Title
Essays in Finance and Environmental Economics

Author
Becker, Thomas A.

Publication Date
2010

Peer reviewed|Thesis/dissertation

Essays in Finance and Environmental Economics

by

Thomas Anthony Becker

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Economics

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Maurice Obstfeld, Chair
Professor Barry Eichengreen
Professor Brad DeLong
Professor Brian Wright

Fall 2010

Essays in Finance and Environmental Economics

Abstract

Essays in Finance and Environmental Economics

by

Thomas Anthony Becker
Doctor of Philosophy in Economics

University of California, Berkeley

Professor Maurice Obstfeld, Chair

This dissertation is comprised of three chapters, each of which contributes to the field of behavioral finance. Two of the chapters focus on topics in environmental economics and the third on U.S. household finance. All of the papers analyze the incentives and behavior of individuals (or firms in the case of the first chapter) to provide insight into macro phenomena.

"The Effects of Carbon Markets on Equity Prices and Volatility" uses a firm-level dataset of carbon assets and liabilities between 2005 and 2007 to examine the impact of volatile carbon prices on equity prices. I find that the changes to firms' market capitalizations during a period of falling carbon prices are explained by the change in the net present value of emission permit holdings. Equity prices respond to changes in the mark-to-market value of firms' carbon permit shortfalls or surpluses and carbon price volatility increases the volatility of equity prices. I also document considerable delays in the equity and options markets responses to developments in carbon markets and attribute these lagged responses to information constraints and the novelty of carbon markets.

"Outstanding Debt and the Household Portfolio," co-written with my classmate Reza Shabani, alters a simple portfolio choice model to allow households to retire outstanding debt and realize a risk-free rate of return equal to the interest rate on that debt. Using the Survey of Consumer Finances we find that households with mortgage debt are 10 percent less likely to own stocks and 37 percent less likely to own bonds compared to similar households with no outstanding mortgage debt. To show that our results are not driven by irrational behavior amongst a subset of households, we construct two proxy variables for financial naivete. Finally we calculate the costs of non-optimal investment decisions in the presence of various forms of household debt including mortgages, home equity loans and credit card debt. We find that 26 percent of households should forego equity market participation on account of the high interest rates that they pay on their debt.

"Crude Drilling: An Analysis of Incentives and Behavior in the Oil Industry During the 1860s" explains why rates of oil extraction in the nascent oil industry far exceeded the profit maximizing levels predicted by the economic theory of non-renewable resources. The analysis combines historical narrative accounts with property sale and lease data and information on oil well owners to explain how individual incentives led to aggregate over-drilling. In particular, I focus on the incentives of under-capitalized wildcat drillers as an explanation for the excessive waste and under-investment that characterized the early oil market. I find that these poorer prospectors were incentivized to extract oil at higher than optimal rates because of the characteristics of their property lease contracts and the low-cost drilling technology they used to bore exploratory wells. Low barriers to entry in the early oil drilling business led to an influx of wildcat drillers into the nascent oil market and delayed the entry of additional well-capitalized drillers. The result was a market characterized by cyclical supply shocks, low levels of investment in storage and conservation, and corresponding price instability.

Professor Maurice Obstfeld
Dissertation Committee Chair

To my wife, Mari Becker

Contents

List of Figures

List of Tables

Acknowledgments

I recall thinking over my reasons for choosing to work toward a Ph.D during the summer before I started at the University of California, Berkeley in 2006. I had been interested in political economy and economic history during my undergraduate studies, but had applied to graduate school eager to shift my focus to international and finance-related topics. I wanted to answer questions in these fields that were novel, timely, and of interest to a broad audience of readers. I had also become increasingly interested in working on environmental topics and was hopeful that I could incorporate these issues into my dissertation.

Reflecting on my experience over the past four years, I feel extremely fortunate to have been able to achieve each of these research goals. But more meaningful has been the opportunity that this degree has afforded me to meet, receive feedback from, and collaborate with so many remarkable economists. This dissertation has benefitted enormously from their advice and input. First, I would like to thank my advisor, Maury Obstfeld, for his support, mentorship, and his help in facilitating and accommodating my relocation to Cambridge for two years. It has been a privilege to learn international economics from you. To my oral committee members, Barry Eichengreen, Brad DeLong, Pierre-Olivier Gourinchas, and Brian Wright, I thank you for your constructive feedback on the research that I had completed prior to my oral defense and for steering me away from a second project that would have been less fruitful than my subsequent research on carbon markets. Barry Eichengreen was especially helpful as the chair of my oral committee, an insightful professor of the history of economics and finance, and a writer of letters of reference on my behalf. I am also deeply indebted Adam Szeidl for teaching me finance theory, offering so many helpful comments on my research on household debt, and for also writing letters of references on my behalf. David Card, you were a great mentor when we worked together at Cornerstone Research, and your open door and words of advice at Berkeley were always appreciated. Michael Greenstone, thank you for informing me of and then helping me to write a strong fellowship application for the EPA STAR Fellowship. I am also indebted to Dan Bergstresser, Mihir Desai, and all of the members of the faculty of Harvard University that gave helpful feedback and accommodated me while I completed my dissertation research from Cambridge.

In addition to these extraordinary colleagues, I am also greatly indebted to a number of organizations and institutions that supported my studies. Foremost, I am grateful for the three years of research funding that I received from the Environmental Protection Agency STAR Fellowship. The fellowship proved to be much more than just a source of funding, as it connected me with a network of interdisciplinary environmental researchers and provided me with an unparalleled level of flexibility to conduct my research. I also feel incredibly fortunate to be earning my degree from the University of California at Berkeley. The institution's devotion to excellence in public education and its remarkable history provided me with a world-class education and was a constant source of inspiration. My first-year funding from Berkeley allowed me to focus my energy on coursework and the accessibility of research funds in subsequent years allowed me to develop a strong foundation for my research. Berkeley's Center for European Studies Pre-Dissertation Fellowship funded a summer travel grant to Germany that was instrumental in my decision to conduct research on European carbon markets. I also benefitted greatly from the visiting scholar program at Harvard University, which allowed me to complete my dissertation from Cambridge while also providing me with complete access to its seminars, libraries, and faculty.

Finally, I would like to thank my family and friends. This degree often proved to be as much a psychological challenge as an intellectual undertaking, and without your help and support I couldn't have done it. To Reza Shabani, who has been both a great friend and colleague, thank you for challenging me throughout graduate school and for being a fantastic collaborator. Thanks to Toby Hyde, for all of your help with the research for the carbon markets paper and your continued interest in my research. To my mom, dad, Anna, and Billy, I could have never gotten this far without your unwavering support and faith in my abilities and I love you all very much. And of course I give special thanks to my wife, Mari Becker, for your love and dedication that made this possible. You are a constant source of inspiration and my most trusted friend and I will always look back fondly on all of the time that we spent together during these four years.

Introduction

Upon first glance, there might appear to be few similarities between the chapters of this dissertation. They examine diverse topics that include the reaction of equity and derivatives markets to a 2006 decline in carbon prices, the financial decisions of U.S. households in the 1990s and early 2000s, and the behavior of oil prospectors in the 1860s. Each chapter employs a different set of analytical tools, ranging from the classification of historical narrative accounts to the complex econometric analysis of detailed household financial surveys, to answer a diverse set of questions. But what unites this dissertation is that all of my research seeks to gain insight into macro phenomena by analyzing the situations and incentives of individuals (or firms in the case of the first chapter). The papers can also all be broadly classified as making contributions to the field of behavioral finance with two of them focusing on environmental topics, namely oil and carbon markets.

This dissertation also traces my graduate school progress in reverse chronological order: the third chapter contains my oldest work and the first chapter my most recent. The third chapter, written largely during my first year of graduate school, marked the start of my research on environmental topics and my interest in looking at individuals' incentives to understand aggregate behavior. The second chapter, co-written with my classmate Reza Shabani between the second and third year of the program, expanded my econometric toolkit while also addressing an interesting and current topic in behavioral finance. The first chapter, which has been written over the course of the last two years, represents the culmination of my graduate studies and builds upon the tools and insights gained from my earlier research. It examines an issue at the intersection of financial and environmental economics and provides important new insights into the spillover effects of emissions trading into other financial markets.

Chapter 1, entitled "The Effects of Carbon Markets on Equity Prices and Volatility," uses a new dataset of firms regulated by the E.U. cap-and-trade system to analyze how changes in the value of firms' carbon balance sheet impacts equity and derivatives prices. The idea for this project originated while reading an *Economist* article about the prolonged period of carbon price volatility surrounding the first information release in the spring of 2006. Though I suspected that this event could hold promise for an interesting research project, it was only during my second year of graduate school that I began to formulate a research agenda around this topic. I began by constructing the dataset during my third year and have spent the better part of this past year completing the analysis.

The paper uses a firm-level dataset of carbon assets and liabilities between 2005 and 2007 to examine the impact of volatile carbon prices on equity prices. I find that the changes to firms' market capitalizations during a period of falling carbon prices are explained by the change in the net present value of emission permit holdings. Equity prices respond to changes in the mark-to-market value of firms' carbon permit shortfalls or surpluses and carbon price volatility increases the volatility of equity prices. I also document considerable delays in the equity and options markets responses to developments in carbon markets and attribute these lagged responses to information constraints and the novelty of carbon markets.

Chapter 2, entitled "Outstanding Debt and the Household Portfolio," examines the effect of household debt on investment decisions. This chapter was co-written with my classmate Reza Shabani and published in 2010 as an article in the *Review of Financial Studies*. The project arose out of our 2007 lunchtime discussions of the inadequacies of the existing explanations of transaction or information costs to explain the equity non-participation puzzle. We thought that a simpler and more plausible explanation for non-participation was the debt repayment option available to the majority of households with mortgage and consumer debt. Given the explosion of household debt and mortgage debt during the previous decade, we felt the project merited further investigation and began to examine household-level financial data. The result of this collaboration was a timely empirical analysis connecting the liabilities side of U.S. households' balance sheets with their asset market participation decisions and portfolio choices.

The paper combines a behavioral finance model with empirical methods from labor economics to derive its results. We alter a simple portfolio choice model to allow households to retire outstanding debt and realize a risk-free rate of return equal to the interest rate on that debt. Using the Survey of Consumer Finances we find that households with mortgage debt are 10 percent less likely to own stocks and 37 percent less likely to own bonds compared to similar households with no outstanding mortgage debt. To show that our results are not driven by irrational behavior amongst a subset of households, we construct two proxy variables for financial naivete. Finally we calculate the costs of non-optimal investment decisions in the presence of various forms of household debt including mortgages, home equity loans and credit card debt. We find that 26 percent of households should forego equity market participation on account of the high interest rates that they pay on their debt.

Chapter 3, entitled "Crude Drilling: An Analysis of Incentives and Behavior in the Oil Industry During the 1860s," examines the drilling behavior of oil prospectors in Western Pennsylvania during the first oil boom. The project was inspired by my wife's grandfather, Robert O. Anderson, who was a wildcat oilman in the second-half of the 20^{th} century. Interested in learning more about the motives and economic consequences of this type of prospecting, I decided to investigate the nations' first wildcatters to fulfil my second year economic history paper requirement. What began in the historical archives of the Berkeley library in 2007 culminated with a research trip to the oil region of Venango County, Pennsylvania this past summer. The project allowed me to synthesize a great deal of historical information concerning the history of the oil industry and to offer new economic insights into the sources of the volatility that characterized early oil markets.

This paper explains why rates of oil extraction in the nascent oil industry far exceeded the profit maximizing levels predicted by the economic theory of non-renewable resources. The analysis combines historical narrative accounts with property sale and lease data and information on oil well owners to explain how individual incentives led to aggregate over-drilling. In particular, I focus on the incentives of under-capitalized wildcat drillers as an explanation for the excessive waste and under-investment that characterized the early oil market. I find that poorer prospectors, whom I classify as wildcat drillers, were incentivized to extract oil at higher than optimal rates because of the characteristics of their property lease contracts and the low-cost drilling technology they used to bore exploratory wells. Low barriers to entry in the early oil drilling business led to an influx of wildcat drillers into the nascent oil market and delayed the entry of additional well-capitalized drillers. The result was a market characterized by cyclical supply shocks, low levels of investment in storage and conservation, and corresponding price instability.

Chapter 1

The Effects of Carbon Markets on Equity Prices and Volatility

Data about the actual carbon emission levels by companies began filtering out to the markets on April 25, more than two weeks ahead of the May 15 release date. Traders called the situation "chaotic."
–New York Times, May 16, 2006

Introduction

This paper examines the effects of carbon markets on the equity prices of regulated firms. The 1997 Kyoto Protocol, which was the first international agreement to coordinate the reduction of greenhouse gases, permitted signatories to choose the mechanism by which they reduced emissions. In 2005, the member countries of the European Union (E.U.) chose to fulfil their treaty obligations by launching the world's first market-based system to limit greenhouse gas emissions. Often referred to as a cap-and-trade system, the European Emission Trading System (ETS) places limits on the amount of carbon dioxide (CO_2) that can be emitted from industrial facilities throughout the E.U. by issuing carbon permits to regulated facilities. The system reduces emissions over the course of decades by gradually lowering the aggregate number of permits issued and allowing firms to purchase permits from one another. This paper uses the data from the first round of the E.U. system to develop an understanding for how carbon markets impact existing capital markets. In particular, it examines how a period of high carbon price volatility and a sharp decline in carbon prices affected the equity volatility and valuations of regulated firms.

Viewed from a corporate finance perspective, carbon markets are a regulatory change that require regulated firms to manage a new asset, carbon permits, along with the corresponding liability of future carbon emissions. Fluctuations in the price of carbon permits should therefore impact firm value based on the present value of their net carbon position. In order to test this theory, I use a new dataset that combines E.U. carbon registry data, which are collected by member countries on a facility-by-facility basis, with firm-level equity data. The registry data contain both the allocated emissions permits and the verified annual emissions for over 12,000 facilities in the E.U. between 2005 and 2007. By matching these carbon data to each owner's equity data, I am able to examine the impact of carbon prices on firm value. The data cover 347 public firms across a broad range of industries with a total 2005 market capitalization of over 5.5 trillion euros.

The results show that a period of high carbon price volatility in April and May of 2006 was associated with an approximately 8% increase in the market betas of equities regulated by the ETS. I also find that the decline in carbon permit prices during that period was associated with economically and statistically significant abnormal equity returns. The abnormal returns are explained by the net carbon positions, allocated permits less verified emissions, of individual firms. Firms with a net long (short) position experienced negative (positive) returns in response to lower carbon prices. Changes in firms' market capitalizations are correlated with the change in value of firms' permit holdings. The overall 14 euro decline in carbon permit prices was associated with an over 90 billion euro decline in the market capitalizations of regulated firms.

The analysis also finds that the valuation impact of carbon price fluctuations was not immediately incorporated into equity and options prices. Carbon prices fell steeply at the end of April, but equity prices did not fully respond until the middle of May when official emissions data were published. Options markets also had a delayed price response. The implied volatilities of equity call options show that forward-looking volatility estimates failed to incorporate the effects of carbon market volatility into options prices. The delayed response of equity and options markets to carbon price changes can be attributed to the newness of carbon markets and informational challenge of incorporating carbon positions into firms' equity valuations.

These findings have important implications for firms and investors. Greenhouse gases are emitted by industrial processes undertaken by every sector of the economy. A cap-and-trade system that regulates these emissions has a much broader impact on the economy than an emissions market focused only on power plant emissions.[1] Under a cap-and-trade system firms face both the long-run cost associated with emissions abatement as well as a short-term cost related to their exposure to carbon price variability. To the extent that firms seek to limit stock price fluctuations stemming from dislocations to carbon markets, they must use financial instruments to hedge their carbon exposure. For investors, the findings show that carbon permit prices can have material effects on equity prices. Though financial reporting requirements do not yet require firms to report their carbon permit holdings, these assets have large

[1] The first cap-and-trade system was created in 1990 by Title IV of the U.S. Clean Air Act. It capped the emission of sulfur dioxide and nitrous oxide from 110 coal-fired power plants to reduce the incidence of acid rain.

cash and earnings effects for many regulated firms. The finding that markets were slow to incorporate carbon price impacts into equity and options prices shows that investors did not anticipate the extent to which carbon markets can impact other capital markets.

The findings of this paper are related to the findings of Bartov and Bodnar (1994). They examine how changes to the value of firms' foreign exchange holdings can impact firm value and find that lagged changes in exchange rates explain abnormal returns for U.S. firms with large foreign exchange holdings. They attribute the delayed equity response to the complexity of modelling firms' exchange rate exposure and the short history of floating exchange rates in their dataset. My finding that equity and options markets were slow to react to developments in carbon markets is supported by models of gradual information diffusion and information capacity constraints developed by Hong and Stein (1999) and Peng (2004), and tested empirically by Kewei and Moskowitz (2005). This paper also adds to the extensive literature, documented in Kothari and Warner (2007), that uses event studies to uncover relationships between economic variables and equity returns.

My findings also contribute to the growing environmental economics literature. The use of facility-level data to uncover aggregate firm and industry effects is similar to the approach used by Greenstone (2002). His analysis of the Clean Air Act finds that environmental legislation can cause significant changes in firm output. Veith, Werner, and Zimmermann (2009) examine the equity price response of a portfolio of European power producers to fluctuations in carbon prices throughout the first phase of the ETS and find that the stock returns of utilities were positively correlated with carbon prices. Bushnell, Chong, and Mansur (2009) also examine the equity response of firms to the first release of carbon data in 2006. Their study examines the equity returns of the component companies of the Dow Jones EuroStoxx 600 index between April 26-28, 2006. They find that firms in industries with a high ratio of emissions permits to market capitalization experienced price declines following the first release of carbon data. My paper, by examining both the valuation and volatility impacts of carbon price changes for all of the firms regulated by the ETS, provides a more complete analysis of the interaction between carbon markets and equity prices.

The remainder of the paper is organized as follows. Section 1.1 contains an overview of the European cap-and-trade system, timeline of events during the event period, and a model for how carbon prices are incorporated into equity valuations. Section 1.2 provides an overview of the data. Section 1.3 presents the volatility and valuation impacts of the decline in carbon prices. Section 1.4 shows that changes in firms' market capitalizations reflected changes to the value of firms' total carbon exposure. Section 1.5 contains an analysis of the timing of the equity and options market reactions to developments in carbon prices. The final section concludes.

Figure 1

Price and Volume of European Carbon Futures Contracts

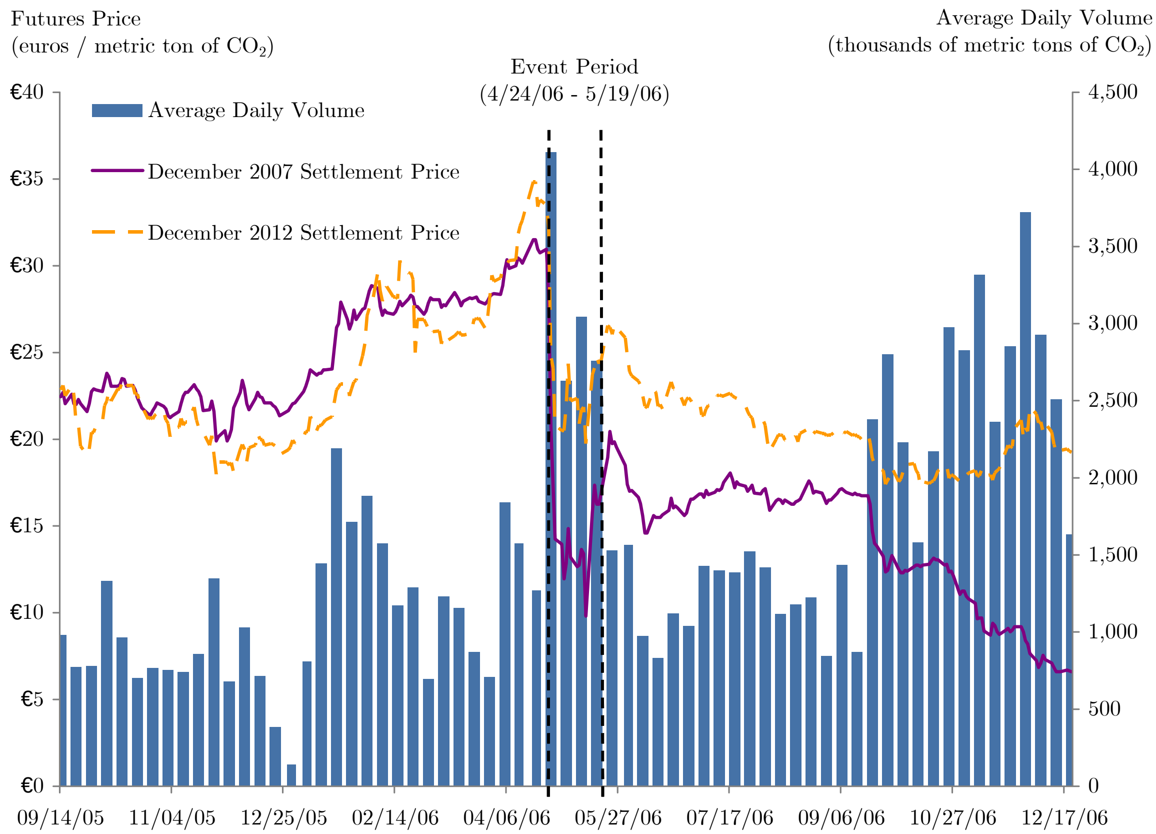

Data are from the European Climate Exchange (ECX). Futures prices for December 2007 and December 2012 are for carbon permits delivery in the final months of Phase I and Phase II of the European Union cap-and-trade system, respectively. The event period spans from one day prior to the first leak of emissions data to the end of the week after the official release of 2005 emissions data. Volumes are daily averages during each trading week.

1.1 Background and Motivation

1.1.1 Characteristics of the Market

The Emissions Trading System (ETS) was created by the European Union in 2003 as means of reducing greenhouse gas emissions amongst its member states. It is structured as a decentralized market whereby the 27 E.U. member countries allocate emissions permits across industries and amongst individual facilities under the coordination of the European Commission.[2] The ETS was set up in multiple phases with

[2]Kruger, Oates, and Pizer (2007), Buchner and Ellerman (2007b), and Ellermann and Joskow (2008) all provide detailed studies of the E.U. ETS and the permit allocation procedures for the initial

the initial three year phase beginning in January 2005.[3] Regulated facilities were allocated a total of 6.1 billion emissions permits to meet their emissions needs over the entire first phase of the system. Since the first phase of the ETS was structured as a trial phase, permits allocated for the 2005-2007 period were only valid for emissions during that three year period. Firms with emissions in excess of their allocation of permits had to purchase permits equal to their shortfall from firms with a surplus of permits.[4]

Between 2005 and the spring of 2006 carbon prices rose steadily on projections that the market had an aggregate permit shortfall of 1-2%.[5] By early April of 2006, futures prices for carbon permits had risen to over 30 euros per metric ton.[6] The first public release of 2005 emissions data was scheduled for May 15, 2006 on a centralized E.U. website known as the Community Independent Transaction Log (CITL).

Starting on April 25, 2006 some countries began to release their national data ahead of schedule. These data roiled carbon markets as investors tried to discern the overall state of the market from partial data. Carbon permit prices fell sharply and remained volatile during the subsequent weeks. Following the eventual release of E.U.-wide data on May 15, permit prices stabilized around 15 euros per metric ton. The drop in carbon prices reduced the value of Phase 1 permits by over 100 billion euros. Figure 1 shows the prices and volumes of carbon permit futures contracts in the months surrounding the release of 2005 emissions data.[7] The two futures prices are for

---

phase of the market. All facilities within the nine sectors designated by the European Commission– electric power, refineries, iron & steel, cement, paper & pulp, coke ovens, glass, and ceramics– participate in the system.

[3]The second phase, between 2008–2012, was scheduled to coincide with the timeline for emissions reductions agreed to by the EU member states under the Kyoto Protocol.

[4]The penalty for verified emissions in excess of surrendered permits, which establishes the *de facto* price ceiling for emissions in each phase, was 40 euros in the first phase of the ETS. The penalty price was scheduled to increase to 100 euros for the 2008-2012 phase of the ETS. In the second round of the ETS, firms would also be able to use certified emissions reductions (CERs) and Emissions Reduction Units (ERUs), which are credit for emission reductions in devloping countries, to offset a portion of their permit shortfall. CERs and ERUs were not part of the first phase of the ETS.

[5]Investment bank analyst reports written in early 2006 forecasted rising permit prices through 2008 based on an estimated permit shortfall and fuel-switching costs between coal and natural gas. A March 6, 2006 JPMorgan Chase analyst report discussed the shortfall as follows: "In early January we estimated, based on PointCarbon data, that the shortfall for 2005 would be around 40mt, with an aggregate phase 1 shortfall of c260mt. CERA estimates that the 2005 shortfall will be around 60mt, with an aggregate phase 1 shortfall of 190mt...We use these as the base lines for high- and low-shortfall scenarios."

[6]At 30 euros per metric ton, the notional value of the emissions permits allocated in the first phase of the ETS was 183 billion euros.

[7]The spot market for emissions permits is transacted primarily through the over the counter market. Futures prices from the largest futures exchange, the European Climate Exchange, provide a continuous and more transparent reflection of the value of emissions permits.

carbon permit delivery in the final month of the first two trading periods of the ETS.[8] Carbon prices experienced a large overall decline coupled with high volatility during this period. Average traded volumes were more than three times higher during the event period than during the preceding months. This analysis focuses on the volatile weeks between the initial releases of emissions data and the end of the week after the official European Commission release of emissions data.

Table 1 contains a timeline of news releases concerning carbon markets during the event period. On April 25, Holland and the Czech Republic released their reported surpluses of 17% and 15%, respectively. Markets were not only surprised by the early release of data but also by the large discrepancies between allocated permits and actual emissions. The volume of traded permit futures jumped and prices dropped. The following day, and partially in response to the Dutch and Czech release, France, Walloon Belgium, and Estonia also released their data, which all showed substantial surplus allocations. Prices remained volatile throughout the following weeks, responding to additional emissions data and continued uncertainty. Though the eight countries that released their emissions data prior to the May 15 release data comprised only 22% of the allocated permits for 2005, all reported substantial surpluses, which cast doubt on the whether Phase 1 permits would retain any scarcity value.

On May 12, the Friday before the scheduled release of the 2005 data, the research firm Point Carbon cited a malfunction of the CITL website and reported that the system displayed emissions data indicating that there was a 3.4% surplus of permits across the E.U. in 2005. Carbon prices plunged on this news and analysts began to forecast permit prices dropping to the low single digits. The verified European Commission data that was released on the following Monday showed an aggregate surplus of permits for 2005, but the 2.5% surplus was lower than many had forecasted and than the Point Carbon leak had indicated.[9] Prices rebounded and in the subsequent months remained relatively stable around 15 euros per metric ton.

The spike in volatility and decline of carbon prices during these weeks form the basis for the subsequent analysis. The period does not contain any other large news affecting commodity prices. Oil and natural gas prices, both of which can impact the industries that are regulated under the ETS, exhibited similar behavior during the event period as compared to the surrounding months. The volatility of electricity

[8] Differences in prices for the 2007 and 2012 futures contracts reflect the different penalty prices for non-delivery and the distinct allocation procedures for each round of the ETS. The two futures prices began to diverge during the event period because the emissions data from the first year data cast doubt on whether the first phase of the ETS would have an overall shortfall of permits. Individual member countries also reacted to the 2005 surpluses with pledges to reduce their permit allocations in subsequent phases of the system.

[9] Poland, Cyprus, Luxembourg and Malta did not report their emissions because of administrative delays and were not included in either the Point Carbon news release or the May 15 CITL data. Some residual uncertainty over the overall surplus for 2005 remained because of an anticipated surplus in Poland. The total 2005 permit surplus was 3.8%, which became known in June when Poland released its emissions data.

Table 1

Timeline of News Concerning European Carbon Markets

Date	News	Daily CO_2 Price Change	December 2007 Futures Prices (euros / metric ton of CO_2)	Daily CO_2 Market Volume[a] (metric tons of CO_2)	Cumulative % of 2005 Permit Allocations Reported
4/25/2006	Holland reports a 7% surplus and the Czech Republic reports a 15% surplus of 2005 permits.	-8.6%	28.30	2,702,000	8.9%
4/26/2006	France reports a 12% surplus, Belgium (Walloon region) a 15% surplus, and Estonia a 32% surplus of 2005 permits.	-26.7%	20.75	5,253,000	20.1%
4/27/2006[b]	Global Insight: "We've been taken by surprise." JPMorgan: "It may be, though, that only the environmental 'winners' not the environmental 'sinners' that have reported so far."	-16.6%	17.30	7,477,000	
4/28/2006	Rumors that Germany and U.K. may publish emissions data earlier than May 15 deadline.	-17.6%	14.25	4,224,000	
4/30/2006	European Commission deadline to surrender 2005 allowances.				
5/2/2006	Sweden reports a 11% surplus of 2005 permits.	-14.3%	11.95	3,386,000	21.2%
5/4/2006	Lithuania reports a 46% surplus of 2005 permits.	15.6%	14.85	4,226,000	21.7%
5/5/2006	Slovenia reports a 5% surplus of 2005 permits.	-11.1%	13.20	2,764,000	22.2%
5/12/2006[c]	Point Carbon cites a malfunction on a European Commission website that shows an aggregate surplus of 69.95 MMT, or 3.8%, in 2005. Fortis forecasts a surplus of 75 MMT and that "prices next week may sink to the region of 5 euros."	-26.9%	9.80	3,031,000	
5/15/2006	European Commission reports a European-wide surplus of 44 MMT, or 2.5%, in 2005.	60.7%	15.75	3,856,000	87.3%

News releases are compiled from a search of Bloomberg News articles with Environmental Credits tags.
a) Traded volumes, which are from the European Climate Exchange (ECX) and do not include OTC trades of carbon products.
b) Both quotations from the 4/27/2006 Bloomberg News article "Emissions Prices Head for Biggest Three-Day Drop."
c) Estimates and quotation from the 5/12/2006 Bloomberg News article "EU Granted Surplus of Carbon Permits in 2005, Fortis Says."

futures prices did increase during the event period in response to the changes in carbon futures prices.[10] However, both the volatility and change in electricity futures was minor when compared with carbon prices; electricity futures experienced an overall 6% price decline during the event period.[11]

1.1.2 Model for How Carbon Permits Impact Firm Valuation

The theoretical justification for creating the ETS as a mechanism to limit greenhouse gas emissions is based on the premise that the marginal cost of reducing emissions varies across industries and on a facility-by-facility basis. An emissions market, by assigning a scarcity value to emissions permits issued in lower numbers than the business-as-usual emissions trajectory, will incentivize abatement amongst the regulated firms. The market price for permits will be set equal to the marginal cost of emissions abatement at the firm (facility) with the lowest cost of abatement across all regulated firms (facilities). Newell and Stavins (2003) show that a market-based system offers a theoretically lower aggregate cost of compliance when compared to the compliance costs of a carbon tax or a facility-by-facility emissions limit and Carlson et. al. (2000) estimate the realized gains from emissions trading in the U.S. sulfur dioxide emissions market.

