Innovation, Diffusion, and Regulation in Energy Technologies

by

Theodore Robert Fetter

University Program in Environmental Policy
Duke University

Date: ______________________
Approved:

____
Lori Snyder Bennear, Supervisor

____
William A. Pizer

____
Steven E. Sexton

____
Christopher Timmins

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the University Program in Environmental Policy
in the Graduate School of Duke University
2017

Abstract

Innovation, Diffusion, and Regulation in Energy Technologies

by

Theodore Robert Fetter

University Program in Environmental Policy
Duke University

Date: ____________________
Approved:

____
Lori Snyder Bennear, Supervisor

____
William A. Pizer

____
Steven E. Sexton

____
Christopher Timmins

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy
in the University Program in Environmental Policy
in the Graduate School of Duke University
2017

Abstract

The innovation and diffusion of new technologies is one of the central concerns of economics. New inventions or technological combinations do not spring fully formed into the world; as firms encounter and learn about new technologies they experiment, refine, and learn about them, improving productivity (and sometimes earning economic rents). Understanding the processes by which firms learn, and how these processes interact with regulations, is fundamental to understanding the emergence of new technologies, their contribution to growth, and the interaction of innovation and regulation.

This dissertation addresses how firms learn and respond to regulations in the context of emerging technologies. Within this framework, I address several questions. When production inputs are socially controversial, do firms respond to disclosure laws by voluntarily constraining their inputs? Do these public disclosure laws facilitate knowledge transmission across firms, and if so, what are the implications for public welfare—for instance, do the gains from trade outweigh any effects of reduced incentives for innovation? I study these questions in the context of hydraulic fracturing, though the results offer insight for more general settings. Panning out to a much broader view, I also explore how energy-related technologies—in both generation and consumption—diffuse across national boundaries over time, and whether innovation and diffusion of energy-efficient technologies has led to more or less energy-efficient economic growth.

In my first paper, I contribute to improved understanding of the conditions in which information-based regulations, which are increasingly common in multiple policy domains, decrease externalities such as environmental pollution. Specifically, I test whether information disclosure regulations applied to hydraulic fracturing chemicals caused firms to decrease their use of toxic inputs. Prior to these mandatory disclosure laws, some operators voluntarily disclosed fluid components for some or all of their wells. I compare the chemical mixtures used prior to the mandatory disclosure laws to those used after the laws took effect, using a difference-in-differences method motivated by the difference in timing of state-level disclosure laws. I use voluntary disclosures to measure the toxicity of fluids prior to mandated disclosure, and thus observe a composite effect of both full reporting and disclosure pressure. These effects likely have opposite signs; I employ several methods to tease them apart so that I can separately identify the effect of disclosure pressure. My analysis, which covers over 70,000 wells in seven states, suggests that state disclosure regulations resulted in a large and persistent decrease in the use of toxic and regulated chemicals in fracturing fluids. This is not the first paper to find that disclosure regulations can change firms' behavior, but it demonstrates such an effect in a setting in which consumer or market pressure is minimal or nonexistent: firms that produce undifferentiated products for an intermediate market, and disclosure policies that do not generate readily accessible or interpretable information.

The second paper tests whether disclosure laws facilitated the transmission of useful knowledge across companies. It is well established that economic agents learn about new technologies in part from other adopters, though even sophisticated firms may not take full advantage of social learning. With my co-authors, I examine whether firms took advantage of environmentally-focused disclosure laws to learn from competitors and improve productivity. We find evidence that they did: following mandatory disclosure we observe convergence in productivity per well, in production inputs, and strong evidence of a link between the two. To our knowledge this is the first study to examine this pathway for social learning in an emerging technology. This could also be interpreted as a form of technology diffusion facilitated by environmental regulation.

In my third paper, I address a broader scale of technology change, looking for evidence that improved technologies for energy generation and consumption have allowed less energy-intensive or pollution-intensive growth in developing countries. I analyze panel data on Gross Domestic Product (GDP) and national energy consumption to look for evidence of technology "leapfrogging" (i.e., decreased intensity of energy consumption for a given level of economic growth). I combine 1960-2014 data on energy consumption from the International Energy Agency with historical data that extends back to 1861 for several countries on energy consumption and fuel source, as well as GDP. I compare countries at the same income level and test whether energy consumption and energy intensity are different for today's less-developed countries compared to today's industrialized countries when they had similar income levels. Compared to prior analysis, my much longer time series allows me to test for leapfrogging over a scale appropriate to the pace of widespread technological change.

Contents

List of Tables

List of Figures

List of Abbreviations and Symbols

Abbreviations

BLM	Bureau of Land Management
CAS	Chemical Abstracts Service
DD	Difference in Differences
DEP	Department of Environmental Protection
DOGGR	California Division of Oil, Gas, and Geothermal Resources
EDWIN	Exploration and Development Well Information Network
EPA	Environmental Protection Agency
ft	Feet
gal	Gallons
GDP	Gross Domestic Product
GWPC	Groundwater Protection Council
IC	Industrialized Country
IEA	International Energy Agency
IOGCC	Interstate Oil and Gas Compact Commission
IQR	Inter-quartile range
LDC	Less-Developed Country
MCF	Thousand cubic feet
NGO	Non-Governmental Organization
OECD	Organization for Economic Cooperation and Development

OLS	Ordinary least squares
PADEP	Pennsylvania Department of Environmental Protection
ppm	Parts per million
PPP	Purchasing Power Parity
PTRC	Priority Toxic or Regulated Chemical
QS	Quality-Similarity
RSEI	Risk-Screening Environmental Indicators
SDWA	Safe Drinking Water Act
TEU	Total Energy Use
TFC	Total Final Consumption
TJ	Terajoules
TPES	Total Primary Energy Supply
TRI	Toxics Release Inventory
UIC	Underground Injection Control
USD	United States Dollar
VR75	Voluntary Reporter, 75 percent: Operators that report fracturing fluid contents for at least 75 percent of horizontal wells prior to mandatory disclosure regulations
VR90	Voluntary Reporter, 90 percent: Operators that report fracturing fluid contents for at least 90 percent of horizontal wells prior to mandatory disclosure regulations

Acknowledgements

When my wife came home from campus after defending her Master's thesis and quickly picked up a book in the same field "for fun," I reflected that she was the sort of person who ought to enter a PhD program, because she loved one particular field so much that she would happily spend all of her days and nights learning about it and no other. I had too many disparate interests, I thought, and loved all of them, too much to pick just one to focus on for the time it would take to complete a doctorate. I was reflecting on this memory not long ago as I came home from my first presentation at the National Bureau of Economic Research (NBER) and, for fun, sat down to watch some videos on measure theory, Hilbert spaces, and other exciting mathematical topics. Evidently I've changed quite a bit, in preferences or self-understanding or both: I'm still interested in many things, but I am much more focused on a smaller set. More importantly, I'm convinced I can indulge my most central interests, sequentially or simultaneously, in my new career.

A number of people have been instrumental in supporting me through that evolution. First and foremost, my committee at Duke—Lori Bennear, Chris Timmins, Steve Sexton, and Billy Pizer—have been tremendously helpful throughout my time here. There is no chance I would be here were it not for their guidance and support. A number of other faculty at Duke have offered useful critiques and insightful comments, both on the ideas that ultimately made it into the dissertation and many more that did not. These include Ed Balleisen, Ronnie Chatterji, Bentley Coffey, Marc

Jeuland, Brian Murray, Richard Newell, Subhrendu Pattanayak, Alex Pfaff, Amy Pickle, Marty Smith, Jeff Vincent, Erika Weinthal, Jonathan Wiener, and Daniel Xu. I've also benefited from discussions with, and mentorship from, several faculty at Yale, including Matt Kotchen, Ken Gillingham, Eli Fenichel, and Rob Mendelsohn. Finally, I've had many valuable discussions and suggestions from colleagues at conferences and workshops, including the PhD Academy on Sustainability and Technology at ETH Zurich, and conferences of the Association for Environmental and Resource Economists (AERE), the Alliance for Research on Corporate Sustainability (ARCS), the Academy of Management (AOM), and the NBER Environment and Energy Economics Program.

I also would not have succeeded as I have without generous financial support from several worthy organizations and programs. These include Resources for the Future (for the Joseph L. Fisher Doctoral Fellowship), the Nicholas Institute for Environmental Policy Solutions (for the Duke Environmental Economics Doctoral Scholars Fellowship), the Duke University Energy Initiative, the Yale Center for Environmental Law and Policy, and Equitable Origin. I am deeply grateful to each of these organizations for supporting my research.

I'm grateful, too, to a number of professors who have assisted me in various ways that may not be directly related to the dissertation, but have helped guide me in some substantial element of the process. These include Michael Ash, Jim Boyce, Ben Cashore, Julie Caswell, Jonathan Feinstein, Dan Fetter, Gordon Geballe, Doug Kysar, Dan Lass, Paul Mohai, Manuel Pastor, Ken Small, John Stranlund, Heidi Williams, and Amy Wrzesniewski. I've received similar types of support from fellow students (some of whom are now faculty), including Jesse Burkhardt, Nathan Chan, Jeffrey Chow, Peter Christensen, Katy Hansen, Matt Johnson, David Kaczan, Namrata Kala, Joowon Kim, Justin Kirkpatrick, Honggi Lee, Yating Li, Jenny Orgill, Craig Palsson, Chris Paul, Brian Prest, Andrew Steck, Stephanie Stefanski, and

Faraz Usmani. A number of other friends have offered useful moral support as well. There is no way I can list them all, but I would particularly like to call out Charles Vogl and Socheata Pouev (who served me delicious homemade dinners, dozens of times, when I was working too hard to take good care of myself), as well as Eric, Marikler, and Daniel Toensmeier; Gabriel Grant and Sarah Townsend-Grant; Miriam and Carol Gross; and Jasmine Hyman, Sharon Smith, and Laura Bozzi.

My parents taught me to value education and hard work, but more importantly they taught me to stay engaged and work for positive change in the world. My path has perhaps been less straightforward than they may have hoped, but throughout all of my adventures I have done my best to honor their investment in me, and their wise advice, by engaging with the world in constructive and creative ways.

Finally, I am grateful to Nicole for allowing me the freedom to tear up the floorboards underpinning my life once more, leave a perfectly good job (or two), and go on this seemingly crazy PhD journey. I'm deeply thankful for her sage advice on a daily basis (even the days I don't take it), and for the countless hours she has spent patiently counseling me through the rough parts. I do believe we've come out closer for it all, which is a testament to her wise and flexible thinking and her incredible reserves of kindness, patience, and love. As ever, I'm looking forward to our next adventure.

