before using the NAIC data, we first compare it to the TRACE data both for coverage and

to determine whether insurance companies trade differently than the rest of the corporate bond market. According to the Federal Reserve's Flow of Funds statement, insurance companies own 24.6% of outstanding corporate bonds in 2002Q3-2006Q4.³² While several other papers, notably Bessembinder *et al.* (2006) and Campbell and Taksler (2003), have previously used NAIC data, to our knowledge we provide the first direct comparison of the two databases.³³

The NAIC Data Appendix and Tables B1 and B2 describe the NAIC data and how it compares to the TRACE database. Importantly, in the process of making this comparison, we discovered a systematic error in how NAIC's trades are reported. Many NAIC trades are disaggregated and reported as multiple transactions in the NAIC database. Since previous research on the NAIC database (e.g. Bessembinder *et al.* 2006 do not mention this problem of disaggregation, we assume that they treated these multiple transactions as multiple trades, when they are not. This leads to an over-reporting in the number of trades and an under-reporting of the true price dispersion.³⁴

NAIC's reporting requirements require many individual trades to be split into separate records for reporting purposes. For example, insurance companies must separately report bonds purchased and sold in the same calendar year from bonds purchased and held through the end of the year. This means if an insurance company purchases \$1 million par of a bond on January 1, 2001 and sells \$500,000 of this before December 31, 2001 and the remaining \$500,000 in the following year, under NAIC reporting guidelines, this single purchase would be split into two separate purchases of \$500,000 each. If this is treated as two trades, volume is unaffected, but the number of trades is overstated and price standard deviation is understated. A more complete discussion of the misreporting of trades is explained in the NAIC Appendix.

Table A.3 reports the steps we took to process the raw NAIC file into our cleaned NAIC database. We only use those bonds from the NAIC database that are also in the Cleaned

³²Campbell and Taksler (2003) estimate that insurance companies hold between one-third and 40% of corporate bonds.

³³Bessembinder *et al.* (2006) do divide the NAIC database into TRACE and non-TRACE samples, but do not compare trading by NAIC members to trading by non-NAIC members.

³⁴We do not know trade disaggregation changes Bessembinder *et al.* (2006) results. However, since Edwards *et al.* (2007) results are similar using TRACE data, we assume this issue does not change the results substantially.

Historical TRACE database for our analysis.³⁵ Because of the misreporting issue discussed above, Table A.3 reports the total number of transactions from the NAIC database in the column labeled “Ungrouped Trades.” It also reports an estimate of the true number of trades by grouping transactions with identical CUSIP, date, price, and counterparty into a single record with volume summed for the grouping. These are labeled “Grouped Trades” in a separate column in Table A.3. The NAIC data appendix contains more details on this process. From July 2, 2002 to December 31, 2006, the clean NAIC database contains 14,574 bonds. There are 481,135 ungrouped trades, which correspond to 394,679 grouped trades. This compares to 21,217,807 trades on 30,958 bonds in the Cleaned Historical TRACE database over the same period.

Table A.4 compares the cleaned NAIC and TRACE datasets by Phase and shows that insurance companies trade very differently than the rest of the corporate bond market for the same time period and universe of bonds.³⁶ It compares the number of bonds covered, the trading volume, the number of trades, and the trade sizes in both cleaned databases. A high percentage of Phase 1, Phase 2, and Phase 3B TRACE bonds are contained in NAIC (94.2%, 81.7%, and 72.7% respectively). NAIC contains 42.2% of Phase 3A TRACE bonds. NAIC volume, however, is much smaller percentage of TRACE volume for all Phases. For Phase 1 bonds, during the 90 days after the announcement of the Phase, the NAIC volume is 6.3% of comparable TRACE volume. For Phase 2, 3A, and 3B, NAIC volume is 11.5%, 7.2%, and 4.4% of TRACE volume respectively.

The number of NAIC trades is an even smaller percentage of TRACE trades for Phases 1, 2, and 3A. In Phase 3B, the percentage of trades in TRACE is lower than the percentage of volume. This means for Phases 1, 2, and 3A, TRACE trades are usually larger than rest of the market. Grouped NAIC trades are larger than TRACE trades, on average, by a factor of 4.1

³⁵45,902 bonds in the NAIC database are not in the Cleaned Historical TRACE database. A large fraction of these bonds are SEC Rule 144a bonds. SEC Rule 144A bonds are not covered by TRACE during our sample period.

³⁶In order to examine trading in Phase 1 bonds before the start of Phase 1, we use NAIC data from the period January 1, 2000 until July 1, 2002. We only compare trading activity between the databases during the TRACE period, which starts July 2, 2002.

in Phase 1, 2.1 in Phase 2, 2.2 in Phase 3A. Thus, NAIC is a small share of TRACE's volume and trades, but the average size of NAIC trades is often larger than the average TRACE trade.

Table A.4 also compares price standard deviation between NAIC and TRACE. The standard deviation for NAIC trades is typically much smaller than for TRACE. This is true for each Phase using the ungrouped NAIC trade database, and is true for Phases 1, 2, and 3B using the grouped NAIC trade database. It's worth noting that the NAIC price standard deviation is measured using far fewer CUSIPS and bond-days. In Phase 2 for example, we measure TRACE price standard deviation for 2,130 CUSIPs and 40,713 bond-days, while we only measure it for 261 CUSIPs and 481 bond-days in NAIC using grouped trades. This restricts our ability to draw strong inferences about price standard deviation using the NAIC sample.

We conclude that the NAIC database represents a small fraction of the trading in the corporate bond market covered by TRACE. Summing volume across all four Phases in the 90-days after Phase start, NAIC volume is only 7.6% of total TRACE volume. NAIC trades are also typically larger than those in the TRACE database. It is therefore possible that the effects of transparency may manifest themselves differently in TRACE than in NAIC. As a consequence, conclusions drawn about TRACE from the NAIC dataset may not be representative of the overall corporate bond market.

2.7.1 Trading Activity and Price Dispersion Using NAIC data

Table 2.9 reports volume/issue size and price standard deviation (both grouped and ungrouped) for 90 days before and after each Phase using the NAIC database. It also reports in column (9) coefficients from difference-in-differences regressions similar to those reported in Table 2.4. The Phase 1 difference-in-differences results in column (9) of Table 2.9 are not significant for either volume/issue size or price standard deviation. In addition, the within-bond comparisons, shown in columns (5) and (6), are mixed. That is, the fraction of Phase 1 bonds that experience a decrease in volume/issue size is greater than the fraction experiencing an increase, but the fraction of Phase 1 bonds that experience a decrease in price standard deviation is less.

There are several possible reasons for the lack of significant or consistent results for Phase

Table 2.9: NAIC Trading Activity and Price Dispersion for the 90-Day Window Around Phase Start

	90-Day Mean		90-Day Median		Before > After		Before = After	
	Before (1)	After (2)	Before (3)	After (4)	Before (5)	After (6)	zero (7)	non-zero (8)
A. Trading Activity								
Volume / Issue Size								
Phase 1 (N=334)	0.06%	0.06%	0	0	51.50%	44.00%	4.50%	0.0%
Phase 2 (N=2,294)	0.05%	0.04%	0	0	44.10%	35.30%	20.70%	0.0%
Phase 3A (N=4,983)	0.02%	0.02%	0	0	24.60%	25.30%	50.10%	0.0%
Phase 3B (n=2,319)	0.01%	0.01%	0	0	21.30%	15.70%	63.00%	0.0%
B. Price Dispersion								
Daily Price Standard Deviation (ungrouped)								
Phase 1 (N=253)	0.22	0.26	0.10	0.12	36.0%	59.0%	5.0%	0.0%
Phase 2 (N=392)	0.17	0.26	0.13	0.23	45.0%	36.0%	20.0%	0.0%
Phase 3A (N=464)	0.21	0.39	0.19	0.33	35.0%	35.0%	30.0%	0.0%
Phase 3B (N=109)	0.23	0.04	0.19	0.04	32.0%	25.0%	43.0%	0.0%
Daily Price Standard Deviation (grouped)								
Phase 1 (N=213)	0.36	0.42	0.26	0.27	35.0%	65.0%	0.0%	0.0%
Phase 2 (N=261)	0.24	0.40	0.21	0.37	47.0%	46.0%	8.0%	0.0%
Phase 3A (N=273)	0.43	0.70	0.40	0.67	49.0%	37.0%	14.0%	0.0%
Phase 3B (N=65)	0.38	0.04	0.31	0.04	45.0%	25.0%	31.0%	0.0%

This table reports 90-day before and after comparisons of trading activity and price dispersion using data from the National Association of Insurance Commissioners (NAIC). Average daily volume is averaged over all bond-days in either the 90 calendar days before or after the Phase start. If a bond does not trade, it contributes zero daily volume for that day. For average daily price standard deviation, the sample of bonds is restricted to bonds where there is at least one day in the 90 days before the phase start with at least two trades and there is at least one day in the 90 days after the phase start with at least two trades. After computing the within-day price standard deviation for each bond for all days with at least two trades, we average across these days during either the 90 days before or after the phase start. Reported average daily price standard deviation is the average across these bonds. Ungrouped price standard deviation does not combine trades into a single observation, while grouped price standard deviation is based on combined trades to deal with split reporting issues described in the data appendix. Column (9) reports estimates of the Disseminate x Post in difference-in-difference regressions for 90-days, which parallel those reported in Table 2.4. N refers to the number of bonds that change dissemination status in the Phase. * significant at 10%; ** significant at 5%; *** significant at 1%

1. It may be that TRACE has no effect on Phase 1 bonds. It may also be that the insurance segment of the corporate bond market, which NAIC measures, behaves very differently than the remainder of the market. It may also be because the amount of trading captured by NAIC is so much smaller than the entire corporate bond market covered by TRACE, making it difficult to detect changes due to dissemination.

The results for Phases 2, 3A, and 3B provide further evidence on these alternatives explanations. The NAIC price standard deviation results in column (9) for Phases 2 and 3A in Table 2.9 are not significant, while all of our price dispersion results using the TRACE database found significant declines. Moreover, in Phase 3B, where the TRACE results on trading activity and price standard deviation are strongest, the corresponding NAIC estimates are marginally significant. Thus, the lack of significant Phase 1 NAIC results does not necessarily imply that TRACE did not have an effect in Phase 1.

2.7.2 Dealer Trading Activity

The NAIC database contains additional information not available in TRACE. In particular, it identifies the counterparty dealer opposite the NAIC member for each trade. Bessembinder *et al.* (2006) use this data to examine dealer concentration ratios during Phase 1. Though it only represents dealer trades with insurance companies, this data provides an opportunity to measure how dealers are affected by dissemination. The NAIC Data Appendix describes how we compute trading activity by dealer. We employ our difference-in-differences design to examine dealer volume and the number of trades across all four Phases. We present results for both the top 4 and top 8 dealers. The top 4 dealers cover 37.9% of volume and 32.7% of trades. The top 8 dealers cover 68.4% of volume and 58.6% of trades.³⁷

Table 2.10 reports difference-in-differences estimates of the effect of TRACE on dealer volume and number of trades. For each Phase, we compute the par value of trades and count the number of trades that a counterparty was party to during the 12 weeks before and after

³⁷We explored other divisions such as top 2, top 5, and top 10, but all of the results are qualitatively similar to what we describe herein.

the Phase start. We examine weekly volume and number of trades because NAIC trades less frequently. Across Phases, there is a 15.3% reduction in par volume for each of the top 4 dealers due to TRACE and a 15.6% reduction for each of the top 8 dealers. When examining dealer volume estimates by Phase, there is a significant drop in volume traded for both top 4 and top 8 dealers in Phases 1, 3A, and 3B. There is also a reduction in Phase 2, but it is not significant. The analysis on the number of trades per dealer is similar.

Overall, the results indicate that trading activity between dealers and insurance companies is rebalanced away from the largest dealers due to TRACE. If this result holds for the entire corporate bond market, this would indicate that TRACE, although reducing overall trading activity, also leveled the playing field between the largest dealers and the remaining dealers.

2.8 Conclusions and Implications

The introduction of TRACE, which was implemented in four Phases over a three-and-a-half year period, combined with the availability of trading records before and after dissemination, provides a unique opportunity to study how markets respond to transparency. This paper finds that mandated post-trade transparency in the corporate bond market leads to an overall reduction in trading activity. No sample of bonds in any Phase experiences an increase in trading activity and Phase 3B bonds experience a large and significant reduction. For that group, TRACE reduces trading activity by 41.3% in the 90 days following the dissemination of price and volume information. This finding is robust across different measures of trading activity and alternative regression specifications. Event studies support a causal interpretation of our findings since the decrease occurs immediately after the start of dissemination.

Price dispersion also decreases due to TRACE. This decrease is significant across bonds that change dissemination in Phases 2, 3A, and 3B, but is largest, 24.7%, for Phase 3B bonds. This finding is also robust across different measures of price dispersion and alternative regression specifications. Moreover, event studies show that the fall in price dispersion occurs immediately after the start of dissemination. It is important to note, if the transparency introduced in Phase 1 affects bonds that become transparent in subsequent Phases, our estimates are probably

Table 2.10: Estimates of Transparency Effects on NAIC Dealer Market Share for 12 Week Window Around Phase Start

	Par Volume (\$M)				Number of Trades			
	Weekly Average Before		Estimate		Weekly Average Before		Estimate	
	Top 4 Dealers (1)	Top 8 Dealers (2)	Top 4 Dealers (3)	Top 8 Dealers (4)	Top 4 Dealers (5)	Top 8 Dealers (6)	Top 4 Dealers (7)	Top 8 Dealers (8)
All Four Phases	87.59	81.87	-13.36*** (4.97)	-12.74*** (3.37)	35.04	33.97	-3.08*** (0.96)	-2.87*** (0.792)
			-15.3%	-15.6%			-8.8%	-8.4%
	A. Pooled							
Phase 1	112.64	115.13	-28.06*** (8.19)	-20.22** (8.88)	40.46	41.88	-6.20*** (1.92)	-4.28** (1.82)
			-24.9%	-17.6%			-15.3%	-10.2%
	B. By Phase							
Phase 2	98.50	92.02	-2.648 (12.56)	-5.283 (7.41)	41.06	38.57	-3.66** (1.66)	-1.84 (1.47)
			-2.7%	-5.7%			-8.9%	-4.8%
Phase 3A	126.88	108.31	-20.09*** (7.64)	-22.40*** (4.22)	49.52	46.20	-1.04 (1.95)	-4.39** (1.85)
			-15.8%	-20.7%			-2.1%	-9.5%
Phase 3B	12.34	12.04	-2.66** (1.27)	-3.07*** (0.87)	9.13	9.22	-1.43** (0.72)	-0.97* (0.50)
			-21.5%	-25.5%			-15.7%	-10.5%
H ₀ : Phase effects equal			0.003	0			0.090	0.120
# of Phase 1 counterparties			79	79			79	79
# of Phase 2 counterparties			81	81			81	81
# of Phase 3A counterparties			84	84			84	84
# of Phase 3B counterparties			83	83			83	83
# of counterparty-weeks			7,848	7,848			7,848	7,848

This table reports estimates of Disseminate x Post for par volume and the number of trades for counterparties in the NAIC database in a difference-in-difference regression following Table 2.4. Panel A reports estimates from Phases 1, 2, 3A, and 3B pooled together, while panel B reports estimates for each Phase separately. Robust standard errors clustered by bond and Phase are in parenthesis immediately below the estimates. Top 4 and 8 dealers are computed based on rankings of dealers of the total par volume of all trades between 2000-2001. The dealers in the top 4 and 8 are identical if the ranking is based on number of trades between 2000-2001. The time period is 12 weeks before and after dissemination. Across the sample, there are a total of 87 composite counterparties constructed from the NAIC dataset. Each dependent variable is a counterparty-week, corresponding to all of the trades with the counterparty among bonds in the Phase for the week. Weekly average before corresponds to the 12 week average for par volume or number of trades for counterparties immediately before the Phase start. Percentage effects are computed by dividing the estimate by the prior mean. * significant at 10%; ** significant at 5%; *** significant at 1%

lower bounds on TRACE's overall impact.

To further investigate how bond characteristics affect our results, we examine trading activity and price dispersion for samples with the same credit quality and issue size across Phases. We find that the credit quality is the most consistent factor in explaining the reduction in trading activity. High-yield bonds experience a significantly greater reduction in trading activity than investment grade bonds. Our results confirm the view that transparency has a limited impact on the trading activity of the most liquid and investment grade segment of the market. Moreover, our results show that ignoring the less actively traded and high-yield bonds in Phase 3B leads to an incomplete account of TRACE's effect on trading activity.

One possible reason TRACE has different effects on high-yield bonds is that pre-TRACE trading in high-yield bonds may be relatively more opaque than trading in investment-grade bonds. As a result, TRACE may provide more incremental information and thus cause larger change in the high-yield market. A second possible reason is that the lower trading activity in high-yield bonds post-TRACE may be the result of a supply-side response of dealers. Price dispersion falls more for high-yield bonds post-TRACE. In addition, high-yield bonds trade less frequently than investment grade bonds pre-TRACE. The fact that there is a large reduction of price dispersion for thinly traded high-yield bonds may result in lower spreads and thus cause dealers to hold less inventory. This in turn may result in a decrease in trading activity.

There are several welfare implications of increased transparency in the corporate bond market. One consequence is that it may change the relative bargaining positions of investors and dealers, allowing investors to obtain fairer prices at the expense of dealers. The reduction in price dispersion should allow investors and dealers to base their capital allocation and inventory holding decisions on more stable prices. Therefore, the reduction of price dispersion likely benefits customers and possibly, but not necessarily, dealers.

The implications of a reduction in trading activity are not as clear. Whether a reduction in trading activity is desirable depends on why market participants trade. A decrease in trading activity may be beneficial if much of the trading in a bond is unnecessary "noise" trading. On the other hand, if most trading is information-based, a decrease in trading activity may

slow down how quickly prices reflect new information. In addition, if the decrease in trading activity is the result of dealers' unwillingness to hold inventory, transparency will have caused a reduction in the range of investing opportunities. That is, even if a decline in price dispersion reflects a decrease in transaction costs, the concomitant decrease in trading activity could reflect an increased cost of transacting due to the inability to complete trades.

Our results on the corporate bond market have two major implications for the current and planned expansions of mandated market transparency. The implicit assumption underlying the proposed TRACE extensions and the use of TRACE as a template for regulations such as Dodd-Frank is that transparency is universally beneficial. First, it is not clear that transparency for all instruments is necessarily beneficial. Overall, trading in the corporate bond market is large and active, although, as seen, not comparable across all types of bonds. Many over-the-counter securities are similar to the bonds FINRA placed in Phase 3B. That is, they are infrequently traded, subject to dealer inventory availability, and trading in these securities is motivated by idiosyncratic, firm-specific information. Therefore, the expansion of TRACE-inspired regulations, such as those for 144a bonds, asset- and mortgage-backed securities, and the swap market, may have adverse consequences on trading activity and may not, on net, be beneficial.

Second, our results indicate that transparency affects different segments of the same market in different ways. As a consequence, our results provide empirical support for the view that not every segment of each security market should be subject to the same degree of mandated transparency. Both academic commentators (French *et al.* 2010, Acharya *et al.* 2009) and leading industry associations (e.g., Forum 2013) have articulated this position. Despite these recommendations, the expansion of transparency by the Commodity Futures Trading Commission (CFTC) in various swap markets, i.e. interest rate, credit index, equity, foreign exchange and commodities, in December 2012 and February 2013 was immediate for all swaps in those markets. This stands in sharp contrast to FINRA's cautious implementation of TRACE in Phases. Our results on the effect of transparency in the corporate bond market suggest that the extension of mandatory transparency to all markets may make it more difficult to transact in some of those markets.

Chapter 3

Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale

3.1 Introduction

New technologies are important contributors to economic growth¹, but little is known about how firms learn to profitably use them. While there is longstanding evidence that firms learn from their own experiences (learning-by-doing), and from others (social learning), the specific actions that firms actually take in learning are not well understood. Models of learning predict that firms efficiently analyze information about new technologies, invest in experiments to create new information, and incorporate information generated by other firms.² However, to test these models, it is necessary to observe data on the information that firms have, which is difficult to acquire in many empirical settings. This paper tests predictions of learning models for the first time, using data on oil companies that employ hydraulic fracturing (fracking) in

¹See, for example, Arrow (1962), Romer (1986) and Kogan *et al.* (2012)

²See Aghion *et al.* (1991) in the single agent context and Bolton and Harris (1999) in the multi-agent context.

the North Dakota Bakken Shale. The data covers operational choices, profits, and measures of the information firms had when making choices. The oil companies in this data learn to use fracking more profitably over time, but are slow to respond to new information, avoid experiments and underutilize data provided by their competitors.

Fracking is a useful context to study learning behavior in firms. The profit maximizing choice of fracking inputs may vary across drilling locations in unpredictable ways, so firms must empirically learn this relationship over time and change their behavior accordingly. In North Dakota, firms can learn about fracking from a wealth of publicly available information. Regulators collect and publicly disseminate unusually detailed, well-specific information about oil production and fracking input choices. Moreover, regulators delay dissemination until 6 months after a well is fracked, making it possible to measure differences in knowledge about fracking across firms. The industry is not concentrated, which motivates studying learning as a single agent problem. During the time period I study, there are 70 active firms, the market share of the largest firm is only 13% and the combined share of the five largest firms is under 50%. The two main inputs to fracking, sand and water, are commodities, as is the output of fracking, crude oil. The unique regulation and industry structure make fracking in the Bakken shale an unusually compelling setting for studying learning in firms. Moreover, the stakes in fracking are large. Using a production function, I estimate that the average NPV of profits per well for actual fracking choices is about \$12.8 million, while the average profit for each well's most profitable choice is \$24.5 million. Since the regulator in North Dakota expects that 40,000 wells will eventually be fracked over the next 18 years, the potential for lost profits from inefficient learning is substantial.³

Learning-by-doing and social learning are both important in this context. Between 2005 and 2006, the average well is fracked by a firm that had fracked only a single well before. By 2011, the average well is fracked by a firm that had previously fracked 117 wells. Thus, firms can learn from an increasing amount of their own experience. However, North Dakota's disclosure laws make it possible for firms to study their competitors' data. Between 2005 and

³See <https://www.dmr.nd.gov/oilgas/presentations/NDOGCP091013.pdf>

2006, the average well is fracked by a firm that can observe 10 wells previously fracked by other firms, a number which rises to 1,783 in 2011. As a result, most of the information firms have comes from others, and firms have the ability to socially learn.

The data I collect from the regulator in North Dakota is well suited to estimate the relationship between location, fracking, and oil production. I observe the complete operating history of every firm and every well they frack in the Bakken Shale between January 2005 and December 2011 (70 firms and 2,699 wells), so there is no possibility for survivorship bias. The data contains precise measurements of a well's production, location and most important fracking inputs, so there are no endogenous omitted variables. Moreover, the engineering requirements for wells drilled into the Bakken prevent firms from selecting observed fracking inputs on the basis of information I do not observe. Thus, the standard endogeneity problem in production function estimation is unlikely to be a concern.

Using the data I collect, I semi-parametrically estimate a production function for fracking which represents what firms need to learn. These estimates show that amount of oil in the ground and the sensitivity of its production to fracking both vary over space, a result that is consistent with geological theory and data. Estimates made using subsets of the data that were available to firms when they were fracking have qualitatively similar results, suggesting that firms could have used this data to make informed fracking decisions. The estimated production function fits the data well and is stable across robustness tests.

I use this production function to measure how quickly firms learn. Wells fracked in 2005 capture only 16% of the profits that optimally fracked wells would have produced. However, profit capture grows almost monotonically over time, with firms capturing 68% of maximal profits in 2011. This growth is driven by improved fracking input choices, with firms gradually increasing their use of sand and water towards optimal levels over time. I interpret this upward trend in the profitability of fracking input choices as evidence for learning.

Existing research measures learning from upward trends in *productivity*, or residual production that is not explained by input choices. I test for productivity based learning by analyzing the growth of estimated production function residuals over time. Wells fracked in 2011 are

34% more productive than wells fracked in 2005, suggesting some role for productivity-driven learning. However, the majority of the growth in productivity occurs by 2008, and there is no statistically significant difference in productivity between 2008 and 2011. This contrasts with the fraction of profits captured, which increases monotonically over time, and from 44% to 67% between 2008 and 2011. Thus, during 2008-2011, when 95% of wells in my data are fracked, there is little productivity growth, even though there is substantial growth in the fraction of profits captured. These results help clarify the difference between models of learning in which knowledge is a direct input in the production function, and a model of learning about the production function itself.

To see if firms are using their information to make better fracking choices over time, I estimate *ex ante* production functions for each well, using the subset of the data that firms had when they were making choices. I use these estimates to compute *ex ante* profits. Though firms capture 76% of *ex ante* optimal profits in 2007, they capture only 68% in 2011. The fraction of *ex ante* profits falls because initial fracking input choices are close to the (then) estimated optimal levels, but optimal levels subsequently change more quickly than choices do.

Theory predicts that firms may sacrifice estimated profits in the current period by experimenting in order to generate information for the future. To test if experimenting behavior can rationalize the decline in the fraction of estimated *ex ante* optimal profits captured, I estimate a simple model of fracking input choice under technology uncertainty. In this model, firms have preferences over the expectation and standard deviation of their *ex ante* estimates of profits for a fracking input choice. If firms are experimenting, they should be empirically more likely to choose inputs with higher standard deviations of profit. I do not find support for this theory. Firms are more likely to select fracking designs with higher expected profits and *lower* standard deviation of profits. Firms are indifferent between a \$0.60-\$0.98 increase in expectation of profits and a \$1 reduction in the standard deviation of profits.