A model for how short-term carbon price fluctuations impact firm value is best understood within an accounting framework. Though financial reporting rules do not require carbon permits to be reported in firms' financial disclosures, changes in their prices impact the balance sheets of regulated firms. Nearly all of the emissions permits issued during the first phase of the ETS were grandfathered to existing industrial facilities based on their self-reported historical emissions.[12] Since firms did not pay for the permits, their distribution immediately increased the size of the balance sheets of regulated firms. The asset side of the balance sheet increased by the market value of the permits and the liabilities side increased by the expected value of the emissions required for production. In 2005, the European Commission distributed carbon permits to cover emissions between 2005 and 2007.[13] An *ex ante* valuation

[10] Sijm, Neuhoff, and Chen (2006) show that the price of carbon permits directly impacts electricity prices, as it changes the switching cost between cheaper, but more carbon intensive, coal and more expensive, but less carbon intensive, natural gas.

[11] Electricity prices differ across countries based on differences in regulations and generation. The European benchmark is German baseload electricity for delivery in one year. Its price declined from 58.6 euros per Megawatt hour on April 25 to 55.2 euros on May 22, as reported by the electricity broker the GFI Group.

[12] Auctioning for the first ETS phase was mandated by the European Commission to account for no more than five percent of the allocation in any country. Only Denmark, Hungary, Lithuania, and Ireland auctioned 5, 2.5, 1.5, and 0.75 percent of their respective permits for the 2005-2007 phase. This accounted for 0.13% of the first phase permits. Auctioned permits in the second phase were expected to account for far less than the ten percent limit.

[13] The permit allocations for each facility were based on the self-reported emissions from 2004 and

for the permits issued for the first phase of the ETS would assign a value $V_{i,0}$ to the permits for firm i at time 0 as follows,

$$V_{i,0} = \sum\nolimits_{f=1}^{M_i} \left[q_0 A_{f,i,0} - \sum\nolimits_{j=1}^{3} (1+r)^{-(j-1)} q_0 \mathbb{E}_0(E_{f,i,j}) \right], \tag{1.1}$$

where:
M_i = the number of regulated facilities owned by firm i;
q_0 = the price of carbon permits at time 0;
$A_{f,i,0}$ = the number of allocated permits and emissions at facility f owned by firm i at time 0;
r = the discount factor used to discount firm cash flows; and
$\mathbb{E}_0(E_{f,i,j})$ = the expectation at time 0 of the emissions required for production at facility f owned by firm i in in year j.

This equation values Phase 1 permits based on the assumption that emissions permits and liabilities are marked-to-market and that permit prices follow a random walk. To arrive at an expression for how the ETS impacts the valuation of participating firms I make the following three assumptions:

1. That the market is able to forecast approximate carbon positions for individual companies after the submission of the allocation plans in 2005. This can be represented by $(A_{f,i,1} - \mathbb{E}_0 E_{f,i,1}) \approx (A_{f,i,1} - E_{f,i,1})$. This assumption is supported by the numerous equity analyst reports issued in the second half of 2005 and early 2006 that estimated the net balance of carbon permits for companies across all of their regulated facilities.

2. That $(A_{f,i,1} - E_{f,i,1}) \approx (A_{f,i,j} - E_{f,i,j}), \forall j > 1$. Namely that the reported surplus/shortfall of emissions permits at each facility in the initial year of the system is a good predictor of future surpluses/shortfalls. Nearly all of the regulated facilities received an allocation of permits for the first phase of the ETS that were divided evenly across the three years. Since firms knew of the ETS in advance of the first phase of trading, there was an incentive to make the least costly emissions reductions prior the start of the first phase.[14]

3. That firms do not hedge their carbon exposure. Since some firms do hedge a portion of their carbon exposure, the value of permits can be thought of an upper bound of the potential impact of carbon market exposure on firm value.

With these assumptions, I derive the following expression for the net present value of Phase 1 emissions permits on firm value:

---

the allocation procedures established by each member country.

[14]See Buchner and Ellerman (2007a) for a discussion of the emissions abatement that occured prior to the onset of the ETS.

$$V_{i,t} = \sum\nolimits_{j=1}^{3} \sum\nolimits_{f=1}^{M_i} (1+r)^{-(j-1)} \left[q_t \left(A_{f,i,1} - E_{f,i,1} \right) \right]. \qquad (1.2)$$

If the number of allocated permits perfectly matches emissions, both the asset and the liability sides of the equation move together and changes in carbon prices should have no effect on firm value. For firms with a surplus or shortfall of carbon permits relative to their emissions requirements, changes in carbon permit prices should be associated with a one-for-one change in the market capitalization of regulated firms. The stock prices of firms with a surplus of permits should be positively correlated with carbon permit prices. The stock prices of firms with a shortfall of permits should be negatively correlated with carbon permit prices. Also, more volatile carbon permit prices should increase the volatility of the equities with large net positions in carbon permits.

This type of analysis was undertaken by a number of equity analysts at leading investment banks in late 2005 and early 2006 in an attempt to estimate the impact of the first phase of the ETS on equity prices. Analyst reports for European utilities were the most focused on carbon prices, as it impacted both balance sheets and fuel switching costs. Most concluded that active hedging within the utilities industry mitigated the short term impact of carbon prices on utilities' valuations. Analyst reports covering the cement, chemicals, pulp and paper, and mining industries also addressed the potential earnings effects of estimated shortfalls or surpluses of permits for European and international firms with facilities regulated by the ETS.

1.2 Data Sources

To test whether equity markets are affected by the carbon positions of firms, I combine the carbon registry data from the ETS with the equity data of regulated firms. The carbon data, which are from the national emissions registries aggregated by the CITL, are the same facility-level data that were released to the public during the event period. The registry data contain 9,824 industrial facilities that were allocated emissions permits in the 2005 calendar year.[15] Table 2 contains a summary of the 2005 registry data, organized by ownership type.

The first row of Table 2 contains the 347 publicly-listed firms that are used for this analysis. These firms, though they owned slightly less than 40% of the regulated sites, accounted for over two-thirds of the allocated permits in 2005. On average, the publicly-listed owners own a greater number of larger sites than either private firms or municipal owners.[16] The sites with 'Other Publicly-Listed' owners are firms with a publicly listed equity, but with a permit allocation of less than 10,000 metric tons,

[15] The first phase covered a total of 12,602 facilities. Some facilities were not allocated emissions for 2005, as they were not yet operational, or located in Bulgaria or Romania, which did not being participation in the ETS until 2007.

[16] The low thermal threshold mandated by the ETS, at 20 MW, leads to many small, privately

Table 2

Ownership Statistics from the 2005 E.U. Community Independent Transaction Log

Ownership Type	Number of Companies	Allocated Emissions (thousands of metric tons of CO_2)	% of Allocation	Number of Sites[a]	% of Sites	Average Sites per Owner	Average Allocation per Site (thousands of metric tons of CO_2)	Total 2005 Mkt. Cap. (€ millions)
Publicly-Listed	347	1,401,506,114	68%	4,020	39%	11.6	4,038,923	5,532,151
Privately-Owned	3,512	451,595,185	22%	5,205	50%	1.5	128,586	
Municipally-Owned	346	163,561,350	8%	796	8%	2.3	472,721	
Other Publicly-Listed[b]	146	37,736,847	2%	309	3%	2.0	249,447	

This table contains the breakdown of ownership for the 9,894 sites with a positive allocation of pollution permits in the 2005 Community Independent Transaction Log (CITL) published by the European Commission.
a) The emissions of 624 sites with multiple owners were apportioned according to ownership percentages.
b Other publicly-listed companies include 146 companies with aggregate emissions of less than 10,000 metric tons in 2005, firms in non-carbon intensive industries, or insufficient financial data during the event period or estimation period. See Appendix A.1 for a description of the sample of firms.

in industries with limited exposure to commodity prices, or with insufficient financial data during the event period. Appendix A.1 contains a detailed explanation of the matching procedure to link carbon data to equity data and the construction of the sample of publicly-listed firms.

The equity data for these 347 publicly-listed firms comes from DataStream. It includes daily stock price data as well as annual financial data for each of the regulated firms. The combined dataset provides a complete picture of the characteristics of the firms and industries regulated by the ETS. Table 3 contains a summary of the industries, the emissions statistics, and the country of incorporation for the sample of firms. Most of the firms are large, multi-national corporations with operations across multiple nations in the European Union. Over 70% of the firms are incorporated in Europe, but U.S. and other foreign firms are also well represented in the data. Although utilities own the largest number of sites, other industries are also well represented, especially energy-intensive industrial sectors such as mining, resource extraction, and materials processing.

The middle columns in Table 3 contain the summary statistics for firms' reported net position in carbon permits for 2005. The data show that publicly-listed firms had an overall surplus of 26 million permits in 2005. This is equivalent to 1.8% of their allocated permits, which is lower than the 3.8% surplus for the ETS system as a whole. Panel B shows how this overall surplus of permits is distributed across industries. Utilities received an aggregate under-allocation of permits and all other industries received an over-allocation.[17] However, at the firm-level there is substantial heterogeneity in net positions as evidenced by the maximum and minimum values. Within every industry there are firms with both shortfalls and surpluses of permits. Table 3 also contains the means and standard deviations of firms' 2005 asset holdings and market capitalization in the rightmost columns.

Figures 2 and 3 reinforce the heterogeneity of permit allocations and net positions across and within industries. The data in Figure 2, which are shown in log scale, show that while utilities tended to be the largest emitters of greenhouse gases, permit allocations were broadly distributed across all industries. There is also heterogeneity in the net carbon position of firms, as can be seen in Figure 3. Though utilities are the most under-allocated, there are an equal number of under-allocated firms from other industries. There is also a representative sample of industries within the most over-allocated quartile. Figure 3 shows that while many firms had permit allocations roughly equal to their first year emissions, a sizeable number had substantial surpluses or shortfalls of permits.

---

owned facilities being included in the data. For example many hospital backup generators are regulated by the system. Also, the requirement that all brick kilns be regulated adds over 1,100 privately owned facilities with very low aggregate emissions.

[17] Trotignon and Delbosc (2008) analyze the shortfall/surplus of permits at the facility-level and characterize the net-sellers and net-buyers of permits across industries and countries using the CITL data.

Table 3

Summary Statistics of Equities with Carbon Exposure in the 2005 E.U. CITL

			2005 Allocated Permits (thousands)			2005 Permit Shortfall / Surplus (thousands)					2005 Assets (€ billions)		2005 Mkt. Cap. (€ billions)	
	Firms	Sites	Total	Mean	Std. Dev.	Total	Mean	Std. Dev.	Min.	Max	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Total Sample														
All firms	347	4,020	1,401,506	4,039	12,769	26,213	76	1,812	-8,612	24,856	18.94	43.07	21.36	47.82
Panel B: By Industry														
Utilities	45	1,809	780,657	17,348	27,430	-52,801	-1,173	2,874	-8,612	4,813	25.20	34.52	19.74	21.57
Oil & Gas	32	415	186,204	5,819	7,923	13,083	409	782	-220	3,467	37.02	55.14	44.88	70.22
Construction & Materials	29	506	142,070	4,899	6,414	10,692	369	779	-352	3,796	9.17	10.70	8.63	9.53
Basic Resources	48	379	197,838	4,122	13,952	40,501	844	3,633	-82	24,856	6.27	8.16	7.22	13.51
Chemicals	49	215	37,649	768	1,079	6,427	131	226	-94	1,229	6.42	9.03	7.33	10.86
Industrial Goods & Services	48	210	35,241	734	3,090	4,548	95	422	-375	2,828	23.70	82.49	23.43	90.22
Automobiles & Parts	17	107	6,623	390	704	824	48	58	-9	182	55.22	71.76	36.68	45.09
Food & Beverage	35	256	10,697	306	612	1,978	57	133	-4	734	10.46	14.24	15.26	23.07
Health Care	21	73	2,656	126	191	635	30	70	-13	306	26.31	26.97	57.92	53.48
Personal & Household Goods	13	37	1,479	114	160	284	22	38	-3	141	17.80	20.74	24.81	37.44
Technology	10	13	394	39	26	44	4	5	-6	10	18.26	18.84	31.35	41.42

			2005 Allocated Permits (thousands)			2005 Permit Shortfall / Surplus (thousands)					2005 Assets (€ billions)		2005 Mkt. Cap. (€ billions)	
	Firms	Sites	Total	Mean	Std. Dev.	Total	Mean	Std. Dev.	Min.	Max	Mean	Std. Dev.	Mean	Std. Dev.
Panel C: By Allocation Quartile														
Largest Allocation	87	2,998	1,324,816	15,228	22,055	16,092	185	3,627	-8,612	24,856	30.21	43.11	29.43	47.02
Third Quartile	87	621	64,625	743	488	7,815	90	194	-862	962	23.25	68.67	25.35	71.99
Second Quartile	87	262	9,568	110	43	1,831	21	24	-19	110	13.26	19.51	16.46	29.22
Smallest Allocation	86	139	2,497	29	12	475	6	9	-17	37	8.81	14.52	13.96	27.68
Panel D: By Surplus/Shortfall Quartile														
Largest Surplus	87	1,829	590,832	6,791	12,008	92,288	1,061	2,764	119	24,856	31.96	71.95	34.13	80.03
Third Quartile	87	400	44,925	516	1,170	4,361	50	29	16	117	11.12	27.11	13.45	27.86
Second Quartile	87	257	19,198	221	1,255	576	7	4	1	12	11.12	15.14	15.51	25.99
Largest Shortfall	86	1,534	746,551	8,681	21,394	-71,012	-826	1,953	-8,612	1	21.52	31.82	22.32	32.56
Panel E: By Country of Incorporation														
Europe	246	3,578	1,275,173	5,184	14,806	11,444	47	2,123	-8,612	24,856	16.01	30.32	16.47	29.83
United States	71	347	84,290	1,187	3,156	7,596	107	263	-166	1,350	31.76	74.89	42.17	86.40
Other	30	95	42,043	1,401	5,571	7,174	239	950	-153	5,182	12.48	16.06	11.81	16.63

The sample consists of 347 firms covered by the first phase of the 2005-2007 of the European Union Emissions Trading System (ETS) with at least one facility in 2005 and total allocated emissions across all facilities in excess of 10,000 metric tons of carbon dioxide. The firm must have an actively traded equity between April 25 - May 16, 2006 and at least 70 trading days of returns during 150 trading day estimation period. The emissions permit data are from the European Commission Community Independent Transaction Log (CITL) for 2005. Industry classifications are based on the Dow Jones ICB system.

Figure 2

Carbon Permit Allocation By Industry for 2005 (Log Scale)

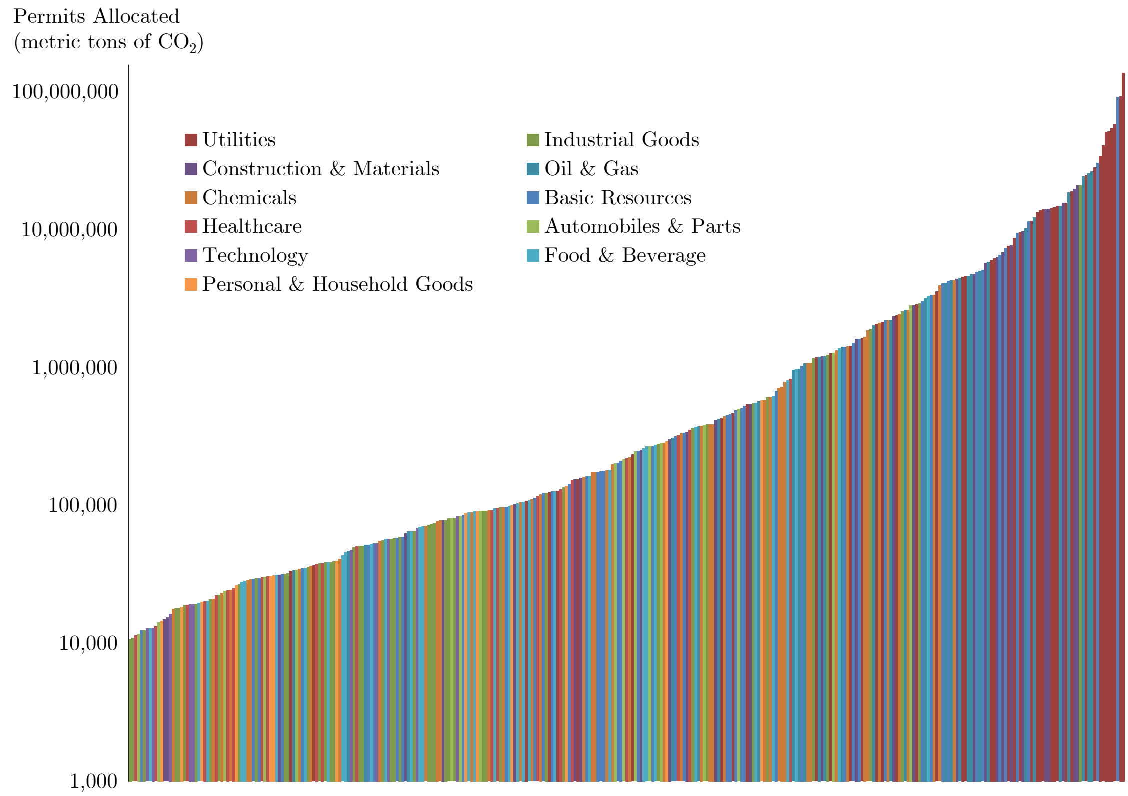

The sample consists of the 347 publicly-listed firms covered by the first phase of the 2005-2007 of the European Union Emmissions Trading System (ETS) with at least one facility with a positive allocation in 2005 and total allocated emission across all facilities in excess of 10,000 Metric Tons of carbon dioxide. The firm must have an actively-traded equity between April 25 - May 16,2006 and at least 70 trading days of returns during both the estimation period. Firms are ranked in order of their observed shortfall or surplus in emissions credits in 2005 and color-coded by their ICB classification.

Figure 3

Net Carbon Permit Shortfall/Surplus By Industry for 2005

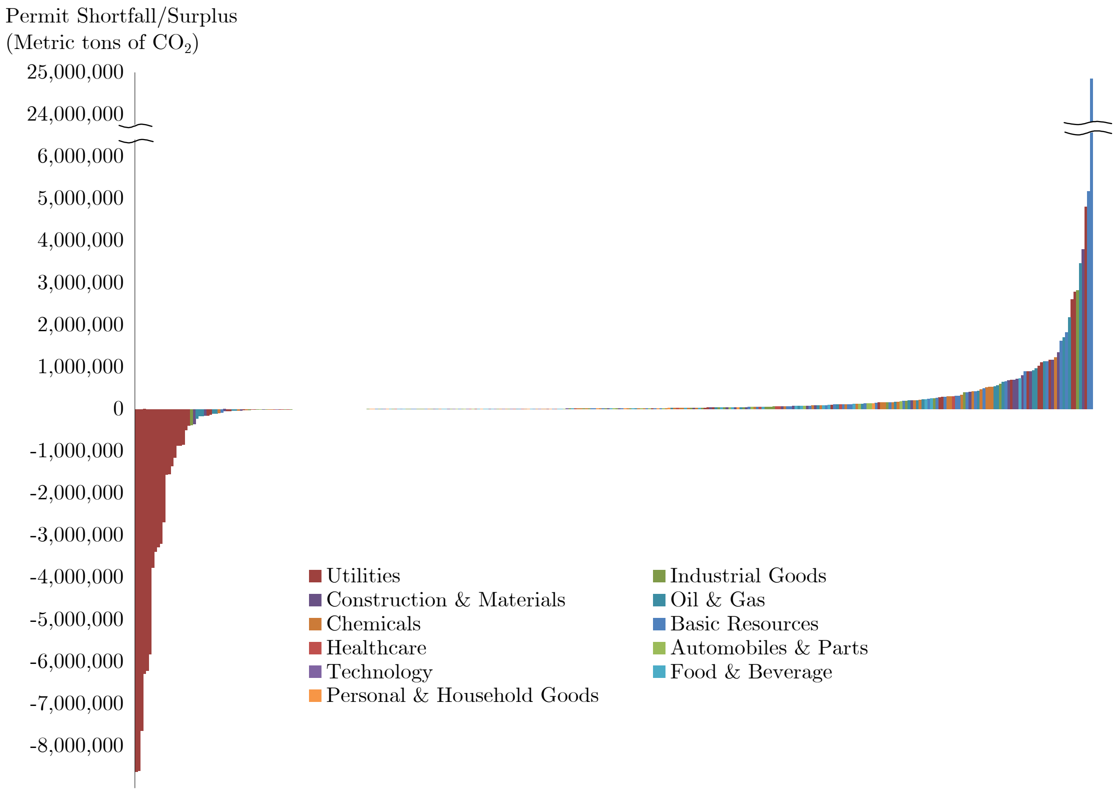

The sample consists of the 347 publicly-listed firms covered by the first phase of the 2005-2007 of the European Union Emmissions Trading System (ETS) with at least one facility with a positive allocation in 2005 and total allocated emission across all facilities in excess of 10,000 Metric Tons of carbon dioxide. The firm must have an actively-traded equity between April 25 - May 16, 2006 and at least 70 trading days of returns during both the pre- and post- event periods. Firms are ranked in order of their observed shortfall/surplus in emissions credits in 2005 and color-coded by their ICB classification. Firms with Financial, Insurance, Real Estate (FIRE), and Media ICB classification codes are excluded.

Permit allocations to industries and individual facilities are determined by the individual E.U. member states resulting in variation in stringency of caps across countries and allocation procedures across firms.[18] Variation in firms' permit holdings result from differences in allocation procedures across countries. Over 47% of the firms in the sample own facilities in more than one country, which makes their total allocation the sum of two distinct allocation procedures. Since the net permit position is the difference between allocated permits and actual emissions, differing abatement costs across facilities based on their age and efficiency further increases the heterogeneity in net positions. The institutionally-driven heterogeneity of permit allocations to the firms regulated by the ETS provides an ideal setting to test how exogenous changes to carbon prices impact equity valuations.

1.3 Analysis

The turmoil in carbon markets during the spring of 2006 provides an excellent opportunity to study the effects of carbon prices on equity valuations. The early release of emissions data by a number of E.U. countries was unanticipated and exogenous to other factors that impact equity valuations. Firm-specific carbon exposures had been analyzed by equity analysts across many of the regulated industries and a number of firms had reported their carbon allocations and emissions in their 2005 annual reports. These factors allow me to isolate the impact of short-term disruptions to carbon markets on equity markets.

For the empirical analysis, I define three time intervals. The 150-day pre-event (or estimation) period, which consists of trading day -155 through trading day -6 relative to April 24, the day before the first release of country-level emissions data. The event period, which is a 19-day window beginning the day before the Dutch and Czech releases of emissions data and ending at end of the week after the final release of E.U. emissions data.[19] The 150-day post-event period spans from trading day +6 after the May 19, 2006 end of the event period to trading day +155.

Since the ETS is a European regulatory system and the prices of its associated carbon assets are euro-denominated, the analysis is performed on euro-denominated assets. The majority of firms in the sample have their primary equity listed on a euro-denominated exchange. Performing the analysis in euros also allows for a uniform interpretation of changes to carbon permit asset holdings and firm market capitalizations. For foreign listed firms, I use both the euro-denominated equity return

[18]Phase I allocations were based on firms' self-reported 2004 emissions at each facility.

[19]The choice of an appropriate event period is complicated by the incomplete data released on May 15, 2006, representing only 87.3% of the allocated permits, and the challenge of aggregating facility-level emissions data to the corporate level. The results are robust to using a shorter 16-day event window ending the day after the May 15, 2006 data release. See Kothari and Warner (2007) for a discussion of the efficiency of long-horizon event studies.

and euro-denominated home market return to calculate betas and abnormal returns. The results are robust to the use of local currency returns.

1.3.1 Variance of Equity Returns

The theory for how short-term fluctuations in carbon prices should impact equity prices, captured by equation (2), predicts that more volatile carbon prices should increase the volatility of equity valuations. I test for changes in the equity volatility of regulated firms using three measures of volatility: the market betas of firms calculated using the DataStream market-weighted index for their home market, the mean squared daily equity returns, and the standard deviation of the equity returns. By comparing these measures across each time interval, I am able to calculate the degree to which equity volatility increased in response to higher carbon price volatility.

Table 4 presents the mean and median betas, mean squared returns, and standard deviation of returns for the sample of 347 regulated firms during each time interval. The first row of Panel A quantifies the increase in carbon market volatility during the event period, as compared with both the pre- and post-event periods. The standard deviation of daily CO_2 returns increased by an order of magnitude during the event period. The comparisons of means and medians for each of the equity volatility measures shows that the volatilities of regulated firms were higher during the event period than the pre- and post-event periods. Both the mean and median betas increased by approximately 8% between the pre-event and event-periods. Ratios of mean squared daily returns and standard deviations of returns increased by an even greater percentage during the event period. All three of these volatility measures declined during the post-event period with the resolution of uncertainty over first year carbon positions.

Panel B contains the statistical tests of the difference in volatility during the pre- and post-event period relative to the event period. The paired t-test compares differences in means across intervals and a Wilcoxon sign-ranked test compares sample medians across intervals.[20] With a high degree of statistical certainty, I reject the null hypothesis that the volatility in the pre-event period was equal to the volatility during the event-period. Each of the measures was higher during the event period. I also reject the null hypothesis that the event period mean squared daily returns and standard deviation of returns was equal to the post-event period at the 1% confidence level. I am unable to reject the null hypothesis that the average beta of regulated firms was the same during the event period and the post-event period, but reject the null that the medians were the same at the 10% confidence level.

The results in Table 4 indicate that the equity volatility of regulated firms increased in response to higher carbon price volatility. The sample of regulated firms, many with stable and predictable cash flows such as utilities and industrials, is pre-

[20] The Wilcoxon sign-ranked test is a non-parametric test of medians that does not require distributional assumptions for the variable of interest.

Table 4

Changes in Stock Price Volatility in Response to Carbon Price Uncertainty

Panel A: Risk Measures

	Sample	Pre-event Period[a]		Event Period[b]		Post-event Period[c]	
	Size	Mean	Median	Mean	Median	Mean	Median
Std. Dev. of CO_2 Returns		0.0188		0.1915		0.0288	
Beta from 1-Factor Model	347	0.9126	0.8913	0.9813	0.9685	0.9569	0.9478
Mean Squared Daily Returns	347	0.00034	0.00024	0.00045	0.00031	0.00035	0.00026
Std. Dev. of Equity Returns	347	0.0172	0.0156	0.0190	0.0175	0.0173	0.0162

Panel B: Hypothesis Tests

	Hypothesis One[d]		Hypothesis Two[e]	
	Paired t-Test Statistic	Wilcoxon Signed-Rank Test Z-Statistic	Paired t-Test Statistic	Wilcoxon Signed-Rank Test Z-Statistic
Beta from 1-Factor Model	-2.4995***	-2.250**	-0.9595	-1.652*
Mean Squared Daily Returns	-5.9072***	-5.844***	-4.2896***	-6.333***
Std. Dev. of Equity Returns	-5.4253***	-4.864***	-4.7634***	-5.106***

Carbon returns are calculated using the December 2007 futures prices from the European Climate Exchange. Beta estimates are from 1-factor regressions of logarithmic returns for the euro-denominated return index of each equity on the euro-denominated DataStream market index in the home market of the equity during the each period. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

a) The pre-event period consists of 150 trading days beginning 155 trading days before the the event period to five trading days before the event period.

b) The event period consists of the 19 trading days between the day prior to the first leak of country-level emissions news on April 25, 2006 to the end of the week after the release of system-wide 2005 emissions data by the European Commission on May 15, 2006.

c) The post-event period consists of 150 trading days beginning five trading days after the end of the event period through 155 trading days after.

d) Hypothesis One tests whether the pre-event-period betas/returns/standard deviations are the same as the event-period betas/returns/standard deviations. e) Hypothesis Two tests whether the post-event-period betas/returns/standard deviations are the same as the event-period betas/returns/standard deviations.

disposed to have a low beta, as can be seen during the pre-event period average beta of 0.91. During the weeks of carbon price uncertainty, that beta, as well as the volatility of returns, increased substantially. The mean and median betas were lower in the post-event period, but only the median is statistically different from zero. This may be because of the residual uncertainty in the carbon market stemming from the lack of Polish emissions data, or market uncertainty over Phase 2 allocations, which were submitted in the summer of 2006. Note that CO_2 returns were nearly twice as volatile during the post-event period as compared with the pre-event period. To better understand the mechanism that connects carbon markets to equity prices, the next subsection analyzes the direction and magnitude of abnormal returns during the event period.

1.3.2 Cumulative Abnormal Returns

Equation (2) predicts that the stock prices of firms with a net shortfall (surplus) of permits should rise (fall) in response to a drop in carbon prices. This section tests this prediction using an event study methodology. To do so, I estimate the cumulative abnormal stock returns of regulated firms during the event period.

Abnormal returns are computed using a 1-factor market model.[21] The abnormal return on each day during the event period is,

$$\widehat{\epsilon}_{i,t} = r_{i,t} - \widehat{\alpha}_i - \widehat{\beta}_i r_{m,t}, \tag{1.3}$$

where:

$r_{i,t}$ = the euro-denominated, logarithmic return for security i from day $t-1$ to day t;

$r_{m,t}$ = the euro-denominated, logarithmic return of the DataStream return index for the home market of security i from day $t-1$ to day t;

$\widehat{\alpha}_i, \widehat{\beta}_i$ = are the ordinary least squares (OLS) estimates of the intercept and market beta for security i computed during the 150-trading-day estimation period.

Since the event period is the same for all regulated firms and many of the firms come from overlapping industries, there may be correlation in the cross-sectional returns of firms. This correlation can lead to non-zero covariances across abnormal returns which effects the efficiency of OLS hypothesis testing. Campbell, Lo, and MacKinlay (1997) show that aggregating abnormal returns into a portfolio and then testing the cumulative abnormal returns (CAR) of the portfolio mitigates a non-zero covariance of abnormal returns across firms. Portfolios of regulated firms are constructed by industry, by carbon permit allocations and net positions, and by

[21] To check the robustness of the results, I also estimate abnormal returns using a 2-factor model with industry returns and find similar results. The results are also robust to the choice of currency (euro-denominated or local currency) and the calculation of returns (logarithmic or arithmetic).

country of incorporation. Cumulative abnormal returns for a portfolio of regulated firms n are calculated as follows,

$$CAR_n = \frac{1}{T}\sum\nolimits_{k=0}^{T}\frac{1}{N}\sum\nolimits_{i=0}^{N}\widehat{\epsilon}_{i,k}, \tag{1.4}$$

where T is the number of days in the event period and N is the number of regulated firms in the portfolio.

Testing the statistical significance of cumulative abnormal returns is complicated by the results in Section 1.3.1 that show that the volatility of regulated firms increased during the event period. The Brown and Warner (1985) test statistic, $\frac{CAR_n}{se(\widehat{\epsilon}_n)\sqrt{T}}$, which is the ratio of cumulative abnormal returns during the event period to the time-weighted standard error of returns during the estimation period can lead to Type I errors resulting from a misestimation of the true standard deviation of abnormal returns. To overcome this problem, I compute both a parametric t-test statistic as well as a non-parametric Wilcoxon sign rank statistic that tests whether the median abnormal return of each portfolio is statistically different from zero. Boehmer, Masumeci, and Poulsen (1991) show that the combined use of the standard t-test statistic along with a sign rank test provides a more robust test of the null hypothesis in the presence of event-induced variance changes. The statistical significance estimates of the abnormal returns yields similar results across both tests.