1

Introduction

The innovation and diffusion of new technologies is one of the central concerns of economics. New inventions or technological combinations do not spring fully formed into the world; as firms encounter and learn about new technologies they experiment, refine, and learn about them, improving productivity (and sometimes earning economic rents). Understanding the processes by which firms learn, and how these processes interact with regulations, is fundamental to understanding the emergence of new technologies, their contribution to growth, and the interaction of innovation and regulation.

This dissertation addresses how firms learn and respond to regulations in the context of emerging technologies. Within this framework, I address several questions. When production inputs are socially controversial, do firms respond to disclosure laws by voluntarily constraining their inputs? Do these public disclosure laws facilitate knowledge transmission across firms, and if so, what are the implications for public welfare—for instance, do the gains from trade outweigh any effects of reduced incentives for innovation? I study these questions in the context of hydraulic fracturing, though the results offer insight for more general settings. Panning out to a much broader view, I also explore how energy-related technologies—in both generation and consumption—diffuse across national boundaries over time, and whether innovation and diffusion of energy-efficient technologies has led to more or less energy-efficient economic growth.

1.1 Summary of papers

In Chapter 2, "Fracking, Toxics, and Disclosure," I test whether firms engaged in hydraulic fracturing for oil and gas responded to information disclosure regulations in part by reducing their use of toxic and regulated chemical additives. The chemical additives used in hydraulic fracturing have inspired substantial public concern, as well as some regulatory concern, especially because in the early years of the emerging technology companies withheld information about chemical formulas from regulators and the public—as well as from one another. I exploit differences in state-level regulatory timing to estimate a causal effect of disclosure regulations, inferring pre-regulation chemical use from voluntary reports that companies made prior to mandatory disclosure. I find that firms did reduce their use of toxic chemicals in response to the mandatory disclosure regulations; although the effect takes time to manifest, about nine months to a year, it is both persistent and large in magnitude. I also observe a simultaneous increase in the quantity of chemicals that companies declare as "proprietary," although this appears to be a secular trend and not necessarily caused by the mandatory disclosure regulation. This paper adds to our understanding of how firms respond to information disclosure laws in non-consumer-facing settings, which has rarely been studied using methods that can accurately assess causal effects.

In Chapter 3, "Learning by Viewing? Social Learning, Regulatory Disclosure, and Firm Productivity in Shale Gas," I examine whether shale gas operators took advantage of these same information disclosure laws to learn from their competitors and improve productivity—that is, the extent to which disclosure facilitated knowledge transfer across companies. With my coauthors, I exploit an unusual episode in Pennsylvania in which the regulatory body collected chemical input data for a 14-month period (the data were technically public, but quite difficult to access). Using data collected from requests under the state Right-To-Know Laws, as well as additional data from public disclosures later, we study how the change in disclosure regime affected operators chemical use and well productivity. Our results suggest the operators who were least productive prior to the public disclosure period did take advantage of the social learning opportunity: they improved productivity faster than the firms that were initially more productive, and their production inputs grew more similar to the firms that were more productive ex ante. This finding suggests information disclosure regulations created a pathway for social learning—a form of technology diffusion facilitated by regulation. It also suggests firms may have been accurate in claiming that mandatory information disclosure could threaten the competitive advantage of technology leaders.

In my final chapter, "Energy Transitions and Technology Change: 'Leapfrogging' Reconsidered," I address a broader scale of technology change, and test whether improved technologies for energy generation and consumption have allowed less energy-intensive and less pollution-intensive growth in developing countries. This paper builds on a recent analysis of "energy leapfrogging" by Arthur van Benthem, analyzing a data series on national GDP and energy consumption extending back to 1861 for some countries (improving on the 47-year series from the earlier paper, and also using about 50 more countries). My extended time series allows me to test for efficiency improvements over a scale that is more appropriate to the pace of widespread technological change, and in my analysis I find evidence of "leapfrogging" where van Benthem does not. This suggests that over the long run, efficiency improvements dominate rebound effects and industrial outsourcing; this, in turn, has implications for long-run energy demand forecasting.

1.2 Common themes

Taken together, these papers address important gaps in our understanding of how firms learn, innovate, and adapt to regulatory pressure in a dynamic energy industry. Chapters 2 and 3 share an empirical context (hydraulic fracturing for shale gas), and there is a fairly direct line between them (e.g., both address the effects of information-based regulations on firms' decisions). Though it seems quite different in some ways, Chapter 4 also has some shared topical themes—technological change and innovation, particularly in the domain of energy technologies—but addresses, obviously, a very different scale both chronologically and in terms of the number of agents involved. Methodologically, all three chapters have offered learning opportunities for me, allowing me to build an analytical toolbox of quantitative tools useful for insightful characterization of problems that are of interest to both researchers and policy makers.

Within the context of the emerging technology of hydraulic fracturing, both Chapters 2 and 3 address the impacts of disclosure regulations. Although often motivated primarily by the public "right to know" about risks that arise from the storage, use, and disposal of toxic chemicals—or, for that matter, risks such as explosion hazards, possible safety hazards from consumer products, or other risks that may not be immediately apparent—disclosure laws have been found to affect firms' behavior in other aspects of their decision-making. Prior literature in economics and policy has focused primarily on voluntary self-regulation by firms, and has focused primarily on firms that have a direct line to consumers. My setting is unusual in that these oil and gas firms do not sell differentiated products to consumers, and thus generally do not experience pressure from a consumer channel. Furthermore, the nature of the specific information disclosure laws I study does little to facilitate access or comprehension by a non-technical audience. This sets the stage for both of the papers on hydraulic fracturing: the first paper analyzes how firms in non-consumer-facing industries respond in terms of reducing public bads, supplementing our knowledge of how disclosure laws alter firm behavior. The second paper opens an entirely new line of inquiry, regarding the potential for these laws to create pathways for knowledge transmission that were previously inaccessible or overly costly.

1.3 Policy implications

Chapter 2 may have the most direct policy implications. Though it is set in the context of one particular new technology, information-based regulations are common in many domains—and can be especially valuable for emerging technologies in which command-and-control regulation may hinder innovation, market-based regulations are impractical, and the risks of adverse consequences may be uncertain or poorly characterized. Some popular media have called out hydraulic fracturing as a highly risky technology, and have often suggested that the brew of toxic chemicals in fracturing fluid is among the riskiest elements. In my reading of the chain of events, this concern may have been more applicable in the earlier days of hydraulic fracturing, when exploration and production firms themselves were sometimes in the dark regarding the chemical mixtures that were used in their wells. In the context of prior secrecy, the mandatory disclosure regulations that states passed starting in 2010 represented a substantial change for operators, regulators, and the public. This is especially true given widespread public fear about unknown chemicals leaking into surface water and groundwater, and the industry's traditional opposition to public disclosure in general.

I have discussed my findings in Chapter 2 with a number of stakeholders, including regulators, operators, contractors, and third parties seeking to set private standards for responsible operation. The regulators and operators say that my findings corroborate their intuition and experience. One operator provided a quite informative summary of their policy regarding the use of toxic chemicals, documenting a tiered approval system that requires increasingly higher levels of management permission for the use of substances based on their potential to have adverse effects on the environment or human health. In conversations with this operator, it became clear that mandatory disclosure regulations were a critical motivation for this operational policy—in combination with a feeling that it was "the right thing to do," to be sure, but the regulation gave managers both an extra impetus to implement the policy and a tool by which to force their suppliers and contractors to provide substitute chemicals that serve the same purposes but carry less environmental or occupational risk.[1]

Thus, the paper in Chapter 2 has direct implications for regulators in the specific setting of hydraulic fracturing: mandatory information disclosure works to change operators' behavior. Having worked in close proximity to the oil and gas industry for several years prior to my doctoral program, I find this a surprising result.[2] The oil and gas industry is among the world's technically, politically, and legally most sophisticated industries, and the disclosure regulations I study are not particularly conducive to behavioral change. The latter observation is based in part on the characteristics that Fung et al. (2007) suggest are most likely to promote behavior change: policies that, among other things, encourage disclosure of information that is readily accessible and can be integrated into consumers' decision-making. For the regulations I study in Chapter 2, consumers had to download reports individually, so information was difficult to assemble into a database or compare over space or time. In addition, the overwhelming diversity of substances (approximately 2,000 distinct chemicals) would present challenges for interpretation, even for a relatively well-

[1] Thus, in the three-mechanism framework of Bennear and Olmstead (2008), this operator's behavior change seems to be motivated most substantively by the managerial channel.

[2] Indeed, I was initially reluctant to pursue this research topic because I believed I would spend many hours cleaning data only to find no response to the regulation.

educated audience. Furthermore, given their strong political power and longstanding relationships with regulators, oil and gas firms may reasonably have expected little future regulation to arise after providing the information demanded by these policies. In light of all this, the finding that the disclosure regulations caused a large and persistent drop in toxics use is potentially quite valuable for policymakers. It also addresses a research gap pertaining to the conditions under which "disclosure works."

The paper on "Learning by Viewing" carries a different set of implications for policy. In general, our findings suggest that policy makers may face a tradeoff between satisfying the public's right to know and firms' right to secrecy, which may be an important element of the regime by which firms appropriate returns from their investments in innovation. At the same time, it may be that the gains from trade (regarding knowledge of effective chemical formulas) outweigh the losses from failure to innovate. In ongoing developments, my co-authors and I are investigating these questions so as to sharpen policy recommendations.

The paper on energy transitions, as it stands, carries some important information for energy demand forecasters, especially in low-income and developing countries. The paper on which mine builds (van Benthem, 2015) suggests there are important technology rebound effects that may drive today's developing countries to use more energy per unit of economic growth than did developing countries of the past, which bodes ill for forecasters who are assuming some level of energy savings or "leapfrogging" due to the availability of more efficient technologies. My paper offers additional insight and notes that while this technology rebound effect may be of concern for some countries, on average it is not, when considering a broader range of developed countries, a longer timeline for analysis, and a more complete set of energy technologies—especially on the generation and distribution side. That is, the forecasting agencies may be correct after all in assuming some level of energy efficiency available to today's developing countries that did not exist in previous decades. That said, the more valuable policy implications may be those that come from ongoing development of this paper, including exploration of heterogeneous effects and interactions with historical and current institutions.

2

Fracking, Toxics, and Disclosure

2.1 Introduction

Information-based regulations are increasingly common in regulatory policy. In contrast to command-and-control regulation that prescribes or proscribes particular technologies or practices, and market-based regulations that directly modify agents' incentives via price effects, these policies require certain entities to disclose elements of their production process, by-products, or other information that may affect the welfare of stakeholders but would be unobservable to them in the absence of mandatory disclosure. Such policies are often motivated by the notion that the public has an inherent right to know elements of government and private decision making when those elements could conceivably affect public welfare. Perhaps surprisingly, disclosure regulations may also influence regulated entities to change their behavior, even absent other policies.