My calculation of the expectation and standard deviation of profits assumes that firms equally learn from their own and others' experiences. However, firms may treat the social portion of their data differently than the data they directly experience, and in the process

form different estimates of profits than what I calculate. To account for this possibility, I modify my fracking input choice model to allow for weighted production function estimates. I use this model and data on firms' choices to estimate the weight they place on their own experiences relative to their competitors' experiences. Most firms place more weight on their own experiences than their competitors' experiences. Even after controlling for weighted estimates, firms still prefer fracking choices with lower standard deviations and higher means.

This paper finds that firms are reluctant to experiment and ignore valuable data generated by their competitors. These firms are not unsophisticated or under-incentivized. They have access to capital markets, are managed by executives with engineering and business education and are the primary equity holders in the wells they frack. These findings stand in contrast to some theories of efficient learning behavior by rational agents, which predict that firms will take experimental risk and learn from all the information they have.

In addition to its usefulness as a laboratory to study learning, fracking plays a prominent role in current public policy debates about growing oil production and its effects on the environment. The US EIA reports that fracking has caused national oil production to grow 22% since 2009, reversing almost two decades of declines.⁴ There is early evidence that fracking-driven resource booms have affected housing prices⁵ and local banking markets.⁶ However, there are growing concerns about the potential for fracking to negatively affect the quantity and quality of local ground water supplies,⁷ which the US EPA is currently studying.⁸ In response to these concerns, federal regulators have proposed significant increases to disclosure requirements for fracking operations.⁹ Though this push for increased transparency around

⁴<http://www.eia.gov/todayinenergy/detail.cfm?id=13251>

⁵Muehlenbachs *et al.* (2012) find that housing prices increase after the introduction of fracking to a community, except for houses that depend on groundwater.

⁶See Gilje (2012)

⁷See Vidic *et al.* (2013) for an overview

⁸See <http://www2.epa.gov/hfstudy>

⁹See Deutsch (2011).

fracking is driven by environmental concerns, new disclosure regulations may also have an impact on learning by increasing the availability of data.

Finally, the Bakken Shale unlikely to be the last oil and gas formation where fracking and the learning it requires play an important role. Fracking is currently in use in the Eagle Ford and Barnett Shales in Texas, the Woodford Shale in Oklahoma, and several locations in Canada. International oil companies are now developing shale resources in Argentina, Poland and China. The results of this paper may be useful to both policy makers and oil & gas companies alike in regulating access to information and understanding the benefits of more efficient learning behavior.

3.1.1 Related literature

Firms in many industries and time periods have become more productive by learning from their own experiences. Researchers studying the manufacturing of World War II ships (Thornton and Thompson 2001), aircraft (Benkard 2000) and automobiles (Levitt *et al.* 2012) have documented an important empirical regularity: with the same inputs, firms are able to produce more output as they accumulate experience in production.¹⁰ That is, they learn by doing (LBD). The LBD result that productivity is correlated with experience suggests that the knowledge embedded in this experience is a direct input to the production function. Changes over time in capital, labor and materials are thus interpreted as profit-maximizing responses to increases in productivity, not changes in specific knowledge. In this paper, I instead assume that the production technology itself is initially unknown and that experience has no direct impact on production. As firms accumulate experience in fracking, they acquire more data about the fracking production function, perform inference on this data, and make more profitable input choices on the basis of their inference. This is similar to the approach taken by Foster and Rosenzweig (1995) and Conley and Udry (2010) in the development literature.

Economic theory predicts that when firms are learning about a new technology, they face a tradeoff between “exploration” and “exploitation” (or experimentation). Firms may actively

¹⁰This phenomenon has also been observed by Anand and Khanna (2000) in the corporate strategy setting.

learn by experimenting with fracking input choices that have highly uncertain profits or passively learn by exploiting choices with high expected profits. Except in the simplest theory models, the optimal amount of experimentation and exploitation is a challenging problem to solve. However, most models of learning predict that forward-looking firms will always do some experimenting. In the single agent context, Aghion *et al.* (1991) show that forward-looking firms will almost always do some exploration. Bolton and Harris (1999) find a similar result in the multi-agent context. Wieland (2000) employs computational methods to characterize the costs and benefits of exploration, finding that firms who only exploit can get stuck, and repeatedly choose suboptimal actions. To my knowledge, this paper is the first to empirically measure the amount of experimenting that firms perform in a learning situation.

This paper adds to a wide literature documenting the existence and importance of social learning between firms. Most of this evidence is in agricultural settings. Ryan and Gross (1943), Griliches (1957) and Foster and Rosenzweig (1995) demonstrate that farmers learn about the benefits of adopting new technologies from the experiences of their neighbors. Conley and Udry (2010) show that farmers in Ghana learn about the efficient use of fertilizer from other farmers in their social networks, demonstrating that social learning in agriculture is not limited to the adoption decision. Social learning has also been observed in manufacturing. During the construction of WWII ships, Thornton and Thompson (2001) find that firms benefited from accumulated experience by other firms. Similarly, Stoyanov and Zubanov (2012) find evidence that firms in Denmark became more productive after hiring workers away from their more productive competitors.

Finally, this paper is complementary to the existing literature on learning behavior by oil and gas companies. Levitt (2011) shows that the observed temporal and spatial patterns of the oil exploration process match the predictions of a forward-looking learning model. In a study of offshore drilling, Corts and Singh (2004) show that as oil companies gain experience with their service contractors, they learn to trust them and tend to select low-powered contracting terms. Kellogg (2011) studies this phenomenon in the on-shore setting and shows that oil companies and their service contractors jointly learn to be more productive in drilling as they

accumulate shared operating experience.

The remainder of the paper is as follows. In Section 3.2, I provide institutional background on fracking in North Dakota and describe the data I have on operational choices, production results and information sets. Next, in Section 3.3, I estimate a production function model of fracking and evaluate its ability to predict oil production. In Section 3.4, I use the production function estimates to test if firms learned to make more profitable fracking choices over time. In Section 3.5, I specify and estimate the model of fracking input choice under technology uncertainty. Finally, I conclude in Section 3.6.

3.2 Institutional Background and Data

3.2.1 Fracking and US Oil Production

The hydraulic fracturing of shale formations, like the Bakken, has had a profound impact on the fortunes of energy producing states and the US as a whole. In 2009, the US Energy Information Administration reported that national oil production grew 6.8% year-over-year, the first increase in over two decades.¹¹ This trend has continued and between 2009 and 2012, national oil production increased 21.7%. Three states represent the majority of this growth: Texas, Oklahoma and North Dakota. This paper focuses on what has happened in North Dakota.

In March 2012, North Dakota surpassed Alaska to become the second most prolific oil producing state in the US, after Texas. Between January 2005 and July 2013, oil production in North Dakota increased from 93,000 barrels (bbl) per day to 874,000 bbl per day. During the same time period, total US oil production increased from 5.63 million bbl per day to 7.48 million bbl per day, meaning that increased production in North Dakota amounted to 42% of the net increase in total production. Though production increased in Texas and Oklahoma as well, it is striking that North Dakota went from producing less than 2% of national oil

¹¹See the EIA Annual Energy Review, 2009. <http://www.eia.gov/totalenergy/data/annual/archive/038409.pdf>

production to almost 12% in the span of 8 years.¹² This vast expansion in North Dakotan oil production coincided with the introduction of fracking to the Bakken Shale formation.

3.2.2 The Bakken Shale and Hydraulic Fracturing

The Bakken Shale spans 200,000 square miles in North Dakota, Montana and Saskatchewan.¹³ It lies 10,000 feet underground and contains 3 distinct layers: the upper Bakken member (a shale layer), the middle Bakken member (a layer of sandstone and dolomite), and the lower Bakken member (also a shale layer). The US Geological Survey estimates that the upper and lower shales together contain 4.6 billion bbl of recoverable oil.¹⁴ Though the middle Bakken member is not formed from organic material and as such does not generate any oil of its own, firms typically drill horizontally through it and use hydraulic fracturing, or “fracking”, to make contact with the oil bearing shales above and below, as shown in Figure 3.1.

Fracking is the process of pumping a mix of water, sand and chemicals into a well at high pressures. The high pressure of the mix fractures the surrounding rock and the sand in the mix props those fractures open.¹⁵ The fractures created by fracking the middle Bakken radiate outwards into the upper and lower Bakken shales, as shown in Figure 3.1. These fractures both serve as a conduit between the wellbore in the middle Bakken and the upper and lower shales, and also increase the permeability of the upper and lower shales.

Permeability is a geological measure of the ease at which oil naturally flows through rock. The upper and lower shales are unusually impermeable, making it impossible for the oil they contain to naturally reach a wellbore drilled through the middle member. Without fracking,


¹²Texas also experienced production significant production increases during that same time period, though from a much higher base level (from 1.08 million bbl per day to 2.62 million bbl per day, a 143% increase). Much of this increase can also be attributed to the technology changes described here. Operators applied fracking technology successfully to the Eagle Ford, Permian and Barnett shales.

¹³See Gaswirth (2013)

¹⁴See Gaswirth (2013)

¹⁵Chemicals reduce mineral scaling, inhibit bacterial growth, reduce wear and tear on fracking hardware and increase the buoyancy of sand in the fracking mixture. See <http://www.fracfocus.org> for an overview.

Figure 3.1: Diagram of a Hydraulically Fractured Bakken Shale well



Adapted from Hicks (2012)

wells drilled into the middle member will not produce profitable quantities of oil.¹⁶ After fracking, oil inside the lower and upper shales can more easily travel through the new fractures into the wellbore in the middle member.

Firms choose how much water and sand to use in fracking and this choice can have a large impact on the profitability of a well. Wells fracked with more sand and water may produce more oil than wells fracked with less, but fracking is expensive, and water and sand represent the bulk of this expense. In 2013, the reported costs of fracking range from \$2-5 million per well, out of total well costs of \$9 million.¹⁷ Thus, to maximize profits, firms must balance the benefits of sand and water use in fracking with their costs. This requires firms to understand the relationship between oil production and fracking inputs, and it is unlikely that firms initially knew this relationship. The first Bakken wells to be developed with fracking were not drilled until 2005, and at the time, the firms developing those wells had limited experience in fracking shale formations.¹⁸ Without prior experience, firms had to learn how to use fracking by doing it themselves or by studying their competitors.

There is now a growing literature about best practices in fracking. Petroleum engineers have found that wells fracked with more water and sand are often more productive than similar wells with less aggressive fracking treatments.¹⁹ However, there is also evidence that the relationship between oil production and fracking inputs is not necessarily monotonic and that it varies over drilling locations.²⁰ Research documenting these results was not publicly available to firms during the time period I study, which means that firms faced a complicated

¹⁶See Hicks (2012)

¹⁷See Hicks (2012)

¹⁸Fracking was first successfully used in shale formations in the 1990s. Under the hunch that permeability issues could eventually be resolved through the use of fracking, Mitchell Energy worked for years on its own and with the help of the US Department of Energy to learn how to apply fracking technology to the Barnett shale in Texas. They succeeded in 1997. See Shellenberger *et al.* (2012). Two firms active in North Dakota, EOG and XTO, were active in the Barnett as well. However, the Barnett Shale is different from the Bakken. Barnett wells are drilled directly into the shale layer, and produce natural gas instead of oil. It is unlikely that any knowledge that these firms may have had about fracking in the Barnett was useful in the Bakken.

¹⁹See Shelley *et al.* (2012)

²⁰See Baihly *et al.* (2012)


learning problem.

3.2.3 The Information Environment in North Dakota

Firms in North Dakota can learn about the relationship between oil production, location, and fracking inputs from the past experiences of other firms. After a firm fracks a well, the oil and gas regulator in North Dakota requires the firm to submit a well completion report, detailing the well's horizontal length, location and fracking inputs. Additionally, the regulator and tax authorities require the firm to submit audited production records on a monthly basis. The regulator publishes this information on the internet, making it easy for firms to learn information about every previously fracked well in the state, including information about wells that they took no part in developing.

North Dakota's well confidentiality laws generate a 6 month delay between when firms submit well completion reports and when the regulator makes them public. This delay creates differences across firms in what wells they can learn from at each point in time, as the operating firm of a well has a temporary knowledge advantage over other firms. However, the ownership structure of mineral rights in a well mitigates some of these differences. Mineral rights for a well are often owned by many separate firms. Every firm that owns mineral rights in the area spanned by a well is entitled to pay a share of the capital expenditures needed to develop the well in exchange for a share of the revenue generated by the well. The firm with the largest mineral rights claim in a well is called the "operator", and it retains all control rights, including the choice of the well's fracking inputs. The remaining owners of mineral rights are called "non-operating participants". Figure 3.2 depicts a hypothetical ownership situation for a well in the Bakken. The land spanned by the well is a 2 mile by 1 mile rectangle, called a "spacing unit". Within this spacing unit, Firm A has the largest mineral rights claim, followed by firms B and C. The wellhead enters the ground in A's claim and the horizontal segment passes through B's claim. Though the well does not directly pass through C's claim, it is close enough to C's claim that it may be drawing oil from the claim. While A retains control rights,

Figure 3.2: Diagram of a hypothetical spacing unit



B and C must pay their respective share of capital expenditures.²¹

Non-operating participants have immediate access to a well's completion report.²² This means that non-operating participants in a well are not subject to well confidentiality rules and thus observe information regarding a well before the public does.

3.2.4 Data

Well Characteristics and Production History

I have collected operating and production data for every well targeting the Bakken shale formation in North Dakota that was fracked between January 1, 2005 and December 31, 2011. This data is reported by oil companies to the North Dakota Industrial Commission (NDIC), and the NDIC publishes their submissions on the internet. For each well i , I observe the

²¹Firms can choose to opt out of a spacing unit, but that does not allow them to operate another well within the spacing unit, so opt outs are rare.

²²See Larsen (2011)

location of its wellhead in latitude lat_i and longitude lon_i coordinates, its horizontal length H_i , the mass of sand S_i and volume of water W_i per foot of horizontal length used in fracking and the identity of the operating firm f_i . Additionally, I observe oil production Y_{it} for well i in its t -th month of existence and the number of days D_{it} during that month that the well was actually producing. Let X_{it} denote the set (H_i, f_i, D_{it}) and let Z_i denote the set (S_i, W_i, lat_i, lon_i) . Then the dataset (Y_{it}, X_{it}, Z_i) has a panel structure, where i indexes wells and t indexes well-specific timing. Though I only study wells fracked during 2005-2011, I have production data through February 2013, making it possible to study the performance of all wells for at least a year. While the production history is reported electronically on the NDIC website, the static well characteristics are stored in PDF format, so much of this dataset was entered into the computer manually. I also observe the “township” τ_i that the wellhead lies in. Townships are 6 mile by 6 mile squares, defined by the US Geological Survey and are a standard measure of location in the oil & gas business. There are 272 townships in North Dakota with Bakken wells during 2005-2011. I have also collected the geographic boundaries of the spacing units for every well. This data comes from various portions of the NDIC website.

Though most of the data I collect from the NDIC is self reported by firms, there are two reasons why it is likely to be truthfully reported. First, oil and gas regulations in North Dakota specify explicit penalties for failure to report required information and false reporting, including fines of up to \$12,500 per day per offense and felony prosecution.²³ Second, because operators wish to collect payment for capital expenditures from their non-operating partners, they must share the documentation and billing they receive from their service contractors. If operators were to report data to the NDIC that was at odds with what they had shared with their non-operating partners, they might jeopardize their ability to collect payment.

Table 3.1 reports the cross-sectional distribution of well characteristics and oil production in the first year. There is substantial variation across wells in both fracking input use and oil production. The 75th percentiles of sand, water and oil production are more than double their respective 25th percentiles. This variation will be important later on in estimating the

²³See Section 38-08-16 in the NDIC Rulebook.

Table 3.1: *Summary Statistics*

Variable	Mean	Std. Dev	P25	P50	P75	N
lbs sand per foot	265.02	138.68	158.27	264.53	378.66	2,699
gals water per foot	188.87	110.73	100.31	181.70	249.52	2,699
horizontal feet in length	8,040	2,138	5,600	9,135	9,518	2,699
avg producing days per month	26.80	2.99	25.90	27.56	28.67	2,699
oil production per foot in first year	10.86	8.95	5.38	8.39	12.99	2,699
# non-operating participants	3.00	2.50	1.00	3.00	4.00	2,699
# past wells fracked by operator	80	82	16	49	125	2,699
# past wells fracked by others	1,089	658	511	1,062	1,698	2,699

relationship between oil production and fracking inputs. Most wells have horizontal segments that are 9,000 feet or longer. The length of a well's horizontal segment is determined by the size of its spacing unit. Though not shown in the table, approximately 75% of wells have rectangular spacing units that are two miles wide and one mile tall. The remaining 25% have 1 mile square spacing units. The average well produces almost 11 bbl per foot of horizontal length in its first year. Since the price of oil averaged \$76 per bbl during 2005-2011, the value of production in the first year for the average well is worth \$6.6 million. Most wells tend to produce on the majority of days during a month, and though not shown in the table, only 93 wells have fewer than 20 average producing days. The bottom rows of Table 3.1 show the distribution of non-operating participants and past experience across wells. In the average well, 3 other firms obtain knowledge about a well at the same time as the well's operator. The average well is fracked by a firm that has previously fracked 80 of its own wells, and can observe the data on 1,089 wells fracked by others.

Table 3.2 shows the distribution of well characteristics and oil production. The number of wells fracked and the number of active townships and firms all increase over time. More than 65% of all wells are fracked during the last two years, and in 2011, wells are fracked in 85% of townships by 70% of all firms. Over time, firms frack longer wells, using more sand and more water. Firms operating in 2011 use more than three times as much sand and four times as much water per foot of horizontal length, on average, as firms in 2005. However, average oil production does not rise monotonically, reaching its peak in 2008 and then falling thereafter.

Table 3.2: *Summary Statistics by Year*

		2005	2006	2007	2008	2009	2010	2011
# wells fracked		10	20	94	352	463	691	1,069
# active townships		9	17	37	102	132	179	231
# active firms		5	11	17	28	34	47	49
Sand	Average	94.50	136.88	134.64	180.00	212.75	308.82	302.79
	Std. Dev	22.01	152.43	143.15	146.79	145.32	121.68	110.85
Water	Average	49.53	64.29	95.67	108.28	137.08	215.14	232.68
	Std. Dev	25.03	61.87	83.72	59.90	88.36	99.88	111.28
Length	Average	6,883	6,062	7,017	7,283	7,238	8,006	8,795
	Std. Dev	1,679	2,001	2,048	2,233	2,316	2,144	1,715
Oil	Average	3.08	4.85	10.76	13.41	11.55	11.15	9.73
	Std. Dev	1.94	7.59	13.72	15.16	9.83	6.72	5.78

Table 3.3 reports average oil production per foot by quintiles of sand and water use per foot.²⁴ Across both sand and water use, the highest input levels are associated with higher oil production. For every quintile of water use (columns), the top quintile of sand use has higher production than the bottom quintile. For all but the second quintile of sand use (rows), the top quintile of water use has higher production than the bottom quintile. Thus the data shows that sand and water use affect oil production, though not strictly monotonically.

To verify the importance of spatial heterogeneity in the relationship between fracking inputs and oil production, I estimate a simple Cobb-Douglas production function for fracking, with and without township fixed effects. I regress the log of first years oil production per foot of horizontal length on the well's log sand use and log water use:

$$\log \text{ oil per foot}_i = \alpha_0 + \alpha_S \log S_i + \alpha_W \log W_i + \tau_i + \epsilon_i$$

Table 3.4 reports coefficient estimates for this regression. The first column shows estimates without fixed effects, and the second column shows estimates with fixed effects. Consistent with the results in Table 3.3, higher sand and water use are associated with higher production. This is true with and without fixed effects. However, the inclusion of township fixed effects

²⁴To control for the effects of location, I first subtract the average levels of oil production and input use per township from actual production and input use. Then, I add back the overall average levels, creating township fixed effects.

Table 3.3: *Average First Year's Oil Production per Foot of Horizontal Length by Quintiles of Sand and Water Use*

		Quintiles of Water Use				
		First	Second	Third	Fourth	Fifth
Quintiles of Sand Use	First	8.09 (0.33)	8.27 (0.44)	6.77 (0.73)	8.82 (2.20)	10.16 (0.81)
	Second	9.53 (0.37)	10.50 (0.35)	9.27 (0.33)	10.50 (0.56)	9.37 (1.37)
	Third	10.25 (0.52)	11.51 (0.36)	10.91 (0.29)	10.56 (0.35)	10.81 (0.76)
	Fourth	10.71 (0.52)	10.48 (0.58)	13.24 (0.55)	11.46 (0.40)	11.87 (0.48)
	Fifth	10.80 (1.06)	12.24 (0.83)	13.37 (0.98)	13.19 (0.52)	13.85 (0.37)

Net of township fixed effects. Standard errors in parentheses.

decreases the coefficient on sand use and increases the coefficient on water use, suggesting the existence of spatial heterogeneity in oil production and the possibility that firms make different input choices in different locations.

Oil Prices

I collect the daily spot prices for West Texas Intermediate crude oil at the Cushing, Oklahoma oil trading hub from the US Energy Information Administration. The Cushing price is the reference price for oil futures traded on the NYMEX commodity exchange, and the Cushing hub is connected to North Dakota through the Keystone and Enbridge pipeline systems. Figure 3.3 plots quarterly average oil prices at the Cushing hub. Between 2005-2011, there was a boom and bust in oil prices, with prices climbing from approximately \$60 per bbl in early 2007, reaching more than \$120 per bbl in mid 2008 and falling to \$45 per bbl in early 2009. In 2010-2011, when more than 65% of the wells are fracked, oil prices average \$87 per bbl.

3.2.5 Drilling and Fracking Costs

Though the NDIC does not require firms to report their costs, the legal process in North Dakota occasionally makes this information public. In particular, when a non-operating


Table 3.4: *Spatial Heterogeneity in the Relationship Between Sand, Water and Oil Production*

	(1) Log Oil per foot	(2) Log Oil per foot
α_0	-0.0280 (0.104)	0.319 (0.0948)
α_S	0.352 (0.0211)	0.208 (0.0183)
α_W	0.0512 (0.0228)	0.137 (0.0185)
Township FE		X
N	2,698	2,698
R^2	0.159	0.618

Standard errors in parentheses. OLS estimates of

$$\log \text{ oil per foot}_i = \alpha_0 + \alpha_S \log S_i + \alpha_W \log W_i + \tau_i + \epsilon_i$$

Figure 3.3: *Quarterly Average Cushing Oil Prices*



mineral rights owner decides not to participate in a well, the operator can ask the NDIC to impose a “risk penalty”, which temporarily prevents the non-participant from earning revenue from its mineral rights.²⁵ In order to make this request, the operator must legally submit its estimate of the cost of drilling and fracking the well, and this information is publicly recorded by the NDIC. Of the 2,699 wells in this dataset, the cost records for 90 are in the public domain for this reason.

These wells span several years, so to make their costs comparable, I normalize them using a cost index. There is no single publicly available cost index that is both specific to the Bakken and available for all of 2005-2011, so I construct one by combining several other indices. Between the first quarter of 2005 and the fourth quarter of 2007, the index grows at the rate of the BLS Producer Price Index for oil & gas extraction. Between the the first quarter of 2008 and the fourth quarter of 2009, the index grows at the rate of a cost index for vertical wells drilled in North Dakota, published by Spears & Associates, a private consulting firm.²⁶ Finally, starting in the first quarter of 2010, the index grows at the rate of the Spears & Associates cost index for horizontal wells drilled in North Dakota. I fix the cost index to 1 in the first quarter of 2005 and define “normalized costs” as reported costs divided by the cost index. Figure 3.4 plots the cost index over time.

To estimate the individual components of costs, I regress normalized costs for these 90 wells onto a constant, lateral length, total sand use, total water use and year-quarter fixed effects. The adjusted R-squared of this regression is 0.54, and the coefficients on lateral length, sand and water are all significantly different from zero at the 5% level. I define the fixed drilling and fracking cost as the sum of the constant and the year-quarter fixed effects, the variable

²⁵A non-participating mineral rights owner faced with a risk penalty forfeits a significant portion of its share of the well’s revenue. In North Dakota, risk penalties are set to 200% of a non-participant’s share of capital expenditures. This means that non-participants do not earn any revenue from a well in which they own mineral rights until the well has generated 200% of its capital expenditures in oil production.