Table 5 contains the mean and median cumulative abnormal returns during the event period for the regulated firms in the sample. The results show that regulated firms experienced a statistically significant average CAR of -1.65% during the event period. Since firms had an aggregate oversupply of permits, this is the result that one would expect from the declining carbon prices during the event period. For the sample of regulated firms his translates to a 90 billion euro decrease in market capitalization. Lower carbon permit prices decreased the value of the net surplus of carbon permits allocated during the first phase of the ETS. The CAR for the median firm was -1.48%, which was statistically significant at the 1% confidence level.

The aggregate negative returns of the ETS are driven by the four industries (oil & gas, basic resources, construction, and chemicals) with the largest net surplus of permits. As is shown in Table 3, these four industries account for over 40% of the allocated permits and nearly 90% of the surplus permits in the first phase of the ETS. Their average CARs, shown in Panel B of Table 5, ranged between -3% and -4.4% during the event period. All of the CARs were statistically different from zero according to both the parametric and non-parametric tests, except for the oil and gas industry, whose -3.98% mean CAR is statistically indistinguishable from zero.[22]

Panel D shows that the quartile of firms with the largest permit surplus had abnormal returns that diverged from the quartile of firms with the largest permit

[22]The oil and gas industry has the highest abnormal return variance during the estimation period. This causes the economically significant -3.98% mean CAR during the event period to be statistically indistinguishable from zero.

Table 5

Cumulative Abnormal Returns (CAR) of Equities with Carbon Exposure

		Event Period[a]	
	Firms	Mean	Median
Panel A: Mean CAR (%)			
Full Sample	347	-1.65**	-1.48***
Panel B: Mean CAR (%) By Industry			
Utilities	45	-0.26	0.41
Oil & Gas	32	-3.98	-3.65***
Construction & Materials	29	-4.00**	-3.65***
Basic Resources	48	-4.38**	-4.05***
Chemicals	49	-2.96**	-2.30**
Industrial Goods & Services	48	-0.45	-1.62
Automobiles & Parts	17	2.85	0.85
Food & Beverage	35	0.32	-0.07
Health Care	21	1.48	1.21
Personal & Household Goods	13	-0.31	0.50
Technology	10	-2.71	-2.02
Panel C: Mean CAR (%) By Allocation Quartile			
Largest Allocation	87	-2.99***	-2.94***
Third Quartile	87	-0.95	-1.19
Second Quartile	87	-0.83	-0.82
Smallest Allocation	86	-1.82**	-1.74**
Panel D: Mean CAR By Surplus/Shortfall Quartile			
Largest Surplus	87	-3.29***	-3.30***
Third Quartile	87	-0.12	-0.52
Second Quartile	87	-2.73***	-2.44***
Largest Shortfall	86	-0.44	-1.25
Panel E: Mean CAR (%) By Country of Incorporation			
Europe	246	-2.03***	-1.55***
United States	71	-0.90	-0.10
Other	30	-0.32	-2.82

The sample consists of 347 firms covered by the first phase of the 2005-2007 of the European Union Emissions Trading System (ETS) with at least one facility in 2005 and total allocated emission across all facilities in excess of 10,000 Metric Tons of carbon dioxide. The Cumulative Abnormal Return (CAR) is the sum of residuals from a 1-factor regression of the logarithmic return of the euro-denominated Total Return Index on the euro-denominated, market-weighted DataStream Return Index for the home market of the equity using estimates from the 150 day pre-event estimation period. Industry classifications are based on the Dow Jones ICB system. Statistical significance for the mean CARs use a t test. Statistical significance for the median CARs use a Wilcoxon sign rank test. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

a) The event period is the nineteen trading days between the day prior to the April 25, 2006 leak of data and the end of the week following the May 15, 2006 release of EU data.

deficit. Most of the firms in the surplus quartile are from the four industries with statistically significant negative CARs. This quartile had an average and median cumulative abnormal return of -3.3% during the event period. Both the mean and median abnormal returns are statistically significant at the 1% confidence level.

Firms in the industry with the largest shortfall of permits, the utilities industry, experienced only a slightly negative average abnormal return and a positive median return, both of which were statistically indistinguishable from zero. The theory presented in Section 1.1.2 predicts that firms with a negative net position should experience positive returns in response to a fall in carbon permit prices. This is not evidenced by the data. Panel D shows that the quartile of firms with the largest shortfall of permits, many of which are from the utility industry, had a mean CAR of -0.44 that is statistically indistinguishable from zero.

Carbon Price Hedging

There are two factors that might explain the lack of a positive abnormal returns for both utilities' and the quartile of firms with the largest shortfall of permits. First, carbon prices not only impact the value of carbon permits on utilities' balance sheets, but they also impact the price of electricity, which has a direct effect on their revenue. During the event period, the fall in carbon prices led to a 6% decline in electricity futures prices. This decline in future revenue exerted negative pressure on utilities' equity prices, which would counteract a balance sheet-driven stock price increase.

A second explanation is that firms with large carbon liabilities more actively hedged their exposure to fluctuations in carbon prices. Froot, Scharfstein, and Stein (1993) show that firms have an incentive to hedge their exposure to price changes to the extent that they seek to limit cash flow variability from existing assets. Utilities were the industry issued the fewest permits relative to historical emissions and also the industry that is most engaged in hedging revenues and costs.[23] The valuation impact of changes to carbon prices discussed in Section 1.1.2 assumes a lack of carbon price hedging because of the dampening effect that such hedges would exert on the balance sheet valuation of carbon permits. If the drop in electricity prices had an offsetting effect and the hedged position of utilities had a dampening effect, together these effects explain the non-positive CAR of both utilities and the firms in the quartile with the largest carbon permit shortfall. To the extent that firms' seek to minimize unexpected changes to their cash flows and equity prices, this finding shows that firms regulated by cap-and-trade systems have an incentive to hedge their exposure to carbon prices.

Though firms with the largest permit shortfall did not experience positive abnormal returns, their stock price behavior during the event period is sharply contrasted by the very negative CARs of the surplus firms. Since the non-parametric analysis

[23] Analyst reports written in early 2006 on the European utilities sector noted the extent to which these firms had hedged their carbon exposure.

of CARs does not account for differences in firm characteristics across portfolios of firms, the next section uses regressions of individual carbon positions and firm characteristics to isolate the marginal contribution of the net carbon position of firms on their abnormal returns.

1.4 Factors Influencing Cumulative Abnormal Returns

This section examines whether the carbon balance sheet positions of individual firms explain abnormal returns. The analysis in the first subsection uses carbon data and firm-specific financial data to isolate the effect of permit shortfalls or surpluses on abnormal returns. The second subsection calculates the degree to which changes in firms' total carbon balance sheet value pass through to changes in market capitalization during the event period. Both sets of analysis show that firms' net carbon positions dictate the degree to which changes in carbon prices affect firms' equity valuations.

1.4.1 Net Positioning

Table 6 shows the results of regressions of firms' net permit holdings on their event period CARs. All of the specifications include industry fixed effects to account for the cross-industry differences in abnormal returns shown in Table 5. The percentage impact columns calculate the percentage point change in the event period cumulative abnormal return associated with a one standard deviation change in the explanatory variable.

The coefficient on firms' shortfall or surplus in 2005 carbon permits is negative and statistically significant across all of specifications. A one standard deviation increase in net permit holdings is associated with an approximate negative four tenths percentage point CAR during the event period. The magnitude is stable when controlling for 2005 total permit allocations, firm asset holdings, market capitalization, and country of incorporation. Firms with larger permit allocations experienced more negative returns in response to falling carbon prices. However, the coefficient on the permit allocation variable decreases in magnitude and becomes statistically indistinguishable from zero after controlling for firm size using market capitalization. Though statistically insignificant, the very negative coefficient in specification (5) on the indicator variable for firms incorporated in Europe confirms the results in Table 5 showing that these firms had more negative returns than foreign firms. Though the goodness of fit across all of the specifications is relatively low, which is typical for regressions explaining abnormal returns, specification (4) explains nearly 2.5% of the total variation in CARs.

Table 6

Regressions of Emissions Permit Holdings on Equity Cumulative Abnormal Returns

	Dependent Variable is Event Period CAR (in percent)[a]									
	(1)		(2)		(3)		(4)		(5)	
	Coefficient	Percentage Impact	Coefficient	Percentage Impact	Coefficient	Percentage Impact	Coefficient	Percentage Impact	Coefficient	Percentage Impact
2005 Permit Shortfall /	-0.2209*	-0.40	-0.2043**	-0.37	-0.1893*	-0.34	-0.2282**	-0.41	-0.2273**	-0.41
Surplus (millions)	(0.1323)		(0.1081)		(0.1063)		(0.1080)		(0.1086)	
2005 Allocated			-0.0316	-0.40	-0.0432**	-0.55	-0.0311	-0.40	-0.0257	-0.33
Permits (millions)			(0.0196)		(0.0200)		(0.0217)		(0.0223)	
Total Assets					0.0116*		-0.0310		-0.0257	
in 2005 (€ billions)					(0.0066)		(0.0260)		(0.0255)	
Market Cap.							0.0403*		0.0374*	
in 2005 (€ billions)							(0.0215)		(0.0211)	
Incorporated in Europe									-1.297	
									(0.9719)	
Constant	-1.6406***		-1.5143***		-1.696***		-1.7866***		-0.8383	
	(0.3831)		(0.4044)		(0.4471)		(0.4522)		(0.8292)	
Industry Fixed Effects	Yes		Yes		Yes		Yes		Yes	
Observations	347		347		347		347		347	
R-squared	0.010		0.012		0.022		0.024		0.024	

The Cumulative Abnormal Return (CAR) is the cumulative residual from a 1-factor regression of the logarithmic return of the euro-denominated Total Return Index on the euro-denominated, market-weighted DataStream Return Index for the home market of the equity using estimates from the 150 day pre-event estimation period. The "Percentage Impact" column reports the percentage point change in the CAR associated with a one standard deviation increase in the explanatory variables. Robust standard errors are reported in parentheses below the coefficient estimates. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

a) The event period consists of the 19 trading days between the day prior to the first leak of aggregate emissions news on April 15, 2006 to the end of the week after the release of first year emissions data by the European Commission on May 15, 2006.

The results in Table 6 establish a relationship between firms' net carbon position and their abnormal returns in response to falling carbon prices. However, the relationship established in equation (2) shows that it is not only current shortfalls or surpluses that should impact firm valuations, but also the discounted value of future permit holdings. Since firms were allocated three years worth of permits at the outset of Phase 1 of the ETS, markets should have incorporated the total change in value of all three years worth of net permit holdings. Multiplying the point estimates on the shortfall/surplus coefficient by three provides an approximate magnitude of the total Phase 1 net permit holdings on abnormal returns. A one standard deviation change in Phase 1 net permit holdings yields a -1.23% CAR during the event period. A two standard deviation change in Phase 1 net permit holdings can explain nearly all of the difference between the quartiles of firms with the largest surplus of permits and the quartile with the largest shortfall of permits shown in Figure 4.

1.4.2 Market Capitalization Impact of Lower Carbon Prices

Specification (4) in Table 6 shows the marginal impact of net carbon positions on CARs when controlling for firms' market capitalization as a separate explanatory variable. This subsection explicitly tests whether the euro-denominated changes to firms' market capitalization during the event period reflected the change in the value of firms' net carbon position.

To test the market capitalization impact of changes to carbon prices, I run the following regression,

$$\Delta MarketCap_{event} = a + b\left[\Delta q_{event} \cdot 3\left(A_1 - E_1\right)\right], \tag{1.5}$$

where:

$\Delta MarketCap_{event}$ = market capitalization of each firm at the start of the event period multiplied by (1+ CAR), and

Δq_{event} = the 14 euro per ton price decline of December, 2007 carbon permit futures during the event period.

The net carbon position of each firm in 2005 is multiplied by three to approximate the total shortfall/surplus in the Phase 1 permits. As shown in equation (2), the current value of a firms' net carbon position should account for the discounted value of future carbon net positions, so multiplying 2005 net positions by a factor of three should lead to a slightly downwardly biased estimate of the true effect. If the statistically significant CARs shown in Section 1.3.2 were due to changes in firm value attributable to the total change in Phase 1 permit values, one would expect the b in equation (5) to be approximately equal to one.

Table 7 shows the regression results of Equation (5) for both the entire sample of regulated firms and the sample of the five industries that accounted for 96% of the allocated 2005 permits. The results show that changes in the market capitalization

Table 7

Market Capitalization Impact of Carbon Price Decline

	Change in Market Capitalization[a] (in billions of euros)	
	Full Sample	Top 5 Industries[b]
	(1)	(2)
Change in Value of Phase 1 Net Permit Holdings During Event Period[c]	1.5470**	1.6800***
(€ billions)	(0.7353)	(0.6926)
Constant	-0.0675	-0.3163***
	(0.0769)	(0.0784)
Industry Fixed Effects	Yes	Yes
Observations	347	203
R-Squared	0.008	0.025

Robust standard errors are reported in parentheses below the coefficient estimates. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

a) The change in market capitalization is the product of the euro-denominated market capitalization on April 21 and one minus the cumulative abnormal return during the event period.

b) The top 5 industries are: utilities, oil & gas, construction & materials, basic resources, and chemicals. These industries account for 96% of the allocated 2005 permits.

c) The change in the value of net permit holdings is equal to the price change of December, 2007 carbon permit futures prices during the event period (decline from 30.95 euros to 16.95 euros) multiplied by three times the net permit holdings of each firm in 2005.

of regulated firms were positively correlated with the change in value of Phase 1 ETS permits during the event period. However, the average effect was substantially greater than one-for-one, as evidenced by the coefficients of 1.54 and 1.68 on the net permit holdings coefficient. The coefficient in specification (1) is statistically significant at the 5% level and the restricted sample is statistically significant at the 1% level.

The coefficients on the net permit holding variables in Table 7 indicate that the abnormal changes in regulated firms' equity value during the event period were greater than the total change in value of Phase 1 permit holdings. Since the ETS was structured as a multi-phase system, with the second phase scheduled to begin in 2008, the greater responsiveness of regulated firms' market value may incorporate expectations of lower asset values for yet unallocated phase two permits.[24] Since the market value of a publicly traded company can be calculated as the net present value of all future price changes and dividends, so too should it reflect expected changes in value for future carbon permit holdings. Figure 1 shows that while the price decline for Phase 2 ETS permits was less severe than the decline in phase one permit prices, these permits prices also declined on the news of the surplus of permits in 2005. Though the

[24]The E.U. member countries were in the process of establishing their Phase 2 National Allocation Plans in May of 2006.

exact phase two allocations remained uncertain at the time, the allocation procedures in each country could reasonably be assumed to remain comparable to the Phase 1 allocation procedures. This means that firms net permit holdings in 2005 are a strong predictor for future permit surpluses or shortfalls and that the carbon price decline would impact the value of these future holdings.

Timing of Market Reactions

This section examines the timing of the equity market reaction to changes in carbon prices. It also analyzes data from options markets to understand how carbon market volatility was incorporated into forward-looking equity volatility estimates. The findings show that the stock and options markets responses lagged the price changes and volatility in carbon markets. I attribute these results to the short history of carbon markets and informational constraints.

1.4.3 Equity Market Reaction

The early release of emission data by a subset of countries before the official May 15 release date caused the 19-day period of carbon market volatility and price declines that is examined in the proceeding sections. The results in Sections 1.3 and 1.4 showed that firms with large net carbon exposures had economically and statistically significant returns over the course of the entire event period. An analysis of the daily abnormal returns of the firms with the largest net positions in carbon permits provides a more detailed understanding of when the carbon price changes were incorporated into equity valuations.

Figure 4 charts the daily abnormal returns for the quartiles of firms with the largest surplus and largest shortfall of permits. The abnormal returns of both groups moved in tandem during the week prior to the first release of emissions data and in the weeks between the first release of country-level data and the official release of emissions data. The over three percentage point divergence of CARs between the two groups during the event period, shown in Panel D of Table 5, occurred during two sub-periods. The two groups first diverged during the three days following the initial country-level releases of carbon data when carbon permit prices dropped 14 euros.[25] During this period the surplus quartile had a statistically significant abnormal return of -1.24%, whereas the quartile of firms with the largest permit shortfall experienced a statistically insignificant abnormal return of -0.26%.[26] The

[25] Table II contains the complete timeline of carbon-market news releases during the event period. The five countries that reported emissions between April 25-27 accounted for 20.1% of allocated permits.

[26] The quartile of firms with the largest permit surplus had a cumulative abnormal return of -1.24 from April 25-27 that is statistically significant at the 1% level.

Figure 4

Daily Abnormal Returns Aggregated By 2005 Permit Shortfall/Surplus

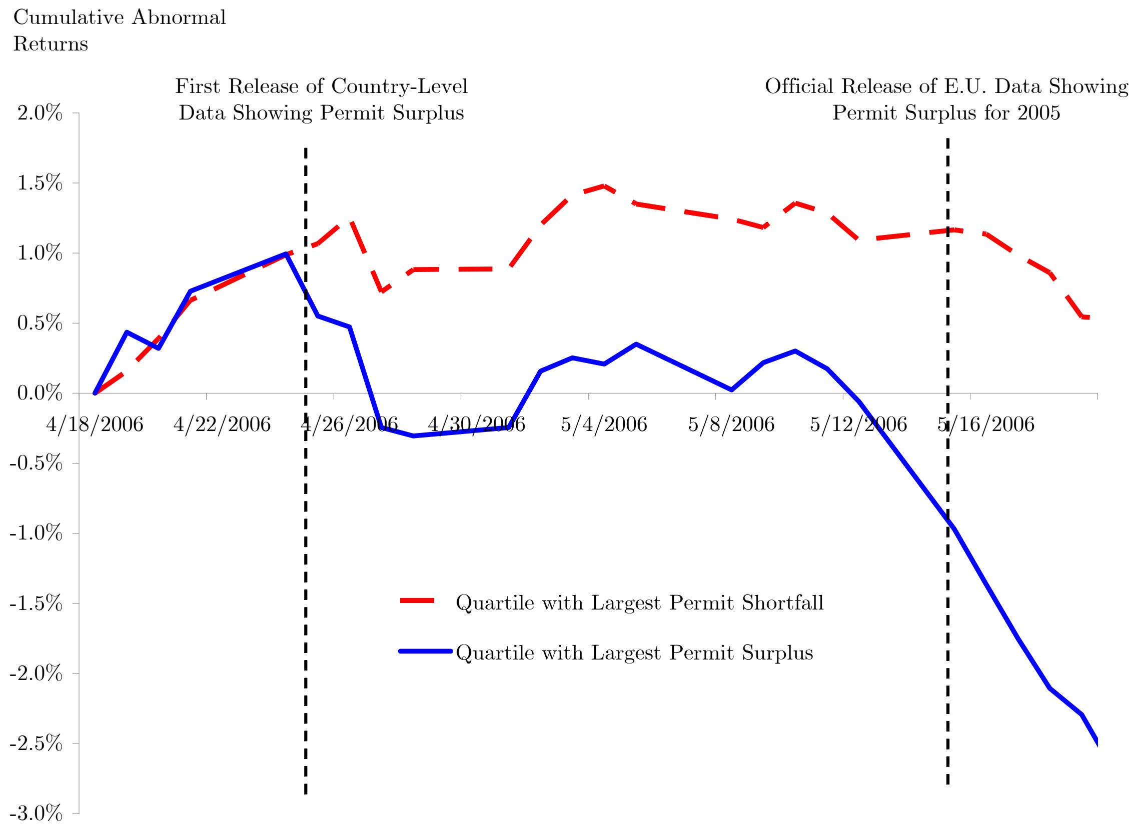

Abnormal returns are the daily residuals from a 1-factor regression of the logarithmic return of the euro-denominated security on the euro-denominated, market-weighted DataStream market return for the home market of the equity. Intercepts and betas are estimated over the 150-trading-day beginning 155 trading days prior to April 25, 2006. Abnormal returns are aggregated by realized 2005 permit shortfall/surplus quartiles.

second divergence occurred during the week following the official release of emissions data for the E.U. on Monday, May 15. The two groups CARs diverged by 1.39% between May 12-16 and another 0.34% between May 17-19. Though carbon prices had been extremely volatile in the intervening weeks, the carbon futures price on May 19 was the same as it had been on April 27 after the first release of carbon data.

The time path of abnormal returns during the event period shows that markets did not fully incorporate the impact of lower carbon prices into their equity valuations until after the official release of emission data. This delayed response can be attributed to two possible factors. First, since this period was only a little more than a year after carbon markets had begun active trading, the full impact of these prices on equity values may not have been fully understood by investors. Though a number of analyst reports had calculated the earnings impact under different carbon price scenarios for many of the regulated firms, the large carbon exposures of certain firms may have been underestimated. Second, investors may have been unable to accurately forecast individual firms' expected emissions based on the historical 2004 emissions data and 2005 allocated permits. Without an expectation for firms' shortfall or surplus, an investor would not be able to incorporate the impact of changes in carbon permit holdings on earnings until the complete 2005 emissions data was released on May 15.

1.4.4 Options Market Reaction

The findings in Section 1.3.1 that showed that the volatility of regulated equities increased in response to fluctuating carbon prices should also be observable in options markets. A large body of research, including Christensen and Prabhala (1998) and Britten-Jones and Neuberger (2000), shows that implied volatility is a better predictor than past volatility in forecasting future volatility in equity markets. An analysis of implied volatilities and new releases by Ederington and Lee (1996) shows that implied volatilities decline after scheduled economic data releases and rise after unscheduled information releases. Based on these findings, one would expect the implied volatilities for the equity options of regulated firms to have risen after the unscheduled news release in late April and fallen after the scheduled release of carbon data on May 15.

To examine the response of implied volatilities to news from the carbon markets, I use call options data from DataStream for the 77 of the 347 firms that have options traded on exchanges during the event period.[27] The 3 month constant-maturity implied volatility estimates are calculated using the binomial pricing model of Cox, Ross, and Rubinstein (1979). The implied volatilities are averaged daily across the firms in the sample.

Figure 5 shows the time series of the implied volatilities of regulated firms. Confirming the results from the previous subsection that showed a delayed price response in equity markets, options prices exhibited an even more pronounced lagged response to the volatility in carbon markets. Instead of rising after the unscheduled release of the first country-level emissions data on April 25, the equity implied volatilities declined during the subsequent two weeks when carbon prices were at their most volatile. Only after the scheduled release of E.U. emissions data on Monday, May 15 do the implied volatilities of regulated equities rise above their pre-event period average. These data tell a story of options markets neglecting to incorporate the impact of carbon permit prices on firm value. Instead of incorporating the valuation effects of falling carbon prices into the equity options prices after the initial country-level surpluses were reported, they only responded after the complete emissions data were released. Possibly in further anticipation of the effects of Polish emissions data on

[27] The 77 firms in the sample are those with at least thirty days of call options data in the pre- and post-event periods.

Figure 5

Implied Volatilities of Regulated Equities

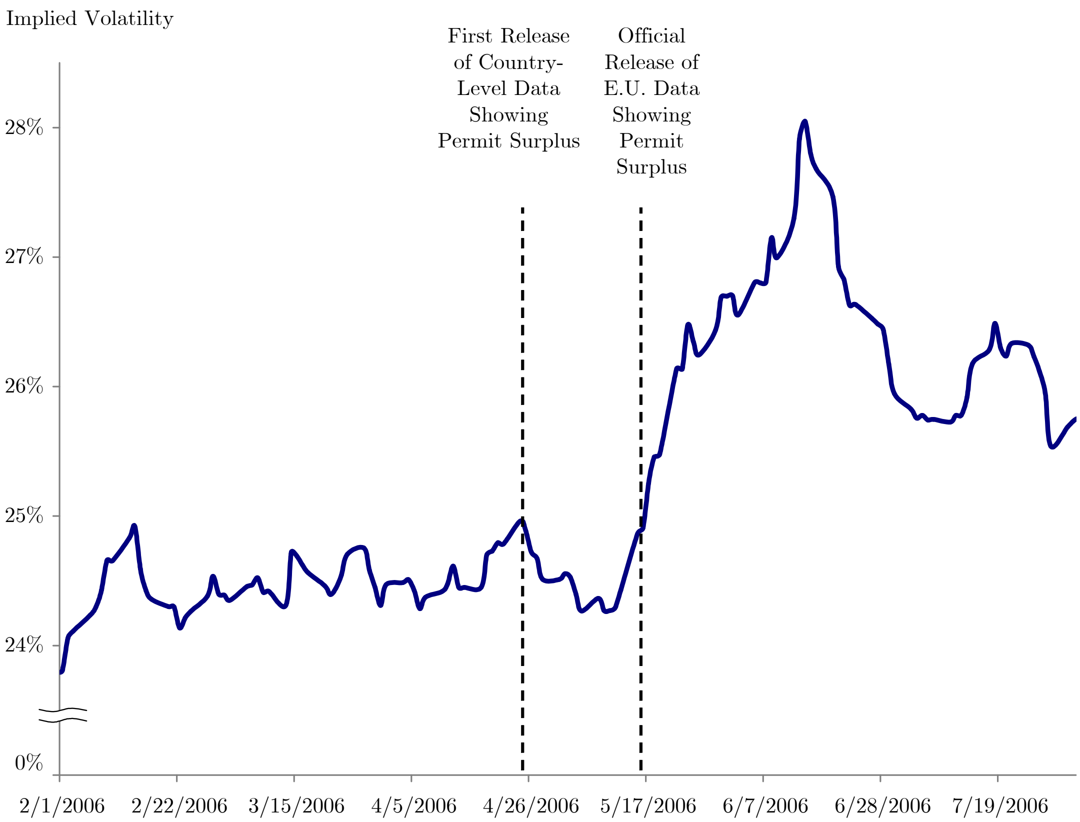

The average daily implied volatility is calculated using the prices of exchange-traded call options from DataStream. Implied volatilites are calculated using the binomial options pricing model of Cox, Ross, and Rubinstein (1979) for a constant maturity 90 day call option. The sample includes the 77 firms regulated by the ETS with at least thirty days of options data in the pre- and post-event periods.

carbon valuations, the implied volatilities continued to increase until June 14 when Poland released its emission data. The behavior of implied volatilities during the event period shows that options prices failed to incorporate carbon price changes prior to the spike in realized volatility when equity prices exhibited their delayed response on May 15.

1.5 Conclusion

This paper shows that price fluctuations in carbon markets can have large price and volatility impacts on the equity prices of regulated firms. Using a new dataset on the first phase of the European Trading System, I show that a period of high carbon price volatility in 2006 led to higher equity volatility for regulated firms and abnormal returns consistent with a balance sheet valuation of carbon assets or liabilities of over- and under-allocated firms. Changes in carbon prices have large impacts on the net present value of future net permit positions and thereby cause large changes in firms' market capitalization.

Instead of anticipating and immediately incorporating price changes in carbon markets into stock and option prices, equity and derivatives markets were slow to respond to carbon price changes. I attribute these delayed responses to the newness of carbon markets and the failure to recognize the large number of firms with significant carbon price exposure. To the extent that the delayed response of equity and options markets can be attributed to investors learning how to incorporate carbon prices into their models, future research examining the market responses during subsequent reporting periods in the first phase of the ETS as well as the second phase of the ETS will provide a more complete understanding of the interactions between carbon markets and existing capital markets. As financial reporting standards begin to require carbon asset disclosure in corporate financial statements, investors will be better able to react to carbon-related news.

The paper also shows that permit allocation procedures matter. One of the main drivers of the results were the large net positions in carbon permits that were granted to many firms regulated by the ETS. Though the differences between allocated permits and business as usual emissions may not be large on a facility-by-facility basis, many regulated firms own numerous facilities which can compound allocation inequities. Though the permit allocation procedures differed across member states, preference was given to facilities in industries deemed to be trade sensitive, such as steel and aluminum production. Governments should recognize that conducting industrial policy through cap-and-trade systems exerts costs on industries and firms. The findings also indicate that more unequal allocations will lead to greater equity market volatility attributable to fluctuations in carbon prices.

Since it was the uncoordinated release of carbon data across different countries that caused the period of volatility analyzed in this paper, governments should guard carbon emissions data as they do other market sensitive data such as employment or growth figures. Another institutional change that would help to limit the volatility of carbon prices stemming from information releases is the more frequent reporting of emissions data. More frequently collected emissions data would smooth the price adjustment process and decrease aggregate uncertainty in these markets.

The findings attribute a portion of the muted reaction of utility equity prices to volatile carbon prices to their largely hedged position in the carbon market. The list of publicly-traded companies regulated by the ETS shows that there are many companies not traditionally associated with high frequency commodity price exposure. Firms that fail to hedge a large net carbon positions face the prospect of higher volatility in their cost of capital stemming from changes to carbon prices. Firms that hedge their carbon exposure less efficiently than their competitors may find themselves at a competitive disadvantage. A certain outcome of a growth and expansion of greenhouse gas cap and trade systems is that they create large business opportunities for banks and exchanges to create and intermediate the trading of carbon derivative products.

Understanding the impact of emissions market on equity markets is important for the ongoing policy discussions concerning cap-and-trade and the regulations of greenhouse gases in the United States and in other countries. As policymakers craft legislation to stem the growth of carbon emissions, regulation should be undertaken with a full understanding of the different impacts that different forms of regulation exert on existing markets. This paper shows that a cap-and-trade mechanism, though allowing for abatement efficiency, may exert costs on regulated firms that include increases in the volatility of their cost of capital. Data from the E.U. cap-and-trade system allows these types of effects to be quantified and discussed outside of a purely theoretical framework.

Bibliography

[1] Bartov, Eli, and Gordon M. Bodnar, 1994, Firm Valuation, Earnings Expectations, and the Exchange-Rate Exposure Effect, *Journal of Finance* 49, 1755–1785.

[2] Boehmer, Ekkehart, Jim Masumeci, and Annette B. Poulsen, 1991, Event-study Methodology Under Conditions of Event-Induced Variance, *Journal of Financial Economics* 30, 253–272.

[3] Brown, Stephen J., and Jerold B. Warner, 1985, Using Daily Stock Returns: The Case of Event Studies, *Journal of Financial Economics* 14, 3–31.

[4] Britten-Jones, Mark and Anthony Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, *The Journal of Finance* 55, 839–866.

[5] Buchner, Barbara K., and Denny A. Ellerman, 2007a, Over-Allocation or Abatement? A Preliminary Analysis of the EU ETS Based on the 2005 Emissions Data, Working paper 90, Fondazione Eni Enrico Mattei.

[6] Buchner, Barbara K., and Denny A. Ellerman, 2007b, The European Union Emissions Trading Scheme: Origins, Allocation, and Early Results, *Review of Environmental Economics and Policy* 1, 66–87.

[7] Burtraw, Dallas, and Karen L. Palmer, 2008, Compensation Rules for Climate Policy in the Electricity Sector, *Journal of Policy Analysis and Management* 27, 819–847.

[8] Bushnell, James B., Howard Chong, and Erin T. Mansur, 2009, Profiting from Regulation: An Event Study of the European Carbon Market, Working paper 15572, National Bureau of Economic Research.