In explaining the increased use of information-based regulation, scholars have identified several factors. One of these is that technological advances (e.g., cheaper computing power and data storage) makes information easier to collect, store, process, and provide to the public. Another is political in nature: in a contentious political environment, elected officials may face lower political costs if they adopt regulations that merely require disclosure, compared to more intrusive or costly requirements. Furthermore, prescriptive regulation (such as bans on specific substances or processes) can suppress valuable innovation, especially in the context of emerging technologies, or when the magnitude of potential external harms is uncertain. Disclosure regulations, on the other hand, offer an opportunity to "wait and see" while also allowing the public and regulators to gather more information about issues of concern.

These factors have led to the use of information-based regulations in a wide range of industries, from financial regulation to food safety, in the US and many other countries (Dasgupta et al., 2007). They are most commonly used in consumer-facing industries and the financial industry (Fung et al., 2007), consistent with the underlying motivation of protecting the public's inherent right to know certain operational details. Most research on the effects of disclosure laws, too, has focused on consumer-facing industries. In general, empirical evaluation on the effects of disclosure regulations on regulated entities' behavior suggests that even profit-maximizing firms, under certain conditions, change their behavior in ways that seem to enhance public welfare, such as reduced environmental pollution or improved labor practices.

This basic story leaves some questions unanswered, however. Most empirical studies have been set in consumer-facing industries, and firms that provide outputs in commoditized intermediate product markets may not respond to (or experience) external pressure in the same way. Also, most prior empirical studies document the effects of policies that make information relatively accessible to outside parties, and in many cases also policies that result in one big disclosure event, such as an annual inventory of toxic chemical releases representing a large number of firms at once. These program features and others, such as how easily outsiders can interpret any information disclosed, affect how firms respond to disclosure laws (Fung et al., 2007). In this context, analyzing the effectiveness of disclosure regulations in a new context contributes to understanding the conditions in which they can influence firms' behavior, and how.

In this paper I examine the effects of mandatory disclosure regulations on the chemicals that firms use in hydraulic fracturing for producing oil and gas from shale formations. This technology, which is currently conducted primarily in the United States (but is rapidly diffusing to many other countries), combines a suite of technological innovations that allows firms to produce oil and gas from geologic formations that have previously been considered unproductive or uneconomical. Although some of the technologies have been in use for decades, recent innovations since the late 1990s and early 2000s have led to a rapid and widespread expansion of the technology, which in turn has led to the development of oil and gas wells in areas that have not experienced such development in many years or, in some cases, ever. The technology involves high-pressure injection of millions of gallons of fluid down a wellbore, including 50,000 to 100,000 gallons of chemicals, some of which are or could be toxic to humans and ecosystems (Stringfellow et al., 2014). Although major spills are infrequent, regulators and the public have expressed concern about the potential toxicity of the chemicals used.

Since 2010 twenty-eight US states have passed laws requiring oil and gas operators to disclose the chemical components of their fluids. These laws are nearly identical across states, with many using virtually identical text. Prior to these laws, some operators also voluntarily disclosed chemicals used in some or all of their fractured wells, predominantly through a web-based tool created by an industry council. The voluntary disclosures contain essentially the same information as the legally required reports. This includes some firms' choice to declare chemical identities as proprietary trade secrets—thus not revealing the chemical identity. All state laws permit this "proprietary" declaration, and some firms also used the proprietary declaration in their voluntary reporting as well.

I compare the chemical mixtures used prior to the mandatory disclosure laws to those used after the laws took effect, using a difference-in-differences method motivated by the differences in state-level regulatory timing. Many operators work in multiple geologic plays and multiple states; based on conversations with operators, I assume that firms implement any toxics-reduction policies at the level of geologic play, implying some leakage outside of the state boundary (i.e., for geologic plays that cross state borders). I control for this leakage by incorporating fixed effects for the interaction of time and geologic play.

I use the voluntarily reported data to measure chemical use prior to the regulation. However, the voluntary reports do not represent all wells fractured prior to mandatory disclosure. Thus, with the passage of mandatory disclosure regulations, I observe a composite effect of full reporting (i.e., because all firms are required to disclose chemicals for all wells) and disclosure pressure (the effect, if any, of the mandatory disclosure regulation on firms' choice of chemicals). The full reporting effect is likely positive, if voluntary reports are cleaner on average, whereas the disclosure pressure effect would be negative, if firms reduce reducing their use of toxics in response to regulations.

I use two methods to distinguish the full reporting effect from the disclosure pressure effect. First, I run a separate analysis on that is limited to wells operated by "frequent voluntary reporter" firms—that is, those firms that voluntarily reported a large proportion of their wells before disclosure was mandatory. By construction the full reporting effect is lower for these firms, so the composite effect is more strongly weighted toward the disclosure pressure effect. Second, in one state, I exploit an unusual period in which all firms were required to report full information to the state regulator—but not in a way that was readily available to the public at the time—

even as some firms voluntarily reported information to a public website. The average voluntary public report in that period shows concentrations of toxic and regulated chemicals that are significantly lower than the average of the regulator-only reports, suggesting the full reporting effect is positive and large.

The analysis suggests that firms reduced their use of toxic chemicals in response to the mandatory disclosure regulations. Though the effect does not occur immediately after the regulation, it does persist over time, for at least three years after the regulations come into effect. My analysis also suggests that firms' use of proprietary chemicals increased over the same time period, although this trend does not appear to be causally related to the mandatory disclosure regulations.

This paper is among the first empirical analyses of the effects of mandatory disclosure regulations in a non-consumer-facing industry that uses methods that can accurately assess causal effects of those regulations. As such, it provides useful insight for both scholars and practitioners regarding how firms respond to mandatory disclosure regulations when consumers have little or no direct influence on firms' profit-maximizing abilities. The analysis also demonstrates a method, perhaps useful in other contexts, to discern the effects of mandatory disclosure regulations when data are available from a pre-existing voluntary reporting scheme. By showing that disclosure regulations can influence firms' behavior even when firms sell into intermediate markets, the timing of disclosures is diffuse, and reports are difficult for outsiders to access or interpret, this paper helps to elucidate the mechanisms by which information-based regulations operate. Furthermore, these findings may help policymakers to choose among alternative regulatory instruments or specific policy design elements.

The remainder of the paper proceeds as follows: Section 2.2 provides background on the empirical setting and discusses related literature. In Section 2.3 I document the data I use, and in Section 2.4 I describe the empirical method. In Section 2.5 I summarize results and document robustness checks. Section 2.6 concludes.

2.2 Background

2.2.1 Hydraulic fracturing

The shale gas and oil boom in the US has dramatically altered global energy markets and has brought jobs, royalties, and tax revenues to nearby communities (Hausman and Kellogg, 2015). At the same time, environmental groups have raised concerns about aspects of the production process. Among these concerns is the contents of the fracturing fluid, which usually consists of water, sand or another granular material, and a cocktail of chemical additives. The process of hydraulic fracturing involves the injection of this fluid into rock formations hundreds to thousands of meters below the surface of the earth. The chemicals serve a number of purposes that enhance the productivity of water and sand. These include reducing the viscosity of water to allow faster pumping and induce higher pressures, enhancing natural fractures in the substrate, carrying the proppant farther into these fractures, helping to build up a solid barrier on the formation face, and minimizing the growth of bacteria that might interfere with metal casing or cause other problems. An additive that enhances performance of one objective may degrade performance of another: e.g., a more viscous fluid can carry proppant farther into fractures, but may require additional additives to break the viscosity later in the fracturing operation. Engineers have invested considerable research into optimal design characteristics of a fracture operation, with environmental toxicity of the fluid as one consideration (Montgomery, 2013; Gulbis and Hodge, 2000). Industry practitioners indicate that the choice of specific individual components is rarely if ever driven by cost, because the cost of the chemicals themselves is small in comparison to the overall cost of the fracturing and stimulation operation.[1]

[1] Personal communication with Mark Boling, Southwestern Energy, 2013.

Many of the chemicals used can be toxic to human or ecological health. One of the earliest concerns to arise about fracturing technology was the identity and toxicity of the chemicals in fracturing fluid and the possibility they might migrate or be accidentally released into ground water or surface water. While these chemicals represent a small proportion of fracturing fluid (usually on the order of two percent of total volume), this fraction represents 60,000 to 100,000 gallons for a typical operation that uses three to five million gallons of fluid. Public and regulatory concerns about risks associated with fracturing seem to have arisen due to the proximity of wells to non-industrial land uses, the large number of wells developed in a short period of time, and a few high-profile incidents of water pollution that some media reports attributed to spills or methane intrusion. Media coverage of fracturing chemicals highlighted both the toxicity of some chemicals and the industry's desire to keep chemicals secret (e.g. Elgin et al., 2012; Haas et al., 2012).

2.2.2 Regulatory setting

The 2005 Energy Policy Act exempted hydraulic fracturing from the Underground Injection Control (UIC) provision of the Safe Drinking Water Act (SDWA), which would otherwise have regulated operators' choice of chemicals injected underground in the chemical slurry.[2] As a result, except for fractures that used diesel fuel, operators did not need to disclose to regulators or the public the compounds they used in high-pressure slurries. Companies were initially reluctant to disclose the contents of chemical fluids, citing proprietary concerns and also claiming that the compounds involved are not harmful when properly handled. However, a Congressional investigation initiated in 2010 and published in 2011 identified many instances in which companies did not know the chemical makeup of compounds they were using (Wax-

[2] Some observers referred to this provision as the "Halliburton Loophole," although others argue it was merely a clarification of whether SDWA applied.

man et al., 2011). In the wake of this report and other developments, and perhaps due to industry concern that the practice could face more severe regulation or a moratorium, individual states began to pass legislation that required companies to disclose the chemical additives used in their formulas.

No federal or state policies regulate chemicals used in fracturing, other than the information disclosure laws analyzed here and generic regulations such as the Emergency Planning and Community Right-to-Know Act (EPCRA) and the Occupational Safety and Health Act (OSHA). There were no changes in EPCRA or OSHA that would affect fracturing fluid contents over the period of this study. As noted above, the only provision of the UIC laws that applies to fracturing fluids is that operators wishing to use diesel fuels as part of the fracturing fluid must obtain a permit in advance. The definition of "diesel fuel" for this purpose remained largely constant over the period I study: EPA issued draft guidance for UIC permit writers on this definition in May 2012 to clarify the definition, but noted that permit writers should continue to use existing regulations until final guidance was issued, which was in February 2014 (USEPA, 2012a, 2014).

Thus, during the time period of this study, other than the definition of diesel fuel there was no change in federal regulations that affected firms' choices of fracturing chemicals (though there were some proposed changes, as documented in Section 2.4.2). The only significant regulatory changes at the state level was the passage of the laws I study. The policies are quite similar across states.