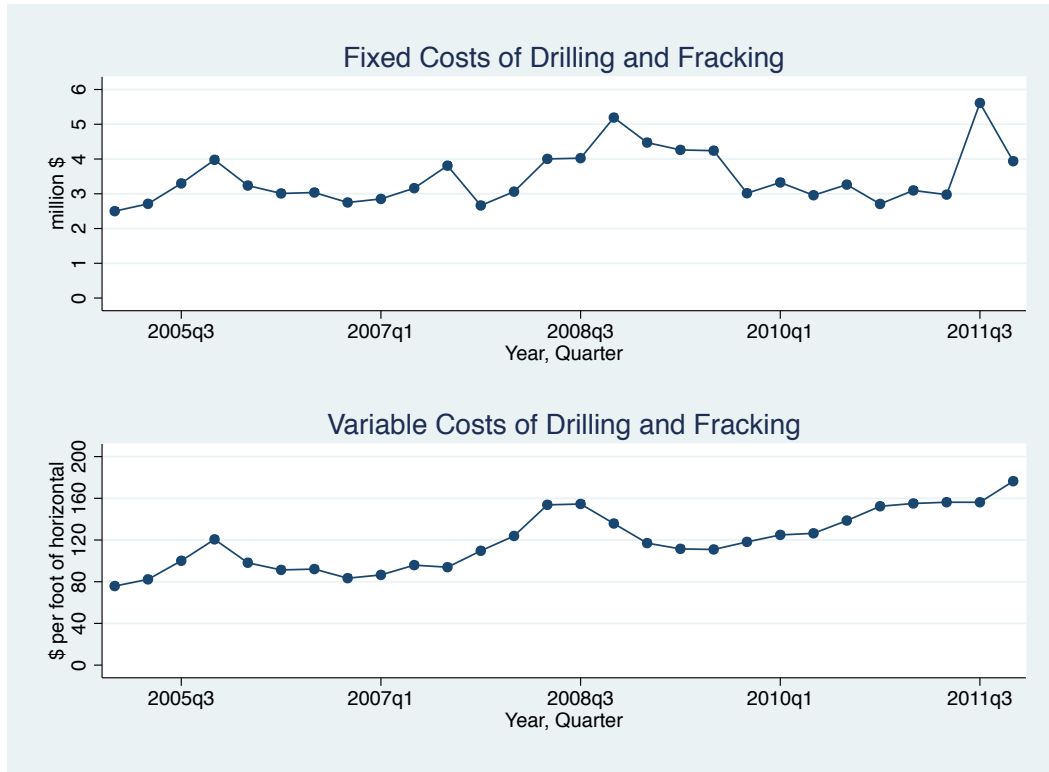
²⁶Spears & Associates surveys independent engineers in North Dakota quarterly, asking them to estimate the cost of a reference well. The cost estimates are divided into 14 categories, of which 4 are fracking related and 10 are drilling related. The data is separately available for a vertical reference well design, which begins in the first quarter of 2008 and a horizontal reference well design, which begins in the first quarter of 2010. The vertical reference design does not include a fracking treatment. The characteristics of the reference wells stay constant over time, so the changes in estimated costs are due to changes in prices, not quantities.

Figure 3.4: *Fracking Cost Index*



The cost index is computed from the BLS Producer Purchasing Index (PPI) for the Oil & Gas Extraction industry from the first quarter of 2005 to the fourth quarter of 2007. Then, from the first quarter of 2008 to the fourth quarter of 2009, it is calculated from the Spears & Associates data for vertical wells in North Dakota. Finally, from the first quarter of 2010 to the fourth quarter of 2011 it is calculated from the Spears & Associates data for horizontal wells in North Dakota.

Figure 3.5: *Fixed and Variable Costs of Drilling and Fracking*




The variable costs of using sand and water in fracking are estimated from a regression of the normalized total drilling and fracking costs for 90 wells with cost data in the public domain onto a constant, lateral length, total sand use, total water use and year-quarter fixed effects. The estimated fixed cost of drilling and fracking is equal to the constant plus the year-quarter fixed effect, divided by the cost index. The estimated variable cost of drilling and fracking is equal to the coefficient on lateral length, divided by the cost index.

drilling and fracking cost as the coefficient on lateral length, and the sand and water costs as the coefficients on sand and water use. Finally, I generate time-specific costs by multiplying these estimates by the cost index. Figures 3.5 and 3.6 plot these costs over time.

Information Sets

At time t , firm f can learn about fracking from three sets of wells. First, f can observe all wells that the regulator has made public by time t . This public knowledge includes wells that f operated and wells that other firms operated. Second, f can observe its own wells which are not yet public knowledge, due to well confidentiality. Third, f can observe other firms' wells in which it is a non-operating participant. I can compute the first two sets of information from

Figure 3.6: *Variable Costs of Using Sand and Water in Fracking*



The variable costs of using sand and water in fracking are estimated from a regression of the normalized total drilling and fracking costs for 90 wells with cost data in the public domain on a constant, lateral length, total sand use, total water use and year-quarter fixed effects. The estimated cost of pumping 1 pound of sand is equal to the coefficient on sand use, divided by the cost index, while the estimated cost of pumping 1 gallon of water is equal to the coefficient on water use, divided by the cost index.

well completion reports alone. To compute the third set, I must identify the mineral rights owners in each well's spacing unit.

I collect mineral rights lease data from DrillingInfo.com, which digitally records the universe of mineral rights transactions filed in county registries of deeds. These leases are often between a surface owner and an intermediary lease broker operating on behalf of an oil company. Once the broker acquires a lease, it assigns this lease back to its client, a transaction which is not recorded by DrillingInfo.com. To capture the information in the lease assignment process, I also scrape the website of the North Dakota Registry Information Network (www.ndrin.com), which electronically records lease assignments. I combine this lease and lease assignment data into a single dataset identifying the names of any firm that has mineral rights in a spacing unit. I assume that all firms with mineral rights in a well's spacing unit that are not the well's operator are non-operating participants.²⁷

Outside Experience

Throughout the paper, I assume that the only knowledge firms have about fracking comes from the wells fracked in North Dakota during 2005-2011. To assess the validity of this assumption, I collect firm-specific drilling history from IHS International for the 8 most active firms in my data, which I report in Table 3.5. In the first column, I list the number of wells each firm completed in the Bakken during 2005-2011. These 8 firms frack 60% of the wells in the dataset. During the time period I study, these firms are all publicly held, either as independent firms (Brigham, Continental Resources, EOG, Hess, Marathon and Whiting) or as subsidiaries of larger oil companies (Burlington is owned by Conoco Phillips, XTO is owned by Exxon Mobil).

On the right hand side of Table 3.5, I list the US operating history of these firms outside of North Dakota. In the 10 years prior to the period I study, these firms collectively completed tens of thousands of vertical wells, which are typically drilled into conventional formations,

²⁷That is, I assume that no mineral rights owners are non-participants. Since only 90 out of 2,699 wells in this time period had risk penalty challenges, this is a reasonable assumption.

Table 3.5: *Wells Completed by the 8 Most Active Firms, by Location, Time and Well Characteristics*

Firm	North Dakota		Outside North Dakota		
	2005-2011	1995-2004		2005-2011	
	Bakken Shale	Vertical	Horizontal	Vertical	Horizontal
Brigham	113	161	0	93	0
Burlington	105	3,826	26	2,792	532
Continental Resources	313	597	3	657	167
EOG	354	4,659	91	6,566	2,914
Hess	165	639	2	219	15
Marathon	223	2,221	4	813	87
Whiting	247	131	0	1,150	11
XTO	101	2,349	53	7,749	2,801
Rest of industry	1,078				

without frack jobs. However, they only completed 179 horizontal wells, suggesting that they had very little experience with the technology necessary to develop wells in the Bakken Shale. Only three firms had previously completed more than ten horizontal wells, and two had done none. During 2005-2011, all eight firms are active outside North Dakota, with four firms completing more than a thousand wells each. Except for EOG and XTO, the vast majority of contemporaneous operational experience outside North Dakota is in vertical wells, though seven of the eight firms do complete horizontal wells. Thus, there is limited scope for these firms to learn about fracking from experience outside of the Bakken.

3.3 The Fracking Production Function

To quantify what knowledge firms learn about fracking, it is necessary to measure the empirical relationship between oil production, location and fracking input choices. I do this by estimating a production function for fracking. This production function accounts for variation in oil production across a well's life and variation between wells in average production levels.

A well's production changes over time due to age and maintenance-driven downtime. I measure the impact of these factors on oil production using a simple model common in the petroleum engineering literature. Because a well's age is outside the firm's control and because maintenance needs are both similar across wells and scheduled in advance, I argue that the

time-varying error in production is plausibly exogenous.

Wells have different average production levels due to differences in their horizontal lengths, locations and fracking inputs. Location and fracking inputs may nonlinearly affect production, so I measure their impact non-parametrically, using Gaussian process regression (GPR), which I describe in detail below. The well-specific error in average production includes the effects of unobserved inputs, such as chemicals, the unobserved amount of oil that can be recovered and its sensitivity to fracking. I argue that chemical choices are independent of sand and water choices for engineering reasons, and that the information which only firms observe about the well's specific geological properties while drilling is unlikely to be correlated with production outcomes.

In the next two sections, I explain this production function model in further detail.

3.3.1 The Time Series of Oil Production

Per unit of time, wells of all kinds (including non-fracked wells in conventional formations) tend to produce more oil when they are younger and less oil when they are older. This decline in performance over time is not surprising, because the amount of oil that can be recovered is finite and as more of it is pumped out of the ground, the rest becomes more difficult to recover. For nearly 70 years, petroleum engineers have used the simple "Arps" model to illustrate this basic phenomenon (see Fetkovich 1980). The Arps model states that oil production in the t -th month of well i 's life is:

$$Y_{it} = Q_i t^\beta \exp(\nu_{it})$$

where Q_i is the *baseline* level of production, $\beta < 0$ is a constant governing the production decline of the well and ν_{it} is a mean-zero production shock. In log terms, this is

$$\log Y_{it} = \log Q_i + \beta \log t + \nu_{it}$$

meaning that a 1% increase in a well's age should decrease per period production by $-\beta\%$, on average.

The operator of a well chooses D_{it} , the number of days during month t that well i is

producing. Unless the well needs maintenance, there is no reason the operator would choose to produce for fewer than the full number of days during a month. All wells experience two routine maintenance events: the installation of external pumping hardware, and the connection of the well to a gas pipeline network. During maintenance, the operator must shut the well down, reducing D_{it} . My data does not indicate whether maintenance occurs in a month, but it does report the number of producing days D_{it} , which I incorporate in the model:

$$\log Y_{it} = \log Q_i + \beta \log t + \delta \log D_{it} + \nu_{it}$$

The time-varying shock to log production, ν_{it} , is the result of unobserved geological variation and deviations from the Arps model. Firms cannot control t , the age of a well, and it is unlikely that firms observe anything correlated with ν before choosing to do maintenance. Even if they did, firms would rather have the well producing on more days than fewer days, independent of ν . Moreover, firms cannot predict ν when fracking the well, which happens before production starts. For these reasons, I assume that ν is exogenous:

$$\mathbb{E}[\nu_{it} \mid t, H_i, D_{it}, S_i, W_i, lat_i, lon_i] = 0$$

3.3.2 The Cross Section of Oil Production

I specify a semi-parametric model for $\log Q$, the log of baseline production:

$$\log Q_i = \alpha + \eta \log H_i + f(S_i, W_i, lat_i, lon_i) + \epsilon_i$$

The parametric part of this model, $\alpha + \eta \log H_i$, is a Cobb-Douglas production function relating the horizontal length of a well to its baseline production. Though it may seem natural that η should equal one, there are practical reasons why this may not be true. Fracking applied to the furthest away points of the horizontal segment of a well may not always perform as well as fracking applied to the closest points. If this decline in effectiveness is nonlinear, wells with longer horizontal segments may not proportionally outperform wells with shorter horizontal segments. The Hicks-neutral productivity α measures the average log baseline production across wells. I discuss the well-specific productivity shock ϵ_i in more detail below.

The function $f(S_i, W_i, lat_i, lon_i) = f(Z_i)$ captures the relationship between baseline production, location and fracking choices. Table 3.4 in the data section suggests that this relationship differs across locations, and current petroleum engineering suggests that it may be nonlinear. For this reason, I estimate $f(Z_i)$ non-parametrically, using Gaussian process regression, or GPR. GPR makes kernel regression techniques available within a panel data framework. Because there are few examples of GPR in applied economic settings, I provide a basic overview of its application here.

Gaussian process regression

A *Gaussian process* G is a probability distribution over continuous real functions. Gaussian processes are defined by two functions: a mean function $m(Z)$ and a positive definite covariance function $k(Z, Z')$. The mean function is the expectation of the value of a function f drawn at random from G at the point Z . The covariance function is the covariance between $f(Z)$ and $f(Z')$. In mathematical terms, the mean and covariance functions satisfy:

$$m(Z) = \int f(Z)dG(f)$$

$$k(Z, Z') = \int (f(Z) - m(Z))(f(Z') - m(Z'))dG(f)$$

A Gaussian process is “Gaussian” because the joint distribution of the values $f(Z_1)...f(Z_N)$ is multivariate normal, with a mean vector μ and covariance matrix Σ given by:

$$\mu = (m(Z_1)...m(Z_N))^T$$

$$\Sigma_{i,j} = k(Z_i, Z_j)$$

This implies that the distribution of $f(Z)$ is also normal with mean $m(Z)$ and variance $k(Z, Z)$. The normality property makes it easy to compute the likelihood that a dataset $(g_i, Z_i)_{i=1}^N$ is generated by the relationship $g = f(Z)$ for a function f drawn from a Gaussian process with mean $m(Z)$ and covariance $k(Z, Z')$. By selecting mean and covariance functions from parametric families, the parameters that best fit the dataset can be estimated using maximum likelihood.

To estimate the function $f(Z_i)$ above, I assume $m(Z) = 0$ due to the presence of the constant term, α , in the parametric portion of the production function. I assume that $k(Z, Z')$ takes the form of a multivariate normal kernel:

$$k(Z_i, Z_j | \gamma) = \exp(2\gamma_0) \exp\left(-\frac{1}{2} \sum_{d \in S, W, lat, lon} \frac{(Z_{i,d} - Z_{j,d})^2}{\exp(2\gamma_d)}\right)$$

The first parameter, γ_0 , measures the variance of the unknown function $f(Z)$. As points (Z_i, Z_j) become arbitrarily close to each other, the covariance function approaches the variance of f , and its formula collapses to $\exp(2\gamma_0)$. The remaining parameters $\gamma = (\gamma_S, \gamma_W, \gamma_{lat}, \gamma_{lon})$ measure how smooth f is in each dimension.

If the mean function is 0 and the covariance function parameters are γ , then the log likelihood of the data $(g_i, Z_i)_{i=1}^N$ is:

$$\log \mathcal{L}(\gamma) = -\frac{1}{2} g^\top K(\gamma)^{-1} g - \log |K(\gamma)| - \frac{N}{2} \log(2\pi)$$

where $g = (g_1 \dots g_N)^\top$ and $K(\gamma)_{i,j} = k(Z_i, Z_j | \gamma)$. The process of maximizing this likelihood over γ is called *Gaussian process regression*, or GPR. Conditional on γ and the data (g, \mathbf{Z}) , the distribution of f evaluated at an out-of-sample point \tilde{Z} is normal, with mean and variance given by:

$$\begin{aligned} \mathbb{E} \left[f(\tilde{Z}) | g, \mathbf{Z}, \gamma \right] &= k(\tilde{Z} | \gamma)^\top K(\gamma)^{-1} g \\ \mathbb{V} \left[f(\tilde{Z}) | g, \mathbf{Z}, \gamma \right] &= k(\tilde{Z} | \gamma)^\top K(\gamma)^{-1} k(\tilde{Z} | \gamma) \end{aligned}$$

where $k(\tilde{Z} | \gamma) = (k(Z_1, \tilde{Z} | \gamma) \dots k(Z_N, \tilde{Z} | \gamma))^\top$. Note that the formula for the mean of $f(\tilde{Z})$ is similar to the formula for the estimated regression function in kernel regression.²⁸ However, the additional assumptions about the distribution of possible regression functions in GPR make it possible to select smoothing parameters γ using likelihood techniques, which is not possible in kernel regression. Moreover, since GPR can be defined in terms of a likelihood

²⁸In kernel regression, the term $k(\tilde{Z} | \gamma)^\top K(\gamma)^{-1}$ in the estimated regression function is replaced with $\frac{k(\tilde{Z} | \gamma)^\top}{\sum_i k(Z_i, \tilde{Z} | \gamma)}$. However, the estimates of variance in kernel regression are not directly comparable to the variance formulas in GPR.

function, it can easily be incorporated into panel data methods, something which is challenging in standard kernel regression.

Gaussian processes are commonly used in the artificial intelligence and operations research literatures, though their application in economics is so far limited to econometric theory.²⁹ For a detailed treatment of Gaussian processes, see Rasmussen and Williams (2005).

The Well-Specific Shock ϵ_i

The well-specific shock to log baseline production, ϵ_i , contains unobserved inputs to the fracking process and unobservable variation in geology. Fracking chemicals are the main unobserved input.³⁰ Firms primarily use chemicals to inhibit bacterial growth in the fracking mixture, to provide lubrication for the pumping units used in fracking and to prevent corrosion and mineral scaling in the well pipe.³¹ There is evidence in the petroleum engineering literature that an operator's choice of chemicals does not directly affect the efficiency of its sand and water choices, so I assume that sand and water choices are independent of chemical choices.³²

The petroleum engineering literature predicts that different parts of the Bakken contain different amounts of oil and respond to fracking inputs differently.³³ In particular, wells that are drilled into parts of the Bakken which are thicker, contain more organic material or are more thermally mature have more oil to draw from, and as a result, fracking inputs may be more productive. Similarly, fracking inputs may generate more extensive fracture networks in wells drilled into more permeable parts of the Bakken than wells in less permeable parts. However, aside from the location-specific nature of the production function, I do not have

²⁹See Kasy (2013) for a recent example.

³⁰Another unobserved input is the characteristics of the piping and fracking hardware that firms use to implement frack jobs. This hardware determines the number of fracture initiation points, their distribution across the lateral segment and the level of pressure inside the wellbore.

³¹See <http://www.fracfocus.org> for further details on the chemicals used in fracking.

³²See, for example, Jabbari *et al.* (2012)

³³See Baihly *et al.* (2012), Jabbari *et al.* (2012) and Saputelli *et al.* (2014)

data to control for geological variation in the Bakken.³⁴ If firms have geological data that may be indicative of how much oil a well contains or how amenable it is to fracking, they may adjust their fracking inputs in response and ϵ_i will not be independent of these choices. Unfortunately, I do not have instruments for fracking input choices, so it is important to consider what additional information firms could have about the wells they are fracking and whether they use it to make fracking decisions.

For the vast majority of wells, firms do not have well-specific information about the thickness, organic content, thermal maturity or permeability of the rock they drill into. To get this information, firms must perform expensive and time-consuming geological tests, the results of which are publicly documented by the NDIC.³⁵ These tests are only possible if firms elect to drill the vertical portion of the wellbore all the way through the entire Bakken formation, which they rarely do.³⁶

Firms do have a potentially useful source of information about well quality in the samples of rock that they collect during drilling, called “cuttings”. As the drill bit passes through the upper Bakken shale on its way into the middle Bakken, firms can analyze the returned rock, which may be indicative of the amount of the oil and the level of permeability in the upper Bakken shale at the location where the horizontal segment starts. However, since the goal in horizontal drilling is to stay inside the middle Bakken, firms receive no additional information about the upper Bakken shale and receive no information at all about the lower Bakken shale during the course of drilling. Moreover, the characteristics of the upper Bakken shale can change over the length of the horizontal segment, and there is no guarantee that the lower Bakken shale has the same characteristics at a point as the upper Bakken shale. During the time period I study, laboratory tools to infer rock properties like permeability from cuttings

³⁴In the appendix, I analyze the (limited) publicly available data on thickness, organic content and thermal maturity. Broadly speaking, this data is not well-specific (it is spatially interpolated from a small number of wells) and does not explain much variation in production after conditioning on location.

³⁵Specifically, firms use gamma ray well logs to determine thickness, rock evaluation pyrolysis of cuttings or well cores to measure organic content and thermal maturity and drill stem tests or MRI/NMR tests to measure permeability.

³⁶For example, Sitchler *et al.* (2013), a recent petroleum engineering study of well performance, fracking inputs, and geology characteristics, has the necessary data for just seven wells.

data had not yet been developed.³⁷ Thus, the information firms can acquire during drilling is unlikely to be helpful in choosing fracking inputs, and in practice may not be used at all.

For these reasons, I argue that ϵ_i is exogenous to firm choices and other well characteristics:

$$\mathbb{E}[\epsilon_i \mid t, H_i, D_{it}, S_i, W_i, lat_i, lon_i] = 0$$

Combining everything together, the whole production function model is:

$$\log Y_{it} = \alpha + \beta \log t + \delta \log D_{it} + \eta \log H_i + f(Z_i) + \epsilon_i + \nu_{it}$$

Since Gaussian process regression generates a normal likelihood for $f(Z_i)$, I assume that ν_{it} and ϵ_i are both normal, with zero mean and variances σ_ν^2 and σ_ϵ^2 , respectively.

3.3.3 Likelihood

I compute the likelihood function in two steps. In the first step, I treat the unobserved effect of fracking and location $f(Z_i)$ as observed and compute the likelihood of (Y_{it}, X_{it}) conditional on $f(Z_i)$ and the parameters. In the second step, I integrate out the unobserved values of $f(Z_i)$ using the likelihood function for $f(Z_i)$ generated by GPR. I describe the likelihood calculation in detail in the appendix.

3.3.4 Production Function Estimates

Table 3.6 shows maximum likelihood estimates of the semi-parametric production function described above in addition to a simpler parametric specification. The parametric specification replaces $f(S_i, W_i, lat_i, lon_i)$ with township fixed effects, τ_i , and a Cobb-Douglas production technology in sand and water, $\kappa_S \log S_i + \kappa_W \log W_i$.

All of the parametric model coefficients are statistically significantly different from zero in both specifications and the coefficients common to both have similar estimates. As expected, wells produce less oil per month as they age, with an estimated log decline rate of -0.56 .³⁸

³⁷See, for example, Ortega *et al.* (2012), who note that “Cuttings have not been used in the past quantitatively for optimization of hydraulic fracturing jobs.”

³⁸Current geophysics research on the Bakken has found similar decline rates. Hough and McClurg (2011),

The coefficient on days producing is 1.75, suggesting that when wells undergo maintenance, production per day is lower than when wells do not have maintenance issues. Wells with longer horizontal segments produce more oil than wells with shorter segments, but the effect is not linear. Doubling the horizontal length of a well increases production by 80% in the Cobb-Douglas specification and 85% in the Gaussian process. The variance of ϵ is larger in the Cobb-Douglas specification than in the Gaussian process, suggesting that the flexibility of the Gaussian process explains more of the variation in baseline oil production than Cobb-Douglas and location fixed effects do. The estimated Cobb-Douglas marginal productivities of sand and water are precisely estimated and are smaller than the preliminary estimates in Table 3.4. Sand and water both increase oil production, with decreasing returns to scale.

The estimated GPR smoothing parameters do not have an intuitive interpretation, so I illustrate the estimated production relationships graphically in Figure 3.7. The top panel is a contour plot of the non-parametrically estimated function $f(S_i, W_i, lat_i, lon_i)$, evaluated at the geographic centroid of the most active township during this time period. The lines are iso-production curves, which are combinations of sand and water choices with the same estimated value of f . Across all levels of water use, greater sand use is associated with higher oil production, while greater water use is only associated with higher production at the highest level of sand use, and only in a limited range. The middle panel shows contour lines for the Cobb-Douglas specification. The Gaussian process and Cobb-Douglas specifications make starkly different predictions about the impact of fracking inputs and location on oil production. At the average sand and water choices for this township, 266 lbs and 131 gals per foot, respectively, the Gaussian process predicts -3.5 log points of baseline production, while Cobb-Douglas predicts -3.1, meaning that the predictions of the two models differ by 40%. Additionally, the non-parametric specification makes different predictions in different locations. The bottom panel shows contour lines for the production function evaluated at the centroid of a nearby township. The location of the most productive sand and water choices differ across the two townships. In the top panel, the maximal choice is approximately 600 lbs sand and

for example, estimates the decline rate to be -0.5 .

Table 3.6: *Production Function Model Estimates*

Coefficient	Cobb-Douglas		Gaussian Process	
	Estimate	Std. Error	Estimate	Std. Error
α			-4.4152	(0.3278)
β	-0.5576	(0.0024)	-0.5570	(0.0024)
δ	1.7543	(0.0035)	1.7549	(0.0035)
η	0.7977	(0.0363)	0.8479	(0.0357)
γ_0			-0.3945	(0.0572)
γ_S			6.1757	(0.1343)
γ_W			5.9467	(0.1232)
γ_{lat}			-2.4702	(0.0539)
γ_{lon}			-2.2376	(0.0609)
κ_S	0.1582	(0.0157)		
κ_W	0.1148	(0.0159)		
$\log \sigma_\epsilon$	-0.9086	(0.0147)	-1.0591	(0.0187)
$\log \sigma_\nu$	-0.4898	(0.0024)	-0.4897	(0.0024)
Township Fixed-effects	X			
Overall R^2	0.783		.811	
Between R^2	0.813		.882	
Within R^2	0.764		.764	
# Wells	2,699			
# Well-months	91,783			

Maximum likelihood estimates of the Cobb-Douglas production function model:

$$\log Y_{it} = \beta \log t + \delta \log D_{it} + \eta \log H_i + \kappa_S \log S_i + \kappa_W \log W_i + \tau_i + \epsilon_i + \nu_{it}$$

and the Gaussian process production function model:

$$\log Y_{it} = \alpha + \beta \log t + \delta \log D_{it} + \eta \log H_i + f(Z_i | \gamma) + \epsilon_i + \nu_{it}$$

Y_{it} is oil production for well i when it is t months old, D_{it} is the number of days producing, H_i is the horizontal length, and Z_i is the vector of sand use S_i , water use W_i , latitude lat_i and longitude lon_i . τ_i is a set of township fixed effects. “Between” R^2 is the R^2 for the average predicted log baseline production. “Within” R^2 is the R^2 for the predicted time series of production.

200 gals water, per foot, while in the bottom panel it is 400 lbs sand and 500 gals water, per foot. This variation across townships in the relationship between oil production and inputs is not possible with the Cobb-Douglas specification, so for the rest of the paper, I focus on the Gaussian process specification.