[9] Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer, 2000, Sulfur Dioxide Control By Electric Utilities: What Are the Gains From Trade? *Journal of Political Economy* 108,1292–1326.

[10] Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press, Princeton, NJ.).

[11] Christensen, Bent J., and Nagpurnanand R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.

[12] Cox, John C., Stephen A. Ross, and Mark Rubinstein, 1979, Option Pricing: A Simplified Approach, *Journal of Financial Economics* 7, 229–263.

[13] Dubofsky, David A., 1991, Volatility Increases Subsequent to NYSE and AMEX Stock Splits, *Journal of Finance* 46, 421–431.

[14] Ederington, Louis H., and Jae Ha Lee, 1996, The Creation and Resolution of Market Uncertainty: The Impact of Information Releases on Implied Volatility, *The Journal of Financial and Quantitative Analysis* 31, 513-539.

[15] Ellerman, A. Denny, and Paul L. Joskow, 2008, The European Union's Emissions Trading System in Perspective, Working paper, Pew Center on Global Climate Change.

[16] Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein, 1993, Risk Management: Coordinating Corporate Investment and Financing Policies, *Journal of Finance* 48, 1629–1658.

[17] Greenstone, Michael, 2002, The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures, *Journal of Political Economy* 110, 1175–1219.

[18] Hong, Harrison, and Jeremy C. Stein, 1999, A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets, *Journal of Finance* 54, 2143–2184.

[19] Hou, Kewei, and Tobias J. Moskowitz, 2005, Market Frictions, Price Delay, and the Cross-Section of Expected Returns, *Review of Financial Studies* 18, 981–1020.

[20] Jaraite, Jurate, Frank Convery, and Corrado Di Maria, 2009, Assessing the Transaction Costs of Firms in the EU ETS: Lessons from Ireland, Working paper, University College Dublin.

[21] Joskow, Paul L., Richard Schmalensee, and Elizabeth M. Baily, 1998, The Market for Sulfur Dioxide Emissions, *American Economic Review* 88, 669–685.

[22] Kothari, S. P., and Jerold B. Warner. 2007. Econometrics of Event Studies, in B. Espen Eckbo (ed.): *Handbook of Corporate Finance: Empirical Corporate Finance,* vol. 1, pp. 3–37. (Elsevier Science, Amsterdam, The Netherlands).

[23] Kruger, Joseph, Wallace E. Oates, and William A. Pizer, 2007, Decentralization in the EU Emissions Trading Scheme and Lessons for Global Policy, Discussion paper, Resources for the Future.

[24] Newell, Richard G., and Robert N. Stavins, 2003, Cost Heterogeneity and the Potential Savings from Market-Based Policies, *Journal of Regulatory Economics* 23, 43–59.

[25] *The New York Times*, 2006, Data Leaks Shake Up Carbon Trade, by Heather Timmons, May 16, C1.

[26] Peng, Lin, 2005, Learning with Information Capacity Constraints, *Journal of Financial and Quantitative Analysis* 40, 307–329.

[27] Sijm Jos, Karsten Neuhoff, and Yihsu Chen, 2006, CO_2 cost pass-through and windfall profits in the power sector, *Climate Policy* 6, 49–72.

[28] Trotignon, Raphaël, and Anaïs Delbosc, 2008, Allowance Trading Pattersn During the EU ETS Trial Period: What Does the CITL Reveal? Climate Report No. 13, Caisse des Depots.

[29] Veith, Stefan, Jörg G. Werner, and Jochen Zimmermann, 2009, Capital Market Response of Emission Right Returns: Evidence from the European Power Sector, *Energy Economics* 31, 605–613.

Chapter 2

Outstanding Debt and the Household Portfolio

I hope that after I die, people will say of me: 'Wow, that guy sure owed me a lot of money!'

–Jack Handey, Saturday Night Live

Introduction

This paper examines the effect of debt on the household portfolio. Whereas standard portfolio choice models focus primarily on the asset side of the household balance sheet, we examine the effects of liabilities on investment decisions. Throughout the life cycle, many households accumulate debt from a variety of sources including mortgages, student loans, and consumer debt. Retirement of this debt offers households a return equal to the interest rate on their loan, which is almost always greater than the return to investing in the risk-free asset. Higher interest rates on household debt thereby decrease the expected excess return to investing in the risky asset and reduce the benefit to equity participation. By developing a framework for how liabilities are incorporated into the financial decisions of households, this paper provides an understanding of the effects of the current financial crisis on future asset demand. In particular, our analysis offers insights into the long-term effects of current efforts by the Federal Reserve Bank and Government Sponsored Enterprises (GSEs) to refinance a large number of households into mortgages with historically low interest rates.

Our analysis is particularly relevant given the amount of debt currently held by U.S. households and the role that debt has played in instigating the current financial crisis. Between 1985 and the third quarter of 2008, the inflation-adjusted level of household debt increased from $1.4 trillion to $16.8 trillion, a twelve-fold increase.

Table 1
Characteristics of debt held by U.S. households

	% of	Median	Mean	Annual After-Tax Interest Rates (%)		
Type of Debt	Households	Balance ($)	Balance ($)	25 Percentile	Median	75 Percentile
Mortgage	40.8	89,000	115,650	4.6	5.4	6.5
Credit Cards with Balance	32.6	2,800	5,410	9.0	13.9	18.0
Home Equity Line of Credit	4.8	18,549	31,641	3.8	5.4	6.8
Other Home Equity Loans	1.7	13,195	32,006	4.7	6.0	7.4
Vehicle Loans	33.1	9,094	11,462			
Other Installment Loans	14.6	1,900	9,551			
Student Loans	11.4	6,000	12,809			

This table reports the debt characteristics of all households from the triennial 1989-2004 Survey of Consumer Finances (SCF). Estimates for credit card rates are from the years 1995-2004 because credit card rates were not surveyed before 1995. The interest rates for vehicle loans, other installment loans, and student loans are not surveyed by the SCF. All estimates are weighted using the population weights from the SCF. All dollar amounts are reported in 2004 dollars. See Appendix A.3 for a detailed explanation of how after-tax interest rates are constructed.

During the same period, aggregate household disposable income increased by less than seven-fold.[1] Table 1 presents the most common types of household debt, along with the mean and median balances and, when available, their after-tax annual interest rates.[2] It shows that mortgage debt is the most common form of debt and the largest liability on the household balance sheet. Mortgage debt has also grown at a faster rate than disposable income, home prices, and total assets over the past 25 years. Figure 1 uses data from the Federal Reserve Flow of Funds to show how ratios of mortgage debt to household income and assets have changed over time. Each of the three ratios shown increased by over 50% between 1985 and 2008, and in 2006 the level of household mortgage debt grew to exceed the aggregate disposable income of U.S. households.

In this study, we test implications from a simple portfolio choice model in the presence of mortgage debt. By substituting individual mortgage rates for the risk-free rate of return, our theoretical framework produces testable predictions for stock and bond participation, as well as optimal portfolio shares. We use the Survey of Consumer Finances (SCF) from 1989 to 2004 to test these predictions, and find that households with mortgage debt are 10% less likely to own stocks and 37% less likely to own bonds compared to similar households with no outstanding mortgage debt. To show that our results are not driven by irrational behavior amongst a subset of households, we construct two proxy variables for financial naivete. We then incorporate additional forms of debt into our analysis and estimate the welfare

[1]Numbers are from the December 2008 Federal Reserve Flow of Funds. Disposable household income is gross income less taxes. All amounts are in 2004 dollars.

[2]See Appendix A.3 for a detailed explanation of how we construct after-tax interest rates.

costs of sub-optimal portfolio composition amongst households that borrow at high interest rates and simultaneously hold low yielding investment assets. The majority of households behave in accordance with our model and for most of those who do not, the costs incurred are quite low.

Figure 1

Ratios of U.S. mortgage debt to income and assets 1985-2008Q3

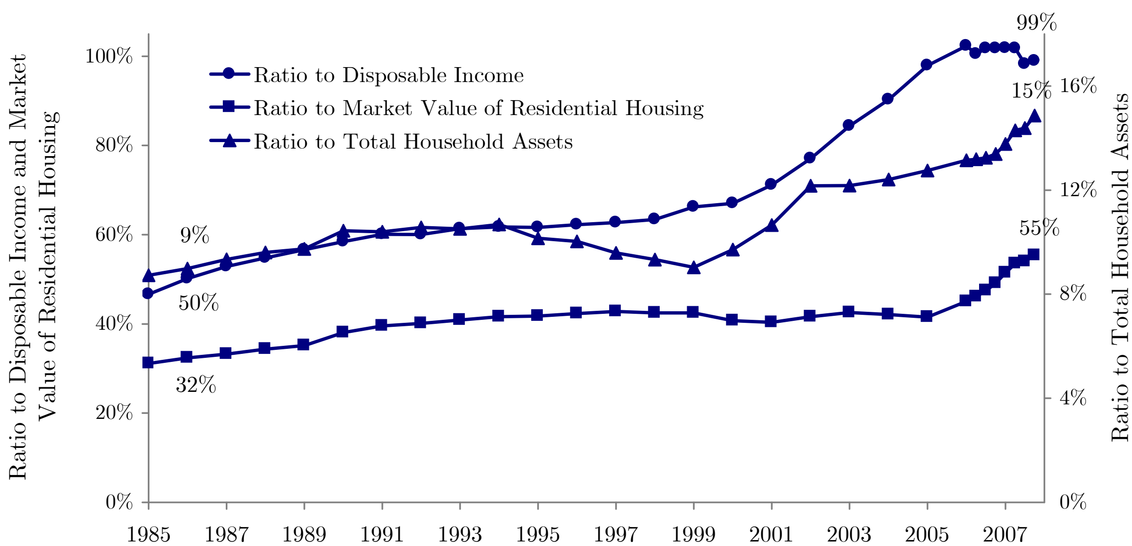

This figure shows changes in household mortgage debt relative to income and assets over time. The data are from the December 2008 Federal Reserve Flow of Funds. Mortgage debt includes home equity lines of credit and other forms of home equity borrowing. Household disposable income is gross income less taxes.

Our work contributes to several areas of the household finance literature. It adds to the existing explanations for limited equity market participation. Financial theory predicts that all households will take some amount of risk as long as it offers a positive expected return. Yet as Campbell notes in his 2006 Presidential Address to the American Finance Association, "limited participation [even among quite wealthy households] poses a significant challenge to financial theory and is one of the main stylized facts of household finance." One contribution of this extensive literature that is similar to our analysis is that of Heaton and Lucas (2000).[3] The authors use the SCF to show that entrepreneurial risk is an important determinant of portfolio choice. Their work provides a plausible explanation for limited participation, particularly amongst wealthier households.

This paper also adds to a relatively sparse literature examining the effects of housing on the household portfolio. Our portfolio choice setup has similarities to the

[3]See Section 2 of Campbell (2006) for an overview of the literature on limited participation.

mean-variance optimization framework used by Flavin and Yamashita (2002), though our findings are very different. We find that mortgage debt decreases the benefit to equity participation, whereas the estimates of asset returns used by Flavin and Yamashita suggest that it is optimal for households to subsidize stock ownership using mortgage debt. Our finding that mortgage debt reduces stock ownership is similar to the work of Chetty and Szeidl (2009), although both the economic channels and the empirical specifications are different. Their paper studies the effect of housing and mortgage debt on household portfolios, focusing on home price risk and the consumption commitment of housing as the underlying mechanisms. Their central finding is that an exogenous increase in the *amount* of mortgage debt reduces the portfolio share of stocks. In contrast, we investigate how the *presence* of mortgage debt and the mortgage interest rate influence financial portfolios through the debt retirement channel.

Other papers in this literature, such as Grossman and Laroque (1991), Fratantoni (2001), Cocco (2005), and Yao and Zhang (2005), focus on portfolio choice over the life-cycle. In contrast to our study, these papers solve well-defined dynamic optimization problems that include various elements of housing. They calibrate the parameters of their model and use results from numerical simulations to explain relationships observed in empirical data.

The two papers from this literature that are most relevant for our study are Cocco (2005) and Yao and Zhang (2005). Cocco (2005) finds that investment in housing keeps liquid assets low early in the life-cycle and reduces the willingness of younger households to pay the fixed cost required for equity market participation. Yao and Zhang (2005) include the decision of whether to rent or own, which allows investors to separate their choice of housing consumption from their choice of housing investment. They find that when indifferent between renting and owning a home, investors that own a home hold a lower equity share in their total wealth and a higher equity share in their liquid wealth. Whereas these papers include the consumption value and price risk associated with housing as important determinants of portfolio choice, our paper evaluates the impact of mortgage debt on the household portfolio. By using micro data to study how variation in debt interest rates affects household portfolio composition, our paper presents novel empirical results and identifies a mechanism for non-participation that is not addressed in previous work.

Finally, our paper relates to existing work that studies the effects of a wedge between borrowing and lending rates. Davis, Kubler, and Willen (2006) use a life-cycle model of portfolio choice to analyze the effects of such a wedge on the demand for equity. They conclude that the demand for equity is minimized when the rate at which households can borrow is equal to the expected return on equity. Zinman (2007a) uses household-level data from the SCF to determine whether households tend to borrow at high rates while simultaneously lending at lower rates. We perform a similar analysis towards the end of our paper to determine the prevalence of non-optimal bond market participation and its associated welfare costs.

The remainder of this paper is organized as follows. In Section 2.1, we discuss implications of a portfolio choice model, which includes the option to retire mortgage debt. Section 2.2 describes the data sources used in our analysis. Section 2.3 presents the results of regression analysis used to test the predictions of our model. In Section 2.4, we incorporate additional forms of household debt and calculate the costs of non-optimal bond market participation for households that appear to borrow at high interest rates and simultaneously lend at low interest rates. The final section concludes.

2.1 Portfolio Choice and Mortgage Debt

We consider a household with an outstanding fixed-rate mortgage facing the decision of how to allocate disposable income between stocks, bonds, and repayment of mortgage debt. We begin with a frictionless world with no liquidity issues, taxes or other impediments to optimizing between investments, and we later incorporate such factors into our empirical analysis.[4] The decision of the household to pay an amount λ in excess of the required mortgage payment results in a return characterized by the cash flows shown in Figure 2. In the absence of any early payment (i.e., $\lambda = 0$), the future required monthly payments, m, amortize the remaining principal of the mortgage to zero by date T. Yet when a portion of the debt is retired early, the remaining principal is immediately reduced by λ, and the required monthly payments now amortize the principal to zero at time T'. The household realizes its cash flows in the form of foregone monthly payments between T' and T or upon the sale of the house or refinancing of the mortgage contract prior to date T. Thus an early payment in the amount of λ is equivalent to the purchase of a bond with face value λ, a yield-to-maturity equal to the mortgage interest rate r, and a duration D, where $T' < D < T$.[5]

The equivalence of mortgage debt retirement to bond investment has implications for asset market participation and optimal portfolio allocation. We consider a simple one-period portfolio choice model in which households with power utility and constant relative risk aversion choose the proportion of their wealth to allocate to the risky asset. As in the standard theory of portfolio choice, households without a mortgage will choose an equity share equal to the ratio of the expected excess return to the price and quantity of risk. On the other hand, households with the option to retire their mortgage debt early in exchange for the payoff structure shown in Figure 2 will fall into one of three categories depending on their mortgage interest rate.

Households with a mortgage interest rate below the return on risk-free bonds have no incentive to retire their debt ahead of schedule and thus hold a portfolio identical to that of households without a mortgage. Households with a mortgage rate between

[4] We define and discuss the assumptions of our model in detail in Appendix A.1.

[5] A formal derivation of this result is available in an online appendix to this paper.

Figure 2

Timeline of cash flows from repayment of mortgage debt

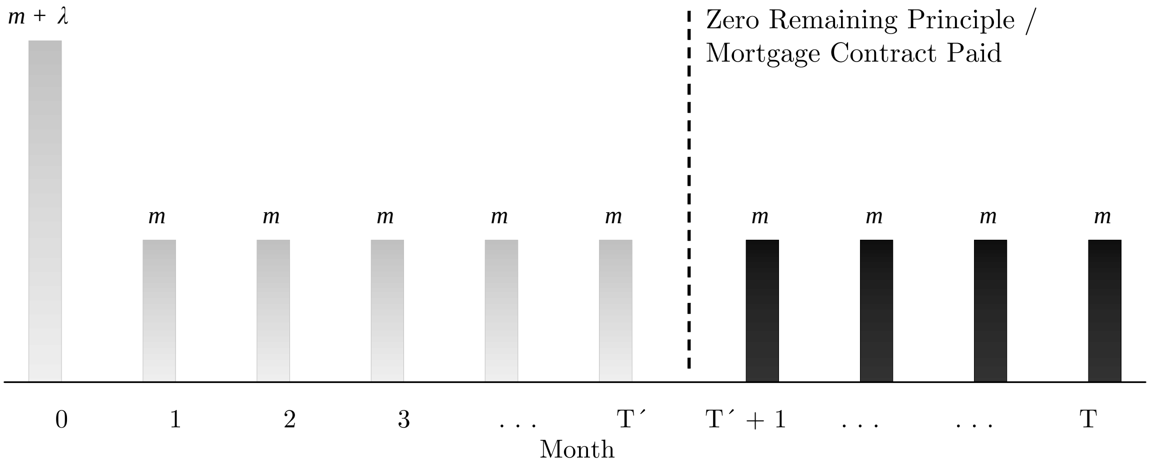

This figure shows the cash flows that result from the early repayment of mortgage debt with an original term of T months. An early payment of delta immediately reduces the remaining principle and the remaining monthly mortgage payments amortize the principle to zero at time T'<T. Cash-flows are realized in the form of foregone monthly mortgage payments between time T' and T or upon prepayment of mortgage principle through sale or refinancing.

the risk-free rate and the expected market return will invest a smaller share of their wealth in the risky asset because of the diminished expected excess return it offers. Finally, those with a mortgage rate greater than the expected market return will not allocate any of their wealth to the risky asset. These are households that in the absence of any short selling constraints would find it optimal to short the stock market in order to repay their mortgage debt. Thus the optimal share of wealth invested in the risky asset by household i can be summarized as:

$$\alpha^i = \begin{cases} \frac{\mathbb{E}[\tilde{R}]-R_f}{\gamma\sigma^2} & \text{if no mortgage} \\ \alpha^i_d & \text{if mortgage} \end{cases} \tag{2.1}$$

where:

$$\alpha^i_d = \begin{cases} \frac{\mathbb{E}[\tilde{R}]-R_f}{\gamma\sigma^2} & \text{if } R^i_d < R_f \\ \frac{\mathbb{E}[\tilde{R}]-R^i_d}{\gamma\sigma^2} & \text{if } R_f \leq R^i_d \leq \mathbb{E}[\tilde{R}] \\ 0 & \text{if } \mathbb{E}[\tilde{R}] < R^i_d \end{cases} \tag{2.2}$$

and R_f is the net return on the risk-free asset, R^i_d is the mortgage interest rate of household i, and $\tilde{R}$ is the net return on the risky asset where $\log(1+\tilde{R}) \sim N(\mu, \sigma^2)$.

2.1.1 Implications for asset market participation

Stock market participation

The effect of mortgage debt on stock ownership is best captured by quantifying the benefit to equity participation. This section uses Equations 2.1 and 2.2 along with household-level data on mortgage interest rates to provide a sense of how mortgage debt affects the payoff to stock ownership. The benefit to equity participation can be quantified as:

$$\alpha W \left(R^{ce} - R_f\right), \tag{2.3}$$

where W represents initial wealth and R^{ce} is the net certainty equivalent return.[6] Equation 2.3 represents the optimal dollar amount of wealth invested in the risky asset multiplied by the risk-adjusted excess return earned on each dollar. For households with a mortgage, our model reduces the benefit to participation to:

$$\alpha_d^i W \left(R^{ce} - R_d^i\right). \tag{2.4}$$

The benefit of stock ownership is lower for households with mortgages and sufficiently high mortgage rates because they have less wealth invested in the risky asset ($\alpha_d^i W \leq \alpha W$) and earn less on each dollar invested in the risky asset ($R^{ce} - R_d^i < R^{ce} - R_f$). Therefore these households will be less likely to own stocks. This implication is stated formally in Proposition 1.

Proposition 1 *The effects of mortgage debt on stock market participation are:*

1. *Having a mortgage should decrease the probability of equity market participation.*
2. *Conditional on having a mortgage, a higher mortgage interest rate should decrease the probability of equity market participation.*

To get a sense of the degree to which mortgage debt reduces the benefit to equity participation, we construct a graphical representation similar to the one used by Vissing-Jorgensen (2002) to study the effect of information and transaction costs on the benefit to participation. Figure 3 shows the annual benefit to equity participation for non-participants using household-level data from the SCF. The vertical axis represents the annual dollar benefit for non-participants calculated using Equations 2.3 and 2.4. For households without a mortgage, we follow Vissing-Jorgensen and assume a value of 0.04 for $R^{ce} - R_f$. This value is approximated using an historical tax-adjusted annual equity premium of 5.6% and an arbitrary risk-adjustment down

[6] The certainty equivalent return is the rate of return that satisfies the condition $\mathbb{E}U[W(1+\tilde{R})] = U\left[W\left(1 + R^{ce}\right)\right]$. In other words, the investor is indifferent between investing W in a risky asset with stochastic return $\tilde{R}$ and investing it in a risk-free asset with return R^{ce}.

Figure 3
Benefit to stock market participation for non-participants

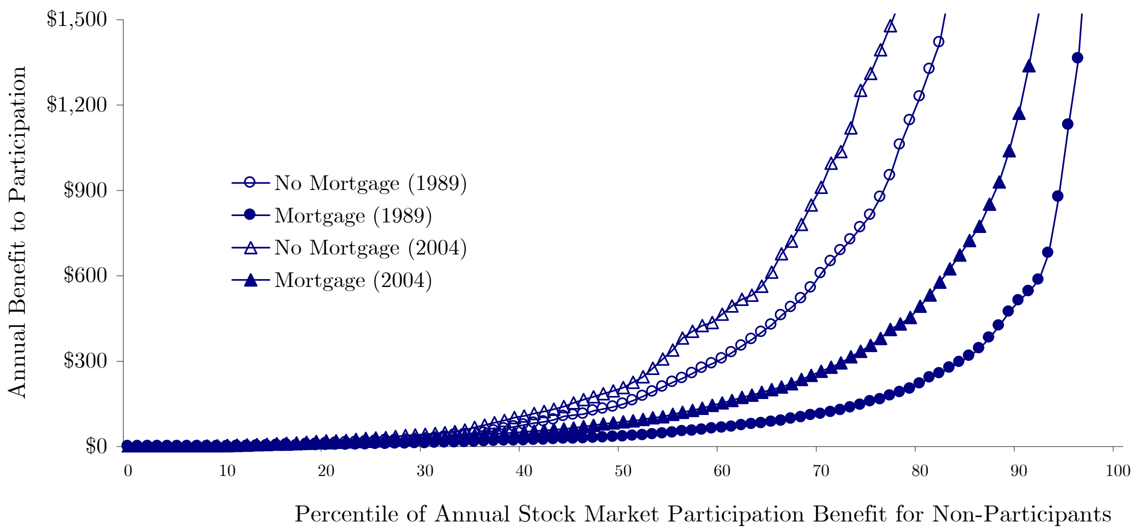

This figure shows the annual benefit to stock market participation for non-participants. The benefits are estimated using Equations 3 and 4 and data from the triennial 1989-2004 Survey of Consumer Finances (SCF). The benefits are weighted using population weights from the SCF. The sample is limited to homeowners who do not live on a farm or in manufactured housing. Dollar amounts are reported in 2004 dollars.

to 4%.[7] We also set the value of α equal to the mean equity share of participants without a mortgage in a given year. For W, we substitute in liquid wealth defined as the sum of checking, savings, and money market accounts, certificates of deposit, savings bonds, and all stocks and bonds held outside of retirement accounts. Like Vissing-Jorgensen, we use liquid wealth, although in the portfolio choice model the optimal equity share is a fraction of total wealth. We discuss the difference between these two measures of wealth in Section 2.1.2.

The curves that represent households with mortgages are calculated using Equation 2.4. We calculate α_d as the mean equity share of participants with mortgages in a given year. For $R^{ce} - R_d^i$, we use an annual equity premium equal to the after-tax return on equity, as calculated by Poterba (2002), minus the after-tax mortgage rate of household i. We then multiply this premium by 0.71, which is the factor used by Vissing-Jorgensen to adjust the equity premium for risk. We also topcode the value of $R^{ce} - R_d^i$ at 0.04 to account for implausibly low mortgage rates and set a lower

[7] The historical tax-adjusted annual equity premium is approximated using an historical equity premium of 7% and a tax rate of 20%. Though Vissing-Jorgensen does not explicitly name a source for these figures, the approximated value of 5.6% closely matches the 5.77% reported by Poterba (2002) using data from 1926 to 1996.

limit at zero to enforce our assumption of no short positions in equity.

The horizontal axis of Figure 3 represents the percentile of benefits from the cross-sectional distribution of non-participants and thus gives the dollar amount of annual costs necessary to explain the decision to forego equity participation for different percentages of non-participants. The figure shows that in 1989 an annual participation cost of $300 is sufficient to explain non-participation in equity markets for 60% of households without a mortgage. Yet that same cost of $300 can explain 85% of non-participation amongst households with mortgages. The figure also shows that the difference in benefits was greater in 1989 than it was in 2004, reflecting the fact that mortgage rates were much higher in 1989 compared to 2004.

Bond market participation

The effect of mortgage debt on bond ownership is more straightforward. As long as the mortgage interest rate is greater than the interest rate on risk-free bonds, households will forego bond ownership in favor of paying down mortgage debt. This result is in contrast to stock ownership in which households with a mortgage and a mortgage rate greater than the risk-free rate may still find it optimal to invest in the risky asset, albeit in smaller amounts. Therefore, we should find that having a mortgage has a greater effect on bond ownership than it does on stock ownership. This implication is stated formally in Proposition 2.

Proposition 2 *Having a mortgage should decrease the probability of bond market participation by an amount greater than the decline in the probability of equity market participation.*

Cocco (2005) and Yao and Zhang (2005) also find that debt repayment and bonds serve as substitute assets; with costless refinancing, investors will never hold bonds and mortgage debt simultaneously. By introducing a wedge between borrowing and lending rates, Davis, Kubler, and Willen (2006) also find that households will rarely hold bonds. Our analysis uses household-level debt interest rates to provide additional insight on bond market participation in two ways. First, we explicitly test the degree to which the effect of debt on bond market participation is different from the effect of debt on stock market participation. Second, we estimate the costs of non-optimal investment for households whose behavior is not consistent with the predictions of these models in that they simultaneously hold low yielding bonds and high interest rate debt.

2.1.2 Implications for portfolio shares

We now discuss implications for optimal portfolio allocation. In the basic portfolio choice model described at the beginning of Section 1, W represents total wealth, which includes discounted future labor income. Therefore the optimal equity shares used for the participation decision are shares of total wealth and not liquid wealth. In fact, the effects of mortgage debt on the portfolio share of liquid wealth are quite different. Models that examine portfolio choice in the presence of non-tradable labor income, including Heaton and Lucas (1997) and Viceira (2001), find that equity shares ought to decline throughout the life cycle. This is because households initially choose an optimal share of wealth to invest in the risky asset while considering their future labor income as a safe asset. As the life cycle progresses and future labor income is realized, it is substituted with bonds, which are a tradeable form of safe assets.

Empirically however, the opposite is true; equity shares tend to increase over the life cycle. Cocco (2005) uses a life cycle model with housing included in the utility function to explain this empirical finding. He finds that younger households are highly invested in housing and thus have limited wealth to invest in stocks, which reduces the benefit to stock ownership. His conclusion is similar to our model's prediction that mortgage debt lowers the return to equity investment. Furthermore, our model predicts that while households with a mortgage are less likely to hold stock, conditional on participation their equity share of liquid wealth is actually higher. This is because households chose to pay down mortgage debt rather than hold bonds, assuming their mortgage rate is greater than the risk-free rate on bonds. This prediction of our model is consistent with the Davis, Kubler, and Willen (2006) finding that "when households cannot borrow at the risk-free rate they invest nothing in bonds and equity holdings equal financial wealth."

Figure 4 shows the evolution of optimal portfolio shares over the life cycle as predicted by our theoretical framework. The horizontal axis represents time and the vertical axis represents the share of total lifetime income. We assume that all households are born with only future labor income and use a mortgage to finance the purchase of their home. As households realize their labor income, they allocate it between the risky asset, which is stocks, and the safe asset, which is initially repayment of mortgage debt. So long as the household has mortgage debt outstanding, it prefers to repay it rather than invest in bonds that offer a lower rate of return. After the mortgage is paid off, the household begins to invest its realized labor income in bonds.

The evolution shown in Figure 4 offers predictions for the optimal equity share of liquid wealth for the categories shown in Equations 1 and 2. Households without a mortgage or with a sufficiently low mortgage interest rate allocate a portion of their wealth to bonds and therefore have an equity share of liquid wealth that is less than one. Households with a mortgage rate between the risk-free rate and the expected market return hold all of their liquid wealth in stocks and have an equity share equal to one. Households with an interest rate greater than the expected return on the risky asset will not hold any stocks. These implications are stated formally in Proposition 3.