Table 2.1 provides basic information about state laws requiring disclosure of chemical additives in hydraulic fracturing fluid. Of the 18 states with significant fracturing activity and disclosure laws, six require operators to report information to the FracFocus registry, a web-based database created by the Groundwater Protection Council and the Interstate Oil and Gas Conservation Commission (GWPC and IOGCC, 2015). Five states, including several of the states that passed the earliest disclosure rules, require operators to report information to a state regulatory agency or commission. The seven remaining states allow operators to choose their reporting location (i.e., to FracFocus or the state), although one (Oklahoma) notes that the state regulator will upload to FracFocus any information it receives. Despite these differences in reporting registry, state laws generally require similar information about chemical additives, especially among states that require or allow reporting to FracFocus. The required information generally includes ingredient name, chemical abstract service (CAS) number, concentration in fracturing fluid (typically the maximum concentration used in any fracturing stage), supplier name, and trade name if applicable. When operators upload information to FracFocus they are also asked to provide data on well location and characteristics including vertical depth, water volume used, latitude and longitude, and well name.

All states allow exemptions for the disclosure of additives considered to be confidential business information that firms believe gives them competitive advantage. Approaches to accommodating trade secrets are broadly similar across states, in part because of the Uniform Trade Secrets Act (promulgated by the Uniform Law Commission in 1979 and passed by 46 states), which aims to harmonize standards for trade secret protection. Operators must declare an exemption for individual chemical ingredients for which they claim trade secret status, and reports uploaded to FracFocus typically include the concentration of the chemical used but not its name or chemical identification number. Some states also require operators to report the chemical family to which the proprietary substance belongs. Thus, reports typically show the total quantity of proprietary chemicals used, but little additional information about the nature of these chemicals.

2.2.3 Related literature

Bennear and Olmstead (2008) identify three mechanisms through which disclosure

Table 2.1: State disclosure laws

State	Effective date	Basis for effective date	Reporting location
Wyoming	5-Sep-10	Frac job	State agency
Arkansas	15-Jan-11	Drilling permit	State agency
Michigan	22-Jun-11	Frac job	State agency
Montana	27-Aug-11	Frac job	FracFocus or state agency
West Virginia	29-Aug-11	Frac job	State agency
Louisiana	20-Oct-11	Drilling permit	FracFocus or state agency
Texas	1-Feb-12	Drilling permit	FracFocus
New Mexico	15-Feb-12	Frac job	State agency
Colorado	1-Apr-12	Frac job	FracFocus
North Dakota	1-Apr-12	Frac job	FracFocus
Pennsylvania	16-Apr-12	Frac job	State (2/5/11); FracFocus (4/16/12)
Ohio	11-Jun-12	Frac job	FracFocus or state agency
Utah	1-Nov-12	Frac job	FracFocus
Oklahoma	1-Jan-13	Frac job	FracFocus or state agency
Mississippi	4-Mar-13	Frac job	FracFocus or state agency
Alabama	9-Sep-13	Frac job	FracFocus
Kansas	2-Dec-13	Frac job	FracFocus or state agency
California	1-Jan-14	Frac job	FracFocus or state (1/1/14); State (1/1/16)

Excludes some states with little or no fracturing activity.
1. Pennsylvania required reporting to the state in February 2011, but changed the reporting location to FracFocus in April 2012.
2. Oklahoma's regulations note the state regulator will report to FracFocus any information it receives.
3. In California, reporting to the state (rather than FracFocus) became mandatory on January 1, 2016.

might result in abatement. The first is through the market: For instance, if consumers have information about firms' environmental performance and prefer greener goods, then they can exert market pressure in hopes of inducing firms to improve. Market mechanisms could also operate through other channels, such as by higher financial or legal liabilities for firms engaged in dirtier production practices. The second mechanism is political: information may increase the ability of a concerned public to lobby for stronger regulation. Finally, disclosure may affect an organization's internal decision making, as individuals within the firm change their behavior as a result of measuring and reporting data.

As noted, some empirical analyses have been conducted on the effectiveness of disclosure policies using methods that can accurately identify causal effects. Most of this evidence is from industries where the information disclosed is of relatively high visibility to consumers, including electricity (Delmas et al., 2010; Kim and Lyon, 2011), drinking water (Bennear and Olmstead, 2008), and restaurant hygiene (Jin and Leslie, 2003). Analyses within industries that are less consumer-facing, such as manufacturing facilities that report to the Toxics Release Inventory (TRI), find that releases of reportable toxic chemicals have declined over the course of the program by as much as 50 percent (Bennear and Coglianese, 2005). However, because no data were available on toxic releases prior to the start of the program, this decrease cannot be definitively attributed to the public disclosure requirement (Bennear and Olmstead, 2008).

There has been no comprehensive empirical investigation of which mechanisms are most powerful or effective. Bennear and Olmstead (2008) note that in their setting, the market mechanism is likely not relevant since water suppliers are essentially monopolists, and the internal mechanism is unlikely to play a large role because of pre-existing monitoring requirements. Thus, they conclude the political mechanism drives the results in their study. Doshi et al. (2013) investigate the internal mechanism in depth by identifying what characteristics of firms influence or moderate their response to disclosure regulations, but they do not rule out the possibility that consumer or political pressure are also acting simultaneously to influence firms' decisions.

Indeed, cleanly isolating the operation of one particular mechanism is likely not possible without a case study of a particular (small) set of firms. For analysis on a larger group, the relative importance or effectiveness of different mechanisms is best teased out by analyzing the effects of disclosure regulations that operate in different settings, and noting which mechanisms are most likely to be relevant within those settings. This will help public decision makers to identify the potential for information-based regulations to have desirable results, and conceivably to design effective public policy. This should be especially important as disclosure regulations become more prevalent in non-consumer-facing industries, where the market mechanism is less influential. It may also be more important in contentious or fractured political environments in which political channels are less effective because the threat of regulation is less credible.

2.3 Data

I create a novel data set with information about well completions, including chemical constituents of fracturing fluids, for 73,211 wells in seven states; all wells were hydraulically fractured between 2011 and 2015. I combine three types of data: chemicals used in each fracture; state-level censuses of well completions; and information about individual chemical characteristics, including toxicity and other aspects of interest. The following sections describe each of these three data types.

2.3.1 Chemical additives

I assemble data on chemical additives from the FracFocus database (GWPC and IOGCC, 2015) and one state regulatory agency (California) that has comparable and accessible data. FracFocus is the legally required reporting registry for operators to report chemicals in five states in my analysis (Colorado, North Dakota, Pennsylvania, Texas, and Utah).[3] I also include FracFocus reports from Oklahoma, because while operators can choose whether to report to FracFocus or the state, that state's disclo-

[3] FracFocus is also the reporting registry for wells fractured in Alabama, but there were very few wells in that state during the analysis period (126) and even fewer wells that meet the criteria for usable observations. Thus, I exclude Alabama from the analysis.

sure law indicates that the regulatory agency will upload to FracFocus all information it receives. I also collect information on chemical additives from the California Division of Oil, Gas and Geothermal Resources (DOGGR). California used FracFocus as an alternative reporting site prior to January 1, 2016, and allowed operators to report to either its registry or FracFocus. Thus, operators would report chemicals for any well fractured in California to either FracFocus or the California state registry, and both registries provide comparable data—the location of the wellhead, depth, water volume, operator, and fracture date. For each chemical, in FracFocus (for all states) as well as the California registry, I observe the chemical name, identification number, supplier, purpose, and maximum additive concentration in the fracturing fluid.

I attempted to compile chemical reports from other state regulators. However, the data available and the accessibility of those data vary widely. Some regulators allow inspection of chemical disclosure reports only in person at state offices. Although some disclosure reports are posted online for some states, the reporting format is highly variable, sometimes even within the same state, and chemical information is frequently provided along with other aspects of well completion in documents that can number close to a hundred pages per well. In addition, the data items available on these disclosure reports generally differ from the data available on FracFocus. Thus, this analysis focuses on states that use FracFocus or have comparable state registries (i.e., California). This increases the internal validity of the results reported here, but at the cost of some external validity: the effects I observe may not hold in states where chemical disclosures are less accessible.[4]

[4] Indeed, with respect to the three mechanisms by which disclosure may alter firms behavior—market, political, or internal managerial channels—testing whether firms reduce toxics use in states where reports are less accessible to the public would help to discern the relative importance of the political channel in relation to the internal managerial channel. For instance, in these states the internal managerial channel might be the primary mechanism by which disclosure would have an effect, since limiting public accessibility may limit the influence of the political channel. Future work could address this issue by considering effects of disclosure within states that require reporting

When FracFocus was launched, GWPC and IOGCC provided fracturing fluid chemical reports only as individual PDF files that were easy to download individually but challenging to compile en masse. At the time, GWPC and IOGCC specifically stated their intent to provide a forum where the public could view individual reports but not look at many reports at once. At least two entities successfully scraped the entire FracFocus database in late 2012 or early 2013; one of these, the environmental NGO Skytruth, provided the resulting data online for public download (Skytruth, 2013). In summer 2015, FracFocus provided a dataset for public download; however, that data release excluded many of the earliest reports (wells fractured prior to April 2013). My dataset combines information from both sources, including all wells available on FracFocus with completion dates between February 2011 and June 2015.

The resulting dataset includes 80,190 unique observations of wells with chemical information disclosed either voluntarily or in compliance with state regulations. I drop 3,244 of these wells for one of three reasons: (i) there is not sufficient information to classify them as voluntary or mandatory reports, (ii) there is no information on basic well parameters that I use as covariates (water volume, depth, oil/gas classification), or (iii) the calculated sum of fracturing fluid concentrations exceeds 200 percent. Nearly 40 percent of the observations have a calculated sum over 100 percent, partly due to rounding error and partly because disclosure laws require operators to report maximum concentrations across all fracturing stages rather than actual average concentrations. However, the calculated sum drops off quickly after 101 percent; for instance, just 4.5 percent of wells have a calculated sum over 103%. Visual inspection of FracFocus reports for wells with a calculated sum over 200% suggest most arise from either operator error in completing forms, or errors generated by the scraping procedure (e.g., errors that resulted in transposing two data columns). I ran the analysis using alternative thresholds for this cutoff ranging from

to locations other than FracFocus.

101% to 200%, and the results are essentially identical to those reported here.

Finally, I remove 3,735 observations of wells fractured before or after the time cutoff values shown in equation (2.3), the time-varying difference-in-differences econometric specification (i.e., more than 6 quarters prior to the regulation, or more than 12 quarters after). These observations would drop out of the estimation of equation (2.3) in any case, and excluding them from the estimation of (2.1) and (2.2) makes results from these specifications more comparable. I have also run the analyses including these observations, and the results are essentially identical to those reported in this paper.