The fit of both models is high, with R^2 's of 78% for the Cobb-Douglas model and 81% for the Gaussian process model. The “between” R^2 's, which measure the correlation of predicted baseline production and actual baseline production, are higher, at 81% and 88%, respectively. The production function models fit the data well for several reasons. Both the inputs to fracking, sand and water, and the single output of fracking, crude oil production, are precisely measured. The main unobserved input, fracking chemicals, does not directly affect production or observed input choices, and Gaussian process regression flexibly controls for spatial heterogeneity. Moreover, the production function for fracking is an approximation to a true physical relationship between sand, water, location and oil production. However, since I estimate this approximation non-parametrically, there is the possibility that the estimated smoothing parameters are too narrow, leading to over-fitting.

To check for this, I perform a cross-validation test of the model estimates. For each of 25 test runs, I randomly split the wells into two separate datasets: a training dataset containing 90% of the wells, and a validation dataset containing the remaining 10%. I re-estimate the production function on the training dataset and use the estimates to predict production in the validation dataset. I save the estimated production function coefficients, the R^2 values generated by the training data and the R^2 values generated by the validation data, and report their distribution across test runs in Table 3.7. The parametric components of the production function model are quite stable across runs, with the average model estimates being similar to the full dataset maximum likelihood estimates. The standard deviations across runs are smaller than the maximum likelihood standard errors for the full dataset. Though the R^2 values for validation samples are lower than for training samples, they are still quite high, with the average overall R^2 for validation samples at approximately 78%, compared to 81% in the training samples. To complement these checks, I provide a series of robustness checks of the

Figure 3.7: Contour Plots of Production Function Estimates

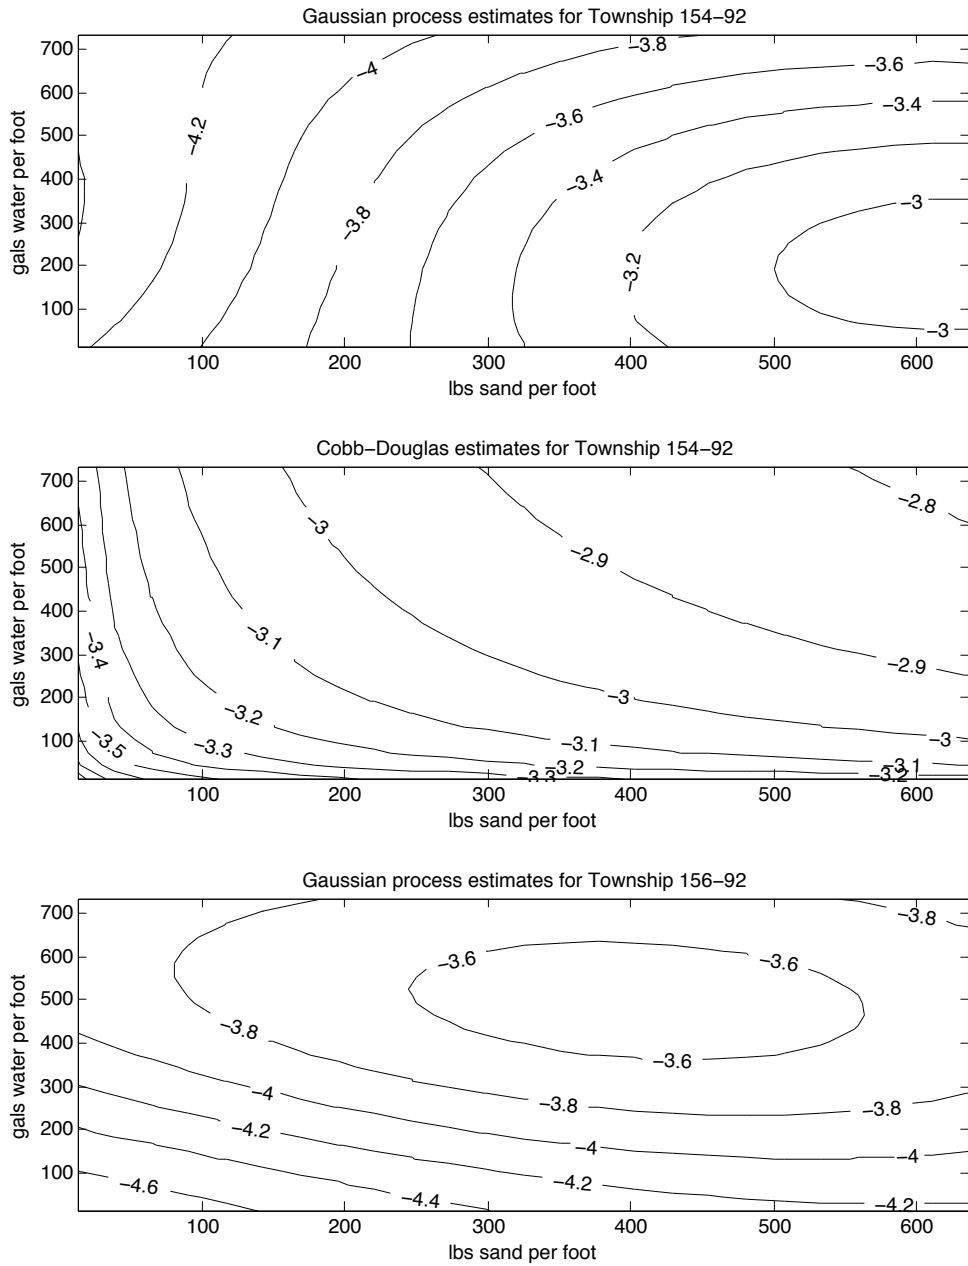


Table 3.7: *Production Function Model Cross-Validation Statistics*

Coefficient	Average Estimate	Std. Dev. of Estimate
α	-4.3388	0.1295
β	-0.5570	0.0016
δ	1.7533	0.0055
η	0.8404	0.0136
γ_0	-0.4046	0.0290
γ_S	6.1659	0.0532
γ_W	5.9278	0.0616
γ_{lat}	-2.4454	0.0236
γ_{lon}	-2.2211	0.0392
$\log \sigma_\epsilon$	-1.0545	0.0109
$\log \sigma_\nu$	-0.4916	0.0063
R^2 comparisons		
R^2 type	Avg. in training	Avg. in validation
Overall R^2	0.8116	0.7835
Between R^2	0.8826	0.8098
Within R^2	0.7652	0.7594
# Wells	2,699	
# Well-months	91,783	
# Cross validation samples	25	

Maximum likelihood estimates of the production function model:

$$\log Y_{it} = \alpha + \beta \log t + \delta \log D_{it} + \eta \log H_i + f(Z_i | \gamma) + \epsilon_i + \nu_{it}$$

Y_{it} is oil production for well i when it is t months old, D_{it} is the number of days producing, H_i is the horizontal length, and Z_i is the vector of sand use S_i , water use W_i , latitude lat_i and longitude lon_i .

stability of the production function across well cohorts in the appendix.

The consistency of the coefficient estimates across cross-validation tests and the high goodness-of-fit measures in validation samples suggest that the maximum likelihood estimates in Table 3.6 do not suffer from over-fitting and represent a stable and causal relationship between inputs and production.

3.4 Evidence for Learning

As firms learn to use fracking technology more efficiently, they should make more profitable fracking design choices. If oil prices, input costs and the quality and size of drilling locations were constant over time, I could test this prediction by extrapolating future production from current production and simply check if average expected discounted profits per well increased over time. However, oil prices, input costs and locations do vary over time, so I control for this variation by examining trends in the ratio of actual profits to counterfactual maximal profits. That is, I compute a profitability measure which compares the profits firms earned with the highest amount of profits they could have earned with the best fracking design for each well.

I use the fracking production function to compute these profits. The profits to well i fracked using design j are

$$\Pi_{ij} = \phi P_i \mathbb{E} \left[\sum_{t=1}^T \rho^t \tilde{Y}_{ijt} \right] - c_i(S_j, W_j)$$

where ϕ is the fraction of oil production the firm keeps for itself, P_i is the price the firm will receive for its oil production, T is the number of periods the well is expected to produce for, ρ is the per-period discount rate, \tilde{Y}_{ijt} is the realization of the level of oil production for well i under fracking design j at age t , and $c_i(S_j, W_j)$ is the total cost of drilling and fracking that design.³⁹ The main empirical object needed in the calculation of Π_{ij} is the expected present value of discounted oil production, $\mathbb{E}[DOP_{ij}]$:

$$\begin{aligned} \mathbb{E}[DOP_{ij}] &= \mathbb{E} \left[\sum_{t=1}^T \rho^t \tilde{Y}_{ijt} \right] \\ &= \sum_{t=1}^T \rho^t \mathbb{E} \left[\tilde{Y}_{ijt} \right] \end{aligned}$$

I compute this expectation conditional on two different information sets: the full data that

³⁹I assume firms believe oil prices follow a martingale process, and thus use a single price, P_i for all future revenues. Additionally, I assume that the fraction of oil revenue that accrues to the firms is 70%, based on typical royalty rates of 16.5%, state taxes of 11.5% and ongoing operating costs of 2%. I set $T = 240$ months, though the NDIC expects Bakken wells to produce for 540 months, making these profit calculations an underestimate. I set $\rho = .9$, which is the standard discount rate use in oil & gas accounting. At this rate, the difference between 540 months and 240 months is only 2.6% in present value terms.

I have, and the data each firm had when it made a fracking design decision. The first case represents an *ex post* expectation, and provides a way of asking whether firms made better fracking design decisions over time, given today’s knowledge. The second case represents an *ex ante* expectation, and provides a way of asking whether firms’ choices were consistent with static profit maximization, given my measures of their information sets.

In both cases, I combine the production function parameter estimates in Table 3.6 with the normality assumptions on the unobserved terms to compute a probability distribution over oil production. Since the production function estimates depend on the full dataset, this means that I am computing *ex ante* expectations under the assumption that firms had the same beliefs about the production function parameters as I do now. This is a strong assumption. The *ex ante* calculation of expected oil production will be biased if firms had different beliefs than I do about the decline rate β , the productivity of producing days δ and horizontal length η , the bandwidth parameters γ and the variances σ of the unobservable production shocks. I assume that these biases are small, as decline rates and productivity parameters can be predicted using geophysical models⁴⁰, and bandwidth and variance parameters do not affect the asymptotic properties the production function estimate.⁴¹ Moreover, the impact of fracking design and location $f(Z)$ is computed nonparametrically from both the bandwidth parameters γ and the information set. Thus firms with different information sets will have different beliefs about $f(Z)$, and these beliefs will differ from the *ex post* beliefs as well.

I present the full calculation of expected discounted oil production in the appendix.

3.4.1 *ex post* Comparisons

Over time, firms choose fracking designs with higher *ex post* expected profits. The top half of Figure 3.8 plots the *ex post* ratio of actual profits to maximal profits per well.⁴² The average fraction of profits captured increases nearly monotonically over time, from 15.7% in 2005 to

⁴⁰See Fetkovich (1980).

⁴¹See section 7.1 in Rasmussen and Williams (2005).

⁴²I only include wells in this calculation that have both positive actual profits and positive maximal profits. Over the entire sample, 5.2% of wells have either negative actual profits or negative maximal profits.

Figure 3.8: *Fraction of Positive Profits Captured and Maximal Profits by Year, ex post*





Figure 3.9: Average Profit Maximizing Sand Use and Actual Sand Use Per Well, ex post



67.6% in 2011. Much of this growth happens in two phases. Between 2005 and 2007, the fraction increases from 15.7% to 43.9%, and between 2009 and 2010, the fraction increases from 44.8% to 65.5%. By 2011, firms earn an average of 67.6% of the maximum profits they could have earned with optimal fracking input choices.

The bottom half of Figure 3.8 shows how these maximal profits evolve over time. When oil prices were at their peak in 2008, the profit maximizing input choice for the average well would have generated \$36.1 million in profits, meaning that in 2008, foregone profits from inefficient fracking choices averaged \$21.3 million per well. By 2011, lower oil prices reduced these maximal profits to \$25.6 million per well. Combined with the higher fraction of profits captured, firms in 2011 left only \$9.9 million on the table.

Firms captured more profits by selecting more profitable fracking designs over time. In Figures 3.9 and 3.10, I plot average profit maximizing and actual input use per well over time. Though firms use less sand in fracking than the estimated profit maximizing levels, starting

Figure 3.10: *Average Profit Maximizing Water Use and Actual Water Use Per Well, ex post*





Figure 3.11: Gaussian Process Year Effects



in 2009, actual choices approach optimal choices. In 2005 and 2006, the average well was fracked with approximately 275 lbs sand per foot less than the profit maximizing level. This difference in sand use doesn't meaningful fall until reaching 132 lbs per foot in 2010. By 2011, the difference between optimal sand use and actual sand use is only 120 lbs per foot.

Though the differences in actual and optimal water use start out considerably larger than the differences in sand use, actual water choices get closer to optimal water choices in almost every year. In 2005, firms fracked the average well with 300 gals per foot less water than the water use in the optimal well. By 2011, the difference is only 98 gals per foot. These trends in actual input use towards optimal input use are consistent with the idea that firms are learning about the efficient use of fracking inputs as they observe more data, and with this knowledge they make more profitable choices.


3.4.2 Profitability vs. Productivity

The existing literature on learning in firms focuses on *productivity* instead of *profitability*. Scholars in this literature measure learning by comparing estimates of the time-varying component of Hicks-neutral productivity with the amount of experience a firm has in producing.⁴³ This approach to studying learning does not treat the production function as an object for firms to learn. Rather, the knowledge from accumulated experience serves as an *input* to the firm's production function, in the same way that labor, capital and materials do.

To determine if firms in this dataset became more productive, in addition to more profitable, I add year fixed effects to the Gaussian process production function specification, and plot their estimated values and confidence intervals in Figure 3.11.

Wells fracked in 2005 are actually 13.7% *more* productive than wells fracked in 2006. However, the confidence interval around this estimate is wide enough to include zero, as there are only 10 wells in 2005 and 20 wells in 2006. Wells fracked in later years are more productive than wells fracked in 2005 or 2006. For example, wells fracked in 2009 are 35.6% more productive than those in 2005, and 49.3% more productive than those in 2006. Again, the confidence intervals around these estimates are wide, and I cannot reject the hypothesis that there is no change in productivity between 2006 and 2009. In each of the next 2 years, productivity falls slightly, though the differences are not statistically significant. Overall, wells fracked between 2008-2011 cohorts are more productive than the earliest wells, but there is no productivity growth during 2008-2011. Since this time period covers 95% of the wells studied in this paper, I interpret this as evidence that firms learned to be more productive only in the earliest years. In contrast, the results in the previous section show that firms learned to be more profitable in all years, and especially during 2008-2011.

Figure 3.12: Fraction of Positive Profits Captured and Maximal Profits by Year, ex ante



3.4.3 *ex ante* Comparisons

Though firms make choices which approach the *ex post* estimates of optimal choices over time, those choices do not always maximize the *ex ante* estimates of expected profits. The top half of Figure 3.12 plots the ratio of actual profits to maximal profits per well using *ex ante* expectations.⁴⁴ Firms initially make fracking input choices with expected profits that are close to the optimal choices, capturing 76.0% of potential *ex ante* profits in 2007. However, profit capture actually falls over time, reaching 67.8% in 2011, approximately the same level as the *ex post* case in 2011.

While the fraction of profits captured falls, *ex ante* expectations of maximal profits rise from 2009-2011, as show in the bottom half of Figure 3.8. Unlike the *ex post* case, where the highest level of maximal profits coincides with the 2008 peak in oil prices, *ex ante* maximal profits are highest in 2011, reaching \$28.8 million per well. Though average oil prices are similar in 2008 (\$100 per bbl) and 2011 (\$95 per bbl), firms have much more information about fracking in 2011 and this information generates more optimistic expectations. The combined effect of falling *ex ante* profit capture and rising maximal profits increases foregone *ex ante* profits from \$3.1 million in 2007 to \$10.6 million in 2011.

Firms capture a shrinking fraction of *ex ante* profits over time because their actual sand use grows more slowly than the expected profit maximizing sand use does. Figure 3.13 plots average profit maximizing and actual sand use per well over time. In 2007, actual sand use is quite similar to *ex ante* optimal sand use. However, as the data firms have to learn from accumulates, optimal sand use increases faster than actual sand use, and by 2011, the difference between optimal and actual sand use reaches 131 lbs per foot. Though this difference is similar to the difference in the *ex post* case during 2011, it is striking that the differences in actual and

⁴³For example, Benkard (2000) correlates log labor requirements per unit of production with measures of experience (and forgetting), and Thornton and Thompson (2001) estimate a semi-parametric production function model in which various measures of experience are direct inputs to production.

⁴⁴As in the *ex post* case, I only include wells in this calculation that have both positive actual profits and positive maximal profits. Over the entire sample, 6.1% of wells have either negative actual profits or negative maximal profits. Half of these wells are fracked in 2009. Moreover, I further limit the set of wells by computing expected profits for the subset of wells that are fracked by firms which can observe 50 wells and 300 well-months of production history. The first wells that satisfy this criteria are not fracked until 2007.

Figure 3.13: *Average Profit Maximizing Sand Use and Actual Sand Use Per Well, ex ante*





Figure 3.14: *Average Profit Maximizing Water Use and Actual Water Use Per Well, ex ante*



optimal sand use increase over time in the *ex ante* case while decreasing in the *ex post* case.

Figure 3.14 plots average *ex ante* optimal and actual water use per well is similar to the *ex post* case in Figure 3.10: on average, firms use less than the *ex ante* optimal amount of water in fracking, but make improved water choices over time. In 2007, firms use 385 gals per foot less water than the optimal level. This difference shrinks in each year, and by 2011, it is only 107 gals per foot.

3.5 Fracking input choice model

Though firms do learn over time, many of their choices do not coincide with the predicted optimal choices, even on an *ex ante* basis. I consider two possible explanations for this phenomenon based firm preferences. First, firms may care about the uncertainty in their estimates of the profits of a fracking design. Second, in estimating the profits of a fracking design, firms may weigh their own data differently than the data generated by their competitors.

3.5.1 Preferences Over Uncertainty

In comparing the expected profits a firm earned to the maximal expected profits a firm could have earned, I have implicitly assumed that the *correct* strategy is for firms to select fracking designs solely on the basis of expected profits, without regard to the uncertainty of profits across designs. There are two potential problems with this assumption. First, viewing fracking design as an investment project selection problem, there may be financial or organizational factors that cause firms to have preference over uncertainty. Second, when learning about the performance of different fracking designs, firms may care about uncertainty through the *explore vs. exploit* tradeoff that exists in all learning problems.

Though it is appropriate for firms to ignore uncertainty in simple and frictionless models of investment project selection, there are practical reasons why uncertainty may also matter. Firms raise outside capital to finance operations and the presence of debt capital can lead firms to select fracking designs with higher uncertainty, as bond holders will bear the downside risk. On the other hand, capital constrained firms may not necessarily have the option of

selecting fracking designs with higher uncertainty if they are more expensive to implement. Financial considerations can thus push firms towards or away from fracking designs with more uncertain profits. Firms must also hire and incentivize potentially risk averse engineers, who select fracking designs. Depending on the extent of their career concerns and the structure of their compensation, engineers themselves may have preferences over uncertainty.

The prescribed learning strategies in most theoretical models of learning involve uncertainty seeking behavior. Analyses of the *explore vs. exploit* tradeoff in learning predict that agents should always do some amount of exploration, by selecting actions with more uncertain payoffs. This tradeoff will frequently require agents to sacrifice expected payoffs in the present in order to acquire uncertainty resolution in the future. Since actions with the more uncertain payoffs can resolve more future uncertainty, experimenting agents should have a positive taste for uncertainty.

Most theory models predict that agents will experiment, at least initially. In most of the settings studied by Aghion *et al.* (1991), a fully rational, expected present discounted value maximizing agent will do some amount of exploring forever and a similar result obtains in the multi-agent context studied by Bolton and Harris (1999). The implied preferences for uncertainty in both of these models arise out of the natural dynamics of learning problems. Agents are still risk neutral over their payoffs, but because there is present value to better information in the future, they prefer those actions with uncertain payoffs which can produce more future information.

Empirically, oil companies exhibit both risk seeking and risk averse behavior. The process of acquiring mineral rights for new drilling prospects and establishing the existence of oil within those prospects is an especially risky one (see, for example Walls and Dyer 1996 and Reiss 1989). However, oil companies are price takers in the world market for oil, and many use financial markets to hedge some or all of their future oil production, suggesting that firms may wish to avoid risks associated with future price fluctuations (see Haushalter 2000).

Whether the companies I study here prefer fracking input choices with more or less uncertain production is an empirical question. I estimate firm preferences over expectations

and variance of fracking designs by analyzing realized choices. To do this, I fit a multinomial logit preference model of fracking design choice in which the “utility” a firm has for fracking design j applied to well i is:

$$\begin{aligned} u_{ij} &= \xi_m (\phi P_i \mathbb{E}[DOP_{ij}] - c_i(S_j, W_j)) + \xi_s \phi P_i (\mathbb{V}[DOP_{ij}])^{\frac{1}{2}} + \epsilon_{ij} \\ &= \tilde{u}_{ij}(\xi_m, \xi_s) + \epsilon_{ij} \end{aligned}$$

where ϕ is the fraction of oil revenues firms keep, P_i is the price of oil for well i , $c_i(S_j, W_j)$ is the cost of fracking design j for well i , and ϵ_{ij} is an iid logit error. The parameters (ξ_m, ξ_s) represent the firm’s preference over expected present discounted revenues and the standard deviation of present discounted revenues, conditional on the data they have. Under this preference specification, the probability that a firm selects design j for well i is given by the standard logit formula:

$$p_{ij} = \frac{\exp(\tilde{u}_{ij})}{\sum_k \exp(\tilde{u}_{ik})}$$

The mean utilities in this preference model are linear in the expectation and standard deviation of profits to a fracking design. Preferences of this type have precedence in the theoretical learning literature. Brezzi and Lai (2002) show that a linear combination of the expectation and standard deviation of the payoff to a choice can represent a simple and efficient approximation to the Gittins index value for the choice, if the choices have independently distributed payoffs. Since Gittins and Jones (1979) show that ordinal preferences over Gittins indices result in dynamically efficient learning behavior, agents that utilize these linear approximations attain near-optimal learning. Though the profits to fracking input choices are not distributed independently, authors in the computer science and operations research literatures have found that these learning strategies also perform well in the general case. In those literatures, learning strategies which select the choice with the highest value of a linear combination of the expectation and standard deviation of payoffs are called “upper confidence bound”, or UCB strategies. Rusmevichientong and Tsitsiklis (2010) and Srinivas *et al.* (2012) have established that UCB strategies quickly identify the highest performing choice, and do so in a way which minimizes an agent’s *ex post* cumulative regret over its past choices. UCB

strategies are also reported to be in use at major technology companies, like Yahoo, Microsoft and Google (see Chapelle and Li 2011, Graepel *et al.* 2010 and Scott 2010). In all of the existing literature which utilizes UCB learning strategies, the weight on the standard deviation of the payoffs to a choice is positive, hence the “upper” in upper confidence bound strategies. This paper is not the first in economics to utilize UCB learning strategies in an empirical context. Dickstein (2013) estimates the parameters of a UCB learning strategy in a study of learning behavior by physicians.

With data on the choices firms made, expectation and standard deviation calculations made using their information sets, and oil price and fracking cost data, I estimate the parameters (ξ_m, ξ_s) using maximum likelihood. I estimate separate values of (ξ_m, ξ_s) for each of the 8 most active firms, and also estimate a pooled value of (ξ_m, ξ_s) for the industry as a whole. Table 3.8 reports these coefficient estimates, standard errors, and several measures of goodness-of-fit. All firms and the pooled industry have positive “taste” for the expectation of profits of a fracking design and negative “taste” for the standard deviation. That is, every firm appears to avoid fracking input choices with high uncertainty. I can reject risk-neutrality for all firms and for the pooled industry. In dollar terms, firms make choices as if they are willing to accept a reduction in expected profits of \$0.60 to \$0.98 for a reduction of \$1 in the standard deviation of profits.