Proposition 3 *The effects of mortgage debt on the optimal portfolio shares of liquid*

Figure 4
Life-cycle evolution of portfolio shares

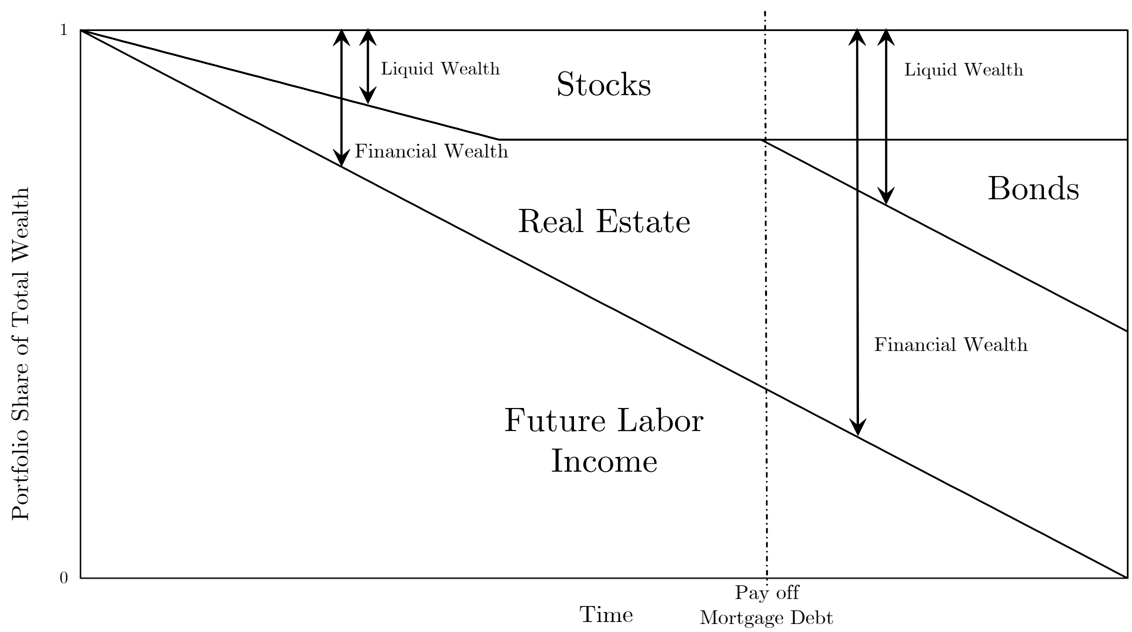

wealth are:

1. *Conditional on equity participation, having a mortgage should increase the equity share of liquid wealth.*

2. *Conditional on having a mortgage with an interest rate greater than the risk-free rate of return, a higher mortgage interest rate should decrease the equity share of liquid wealth.*

Table 2 shows the mean and median equity shares of liquid wealth from the SCF based on the mortgage classes in Equations 2.1 and 2.2. As predicted, the average equity share is smaller for households without a mortgage relative to households with a mortgage rate between the risk-free rate and the expected market return. Also, households with a mortgage rate greater than the expected market return have a median equity share equal to zero. However, households with a mortgage rate lower than the risk-free rate have a greater equity share than that of any other group. This suggests the presence of other factors that we have yet to consider. Households with and without mortgages differ in many important ways, as do households with high mortgage rates and those with low rates. Many of these differences, such as age and

Table 2

Portfolio equity share by mortgage class

Mortgage Class	% of Total	Equity Share of Liquid Wealth	
		Median (%)	Mean (%)
No Mortgage	36.8	13.3	36.4
Mortgage			
Mort Rate < Risk-Free	5.4	82.0	56.9
Risk-Free ≤ Mort Rate ≤ Equity	57.1	50.0	47.8
Equity < Mort Rate	0.7	0.0	30.7

This table uses data from the triennial 1989-2004 Survey of Consumer Finances (SCF) to report average equity shares of liquid wealth based on four separate mortgage classifications. All estimates are weighted using population weights from the SCF. The sample is limited to homeowners who do not live on a farm or in manufactured housing. The risk-free rate is 3.39between 1926 and 1996. The return on equity is 9.16after-tax return on stocks between 1926 and 1996. Both of these estimates are from Poterba (2002). The mortgage rate is adjusted for taxes, as described in Appendix A.2.

wealth, affect both asset market participation and optimal portfolio shares.[8] For this reason, we turn to regression analysis to isolate the ceteris paribus effects of mortgage debt on asset ownership and the portfolio share.

2.2 Data Sources

Our primary data source is the Survey of Consumer Finances (SCF) which is conducted on a triennial basis by the Federal Reserve Board. The SCF is a dual-frame, cross-sectional survey in which two-thirds of respondents comprise a representative sample of U.S. households based on geographic and income information and the remainder of respondents are oversampled from wealthy households. It is unmatched in its level of detail with regard to both household asset allocation and debt obligations. The dataset also includes demographic and survey information for each of the roughly 4,000 households sampled every three years between 1989 and 2004. Sample weights are used to correct for survey non-response and allow us to approximate a representative sample of U.S. households in each year.[9] A multiple imputation methodology is used to decrease the sampling variance of the data.[10] Because we want to isolate the effects of mortgage debt on the household portfolio, we must compare households that

[8]Ameriks and Zeldes (2004) examine the investment advice provided by Vanguard, which varies widely according to age. Wachter and Yogo (2008) examine how portfolio shares vary with wealth.

[9]Kennickell and Woodburn (1999) provides a detailed description of how sample weights in the SCF are used to create a population-weighted sample of U.S. households.

[10]Each observation has five implicates constructed to lower the sampling variance that results from missing data. See Montalto and Sung (1996) for a detailed description of multiple imputation in the SCF.

Table 3
Summary statistics from the Survey of Consumer Finances

	Obs.	% of HHs	Median	Mean	Std. Dev.
Income	16,313	100	$53,438	$81,767	$210,437
Net Worth	16,313	100	$162,338	$498,568	$2,148,430
Stocks and Stock Mutual Funds (Non-Retirement)	8,268	33.8	$25,037	$177,000	$1,209,695
Bonds and Bond Funds (Non-Retirement)	8,821	45.1	$9,282	$71,173	$562,149
Mortgage Balance	9,833	64.0	$73,852	$95,570	$97,792
Mortgage Interest Rate	9,833	64.0	7.75%	7.93%	2.08%

Mortgage Characteristics	% of HHs	Financial Naivete	% of HHs
Adjustable Rate Mortgage (ARM)	9.8	Credit-Adjusted Rate is Too High	11.7
Refinanced Mortgage Contract	30.8	Does Not Know Mortgage Rate	6.2

Risk Preferences	% of HHs	Savings Horizon	% of HHs
Not Willing to Take Risks	37.9	Next Few Months	15.9
Average Risk	42.2	Next Year	12.1
Above-Average Risk	16.2	Next 2-5 Years	25.8
High Risk	3.6	Next 5-10 Years	28.9
		Longer Than 10 Years	17.3

This table reports selected summary statistics for the triennial 1989-2004 Survey of Consumer Finances (SCF). All estimates, including percentage of households, are weighted using population weights from the SCF. The sample is limited to homeowners who do not live on a farm or in manufactured housing. We truncate income and wealth at the 1st and 99th percentiles of the cross-sectional distribution to limit the influence of outliers. All dollar amounts are reported in 2004 dollars.

have paid off their mortgage to those who haven't. To obtain a more uniform sample of households who have had or currently have the opportunity to repay their mortgage debt, we limit our sample to homeowners whose primary residence is not a farm or in manufactured housing. Summary statistics of selected variables are reported in Table 3.

Using the SCF, we define *stocks* as the sum of any individual stocks and stock mutual funds held outside of retirement accounts. Our definition of *bonds* includes certificates of deposit, U.S. savings bonds, Treasury, municipal, corporate, foreign, and mortgage-backed bonds, cash-value of life insurance, and any other interest-bearing managed assets and trusts held outside of retirement accounts.[11] We focus on assets held outside of retirement accounts because households are free to reallocate these assets towards other investments, including debt retirement, without incurring the penalties imposed on the early liquidation of retirement accounts.

The SCF has several distinct advantages over other household-level datasets. Most importantly, it is unmatched in its detail of both the assets and liabilities that make up the household portfolio. This makes the SCF the only dataset that provides detailed and consistent information on both portfolio shares and the characteristics of mortgage debt. The high level of detail in a households' asset holdings allows us to distinguish between securities held within retirement accounts and those held in discretionary accounts. The SCF is also the only survey that contains information on the household's savings horizon and self-described tolerance for investment risk. These data allow us to control for both differences in risk aversion across households and the role of liquidity in a household's financial decision making. As described in Section 2.4, households may choose to forego repayment of mortgage debt despite the high return it offers if they need liquid funds for upcoming expenses. Controlling for a household's savings horizon allows us to separate the asset allocation decision from the decision of how much to hold in liquid assets.

One shortcoming of the SCF is that a cross-sectional dataset does not allow us to observe changes to the portfolio composition of a particular household in the years following the full repayment of mortgage debt. To do so requires a panel dataset with detailed data on both portfolio composition and mortgage characteristics.[12] The Panel

[11] We characterize certificates of deposit and U.S. savings bonds as bonds rather than safe assets because these assets serve an investment purpose and are not liquid securities that households may use for daily transactions or to serve as a buffer against unforeseen expenses.

[12] The panel datasets containing information necessary for our study include the SCF 1983-1989 Panel Survey, the Tax Model Data maintained by the National Bureau of Economic Research, the Survey of Income and Program Participation (SIPP), and the Panel Survey of Income Dynamics (PSID). The SCF 1983-1989 Panel Survey is insufficient as it has a very low number of observations and variables that are coded inconsistently across years. The Tax Model data used in Heaton and Lucas (2000) are also insufficient, as the tax filings do not cleanly identify asset holdings or mortgage characteristics. The SIPP contains five separate panel surveys of households in 1984-1986, 1996, and 2001. These data provide a relatively detailed picture of household asset holdings but follow each household over a short interval of time thereby diminishing much of the benefit of a panel dataset.

Survey of Income Dynamics (PSID) provides the best panel data for our analysis. The dataset lacks some important variables including risk-tolerance, savings horizon, bond holdings, mortgage interest rate, and mortgage type (i.e., adjustable-rate versus fixed-rate). Yet it has a significant advantage over other panel datasets because it tracks households over an extended period of time. The variables necessary for our analysis are available in the Wealth Waves conducted in 1984, 1989, 1994 and every other year between 1999 and 2005. This allows us to track the evolution of the household portfolio for up to 21 years. We use the PSID panel data in Section 2.3 to examine whether households that pay off their mortgage are subsequently more likely to own stock.

2.3 Regression Analysis

2.3.1 Empirical model and results

In this section, we test the predictions of our model for asset market participation, as summarized in Propositions 1 and 2, and the predictions for portfolio shares, as summarized in Proposition 3. To isolate the effect of mortgage debt and mortgage interest rates on asset market participation, we model the unobservable benefit to household i from owning stocks or bonds as:

$$y_i^* = \mathbf{x}_i'\boldsymbol{\beta} + \delta\theta_i + \varepsilon_i \ , \tag{2.5}$$

where $\mathbf{x}_i$ is a vector containing household characteristics relevant for asset market participation, such as age, education, and wealth, θ_i contains the mortgage variable of interest for household i, and ε_i is an error term drawn from a logistic distribution with mean zero and known variance $\pi^2/3$.[13] Though we do not observe y_i^* directly, we do observe whether or not a household owns stocks or bonds. We characterize this variable as y_i, where:

$$y_i = \begin{cases} 1 & \text{if } y_i^* > 0 \\ 0 & \text{if } y_i^* \leq 0 \end{cases} \tag{2.6}$$

Then it is straightforward to show that:

$$\Pr\left(y_i = 1|\mathbf{x}_i, \theta_i\right) = G\left(\mathbf{x}_i'\boldsymbol{\beta} + \delta\theta_i\right) \ , \tag{2.7}$$

where $G(\cdot)$ is the cumulative distribution function of ε. We then use maximum likelihood to estimate the effect of mortgage debt and the mortgage interest rate on the probability of stock or bond ownership, $\partial \Pr\left(y_i = 1|\mathbf{x}_i, \theta_i\right)/\partial\theta_i$.

We also use two types of regressions to estimate the effects of mortgage debt on the portfolio share. In estimating the effects of the mortgage interest rate on the

[13]The assumption of a known variance is a simple normalization. See Chapter 21 in Greene (2003).

equity share, we use an unconditional Tobit model with a mass point at zero because many households with mortgages do not participate in equity markets and thus have a portfolio share equal to zero. We also use an ordinary least squares specification to estimate the effects of mortgage debt on portfolio shares amongst households who participate in equity markets.

Stock market participation

Table 4 shows the results of unweighted logit regressions of stock and bond market participation on mortgage characteristics. The regressions include year fixed-effects and control for household characteristics, such as demographics, education, age, income, wealth, risk tolerance, savings horizon, and private business ownership. The reference household is headed by a white male between the ages of 35 and 45, with a college degree, an average risk tolerance, a savings horizon of less than one year, and income and wealth equal to their respective sample means. The probability estimates are computed using the odds ratio associated with each logit coefficient estimate. They show that our control variables have the expected marginal effects on asset market participation.

Column 1 of Table 4 shows that, all else equal, having a mortgage decreases the probability of stock ownership by 21%. For the reference household this corresponds to a decrease in the probability of stock market participation from 42.7% to 33.7%, a decline of 9 percentage points. Column 2 estimates the effect of the mortgage interest rate on participation. It shows that conditional on having a mortgage, an increase in the annual mortgage interest rate of one standard deviation—or 2 percentage points—decreases the probability of stock market participation by 14.2%. This suggests that, as our model predicts, the mortgage interest rate is the mechanism that reduces the probability of equity participation. These findings are consistent with the predictions of Proposition 1 and are statistically significant at the 1% level.

Bond market participation

Column 3 of Table 4 shows that mortgage debt has an even greater negative impact on bond market participation. Having a mortgage decreases the likelihood of bond ownership by 40.1%. To test whether this decline is significantly greater than the 21% decline in the probability of stock market participation, we use the Seemingly Unrelated Regressions (SUR) framework of Zellner (1962).[14] We write a two-equation system for the benefit to stock and bond ownership as:

$$y_{is}^* = \mathbf{x}_i' \boldsymbol{\beta}_s + \delta_s \theta_i + \varepsilon_{is} \tag{2.8}$$

$$y_{ib}^* = \mathbf{x}_i' \boldsymbol{\beta}_b + \delta_b \theta_i + \varepsilon_{ib} \tag{2.9}$$

[14] We thank an anonymous referee for suggesting such a test.

Table 4

Equity and bond participation logit regressions

	Household Owns Equities				Household Owns Bonds			
	(1)		(2)		(3)		(4)	
	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates
Reference Household	-0.2948***	42.7%	-0.1759	45.6%	-0.3137***	42.2%	-0.5349***	36.9%
	[0.1129]		[0.1616]		[0.1016]		[0.1479]	
Household with Mortgage	-0.2360***	-21.0%			-0.5119***	-40.1%		
	[0.0478]				[0.0439]			
Mortgage Interest Rate			-0.0775***	-14.2%			-0.0606***	-11.3%
			[0.0157]				[0.0142]	
ARM Mortgage	-0.0615	-6.0%	-0.0785	-7.5%	-0.1279**	-12.0%	-0.1476***	-13.7%
	[0.0621]		[0.0620]		[0.0562]		[0.0566]	
Have Private Business	0.2412***	27.3%	0.2258***	25.3%	-0.0482	-4.7%	-0.1021**	-9.7%
	[0.0440]		[0.0537]		[0.0415]		[0.0508]	
Income (in $ millions)	0.5318***	26.0%	0.7394***	24.3%	0.2765***	8.7%	0.3706***	13.9%
	[0.0740]		[0.1196]		[0.0510]		[0.0754]	
Income^2 (in $ millions)	-0.0419***		-0.0634***		-0.0245***		-0.0294***	
	[0.0070]		[0.0112]		[0.0049]		[0.0075]	
Net Worth (in $ millions)	0.0290***	9.7%	0.0280***	3.7%	0.0131***	6.2%	0.0135**	0.2%
	[0.0048]		[0.0078]		[0.0036]		[0.0057]	
Net Worth^2 (in $ millions)	-0.0002***		-0.0002***		-0.0001***		-0.0001***	
	[0.0000]		[0.0001]		[0.0000]		[0.0000]	

	Household Owns Equities				Household Owns Bonds			
	(1)		(2)		(3)		(4)	
	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates
Age Under 35	-0.3133***	-26.9%	-0.2571***	-22.7%	-0.2467***	-21.9%	-0.1803**	-16.5%
	[0.0727]		[0.0770]		[0.0672]		[0.0711]	
Age 45-54	0.1796***	19.7%	0.1633***	17.7%	-0.0026	-0.3%	-0.0202	-2.0%
	[0.0542]		[0.0602]		[0.0506]		[0.0560]	
Age 55-64	0.3975***	48.8%	0.3305***	39.2%	0.0955	10.0%	0.0337	3.4%
	[0.0640]		[0.0766]		[0.0588]		[0.0708]	
Age Over 65	0.6702***	95.5%	0.4000***	49.2%	0.5351***	70.8%	0.2945***	34.2%
	[0.0714]		[0.1023]		[0.0653]		[0.0923]	
Educc - Grade School	-1.6006***	-79.8%	-1.3542***	-74.2%	-1.1240***	-67.5%	-1.1755***	-69.1%
	[0.0838]		[0.1328]		[0.0681]		[0.1117]	
Educc - High School	-0.9190***	-60.1%	-0.8633***	-57.8%	-0.4895***	-38.7%	-0.5098***	-39.9%
	[0.0489]		[0.0633]		[0.0454]		[0.0578]	
Educ. - Some College	-0.5516***	-42.4%	-0.5425***	-41.9%	-0.3128***	-26.9%	-0.3904***	-32.3%
	[0.0522]		[0.0643]		[0.0494]		[0.0609]	
Risk Pref. None	-1.1896***	-69.6%	-1.1366***	-67.9%	-0.5030***	-39.5%	-0.4077***	-33.5%
	[0.0505]		[0.0696]		[0.0446]		[0.0592]	
Risk Pref. > Average	0.4225***	52.6%	0.4474***	56.4%	-0.0910**	-8.7%	-0.0184	-1.8%
	[0.0483]		[0.0559]		[0.0450]		[0.0534]	
Risk Pref. High	0.0566	5.8%	0.0976	10.2%	-0.4813***	-38.2%	-0.4156***	-34.0%
	[0.0848]		[0.1018]		[0.0789]		[0.0970]	

	Household Owns Equities				Household Owns Bonds			
	(1)		(2)		(3)		(4)	
	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates
Saving - 1 Year	0.1175	12.5%	0.1882*	20.7%	0.1688**	18.4%	0.1153	12.2%
	[0.0833]		[0.1083]		[0.0726]		[0.0964]	
Saving - Few Years	0.3581***	43.1%	0.4044***	49.8%	0.3213***	37.9%	0.2718***	31.2%
	[0.0689]		[0.0901]		[0.0607]		[0.0810]	
Saving - 5-10 Years	0.5063***	65.9%	0.5759***	77.9%	0.4378***	54.9%	0.4065***	50.2%
	[0.0667]		[0.0865]		[0.0593]		[0.0776]	
Saving - >10 Years	0.6165***	85.3%	0.6374***	89.2%	0.6249***	86.8%	0.5799***	78.6%
	[0.0710]		[0.0905]		[0.0637]		[0.0816]	
Demographic Controls	Yes		Yes		Yes		Yes	
Year Fixed Effects	Yes		Yes		Yes		Yes	
Observations	16,313		9,833		16,313		9,833	
Pseudo R-Squared	0.24		0.20		0.11		0.09	
Mean Dep. Var.	0.507		0.477		0.541		0.486	

This table contains the results of logit regressions of equity and bond ownership on household characteristics and mortgage debt using data from the triennial Survey of Consumer Finances (SCF). The sample is limited to homeowners who do not live on a farm or in manufactured housing. Columns 2 and 4 are restricted to households with a mortgage. We truncate income and wealth at the 1st and 99th percentiles of the cross-sectional distribution to limit the influence of outliers. All regressions control for gender, race, marital status, and children. Robust standard errors have been adjusted for variation between implicates and are reported below the coefficients in brackets. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***. The reference household is headed by an unmarried, 35-44-year-old white male with a college degree, average risk preference, and income and wealth equal to their respective sample means. The "Probability Estimates" column reports the probability of stock or bond ownership for the reference household and the change in these probabilities associated with a unit increase in a binary variable and a one standard deviation increase in a continuous variable. Our results are robust to using decile dummies and splines to control for income and wealth.

where the subscripts s and b denote stocks and bonds, respectively, and θ_i is a dummy variable equal to 1 if household i has a mortgage and 0 otherwise. For testing hypotheses across equations, the SUR system in Equations 2.8 and 2.9 allows for non-zero covariance between the error terms ε_{is} and ε_{ib} for household i. However, the use of standard inference theory requires the additional condition of multinormality of the dependent variables, which makes the SUR framework ideal for Ordinary Least Squares (OLS) and probit regressions, but not logit regressions.[15] To overcome this limitation, we employ a regression specification designed to provide a 'sandwich' estimate of the covariance matrix of two logit regressions within a single estimating equation. This method, first proposed by Vella (1992), mimics estimation of the covariance matrix in a SUR system. It allows us to test the null hypothesis that the effects of mortgage debt on stock and bond ownership are statistically equivalent while maintaining the assumption that our error terms follow a logistic distribution.

To implement the 'sandwich' method, we first stack our stock and bond regression data sets on top of one another to obtain a single vector $\mathbf{y}^{sb}$, which contains y_{is} followed by y_{ib} for each household i.[16] Our regression then takes the functional form:

$$\Pr(y_i^{sb} = 1|\mathbf{x}_i, \theta_i, b_i) = G\left(\mathbf{x}_i'\boldsymbol{\beta} + \alpha b_i + \delta_0\theta_i + \delta_1\theta_i b_i\right) \ , \tag{2.10}$$

where y_i^{sb} is a binary variable indicating either stock or bond ownership, and b_i is an indicator variable equal to 1 when y_i^{sb} corresponds to bond ownership and 0 when it corresponds to stock ownership. The coefficient δ_0 represents the effect of mortgage debt on stock ownership while δ_1 represents the effect of mortgage debt on bond ownership *that is in addition to* the effect on stock ownership. Thus a statistically significant estimate of δ_1 implies that mortgage debt has a significantly greater effect on bond market participation than it does on stock market participation.

Column 1 of Table 5 shows estimates of coefficients δ_0 and δ_1 from Equation 2.10. It shows that the estimate of δ_1 has high statistical significance as measured by its Wald Statistic, so we are able to reject our null hypothesis. To check robustness, Table 5 also shows results from the SUR system in Equations 2.8 and 2.9 using both OLS and probit specifications. In both specifications, we are again able to reject the null hypothesis that a mortgage reduces the probability of stock and bond ownership by equal amounts. The results in Tables 4 and 5 are thus consistent with Proposition 2. In particular, they show that mortgage debt has a greater negative impact on bond ownership than on stock ownership, and that this difference is statistically significant at the 1% level.

[15]We use logit specifications as opposed to probits because a number of both theoretical and practical considerations can make the logistic distribution more appropriate than the normal distribution for use in population studies. See Borooah (2001) and Hahn and Soyer (2005) for more details.

[16]We cluster the standard errors by household to nullify the artificial increase in precision that results from each household appearing twice in a single regression.

Table 5
Differences in effect of mortgage debt on stock and bond ownership

		OLS SUR		Probit SUR	
	Logit	Stocks	Bonds	Stocks	Bonds
	(1)	(2)	(3)	(4)	(5)
Effect of Mortgage on Stock Ownership	-0.2078***				
	[0.0400]				
Additional Effect on Bond Ownership	-0.3448***				
	[0.0457]				
Effect of Mortgage on Participation		-0.0447***	-0.1100***	-0.1388***	-0.3119***
		[0.0083]	[0.0093]	[0.0284]	[0.0268]
Null Hypothesis: Stocks = Bonds					
Wald Statistic	**7.64**				
Chi-squared Statistic		**32.20**		**23.92**	
P-Value	**<0.001**	**<0.001**		**<0.001**	
Observations	32,626	16,313	16,313	16,313	16,313
Clustered by Household	Yes	No	No	No	No
Pseudo R-Squared	0.15	0.29	0.14	0.24	0.11
Mean Dep. Var.	0.524	0.507	0.541	0.507	0.541

This table uses data from the triennial 1989-2004 Survey of Consumer Finances to test the null hypothesis that mortgage debt reduces the probability of stock and bond market participation by equal amounts. Column 1 shows regression results from Equation 10. Columns 2-3 and 4-5 show results from the Seemingly Unrelated Regression (SUR) system in Equations 8 and 9 for OLS and probit specifications, respectively. All regressions include the control variables listed in Table 4. The sample is limited to homeowners who do not live on a farm or in manufactured housing. We truncate income and wealth at the 1st and 99th percentiles of the cross-sectional distribution in order to limit the influence of outliers. Robust standard errors have been adjusted for variation between implicates and are reported below the coefficients in brackets. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

Population-weighted estimates

The regressions shown in Table 4 use sample weights from the SCF and thereby over-represent wealthy households. In Table 6, we use population weights to estimate the effects of mortgage debt on the household portfolio for a sample of homeowners that is representative of the U.S. population. Table 6 shows that, in the population-weighted sample, having a mortgage reduces the probability of stock ownership by 9.8% and of bond ownership by 37.3%. These effects of mortgage debt on asset market participation are smaller than the effects estimated using an unweighted sample, which suggests that our results are strongest amongst more wealthy households.[17] This

[17]The weighted estimates are less statistically significant than the unweighted estimates because sample weights dramatically reduce the efficiency of our regression estimates. This does not diminish

Table 6

Population-weighted equity and bond participation logit regressions

	Household Owns Equities				Household Owns Bonds			
	(1)		(2)		(3)		(4)	
	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates
Reference Household	-0.8003***	31.0%	-0.6179***	35.0%	-0.5898***	35.7%	-0.8987***	28.9%
	[0.1445]		[0.1987]		[0.1289]		[0.1820]	
Household with Mortgage	-0.1034*	-9.8%			-0.4676***	-37.3%		
	[0.0632]				[0.0586]			
Mortgage Interest Rate			-0.0545***	-10.7%			-0.0463***	-9.2%
			[0.0186]				[0.0173]	
ARM Mortgage	-0.0509	-5.0%	-0.0579	-5.6%	-0.1596**	-14.8%	-0.1763**	-16.2%
	[0.0861]		[0.0852]		[0.0779]		[0.0774]	
Have Private Business	-0.0568	-5.5%	-0.0907	-8.7%	-0.1362**	-12.7%	-0.2054***	-18.6%
	[0.0692]		[0.0799]		[0.0633]		[0.0741]	
Income (in $ millions)	3.4425***	21.7%	3.6381***	26.6%	1.1432***	6.1%	1.2382***	8.1%
	[0.5254]		[0.5755]		[0.2755]		[0.3297]	
$Income^2$ (in $ millions)	-0.3469***		-0.3705***		-0.1192***		-0.1282***	
	[0.0460]		[0.0514]		[0.0263]		[0.0362]	
Net Worth (in $ millions)	0.3491***	29.1%	0.3595***	22.5%	0.0971***	6.1%	0.0997***	4.5%
	[0.0546]		[0.0695]		[0.0200]		[0.0261]	
Net $Worth^2$ (in $ millions)	-0.0026***		-0.0028***		-0.0008***		-0.0009***	
	[0.0003]		[0.0007]		[0.0001]		[0.0003]	

	Household Owns Equities				Household Owns Bonds			
	(1)		(2)		(3)		(4)	
	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates	Coefficients	Prob Estimates
Year Fixed Effects	Yes		Yes		Yes		Yes	
Demographic Controls	Yes		Yes		Yes		Yes	
Education Controls	Yes		Yes		Yes		Yes	
Risk Preference Controls	Yes		Yes		Yes		Yes	
Savings Horizon Controls	Yes		Yes		Yes		Yes	
Observations	16,313		9,833		16,313		9,833	
Pseudo R-Squared	0.19		0.17		0.10		0.09	
Mean Dep. Var.	0.338		0.348		0.450		0.423	

This table contains the results of logit regressions of equity and bond ownership on household characteristics and mortgage debt using data from the triennial 1989-2004 Survey of Consumer Finances (SCF). All estimates are weighted using population weights from the SCF. The sample is limited to homeowners who do not live on a farm or in manufactured housing. Columns 2 and 4 are restricted to households with a mortgage. We truncate income and wealth at the 1st and 99th percentiles of the cross-sectional distribution to limit the influence of outliers. Robust standard errors have been adjusted for variation between implicates and are reported below the coefficients in brackets. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***. The reference household is headed by an unmarried, 35-44-year-old white male with a college degree, average risk preference, and income and wealth equal to their respective sample means. The "Probability Estimates" column reports the probability of stock or bond ownership for the reference household and the change in these probabilities associated with a unit increase in a binary variable and a one standard deviation increase in a continuous variable. Our results are robust to using decile dummies and splines to control for income and wealth.

finding is consistent with the benefits to participation as plotted in Figure 3. The figure shows that for households with and without mortgages, sizeable differences in the benefit to equity participation begin to emerge towards the upper end of the wealth distribution. At lower levels of the wealth distribution, the benefit to participation is low enough that having a mortgage makes little difference.

2.3.2 Portfolio share

We now test our model's predictions for the effects of mortgage characteristics on the equity share of liquid wealth. Table 7 shows the results of both conditional OLS and unconditional Tobit regressions of equity share on mortgage characteristics. The equity share is defined as the ratio of stocks to the sum of stocks and bonds.[18] These regressions also include control variables that measure the proportion of assets and debt relative to household net worth. We include these additional controls because a household's allocation of its overall portfolio has implications for the allocation of its liquid portfolio between stocks and bonds.

The first prediction of Proposition 3 is that, conditional on equity participation, having a mortgage should increase the equity share of liquid wealth. Column 1 of Table 7 shows the results of an OLS regression of equity share on mortgage debt, conditional on equity participation. The results show that having a mortgage increases the equity share by roughly 6%. It may seem counter-intuitive that mortgage debt reduces stockholdings but increases the equity share. Yet this finding is consistent because though having a mortgage causes stock holdings to decline by 21%, it causes bond holdings to decline by an even greater 40%. While households may hold fewer stocks due to the smaller expected excess return they offer, they are likely to hold even fewer bonds as these offer a rate of return that is almost certainly lower than the household's mortgage interest rate.

The second prediction of Proposition 3 is that, conditional on having a mortgage, higher mortgage rates should decrease the equity share of liquid wealth. Because many households with mortgages do not participate in equity markets, our underlying sample contains a large number of households with an equity share equal to zero. Therefore, we use a Tobit regression model to estimate the effect of the mortgage interest rate on the equity share. Estimating a Tobit regression accounts for the fact that the large number of observations with an equity share equal to zero do not represent independent realizations drawn from an identical distribution. Column 4 of Table 7 contains the results from a Tobit regression of equity share on the mortgage interest rate. It shows that conditional on having a mortgage, an increase in the

the statistical significance of the relationship between mortgage debt and asset holdings, but instead suggests the effects are strongest amongst the oversampled group, which in this case are wealthy households.

[18]We have also conducted this analysis with equity share defined as the sum of stocks, bonds, and safe assets (which includes cash holdings and demand deposits). We find very similar results.

Table 7
Conditional OLS and tobit regressions of portfolio shares

	Equity Share of Liquid Wealth			
	Conditional OLS Regression		Tobit Regression	
	(1)	(2)	(3)	(4)
Household with Mortgage	0.0607***		0.0460***	
	[0.0079]		[0.0117]	
Mortgage Interest Rate		-0.0051		-0.0126**
		[0.0034]		[0.0053]
ARM Mortgage	-0.0021	-0.0020	0.0161	0.0119
	[0.0111]	[0.0107]	[0.0164]	[0.0167]
Relative Safe Assets	-0.5169***	-0.3749***	-0.6981***	-0.4692***
	[0.1366]	[0.1350]	[0.1683]	[0.1571]
Relative Real Estate	-0.0628	-0.0316	-0.1538*	-0.1041*
	[0.0417]	[0.0246]	[0.0804]	[0.0538]
Relative Bonds	-0.4818***	-0.3039***	-0.3619***	-0.2031**
	[0.1079]	[0.0943]	[0.0975]	[0.0843]
Relative Private Business	-0.0692	-0.0241	-0.1276*	-0.0428
	[0.0490]	[0.0322]	[0.0678]	[0.0535]
Relative Debt	0.0685*	0.0378*	0.1447**	0.0973**
	[0.0353]	[0.0211]	[0.0678]	[0.0462]
Relative Retirement Equity	0.0629	0.0706	-0.0398	-0.0124
	[0.0683]	[0.0522]	[0.0861]	[0.0679]
Constant	0.9121***	0.9515***	0.7120***	0.7437***
	[0.0452]	[0.0411]	[0.0789]	[0.0771]
Year Fixed Effects	Yes	Yes	Yes	Yes
All Controls	Yes	Yes	Yes	Yes
Observations	8,255	4,682	11,236	6,452
Obs. Left-Censored at Zero	-	-	2,981	1,770
R-Squared	0.24	0.16	-	-
Pseudo R-Squared	-	-	0.15	0.11
Mean Dep. Var.	0.717	0.767	0.527	0.557

This table reports results of conditional OLS and unconditional tobit regressions of the portfolio equity share on mortgage characteristics and other controls using data from the triennial 1989-2004 Survey of Consumer Finances. The sample is limited to homeowners who do not live on a farm or in manufactured housing. Columns 2 and 4 are restricted to households with a mortgage. The equity share of liquid wealth is defined as the ratio of stocks to the sum of stocks and bonds. The relative variables are the ratio of each asset (or debt) to net worth. We truncate income and wealth at the 1st and 99th percentiles of the cross-sectional distribution to limit the influence of outliers. Robust standard errors have been adjusted for variation between implicates and are reported below the coefficients in brackets. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

mortgage interest rate leads to a decline in the equity share of liquid wealth.