2.3.2 Chemical characteristics

Companies use nearly 2,000 distinct additives in fracturing fluid, with the fracturing fluid for each well containing usually 10 to 40 different compounds. I calculate four well-level measures of fluid toxicity, including two that directly address toxicity or regulated status.

The first measure, which I label "priority toxic and regulated chemicals" (PTRCs), is the total concentration of chemicals that fall into any of four groups. Three of these groups correspond to regulatory classification: (i) regulated as primary contaminants by the Safe Drinking Water Act; (ii) regulated as Priority Toxic Pollutants for ecological toxicity under the Clean Water Act; or (iii) classified as diesel fuel under EPA guidance on fracturing operations (USEPA, 2012a, 2014). The fourth group includes chemicals listed in the USEPA Risk-Screening Environmental Indicators (RSEI) database (USEPA, 2012b) as having a relatively high risk value for chronic human health effects. The RSEI database provides peer-reviewed relative risk values for nearly 500 chemicals covered by the Toxics Release Inventory, with chemical-specific measures based on detailed toxicity evaluations from both oral and inhalation exposure pathways. Among chemicals used in fracturing fluid, RSEI scores range from 0.5 (oral exposure to formic acid) to 3,000,000 (oral exposure to hydrazine and quinolone); these are unitless scores that are intended to provide order-of-magnitude comparisons. My measure of PTRCs includes chemicals that have an RSEI score of at least 200.

The second measure is the weighted sum of concentrations of individual chemicals in fracturing fluid, with weights corresponding to chronic human health scores from the RSEI database. Unlike the first measure, which is simply a sum of concentrations, this provides a toxicity-adjusted measurement of chemicals of potential concern within the fluid. Still, it is limited to chemicals that have been analyzed for toxicity with sufficient rigor and peer-review that they can be included in the RSEI database. Some chemicals of potential concern that are used in fracturing fluid may not be included in RSEI or covered by water quality regulations, an issue that Stringfellow et al. (2014) discuss in some detail.

The third measure addresses the use of chemicals that are frequently cited in media reports in association with potential dangers of fracturing, regardless of regulatory status or scientific evidence of toxicity. I use this measure to test for the possibility that companies respond more to public concern or media attention than to toxicity or regulatory status. I identify "high media profile" chemicals by searching for media reports that mention particular chemical names in conjunction with fracturing and words indicating danger or hazard, and designate chemicals as high-profile if they were mentioned in at least 100 unique media stories over a five-and-half-year period.[5] Like the first metric, this measure is the sum of reported concentrations of these "high media profile" chemicals.

Finally, I calculate a fourth measure that is the sum of concentrations of chem-

[5] These stories typically describe the rapid spread of fracturing, the proximity of non-industrial land uses, offer a few specific environmental concerns, and list one or more chemicals in use in fracturing fluid. The substance mentioned most often is diesel (1,816 stories), followed by benzene (1,230) and toluene (491 stories).

icals designated as proprietary or confidential business information. This measure allows me to test whether disclosure laws resulted in an increase in companies' use or declaration of proprietary chemicals. This could arise, for instance, if operators were concerned that disclosure laws would imperil confidential business information in a way that voluntary reporting did not. Another possibility is that some operators could continue using regulated, toxic, or publicly controversial chemicals, but "hide" those substances by declaring them as part of proprietary additives. However, caution is warranted in interpreting results for proprietary declarations. While all state disclosure regulations contain specific provisions for the reporting of proprietary chemicals, in a voluntary reporting regime some operators may have chosen not to list proprietary chemicals at all in a voluntarily submitted disclosure form. That is, we might expect to observe an increase in chemicals declared proprietary after disclosure laws, simply because the disclosure laws formalized the necessity to report proprietary chemicals as part of a fracturing fluid. Indeed, my conversations with operators suggest there is heterogeneity in how operators approached the question of proprietary chemical disclosures when reporting was voluntary.

Overall, I consider 790 distinct chemicals as "of potential concern," and I observe 64 of these in fracturing fluid. This includes 46 PTRCs, 13 high-media-profile chemicals, and 58 chemicals with positive RSEI scores.

2.3.3 State permitting databases

I supplement FracFocus data with information from state permitting databases (drawn primarily from DrillingInfo) for two reasons. First, in Texas, the disclosure law applies based on the issue date of the initial drilling permit, rather than the date of fracturing. The FracFocus data do not include initial drilling permit date, so I obtain these dates from the state regulatory database.

Second, I use state permitting information to develop a census of wells that are "completed," representing the universe of wells that are either ready to fracture, in process of fracturing, or where fracturing has been performed, such that the well is ready to produce gas and/or oil. Developing this universe of wells allows me to estimate the amount of voluntary reporting that occurred prior to the mandatory disclosure period. I calculate an operator-level measure of voluntary reporting and use this measure to distinguish "voluntary reporter" firms—those that reported chemical use for a high percentage of their wells before disclosure became mandatory. This in turn helps me tease out the effect of interest (i.e., the effect of disclosure regulations) from the "full reporting" effect. The latter effect is minimized for voluntary reporter firms, since these operators already reported chemical use for a large proportion of their wells in the period before disclosure was mandatory.

I classify operators as voluntary reporters if they provided chemical reports for at least 75% of their horizontal wells in the 12 months prior to the regulation effective date in the relevant state. For instance, Range Resources completed 171 horizontal wells in Pennsylvania, Texas and Oklahoma in the 12 months leading up to the respective disclosure laws in those states, and filed chemical disclosure reports for 153 (89%).[6] Thus, Range Resources would be considered a voluntary reporter (at the 75% threshold) for the purposes of this analysis. I have also run the analysis at other thresholds lower and higher than 75%, and the qualitative results are robust.

2.3.4 Descriptive statistics

Table 2.2 provides a summary of locations and other information for the wells used in this analysis. For about one-quarter of the wells chemical information is voluntarily reported, while information is reported by law for the remaining three-quarters. The

[6] The focus on horizontal wells means that operators who drilled no horizontal wells in the twelve months prior to the regulation cannot be classified as to voluntary reporter status. I exclude these operators from the set of voluntary reporters. Because these are generally small firms, often operating in only one state, and focused on vertical wells, this should not have a significant impact on the results.

proportion of wells that are voluntarily reported is relatively stable across states, although somewhat lower in North Dakota (14%) and higher in California (54%) than other states, which range from 23% to 39%. Over half the wells are in Texas; Colorado contributes the next largest share of wells for an individual state, at 12%. Over half the operators work at least partially in Texas.

Texas also passed the earliest mandatory disclosure law requiring reporting to FracFocus. Unlike other states, however, Texas used the date of issue for the drilling permit, rather than the date of the fracturing job, to determine legal reporting requirements. Thus, many of the 9,341 voluntarily reported wells (nearly half of those in my dataset) were fractured after February 1, 2012, when the Texas law came into effect. Because the operators of these wells obtained drilling permits prior to February 1, any chemical fluid reports I observe were voluntarily submitted.

Table 2.3 shows similar information but focuses on the subset of operators that I classify as "voluntary reporters at the 75% level" (VR75), based on the criterion of voluntarily reporting chemicals for at least 75% of their horizontal wells in the twelve months leading up to mandatory regulation in the relevant state. About half the wells in my sample are fractured by these "VR75" operators, though the percentage of these wells varies by state and ranges from just 1% in California to 77% in Pennsylvania. Relatively few operators, just 8%, fall into the VR75 category; again, this percentage varies by state, from 7% in Oklahoma and 9% in Texas to 49% in Pennsylvania. Comparing the percentage of VR75 operators to wells fractured by these operators also suggests that these operators are also more likely to include the largest operators (because they represent just 8% of operators but fracture 47% of wells). This is also true in each state except California, where it is apparent that the largest operators do not fall into the VR75 category.

Table 2.2: States, wells, and operators (all operators)

	Effective date of disclosure regulation	Wells (voluntarily reported)	Wells (mandatory)	Wells (total)	Percent wells voluntarily reported	Percent wells in state	Number of operators
Texas	2/1/2012	9,341	31,228	40,569	23%	55%	518
Colorado	4/1/2012	3,409	5,393	8,802	39%	12%	95
North Dakota	4/1/2012	949	6,062	7,011	14%	10%	53
Pennsylvania	4/16/2012	1,583	3,462	5,045	31%	7%	45
Utah	11/1/2012	1,149	2,112	3,261	35%	4%	32
Oklahoma	1/1/2013	1,629	4,884	6,513	25%	9%	198
California	1/1/2014	1,082	928	2,010	54%	3%	14
Total		19,142	54,069	73,211	26%	100%	767

In Texas, chemical disclosure was not required for wells that received initial drilling permits prior to the effective date of the regulation, even if the fracturing job happened after that date. Thus, although the Texas regulation came into effect before that of any other state in the analysis, in some cases Texas wells not subject to the regulation may have been fractured after wells in other states that were subject to the regulations in those states.
Total number of operators is less than the sum of the column because some operators work in more than one state.

Table 2.3: States, wells, and operators (voluntary reporter [VR75] operators)

State	Wells (voluntarily reported)	Wells (mandatory)	Wells (total)	Percent wells fractured by VR75 operators	Number VR75 operators	Percent operators among VR75
Texas	5,461	12,840	18,301	45%	46	9%
Colorado	3,005	3,683	6,688	76%	14	15%
North Dakota	498	1,635	2,133	30%	8	15%
Pennsylvania	1,376	2,513	3,889	77%	22	49%
Utah	676	819	1,495	46%	6	19%
Oklahoma	771	1,089	1,860	29%	14	7%
California	12	4	16	1%	3	21%
Total	11,799	22,583	34,382	47%	58	8%

Total number of operators is less than the sum of the column because some operators work in more than one state.

The relationship between operator size and voluntary reporting is worth investigating because my technique to distinguish the full reporting effect from the disclosure pressure effect relies on running the difference-in-differences analysis separately for the VR75 operators. If these are also the largest operators, I may not be able to distinguish a finding of a significant effect of disclosure regulations for VR75 operators from the effect on large operators. Table 2.4 provides some insight into the correlation of size and VR75 status, showing that of the 25 largest operators in the dataset (defined by number of wells), only half (12) fall into the VR75 category, and of the top 10 only 6 (albeit, including the top 5) are classified as VR75. Thus, some but not all of the largest operators are classified as VR75. Table 2.4 also shows which of the largest operators meet another VR-related classification, VR90, identifying firms that met a 90% (rather than 75%) threshold for voluntary reporting.[7]

Table 2.5 provides descriptive statistics for the explanatory and dependent variables I use in my analysis. The average well in my sample is about 8,500 feet deep and uses slightly over 3 million gallons of water. About 61% of the wells in the sample are oil wells (or combined oil and gas wells), and 39% are gas wells only. Of the 71,989 wells for which the well direction is known, 69% are horizontal or directional wells while 31% are vertical wells. The well direction is not reported for 1,222 wells.