I report three goodness-of-fit statistics. The likelihood based pseudo- R^2 , which I refer to as *LLPR*, is defined as 1 minus the ratio of the optimized log-likelihood over the log-likelihood evaluated at the null hypothesis:

$$LLPR = 1 - \frac{\log \mathcal{L}(\hat{\xi}_m, \hat{\xi}_s)}{\log \mathcal{L}(0, 0)}$$

This statistic is similar to a real R^2 in that it varies between 0 and 1, with 0 indicating that the model does not fit any better than no model and 1 indicating that the model fits the data perfectly (see Train 2009). This measure of fit indicates how far from “perfect” the fit actually is, but it does not have a “fraction of variance explained” interpretation the way a true R^2 does. I also compute the correlation between the expected input use implied by the model’s

Table 3.8: *Uncertainty Preference Model Estimates*

Firm	$\widehat{\xi}_m$	$se(\widehat{\xi}_m)$	$\widehat{\xi}_s$	$se(\widehat{\xi}_s)$	# Wells	<i>LLPR</i>	ρ_S	ρ_W
Brigham	11.05	1.05	-11.30	1.16	111	0.24	0.00	0.15
Burlington	12.02	1.25	-15.39	1.57	102	0.34	0.55	0.47
Continental	13.54	0.83	-17.26	1.05	313	0.33	0.53	0.50
EOG	5.88	0.39	-7.75	0.57	339	0.17	-0.18	0.33
Hess	10.69	0.96	-13.10	1.10	143	0.30	0.60	0.45
Marathon	15.52	1.24	-21.99	1.67	209	0.44	0.61	0.30
Whiting	10.25	0.74	-16.97	1.21	247	0.36	-0.02	0.05
XTO	11.56	1.22	-14.36	1.45	101	0.32	0.50	0.51
All	7.46	0.17	-10.39	0.23	2,605	0.23	0.50	0.40

Maximum likelihood estimates of the uncertainty preference model:

$$u_{ij} = \xi_m (\phi P_i \mathbb{E}[DOP_{ij}] - c_i(S_j, W_j)) + \xi_s \phi P_i (\mathbb{V}[DOP_{ij}])^{\frac{1}{2}} + \epsilon_{ij}$$

P_i is the price of oil for well i , $\mathbb{E}[DOP_{ij}]$ is the expectation of the present discounted value of oil production for well i when it is fracked using design j , $\mathbb{V}[DOP_{ij}]$ is the variance of the present discounted value of oil production for i under design j , $c_i(S_j, W_j)$ is the cost of implementing design j on well i , and ϵ_{ij} is an iid logit shock. *LLPR* is a likelihood-based pseudo- R^2 :

$$LLPR = 1 - \frac{\log \mathcal{L}(\widehat{\xi}_m, \widehat{\xi}_s)}{\log \mathcal{L}(0, 0)}$$

where $\mathcal{L}(\widehat{\xi}_m, \widehat{\xi}_s)$ is the likelihood of the model evaluated at the MLE and $\mathcal{L}(0, 0)$ is the likelihood of the model evaluated at the null hypothesis. ρ_S and ρ_W are the correlations of actual sand and water use decisions with their predicted values from the model.

estimated choice probabilities and actual input use, for both sand and water. If expected input use is similar to what is observed in the data, these correlations should be positive and (ideally) close to 1.

The fit of this model varies a fair amount across firms, but is generally modest. The pseudo- R^2 measures are less than 50% for all firms and for the pooled industry, suggesting that the best fitting values of the model's parameters still require a lot of support from the logit errors to rationalize firm behavior. For 6 of the 8 firms, the correlation between predicted sand use and realized sand use is positive, and for 5 it is at least 50%. The correlations between predicted and realized water use are smaller, with only 2 firms having correlations at or above 50%, but no firms have negative correlations. Though the coefficient estimates are all significantly different from zero, the low fit statistics suggest that preferences that are linear in the mean and standard deviation of profits only explain a small portion of observed behavior.

I also estimate a version of this model which includes an interaction term between expected profits and the standard deviation of profits. While learning rules which are nonlinear in the mean and standard deviation do not appear in the existing learning literature, it is possible that true firm preferences over risk and reward are more complicated than a linear model can capture. By including an interaction between expected profits and the standard deviation of profits, I allow for risk preferences that may vary with the mean. Table 3.9 reports estimates of these models. The results are qualitatively the same as Table 3.8, with all firms showing risk aversion and all but one firm showing increasingly negative taste for risk as reward increases. Goodness-of-fit measures are slightly better for these models than for the standard mean/variance models, though this is to be expected from the inclusion of an additional covariate.

Overall, Tables 3.8 and 3.9 provide evidence that firms tend to select fracking designs with higher expected profits and avoid fracking designs with higher standard deviation of profit. This behavior is not consistent with the notion that firms are actively exploring uncertain fracking designs, but it is consistent with passively learning firms that are constrained by organizational or financially motivated variance aversion.

Table 3.9: *Uncertainty Preference Model Estimates, With Interaction*

Firm	$\widehat{\xi}_m$	$se(\widehat{\xi}_m)$	$\widehat{\xi}_s$	$se(\widehat{\xi}_s)$	$\widehat{\xi}_I$	$se(\widehat{\xi}_I)$	# Wells	<i>LLPR</i>	ρ_S	ρ_W
Brigham	20.12	1.95	-8.37	1.27	-2.86	0.44	111	0.31	-0.03	0.16
Burlington	14.08	1.66	-14.81	1.58	-0.70	0.36	102	0.35	0.57	0.51
Continental	21.26	1.29	-17.25	1.12	-2.25	0.27	313	0.37	0.60	0.53
EOG	6.10	0.41	-7.49	0.58	-0.07	0.03	339	0.17	-0.17	0.35
Hess	9.80	1.16	-13.37	1.12	0.38	0.29	143	0.30	0.60	0.45
Marathon	20.55	1.72	-22.36	1.75	-1.44	0.33	209	0.45	0.67	0.25
Whiting	11.58	0.86	-17.03	1.24	-0.24	0.07	247	0.37	0.03	0.05
XTO	13.87	1.70	-14.59	1.48	-0.50	0.24	101	0.32	0.48	0.53
All	9.06	0.21	-10.59	0.23	-0.29	0.02	2,605	0.24	0.53	0.44

Maximum likelihood estimates of the uncertainty preference model:

$$u_{ij} = \xi_m (\phi P_i \mathbb{E}[DOP_{ij}] - c_i(S_j, W_j)) + \xi_s \phi P_i (\mathbb{V}[DOP_{ij}])^{\frac{1}{2}} + \xi_I (\phi P_i \mathbb{E}[DOP_{ij}] - c_i(S_j, W_j)) \left(\phi P_i (\mathbb{V}[DOP_{ij}])^{\frac{1}{2}} \right) + \epsilon_{ij}$$

P_i is the price of oil for well i , $\mathbb{E}[DOP_{ij}]$ is the expectation of the present discounted value of oil production for well i when it is fracked using design j , $\mathbb{V}[DOP_{ij}]$ is the variance of the present discounted value of oil production for i under design j , $c_i(S_j, W_j)$ is the cost of implementing design j on well i , and ϵ_{ij} is an iid logit shock. *LLPR* is a likelihood-based pseudo- R^2 :

$$LLPR = 1 - \frac{\log \mathcal{L}(\widehat{\xi}_m, \widehat{\xi}_s, \widehat{\xi}_I)}{\log \mathcal{L}(0, 0, 0)}$$

where $\mathcal{L}(\widehat{\xi}_m, \widehat{\xi}_s, \widehat{\xi}_I)$ is the likelihood of the model evaluated at the MLE and $\mathcal{L}(0, 0, 0)$ is the likelihood of the model evaluated at the null hypothesis. ρ_S and ρ_W are the correlations of actual sand and water use decisions with the predicted values from the model.

3.5.2 Own-data bias

A different explanation for firms' apparent unwillingness to select the fracking design with the largest expected profits is that I am computing expectations with respect to different beliefs than those held by firms. There are many ways that a firm's beliefs may be different than the ones I calculate: firms may have biased prior beliefs about the role of fracking design and location, they may have simpler beliefs about the functional form relating fracking design and location to production, or my fracking cost and oil price data could be different from the costs and prices firms experience. However, using the data that I have, I am only able to test a simpler explanation. I assume that firms do have the belief structure I have described here, but do not necessarily treat all of the data available to them equally. In particular, firms may weigh data from their own experiences differently than data from the experiences of other firms that they observe through the public disclosure process. I refer to this explanation as "own-data bias".

To test for this phenomenon, I introduce a new parameter, $\lambda \in (0, 1)$, which represents the firm's relative weighting scheme. If $\lambda = 0$, the firm places no weight on the data generated by other firms and if $\lambda = 1$, the firm places no weight on its own data, relying entirely on outside data to learn. At $\lambda = \frac{1}{2}$, the firm puts equal weight on its own data and the data generated by others, which gives the preference model described in the previous section. For each value of λ , I can compute the expectation and standard deviation of *weighted* discounted profits for well i with fracking design j , for which I provide a calculation in the appendix. I then use these weighted profits in the same multinomial logit choice model described in the previous section, and refer to the choice model with weighted estimates as the weighted preference model.

In Table 3.10, I report maximum likelihood estimates of λ , as well as the other preference model coefficients, for the same specification in Table 3.8. The estimated value of λ is less than $\frac{1}{2}$ for all individual firms, and for 5 firms, the 95% confidence intervals do not include $\frac{1}{2}$. The pooled estimate is also less than $\frac{1}{2}$ and its 95% confidence interval does not include $\frac{1}{2}$. Comparing Tables 3.8 and 3.10, the preference model coefficients do change slightly, but allowing for weighted beliefs does not affect the previous conclusion that all firms dislike

Table 3.10: *Weighted Uncertainty Preference Model Estimates*

Firm	$\widehat{\xi}_m$	$se(\widehat{\xi}_m)$	$\widehat{\xi}_s$	$se(\widehat{\xi}_s)$	$\widehat{\lambda}$	$se(\widehat{\lambda})$	# Wells	<i>LLPR</i>	ρ_S	ρ_W
Brigham	14.37	1.33	-17.27	1.71	0.12	0.04	111	0.30	0.04	0.16
Burlington	13.95	1.46	-17.98	1.84	0.41	0.05	102	0.36	0.55	0.50
Continental	18.26	1.10	-22.38	1.36	0.36	0.02	313	0.37	0.63	0.55
EOG	7.01	0.45	-10.65	0.75	0.15	0.04	313	0.20	-0.12	0.36
Hess	10.87	1.00	-14.13	1.19	0.46	0.05	143	0.31	0.62	0.46
Marathon	21.77	1.72	-27.79	2.12	0.34	0.03	209	0.47	0.68	0.24
Whiting	9.82	0.70	-16.00	1.14	0.00		247	0.36	0.04	0.06
XTO	12.98	1.52	-15.84	1.63	0.44	0.04	101	0.32	0.48	0.53
All	8.22	0.19	-11.74	0.26	0.38	0.01	2,605	0.24	0.54	0.45

Maximum likelihood estimates of the uncertainty preference model:

$$u_{ij} = \xi_m (\phi P_i \mathbb{E}[DOP_{ij} | \lambda] - c_i(S_j, W_j)) + \xi_s \phi P_i (\mathbb{V}[DOP_{ij} | \lambda])^{\frac{1}{2}} + \epsilon_{ij}$$

P_i is the price of oil for well i , $\mathbb{E}[DOP_{ij} | \lambda]$ is the expectation of the present discounted value of oil production for well i when it is fracked using design j , $\mathbb{V}[DOP_{ij} | \lambda]$ is the variance of the present discounted value of oil production for i under design j , λ is the weighting parameter, $c_i(S_j, W_j)$ is the cost of implementing design j on well i , and ϵ_{ij} is an iid logit shock. *LLPR* is a likelihood-based pseudo- R^2 :

$$LLPR = 1 - \frac{\log \mathcal{L}(\widehat{\xi}_m, \widehat{\xi}_s, \widehat{\lambda})}{\log \mathcal{L}(0, 0, \frac{1}{2})}$$

where $\mathcal{L}(\widehat{\xi}_m, \widehat{\xi}_s, \widehat{\lambda})$ is the likelihood of the model evaluated at the MLE and $\mathcal{L}(0, 0, \frac{1}{2})$ is the likelihood of the model evaluated at the null hypothesis. ρ_S and ρ_W are the correlations of actual sand and water use decisions with the predicted values from the model. Because Whiting's estimate of λ is at the boundary, standard errors are computed with respect to ξ_m and ξ_s only.

uncertainty in the profits of a fracking input choice. Firms are willing to trade \$0.61 to \$0.83 in expected profits for a reduction of \$1 in the standard deviation of profits, which is a similar range to the model estimated in Table 3.8. The fit of the model in Table 3.10 is somewhat better than the model in Table 3.8, but it is still modest.

3.6 Conclusion

This paper provides one of the first empirical analyses of learning behavior in firms using operational choices, realized profits, and information sets. Oil companies in the North Dakota

Bakken Shale learned to more efficiently use fracking technology between 2005-2011, increasing their capture of possible profits from 15.7% to 67.6% by making improved fracking design choices over time. Contrary to the predictions of most theoretical models of learning, I do not find evidence that firms actively experiment in order to learn. Instead, firms prefer fracking input choices with lower variance, and are willing to give up \$0.60-0.98 in expected profits for a reduction of \$1 in the standard deviation of profits. Finally, firms in my data appear to overweight data from their own operations relative to the data they observe from their competitors.

From a neoclassical economics perspective, it is surprising that these firms do not experiment, even though it is valuable to do so. They operate in an industry known for its appetite for risk and use of advanced technology and have access to a wealth of data to learn from. However, they leave money on the table. Across the 2,699 wells in this data, the average well appears to forego \$12.1 million in profits on an *ex post* basis and \$7.6 million on an *ex ante* basis, resulting in \$20-33 billion in lost profits.

These results complement recent work by petroleum engineers on their own failures to learn to use to new technologies in a variety of contexts. Authors in this literature note that explicit learning efforts like experiments do happen, but less frequently and later in the development of a formation than they should.⁴⁵ Much of this research cites two hurdles to learning: a tendency by operators to prematurely focus their optimization efforts on cost reductions instead of improvements in operational choices, and the absence of incentive contracts between operators and their service contractors. The first phenomenon suggests that operators *believe* they know the production function with high certainty, but later discover their beliefs were wrong. In future work, I plan to incorporate this possibility into my model of input choice under uncertainty. The second phenomenon raises important questions about the effects of contractual incompleteness on the oil and gas exploration industry that I hope to study in future work.

⁴⁵For a detailed overview of this literature, see Vincent (2012)

References

- ACHARYA, V., ENGLE, R., FIGLEWSKI, S., LYNCH, A. and SUBRAHMANYAM, M. (2009). Centralized clearing for credit derivatives. *Financial Markets, Institutions+ Instruments*, **18** (2), 168.
- AGHION, P., BOLTON, P., HARRIS, C. and JULLIEN, B. (1991). Optimal learning by experimentation. *The Review of Economic Studies*, **58** (4), pp. 621–654.
- AKINS, B. (2012). Financial reporting quality and uncertainty about credit risk among the ratings agencies. *MIT Sloan School of Management Working Paper*.
- ALEXANDER, G. J., EDWARDS, A. K. and FERRI, M. G. (2000). The determinants of trading volume of high-yield corporate bonds. *Journal of Financial Markets*, **3** (2), 177–204.
- ANAND, B. N. and KHANNA, T. (2000). Do firms learn to create value? the case of alliances. *Strategic management journal*, **21** (3), 295–315.
- ARROW, K. J. (1962). The economic implications of learning by doing. *The Review of Economic Studies*, **29** (3), pp. 155–173.
- ASQUITH, P., AU, A. S., COVERT, T. and PATHAK, P. A. (2013). The market for borrowing corporate bonds. *Journal of Financial Economics*, **107** (1), 155–182.
- , — and PATHAK, P. A. (2006). Short sales constraints. *MIT Sloan School of Management Working Paper*.
- and MEULBROEK, L. (1995). An empirical investigation of short interest. *MIT Sloan School of Management Working Paper*.
- , OMAN, R. and SAFAYA, C. (2010). Short sales and trade classification algorithms. *Journal of Financial Markets*, **13** (1), 157–173.
- , PATHAK, P. A. and RITTER, J. R. (2005). Short interest, institutional ownership, and stock returns. *Journal of Financial Economics*, **78** (2), 243–276.
- BAI, Y., CHANG, E. C. and WANG, J. (2006). Asset prices under short-sales constraints. *MIT Sloan School of Management Working Paper*.
- BAIHLY, J., ALTMAN, R. and AVILES, I. (2012). Has the economic stage count been reached in the bakken shale? In *SPE Hydrocarbon Economics and Evaluation Symposium*.

- BAO, J. and PAN, J. (2013). Bond illiquidity and excess volatility. *Review of Financial Studies*, **26** (12), 3068–3103.
- BENKARD, C. (2000). Learning and forgetting: The dynamics of aircraft production. *The American Economic Review*.
- BESSEMBINDER, H., KAHLE, K. M., MAXWELL, W. F. and XU, D. (2009). Measuring abnormal bond performance. *Review of Financial Studies*, **22** (10), 4219–4258.
- and MAXWELL, W. (2008). Markets: Transparency and the corporate bond market. *The Journal of Economic Perspectives*, pp. 217–234.
- , — and VENKATARAMAN, K. (2006). Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, **82** (2), 251–288.
- BIAIS, B., GLOSTEN, L. and SPATT, C. (2005). Market microstructure: A survey of micro-foundations, empirical results, and policy implications. *Journal of Financial Markets*, **8** (2), 217–264.
- and GREEN, R. C. (2007). The microstructure of the bond market in the 20th century. *Toulouse University Working Paper*.
- BLOOMFIELD, R. and O’HARA, M. (1999). Market transparency: Who wins and who loses? *Review of Financial Studies*, **12** (1), 5–35.
- BOEHMER, E., JONES, C. M. and ZHANG, X. (2008). Which shorts are informed? *The Journal of Finance*, **63** (2), 491–527.
- BOLTON, P. and HARRIS, C. (1999). Strategic experimentation. *Econometrica*, **67** (2), pp. 349–374.
- BRAVO, R. A. (2003). Corporate-bond pricing system may be too open, critics say. *Wall Street Journal*.
- BREZZI, M. and LAI, T. L. (2002). Optimal learning and experimentation in bandit problems. *Journal of Economic Dynamics and Control*, **27** (1), 87–108.
- CHAPELLE, O. and LI, L. (2011). An empirical evaluation of thompson sampling. In *Advances in Neural Information Processing Systems*, pp. 2249–2257.
- CHEN, J., HONG, H. and STEIN, J. C. (2002). Breadth of ownership and stock returns. *Journal of Financial Economics*, **66** (2), 171–205.
- COHEN, L., DIETHER, K. B. and MALLOY, C. J. (2007). Supply and demand shifts in the shorting market. *The Journal of Finance*, **62** (5), 2061–2096.
- CONLEY, T. and UDRY, C. (2010). Learning about a new technology: Pineapple in ghana. *The American Economic Review*, **100** (1), 35–69.

- CORTS, K. and SINGH, J. (2004). The effect of repeated interaction on contract choice: Evidence from offshore drilling. *Journal of Law, Economics, and Organization*, **20** (1), 230–260.
- CORWIN, S. A. and SCHULTZ, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *The Journal of Finance*, **67** (2), 719–760.
- DECKER, M. (2007). Finra’s trace and the u.s. corporate bond market, presentation to Securities Industry and Financial Markets Association.
- DEUTSCH, J. (2011). *Secretary of Energy Advisory Board Shale Gas Production Subcommittee Second Ninety Day Report*. Tech. rep., Department of Energy.
- DIAMOND, D. W. and VERRECCHIA, R. E. (1987). Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics*, **18** (2), 277–311.
- DICKSTEIN, M. (2013). *Efficient provision of experience goods: Evidence from antidepressant choice*. Working paper, Stanford University.
- DIETHER, K. B., LEE, K.-H. and WERNER, I. M. (2009). Short-sale strategies and return predictability. *Review of Financial Studies*, **22** (2), 575–607.
- DUFFIE, D. (1996). Special repo rates. *The Journal of Finance*, **51** (2), 493–526.
- (2011). *Dark markets: Asset pricing and information transmission in over-the-counter markets*. Princeton University Press.
- , GARLEANU, N. and PEDERSEN, L. H. (2002). Securities lending, shorting, and pricing. *Journal of Financial Economics*, **66** (2), 307–339.
- D’AVOLIO, G. (2002). The market for borrowing stock. *Journal of Financial Economics*, **66** (2), 271–306.
- ECONOMIST, T. (2011). Derivatives: Unlucky for some. proposed rules on taming swaps markets are proving controversial. *The Economist*.
- EDWARDS, A. K., HARRIS, L. E. and PIWOWAR, M. S. (2007). Corporate bond market transaction costs and transparency. *The Journal of Finance*, **62** (3), 1421–1451.
- FETKOVICH, M. (1980). Decline curve analysis using type curves. *Journal of Petroleum Technology*, **32** (6), 1065–1077.
- FIGLEWSKI, S. and WEBB, G. P. (1993). Options, short sales, and market completeness. *The Journal of Finance*, **48** (2), 761–777.
- FORUM, F. S. (2013). *Futures Industry Association, International Swaps and Derivatives Association and Securities Industry and Financial Markets Association*. Tech. rep., Comment Letter RE: Phase-In Schedule for Requirement for Title VII of Dodd-Frank Act.
- FOSTEL, A. and GEANAKOPOLOS, J. (2008). Leverage cycles and the anxious economy. *The American Economic Review*, pp. 1211–1244.

- FOSTER, A. and ROSENZWEIG, M. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, pp. 1176–1209.
- FOX, J. and AHLBRANDT, T. (2002). *Petroleum geology and total petroleum systems of the Widyan Basin and Interior Platform of Saudi Arabia and Iraq*. US Department of the Interior, US Geological Survey.
- FRENCH, K., BAILY, M., CAMPBELL, J., COCHRANE, J., DIAMOND, D., DUFFIE, D., KASHYAP, A., MISHKIN, F., RAJAN, R., SCHARFSTEIN, D. *et al.* (2010). The squam lake report: Fixing the financial system*. *Journal of Applied Corporate Finance*, **22** (3), 8–21.
- GASWIRTH, S. B. (2013). *Assessment of Undiscovered Oil Resources in the Bakken and Three Forks Formations, Williston Basin Province, Montana, North Dakota, and South Dakota, 2013*. Tech. rep., U.S. Geological Survey.
- GECZY, C. C., MUSTO, D. K. and REED, A. V. (2002). Stocks are special too: An analysis of the equity lending market. *Journal of Financial Economics*, **66** (2), 241–269.
- GILJE, E. (2012). *Does Local Access To Finance Matter?: Evidence from US Oil and Natural Gas Shale Booms*. Working paper, University of Pennsylvania.
- GITTINS, J. C. and JONES, D. M. (1979). A dynamic allocation index for the discounted multiarmed bandit problem. *Biometrika*, **66** (3), 561–565.
- GOLDSTEIN, M. A. and HOTCHKISS, E. S. (2007). Dealer behavior and the trading of newly issued corporate bonds. In *AFA 2009 San Francisco Meetings Paper*.
- , — and SIRRI, E. R. (2007). Transparency and liquidity: A controlled experiment on corporate bonds. *Review of Financial Studies*, **20** (2), 235–273.
- GRAEPEL, T., CANDELA, J. Q., BORCHERT, T. and HERBRICH, R. (2010). Web-scale bayesian click-through rate prediction for sponsored search advertising in microsoft’s bing search engine. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 13–20.
- GREEN, R. C., HOLLIFIELD, B. and SCHÜRHOFF, N. (2007a). Dealer intermediation and price behavior in the aftermarket for new bond issues. *Journal of Financial Economics*, **86** (3), 643–682.
- , — and SCHÜRHOFF, N. (2007b). Financial intermediation and the costs of trading in an opaque market. *Review of Financial Studies*, **20** (2), 275–314.
- GREENSTONE, M., OYER, P. and VISSING-JORGENSEN, A. (2006). Mandated disclosure, stock returns, and the 1964 securities acts amendments. *The Quarterly Journal of Economics*, **121** (2), 399–460.
- GRILICHES, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica, Journal of the Econometric Society*, pp. 501–522.
- HARRISON, J. M. and KREPS, D. M. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, **92** (2), 323–336.

- HAUSHALTER, G. D. (2000). Financing policy, basis risk, and corporate hedging: Evidence from oil and gas producers. *The Journal of Finance*, **55** (1), 107–152.
- HICKS, B. E. (2012). North dakota oil & gas update. In *Presented for Dunn County Oil Day in Killdeer, ND on February 21, 2012*, North Dakota Industrial Commission.
- HOUGH, E. and MCCLURG, T. (2011). Impact of geological variation and completion type in the u.s. bakken oil shale play using decline curve analysis and transient flow character. In *AAPG International Conference and Exhibition, Milan, Italy, October 23-26, 2011*.
- JABBARI, H., ZENG, Z. *et al.* (2012). Hydraulic fracturing design for horizontal wells in the bakken formation. In *46th US Rock Mechanics/Geomechanics Symposium*, American Rock Mechanics Association.
- JAMIESON, D. (2006). Report card for trace a mixed bag. *InvestmentNews*.
- JENSEN, R. (2007). The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The quarterly journal of economics*, **122** (3), 879–924.
- JONES, C. M. and LAMONT, O. A. (2002). Short-sale constraints and stock returns. *Journal of Financial Economics*, **66** (2), 207–239.
- KASY, M. (2013). Why experimenters should not randomize, and what they should do instead. *Working Paper*.
- KELLOGG, R. (2011). Learning by drilling: Interfirm learning and relationship persistence in the texas oilpatch. *The Quarterly Journal of Economics*, **126** (4), 1961–2004.
- KETCHUM, G. R. (2012). *Speech: Bond Dealers Association Annual Conference*. Tech. rep., Bond Dealers Association.
- KOGAN, L., PAPANIKOLAOU, D., SERU, A. and STOFFMAN, N. (2012). *Technological Innovation, Resource Allocation, and Growth*. Working Paper 17769, National Bureau of Economic Research.
- KOLASINSKI, A. C., REED, A. V. and RINGGENBERG, M. C. (2013). A multiple lender approach to understanding supply and search in the equity lending market. *The Journal of Finance*, **68** (2), 559–595.
- LARSEN, L. C. (2011). Horizontal drafting: Why your form joa may not be adequate for your company’s horizontal drilling program. *Rocky Mtn. Min. L. Found. J.*, **48**, 51.
- LEARNER, H. (2011). *An Examination of Transparency in European Bond Markets*. Tech. rep., CFA Institute.
- LEVITT, A. (1999). *Testimony to the U.S. Congress House of Representatives House Subcommittee on Finance and Hazardous Materials*. Tech. rep., FINRA.
- LEVITT, C. J. (2011). Learning through oil and gas exploration. *Working Paper*.