2.3.3 Endogeneity

This section describes two potential sources of endogeneity with respect to our mortgage variables and our efforts to address them. One potential source of endogeneity stems from our inability to observe whether or not a household is financially naive. In Section 2.3.2.1, we discuss the empirical implications of financial naivete and describe our method of controlling for it. The other endogeneity concern is that the regressions do not adequately control for wealth effects and that the negative relationship between mortgage debt and stock ownership could be driven by less wealthy households owning fewer stocks. We address this issue in detail in Section 2.3.2.2.

Financial naivete

In Section 2.3.1, we found that conditional on having a mortgage, a higher mortgage interest rate reduces the probability of equity participation. According to our model, this is because a higher mortgage interest rate leads to a lower expected excess return on the risky asset and thus reduces the benefit to stock ownership. However, this effect could also be driven by our inability to observe whether or not a household is financially naive. If financially naive households tend to have higher mortgage interest rates then our mortgage rate variable is endogenous.[19] This means our explanatory variables are no longer orthogonal to the error term and our estimate of the effect of the mortgage interest rate on equity participation is not consistent. That is to say, it could be that households with high mortgage interest rates do not participate in equity markets simply because they are financially naive, and not because they calculate that it is advantageous for them to repay outstanding debts instead.

We address this issue by adding two proxy variables for financial naivete to our participation regressions. The first proxy variable identifies households as financially naive if they fail to refinance out of a high mortgage interest rate and into a lower one when it is optimal for them to do so. Optimal refinancing is the topic of a large literature with analysis ranging from the complete closed-form solution of Agarwal, Driscoll, and Laibson (2007) to the 150 basis point rule-of-thumb calculated by Schwartz (2007). Our approach combines the rule-of-thumb criteria with a crude adjustment for the credit worthiness of each household. We first calculate the spread

[19] Financially naive households could have higher interest rates for two reasons. First, they could fail to shop around for a low rate and instead accept a higher rate. Second, they could fail to refinance into a new mortgage with a lower rate when doing so is optimal. If financially aware households do not make this same mistake, they will on average have lower mortgage interest rates than households that are financially naive. This is a likely scenario as over half of the households with mortgages in our unweighted sample have refinanced their mortgage at least once.

between each household's mortgage interest rate and the prevailing mortgage rate in the sample year. To this number we add a credit-spread, defined as the difference between each household's mortgage rate and the mean mortgage rate in the year the mortgage was taken out. We then tag a household as financially naive if this credit-adjusted spread exceeds the 150 basis points advocated by Schwartz.[20] Adding the credit spread allows us to account for the credit-worthiness of each household. For example, a household with below average credit that has to borrow at 100 basis points above the prevailing rate in any given year will need the current prevailing rate to drop to over 250 basis points below their mortgage rate for refinancing to be optimal.[21] We are also careful not to tag households as naive if they cannot refinance because they have lost their job in the interview year or choose not to refinance because they have less than five years remaining on their mortgage contract.[22]

Our second proxy variable identifies households as financially naive if they are unable to report their exact mortgage interest rate to the SCF surveyor. The SCF contains a parallel coding system used to track the certainty with which the respondent answers each survey question. Exact answers, answers that fall within a range, and missing answers that are subsequently imputed are all recorded in the dataset. We use these data to flag those respondents that are unable to identify their exact mortgage interest rate.[23] Our reasoning is that if the household is unable to identify their mortgage rate and unable to locate a document containing this information, it is unlikely that the rate is a factor in the household's portfolio allocation decisions.

Table 8 shows that the results of our participation regressions are unchanged once we include our proxy variables for financial naivete. The first four columns show results from regressions of stock ownership on control variables, mortgage characteristics, and our proxy variables for financial naivete. Columns 1 and 2 show the effect of having a mortgage on stock ownership and Columns 3 and 4 show the effect of the mortgage interest rate on stock ownership conditional on having a mortgage. The results show that although financially naive households tend to participate less in equity markets, the coefficients on our mortgage variables remain largely unchanged. The last four columns show similar results for bond ownership. Households that are financially naive are less likely to own bonds, but again the coefficients on our mortgage variables are unchanged. These results suggest that the effect of mortgage debt

[20] We also use a more lenient threshold of 250 basis points to separate households that could most clearly benefit from refinancing and yet fail to do so. We find very similar results.

[21] An implicit assumption here is that each household's credit-worthiness has remained constant since the time they took out their mortgage.

[22] The fixed-cost of refinancing makes it a less financially attractive option toward the end of the mortgage term. The 5-year threshold is chosen arbitrarily. We also conduct the analysis with 2-year and 10-year thresholds and find identical results.

[23] Of the respondents in our sample asked for their mortgage interest rate, over 94% provide an exact answer, 1% provide an answer that falls within a range, and 3% are unable to provide any information on their mortgage rate. Mortgage rates for the remaining 1% of respondents were estimated using other supplied answers.

Table 8

Logit regressions of equity and bond ownership controlling for financial naivete

	Household Owns Equities				Household Owns Bonds			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Reference Household	-0.2937***	-0.2974***	-0.1988	-0.1679	-0.3134***	-0.3219***	-0.6208***	-0.5210***
	[0.1129]	[0.1129]	[0.1658]	[0.1616]	[0.1016]	[0.1017]	[0.1513]	[0.1481]
Household with Mortgage	-0.2086***	-0.2279***			-0.4703***	-0.4918***		
	[0.0491]	[0.0483]			[0.0451]	[0.0443]		
Mortgage Interest Rate			-0.0724***	-0.0780***			-0.0416***	-0.0616***
			[0.0178]	[0.0158]			[0.0160]	[0.0143]
Naïve- Mort Rate 150 bp Above	-0.1702**		-0.0451		-0.2494***		-0.1710**	
	[0.0678]		[0.0770]		[0.0621]		[0.0703]	
Naïve- Do Not Know Mort Rate		-0.1038		-0.1377		-0.2414***		-0.2643***
		[0.0829]		[0.0844]		[0.0757]		[0.0785]
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
All Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	16,313	16,313	9,833	9,833	16,313	16,313	9,833	9,833
Pseudo R-Squared	0.24	0.24	0.20	0.20	0.11	0.11	0.09	0.09
Mean Dep. Var.	0.507	0.507	0.477	0.477	0.541	0.541	0.486	0.486

This table contains the results of a logit regression of equity and bond ownership on household characteristics and mortgage debt using data from the triennial Survey of Consumer Finances (SCF). The sample is limited to homeowners who do not live on a farm or in manufactured housing. Columns 3-4 and 7-8 are restricted to households with a mortgage. We truncate income and wealth at the 1st and 99th percentiles of the cross-sectional distribution to limit the influence of outliers. All regressions control for gender, race, marital status, and children. Robust standard errors have been adjusted for variation between implicates and are reported below the coefficients in brackets. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***. The reference household is headed by an unmarried, 35-44 year-old white male with a college degree, average risk preference, and income and wealth equal to their respective sample means. Our results are robust to using decile dummies and splines to control for income and wealth.

on asset ownership is not driven by households that are financially naive.

Wealth effects

The effect of mortgage debt on stock and bond ownership is difficult to isolate from the effects of wealth because both effects move in the same direction. Homeowners that have paid off their mortgage tend to be older and wealthier and therefore more likely to own stocks and bonds. But according to our model, the influence of mortgage debt on stock participation is distinguishable from wealth effects at the point in time that a household makes its final mortgage payment. Upon paying off its mortgage, the household's implied risk-free rate drops from the mortgage interest rate to the interest rate on bonds and the benefit to equity participation increases. At that point we should expect to see a discrete increase in the probability of both stock and bond market participation. In the absence of any discrete increases in household wealth, the increased likelihood of stock and bond ownership can be attributed to the retirement of mortgage debt. Ideally we would plot a smooth profile of stock and bond participation against a mortgage timeline in order to check for a discontinuity at the point in which a household pays off its mortgage. Unfortunately, limitations of the available data prohibit us from implementing this type of regression discontinuity design.

This discontinuity framework illustrates the nature of our endogeneity problem; if our controls for wealth are insufficient, then our results may be driven by the fact that households with mortgages tend to be less wealthy than those without mortgages. One scenario in which our wealth controls could be insufficient is if there are large unobservable increases in wealth that allow a household to simultaneously pay off its mortgage and to purchase stock. Such increases in wealth could result from either large amounts of labor income or unexpected windfalls. We find evidence suggesting that households with larger amounts of expected future labor income are not likely to pay off their mortgage early and purchase stocks because these households tend to take out larger amounts of mortgage debt. Regression analysis using SCF data shows a significant positive effect of mortgage debt on dollar equity holdings. This accords with the observation made by Cocco (2005) that households with higher expected labor income take out larger mortgages thereby creating a positive cross-sectional correlation between leverage and holdings of risky assets.[24] Yet there is still the potential for large increases in wealth resulting from unexpected windfalls. In this scenario, the discrete increase in a household's benefit to participation is likely to be overshadowed by a discrete increase in wealth that could prompt the household to both pay off their mortgage and purchase additional assets such as stocks and bonds. Though this scenario is likely to be an uncommon one, there is no good solution for it when using cross-sectional data.

[24]We do not present our regressions of dollar asset holdings on mortgage characteristics to conserve space. These results are available from the authors upon request.

Table 9

Fixed-effects logit regressions using Panel Survey of Income Dynamics

	Household Owns Stock			
	Unweighted		Weighted	
	Coefficients	Prob Estimates	Coefficients	Prob Estimates
Household with Mortgage	-0.07566	-7.3%	-0.09905	-9.4%
	[0.11202]		[0.14300]	
Net Worth (in $ millions)	0.69040***	23.7%	0.51618**	20.9%
	[0.09065]		[0.25628]	
Net Worth Squared (in $ millions)	-0.03137***		-0.01889	
	[0.00539]		[0.02119]	
Income (in $ millions)	1.55766***	8.9%	1.39875**	8.7%
	[0.53475]		[0.68651]	
Income Squared (in $ millions)	-0.27509**		-0.22983*	
	[0.13155]		[0.12630]	
Year Fixed Effects	Yes		Yes	
Household Fixed Effects	Yes		Yes	
Age and Demographic Controls	Yes		Yes	
Private Business Control	Yes		Yes	
Education Controls	Yes		Yes	
Observations	10,043		10,043	
No. of Households	2,109		2,109	
Pseudo R-Squared	0.05		0.04	
Mean Dep. Var.	0.471		0.488	

This table reports the results of a fixed-effects logit regression of stockholdings on mortgage debt and other control variables. The Probability Estimates column reports the marginal effect on the probability of stock ownership resulting from a unit increase in a binary variable and a one standard deviation increase in a continuous variable. The sample is from the wealth waves of the PSID in 1984, 1989, 1994, and every other year between 1999 and 2005 and includes all homeowners that do not enter the panel without mortgage debt and later accumulate it. Coefficients and probabilities in the Weighted column are estimated using PSID population weights. Robust standard errors are reported in brackets. Coefficients statistically significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***.

We address the potential for endogeneity with respect to wealth in a number of ways. The cross-sectional results from the SCF are robust to additional income and wealth controls including dummy variables representing deciles of the variables' distributions, as well as splines allowing for differences in slope and intercept parameters across the distributions. We use several different methods and sets of assumptions to estimate and control for total lifetime wealth, which includes future expected labor income. We also restrict our analysis to the wealthiest quartile and decile of U.S. households. In each case, we find results very similar to the ones presented in Tables 4 and 7. Finally, we use PSID panel data to address the potential for significant increases in wealth resulting from unexpected windfalls. We estimate a fixed-effects conditional logit model, which controls for unobservable household characteristics by focusing on variation in a single household's portfolio over time. In order to isolate the effects of the discontinuity described above, we exclude from our regression homeowners who enter the panel without mortgage debt and later accumulate it.[25] Table 9 contains the results of unweighted and population-weighted conditional logit regressions of equity participation on mortgage debt. When our sample is weighted to represent the U.S. population, households that pay off their mortgage are subsequently 9.4% more likely to own stocks. This effect is nearly identical in magnitude to the 9.8% effect estimated using data from the SCF, though our estimate is no longer statistically significant.[26]

2.4 Welfare Implications of Debt Retirement

2.4.1 Credit cards and other sources of debt

Our empirical analysis has focused exclusively on mortgage debt for two reasons. First, mortgage debt is the largest liability on the household balance sheet. Second, regression analysis using mortgage debt is likely to suffer from fewer empirical problems stemming from unobserved household characteristics than an analysis incorporating additional types of debt. For example, households with student loans may be better educated and more financially sophisticated than other households. In contrast, households who consistently carry a credit card balance may be more financially naive or suffer from a lack of self-control.[27] Nevertheless, households do indeed hold many different types of debt, often with interest rates that are significantly higher than their mortgage rate. In this section, we show that many households should forego equity or bond market participation on account of the high interest rates they pay on

[25] These households comprise roughly 12% of the relevant PSID sample.

[26] One potential explanation for the loss of statistical significance is that the SCF oversamples wealthy households while the PSID oversamples less wealthy households. As we point out in Section 3.1.3, our results are strongest amongst more wealthy households. Thus regressions using the PSID sample are likely to be less precise.

[27] See, for example, Laibson, Repetto, and Tacman (2003).

Table 10
Stock and bond holdings by debt characteristics

Category	% U.S. Households	% Own Stocks	Mean Stocks ($)	% Own Bonds	Mean Bonds ($)
Overall	100	25.0	38,481	35.2	21,867
No Mortgage	59.2	19.2	36,856	31.1	26,240
Mortgage	40.8	33.3	40,840	41.1	15,516
Mortgage Interest Rate > LT-Bonds	39.2	32.8	36,253	40.9	14,083
Mortgage Interest Rate > Stocks	1.8	11.6	1,543	20.5	1,461
No Credit Card Balance	67.4	28.7	65,744	35.2	30,539
Credit Card Balance	32.6	23.4	10,773	29.9	4,738
CC Interest Rate > T-Bills	30.8	24.5	11,446	30.6	4,940
CC Interest Rate > Stocks	24.8	23.8	10,672	30.1	4,810
No Debt	41.8	21.9	57,881	31.2	33,837
Any Debt	58.2	30.6	40,582	35.0	13,718
Any Debt with Interest Rate > T-Bills	57.0	31.2	40,978	35.4	13,785
Any Debt with Interest Rate > LT-Bonds	55.1	30.5	36,076	35.0	12,046
Any Debt with Interest Rate > Stocks	25.8	23.5	10,392	29.6	4,662

This table reports stock and bond holdings by debt characteristics for all households in the Survey of Consumer Finances (SCF). Estimates for the Overall and Mortgage categories are from the triennial 1989-2004 SCF. Estimates for the Credit Card and Any Debt categories are from the triennial 1995-2004 SCF because credit card rates are not surveyed before 1995. All estimates are weighted using the population weights from the SCF. All dollar amounts are reported in 2004 dollars. The rates of return used for Long-term Bonds (3.39%) and Stocks (9.16%) are from Poterba (2002). The rates of return on T-bills are from the St. Louis Federal Reserve. We calculate the rate of return for each SCF survey year as the average return on 3-Month Treasury bills in the secondary market during that year. The mortgage interest rate is adjusted for taxes, as described in Appendix A.3.

common forms of debt. We also discuss the implications for asset participation and welfare stemming from these additional sources of debt.

Table 10 shows household participation rates and mean dollar holdings of stocks and bonds grouped by debt holdings and interest rates.[28] Overall, households with credit card debt are less likely to own both stocks and bonds and tend to hold these

[28]This sample includes households that rent instead of own their home. The positive cross-sectional relationship between mortgage debt and participation is due to the fact that households with a mortgage tend to be wealthier than households without a mortgage once we include renters. This is also why the overall stock participation rate is lower than the rate shown in the SCF regression results. The inclusion of renters reduces the participation rate from 33.8% to 25%.

assets in much smaller amounts.[29] The most striking result in Table 10 is the comparison of households with and without any form of debt. Between 1995 and 2004, 25.8% of U.S. households held outstanding debt obligations with an after-tax interest rate higher than the average after-tax return to stock ownership.[30] Another 55.1% of households held debt with an interest rate greater than the average return on long-term government bonds.[31] These findings suggest that even in the absence of any of the information or transaction costs often used to explain household non-participation, more than a quarter of U.S. households have little incentive to participate in equity markets on account of their outstanding debt. The combination of mortgage debt, home equity debt, and credit card debt on the household balance sheet is a perfectly rational explanation for a large portion of the limited participation puzzle.

2.4.2 Costs of non-optimal bond market participation

Table 10 also shows that many households with high interest rate debt do in fact hold stocks and bonds. Between 1995 and 2004, roughly 6% of all U.S. households simultaneously owned stocks and carried a credit card balance with an annual interest rate greater than the long-term return on equities. This suggests that there may be significant foregone benefits to debt repayment and, in particular, the repayment of credit card debt.

In this section, we estimate a lower bound on the annual costs of non-optimal bond market participation as it relates to foregone debt-repayment. We focus on bonds because in the context of our portfolio choice model these assets are assumed to be safe and will not require any adjustments for risk. In each year we estimate the annual foregone benefits to both credit card and mortgage repayment for household i as:

$$B_j^i = \min(bonds^i, balance_j^i) \cdot \left(R_j^i - R_f\right) \quad \text{for } j = credit\ card,\ mortgage \qquad (2.11)$$

where B_j^i is the foregone benefit to household i from repayment of debt type j, $bonds^i$ is the bond holdings of household i, $balance_j^i$ is the outstanding balance of debt type

[29] Credit card debt holders are defined as those households who report that they always or almost always carry a balance on their credit card. This excludes households who use credit cards for liquidity purposes and do not regularly have a high return available to them in the form of paying down their credit card debt.

[30] The long-term after-tax returns on stocks and long-term government bonds are from Poterba (2002). He estimates that between 1926 and 1996, these returns are equal to 9.16% and 3.39%, respectively.

[31] The true proportions of U.S. households with debt and interest rates greater than the returns on stocks and bonds are almost certainly greater than 25.8% and 55.1%, respectively. Due to data limitations, our calculations do not include the interest rates charged on additional forms of debt, such as car loans or student loans. Also, credit card debt is systematically underreported in the SCF. See Section II-A of Zinman (2007b) for details.

Table 11
Annual costs of non-optimal bond market participation

Type of Debt	% of Households	Median ($)	Mean ($)
Mortgage	16.0	60	269
Credit Card	15.1	81	217

This table reports the annual costs of non-optimal bond market participation as defined by Equation (12) for households in the Survey of Consumer Finances (SCF). The medians and means reported are conditional on the cost being positive. Estimates for the Mortgage category are from the years 1989-2004. Estimates for the Credit Card category are from the years 1995-2004 because credit card rates are not surveyed before 1995. All estimates are weighted using the population weights from the SCF. All dollar amounts are reported in 2004 dollars

j for household i, and R_j^i is the annual after-tax interest rate paid by household i on debt type j. The annual benefit is equal to the dollar savings that a household would accumulate through foregone interest charges by using their available bond holdings to retire their high-interest rate mortgage or credit card debt.

To clarify, the existence of 'foregone benefits' as we define them here does not imply the existence of arbitrage opportunities. As Zinman (2007a) and others point out, the repayment of outstanding debt is different from holding other assets, such as demand deposits. Households may prefer the latter as these assets are more liquid and easier to use in daily transactions or as a buffer against emergency expenses. With this in mind, we focus our analysis on the repayment of liabilities using assets that have comparable investment horizons and levels of liquidity.

We first consider the liquidity implications and investment horizons associated with repayment of mortgage debt. As discussed in Section 2.1, the implicit risk-free asset purchased through repayment of mortgage debt has a duration equal to D, which we define explicitly in an online appendix to this paper. Since the typical mortgage held by U.S. households has a term of thirty years, repayment of this debt is equivalent to the purchase of an asset with a duration greater than the savings horizons of many households. However, home equity lines of credit and other types of home equity loans allow households to extract home equity and reverse their decision to retire debt early.

Next we consider early retirement of credit card debt. The repayment of credit card debt corresponds to the purchase of an asset with a shorter duration than the asset associated with mortgage debt repayment. This is why in Table 10 we compare credit card interest rates with the return on 3-month Treasury bills. The retirement of credit card debt is also reversible to the extent that the goods or services a household wishes to purchase can be paid for with the same credit card. To calculate the benefit to debt retirement, we assume that households' checking and savings accounts contain the optimal amount of liquid assets required for daily transactions and emergency savings. The exclusion of cash accounts means bond holdings, such as certificates of deposit and long-term government bonds, are the assets most comparable to the implicit assets purchased through debt repayment.

Table 11 shows the mean and median annual costs of non-optimal bond market participation as calculated using Equation 2.11. It shows that the 16% of U.S. households who hold both bonds and a high interest rate mortgage would save a median of $60 each year by using their bond holdings to repay their mortgage debt. Similarly, 15.1% of U.S. households would have a median annual savings of $81 were they to use their bond holdings to repay credit card debt. The benefits of debt repayment shown in Table 11 are calculated separately for mortgage debt and credit card debt. To estimate the aggregate household benefit to debt repayment, we account for a household's decision of which type of debt it ought to retire first. We assume that households with both bond holdings and outstanding debt first repay the outstanding balance on the account with the highest interest rate. If two accounts have the same interest rate, they would first repay the one with the lowest balance. If the household has any remaining bonds, it then repays the account with the next highest interest rate, and so on until the household runs out of bonds or all outstanding debt is paid.

Figure 5

Annual aggregate costs of non-optimal bond market participation

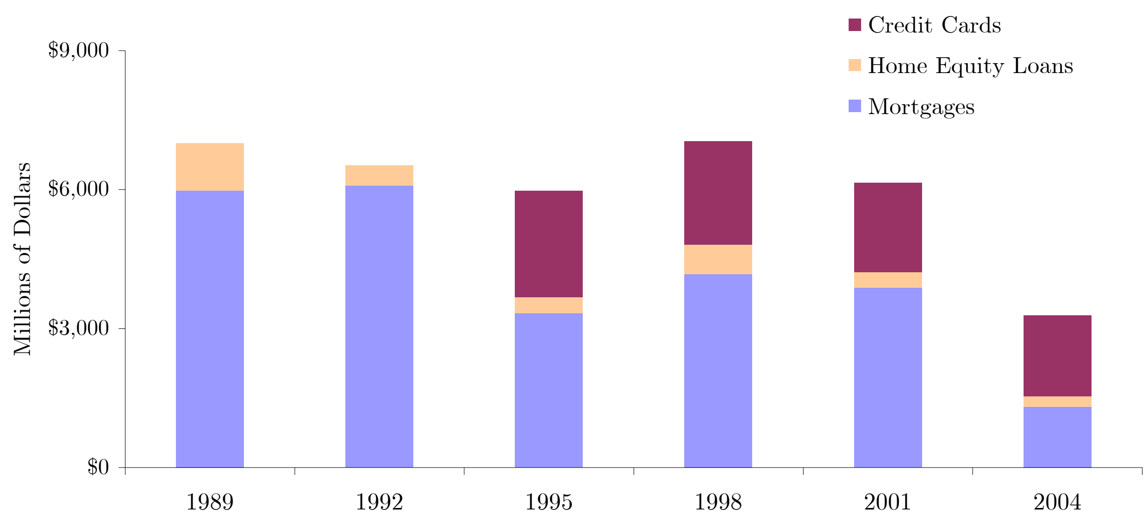

This figure shows annual aggregate dollar costs calculated using the algorithm described in Section 4.2 for the weighted sample of households in each year of the Survey of Consumer Finances. The algorithm calculates the interest savings that would accrue by repaying as much debt as bond holdings allow amongst households that simultaneously hold non-retirement bonds and have outstanding balances on credit cards, home equity loans, or mortgages. Debts on the household balance sheet are repaid in the order of highest interest rate first. Credit card interest rates are first surveyed in 1995. All dollar amounts are reported in 2004 dollars.

Figure 5 shows the annual aggregate costs of non-optimal bond market participation by year. Between 1989 and 2004, the aggregate costs average around $7 billion dollars. In most years, the majority of potential savings come from repaying mortgage debt, followed by credit card debt and then home equity loans. The trend in costs reflects the decline in mortgage interest rates over this time period. Despite the substantial aggregate costs shown in Figure 5, the median costs shown in Table 11 are relatively small. Furthermore, the mean benefits to debt repayment shown in Table 11 are significantly greater than the median benefits. This suggests that the large aggregate costs to non-optimal bond market participation are not a result of rampant financial irrationality by U.S. households, but instead are driven by a small fraction of households who appear to forego large financial gains by failing to use their bond holdings to repay their debts. Using measures of annual income, we find that only 5.1% of these households—or roughly 1% of all U.S. households—have annual costs of non-optimal bond market participation that are greater than or equal to one week's worth of income.

2.5 Conclusion

This paper develops a conceptual framework for the effects of high interest rate debt on asset market participation and portfolio allocation. There are two central themes to our framework. First, households with high interest rate debt have a reduced benefit to equity participation and in many cases should not own stocks. Second, repayment of outstanding debt almost always yields a higher rate of return than many of the safe assets that conventional finance models predict households should hold in large amounts, such as short-term Treasury-bills or long-term government bonds.

Our empirical analysis finds evidence suggesting that households may incorporate the option of debt repayment into their investment decisions. The majority of observed investment choices are consistent with behavior predicted by our theoretical framework, and our regression results highlight the role of debt interest rates as a central mechanism behind participation and allocation decisions. We also show that our results remain unchanged when we account for the potential influence of financial naivete. We use a population-weighted panel regression to account for unobserved wealth effects and find some evidence supporting the results of our cross-sectional analysis. For households whose portfolio decisions are inconsistent with the predictions of our model, we estimate the costs of simultaneously holding high interest rate debt and low yielding fixed-income assets. We find significant aggregate benefits to debt repayment, though they are driven by a small number of households who appear to forego large financial gains by failing to use their bond holdings to repay their debts. For the majority of households who do not behave according to our model, the potential savings from debt repayment are quite low.

Unfortunately, our study offers few predictions for household decisions concerning the trade-off between debt repayment and consumption. Our portfolio choice framework, by relying on a simple and tractable model, neglects the relationship between savings and consumption. We believe the relationship between debt repayment and consumption is an important area of future empirical research, though we acknowledge the data limitations inherent in such research.

Understanding how debt affects household investment is particularly relevant given the events of the recent economic downturn, as well as the high level of household leverage. The past two decades have seen an increase in the number of households with mortgage debt and a decline in average mortgage interest rates. Our analysis offers a few predictions for the effects on household investment and the relative demand for assets. First, the demand for assets is predicted to increase in the upcoming years amongst households that have recently refinanced their mortgage in response to Federal Reserve and GSE intervention in the mortgage markets. Households that refinance into these historically low mortgage rates will have lower returns available to them in the form of early mortgage repayment. Second, the recent turmoil in credit markets and the rethinking of government policies encouraging homeownership are likely to result in the extension of fewer credit contracts. With less newly acquired debt to repay, households will become more likely to hold investment assets. Third, if households have altered their expectations of the return on equity and perceive the risky asset as offering a lower expected excess return, then debt repayment may become more attractive. Finally, to the extent that households will save any forthcoming tax rebates or stimulus checks, as opposed to increasing their consumption, our study predicts that they are likely to choose debt repayment over stock or bond ownership as the preferred method of saving.

Bibliography

[1] Agarwal, S., J.C. Driscoll, and D. Laibson. 2007. Optimal Mortgage Refinancing: A Closed Form Solution. Working Paper W13487, National Bureau of Economic Research.

[2] Ameriks, J., and S. P. Zeldes. 2004. How Do Household Portfolio Shares Vary with Age? Working Paper, Columbia University.

[3] Borooah, V.K. 2001. *Logit and Probit: Ordered and Multinomial Models.* California: Sage Publications.

[4] Campbell, J. Y. 2006. Household Finance. *Journal of Finance* 61:1553–1604.

[5] Canner, N., N. G. Mankiw, and D. N. Weil. 1997. An Asset Allocation Puzzle. *American Economic Review* 87:181–91.

[6] Chetty, R. and A. Szeidl. 2009. The Effect of Housing on Portfolio Choice. Working Paper, University of California, Berkeley.

[7] Cocco, J. F., 2005. Portfolio Choice in the Presence of Housing. *Review of Financial Studies* 18:535–67.

[8] Davis, S. J., F. Kubler, and P. Willen. 2006. Borrowing Costs and the Demand for Equity Over the Life Cycle. *Review of Economics and Statistics* 88:348–62.

[9] Flavin, M., and S. Nakagawa. 2008. A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence. *American Economic Review* 98:474–95.

[10] Flavin, M., and T. Yamashita. 2002. Owner-Occupied Housing and the Composition of the Household Portfolio. *American Economic Review* 92:345–62.

[11] Fratantoni, M. C. 2001. Homeownership, Committed Expenditure Risk, and the Stockholding Puzzle. *Oxford Economic Papers* 53:241–59.

[12] Greene, W. H. 2002. *Econometric Analysis.* New York: Prentice Hall.

[13] Grossman, S. J., and G. Laroque. 1990. Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods. *Econometrica* 58:25–51.

[14] Hahn, E.D., and R. Soyer. 2005. Probit and Logit Models: Differences in the Multivariate Realm. Working paper, George Washington University.

[15] Heaton, J., and D. Lucas. 1997. Market Frictions, Savings Behavior, and Portfolio Choice. *Macroeconomic Dynamics* 1:76–101.

[16] Heaton, J., and D. Lucas. 2000. Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk. *Journal of Finance* 55:1163–98.

[17] Kennickell, A. B., and R. L. Woodburn. 1999. Consistent Weight Design for the 1989, 1992 and 1995 SCFs, and the Distribution of Wealth. *Review of Income and Wealth* 45:193–215.

[18] Laibson, D., A. Repetto, and J. Tobacman. 2003. A Debt Puzzle. In P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (eds.), *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps,* pp. 228–66. Princeton: Princeton University Press.

[19] Monalto, C. P., and J. Sung. 1996. Multiple Imputation in the 1992 Survey Of Consumer Finances. *Association for Financial Counseling and Planning Education* 7:133–41.

[20] Poterba, J. M. 2002. Taxation, Risk-Taking, and Household Portfolio Behavior. In A. J. Auerbach and M. Feldstein (eds.), *Handbook of Public Economics*, *Volume 3*. Amsterdam: Elsevier Science.