The dependent variables, concentrations of various chemicals of interest, exhibit substantial right skew. To reduce the skewness I take logs (adding 0.01 so the log of zero values is defined), but even then there are some outliers which may have an undue influence on OLS coefficients. I therefore take the additional step of winsorizing these outliers, defined as observations in excess of 1.5 times the interquartile range lower than the 25th percentile or higher than the 75th percentile. Table 2.5 shows

[7] Since only one of the ten largest operators is a VR90 firm, analyzing the subset of VR90 firms allows me to better distinguish whether the toxicity-reducing effect of disclosure laws applies only to the largest operators. When I analyze the wells of VR90 operators—using an approach identical to that for VR75 operators—I find results qualitatively identical to those reported in section 2.5.

Table 2.4: Voluntary reporter status of 25 largest operators

Operator	Wells	Operator % of wells	VR75	VR90
Anadarko	5,027	6.9%	yes	yes
Chesapeake Operating, Inc.	3,952	5.4%	yes	
EOG Resources, Inc.	3,405	4.7%	yes	
Pioneer Natural Resources	2,837	3.9%	yes	
Apache Corporation	2,416	3.3%	yes	
XTO Energy (ExxonMobil)	2,318	3.2%		
Devon Energy Production Co., LP	2,044	2.8%		
Sandridge Energy	2,029	2.8%		
Occidental Oil And Gas	1,959	2.7%		
Noble Energy Inc.	1,633	2.2%	yes	
Encana Oil & Gas (USA) Inc.	1,527	2.1%	yes	yes
Aera Energy LLC	1,472	2.0%		
WPX Energy	1,467	2.0%	yes	yes
ConocoPhillips	1,417	1.9%	yes	yes
Marathon Oil	1,348	1.8%	yes	yes
Newfield Exploration	1,304	1.8%		
Chevron USA Inc.	1,139	1.6%		
Energen Resources Corporation	1,123	1.5%		
Continental Resources, Inc.	1,074	1.5%		
Whiting Oil And Gas Corporation	1,051	1.4%		
BHP Billiton Petroleum	918	1.3%		
Hess Corporation	855	1.2%	yes	yes
Range Resources Corporation	750	1.0%	yes	
Laredo Petroleum, Inc.	710	1.0%		
COG Operating LLC	673	0.9%		
Total (top 25)	44,448	60.7%	yes (12)	yes (6)

summary statistics for these winsorized variables as logs; to facilitate interpretation, the table also shows summary statistics for the corresponding levels, though I focus on the logs for the analysis.

Table 2.6 shows the mean and standard deviation for wells reported under each regime type (voluntary or mandatory), and differences between the means for each regime type. Voluntarily reported wells tend to be less deep, use less water, are

Table 2.5: Descriptive statistics

Variable	N	Mean	SD	Minimum	Maximum
Well depth (ft)	73,211	8,507	2,925	100	25,000
Fluid volume (10^6 gal)	73,211	3,130	3,108	1	15,000
Oil well	73,211	0.61	0.49	0	1
Vertical wellbore	71,989	0.31	0.46	0	1
Log relative toxicity score	73,211	-0.4	2.89	-4.61	8.25
Log ppm PTRCs	73,211	-0.65	3.59	-4.61	12.43
Log ppm high-media-profile chemicals	73,211	2.7	3.74	-4.61	11.36
Log ppm proprietary chemicals	73,211	3.42	4.88	-4.61	14.24
Relative toxicity score	73,211	52	387	0	4,667
PTRC chemicals (ppm)	73,211	484	9,305	0	211,263
High-media-profile chemicals (ppm)	73,211	558	5,435	0	98,863
Proprietary chemicals (ppm)	73,211	3,028	43,412	0	4,750,000

Depth is winsorized at a lower bound of 100 feet and upper bound of 25,000 feet. Water volume is winsorized at a lower bound of 1,000 gallons and upper bound of 15,000,000 gallons.
Oil wells include wells that produce oil and gas together.
Log relative toxicity score, log ppm PTRCs, and log ppm high-media-profile chemicals are winsorized at the 75th percentile plus 1.5 times the IQR (for 490, 490, and 190 values, respectively). To accommodate zero values of the underlying levels, 0.01 (ppm) is added to the underlying level for all values. Varying the magnitude of this adjustment does not qualitatively change results.

more likely to be gas wells, and are slightly more likely to be vertical wells. They also tend to have lower log concentrations of PTRCs, high-media-profile chemicals, proprietary chemicals, and lower log relative toxicity scores. All of these differences are statistically significant, but not all are necessarily meaningful; the difference in mean depth, for instance, represents about a 2% shallower well, which is not likely meaningful.

Table 2.6: Comparison of means under voluntary and mandatory reporting regime

	All operators (N = 73,211 wells)			VR75 operators (N = 41,558 wells)		
	Voluntary	Mandatory	Difference	Voluntary	Mandatory	Difference
	Mean	Mean	Mean	Mean	Mean	Mean
Variable	(SE)	(SE)	(SE)	(SE)	(SE)	(SE)
Well depth	0.837	0.856	0.0188***	0.870	0.888	−0.0182***
(10^4 ft)	(0.002)	(0.0013)	(0.0024)	(0.0021)	(0.0016)	(0.0026)
Water volume	2.407	3.386	−0.978***	2.663	4.265	−1.601***
(10^6 gal)	(0.0179)	(0.0140)	(0.0228)	(0.0221)	(0.0219)	(0.0311)
Oil well	0.511	0.641	−0.130***	0.403	0.542	−0.140***
	(-0.004)	(0.002)	(0.004)	(0.005)	(0.003)	(0.006)
Vertical wellbore	0.341	0.306	0.035***	0.308	0.178	0.130***
	(0.003)	(0.002)	(0.004)	(0.004)	(0.003)	(0.005)
Log relative	-0.560	-0.35	−0.21***	-0.7	-0.67	-0.03
toxicity score	-0.02	-0.01	-0.02	-0.03	-0.02	-0.03
Log ppm PTRCs	-0.77	-0.61	−0.17***	-0.87	-1.01	0.14***
	-0.03	-0.02	-0.03	-0.03	-0.02	-0.04
Log ppm high-media-	2.63	2.72	−0.09***	2.74	2.5	0.24***
profile chemicals	-0.03	-0.02	-0.03	-0.03	-0.02	-0.04
Log ppm proprietary	2.6	3.72	−1.12***	2.31	4.2	−1.89***
chemicals	(0.04)	(0.02)	(0.04)	(0.05)	(0.03)	(0.06)

All significance tests allow for unequal variances by group.
* p<0.10, ** p<0.05, *** p<0.01.
Well direction (e.g., vertical wellbore) is known for only 71,989 wells (all operators) and 41,062 wells (VR75 operators).

However, to the extent that public and regulatory concern over fracturing inputs has focused on water and chemical use, the lower reported use of water and chemicals for voluntarily reported wells is consistent with the notion that companies tend to voluntarily disclose operations that may be seen as less controversial. Other explanations exist: other researchers have found companies using greater amounts of water per well as fracturing technology advances (Covert, 2015); thus, the lower water use may be simply due to the fact that voluntary reports tend to occur earlier in time. The greater observed use of chemicals of interest may be attributable to the full-reporting effect discussed previously.

Among the wells that are fractured by VR75 operators, voluntarily reported wells again tend to be less deep, use less water, are more likely gas wells, and are more likely vertical wells. They also tend to have a weakly lower log relative toxicity score and use lower concentrations of proprietary chemicals. However, these voluntarily reported wells have slightly higher concentrations of PTRCs and high-media-profile chemicals. This suggests that restricting the sample to VR75 operators indeed provides a method to distinguish the full-reporting effect from the disclosure-pressure effect.

2.4 Econometric approach

Ideally I would randomly assign mandatory disclosure to some wells and observe the use of toxic chemicals for all wells before and after assignment. Since I cannot manipulate policy, I motivate the empirical model as a natural experiment where states require disclosure at different times. I use a difference in differences specification that exploits variations in the effective dates of state regulations. In this section I document my econometric approach and address empirical models and identifying assumptions.

2.4.1 Empirical models

The basic model for the difference-in-differences approach is

$$Y_{ipjst} = \alpha_s + \lambda_t + \delta D_{st} + \gamma W_i + \theta_j + \psi_p + \sum_{p=1}^{P} \psi_p \times \lambda_t + \epsilon_{ijpst} \tag{2.1}$$

where Y_{ipjst} is the variable of interest (e.g., percent of toxic additives). Subscripts i, j, p, s, and t denote, respectively, well, operator, geologic play, state, and time. Thus α_s, λ_t, θ_j, and ψ_p represent fixed effects for state, time, operator, and geologic play, while W_{it} is a vector of well characteristics. The interaction of geologic play and time period controls for time-varying factors within each geologic play, including whether operators face a mandatory reporting regulation within each play. That is, this interaction term controls for the assumption that operators implement toxics-reduction policies at the level of geologic play (not just within a state jurisdiction), which is consistent with the organizational structures of operators who work over broad geographic areas. The variable for treatment, D_{st}, takes the value one if disclosure was mandatory in state s in period t. Thus, δ represents the average difference-in-differences in toxic chemical use between the treatment and control groups, and if $\delta < 0$ then the use of toxics (or another measure of interest) declined as a result of the mandatory disclosure law.

This specification assumes a limited set of differences between the periods "before mandatory disclosure" and "after mandatory disclosure." These differences include the change in disclosure regulations, of course, as well as different operator decisions over the well parameters in W_{it}, the distribution of wells across geologic plays, and the involvement of individual operators. Still, operators could have experienced changes in managerial structure, or revised company policies in ways that altered the approach to fluid formulation, during the study period, for reasons independent of the change in disclosure laws. There could be also unobserved changes in state regulatory environments over time in ways that cause different responses of firms to existing regulations: for instance, changes in staffing levels, attention or focus of individual agencies or personnel, or updated reporting forms. In order to measure more precisely the effect of the disclosure law conditional on these time-varying, operator-level and state-level characteristics, I incorporate operator-year dummy variables and state-time trends. (Unfortunately using dummy variables at the state-time level would be collinear with the treatment, so instead I interact state dummy variables with an annual time trend.) The econometric specification is then:

$$Y_{ipjst} = \alpha_s + \lambda_t + \sum_{s=1}^{7} \alpha_s \times t + \sum_{j=1}^{J} \theta_j \times \phi_t + \delta D_{st} + \gamma W_i + \theta_j + \psi_p + \sum_{p=1}^{P} \psi_p \times \lambda_t + \epsilon_{ijpst} \quad (2.2)$$

This specification provides an improved focus on the effect of the regulation conditional on potentially confounding variables. However, it does not allow me to measure any differences in the effect of the regulation over time. To address this, I replace the single difference-in-differences variable D_{st} with one that is allowed to vary over time, and measure the effect of the regulation 6 quarters prior and 12 quarters after the date the regulation becomes effective. Thus, this specification is

$$Y_{ipjst} = \alpha_s + \lambda_t + \sum_{s=1}^{7} \alpha_s \times t + \sum_{j=1}^{J} \theta_j \times \phi_t + \sum_{\tau=-6}^{12} \delta_\tau D_{st} + \gamma W_i + \theta_j + \psi_p + \sum_{p=1}^{P} \psi_p \times \lambda_t + \epsilon_{ijpst} \quad (2.3)$$

Equation (2.3) is identical to (2.2) except that D_{st} is an indicator equal to one if disclosure was mandatory in state s at time $t - \tau$. The δ_τ coefficients are the differences in differences corresponding to each time period. By plotting these δ_τ coefficients for each quarter, I can test for differences in the effect of the mandatory disclosure "treatment" over 18 months prior and 36 months following the treatment. I can also verify empirically whether the parallel trend assumption holds: that is, whether the pre-treatment differences between treatment and control are approximately zero.