- LEVITT, S. D., LIST, J. A. and SYVERSON, C. (2012). *Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant*. Working Paper 18017, National Bureau of Economic Research.
- LINTNER, J. (1971). The effect of short selling and margin requirements in perfect capital markets. *Journal of Financial and Quantitative Analysis*, **6** (5), 1173–1195.
- MADHAVAN, A. (1995). Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies*, **8** (3), 579–603.
- MCCARTHY, K., ROJAS, K., NIEMANN, M., PALMOWSKI, D., PETERS, K. and STANKIEWICZ, A. (2011). Basic petroleum geochemistry for source rock evaluation. *Oilfield Review*, **23** (2), 32–43.
- MILLER, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of Finance*, **32** (4), 1151–1168.
- MUEHLENBACHS, L., SPILLER, E. and TIMMINS, C. (2012). *Shale Gas Development and Property Values: Differences across Drinking Water Sources*. Working Paper 18390, National Bureau of Economic Research.
- MULLEN, D. (2004). *Relating to Proposed Amendments to TRACE Rule 6250 and Related TRACE Rules to Disseminate Transaction Information on All TRACE-Eligible Securities and Facilitate Dissemination*. Tech. rep., Letter from Bond Market Association to SEC.
- NAIK, N. Y., NEUBERGER, A. and VISWANATHAN, S. (1999). Trade disclosure regulation in markets with negotiated trades. *Review of Financial Studies*, **12** (4), 873–900.
- NASD (2005a). *NASD’s Fully Implemented ‘TRACE’ Brings Unprecedented Transparency to Corporate Bond Market*. Tech. rep., National Association of Securities Dealers.
- (2005b). *TRACE Factbook*. Tech. rep., National Association of Securities Dealers.
- (2006). *NASD Response to European Commission; Letter in Response to EU Markets in Financial Instruments Directive (MiFiD) Call for Evidence on Pre- and post-trade transparency provisions requested June 12, 2006*. Tech. rep., National Association of Securities Dealers.
- NASHIKKAR, A. and PEDERSEN, L. H. (2007). Corporate bond specialness. *Unpublished working paper*. New York University.
- NAZARETH, A. L. (2004). *Testimony before the United States Senate Committee on Banking, Housing, and Urban Affairs*. Tech. rep., SEC.
- OFEK, E., RICHARDSON, M. and WHITELAW, R. F. (2004). Limited arbitrage and short sales restrictions: Evidence from the options markets. *Journal of Financial Economics*, **74** (2), 305–342.
- ORTEGA, C. E., AGUILERA, R. *et al.* (2012). Use of drill cuttings for improved design of hydraulic fracturing jobs in horizontal wells. In *SPE Americas Unconventional Resources Conference*, Society of Petroleum Engineers.

- PAGANO, M. and RÖELL, A. (1996). Transparency and liquidity: a comparison of auction and dealer markets with informed trading. *The Journal of Finance*, **51** (2), 579–611.
- PIMMEL, A. and CLAYPOOL, G. (2001). *Introduction to shipboard organic geochemistry on the JOIDES Resolution*. <http://www-odp.tamu.edu>.
- RASMUSSEN, C. E. and WILLIAMS, C. K. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- REISS, P. C. (1989). *Economic and Financial Determinants of Oil and Gas Exploration Activity*. Working Paper 3077, National Bureau of Economic Research.
- ROMER, P. M. (1986). Increasing returns and long-run growth. *Journal of Political Economy*, **94** (5), pp. 1002–1037.
- RUSMEVICHIENTONG, P. and TSITSIKLIS, J. N. (2010). Linearly parameterized bandits. *Mathematics of Operations Research*, **35** (2), 395–411.
- RYAN, B. and GROSS, N. C. (1943). The diffusion of hybrid seed corn in two iowa communities. *Rural sociology*, **8** (1), 15–24.
- SAPUTELLI, L. A., SOLIMAN, M. Y., LOPEZ, C. M., CHACON, A. J. *et al.* (2014). Design optimization of horizontal wells with multiple hydraulic fractures in the bakken shale. In *SPE/EAGE European Unconventional Resources Conference and Exhibition*, Society of Petroleum Engineers.
- SCHULTZ, P. (2012). The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors. *Journal of Financial Economics*, **106** (3), 492–512.
- SCOTT, S. L. (2010). A modern bayesian look at the multi-armed bandit. *Applied Stochastic Models in Business and Industry*, **26** (6), 639–658.
- SHELLENBERGER, M., NORDHAUS, T., TREMBATH, A. and JENKINS, J. (2012). *Where the Shale Gas Revolution Came From*. Tech. rep., The Breakthrough Institute.
- SHELLEY, R., GULIYEV, N. and NEJAD, A. (2012). A novel method to optimize horizontal bakken completions in a factory mode development program. In *SPE Annual Technical Conference and Exhibition*.
- SHENN, J. and SCHEER, D. (2009). Finra seeks to amass asset-backed bond data on trace. *Bloomberg*.
- SIFMA (2009). *Report on Outstanding U.S. Bond Market Debt*. Tech. rep., Securities Industry and Financial Markets Association.
- (2013). *Table 1.1 US Corporate Average Daily Trading Volume by Type and Rating*. Tech. rep., Securities Industry and Financial Markets Association.
- SITCHLER, J. C., CHERIAN, B. V., PANJAITAN, M. L., NICHOLS, C. M., KRISHNAMURTHY, J. K. *et al.* (2013). Asset development drivers in the bakken and three forks. In *SPE Hydraulic Fracturing Technology Conference*, Society of Petroleum Engineers.

- SRINIVAS, N., KRAUSE, A., KAKADE, S. M. and SEEGER, M. (2012). Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *Information Theory, IEEE Transactions on*, **58** (5), 3250–3265.
- STIGLER, G. J. (1964). Public regulation of the securities markets. *Journal of Business*, pp. 117–142.
- STOYANOV, A. and ZUBANOV, N. (2012). Productivity spillovers across firms through worker mobility. *American Economic Journal: Applied Economics*, **4** (2), 168–198.
- THORNTON, D. L. (2009). What the libor-ois spread says. *Economic Synopses*, **2009** (2009-05-11).
- THORNTON, R. and THOMPSON, P. (2001). Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding. *American Economic Review*, pp. 1350–1368.
- TRAIN, K. (2009). *Discrete choice methods with simulation*. Cambridge university press.
- VAMES, S. (2003). Nasd expands bond-price reports – data on corporate issues becomes more transparent. *New York Times*.
- VIDIC, R., BRANTLEY, S., VANDENBOSSCHE, J., YOXTHEIMER, D. and ABAD, J. (2013). Impact of shale gas development on regional water quality. *Science*, **340** (6134).
- VINCENT, M. (2012). The next opportunity to improve hydraulic-fracture stimulation. *Journal of Petroleum Technology*, **64** (3), 118–127.
- WALLS, M. R. and DYER, J. S. (1996). Risk propensity and firm performance: A study of the petroleum exploration industry. *Management Science*, **42** (7), 1004–1021.
- WIELAND, V. (2000). Learning by doing and the value of optimal experimentation. *Journal of Economic Dynamics and Control*, **24** (4), 501–534.

Appendix A

Appendix to Chapter 2

A.1 TRACE Data Appendix

The TRACE dataset we use was purchased directly from FINRA. We refer to this data as the Historical FINRA TRACE dataset since it contains both disseminated and non-disseminated trades, indicated by the Dissemination Flag field (DISSEM_FL). There are also two TRACE datasets available on WRDS, one from FINRA and one provided by Mergent FISD. These databases only contain disseminated trade records that were available to market participants in real time and do not contain any non-disseminated trades.

The Historical FINRA TRACE dataset contains 35,284,669 unique trade records, on 35,695 different CUSIPS, for July 1, 2002 until December 31, 2006. All FINRA trade records are self-reported by FINRA members. To each self-reported transaction report, FINRA adds the time and date that it received the report and a flag indicating whether or not the report was disseminated to the public. Then, FINRA generates a message sequence number that is unique within the reporting day. For transactions that are modifies or cancels, the message sequence number of the original trade is included as a separate field called the “original” message sequence number.

We first take the Historical FINRA TRACE dataset and match it to the universe of corporate bonds in the Mergent FISD database, our source for bond characteristics such as

issue size, ratings, maturity, etc. We drop all TRACE bonds that do not match to FISD. We also drop all bonds with equity-like characteristics (convertibles, exchangeables, etc.) since their equity component may be included in the bond price. We next drop all Rule 144a bonds because TRACE does not report trading information on these bonds. Finally, FISD does not report the correct issue size in all cases. For example, there are some bonds with a reported issue size of \$0. After hand checking a number of cases with small issue size, we decided to drop all bonds with reported FISD issue size of less than \$100,000, including those with issue size of \$0. The number of trades eliminated and their corresponding CUSIPs affected by these steps are shown in first section of Table A.1.

FINRA reports prices on bond trades differently for principal and agency trades, denoted by the Buyer Capacity field (BUY_CPCTY_CD) or Seller Capacity field (SELL_CPCTY_CD). Prices reported on principal trades include “any markups or markdowns.” Prices reported on agency trades do not include the commission charged “since commission is reported in a separate field.” (TRACE USER Guide, version 2.2, page 17) To make our prices comparable across all trades, we adjust trade prices for the commission paid whenever the buy or sell commission field is non-empty. A total of 699,833 trades representing 19,999 CUSIPs are modified.

Not all remaining trade records are unique or correct. We eliminate trade records for four main reasons. First, some trades do not take place; they are later modified, revised, or cancelled. Second, some trades are reported more than once. Third, some trade records have erroneous price or volume data. Fourth, some trades have problems with their trade date. Table A.1 reports the number of trade records affected for each of these reasons.

TRACE generates extra trade records for modified, cancelled, or reversed trades. Trades cancelled within a day are marked as cancelled, while trades cancelled on subsequent days are marked as reversals. When identifying trades that are subsequently modified, cancelled or reversed, we rely on three data fields: Trade Status (TRC_ST), Message Sequence Number (MSG_SEQ_NB), and Original Message Sequence Number (ORIG_MSG_SEQ_NB). The first field, Trade Status, has a value of “W” if the trade record is a modification and “C” if it is

Table A.1: Steps from Historical TRACE File to Cleaned Historical TRACE Sample

	CUSIPs (1)	Trades (2)
Source: Historical TRACE File	35,695	35,284,669
Eliminate bonds based on characteristics		
Bonds unmatched to FISD by CUSIP	1,156	200,482
Convertible bonds	1,375	2,297,404
Exchangeable bonds	113	20,420
Other equity-linked bonds	615	221,543
SEC Rule 144a bonds	551	38,583
Bonds with missing issue size or issue size < 100,000	348	95,723
Adjust prices for commissions paid (these transactions are NOT dropped, prices are changed)	19,999	699,833
Eliminate trades which do not take place or when dissemination is delayed		
Modifies: matched to earlier record using sequence number	23,877	883,146
Cancels: matched to earlier record using sequence number	22,557	504,789
Cancels: not matched to earlier record	2,502	4,201
Reversals: matched to earlier record using seven-way match	21,234	837,740
Reversals: matched to earlier record using six-way match	7,507	64,700
Reversals: not matched to earlier record	9,216	44,849
Delayed reversals	2,794	60,698
Delayed disseminations	612	1,382
Eliminate trades which are reported more than once		
Dealer buys (total)	28,417	6,578,859
Dealer buys matched to dealer sells with the same execution time	20,710	590,372
Dealer buys matched to dealer sells with different execution times	26,703	4,468,884
Unmatched dealer buys	26,703	1,519,603
Agency trades (agency buys matched to agency sells)	18,424	563,658
Eliminated trades with price and volume data issues		
Trades with missing price	106	211
Trades with 0 price	0	0
Trades with negative price	0	0
Trades with price greater than 220	497	806
Trades with volume/issue amount ≥ 50% and par value or issue amount is not equal to 0 or 1	2,222	4,597
Trades with volume less than 1000	2,189	5,323
Eliminated trades with timing issues		
Trades executed before bond's offering date	6,654	247,012
Trades executed after bond's maturity date	342	78,052
Trades with different reporting and execution dates	30,736	799,570
Trades that occur on SIFMA Holidays	30,671	581,396
Cleaned Historical TRACE Sample	30,643	21,149,525

Filters are applied sequentially. This table reports the steps from the historical TRACE file to the Clean Historical TRACE file. Other equity-linked bonds have "KNOCK", "REVERSE", "EQUITY", "LINKED", and "TBD" in the bond's FISD issue name. A seven-way match is based on CUSIP, execution date, execution time, price, quantity, buy-sell indicator, and dealer-customer indicator. A six-way match drops the execution time requirement. Price cutoff of 220 is based on computing a bond's value based on its maturity, coupons remaining and lowest value of the treasury yield curve during our entire sample period. The maximum for our sample is 214, which we round to 220. SIFMA holidays correspond to "Recommended Early Close" and "Recommended Full Close" dates listed at <http://www.sifma.org/uploadedfiles/research/statistics/statisticsfiles/misc-us-historical-holiday-market-recommendations-sifma.pdf>

a cancellation of another trade record received by FINRA on the same reporting day. If the trade record is not a cancellation or a modification, Trade Status has a value of “T”. For each reporting day, FINRA assigns every trade report a unique Message Sequence Number. The third field, Original Message Sequence Number, is blank unless the trade is a cancellation or modification. In those instances, the Original Message Sequence Number field contains the Message Sequence Number of the trade record that is being cancelled or modified.

We link together the Original Message Sequence Numbers and the Message Sequence Numbers to create chains of trade records, if such chains exist. Each chain starts with the original trade and ends with the last modified or cancelled trade. The existence of a chain means that some trades will be eliminated. For example, if a trade that is reported with message sequence number “1” is modified in a record with message sequence number “2” and then later modified in a record with message sequence number “3”, we eliminate the first two records. We link these three records together, starting from record 3 and moving backwards until we reach a record with the original message sequence number (the original trade). In a similar manner, if a trade is later canceled, we link the canceled trade record with the original trade record and eliminate both trade records from our sample. Since it is possible for different trades in a chain to have different dissemination statuses, we treat the terminal trade record in a chain as disseminated if any trade record in the chain was disseminated. If a trade is modified, we remove all records except the last one in the chain. For cancelled trades, we eliminate the entire chain of trades. When a cancelled trade cannot be matched to the original trade, it alone is eliminated.

Trade reversals are identified by the As Of Indicator field (ASOF_CD) by “R”. Reversals cannot be tracked using Message Sequence Number and Original Message Sequence Number because the original trade and its reversal are reported to FINRA on different days. Message Sequence Numbers are unique within each reporting day and cannot be linked across days. Therefore, to link a reversal to its original trade, we match based on seven identifying characteristics: CUSIP, execution date, execution time, price, volume, indicator for buyer or seller, and indicator for dealer or customer. Matches with these criteria represent a “seven-way”

match.

Using these seven trade characteristics, however, often leads to a many-to-many match; that is, there is often more than one possible pairing. (In fact, it appears that many reversals are the result of a trade being entered twice and the second record being reversed). After matching reversal and non-reversal reports using these seven trade characteristics, we eliminate the minimum number of exact matches. If there is only one exact match, both the reversal and its matched trade are eliminated. If there is more than one exact match, we eliminate the reversal trade and one of the matching trades. Since, by definition, the trades occur at the same time, date, price, and volume, these characteristics are unaffected by the choice of which matching trade we eliminate. For example, if there are 4 reversals and 5 non-reversals, we drop the 4 reversals and drop the first 4 non-reversals. A total of 837,740 trade records were dropped as part of reversals.

Unfortunately, not all reversals have an exact seven-way match. A large number of the unmatched reversals had a six-way match to another trade if we drop the same execution time requirement. Since execution time is self-reported, we assume these six-way matches were the original trades that were meant to be reversed, and we eliminate the reversal and the matched trade from the sample following the steps above. Even after six-way matches, there are 44,849 records labeled as reversals that we were unable to match to an original trade. We dropped these reversals from the dataset, but were unable to identify the original matched trade. In addition, the As Of Code field can also take on the values “X” corresponding to delayed reversal and “D” corresponding to delayed dissemination. We drop these records as well.

After eliminating modified, cancelled, and reversed records (which represent trades that do not actually take place), we next deal with trades that take place, but are reported more than once. There are two ways for duplicate records to occur in the Historical Trace database. The first involves transactions between two dealers, who both report the trade to FINRA, one as a buy trade and one as a sell trade. Both the buy and sell side of inter dealer trades are included in the Historical TRACE database released by FINRA. In the Mergent FISD

TRACE dataset on WRDS, only the seller's report of an inter dealer trade is included. To mimic this convention, we eliminate trade records for inter dealer trades submitted by the buying dealer. We drop a total of 6,578,859 trade records where the transaction is labeled as a buy (using the Buy/Sell Indicator, RPT_SIDE_CD) and where the transaction is associated with a duplicate sell transaction. In addition, both the buy and sell records must be labeled as inter dealer trades (using the Contra Party Indicator, CNTRA_MP_ID).

As shown in Table A.1, of the 6,578,859 dealer buy records we eliminated, 590,372 correspond to sell records with the same CUSIP, date, price, quantity, inter dealer trade indicator, and execution time. An additional 4,468,884 buy records correspond to sell records that have the same CUSIP, date, price, quantity, inter dealer trade indicator, but a different execution time. As mentioned earlier, execution time is self-reported by both the buyer and seller and we believe that two trades which have a match of all other characteristics other than time are probably duplicates. Finally, 1,519,603 dealer buy records cannot be matched to a dealer sell record. These records are a puzzle given that we have no record of a seller reporting the trade in an inter dealer transaction. To be conservative, we eliminate these remaining unmatched inter dealer buy records.

The second way duplicate trades appear in the database is when the dealer acts as in an agent capacity. If a dealer acts as an agent for a customer, FINRA asks that trade be reported as if the agent "stood between the customer and the contra party" (TRACE USER Guide, version 2.2, page 21). That is, if a dealer sold bonds as an agent for their customer to another party, they would report two records to TRACE: a buy transaction from the customer and a sell transaction to the other party, even though this is a single transaction. We keep only the sell transaction when there is both an agency buy transaction and agency sell transaction with the same price, quantity, execution date, and execution time. This rule leads us to drop 563,658 trade records.

Another reason we eliminate trades from the dataset is that price or volume information appears erroneous. Since, as mentioned above, all trades are self-reported, data entry errors are possible even though FINRA monitors reported trades. We delete records with missing

trade prices. We also drop trade records with unreasonably large prices. To compute our definition of unreasonable large prices, we first calculated the maximum “risk free” price of each bond in the sample. The maximum “risk free” price during our time period is the maximum present value of future coupon and principal payments, discounted using the lowest treasury rate observed across all bonds and all days between July 1, 2002 and December 31, 2006. Across all the bonds in our sample, the maximum risk free price is \$214. To be conservative, we drop all bond trades that take place at price higher than \$220. We also eliminated 4,597 trades where the volume of a single trade was higher than 50% of the issue amount.¹ Finally, we eliminated trades where the volume was reported as less than \$1,000.

The last reason we drop records is trade timing issues. We drop any trade that occurs before its offering date or after its maturity date. We also drop any trade that was reportedly executed on a different day than it was reported. Finally, we drop all records that are reported to have occurred on SIFMA holidays. After these eliminations, we are left with 21,149,525 trades involving 30,643 CUSIPs.

The entire dataset of cleaned bonds is not necessarily useful however to evaluate the effect of TRACE. Our empirical strategy is based on comparing a bond’s trading behavior when it changes from non-disseminated to disseminated. Many bonds will be disseminated for their entire trading history. These include bonds that belong to a FINRA Phase that are issued after the beginning of the Phase date, and bonds that may be issued before a Phase begins but only trade after the dissemination change date for that Phase. There are also bonds that are always non-disseminated. These are bonds that may mature before the beginning of their Phase date as well as bonds that belong to a Phase but never trade after the Phase begins.

Table A.2 outlines the steps from Historical Cleaned Sample in Table A.1 to the Cleaned Phase sample, the sample of bonds which exist and have zero or non-disseminated trading before the start of a Phase and zero or disseminated trading after the start of a Phase. We

¹This may represent a data error in FISD issue size. For example, about 2600 of the eliminated records correspond to one company Alestra, which went through an exchange. FISD reports its issue size as \$83,000, but through press releases we determined it was at least \$400,000,000. Another example is Countrywide CCR.MQ.

begin with a list of all Phase 1, 2, 3A and 3B bonds. There are 26,955 bonds in this list, of which 20,595 exist in the Cleaned Historical TRACE Sample. They have 17,434,020 trades during our sample period. Thus, about two-thirds of the bonds in our Cleaned Historical TRACE Sample are in our Phase list, but this represents 82.4% of the trades.

We obtained Phases 2, 3A, and 3B from FINRA. FINRA provided us with a list of bonds that began being disseminated at the start of each of the three Phases. This list was provided in a non-electronic format where bonds were identified with ticker symbols. Unfortunately, many ticker symbols longer than six characters were truncated. This was a problem for firms with a four character company ticker symbol which also issued bonds with three character security tickers. In particular, many GMAC bonds were truncated. Since FINRA also provided us with coupon and maturity dates for each bond, we were able to hand-match many of the truncated ticker symbols, but not all. The list of Phase 1 bonds was not provided by FINRA, and we generated it ourselves given the criteria listed by TRACE for Phase 1 bonds. That is, in addition to existing before the beginning of Phase 1, bonds had to be investment grade and have an initial issue size of \$1 billion or greater. After determining the set of bonds meeting these criteria, we eliminated all bonds that are on the FINRA lists for Phases 2, 3A or 3B and the bond had to have a disseminated trade before the beginning of Phase 2.

In addition to the four Phases that correspond to the FINRA dissemination dates, FINRA also maintained two other lists of bonds, which we call the FINRA50 and the FINRA120. The FINRA50 represent 50 Non-Investment Grade (High-Yield) securities disseminated under Fixed Income Pricing System (FIPS2). This list of 50 bonds changes over time with bonds both entering and exiting. FINRA lists all of these bonds on their website and there were a total of 149 bonds that were in the FINRA50 at some point during its existence from July 1, 2002 until July 14, 2004. The FINRA120 list is a special set of 120 investment grade rated Baa/BBB that FINRA delayed Phase 2 dissemination for. Phase 2 dissemination started on March 3, 2003 for Phase 2 bonds, but started on April 14, 2003 for the FINRA120. This special sample was created so that FINRA could conduct a controlled experiment to study the effects of dissemination in Phase 2, contained in Goldstein, Hotchkiss, and Sirri (2007).

Table A.2: Steps from FINRA's Phase Listings to Cleaned Phase Sample

	CUSIPs (1)	Trades (2)
Cleaned Historical TRACE Sample (after Table A1)	30,643	21,149,525
Source: FINRA list of Phase 1-3B bonds	26,955	...
TRACE Bonds Never on FINRA phase listing		
Issued between 7/1/02 and 2/7/05 and always disseminated	2,135	1,537,412
Matured before 2/7/05 and never disseminated	1,054	61,300
Matured on or after 2/7/05 and never disseminated	654	68,011
Issued before 7/1/02, always disseminated, trades only after 3/3/2003	99	36,250
Issued on or after 2/7/05 and always disseminated	5,435	860,996
None of the above (disseminated and non-disseminated)	671	1,151,536
Bonds on both FINRA Phase list and Cleaned Historical TRACE Sample	20,595	17,434,020
Phase 1		
list of Phase 1 bonds*	450	4,539,063
bonds in FINRA50 at start of phase	4	32,231
bonds do not exist as of start of phase	60	685,047
bonds do not exist during the period 90 days before until 90 days after start of phase	33	285,253
bonds with non-disseminated trades after start of phase	10	35,034
Cleaned Phase 1 Sample	343	3,501,498
Phase 2		
FINRA's list of Phase 2 bonds	3,747	...
Phase 2 bonds in Cleaned Historical TRACE Sample	3,049	2,934,214
bonds in FINRA50 before or at start of phase	0	0
bonds do not exist as of start of phase	272	21,459
bonds do not exist during the period 90 days before until 90 days after start of phase	229	150,854
bonds with disseminated trades before start of phase	2	4,380
bonds with non-disseminated trades after start of phase	8	25,636
Cleaned Phase 2 Sample	2,538	2,731,885
Phase 3A		
FINRA's list of Phase 3A bonds	16,898	...
Phase 3A bonds in Cleaned Historical TRACE Sample	13,260	8,336,332
bonds in FINRA50 or FINRA120 before or at start of phase	78	603,109
bonds do not exist as of start of phase	983	168,549
bonds do not exist during the period 90 days before until 90 days after start of phase	1,075	259,894
bonds with disseminated trades before start of phase	36	330,722
bonds with non-disseminated trades after start of phase	1	244
Cleaned Phase 3A Sample	11,087	6,973,814
Phase 3B		
FINRA's list of Phase 3B bonds	5,780	...
Phase 3B bonds in Cleaned Historical TRACE Sample	3,678	1,362,059
bonds in FINRA50 or FINRA120 before or at start of phase	26	52,945
bonds do not exist as of start of phase	648	235,319
bonds do not exist during the period 90 days before until 90 days after start of phase	132	46,791
bonds with disseminated trades before start of phase	15	22,135
bonds with non-disseminated trades after start of phase	4	1,738
Cleaned Phase 3B Sample	2,853	1,003,131
Total Cleaned Phase 1-3B Sample	16,821	14,210,328

This table reports the match between the Cleaned Historical TRACE file and FINRA's Phase Listings. Not all bonds in the TRACE Historical Sample are classified in a FINRA Phase. Excluded bonds are those issued after 7/1/02 that are always disseminated and those that mature before 2/7/05 that are never disseminated. We construct the Phase 1 list by including all bonds with disseminated trades before Phase 2 that are not on the FINRA Phase 2, 3A, or 3B lists. The Phase 2, 3A, and 3B lists were obtained directly from FINRA. The FINRA50 and FINRA120 lists are from www.finra.org. Bonds in FINRA's Phase lists that are not in the Cleaned Historical TRACE Sample have either never traded during the sample period or have been eliminated due to cleaning process in Table A.1.