[21] Schwartz, A. 2006. Household Refinancing Behavior in Fixed Rate Mortgages. Working Paper, Harvard University.

[22] Vella, F. 1992. Simple Tests for Sample Selection Bias in Censored and Discrete Choice Models. *Journal of Applied Econometrics* 7:413–21.

[23] Viceira, L. M. 2001. Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income. *Journal of Finance* 56:433–70.

[24] Vissing-Jorgensen, A. 2002. Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures. Working Paper W8884, National Bureau of Economic Research.

[25] Wachter, J. A., and M. Yogo. 2008. Why Do Household Portfolio Shares Rise in Wealth? AFA 2008 New Orleans Meetings Paper.

[26] Yao, R., and H. H. Zhang. 2005. Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints. *Review of Financial Studies* 18:197–239.

[27] Zellner, A. 1962. An Efficient Method of Estimating Seemingly Unrelated Regression Equations and Tests for Aggregation Bias. *Journal of the American Statistical Association* 57:348–68.

[28] Zinman, J. 2007a. Household Borrowing High and Lending Low Under No-Arbitrage. Working Paper, Dartmouth College.

[29] Zinman, J. 2007b. Where is the Missing Credit Card Debt? New Clues and Implications. Working Paper, Dartmouth College.

Chapter 3

Crude Drilling: An Analysis of Incentives and Behavior in the Oil Industry During the 1860s

Prodded by greed and by the landowner, who assumed no risk, but shared handsomely in the gains, and undeterred by prices which were sometimes lower than the cost of production, operators continued to open new territory.

–*The Golden Flood,* Asbury (1941)

Introduction

This paper analyzes how prospectors' initial capital endowments effected the rate of oil extraction during the first decade of oil exploration in the United States. I use a new dataset that combines lease and property data with information on the type drilling technology and the oil flow rates of wells within Venango County, Pennsylvania during the 1860s to compare the extraction rates across different classifications of drillers. The data show that variations in initial capital endowments led prospectors to enter into markedly different property agreements and to utilize different drilling technology. Access to startup capital was inversely related to the rate at which wells were subsequently depleted and an overall under-capitalization of the early oil industry can explain the excessive supply and price volatility that characterized the industry.

The initial decade of oil development in the United States provides a controlled setting in which to examine suppliers' behavior with respect to non-renewable resource extraction. The prevailing opinion at the time was that worldwide oil supplies were limited to the confined geographical region of Western Pennsylvania. Yergin (1991) quotes the State Geologist of Pennsylvania in 1880 remarking that "the amazing exhibition of oil was only a temporary and vanishing phenomenon." The testimonials of early prospectors in the region show that this sentiment of limited oil supplies was shared by market participants. Simultaneously, the demand for petroleum products during 1860s was tightly linked to the demand for illuminants. The industrial uses of oil were initially limited to being a substitute good for the high priced whale oil that dominated the market, as many of its current uses had yet to be discovered.

Though the narrative accounts of oil drillers in Western Pennsylvania during the 1860s indicate an understanding of the inverse relationship between supply and price, their aggregate behavior often contradicted this principle. Overdrilling was rampant in the region. Many rudimentary wells were drilled that could not be controlled once oil was struck. Upon striking oil, prospectors often drilled numerous adjacent wells that increased the immediate supply of oil but thereby reduced the internal pressure of the well and the aggregate amount of recoverable oil. This led to cycles of over- and undersupply with associated fluctuations in price, as can be seen in Figure 1. Drillers almost universally neglected to invest in storage and transportation infrastructure, which meant that excess oil flowed down streams or back into the ground. An explanation for the excessive waste of the industry as a whole can be traced to the incentives of the majority of early oil drillers, wildcatters.

Wildcatters accounted for the majority of drillers in the early years of oil exploration. A group comprised primarily of poor farmers, these prospectors flooded western Pennsylvania following the news of the first oil strike in 1859 in a similar fashion to the California gold prospectors from a decade earlier. They lacked the funds to purchase land and therefore usually leased or sub-leased land that lacked any proven oil deposits, a behavior that gained them the name wildcats. Their lack of financing also led them to rely on less capital intensive drilling technologies. Wildcat wells were often drilled using rudimentary techniques powered by humans and livestock. These had the advantage of low cost and a high degree of mobility, but the disadvantage of resulting in unevenly bored wells of limited depths.

Whereas the behavior of wildcat drillers can be traced to their limited access to capital, merchant and corporate drillers engaged in prospecting operations with increasingly large degrees of initial capitalization. Merchant drillers were predominantly from the region and had access to private capital and some access to bank credit. Corporate drillers created joint-stock companies to pool the capital of investors from New York and Pittsburgh to deploy in the oil region. Both merchant and corporate drillers were more likely to purchase tracts of land outright and thereby limit the problems associated with excessive drilling on a single reservoir. They were also much more likely to invest in the newest drilling technology and thereby gain control of their eventual rate of oil extraction.

These three groups of drillers differed greatly in the types of property arrangements they secured, in their choice of drilling technology, and in their subsequent rates of oil extraction. I propose that these differences arise because the discount rates for future oil revenue amongst each of these groups was inversely related to their initial level of capitalization. These different valuations of future cash flows incentivized the

Figure 1

United States Crude Oil Production and Prices at Well

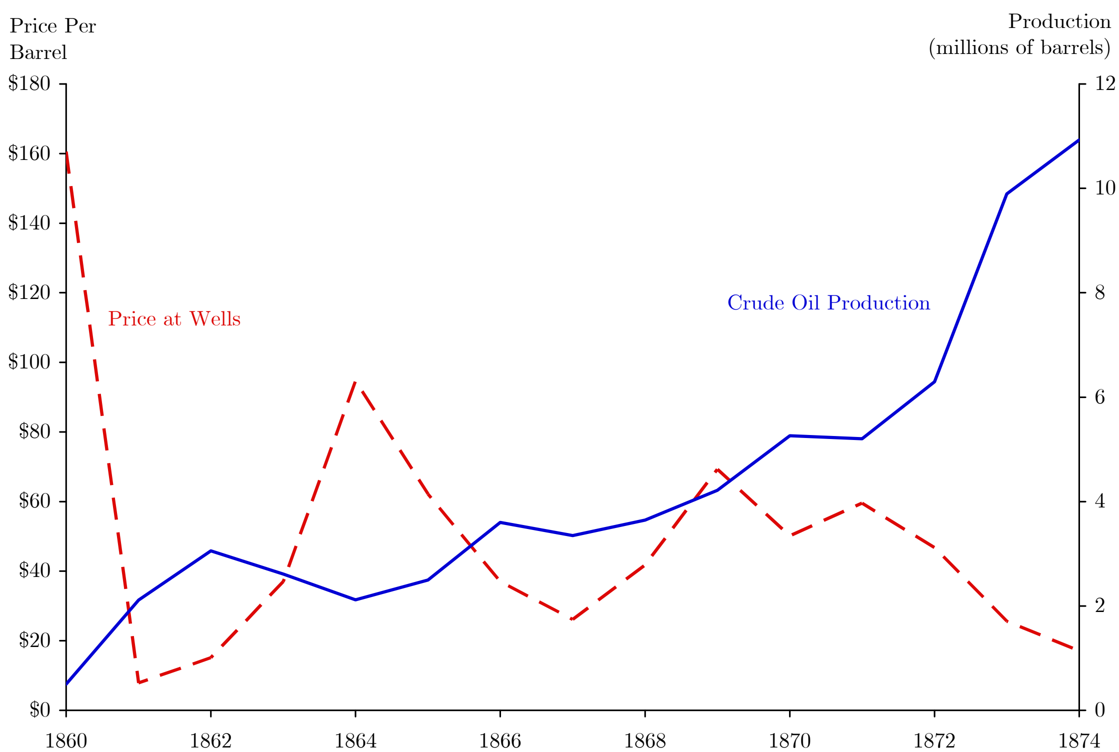

Data are from Logan (1930) and Johnston (2005). Nearly all U.S. production between 1860 and 1874 originated in the Western Pennsylvania region.

different rates of extraction witnessed across the groups of prospectors. By combining narrative accounts with a newly constructed dataset of property contracts, well types, and extraction rates, I show how differences in access to capital changed incentives and drilling behavior.

Economists have been interested in understanding the oil industry and modeling the optimal rate of extraction of non-renewable resources since the early 20th century. The first empirical study of the oil industry was undertaken by Stocking (1928) and characterizes an industry prone to overextraction during most of the first half century of its existence. He cites a 1914 study by the U.S. Bureau of Mines that estimated the economic costs of excessive wells and wasted oil as equal to a quarter of the value of annual U.S. oil production. This empirical finding stands in contrast to the first formal theoretical analysis of non-renewable resource extraction by Hotelling (1931). His static equilibrium model to describe how price and demand dictate the rate of extraction in competitive markets finds that the price of an exhaustible resource should grow at the market rate of interest. This theoretical finding, though often at odds with observed drilling behavior, forms the basis for natural resource economics and shows that rational agents in competitive markets are capable of depleting a non-renewable resource at the socially optimal rate.

The theoretical determinants of oil extraction behavior were not examined again until the oil crisis of the 1970s. Stiglitz (1975) and Arrow and Chang (1978) add information externalities and an examination of high exploration costs to the static Hotelling model to explain instances of sub-optimal behavior by oil suppliers. Gilbert (1979) uses the technique of dynamic programming to expand to understand how uncertainty effects extraction rates. Yuan (2002) applies a game-theoretic approach to the problem to show that it can be rational for a landowner with non-exclusive ownership of a common-pool field to divide and sublease production of the field despite the increase in overall waste from such property arrangements. Each of these theoretical contributions addressed a deficiency of the original Hotelling analysis while also confirming the model's central conclusion that competitive markets should result in rates of extraction that approximate a time-discounted path.

Empirical studies of drilling behavior have continued to find property contracts and oil extraction rates that deviate from the theoretical predictions. Libecap and Wiggins (1985) analyze the failure of private participants to reach optimal outcomes with regard to shared resource pools and find that transaction costs and information asymmetries can explain why these negotiations breakdown and wells continue to be depleted at higher than optimal rates. Porter (1995) examines the bids of firms participating in U.S. Department of the Interior auctions for offshore drilling permits between 1954–1990 and finds that the substantial differences in valuations for wildcat leases (those for drilling rights in regions with no known oil deposits) can be explained by firm-level differences in drilling and exploration costs. This paper confirms each these findings with regard to the early oil industry.

The remainder of the paper is organized as follows. Section 3.1 contains an overview of the impetus for oil drilling in the mid-19th century, the types of technology available to drillers, and the characteristics of the early oil market in Pennsylvania. Section 3.2 describes wildcat, merchant, and corporate drillers and the differing incentives across these groups of prospectors. Section 3.3 examines the lease and property agreements of the three groups of drillers and how different initial capital endowments led to divergence in rates of extraction and different behavior with respect to conservation. The final section concludes.

3.1 Overview of the Early Oil Industry

Crude oil leakages from fissures in the earth have been cited in literature dating back to antiquity. However, between its earliest recorded detection and the mid-19th century crude oil was largely relegated to medicinal uses with the occasional military application.[1] As the demand for products to illuminate and lubricate the machines and industries of the industrial revolution increased, scientists began to recognize crude oil as a prime candidate to replace the dwindling stocks of whale oils that had initially been used to accomplish these tasks. The growing interest in finding ways to extract and refine crude oil during the 1850s was the result of a combination of supply shortages in the market for whale oils and a growth in demand for illuminants and lubricants. This section describes the early market for illuminants as well as the relevant technologies and property laws that played an important role in early oil exploration.

3.1.1 Market for Illuminants

By the mid-1850s, shortages in the supply of whale oils began to drive up prices and elicit concern both within the whaling industry and among consumers. The *New York Times* reported in 1853 that "many [Pacific] whalemen have indulged in a theory that anticipates a constant falling off of the supply of oil from the grounds in the future."[2] This sentiment was aligned with the fact that the rapidly growing demand for illuminants from East Coast and international markets could not be met by decimated whale stocks in the Atlantic Ocean and declining stocks in the Pacific Ocean. The price of whale oil increased steadily beginning in 1840 and by 1850 was selling for upwards of $2.50 per gallon. The search for new whaling grounds had become so critical to U.S. commerce that in 1852 Senator Seward of New York gave a speech on the floor of Congress to support federal funding to survey the waters of the North Pacific for new whaling grounds, describing "the whale fishery as a source of national wealth and an element of national force and strength."[3] As whale oil

[1] The burning of ships using mineral oil is cited in Homer's *Illiad*.
[2] From the January 10, 1853 article "The Whale Fishery."
[3] From the July 31, 1852 *New York Times* article "The North West Whale Fishery."

supplies began to decline in the middle of the 19^{th} century, the search for alternatives to animal-based illuminants began in earnest.

The growing demand for illuminants was a by-product of the economic expansion occurring in the United States. Between 1849-1859, the U.S. economy grew at an annual rate of nearly 6%, with the most rapid growth occurring in the industrial sectors of the economy.[4] Safe, reliable, and affordable sources of illumination were an increasingly important input for industrial production, as they permitted factories and desk workers to work longer hours and to overcome the seasonality of waxing and waning hours of daylight. It was the combination of the steadily rising demand for illuminants and supply constraints that spurred the search for substitutes for dwindling whale stocks.

Prior to the 1850s, animal grease and vegetable oil provided the only substitutes for increasingly expensive whale oil. Each of these alternatives was also constrained in supply, and both suffered from problems of low quality and provided dirtier sources of light. In response to this fuel shortage, leading scientists throughout the United States were employed in the search for a chemical process to distill illuminants from other sources. In 1854 a patent was filed by Dr. Abrahan Gesner for the manufacture of a new liquid hydrocarbon, which he named kerosene. Initially refined from coal, kerosene began to expand its market from industrial applications to home use by the end of the 1850s. Dozens of kerosene refineries were quickly built, but prices remained high as the refining process from coal solids was expensive and yielded relatively low volumes. Furthermore, the quality of light emitted from coal-derived kerosene was smoky and inferior to that produced by whale oil. The growth of the kerosene market and the search for a cleaner alternative to coal encouraged the financial underwriting of a crude oil exploration effort around reported seepages in Western Pennsylvania in the late 1950s.

3.1.2 Oil Drilling Technology in Western Pennsylvania

The first technique used to extract oil from the ground in Western Pennsylvania was the shoveling of pits in the ground close to natural seepages. Local businessmen interested in selling these mineral oils as medicinal elixirs could capture small quantities of the substance by simply digging around natural fissures in the earth and skimming the oil that separated from the water. A group of New York businessmen, under the leadership of George Bissell, recognized that these mineral oils could offer a cheaper and improved alternative to the coal used to derive kerosene. The group formed a joint stock company called the Pennsylvania Oil Company, which was later renamed the Seneca Oil Company, to invest in both the chemical analysis of refining

[4] Data are from the *Historical Statistics of the United States 1789 - 1945*. Annual income growth stemming from agriculture had an annual growth rate that was over half a percentage point lower than the non-agriculture sectors of the U.S. economy between 1849-1859.

the oils and the subsequent drilling of an exploratory well in the oil region of Pennsylvania. The results of their chemical analysis led them to conclude that refined mineral oil was a commercially viable substitute for both whale oil and kerosene processed from coal solids.

The group hired Edwin Drake to purchase a suitable tract of land in Western Pennsylvania and to drill a well to recover large quantities of mineral oil for processing into kerosene. Drake's innovation was to combine the drilling technology used to drill for saltwater deposits with newly available compact steam engines. A derrick, until then only used by salt miners, was combined with a six-horsepower steam engine to create the basic drilling setup that would remain in use for over a century. The merger of these two technologies allowed a deeper and more uniform shaft to be drilled through multiple strata of rock. A uniform shaft diameter allowed tin tubing to be inserted down from the wellhead prevent the borehole from collapsing during drilling. This innovation also limited the amount of oil that was wasted once oil was struck and helped to retain the pressure in the well to ensure a high overall recovery rate. Drake's original setup cost over $2,000 (approximately $33,500 in 2000 dollars) in materials alone and represented a large up-front investment that only the most well financed prospectors could afford.[5]

The alternative, low-cost drilling method of the mid-19th century was to "kick down" a well.[6] The method consisted of grinding rock by means of repeated blows with an iron chisel bit and then removing the particles from the hole with a bailer. This process was slow (on average they would drill less than three feet per day), labor intensive (reliant on primarily human labor and livestock), and unable to drill to deeper depths. Upon striking oil, wells drilled with this method had a coarse borehole that made for problematic oil extraction. Water seepage and paraffin build-up were the two most common problems associated with "kicked down" wells, each of which resulted in a substantial quantity of wasted oil. A loose well casing also led to the leakage of substantial amounts of natural gas, which lowered well pressure and decreased the total amount of recoverable oil from the underground reservoir. However, since these wells did not rely on large derricks or expensive and difficult to transport steam engines, they could easily be moved from one location to another and were relatively cheap to erect.

On August 27, 1859, Edwin Drake, using his innovative steam-powered drilling design, struck oil at a depth of 69 feet in the hills around Titusville, Pennsylvania. The proceeds from the oil sales allowed the Seneca Oil Company to generate substantial returns for all of its investors and led to the founding of new joint-stock companies to invest in the proven oil drilling technology. News of the successful oil strike also caused an influx of oil prospectors to flood the surrounding region of Western Pennsylvania. In the ensuing decade, tens of thousands of prospectors rushed to Western

[5]See Tait (1941) for a complete description of the technological innovations employed in the original Drake well.

[6]This method of drilling a well was also known as the "springpole method."

Pennsylvania in search of fortunes in the oil drilling business.

3.1.3 Free Entry in the Oil Market and Property Law

The principle of free entry characterized the crude oil market throughout the 1860s. The low-technology option to "kick down" wells was sufficient to strike many of the shallow deposits of mineral oil in the region and an abundance of unexplored tracts of land in Pennsylvania lowered the cost of lease and property agreements for oil exploration in these new areas. Low barriers to entry encouraged many individuals to abandon their professions (predominantly in agriculture) and to drill for oil. Tales of Western Pennsylvania farmers becoming millionaires overnight created a stir along the entire eastern seaboard and induced many to enter the oil drilling business.

Property law in the United States also played an important role for encouraging new entrants into the market. Libecap and Smith (2002) describe the early legal arrangements in oil extraction as follows.

> Because of the fugitive nature of subterranean oil and gas, *in situ* property rights were not assigned to surface land owners, as was done with fixed subsurface mineral resources, but instead were granted only upon extraction or capture, as was done with wild animals (*minerals-ferae naturae*). Every surface owner had the right to vigorously extract oil and gas to reduce it to his possession without violating the rights of neighboring surface owners. This practice led to extractive anarchy.

The standard of 'rule of capture' encouraged landowners of properties adjacent to producing wells to offer lease arrangements to newcomers to the industry intent on profiting from the unknown quantities of wealth below the ground. This frequent drilling of adjacent wells drastically reduced overall yields from oil reservoirs by reducing the natural pressure of the oil field and lowering the amount of recoverable oil.

Free entry into any market increases competition, which typically drives down profits. It was an understanding of this concept that induced many of the industrialists of the 19th century to vertically integrate their corporations and to create cartels to serve as barriers to entry into their markets. Throughout the 1860s and 1870s, John D. Rockefeller, the founder of Standard Oil, would implement numerous strategies to limit free entry into the oil refining and transportation businesses but remained hesitant to engage in oil drilling because of the difficulty of erecting barriers to entry in that portion of the supply chain. Asbury (1941) notes that during the 1860s, "producers' organizations tried a score of times to impose restrictions upon drilling... none of these attempts were notably successful; no sooner did drilling stop at one point that unrestrained wildcatting began at another." Oil prospecting was thus the portion of the oil supply chain that remained free of oligopolisitic controls throughout its early decades of development, which also meant that it remained a chaotic and often wasteful enterprise. Cyclic oversupply became the norm within the industry and in the 1860s there were three oil price spikes and plunges of a magnitude not seen again until the oil crisis of the 1970s. An examination of the behavior and incentives of the least financed groups of drillers in the region shows how their lack of profit maximizing behavior combined with free market entry episodically reduced the profitability of the entire industry.

3.2 Types of Drillers

The drillers in the Venango County region of Pennsylvania in the early 1860s can be classified into three groups: corporate drillers, local merchant drillers, and wildcat drillers. The aforementioned Seneca Oil Company was a prime example of the corporate drillers. Well-financed by urban industrialists with a long investment horizon, corporate drillers could invest in surveyors, purchase large tracts of land, and invest in modern drilling technology. At the opposite end of the spectrum were the masses of under-financed prospectors that I classify as wildcatters. These drillers, limited by their financial resources, could neither purchase the best parcels of land nor invest in modern drilling technology. A lack of financing forced wildcat drillers to make decisions based on a much shorter time horizon. Somewhere in between these two extremes was a group that I classify as merchant drillers. These three groups are examined in detail in the following subsections.

3.2.1 Wildcat Drillers

The wildcatters that came to the Titusville region in the early 1860s were made up of primarily poor farmers from the surrounding counties. Wildcat drillers acquired their name based on their mode of operation that centered on leasing properties in areas without any indication that oil was present. Tait (1946) cites a geologist from the 1860s who characterized the psyche of wildcatters as, "if every cent were taken away from them but the money to keep on hunting oil and a bare living while doing it, not one wildcatter would quit the oil game." Wildcatters fundamental deviation from a revenue or profit maximizing principle, principles which form the basis for all of the economic models that followed the Hotelling (1931) model, were a prime contributor to the massive inefficiencies that beset the oil market in its early years. Since these prospectors comprised the overwhelming majority of drillers to descend onto the newly discovered oil region, their non-maximizing behavior dominated the supply side of the oil market.

The behavior of wildcat drillers deviated from the example set by the Seneca joint-stock company in two significant ways. First, lacking the capital to purchase tracts of land, these drillers were forced to sign lease or sub-lease contracts with property owners. After an up-front payment for the initial right to drill that ranged into the hundreds of dollars, these contracts stipulated royalties ranging between one-twelfth and three-fourths of the oil recovered. These onerous property arrangements often consumed most of the small amounts of startup capital that these prospectors owned. Lower initial levels of capitalization also meant that there was infrequently enough cash to purchase the steam-driven boring technology that had been successfully used by Edwin Drake. This meant that the majority of the wildcat wells were "kicked down," which left the wells predestined for inefficient extraction should they successfully strike an oil deposit.

A second characteristic of wildcat behavior in the early decades that also proved to be especially detrimental to the nascent oil market was to lease or to sub-lease lands adjacent to already producing wells. This was done sometimes in the hope of securing a parallel reserve of oil, but was most often done with the intention of piggybacking off of an established reserve. The uncertainty over the size of oil fields and their geological characteristics made these types of property arrangements particularly risky, given the large up-front payments and royalties demanded by landowners. Despite these risks, numerous wildcatters employed this strategy in the oil region of Western Pennsylvania. If these parallel wells were successful, they would often cause the yields from previously productive wells to diminish significantly and sometimes to cease completely. Though the welfare of all parties would be increased through the use of cooperative strategies to avoid this type of parallel drilling, there is no record of successful oil field unitization in the first decade of exploration in Venango County.[7] Instead, wildcat drillers frequently consumed by the desire to strike oil at any cost, repeatedly engaged in drilling strategies that had low and often negative expected returns.

3.2.2 Corporate Drillers

The corporate drillers in the Pennsylvania region are most easily identified as being funded by large sums of outside capital. The main corporate organizations were established in New York City and Pittsburgh as joint-stock companies. These legal structures allowed investors to assume limited liability and simultaneously raise large amounts of working capital to invest in oil drilling ventures. The large sums of capital permitted these groups to hire geologists, invest in the best drilling technology, and to purchase large tracts of contiguous land that was promising for oil development. Upon striking oil, they were amongst the first prospectors to invest in storage and transportation infrastructure.

The Seneca Oil Company, which backed the successful Drake well, was organized

[7] Oil field unitization is the combination or pooling of revenues from multiple wells that produce from a single reservoir. See Libecap and Smith (1999), for a discussion of different types of incentive compatible agreements that are used to maximize the returns to a known oil resevoir.

on March 23, 1858 with an initial capitalization of $300,000.[8] As described in Section 3.1.2, the direction of George Bissell was instrumental in mobilizing these sums of money to direct research, secure land rights in Pennsylvania, and to invest in the development of a new type of drilling technology. The second non-local corporation to invest in the region was the Columbia Oil Company, which was based in Pittsburgh. Its primary land investment in the region was a 500-acre tract of land situated between Drake's Well and Oil City that was purchased in 1859 for $30,000. Bissell had surveyed and attempted to purchase the same piece of property, indicating consistency in the two corporations assessment of promising tracts, but his purchase offer was bested by the Pittsburgh concern. In 1861 the Company was reorganized as a joint-stock company with an initial capitalization of $200,000. The Columbia Oil Company's primary investors were Andrew Carnegie, Thomas A. Scott, and William Coleman.[9] Their goal upon entering the nascent oil drilling trade was to profit by employing the same orderly business practices that had served them well in the railroad and iron industries.

Though many of the large landowners in Venango County were willing to sell their land outright in the initial years of the oil boom, only a few large joint-stock companies invested in Venango County in the 1860s. Asbury (1941) notes that the these non-native investors were not discriminated against by landowners and had the same level of access to the market enjoyed by the other groups of drillers. Despite this equal access to the market and lower cost of capital, additional pools of investors remained hesitant to commit funds. This reticence, which is analyzed in Section 3.3, can initially be attributed to the slowness with which reliable information flowed out of the oil region and thereafter to the anarchy that characterized the Pennsylvania oil market beset by wildcatters.

3.2.3 Merchant Drillers

Straddling the two extremes of the wildcat drillers and corporate drillers is a group of drillers that I classify as merchant drillers. These prospectors had substantially greater access to initial capital than wildcat drillers, but less and more expensive financing than the joint-stock-financed corporate drillers. Most were from nearby counties and were familiar with the oil region. They often had established business connections in the vicinity, which provided them with some insider knowledge and access to materials and supplies. In numbers, there were significantly more of these merchant drillers than there were corporate enterprises, but they too operated as minority participants in a market dominated by the flood of wildcatters in the region.

Since oil prospecting often relied on local knowledge of natural oil seepages and

---

[8]See Giddens (1938) for a full description of the financiers and financing of the Seneca Oil Company.

[9]Carnegie and Scott had collaborated since the early 1850s in the development of the Pennsylvania Railroad and Coleman was a wealthy Pittsburg iron maker.

other geological features that might indicate promising parcels of land, local merchants had a strong advantage in the early decades of the oil market. Many capitalized on this informational advantage by purchasing tracts of land that were especially promising and that subsequently proved to be highly lucrative. Immediately following Drake's oil strike in 1859, a number of small business owners from neighboring towns entered the oil business and were amongst the most successful early participants. As word spread to the surrounding areas, prominent merchants from further afield also moved to Venango County to undertake drilling enterprises in the first few years of the oil boom.

Most of the merchants that entered the oil prospecting trade in Western Pennsylvania had other successful business ventures in the vicinity that aided their entry into the oil business. Though unable to assume limited liability for their ventures, they still possessed ample sums of capital to invest in their oil enterprises. Many purchased steam-engines and used local connections to help to procure materials and supplies. These men all had experience conducting various forms of business in the Pennsylvania region and used these business connections to their advantage. As experienced businessmen, their understanding of the negative price effects of increases to supply is evident in their drilling behavior, which was more akin to corporate drillers than to the underfunded wildcat drillers.

3.3 Drilling Behavior

The historical records used to understand the property arrangements, wells characteristics, and drilling success on the lands along Oil Creek in Venango County, Pennsylvania come from a variety of sources. Primary source documents written by first hand observers and entries from trade journals and local papers usually offer the most detailed information on the events in the region. Bone (1865), Crum (1911), Gidens (1938), and Asbury (1941) contain narrative descriptions of the properties, landowners, and events in the initial years of oil development. Dufur (1884), Boyle (1898), and Whiteshot (1905) each contain monthly chronologies of events and summaries of statistics from trade journals or newspaper articles. McLaurin (1896) contains an excellent landholding map of the Oil Creek region, which is reproduced in Figure 2.

Table 1 contains a summary of the major lease and property agreements available from the historical record organized by driller classification. The property owners correspond with the original farm owners shown in the map of Figure 2, which provides a useful reference for where each property was located along Oil Creek. When data are available, Table 1 also lists the terms of the property arrangements, the technology used, the maximum output of the wells, and the length of time that the property remained productive. The analysis of the drilling behavior of the different groups of prospectors in this section refers to many of these properties either individually or as

Figure 2
Landholding Map of Oil Creek (Venango, PA)

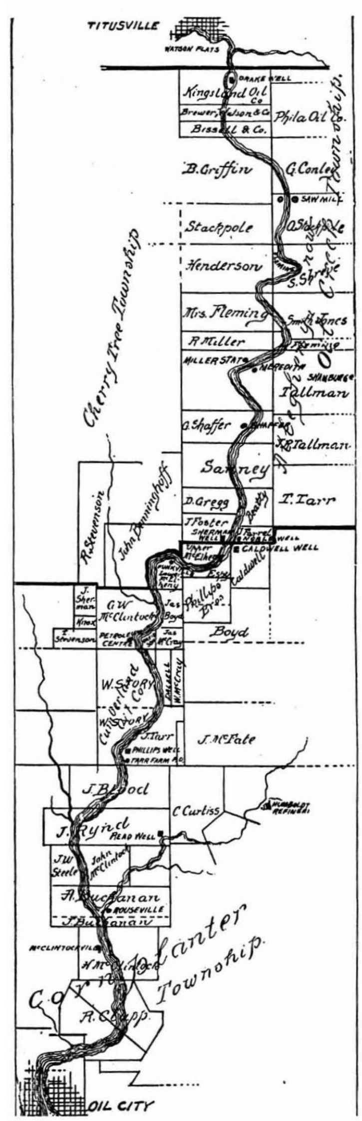

Map is from McLaurin (1896) and shows the approximate location and size of land holdings between 1860-1865.

representatives of the behavior typical of that classification of drillers.

Overall, the data in Table 1 confirm the descriptions of the different driller classification in the previous section. Corporate drillers all purchased large tracts of land outright and drilled their wells using steam engines. Wildcatters tended to lease or sub-lease very small properties and use a mix of kicked-down wells and steam engine drill rigs. Merchants land acquisitions were somewhere between corporate and wildcat drillers; they often drilled on larger parcels of land that were both owned or leased. They predominantly used steam engines to bore their wells, but some resorted to kicked-down wells.

3.3.1 Lease and Property Arrangements

The joint-stock companies, merchants, and wildcatters acquired land and drilling rights in the Titusville region in starkly different manners. After Drake's success, the properties along all of Oil Creek quickly appreciated in value. The rapid growth in land values made it expensive for all but the best financed investors to purchase large tracts of land. The result was the sale of a few large tracts of land to outside investors, almost exclusively corporate drillers, with the remaining tracts of land divided through leasing and sub-leasing contracts.