To calculate standard errors and confidence intervals, I cluster standard errors at the state level, reflecting standard guidance for difference-in-differences (and other) analyses since the model residuals are likely not independent within groups (Angrist and Pischke, 2008; Bertrand et al., 2004). However, the usual procedure for estimating cluster-robust standard errors in OLS (White, 2000) relies on three assumptions for consistent estimates (Mackinnon and Webb, 2016), and two of those assumptions (large number of clusters, and "balanced" clusters with roughly equal numbers of observations) are likely violated in this setting, since I have just seven clusters and one includes about half of my observations. Furthermore, estimated standard errors are almost always too low when these assumptions are violated, leading to over-rejection of the null hypothesis. A number of alternative procedures have been suggested for consistent estimation of standard errors (Cameron et al., 2008; Cameron and Miller, 2015; Mackinnon and Webb, 2016). I implement two of these methods. The main results reported in Section 2.5 use the $G-1$ degrees of freedom correction, which amounts to a t-test with degrees of freedom equal to the number of clusters G minus one. Cameron and Miller (2015) advise this as an easily implemented approach that offers substantial improvements in consistency of estimates. As a robustness check, I also calculate p-values using a wild cluster bootstrap, following the procedures in Cameron and Miller (2015) with one critical improvement (Webb, 2014; Mackinnon and Webb, 2016).[8]

[8] The "regular" bootstrap involves resampling observations from the original sample, with replacement; in a cluster setting, this is applied as a block bootstrap method where the blocks are clusters. However, in Monte Carlo tests, Cameron et al. (2008) find this does not eliminate over-rejection in settings with few clusters. The wild cluster bootstrap performs better with few clusters. Instead of

2.4.2 Identifying assumptions for causal interpretation

The difference-in-differences model is a natural choice suggested by state-level differences in regulatory timing. To interpret δ as the causal effect of the regulation, I must make three identifying assumptions. Two of these are common to difference-in-differences approaches: exogenous timing of the treatment and parallel trends prior to the treatment. The third is peculiar to my setting: the appropriateness of using voluntarily reported data to serve as the pre-treatment measure of behavior.

Exogenous timing of treatment

First, I assume the timing of treatment is exogenous. Given the historically close relationship between the U.S. oil and gas industry and its regulators, it is possible that regulators issued the mandatory disclosure requirement only after operators signaled their readiness, which presumably would occur only when they had minimized its toxic or controversial components. If this were true, it would imply the treatment is not random with respect to the variable of interest, and my estimate of the treatment effect would be biased. However, if there is some endogeneity in the timing of the regulations, meaning that states did not receive the regulatory "treatment" until the industry had reduced the toxicity of their formulas, the impact would be to diminish the effect of interest—that is, a negative bias on the coefficient δ in equations (2.1)

---

sampling observations using pairs of right-hand-side and left-hand-side variables x, y, we calculate a new y^* equal to a predicted value of y plus a weighted residual, where the weights follow a specific random distribution. In the wild cluster, every observation in the original dataset is used exactly once in each bootstrap replication (Webb, 2014). Thus, the randomness that underlies the power of the bootstrap arises from the randomly weighted residual that is added to $\hat{y}$. This is unlike the "regular" bootstrap (better denoted a "pairs cluster," as Cameron et al. (2008) point out), in which the randomness arises from the fact that each observation in the original dataset may be drawn one or more times, or not at all. In a setting with very few clusters—as here—Cameron et al. (2008) and Cameron and Miller (2015) advise weighting the residual by a value chosen from the Rademacher distribution, which takes the values -1 and $+1$ with equal probability. Unfortunately the Rademacher weighting scheme produces at most 2^G different values, which implies that hypothesis testing using the wild cluster bootstrap can be problematic when $G < 10$ or so. Webb (2014) and Mackinnon and Webb (2016) demonstrate why a particular six-point distribution offers better results for small G. I use this "Webb" distribution in implementing the wild cluster bootstrap.

and (2.2) and the coefficients δ_τ (for $\tau \geqslant 0$) in (2.3).

Media reports from the relevant period, and my interviews with industry personnel, provide mixed evidence with respect to industry perspectives regarding chemical disclosure. On balance it seems that early industry opposition to disclosure gave way to nonchalance or acceptance of state laws for disclosure, at least in public statements, but industry has opposed federal disclosure requirements throughout this period. The 2005 Energy Policy Act exempted hydraulic fracturing from EPA regulation under the Underground Injection Control provisions of the Safe Drinking Water Act; there is controversy over whether this represented a clarification or a new loophole, but in any case the industry supported the exemption, which among other things removed any obligation to disclose chemicals. The industry also opposed the Fracturing Responsibility and Awareness of Chemicals Act (FRAC Act), first proposed in the 111th Congress (2009-2010), which would have reversed the SDWA exemption and also required disclosure of fracturing fluid chemicals.

A Congressional investigation, begun in February 2010, provided the first public information on chemical use. The investigators found that many operators did not know what chemicals were used in their operations, nor did their contractors; rather, third-party manufacturers held the information as trade secrets. Waxman et al. (2011) concluded that "... it appears that the companies are injecting fluids containing unknown chemicals about which they may have limited understanding of the potential risks posed to human health and the environment." Also during 2010, New York became the first state to issue a moratorium on permits for hydraulic fracturing, in part due to concerns about chemical use. Perhaps realizing that regulation would come in some form, and evidently preferring state to federal regulation, when the first state regulations were passed in late 2010 operators and industry groups generally expressed nonchalance about the new state laws, and voiced confidence that they could meet disclosure provisions (Sider, 2010).[9]

This expressed acceptance of state disclosure laws continued as more states proposed and passed disclosure regulations. In April 2011 the GWPC and IOGCC opened the FracFocus registry, with support of industry actors, to facilitate disclosure of individual reports (though the registry did not facilitate comparative analysis, as noted in Section 2.3.1). However, industry actors continued to oppose federal regulations that would have required disclosure, including the reintroduction of the FRAC Act in the 112th, 113th, and 114th (2015-16) Congress. Industry representatives also opposed disclosure provisions in two proposed federal regulations. These include a 2014 EPA proposal to require recordkeeping and reporting of fracturing chemicals under the Toxic Substances and Control Act and a BLM rule issued in 2015 (and struck down in 2016) that set standards for fracturing on federal lands.

One theoretical framework that explains the industry behavior with respect to disclosure regulations is that of Lyon and Maxwell (2004), who suggest that corporate environmental strategy is underlain by a dynamic political economy game, in which firms respond to or encourage some forms of regulation in order to forestall other forms that are expected to be more costly or onerous, or less strategically advantageous. In this case, industry may have acceded to state regulation requiring chemical disclosure because they believed it could help to assuage public and regulatory concern. This in turn may help forestall more stringent federal regulation and/or more restrictive local, state or federal regulations, such as bans, moratoria, or large setbacks from incompatible land uses. In this context, the timing of the state regulations does appear to be exogenous, since the industry opposed early attempts to require disclosure, and expressed nonchalance about (state) disclosure regulations

[9] Two industry members I spoke with also indicated that some operators may have welcomed the disclosure requirements because it would force vendors to disclose additional information to operators about chemical use, thus reducing asymmetric information that hinders operators' bargaining with vendors and suppliers. However, this does not explain why operators opposed federal regulation, which would presumably have given them even more power.

only after the Congressional investigation that began in early 2010.

Parallel trends

The second identifying assumption is that events that affect fluid composition in the treatment states also affect the others. For instance, a state law banning the use of certain chemical additives or a price change that affects states differently could violate this assumption. Neither possibility seems plausible here. Chemical-specific bans have not been applied to fracturing fluid in any state I analyze during the period of study. Conditional on time, prices for chemical inputs are likely to be similar since chemicals travel through broad and efficient transportation networks, so the law of one price should hold. I use time fixed effects in all specifications, and I also specifically check whether the parallel trend assumption holds empirically by analyzing the difference in treatment versus control states for six quarters prior to the regulation (see equation (2.3)).

Use of voluntary reports

Since I observe only the fracture fluid data that operators voluntarily disclosed, I must consider if there are selection effects that make the voluntarily disclosed data a nonrandom sample of the full pretreatment population. Theories of corporate environmental strategy generally suggest that when companies have a choice about what kinds of activities to reveal publicly, they will choose to report those that appear more socially or environmentally responsive. That is, the direction of any selection effect is likely toward less toxicity: fractures reported voluntarily are probably those in which operators use fewer toxic chemicals, in expectation, compared to the full (unobserved) population. If voluntarily reported fractures are indeed cleaner than the general population, any bias in the estimate of δ due to this effect is positive. Thus if, for instance, from equation (2.1) we estimate $\hat{\delta} = -0.5$, then the true value is $\delta < -0.5$: that is, the bias due to the voluntary selection effect is positive, and a finding of a significant negative effect of the regulation is conservative.

Nonetheless, it is useful to verify empirically if voluntarily reported fractures are cleaner than those not reported. To test this, I exploit data from an unusual 14-month period in Pennsylvania in which some companies voluntarily reported fracturing fluid contents to a relatively accessible public website (FracFocus.org) while others reported fluid contents to the state regulatory agency in a format that was technically public, but quite difficult to access. Pennsylvania's first disclosure law, effective in February 2011, required companies to report fracturing fluid contents to the state Department of Environmental Protection (DEP). The resulting reports were available for public inspection at regional DEP offices, where individuals interested in accessing them had to identify the permit number of a specific well, contact the appropriate regional DEP office, file a request, schedule an appointment to visit in person (typically three to four weeks in advance), and review a limited number of hard copy documents onsite.[10] Thus, while the Pennsylvania law technically constituted public disclosure, it was primarily geared toward individuals interested in a specific well or small group of wells, and in general operators would not have expected widespread public inspection of their fluid contents. Moreover, mainstream media coverage of the new regulation was minimal: to my knowledge only one mainstream news article, published nearly twelve months after the regulation effective date, mentions the 2011 law (Maykuth, 2012).