Table A.2 explains how we went from FINRA's list of Phase 2, 3A, and 3B bonds, and our list of Phase 1 bonds, to our cleaned Phase samples. For each Phase list, we only use bonds that exist in our Cleaned Historical TRACE Sample. Some bonds on the FINRA lists did not trade during our sample period and thus are not in the Historical TRACE sample. This is shown between lines 1 and 2 under Phases 2, 3A, and 3B.

We next eliminate any bonds that also exist in the FINRA50 or FINRA120 list. Following this, we eliminate bonds that do not exist (i.e., were not issued or matured) during the period 90 days before until 90 days after the start of the Phase. Finally, we dropped some bonds with data problems. There are a few bonds where FINRA report disseminated trades before the start of the Phase, or non-disseminated trades after start of Phase. After applying these steps for each Phase list, what remains is our cleaned sample by Phases. There are a total of 16,825 bonds in our total cleaned Phase sample representing 14,210,328 trades during our time period.

A.2 NAIC Data Appendix

The National Association of Insurance Companies (NAIC) dataset we use is from Mergent FISD available on WRDS. The NAIC requires insurance companies to self-report all securities transactions in their financial statements. There are 63,859 bonds with 1,933,095 reported transactions in the NAIC file over the period January 1, 2000 until December 31, 2006. Schedule D of the annual NAIC filings require insurers to report all bond transactions in one of three categories: bonds added to the portfolio during the calendar year and held through the end of the year, bonds deleted from the portfolio during the calendar year that were not added in the same year, and bonds added and deleted in the same calendar year. For each transaction, the database records the CUSIP, date, par value of the transaction, the actual value of the transaction, if it was an addition or deletion, and a field for the counterparty involved in the transaction. Prices are not reported but can be computed from the ratio of the value received in the transaction to the par value of the bonds in the transaction. Importantly, the names of the insurance companies involved in the transactions are excluded from the data.

To make the NAIC analysis comparable to the TRACE analysis, we first match our sample of NAIC bonds with the Cleaned Historical TRACE sample by CUSIP. The universe of bonds which insurance companies trade is much larger than that reported by FINRA. 45,902 NAIC bonds representing 804,685 reported transactions are not included in our Cleaned Historical TRACE sample and are eliminated. Table A.3 reports the number of transactions and CUSIPs eliminated by this step.

Next we eliminate reported transactions that are not connected to trades. The NAIC database contains all transactions involving insurance companies' bond portfolios, not only buy and sell transactions, but also other transactions such as bond calls and maturities. The type of transaction is coded in the counterparty field. We eliminate all transactions that change bond portfolio holdings that are not buys or sells. These include the following codes: CALL, CANCEL, CONVERT, EXCHANGE, ISSUE, MATURE, PUT, REDEEM, SINKING FUND, TAX-FREE EXCHANGE, TENDER, TRANSFER, PAYDOWN, and REPLACE.

There are two prevalent entries in the counterparty field comprising almost 15% of the cleaned database that required additional attention: DIRECT and VARIOUS. DIRECT may indicate a direct placement, similar to an underwriting, or it may indicate the name of a counterparty in an actual trade. VARIOUS is simply an ambiguous catch-all, where some records may be actual trades and some are not. To check whether DIRECT and VARIOUS represent actual trades, we matched these NAIC records to TRACE using the CUSIP, price, volume, and date of the transaction. For DIRECT, only about 3% of transactions match into the TRACE dataset, while for VARIOUS only about 1% of transactions match. Because of the problems identifying which of the DIRECT or VARIOUS transactions are actual trades, we eliminate them along with the other codes listed above that are not buys and sells. As shown in Table A.3, all such filters eliminate 290,998 reported transactions on 14,095 different bonds.

We eliminate a small number of trades with data issues, i.e. missing prices, negative prices, etc. We next eliminate trades with timing issues, i.e., trades that are executed before or on the bonds' offering or after or on the bonds' maturity date. A large fraction of NAIC transactions take place on the offering and maturity dates. We believe that this is because insurance

Table A.3: Steps from Historical NAIC File to Cleaned Historical Sample

	January 1, 2000 - December 31, 2006		July 1, 2002 - December 31, 2006	
	CUSIPs	Ungrouped Trades	CUSIPs	Ungrouped Trades
Original Source: NAIC Transactions File	63,859	1,933,095	50,968	1,341,471
Match NAIC Bonds with Cleaned Historical TRACE sample				
CUSIP not found in Cleaned Historical TRACE sample	46,060	805,483	33,645	533,092
Eliminate transactions which are not trades				
Non-trade indicated by counterparty field entry (calls, converts, etc.)	13,996	290,802	13,424	233,247
Eliminate trades with data issues				
Missing Price	407	879	407	879
Zero Price (or Zero Par Value)	154	286	153	285
Negative Price (or Negative Par Value Amount)	151	194	132	162
Price greater than 220	53	85	50	82
Trades with volume/issue amount \geq 50% & (ParValue or issue amount not equal to 0 or 1)	359	647	235	421
Trades with volume less than 1000 dollars	140	295	116	249
Eliminated trades with timing issues				
Trades executed on or before bond's offering date	7,371	112,413	5,187	73,700
Trades executed on or after bond's maturity date	925	1,585	925	1,585
Trades executed on weekend or SIFMA Holiday	7,308	26,539	5,404	16,619
Post July 2002 trades executed on days with no TRACE trades**	7	13	7	13
Cleaned Historical NAIC Sample	16,005	693,861	14,573	481,134
				394,678
				1,032,124

Filters are applied sequentially. The CUSIPs column gives total number of CUSIPs eliminated from the database by adding that row's filter. The trades column gives total number of observations eliminated by adding that row's filter. * Price cutoff of 220 based on computing the a bond's maturity, coupons remaining and lowest value of the treasury yield curve during our entire sample period and taking the maximum across bonds. That value of 214 is rounded to 220. **On June 11, 2004, the SEC declared a holiday when because President Reagan died. Grouping is done if the difference in Price is $\leq |0.0|$ and the day, counterparty, insurer type, and buy or sell are equal.

companies are large customers of bond offerings and purchase the bonds at this time. The NAIC rules require its members to list these purchases as a transaction since the bonds are added to their portfolio. Since these transactions are probably part of the underwriting, we do not include them as trades. If an insurance company holds the bond until its maturity, that transaction will also be recorded by NAIC. Finally, we also exclude transactions listed on bond holidays. These screens shown in Table A.3 are similar to those applied to the TRACE dataset in Table A.1.

After the screens and matching, there are 16,006 bonds and 693,862 reported transactions (which we believe to be buys and sells) in our “clean” NAIC sample. Importantly, the NAIC time period in Table A.3 is thirty months longer than the TRACE time period in Table A.1. When we restrict to the time period July 1, 2002 until December 31, 2006, there are 14,574 bonds and 481,135 transactions, as shown in the last three columns of Table A.3.

As mentioned above in Section VII, we believe that many trades in TRACE are disaggregated by the NAIC reporting process. When comparing the NAIC and TRACE databases, there are multiple NAIC transactions that match to a single TRACE trade using CUSIP, date, price and counterparty, but not volume. However, if we group NAIC transactions by CUSIP, date, price and counterparty into a single record with a combined volume, many of these grouped NAIC trades match to a corresponding single trade in TRACE.

There are two reasons that trades are disaggregated in NAIC. The first reason is how NAIC requires transactions to be reported on Schedule D of the annual NAIC filing. Insurers must separately report bonds purchased and sold in the same year from bonds purchased and held through the end of the year. This means if an insurance company purchases \$1 million par of a bond on January 1, 2001 and sells \$500,000 of this before December 31, 2001 and the remaining \$500,000 sometime in the following year, under NAIC reporting instructions, this single purchase would be split into two separate purchases of \$500,000 each, reported in two different sections of Schedule D. One \$500,000 purchase would be reported in the long-term purchase reporting section, and one \$500,000 purchase would be reported in the short-term holding section.

When the NAIC database is compiled, the above trade would appear as two purchases of the same bond occurring on the same day at the same price. In TRACE, however, the dealer who sold the bond would report this as one \$1 million trade. If we aggregate the volume of the NAIC trades that occur in the same bond, on the same day, at the same price, the NAIC transaction would match to the TRACE trade, as a single trade. It's worth noting that since the insurance company sold the bond holdings as two separate pieces of \$500,000 each on two separate days, two distinct sales of \$500,000 would be reported as two sales in both NAIC and TRACE.

A second reason for why a single trade may be reported as multiple trades is that distinct subsidiaries of an insurance company may book portions of a trade to their respective division, and each division makes its own statutory filings to the NAIC. This can occur, for instance, if part of a trade is allocated to the property and casualty group and another portion allocated to the life insurance group. In the NAIC database, this appears as two trades, while in TRACE, it appears as one trade.

We attempt to correct for these two reporting problems by grouping transactions that we believe correspond to the same trade. Any records that share the same date, CUSIP, counterparty, transaction type (buy or sell), and have prices within 1 cent of another are grouped and considered a single trade. We show this grouping in Table A.3. In the cleaned NAIC file, from January 1, 2000 to December 31, 2006, the number of trades reduces from 693,862 to 567,251.

As discussed in Section VII, grouping trades does not affect our NAIC volume analysis. However, the price standard deviation increases when we group trades. We, therefore, report the analysis of NAIC trades both with and without grouping.

To assign the bonds in NAIC to a FINRA Phase, we simply match the cleaned Phase list from TRACE used in Table A.2 to the sample of cleaned NAIC bonds. Table A.4 reports the number of NAIC CUSIPs, and both grouped and ungrouped trades in each Phase. Importantly, in Phase 1, we match 323 CUSIPs out of 343 TRACE Phase 1 CUSIPs.

Table A.4: Comparison of NAIC and TRACE Trading Activity 90 Days After Phase Start

	Phase 1 (1)	Phase 2 (2)	Phase 3A (3)	Phase 3B (4)
A. CUSIPs				
CUSIPs				
Phase CUSIPs in Cleaned NAIC Dataset	323	2,192	4,710	2,076
Phase CUSIPs in Cleaned TRACE Dataset	343	2,682	11,171	2,855
NAIC CUSIPs / TRACE CUSIPs	94.2%	81.7%	42.2%	72.7%
B. Volume and Trades				
Volume				
NAIC volume	15,260,658,392	15,072,130,358	14,470,216,388	1,625,677,320
TRACE volume	243,403,051,641	130,940,914,944	200,109,997,753	37,075,185,871
NAIC volume/TRACE volume	6.3%	11.5%	7.2%	4.4%
Trades				
Ungrouped NAIC trades	6,775	15,056	16,231	3,812
Grouped NAIC trades	5,409	12,236	13,056	3,083
TRACE trades	351,606	221,460	404,035	42,645
Ungrouped NAIC trades/TRACE trades	1.9%	6.8%	4.0%	8.9%
Grouped NAIC trades/TRACE trades	1.5%	5.5%	3.2%	7.2%
Trade Size				
NAIC Ungrouped Average Trade Size	2,252,496	1,001,071	891,517	426,463
NAIC Grouped Average Trade Size	2,821,346	1,231,786	1,108,319	527,304
TRACE Average Trade Size	692,261	591,262	495,279	869,391
NAIC Ungrouped Average/TRACE Average	3.3	1.7	1.8	0.5
NAIC Grouped Average/TRACE Average	4.1	2.1	2.2	0.6
C. Price Standard Deviation				
CUSIPs / Bond-Days used				
Ungrouped NAIC	253 / 1,333	392 / 837	464 / 933	109 / 164
Grouped NAIC	213 / 969	261 / 481	273 / 483	65 / 87
TRACE	340 / 17,087	2,130 / 40,713	6,342 / 70,094	1,129 / 8,786
Price Standard Deviation				
Ungrouped NAIC	0.26	0.26	0.39	0.04
Grouped NAIC	0.42	0.40	0.70	0.04
TRACE	0.88	0.78	0.68	0.45
Ungrouped NAIC Std. Dev./TRACE Std. Dev.	0.30	0.33	0.57	0.09
Grouped NAIC Std. Dev./TRACE Std. Dev.	0.48	0.51	1.02	0.09

This table reports on comparisons between the cleaned NAIC file and the Historical TRACE file for 90 calendar days after the Phase Start.

A.2.1 Market Share Analysis

Unlike TRACE, which does not identify the transacting parties, the NAIC database has two fields which identify the counterparty to the insurance company in an NAIC sell or buy trade. These are: NAME OF PURCHASER (in a sell trade by the insurance company) and NAME OF VENDOR (in a buy trade by the insurance company). NAIC does not identify the name of the insurance company involved in the transaction. This means only one side, the non-insurance company side, is identified for each trade. We use this information, to construct a COUNTERPARTY variable.

Importantly, the counterparty field is not always a trading partner. Insurance companies also use this field to identify transactions that are calls, maturities, conversions, etc. and these were excluded from our sample as described above. Moreover, since each insurance company self-reports the data, there are often name variations in the counterparty field. For example, “J P Morgan”, “J P Morgan & Co”, “J. P. Morgan”, “J. P. Morgan Securities” and “J. P. Morgan Securities, Inc.” are listed as counterparties. We could not classify some counterparty names such as “192” or “9-UNIVERSAL LIFE”, so we group these counterparties together with names that appear infrequently into a LEFTOVER category. We grouped by hand the counterparty names into 106 unique trading partners. These correspond to 105 actual trading partners (originally from 7,319 distinct counterparty names), and the LEFTOVER category (which represents 4,714 distinct counterparty names). Importantly, trades in the LEFTOVER category only represent 11.0% of trades and 9.0% of total volume in the Cleaned Historical TRACE dataset.

In addition, because of mergers, some trading partners, which appear to be listed under separate names, are really part of one entity. For example, Salomon Brothers was acquired by Citigroup in 1998, and in our counterparty fields, the trader is sometimes identified as Salomon Brothers and sometimes as Citigroup, even though they were the same entity for our sample period. We examined all merger and acquisition activity for our 106 counterparties and if a merger took place before January 1, 2000 we combine the trading activity under the successor company’s name.

The following lists counterparties in our dataset that were acquired before our sample period and the successor name:

1. DEAN WITTER was acquired by MORGAN STANLEY on February 05, 1997, so it is called MORGAN STANLEY.
2. On June 30, 1997, MONTGOMERY SECURITIES acquired NATIONSBANK, which was acquired by BANK OF AMERICA on September 30, 1998, so MONTGOMERY SECURITIES is called BANK OF AMERICA.
3. SALOMON BROTHERS was acquired by CITIGROUP in 1998, so it is called CITIGROUP.

If merger activity occurs during the sample period, we keep the successor name. For instance, BANK ONE CORP was acquired by JP MORGAN on July 01, 2004, so it is called JP MORGAN in our sample. There are 22 counterparties that merged during our sample period:

1. ADVEST was acquired by MERRILL LYNCH on December 02, 2005, so it is called MERRILL LYNCH.
2. ALLIANCE CAPITAL MANAGEMENT was acquired by SANFORD C. BERNSTEIN on October 02, 2000, so it is called ALLIANCE-BERNSTEIN.
3. AMSOUTH BANK was acquired by REGIONS FINANCIAL CORP on November 04, 2004, so it is called REGIONS FINANCIAL CORP.
4. AUTRANET INC was acquired by BNY on February 04, 2002, so it is called BNY.
5. BANK ONE CORP was acquired by JP MORGAN on July 01, 2004, so it is called JP MORGAN.
6. FIRST CHICAGO BANK was acquired by NATIONAL BANK OF DETROIT, which was acquired by BANK ONE CORP in April, 1998, so it is called JP MORGAN.

7. BONDS DIRECT was acquired by JEFFRIES AND CO on October 07, 2004, so it is called JEFFRIES AND CO.
8. CHASE was acquired by JP MORGAN on September 13, 2000, so it is called JP MORGAN.
9. CREDIT LYONNAIS was acquired by CREDIT AGRICOLE on March 13, 2003, so it is called CREDIT LYONNAIS-CREDIT AGRICOLE.
10. DAIN RAUSCHER was acquired by RBC on September 28, 2000, so it is called RBC.
11. DONALDSON LUFKIN JENRETTE was acquired by CREDIT SUISSE on November 03, 2000, so it is called CREDIT SUISSE.
12. FIRST UNION was acquired by WACHOVIA on September 01, 2001, so it is called WACHOVIA.
13. FLEETBOSTON FINANCIAL was acquired by BANK OF AMERICA on April 01, 2004, so it is called BANK OF AMERICA.
14. GRUNTAL was acquired by RYAN BECK & CO on April 23, 2002, so it is called GRUNTAL-RYAN BECK.
15. MORGAN KEEGAN was acquired by REGIONS FINANCIAL CORP on December 19, 2000, so it is called REGIONS FINANCIAL CORP.
16. PAIN WEBBER was acquired by UBS on November 03, 2000, so it is called UBS.
17. PRUDENTIAL SECURITIES was acquired by WACHOVIA on July 01, 2003, so it is called WACHOVIA.
18. SPEAR LEADS & KELLOGG was acquired by GOLDMAN SACHS on September 11, 2000, so it is called GOLDMAN SACHS.
19. STANDISH was acquired by MELLON on July 31, 2001, so it is called MELLON.

20. TUCKER ANTHONY was acquired by RBC on March 08, 2002, so it is called RBC.
21. US BANCORP spun off PIPER JAFFRAY on December 31, 2003, so it is called PIPER JAFFRAY as US BANCORP.
22. WASSERSTEIN & PERELLA was acquired by DRESDNER KLEINWORT on January 01, 2001, so it is called DRESDNER KLEINWORT.

When we consolidate counterparties, we have 617,745 trades conducted by 86 unique traders and 76,116 trades for the LEFTOVER category, for a total of 693,861 trades. The total par volume in the LEFTOVER category is 135,375,226,962, which represents 9.0% of the total NAIC trading activity.

Appendix B

Appendix to Chapter 3

B.1 Likelihood Calculation

B.1.1 Step 1

Let $\theta = (\alpha, \beta, \delta, \eta)$ represent the vector of the non-fracking parameters and let $\phi = (\sigma_\epsilon, \sigma_\nu)$ represent the vector of the variance parameters. I compute the pseudo-observation g_i from (Y_{it}, X_{it}) , conditional on θ as

$$\begin{aligned} g_i &= \frac{1}{N_i} \sum_{t=1}^{N_i} (\log Y_{it} - X_{it}\theta) \\ &= \frac{1}{N_i} \sum_{t=1}^{N_i} (g(Z_i) + \epsilon_i + \nu_{it}) \\ &= f(Z_i) + \epsilon_i + \frac{1}{N_i} \sum_{t=1}^{N_i} \nu_{it} \end{aligned}$$

g_i is the sum of the “true” effect of fracking and location on oil production and a normally distributed error with zero mean and variance $\sigma_\epsilon^2 + \frac{1}{N_i}\sigma_\nu^2$.

B.1.2 Step 2

Conditional on the pseudo-observations g_i , the likelihood of (Y_{it}, X_{it}) follows the standard formula for panel data with a random effect on each well. Let $\psi(\cdot \mid \mu, \sigma)$ denote the normal

likelihood with mean μ and standard deviation σ and let $e_{it} = \log Y_{it} - X_{it}\theta$. Finally, let bolded capital letters represent vectors of the time series of a variable. The likelihood of observing $(\mathbf{Y}_i, \mathbf{X}_i)$ conditional on the parameters (θ, ϕ) and the unobserved impact of fracking g_i is

$$\begin{aligned}
\mathcal{L}(\mathbf{Y}_i, \mathbf{X}_i \mid g_i, \theta, \phi) &= \int \psi(\epsilon_i \mid 0, \sigma_\epsilon) \prod_{t=1}^{T_i} \psi(e_{it} - g_i - \epsilon_i \mid 0, \sigma_\nu) d\epsilon_i \\
&= \exp \left(-\frac{1}{2} \left[\frac{1}{\sigma_\nu^2} \left(\sum_{t=1}^{T_i} (e_{it} - g_i)^2 - \frac{\sigma_\epsilon^2}{T_i \sigma_\epsilon^2 + \sigma_\nu^2} \left(\sum_{t=1}^{T_i} e_{it} - g_i \right)^2 \right) \right] \right) \\
&\quad \times \left(\left(T_i \frac{\sigma_\epsilon^2}{\sigma_\nu^2} + 1 \right) (2\pi\sigma_\nu^2)^{T_i} \right)^{-\frac{1}{2}} \\
&= \psi \left(g_i \mid \frac{1}{T_i} \sum_{t=1}^{T_i} e_{it}, \sigma_\epsilon^2 + \frac{1}{T_i} \sigma_\nu^2 \right) \\
&\quad \times \exp \left(\frac{1}{\sigma_\nu^2} \left(\left(\sum_{t=1}^{T_i} e_{it} \right)^2 \left(\frac{\sigma_\epsilon^2}{T_i \sigma_\epsilon^2 + \sigma_\nu^2} - \frac{1}{2T_i} \right) - \frac{1}{2} \sum_{t=1}^{T_i} e_{it}^2 \right) \right)^{-\frac{1}{2}} \\
&= \psi \left(g_i \mid \frac{1}{T_i} \sum_{t=1}^{T_i} e_{it}, \sigma_\epsilon^2 + \frac{1}{T_i} \sigma_\nu^2 \right) J(\mathbf{Y}_i, \mathbf{X}_i, T_i \mid \theta, \phi)
\end{aligned}$$

The first term in this final expression is simply a normal likelihood, evaluated at g_i , the effect of fracking and location for well i . The second term does not depend on g_i . Though g_i is unobserved, by the properties of GPR, the vector \mathbf{g} of g_i 's for all N wells is distributed multivariate normal with mean zero and variance $K(\mathbf{Z} \mid \gamma)$. Thus, I can integrate over the values of g_i to obtain the likelihood in terms of observable data and parameters. Let \mathbf{T} denote the vector of values of T_i , $\Sigma(\mathbf{T}, \phi)$ be a N by N matrix with $\sigma_\epsilon^2 + \frac{1}{T_i} \sigma_\nu^2$ in the i -th diagonal position and zeros elsewhere and let $\mu(\mathbf{Y}, \mathbf{X}, \mathbf{T}, \theta)$ be a vector with $\frac{1}{T_i} \sum_{t=1}^{T_i} e_{it}$ in the i -th

position. Then the full likelihood is:

$$\begin{aligned}
\mathcal{L}(\mathbf{Y}, \mathbf{X}, \mathbf{Z}) &= \int \psi(\mathbf{g} \mid \mathbf{0}, K(\mathbf{Z} \mid \gamma)) \prod_{i=1}^N \mathcal{L}(\mathbf{Y}_i, \mathbf{X}_i \mid g_i, \theta, \phi) dg_i \\
&= \left[\prod_{i=1}^N J(\mathbf{Y}_i, \mathbf{X}_i, T_i \mid \theta, \phi) \right] \int \psi(\mathbf{g} \mid \mathbf{0}, K(\mathbf{Z} \mid \gamma)) \psi(\mathbf{g} \mid \mu(\mathbf{Y}, \mathbf{X}, \mathbf{T}, \theta), \Sigma(\mathbf{T}, \phi)) d\mathbf{g} \\
&= \left[\prod_{i=1}^N J(\mathbf{Y}_i, \mathbf{X}_i, T_i \mid \theta, \phi) \right] \psi(\mu(\mathbf{Y}, \mathbf{X}, \mathbf{T}, \theta) \mid \mathbf{0}, \Sigma(\mathbf{T}, \phi) + K(\mathbf{Z} \mid \gamma))
\end{aligned}$$

where the last line comes as a result of equations A.7 and A.8 from Rasmussen and Williams (2005). Having integrated out the unobserved values g_i , the full likelihood is completely in terms of the observed data $(\mathbf{Y}, \mathbf{X}, \mathbf{T})$, the parameter vectors θ and ϕ , and the covariance matrix $K(\mathbf{Z} \mid \gamma)$ of the nonparametric effect of fracking and location on oil production.

B.2 Expected Present Discounted Value of Oil Production

I compute *ex post* expectations for all wells, and I compute *ex ante* expectations for wells fracked by firms with sufficiently large information sets. I require that a firm's information set has at least 50 wells and at least 300 well-months of production. This limits the set of wells I can analyze, and the earliest wells with information sets this large do not appear until the fourth quarter of 2007.

I compute $\mathbb{E}[DOP_{ij}]$ using both expectation operators for a 10 by 10 grid of possible frack designs j , with sand use between 0 and 650 lbs per foot and water use between 0 and 750 gals per foot. These grid points cover 95% of observed sand choices and 99% of observed water choices. By the normality assumptions in the production function model, the joint distribution of log-production for well i under fracking design j over T months of existence (call this $\log \tilde{Y}_{ij}$) is multivariate normal, with mean μ_{ij} and covariance Σ_{ij} given by:

$$\begin{aligned}
\mu_{ij} &= \tilde{\mathbf{X}}_i \theta + \tilde{g}(Z_{ij}) \\
\Sigma_{ij} &= \tilde{\mathbf{X}}_i \Sigma^\theta \tilde{\mathbf{X}}_i^\top + (\sigma_\epsilon^2 + \sigma_{g,ij}^2) \mathbf{1}_T + \sigma_\nu^2 \mathbf{I}_T
\end{aligned}$$

where $\tilde{\mathbf{X}}_i$ is a matrix of well i 's static characteristics and a vector of log-age values from 1

month to T months, $\tilde{g}(\cdot)$ is the estimated GPR, Z_{ij} is the a vector of design (S_j, W_j) and latitude and longitude for well i , Σ^θ is the covariance matrix for the estimates of θ , $\sigma_{g,ij}^2$ is the estimated variance of the GPR at Z_{ij} , $\mathbf{1}_T$ is a T by T matrix of ones, and \mathbf{I}_T is a T by T identity matrix. With this construction, I am assuming that the variances for ϵ and ν are estimated perfectly (i.e., there is no term in Σ_{ij} that accounts for variance in those estimates).