Sales

The Seneca Oil Company first purchased land in the region around Titusville in 1854. A free flowing oil spring that had been singled out as geologically promising was located on a section of land owned by the lumber firm Brewer, Watson & Company. The 105-acre tract close to the spring was purchased for $5,000 in cash, which was less than $50 per acre. According to Giddens (1938), after the Drake's successful oil strike in 1859, George Bissell proceeded to use as much Seneca Oil Company stock as possible to purchase an additional $200,000 (in 1859 prices) of property in Venango County. In the mid-1850s the Seneca Oil Company's geological surveys of the region had identified a number of promising tracts of land and it was Bissell's intention to purchase as many of them as possible. Upon arriving in the region in late 1859, Bissell noted that "farms that could have been bought for a trifle four months ago, now command $200 & $300 an acre, and when not a drop of oil has been discovered on them... Pittsburgh men consider our property worth millions." Within a few weeks, Bissell had purchased numerous tracts of land and had ordered a second steam engine and pump to drill for oil on these newly acquired properties. The up-front investment in geological surveys, procurement of secured property rights, and acquisition of machinery and equipment were investments in lands that would be systematically developed by the Seneca Oil Company over the next three decades.

The two other large land purchases in the initial years of development along Oil Creek were also undertaken by corporate drillers. In 1859 that the Columbia Oil

Table 1

Property and Extraction Data for Selected Corporate, Wildcat, and Merchant Drillers (1859-1864)

Name	Affiliation	Property Owner	Size	Purchase or Lease	Year	Terms	Drilling Technology	Total Output	Duration
Corporate Drillers									
Seneca Oil	Bissell & Co.	Brewer & Watson	105 acres	Purchase	1856	$5,000	Steam Engine	10 / day	3 years
Penn. Rock Oil	George Bissell	Griffin Farm	~100 acres	Purchase	1859	$20,000	Steam Engine	100s / day	30 years
Columbia Oil	Carnegie & Co.	Story Farm	500 acres	Purchase	1859	$30,000	Steam Engine	100s / day	25 years
Central Petroleum	Bissell & Co.	G.W. McClintock Farm	280 acres	Purchase	1863	$500,000	Steam Engine	1,000s / day	20 year
Wildcat Drillers									
Jonathan Watson	Local Farmer	Hamilton Farm	6 acres	Lease	1859	-	Kicked-down	-	1910
William Phillips	Local Farmer	Tarr Farm	7 acres	Sub-Lease	1860	25% royalty	Steam Engine	4,000 / day	1873
N.D. Woodford	Local Farmer	Tarr Farm	7 acres	Sub-Lease	1860	50% royalty	Steam Engine	2,000 / day	1861
John Fertig	Local Teacher	McElhenny Farm	5 acres	Lease	1861	-	Kicked-down	600 / day	1862
Wildcatters	Adjoining Counties	Numerous Farms	.25 acres-	Lease and	1859-	50-75%	Kicked-down	10-	< 5 years
		(esp. Rynd)	100 acres	Sub-Lease	1864	royalty + bonus		5,000 / day	
Merchant Drillers									
Barnsdall & Rouse	Local Merchants	Parker Farm	small plot	Lease	1859	-	Steam Engine	25-50 / day	-
Barnsdall & Evans	Merchant & Smith	Buchanan Farm	-	Lease	1859	-	Kicked-down	80-100 /day	3 years
William Davidson	Local Merchant	Economite Tract	6,000 acres	Purchase	1860	-	Steam Engine	100 / day	years
John Grandin	Local Merchant	Campbell Farm	500 acres	Lease	1860	$1,000	Kicked-down	300 / day	-
William Abbot	Ohio Merchant	Parker Farm	single well	Purchase	1860	$10,000	Steam Engine	-	-
Abbott & Crossley	Merchant & Smith	Shreve Farm	100 acres	Purchase	1860	$10,000 + 50%	Steam Engine	60 / day	-
Captain A.B Funk	Sawmill Operator	McElhenny Farm	-	Lease	1860	$1,500 + 25%	Small Engine	300 / day	-
Orange Noble	Businessman	Farrell Farm	16 acres	Lease	1860	50% royalty	Small Engine	3,500 / day	6 mo.
Dr. Egbert	Local Physician	Hyde & Egbert Farm	38 acres	Lease	1863	$2,600 + 8%	Kicked-down	3,000 / day	2 years
Jesse Heydrick	Professor	Henry's Bend	200 acres	Owner	1859	NA	Steam Engine	1,800 / day	10 years

Where available, data are from gathered from accounts in Bone (1865), Dufur (1884), Boyle (1898), Giddens (1938), and Asbury (1941).

Company from Pittsburgh purchased a 500-acre tract on the Story farm for $30,000, which amounted to a small per acre premium over the initial Seneca Oil Company purchase price. Then in 1863 the Central Petroleum Company purchased a promising 280 acre tract of land on the farm of G.W. McClintock for $500,000, equivalent to $1,785 per acre. This transaction shows the rapid appreciation of land prices in the region and explains why only the best capitalized organizations were able to purchase land outright.

At first, both the Seneca and Columbia Oil Companies leased their property to merchant and wildcat drillers interested in paying for the right to prospect on parcels of the land. However, by 1860 both companies shifted their land development strategies. Instead of trying to profit on leases to wildcat and merchant drillers, they proceeded to develop the entirety of their properties themselves. Whiteshot (1905) describes the land management of the Central Petroleum Company, "The remarkable results shown by this farm were undoubtedly due quite as much to the excellency of its management as to the superiority of the territory, and stands a bright example of the result of the judicious and economical management of an oil farm." Though there is no record of the discussions that prompted the corporate drillers to shift their land development strategy, the chaos that was developing throughout the region on leased and sub-leased properties offers a potential explanation.

Leases and Sub-Leases

The majority of prospectors did not purchase land outright, but instead leased and sub-leased parcels of land. Boyle (1898) summarizes some of the standard provisions found in the property leases in the Venango County region as follows.

> The characteristics of the leases mostly used during the period extending from 1861 to 1871 inclusive, may be summed up as follows: 1) A great length of term-99 years, or perpetual. 2) A large reservation royalty to be delivered in barrels furnished and paid for usually by the lessee, which royalty shall vest in the lessor as soon as it shall be produced. 3) The right in the lessee to divide the lease into parts and sub-let or assign at his pleasure. 4) The covenant to begin operations within a certain period of time, and in default thereof, or on the cessation of operations at any time for a given period, the lease to become null and void and the land revert to the lessor. 5) The right to abandon, if unprofitable, and thereupon to remove all machinery. 6) The lessee to keep books of account to which the lessor and his agents may have access for inspection. 7) A large cash bonus either paid at the execution of the paper or made payable on some continent happening.

These leases commonly stipulated up-front payments and subsequent royalty payments based on quantities of recovered oil. The following passage from Giddens

(1938) describes the development of lease terms within the initial two years of oil development.

> For a two-acre lease on Oil Creek near Titusville, a landowner received a $200 cash bonus and one-fourth to one-half of the oil. At less desirable points, owners received a bonus of $25 to $50 and one-third to one-eight royalty. Many secured leases during the winters of 1859 for one-tenth of the oil and, in some cases, a still smaller percentage, and no bonus; but such advantageous terms could not be made after the spring of 1860, for prices quickly and steadily advanced in the better localities.

The terms of these agreements varied throughout the 1860s in response to oil spot prices - in times of high prices, royalty percentages were more common, whereas in periods of depressed prices, fixed dollar royalties prevailed.

Wildcatters responded to the incentives inherent in these types of lease arrangements with two strategies. The first strategy was to try to earn back the initial price paid on the lease by pumping a flowing well as quickly as possible. . This meant that wildcat wells were often drilled without regard to the long-term health or productivity of the well. Yergin (1991) describes this aspect of wildcatters behavior as "fueled by the rule of capture - and the race for riches - the wild drive to produce created in the Oil Regions a chaotic scene of heaving populations, of shacks and quick-built buildings." These hastily drilled wells were difficult to control upon striking oil and often impossible to repair in the event of a cessation of oil flow.

The second strategy employed by wildcatters was to move from lease to lease in search of a "gusher," while abandoning active but low productivity wells. Asbury (1941) describes the Oil Creek property market as follows: "prodded by greed and by the landowner, who assumed no risk, but shared handsomely in the gains, and undeterred by prices which were sometimes lower than the cost of production, operators continued to open new territory." Wells with low yielding oil deposits that were abandoned often left flowing streams of oil that polluted the countryside and decreased the remaining natural pressure in the reservoir.

The higher capitalization and local knowledge of merchant drillers provided them with numerous advantages over wildcat drillers. The best capitalized merchants were able to purchase small parcels of land, which allowed them to capture the full upside of the land in the event that they successfully discovered oil. Though it was only a small minority of merchant drillers that were sufficiently well capitalized to make outright land purchases, merchants who leased land could rely on informational advantages in their choice of which property to choose. Though unable to afford the geologists used by corporate drillers, they used their knowledge of the surrounding area to make better informed leasing decisions.

The ability to make larger up-front investments in modern drilling technology also allowed merchant drillers to lease land on better terms. Since property owners recognized merchant drillers' investments in steam-powered derricks as an indication of a higher probability of success, this type of investment would often result in more lucrative lease terms. As owners of established businesses in the region, merchants were less interested in mimicking the wildcat behavior of moving rapidly from region to region in search of new strikes and more willing to deploy fixed capital investments.

An example of the behavior that differentiated merchant drillers from wildcats is evident in the behavior of the group of merchants organized by William Barnsdall that drilled the second producing well in the Titusville region in 1859. Asbury (1941) describes the group that included Barnsdall, a Titusville shoemaker; Henry R. Rouse, a landowner and lumberman; and Boone Mead, a merchant from Warren. Using their knowledge of the area they located a promising piece of land on the farm of James Parker. By pooling their capital, they were able to purchase a steam engine and to drill for oil on this leased plot of land. They first struck oil at a depth of eighty feet. The initial low flow of oil from the well led Rouse and Mead to advocate abandoning the well and searching for another plot of land. However, the patience of Barnsdall and the group's steam-driven drilling apparatus allowed the well's tubing to be withdrawn and the drilling resumed. The well was tubed again at one hundred sixty feet, and on February 19, 1860, began to pump fifty barrels a day. This example of patient development of a tract of land exemplifies how the better-capitalized merchants differed from the wildcat drillers. Their investment in a piece of capital equipment incentivize them to drill their leased plot with greater perseverance than they probably otherwise would have.

3.3.2 Technology and Conservation

The drilling technologies discussed in Section 3.2.2 were not the only technological differences that differentiated drillers of greater and lesser means. Prospectors with larger capital endowments and the ownership of contiguous parcels of land enjoyed economies of scale with regard to storage and transportation infrastructure. The initial costs of these ancillary services could be spread across a number of productive wells. These types of investments allowed better financed drillers to get their product to market more cheaply and with a higher degree of certainty, which provided these drillers with a cost advantage over similarly productive but less integrated wells.

The Columbia Oil Company, under the direction of Andrew Carnegie, took over the development of the Story farm and focused its efforts on building a company town with a systematic strategy for extracting oil that mimicked his methodological expansion of the Pittsburgh Railroads. Tait (1946) notes that the Columbia Oil Company developed the first oil-based company town, paying its employees well and providing "not only homes for workers and their families, but a recreational hall, a reading room, a school house and a brass band." Asbury (1941) identifies this factor as the reason that the property became "the best-managed property on Oil Creek; it was continuously profitable for more than a quarter century. . . in the first six years of its operations the original stockholders received their money back forty-three times."

The systematic development of the Columbia properties stood out amongst all of the producers in the region. Asbury (1941) notes that "none of its wells were among the biggest producers on the creek, although one group of nine yielded a grand total of half a million barrels, but all were carefully and intelligently worked, and no effort was made to rush production and exhaust the field." In order to capture excess flows when there were shortages of barrels or a new well had been struck, the company built a series of large tanks around the property to hold excess oil. A man-made lake of oil that could hold 100,000 barrels of oil subsequently replaced holding tanks and proved to be the best example of conservation during the time period. The lake served the dual purposes of mitigating runoff waste and allowing oil sales to be regulated in response to price changes (Giddens, 1938). These efforts by the Columbia Company resulted in the property remaining profitable and highly productive into the 1890s.

George Bissell developed his lands in a similarly profitable fashion and made capital investments in the first pipelines to move oil from the creek to the transportation hub on the river. By 1865, Bissell's enterprise had also become the largest barrel manufacturer in region. His corporate drilling operations, consolidated in the Seneca Oil Company and Central Petroleum Company, were able to utilize this infrastructure built during periods of high prices to remain profitable through periods of depressed oil prices.

Wildcatters' near uniform lack of investment in technology or conservation measures can be traced through the lease contracts that they signed to their low initial capitalization. Lease arrangements that had lower up-front payments and higher royalty rates incentivize wildcatters to use their inability to purchase heavy machinery to their advantage. Drillers operating manual- or animal-driven jig systems that "kicked down" wells could quickly relocate to what was perceived as a more lucrative location. The low barriers to entry in the early oil market facilitated this strategy. The result was a market in which the majority of participants were interested in the rapid development of a leased property over a short time horizon.

Corporate drillers used their higher capital endowments to increase their initial probability of success and to help ensure long-term profitability of their investments. By employing the most modern drilling technology some merchants were able to survive through periods of especially low prices by slowing extraction and instituting conservation measures while prices remained below profitable levels. However, their experience in the 1860s emphasized that while well-funded investments in the drilling business could be both successful and profitable, the free entry of lesser-capitalized drillers with worse incentives made this venture very risky over the medium-term. Many merchant drillers that had made substantial up-front investments in drilling equipment were bankrupted during the violent price swings that characterized the early market.

Table 2

Number of Pennsylvania Crude Oil Wells and Approximate Output

Date	Shafts Begun	Operating Wells	Former Wells	Present Output (barrels/day)	Total Output to Date (barrels)	Price at Well (per barrel)
December, 1859	6	2	-	10	2,000	$20
November, 1860	300	74	-	50	500,000	$9
August, 1861	800	-	-	-	-	-
June, 1862	495	75	62	5,717	1,000,000	$0.10
April, 1865	2,000	435	~200	3,500	8,784,000	$10

The data are compiled from the Venango County Spectator, Ashbury (1941), Tait (1946), Giddens (1938), Logan (1930), and Dufur (1884).

3.3.3 Market Failure and Wasted Oil

Waste was endemic in the initial years of oil prospecting. Drilling technologies were primitive, geological knowledge was limited, and the infrastructure to store excess supplies of oil was underdeveloped. This section highlights how the low capitalization of wildcatters that led to low rates of technological utilization constituted a form of market failure that incentivized them to extract oil at rates that were detrimental to both their individual profitability as well as the profitability of the more prudent drillers. Much of the waste during the early years of the oil business can be attributed to the numerous wildcatters in a market dominated by free entry.

It is impossible to ascertain exact extraction rates from this time period. Much of the data is qualitative, as instruments to measure flow rates were not often employed and differing sources cite often substantially different flow rates for the same wells. Throughout the 1860s, oil was most often transported in wooden barrels from the derricks in Oil Creek to the transport centers on the nearby Allegheny river. The size of these barrels varied in the early years, ranging between forty gallons and the eventual industry standard of forty-two gallons. As barrels were the predominant method of capture, storage, and transportation from these early wells, it also became the preferred measure of a well's productivity.

Table 2 contains data on the number of wells and their total output during the first five years of development of the oil regions throughout Pennsylvania. It shows that over the first five years of drilling only about a fourth of shafts that were drilled resulted in wells that ever yielded oil. It also shows that a high percentage of wells quickly transitioned from productive to unproductive. The most common reason for these instances of transition was the use of inferior drilling technology. Asbury (1941) cites a statistic that captures the overall inefficiency market that "in 1869 only one well which had been struck in 1859, and ten which had begun to yield in 1860, were still producing."

The most prolific well and one of the best examples of poor reservoir management was struck by William Phillips, a local wildcatter, on a sub-leased portion of the James Tarr farm on November 1, 1860. Its output of four thousand barrels per day single-handedly depressed oil prices in the United States for the next two years. The market was unable to increase demand in response to this jump in output, yet Phillips made no effort to stem the flow of oil. Simultaneously, the owner of the property, in the interest of maximizing his short-term profit, sold scores of leases with no restrictions on subleasing in the months following Phillips' strike. The result was that nearly the entire Tarr property was subsequently subdivided into progressively smaller tracts of land. Within a few years, over two hundred wells were begun on the farm, leading to further overproduction, a decrease in well pressure of the common pool, and massive amounts of wasted oil. Whiteshot (1905) describes the interaction between the Phillips well and the adjacent Woodford well. "They were situated within two rods of each other and the subterranean connection between their source of supply was so manifest that when the Woodford pumped the only remedy left the Phillips was to let the surface water down to shut off the oil from both. Enormous sums were offered by both owners, but as either had it completely in his power to render the property of the other worthless, neither was willing to settle until both wells were nearly ruined by the surface water."

The Tarr property was just one of numerous examples of properties and pools of oil being mismanaged and sometimes sabotaged by irrational neighbors. Giddens (1938) discusses the common occurrence for "those with wells pumping from five to twenty barrels a day found it most discouraging when an adjoining well spouted hundreds of barrels, flooding the market and making the operation of pumping wells unprofitable." In addition to being discouraging, a flood of wildcatters onto an adjacent property often spelled disaster for the common pool because of decreases in well pressure. Yergin (1991) notes that "flush production often damaged the reservoirs, leading to premature exhaustion of gas pressure, and thus far less recovery than would otherwise have been the case."

This lack of coordination among competing drillers led to the sub-optimal equilibrium wherein short-term production was increased at the cost of lower potential medium-term production. This led to fluctuating price dynamics that further depressed large-scale investment in oil production and further increased wildcatters' influence on the oil market. This cyclic volatility in prices, seen in Figure 1, which persisted through the first decades of the oil era resulted in phases when numerous drillers "abandoned leases, and fled in despair, leaving machinery, buildings, wells, in fact, everything" (Giddens, 1938). Along with a high rate of turnover within the industry, this boom-and-bust cycle created a growing reluctance on the part of more rationally-minded groups to reenter the market when prices stabilized. Those groups with the greatest marginal incentive to enter the drilling business, namely corporate drillers, were the most reluctant to do so for fear of another bout of unconstrained production. Outside capital's reluctance to use its market power to coordinate the disorganized drilling market made all market participants worse off.

The persistent threat of unconstrained wildcatting discouraged numerous large investment groups from entering the nascent oil drilling industry. Fearing a glut from nearby producers, additional corporate and merchant drillers were hesitant to invest large sums of capital into new properties. Though the Columbia and Seneca properties were able to remain profitable during periods of depressed prices, that can partially be attributed to their early entry into what had subsequently become an increasingly disorganized and risky enterprise. An aggregate under-investment in supply regulating technologies and persistent property arrangements that led to bad incentives continually encouraged excessive waste and cyclical declines in prices. Poole (1999) cites a famous example of a prominent businessman's discontent with the state of the market, noting that though the Columbia property was "probably the best-managed property on Oil Creek... Andrew Carnegie grew tired of the messy and chaotic state of the oil business and sold out in 1865."

3.4 Conclusion

This paper provides an explanation for why a nearly perfectly competitive market was unable to maintain a profitable level of output for a non-renewable resource. The behavior of the oil market between 1859 and 1870 defied the predictions made by the standard economic theory of Hotelling (1931) that the depletion of a non-renewable resource in a competitive market would occur at a socially optimal rate. Instead production was highly cyclical with decisions to maximize short-term supply simultaneously leading to medium-term undersupply. This paper's explanation for these dynamics is that cohorts of drillers with different levels of capital endowments were incentivized to make decisions to maximize production over different time horizons. In particular, the historical documents show that the incentives of wildcat drillers, who had very low levels of startup capital, distorted the behavior of the entire oil market. Since they drilled rudimentary wells on leased properties, wildcatters had an incentive to maximize short-term production regardless of market prices. Their high mobility and myopia led to massive overproduction across the industry and wasted the supplies of a limited resource.

This description of the dynamics of the early oil industry also explains why the oil drilling business was the last segment of crude oil production to be combined into the vertically integrated oil corporations. By the mid-1870s Standard Oil had used its market power in refining to bring transportation and distribution under its control. Yergin (1991) describes Rockefeller's firm as having "acquired its own tracts of land on which it grew the white oak timber to make its own barrels; it also bought its own tank cars, and its own warehouses in New York, and its own boats on the Hudson River." However, the firm remained reluctant to enter the production portion of the supply chain. One of members of Standard's Executive Committee wrote Rockefeller in 1885, "Our business is that of manufactures, and it is in my judgment, an unfortunate thing for any manufacturer or merchant to allow his mind to have the care and friction which attends speculative ventures" (Yergin, 1991).

It was only against great internal opposition that Rockefeller decided in the 1880s, nearly thirty years after the discovery of oil in Pennsylvania, to enter the business of drilling for oil. The incentives-based explanation for behaviors seen in the 1860s also provides an explanation for why Standard Oil chose to enter oil drilling at the time that it did. Thirty years into the era of oil, much of the readily accessible oil (oil that could be accessed with small operations and limited drilling technological) in the United States had already been located and extracted. The deeper and more challenging oil reservoirs that remained to be discovered in the late 19th- century required initial investments that dwarfed those used in Western Pennsylvania during the 1860s. This was exactly the barrier to market entry that Standard Oil desired. Recognizing that oil drilling was transitioning to an industry dominated by large incumbents, Rockefeller aggressively moved his company into the final portion of the supply chain. "By 1891, though virtually absent from production a few years earlier, Standard was itself responsible for a quarter of America's total output of crude oil" (Yergin, 1991).

Both in describing the behavior of the market in the early years and in providing an explanation for the timing of the entry of large corporations into the market, the model proposed in this paper provides a new insight into the dynamics of the early oil market. It shows how heterogeneity in suppliers' initial capital endowments had far reaching effects on the overall supply and price volatility of the market. A promising avenue of future research would be to investigate the applicability of this model to describe the extraction behavior for other non-renewable resources.

Bibliography

[1] Arrow, Kenneth J. and Sheldon Chang, 1978, Optimal Pricing, Use and Exploration of Uncertain Natural Resource Stocks, Department of Economics Technical Report, Harvard University.

[2] Asbury, Herbert, 1941, *The Golden Flood* (Knopf, New York).

[3] Bone, J. H. A., 1865, *Petroleum and Petroleum Wells* (J. B. Lippincott & Co., Philadelphia).

[4] Boyle, P. C., 1898, *The Derrick's Hand Book of Petroleum: A Complete Chronological and Statistical Review* (Derrick's Publishing, Oil City, PA).

[5] Crum, A. R., 1911, *The Romance of American Petroleum and Gas* (M. M. Marcy, New York).

[6] Devarajan, Shantayanan and Anthony C. Fisher, 1981, Hotelling's "Economics of Exhaustible Resources": Fifty Years Later, *Journal of Economic Literature* 19, 65–73.

[7] Dufur, W. H., 1884, *Handbook of Petroleum* (Derrick Publishing, Oil City, PA).

[8] Giddens, Paul H., 1938, *The Beginnings of the Petroleum Industry Sources and Bibliography* (Pennsylvania Historical Commission, Harrisburg, PA.).

[9] Gilbert, Richard J.,1979, Optimal Depletion of an Uncertain Stock, *The Review of Economic Studies* 46, 47–57.

[10] Harri, Kenneth, 1987, The Wildcatter: A Portrait of Robert O. Anderson (Weidenfeld & Nicolson, New York).

[11] Hoel, Michael, 1978, Resource Extraction, Uncertainty, and Learning, *The Bell Journal of Economics* 9, 642–645.

[12] Homer, The Illiad, translated by Richard Lattimore (University of Chicago Press, Chicago).

[13] Hotelling, Harold, 1925, A General Mathematical Theory of Depreciation, *Journal of the American Statistical Association* 20, 340–353.

[14] Hotelling, Harold, 1931, The Economics of Exhaustible Resources, *The Journal of Political Economy* 39, 137–175.

[15] Johnston, Louis D. and Samuel H. Williamson, 2005, The Annual Real and Nominal GDP for the United States, 1790 - Present, *Economic History Services*, available at (http:// www.eh.net/hmit/gdp/).

[16] Lewis, Tracy R., 1979, The Exhaustion and Depletion of Natural Resources, *Econometrica* 47, 1569–1572.

[17] Libecap, Gary D. and James L. Smith, 1999, The Self-Enforcing Provisions of Oil and Gas Unit Operating Agreements: Theory and Evidence, *Journal of Law, Economics, & Organization* 15, 526–548.

[18] Libecap, Gary D. and James L. Smith, 2002, The Economic Evolution of Petroleum Property Rights in the United States, *The Journal of Legal Studies* 31, S589–S608.

[19] Libecap, Gary D. and Steven N. Wiggins, 1985, The Influence of Private Contractual Failure on Regulation: The Case of Oil Field Unitization, *The Journal of Political Economy* 93, 690–714.

[20] Logan, Leonard M., 1930, *Stabilization of the Petroleum Industry* (University of Oklahoma Press, Norman, OK).

[21] McLaurin, John J., 1896, *Sketches in Crude Oil* (published by the author, Harrisburg, PA).

[22] Pindyck, Robert S., 1979, Uncertainty and Exhaustible Resource Markets, Massachusetts Institute of Technology Working Paper.

[23] Poole, Keith T., 1999, Lecture: Entrepreneurs and American Economic Growth: Andrew Carnegie, available at (http://voteview.com/carnegie.htm).

[24] Porter, Robert H., 1995, The Role of Information in U.S. Offshore Oil and Gas Lease Auction, *Econometrica* 63, 1–27.

[25] Stiglitz, Joseph E., 1975, The Efficiency of Market Prices in Long-Run Allocation in the Oil Industry, in Gerard Brannon: *Studies in Energy Tax Policy* (Bellinger, Cambridge, MA).

[26] Stocking, George W., 1925, *The Oil Industry and the Competitive System; A Study in Waste* (Houghton Mifflin, New York).

[27] Sweeney, James L., 1977, Economics of Depletable Resources: Market Forces and Intertemporal Bias, *The Review of Economic Studies* 44, 125–141.

[28] Tait, Samuel W., 1946, *The Wildcatters, An Informal History of Oil-Hunting in America* (Princeton University Press, Princeton, NJ.).

[29] Tisato, Peter, 1995, Exhaustible Resource Depletion: A Modified Graphical Approach, *The Journal of Economic Education* 26, 51–57.

[30] Whiteshot, Charles A., 1905, *The Oil Well Driller: History of the Oil Industry of the World* (published by the author, Mannington, WV).

[31] Yuan, Lasheng, 2002, Divide and Conquer: Multiple Leasing in Common Pool Oil Fields, *The Canadian Journal of Economics* 35, 36–51.

[32] Yergin, Daniel, 1991, *The Prize: The Epic Quest for Oil, Money & Power* (Free Press, New York).

Appendix A

Appendix

A.1 Construction of the Carbon Dataset

Matching facility-level data to the proper owner is a multistep process. The registry data are organized by country, with separate fields for each facility listing the name of the owner and the name of the installation. An example in the German registry data is a power plant at a BMW factory in Dingolfing with the operator name 'Bayerische Motoren Werke Aktiengesellschaft' and an installation name 'BMW Werk 2.4 Dingolfing Kesselhaus.' This facility is assigned to the German-listed BMW equity data. Facilities with less readily identifiable owner names are matched to firms using the facility's unique permit identification number. This number can be linked to a CITL ownership file listing the name and e-mail address of the account holder. For facilities with multiple owners, which includes many large power generation facilities that are owned by consortia, company annual reports are used to determine ownership percentages.

The 146 firms classified as 'Other Publicly-Listed' are excluded from the analysis for one of three reasons:

1. 49 firms that have a permit allocation of under 10,000 metric tons are excluded because of the low cash implications of low emissions and also the high transaction costs that such firms face when participating in a cap-and-trade system. Jaraite, Convery, and Di Maria (2009) use survey data to estimate the transaction costs incurred by Irish firms during the first phase of the ETS and find that firms with very small allocations often do not have the resources to efficiently manage their permit holdings. The Waxman-Markee legislation passed by U.S. House of Representatives in 2009 also acknowledges the high transaction costs and low abatement potential for small facilities and sets an annual threshold of 25,000 metric tons for regulation.

2. 13 firms have a Level 3 ICB industry classifications that is excluded because those industries are not normally associated with greenhouse gas emissions (as

evidenced by the low number of regulated firms in each industry classification). The excluded industries are Financials, Insurance, Real Estate, Transportation, Telecommunications, Retail, and Media.

3. 84 firms have insufficient returns in either the event period or estimation period, are missing annual financial information, have limited floats, or are traded on pink sheets. Firms without actively traded equities during 70 days of the estimation period and the entire event period are excluded because the abnormal returns cannot be calculated with sufficient certainty. Without annual DataStream financial information it is impossible to calculate market capitalization or control for firm characteristics. Equities with small floats or that are traded on pink sheets often have low or no correlation with the market return during the estimation period and thus have unreliable abnormal return estimates during the event period.

A.2 Assumptions of portfolio choice model

1. *Households have fixed-rate mortgage contracts with no balloon payments.* Many of the model's implications hold for adjustable-rate mortgages (ARMs), though the rate of return on debt repayment of ARMs is uncertain. Balloon payments are assumed away to simplify the algebra and because they are very uncommon.

2. *Households do not default on their mortgage.* It is also sufficient to assume that households believe the probability of default to be zero. In either case, our results remain robust when our sample is restricted to wealthier households for whom default rates are very low.

3. *Households cannot take short positions in equity.* It is sufficient to make this assumption only for households with a mortgage and a mortgage rate higher than the expected return on the risky asset. This assumption is reasonable, given that households with such high mortgage rates are unlikely to have the credit necessary to take short positions in equity.

4. *The refinancing decision is independent of the repayment decision.* That is, households refinance when it is optimal to do so and then make their investment decisions given their new mortgage interest rate. The refinancing option can also shorten the duration of early mortgage repayment.

5. *The inflation rate is equal to zero.* This is done to simplify our calculations. A non-zero rate of inflation does not affect our results.

6. *There are no taxes.* This is done to simplify the algebra. We incorporate the tax benefits of mortgage debt into our empirical analysis. See Appendix A.3 for a detailed explanation of how after-tax mortgage rates are constructed.

A.3 Construction of after-tax interest rates

This appendix describes the construction of after-tax interest rates. For mortgages, home equity lines of credit, and other home equity loans, we use historical marginal tax rate data collected by the Tax Foundation.[1] We merge the federal marginal tax rate τ_f^i using each household's reported marital status and gross income, with the assumption that all married households file jointly. We define the after-tax interest rate for household i as:

$$\left(1 - \tau_f^i - \tau_s\right) \ R^i_{d,\tau=0} \ , \tag{A.1}$$

where τ_f^i is the marginal federal tax rate for household i, τ_s is the marginal state tax rate for which we assume a value of 0.05, and $R^i_{d,\tau=0}$ is the pre-tax mortgage interest rate (or home equity loan interest rate) of household i. For households with multiple mortgages or home equity loans, we report the highest rate since repayment of this account offers the highest rate of return. These estimates of after-tax interests rates are conservative, given that not all households itemize their deductions and that the lower alternative minimum tax (AMT) rate is increasingly the applicable rate for many households.

The after-tax interest rates on credit cards are equal to the nominal interest rates since credit card debt offers no tax advantages. The SCF surveys the interest rate on the credit card account with the highest balance, which may not be the highest interest rate credit card used by the household.

[1] http://www.taxfoundation.org/taxdata/show/151.html.