Because $\log \tilde{Y}_{ij}$ is multivariate normal, the distribution of the *level* of production over time, \tilde{Y}_{ij} , is multivariate log-normal with the same parameters. The mean vector and covariance matrix of this distribution are:

$$\begin{aligned}\tilde{\mu}_{ij} &= \exp\left(\mu_{ij} + \frac{1}{2}\mathcal{D}(\Sigma_{ij})\right) \\ \left[\tilde{\Sigma}_{ij}\right]_{kl} &= \exp\left([\mu_{ij}]_k + [\mu_{ij}]_l + \frac{1}{2}([\Sigma_{ij}]_{kk} + [\Sigma_{ij}]_{ll})\right) (\exp([\Sigma_{ij}]_{kl}) - 1)\end{aligned}$$

where $\mathcal{D}(\cdot)$ represents the diagonal vector of a square matrix and $[M]_{xy}$ is the (x, y) -th entry of a matrix M .¹ Finally, $\mathbb{E}[DOP_{ij}]$ is:

$$\mathbb{E}[DOP_{ij}] = \sum_{t=1}^T \rho^t \tilde{\mu}_{ijt}$$

A similar calculation is available for the variance of present discounted oil production:

$$\begin{aligned}\mathbb{V}[DOP_{ij}] &= \mathbb{V}\left[\sum_{t=1}^T \rho^t \tilde{Y}_{ijt}\right] \\ &= \sum_{t_1=1}^T \sum_{t_2=1}^T \rho^{t_1+t_2} \left[\tilde{\Sigma}_{ij}\right]_{t_1, t_2}\end{aligned}$$

¹By the properties of the log-normal distribution, the mean and standard deviation of production are closely related, with the standard deviation equal to the mean times the exponent of the variance minus 1. This means that the ‘‘correlation’’ between the mean and standard deviation of production, computed across designs j will be positive by construction.

B.3 Weighted Gaussian Process Estimates

Recall that the mean and variance of the Gaussian process estimates of f at the point \tilde{Z} are given by:

$$\begin{aligned}\mathbb{E} \left[f(\tilde{Z}) \mid g, \mathbf{Z}, \gamma \right] &= k(\tilde{Z} \mid \gamma)^\top K(\gamma)^{-1} g \\ \mathbb{V} \left[f(\tilde{Z}) \mid g, \mathbf{Z}, \gamma \right] &= k(\tilde{Z} \mid \gamma)^\top K(\gamma)^{-1} k(\tilde{Z} \mid \gamma)\end{aligned}$$

where $k(\tilde{Z} \mid \gamma) = (k(Z_1, \tilde{Z} \mid \gamma) \dots k(Z_N, \tilde{Z} \mid \gamma))^\top$, $K(\gamma)$ is the matrix of pairwise kernel distances for each point in \mathbf{Z} and $g = (g_1 \dots g_N)^\top$. To compute a *weighted* mean and variance, I introduce a weighting matrix function, $L(\lambda)$, and compute a weighted estimate of the mean and variance:

$$\begin{aligned}\mathbb{E} \left[f(\tilde{Z}) \mid g, \mathbf{Z}, \gamma, \lambda \right] &= k(\tilde{Z} \mid \gamma)^\top L(\lambda)^\top K(\gamma)^{-1} g \\ \mathbb{V} \left[f(\tilde{Z}) \mid g, \mathbf{Z}, \gamma, \lambda \right] &= k(\tilde{Z} \mid \gamma)^\top L(\lambda)^\top K(\gamma)^{-1} L(\lambda) k(\tilde{Z} \mid \gamma)\end{aligned}$$

The weighting matrix function $L(\lambda)$ biases these estimates towards a firm's own experiences when λ is closer to 0 and towards other firms' experiences when λ is closer to 1. In particular, if $(k_0(\gamma), K_0(\gamma), g_0)$ are the subsets of $k(\gamma), K(\gamma), g$ computed using only the firm's own wells, and $(k_1(\gamma), K_1(\gamma), g_1)$ are the subsets computed using only other firms' wells, then the weighted estimates satisfy 3 relationships:

1. At $\lambda = 0$, the weighted estimates are equal to the estimates computed using the subset of wells the firm operated:

$$\begin{aligned}k(\tilde{Z} \mid \gamma)^\top L(0)^\top K(\gamma)^{-1} g &= k_0(\tilde{Z} \mid \gamma)^\top K_0(\gamma)^{-1} g_0 \\ k(\tilde{Z} \mid \gamma)^\top L(0)^\top K(\gamma)^{-1} L(0) k(\tilde{Z} \mid \gamma) &= k_0(\tilde{Z} \mid \gamma)^\top K_0(\gamma)^{-1} k_0(\tilde{Z} \mid \gamma)\end{aligned}$$

2. At $\lambda = \frac{1}{2}$, the weighted estimates are equal to the unweighted estimates:

$$\begin{aligned}k(\tilde{Z} \mid \gamma)^\top L\left(\frac{1}{2}\right)^\top K(\gamma)^{-1} g &= k(\tilde{Z} \mid \gamma)^\top K(\gamma)^{-1} g \\ k(\tilde{Z} \mid \gamma)^\top L\left(\frac{1}{2}\right)^\top K(\gamma)^{-1} L\left(\frac{1}{2}\right) k(\tilde{Z} \mid \gamma) &= k(\tilde{Z} \mid \gamma)^\top K(\gamma)^{-1} k(\tilde{Z} \mid \gamma)\end{aligned}$$

3. At $\lambda = 1$, the weighted estimates are equal to the estimates computed using the subset of wells the firm did not operate:

$$k(\tilde{Z} | \gamma)^\top L(1)^\top K(\gamma)^{-1} g = k_1(\tilde{Z} | \gamma)^\top K_1(\gamma)^{-1} g_0$$

$$k(\tilde{Z} | \gamma)^\top L(1)^\top K(\gamma)^{-1} L(1) k(\tilde{Z} | \gamma) = k_1(\tilde{Z} | \gamma)^\top K_1(\gamma)^{-1} k_1(\tilde{Z} | \gamma)$$

At intermediate values of λ , $L(\lambda)$ interpolates between these extremes. To accomplish this, $L(\lambda)$ takes this form:

$$L(\lambda) = \begin{bmatrix} L_1(\lambda) \mathbf{I}_{n_0} & L_2(\lambda) K_{01}(\gamma) K_{11}(\gamma)^{-1} \\ L_3(\lambda) K_{10}(\gamma) K_{00}(\gamma)^{-1} & L_4(\lambda) \mathbf{I}_{n_1} \end{bmatrix}$$

where n_0 is the number of wells in the firm's information set that it operated, n_1 is the number of wells that other firms operated, the matrices $K_{00}(\gamma)$, $K_{01}(\gamma)$, $K_{10}(\gamma)$, $K_{11}(\gamma)$ are submatrices of $K(\gamma)$:

$$K(\gamma) = \begin{bmatrix} K_{00}(\gamma) & K_{01}(\gamma) \\ K_{10}(\gamma) & K_{11}(\gamma) \end{bmatrix}$$

and the functions L_1, L_2, L_3, L_4 are

$$L_1(\lambda) = 1 + \lambda - 2\lambda^2$$

$$L_2(\lambda) = -\lambda + 2\lambda^2$$

$$L_3(\lambda) = 1 - 3\lambda + 2\lambda^2$$

$$L_4(\lambda) = 3\lambda - 2\lambda^2$$

Thus, $L(\lambda)$ is a quadratic interpolation between $L(0)$, which selects out the firm's own wells, and $L(1)$, which selects out all other firms' wells.

B.4 Geology Covariates

In the production function defined in Section 3.3, the only spatially varying observable characteristics are the well's location and the fracking choices its operator makes. However, the

North Dakota Geological Survey (NDGS) has published maps of potentially relevant geological information. In this appendix, I describe this data and evaluate its ability to explain oil production. The geology data explains a small, but statistically significant amount of variation in production, even after conditioning on a well's location. However, compared to production function models with location fixed effects, the explanatory power of geology data is small and the coefficients do not always have the signs that would be predicted by geology theory.

B.4.1 Available Data

The quantity of oil that a well draws from depends broadly on three geological factors: the thickness of the upper and lower Bakken shales, their total organic content, and their thermal maturity. These three factors describe the quantity of rock in the formation, the fraction of the rock that can generate oil, and the likelihood that oil generation has occurred, respectively. Fortunately, in 2008, the North Dakota Geological Survey (NDGS) published maps and GIS shape files documenting the spatial variation in these characteristics over the area covered by the wells in this paper.² I summarize this data in Table B.1.

As noted in Section 3.2, thicker locations in the Bakken have the potential to contain more oil. Using data from NDGS map GI-59, Panel A of Table B.1 shows the mean, standard deviation and within-township standard deviations of the thickness of the upper, middle, and lower Bakken members across the wells in this paper. The overall Bakken formation averages 86 feet thick, about half of which is the middle member. There is large variation in each of the thickness measures across wells, with the coefficient of variation ranging from 23-40%. However, within a township, the standard deviations of thickness measures are only 22-31% of the overall standard deviations.

In the upper and lower shales, oil can be generated from the fraction of mass that is organic (i.e., containing mostly carbon and hydrogen). All else equal, shale that has a higher organic content has the ability to generate more oil than shale with less organic content. Using data

²These maps are freely available in PDF format at <https://www.dmr.nd.gov/ndgs/bakken/bakkenthree.asp>. The shape files are available for purchase from the NDGS.

Table B.1: *Geology Covariates Summary Statistics*

Variable	Mean	Std. Dev	Min	Max	Within Std. Dev
Panel A: Thickness (ft)					
Bakken Formation	86.05	24.03	5.00	155.00	5.26
Upper Shale	16.50	3.74	1.00	31.00	1.16
Middle Member	42.39	13.10	2.50	82.50	2.86
Lower Shale	27.64	10.93	2.50	57.50	2.90
Panel B: Total Organic Content (%)					
Upper Shale	13.80	2.44	3.00	27.00	1.01
Lower Shale	14.01	2.32	8.50	22.50	1.04
Panel C: Thermal Maturity - Hydrogen Index					
Upper Shale	358.10	179.30	75.00	775.00	32.96
Lower Shale	343.27	181.90	25.00	1125.00	65.72
Panel D: Thermal Maturity - S2-TMAX (degrees celsius)					
Upper Shale	435.68	5.60	417.50	447.50	1.92
Lower Shale	433.89	10.12	387.50	447.50	2.81

$N = 2,699$. Reported Bakken Formation thickness does not exactly add up to the sum of the thickness of the three members in the data. “Within Std. Dev” is the standard deviation of the data after subtracting mean values within townships. Source: NDGS Maps GI-59 and GI-63.

from NDGS map GI-63, Panel B of Table B.1 shows the distribution of organic content in the upper and lower shales. In the average well, approximately 14% of the mass is organic in both members. There is limited variation in organic content, overall and within a township. Though not shown in the table, 99% of wells have 9% or more organic content in the upper shale and, 99% have 9.5% or more in the lower shale. For comparison, the organic content in the Ghawar Field of Saudi Arabia, the most prolific oil field in history, is only 5%.³

Long term exposure to high temperatures converts organic material into oil. The extent of exposure is called thermal maturity, and geologists use three categories to describe the thermal maturity of a rock sample. Thermally immature rock has less exposure than is necessary for the conversion of organic material into oil. Thermally mature rock has enough exposure for the conversion of its organic content into oil. Thermally over-mature rock has too much exposure, and its organic content is converted into natural gas.

In map series GI-63, the NDGS provides two measures of the thermal maturity of the Bakken: hydrogen index and S2-TMAX. Both measures are collected by heating a rock sample to high temperatures and measuring the rate of oil expulsion across temperatures. The maximum rate at which oil is expelled, divided by organic content, gives the hydrogen index. Since hydrogen is one of the two elements contained in all hydrocarbons, more hydrogen indicates higher hydrocarbon generating potential. Potential oil production is higher for larger values of the hydrogen index, with thermally mature rock at values as low as 200.⁴ The temperature of the highest rate of oil expulsion, called S2-TMAX, is the other laboratory measure of thermal maturity. Thermally mature rock corresponds to S2-TMAX values between 435 and 460, with higher values in that range corresponding to higher oil production. Above 460 degrees celsius, oil production is decreasing, and the rock is thermally over-mature.⁵

Panel C of Table B.1 shows the distribution of the hydrogen indices across wells. The average well has a hydrogen index suggestive of thermal maturity for both the upper and lower

³See Fox and Ahlbrandt (2002)

⁴For more information, see McCarthy *et al.* (2011)

⁵For more information, see McCarthy *et al.* (2011)

shales, though approximately 25% of wells are thermally immature. Within a township, the standard deviations of the hydrogen indices are 18-30% of the overall standard deviations. Panel D shows the distributions of S2-TMAX. The average well is just at the start of thermal maturity and no wells are thermally over-mature. Over 80% of wells have S2-TMAX in the range of thermal maturity in the upper shale, and 53% in the lower shale. Within a township, the standard deviations of the S2-TMAX values are 29-34% of the overall standard deviations.

The NDGS developed these maps using the cuttings, cores and well logs that operators are legally required to submit for every well they drill to the NDIC.⁶ Since the NDIC makes these samples and logs available to anyone, the information content in these maps may have been known by market participants before they were published.

Opportunities to measure the thickness, total organic content or thermal maturity of the rock in a specific well are infrequent.⁷ Furthermore, only in the last few years have geologists began to study the use of these cuttings in providing information about well quality.⁸ Even if these techniques had been available (and in widespread use), they would only provide information about the middle Bakken member, as that is the predominant source rock for cuttings.

B.4.2 Explanatory Power

To evaluate the ability of these geology covariates to explain oil production, I estimate Cobb-Douglas production function models with and without them. Table B.2 shows these results. Column 1 is a specification with no township fixed effects and no geology covariates (i.e., it is a simplification of the results in Column 2 of Table 3.6). In Column 2, I add the township fixed effects, increasing the between R-squared from 0.600 to 0.813, suggesting that location-specific

⁶Recall that “cuttings” are the returned rock samples generated during the drilling process. Occasionally operators also preserve contiguous sections of undrilled rock, called “cores”. By North Dakota Century Code 38-08-04, Section 43-02-03-38.1, operators are required to send physical samples of cuttings and cores to the NDGS within 90 days of collection, where they can be publically observed and analyzed by anyone. Additionally, operators are required to submit copies of all well logs and geology tests they perform.

⁷For example, Pimmel and Claypool (2001) notes that “rock eval pyrolysis is not normally used to make real-time drilling decisions because of the lengthy sample preparation, running, and interpretation time.”

⁸See, for example, Ortega *et al.* (2012)

factors explain a large portion of variation in production. Next, column 3 shows a specification with the geology covariates but no township fixed effects. Compared to Column 1, the between R-squared increases by 0.083 to 0.683. The coefficients on 6 of the 8 geology covariates are significantly different from zero and a Wald test rejects the hypothesis that the coefficients on the geology covariates are jointly equal to 0 at the 1% level. However, after conditioning on location, the geology covariates have considerably less explanatory power. Column 4 shows a specification with both township fixed effects and geology covariates. The increase in R-squared values from Column 2 to Column 4 is only 0.003, and only 2 of the 8 coefficients on the geology covariates are statistically significant. Again, a Wald test rejects the hypothesis that the geology covariates are jointly equal to zero. These results show that geology covariates do explain some of variation in oil production, but very little compared to the location fixed effects.

Geology theory predicts that the coefficients on each of these covariates should be positive, as greater thickness, organic content and thermal maturity are all thought to be associated with higher oil production. However, the coefficient on organic content in the upper Bakken shale is negative and statistically significant in both specifications.

The inclusion of geology covariates does not meaningfully change the Cobb-Douglas estimates of the productivity of lateral length, sand or water, as the coefficients in columns 2 and 4 are nearly identical.

B.5 Stability of the Production Function Relationship

In order for firms to empirically learn the production function for fracking, the true relationship between oil production, fracking inputs and location must be stable over time. To verify whether the data is consistent with a stable production function, I examine the performance of wells in similar locations that are fracked with similar designs but in different time periods. If similar wells fracked in different time periods have different performance, on average, then it is possible that the production function is not stable over time.

To implement this test, I estimate two time varying production function specifications and

Table B.2: *Explanatory Power of Geology Covariates*

Coefficient	(1) Log Oil	(2) Log oil	(3) Log Oil	(4) Log Oil
β	-0.557 (0.00237)	-0.557 (0.00237)	-0.557 (0.00237)	-0.557 (0.00237)
δ	1.755 (0.00355)	1.754 (0.00354)	1.755 (0.00355)	1.754 (0.00354)
η	0.436 (0.0370)	0.798 (0.0373)	0.549 (0.0358)	0.795 (0.0374)
κ_S	0.233 (0.0187)	0.158 (0.0161)	0.200 (0.0172)	0.155 (0.0161)
κ_W	0.0521 (0.0200)	0.115 (0.0163)	0.106 (0.0182)	0.115 (0.0163)
κ_{TU}			0.0400 (0.00333)	0.0000848 (0.00744)
κ_{TL}			-0.00205 (0.00114)	-0.00218 (0.00299)
κ_{CU}			-0.0494 (0.00530)	-0.0350 (0.00915)
κ_{CL}			0.0531 (0.00524)	0.00718 (0.00960)
κ_{HU}			0.00129 (0.000115)	0.000441 (0.000268)
κ_{HL}			0.0000825 (0.0000991)	0.0000323 (0.000134)
κ_{SU}			0.00875 (0.00317)	0.0139 (0.00446)
κ_{SL}			0.0153 (0.00144)	0.00212 (0.00337)
Overall R^2	0.690	0.784	0.728	0.785
Between R^2	0.600	0.813	0.683	0.816
Within R^2	0.765	0.765	0.765	0.765
Township Fixed-effects		X		X

Standard errors in parentheses. GLS random effects estimates of the production function model:

$$\log Y_{it} = \alpha + \beta \log t + \delta \log D_{it} + \eta \log H_i + \kappa Z_i + \tau_i + \epsilon_i + \nu_{it}$$

Y_{it} is oil production for well i when it is t months old, D_{it} is the number of days producing, H_i is the horizontal length, and Z_i is the vector of log sand use S_i , log water use W_i , upper Bakken thickness (TU), lower Bakken thickness (TL), upper Bakken organic content (CU), lower Bakken organic content (LU), upper Bakken hydrogen index (HU), lower Bakken hydrogen index (HL), upper Bakken S2-TMAX (SU) and lower Bakken S2-TMAX (SL), and τ_i is a set of township fixed effects. “Between” R^2 is the R^2 for the average predicted log baseline production. “Within” R^2 is the R^2 for the predicted time series of production. Estimated off of all 2,699 wells and 91,783 well-months.

conduct Wald tests of the hypothesis that the time effects are jointly equal to zero. In the first specification, I assume that baseline production is Cobb-Douglas with time-varying coefficients. If the true production function is both Cobb-Douglas and stable, the coefficients should not vary over time. In the second specification, I assume that baseline production is the sum of a year fixed effect and a fixed effect for wells with similar input choices and locations. To do this, I form groups of wells that have the same deciles of sand and water use that are also in the same township. Thus, this specification allows for a non-parametric relationship between baseline production, location and inputs. If the true production function is not Cobb-Douglas, but still constant over time, the time fixed effects should equal zero.

Table B.3 shows the results of these tests. The specifications in columns 1-3 are Cobb-Douglas in lateral length, sand use and water use, with time fixed effects and time fixed effects interacted with the sand and water use coefficients.⁹ Column 1 shows estimates computed from the whole sample. Wells in the 2009 and 2010 cohorts are significantly more productive than wells in the earlier cohorts, and wells in the 2009 cohort are significantly less sensitive to water use than wells in earlier cohorts. Column 2 shows estimates computed from the set of wells that are in bins with 2 or more wells. In this specification, wells in 2009 and 2010 are also more productive than earlier wells, while wells in 2011 are less productive. Wells in 2010 and 2011 are less sensitive to sand use than earlier wells, and wells in 2011 are more sensitive to water use. Column 3 shows estimates computed from the set of wells that are in bins with 2 or more wells fracked in two or more years. Wells in 2010 are more productive than earlier wells, while wells in 2011 are less productive. In this specification, none of the interaction terms are significantly different from zero. In all 3 specifications, a Wald test of the hypothesis that the year effects and their interactions are jointly equal to zero is rejected at the 1% level. These parametric results suggest that if the true production technology is similar to Cobb-Douglas, then its parameters may not be constant over time.

Next, columns 4-6 show estimates for the non-parametric specification. Again, column 4 is

⁹Because there are only 124 wells fracked between 2005 and 2007, I include them in the 2008 cohort, and specify year dummies for the 2009, 2010 and 2011 cohorts.

Table B.3: *Stability of Production Function Estimates Over Time*

Coefficient	(1) Log Oil	(2) Log Oil	(3) Log Oil	(4) Log Oil	(5) Log Oil	(6) Log Oil
β	-0.557 (0.00237)	-0.557 (0.00297)	-0.573 (0.00405)	-0.557 (0.00237)	-0.557 (0.00297)	-0.573 (0.00405)
δ	1.754 (0.00354)	1.797 (0.00452)	1.846 (0.00633)	1.753 (0.00355)	1.796 (0.00453)	1.846 (0.00633)
η	0.761 (0.0399)	0.734 (0.0561)	0.708 (0.0728)	0.850 (0.0680)	0.847 (0.0683)	0.801 (0.0825)
κ_S	0.159 (0.0318)	0.177 (0.0390)	0.171 (0.0621)			
κ_W	0.150 (0.0386)	0.114 (0.0446)	0.107 (0.0690)			
κ_{09}	0.785 (0.218)	0.721 (0.282)	0.645 (0.354)	-0.0236 (0.0473)	-0.0224 (0.0475)	-0.0232 (0.0476)
κ_{10}	0.706 (0.263)	0.732 (0.343)	1.005 (0.412)	-0.124 (0.0589)	-0.121 (0.0592)	-0.119 (0.0597)
κ_{11}	0.0933 (0.227)	-0.892 (0.334)	-1.602 (0.490)	-0.161 (0.0645)	-0.155 (0.0648)	-0.153 (0.0657)
$\kappa_{S,09}$	0.0285 (0.0408)	0.0253 (0.0496)	-0.0186 (0.0644)			
$\kappa_{S,10}$	-0.0621 (0.0456)	-0.137 (0.0670)	-0.0851 (0.0910)			
$\kappa_{S,11}$	-0.0190 (0.0423)	-0.0235 (0.0776)	0.147 (0.110)			
$\kappa_{W,09}$	-0.182 (0.0518)	-0.177 (0.0690)	-0.121 (0.0920)			
$\kappa_{W,10}$	-0.0577 (0.0523)	0.000315 (0.0830)	-0.118 (0.114)			
$\kappa_{W,11}$	0.00842 (0.0469)	0.177 (0.0838)	0.116 (0.116)			
# Well-months	91,783	50,866	25,939	91,783	50,866	25,939
# Wells	2,699	1,399	708	2,699	1,399	708
Overall R^2	0.785	0.805	0.808	0.836	0.827	0.823
Between R^2	0.816	0.828	0.846	0.952	0.894	0.888
Within R^2	0.765	0.792	0.800	0.765	0.792	0.800
Fixed-Effects	Township	Township	Township	Bins	Bins	Bins
Sample	All	Bins 1	Bins 2	All	Bins 1	Bins 2

Standard errors in parentheses. GLS random effects estimates of the production function model:

$$\log Y_{it} = \alpha + \beta \log t + \delta \log D_{it} + \eta \log H_i + \kappa Z_i + \tau_i + \epsilon_i + \nu_{it}$$

Y_{it} is oil production for well i when it is t months old, D_{it} is the number of days producing, H_i is the horizontal length, and Z_i is the vector of log sand use S_i , log water use W_i , dummies for the 2009, 2010 and 2011 cohorts, and interactions between the dummies and log sand use and log water use. τ_i is a set of fixed effects for townships or bins. “Between” R^2 is the R^2 for the average predicted log baseline production. “Within” R^2 is the R^2 for the predicted time series of production. “Bins 1” is the sample of wells in bins with 2 or more wells, while “Bins 2” is wells in bins with 2 or more wells, fracked in 2 or more years.

estimated on the entire sample, column 5 is estimated on the sample of wells in bins with 2 or more wells, and column 6 is estimated on the sample of wells in bins with 2 or more wells fracked in 2 or more years. In these specifications, the later cohorts tend to be less productive than the earlier cohorts. Wells in the 2010 and 2011 cohorts are significantly less productive than wells in the 2008. However, a Wald test of the hypothesis that the year effects are jointly equal to zero cannot be rejected at the 5% level in any of the nonparametric specifications, providing some support to the idea that the production function is stable over time.

Since the true production function is unlikely to be spatially homogenous or monotonic in sand and water or, the non-parametric results here may be more relevant.