At the same time, many operators chose to report chemical additives to the national web-based registry FracFocus, where any individual with access to the internet

[10] Individuals could also schedule an appointment at one of two Pennsylvania locations to use the Integrated Records and Information System (IRIS), in which some chemical disclosure forms were available as scanned PDF documents. The IRIS system is also available on a subscription basis for a substantial fee. However, there was a long wait time for reports to be scanned and uploaded to this system, especially in the height of the fracturing boom in that state. Having used IRIS over an extensive period, I found the wait time was highly variable but could be on the order of 18 to 24 months for some reports.

could quickly download the same information provided they had the well location (e.g., state and county) or identification number. Operators made this decision on a well-by-well basis. Although an industry coalition encouraged its members to report to FracFocus starting in January 2012, reporting to FracFocus was entirely voluntary in Pennsylvania until April 2012. On April 14 of that year a new Pennsylvania law took effect, requiring operators to disclose chemical information directly to FracFocus and replacing the requirement to report to the state.

Using the inspection procedures described above, I obtained information on chemicals used in fracturing fluid for 344 unconventional wells in Pennsylvania for which operators revealed chemical additives only to regulators under the 2011 law. At the same time, operators published information on chemical additives to FracFocus for 1,527 wells fractured between the effective date of the first law (February 5, 2011) and the second law that required disclosure to FracFocus (April 14, 2012). Table 2.7 compares the log concentration of PTRCs and high-media-profile chemicals between the two sets of reports. The public reports disclosed to FracFocus have significantly lower concentrations of PTRCs, and this difference is statistically significant ($p<0.01$). However, there is no difference between the log concentrations of high-media-profile chemicals.

This supports the idea that the full-reporting effect is large and positive. This is consistent with standard theories of corporate environmental strategy: voluntarily reported fractures are indeed cleaner, on average, than those in the general population. Thus, the use of voluntarily reported data as the pre-treatment measure of toxic chemical use in fracturing fluid is likely to result in a lower-bound estimate of the effects of the disclosure regulations.

Table 2.7: Comparison of chemical use in FracFocus and PADEP reports during "semi-public" disclosure period

Measure	FracFocus Mean (SE)	PADEP Mean (SE)	Difference in means (SE)
Log ppm PTRCs	-1.65	-1.09	-0.55***
	(0.07)	(0.19)	(0.20)
Log ppm high-media-profile chemicals	1.02	1.01	-0.01
	(0.07)	(0.22)	(0.23)

Includes reports for 1,527 wells disclosed to FracFocus and 344 wells disclosed to PADEP.
Significance test allows for unequal variances by group. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

2.5 Results

I report regression results separately for each of the dependent variables. In each set of results, I distinguish between the full set of fractures and fractures conducted by VR75 operators. By construction, the magnitude of the full-reporting effect is smaller for VR75 operators; thus, the results of analyses limited to the VR75 operators identify more clearly the disclosure-pressure effect.

2.5.1 Relative toxicity

Table 2.8 reports results from the estimation of models in which the dependent variable is the log of relative toxicity. Column 1 shows the results of estimating equation (2.1), regressing log relative toxicity on available well characteristics with fixed effects for state, year, geologic play, and operator. In this model the effect of disclosure laws is positive, though not significant. However, the observed effect is a composite of the full reporting effect (which we expect is positive and large, as suggested by the analysis in Section 2.4.2) and the disclosure pressure effect. Column 4 shows the result for VR75 operators. The effect of mandatory disclosure for this subset of operators is still positive, but barely. This is consistent with the expectation that the full reporting effect is smaller in magnitude for operators who voluntarily report a large proportion of wells.

Table 2.8: Regression results for log relative toxicity score

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
Depth (10^4 ft)	0.09	0.10	0.10	-0.14	-0.03	-0.04
	(0.25)	(0.21)	(0.21)	(0.34)	(0.46)	(0.45)
Fluid volume	-0.06**	-0.06***	-0.06***	-0.11**	-0.11***	-0.11***
(10^6 gal)	(0.02)	(0.01)	(0.01)	(0.04)	(0.02)	(0.02)
Oil well	0.29*	0.19*	0.21**	0.61	0.32**	0.33**
	(0.14)	(0.09)	(0.08)	(0.33)	(0.11)	(0.11)
Vertical	-0.09	-0.02	-0.02	-0.16	0.06	0.06
	(0.14)	(0.06)	(0.06)	(0.14)	(0.09)	(0.09)
Mandatory	0.09	0.16		0.02	0.08	
disclosure	(0.16)	(0.09)		(0.11)	(0.10)	
6 qtrs before reg.			0.01			-0.10
			(0.51)			(0.66)
5 qtrs before reg.			0.37			0.20
			(0.20)			(0.15)
4 qtrs before reg.			0.16			0.04
			(0.14)			(0.06)
3 qtrs before reg.			0.26			0.10
			(0.18)			(0.14)
2 qtrs before reg.			0.03			-0.20**
			(0.05)			(0.07)
0 qtrs after reg.			0.11			-0.05
			(0.08)			(0.06)
1 qtr after reg.			0.18			0.22
			(0.15)			(0.22)
2 qtrs after reg.			0.26*			0.15
			(0.11)			(0.11)
3 qtrs after reg.			0.10			-0.47*
			(0.25)			(0.23)
4 qtrs after reg.			-0.04			-0.81*
			(0.22)			(0.33)
5 qtrs after reg.			-0.11			-0.69**
			(0.21)			(0.27)
6 qtrs after reg.			-0.05			-0.87**
			(0.31)			(0.28)

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
7 qtrs after reg.			0.13			-0.73**
			(0.49)			(0.26)
8 qtrs after reg.			0.02			-0.84**
			(0.51)			(0.29)
9 qtrs after reg.			0.16			-0.78*
			(0.48)			(0.32)
10 qtrs after reg.			0.16			-0.89**
			(0.45)			(0.31)
11 qtrs after reg.			0.34			-0.81*
			(0.54)			(0.34)
12 qtrs after reg.			0.79			-0.64
			(0.48)			(0.62)
FEs: Oper., state, year, play	✓	✓	✓	✓	✓	✓
FEs: Play × year	✓	✓	✓	✓	✓	✓
FEs: Oper. × year		✓	✓		✓	✓
State-year trends		✓	✓		✓	✓
R^2	0.31	0.41	0.41	0.35	0.44	0.44
N	71,989	71,989	71,989	33,914	33,914	33,914

Standard errors are clustered by state, and adjusted for "few clusters" by using $T(G-1)$ critical values (Cameron and Miller, 2015).
* p<0.10, ** p<0.05, *** p<0.01.

Column 2 shows the results of estimating equation (2.2), which includes all of the variables from equation (2.1) but adds a state time trend and fixed effects for operator-by-year. The effect of mandatory disclosure in this model is still positive, and is larger in magnitude. Column 5 shows the same model for the subset of voluntary reporters; once again the effect of mandatory disclosure appears positive though not significant.

Column 3 of Table 2.8 shows the results of estimating equation (2.3), which is identical to equation (2.2) but demonstrates the effect of the disclosure regulation over time both before and after the effective date. This result is also shown visually in Figure 2.1. The corresponding results for VR75 operators are shown in column 6 of Table 2.8, and Figure 2.2. For both sets of operators, no distinct pre-trend

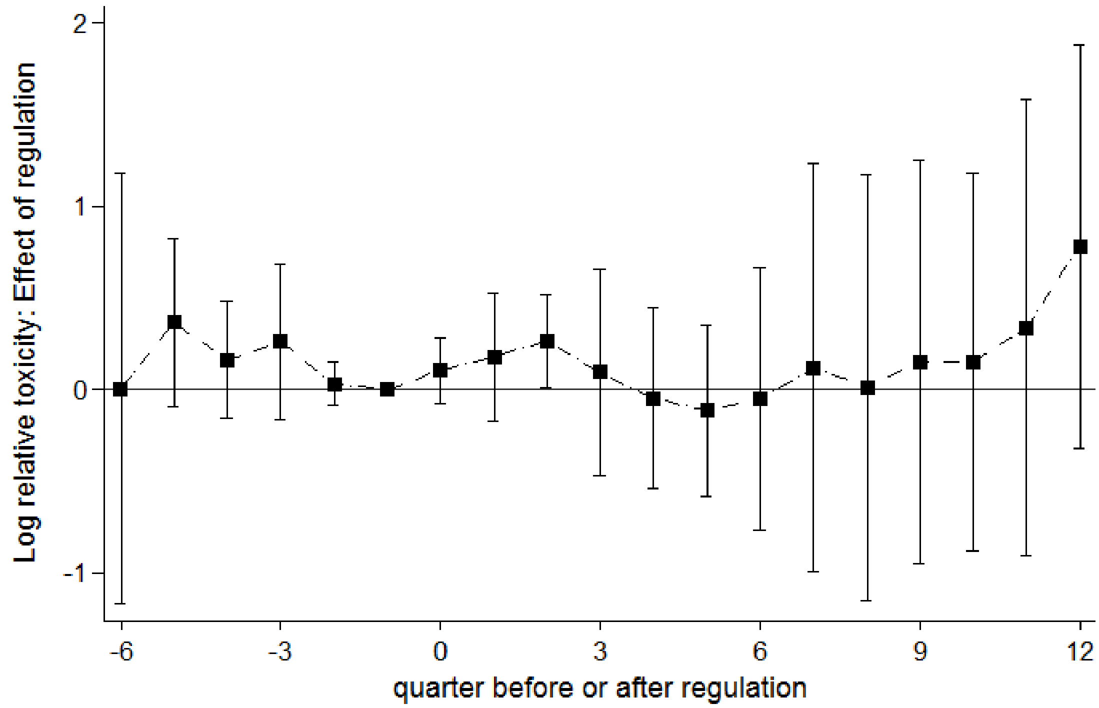

FIGURE 2.1: Difference in differences for log relative toxicity score, all operators

N=71,989 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

appears for the difference between control and treatment, validating the assumption of parallel trends. Figure 2.1 does not suggest the disclosure rule had any effect, based on the composite effect of full reporting and disclosure pressure. However, Figure 2.2 and column 6 of Table 2.8, isolating more clearly the disclosure-pressure effect, do suggest the regulation had a negative effect on the relative toxicity of fracturing fluids. Furthermore, the effect is statistically significant ($p<0.05$) starting about three quarters after the regulation.

2.5.2 Priority toxic and regulated chemicals

Table 2.9 shows the results for priority toxic and regulated chemicals, organized identically to those in the prior table. The results for all operators in columns (1) and

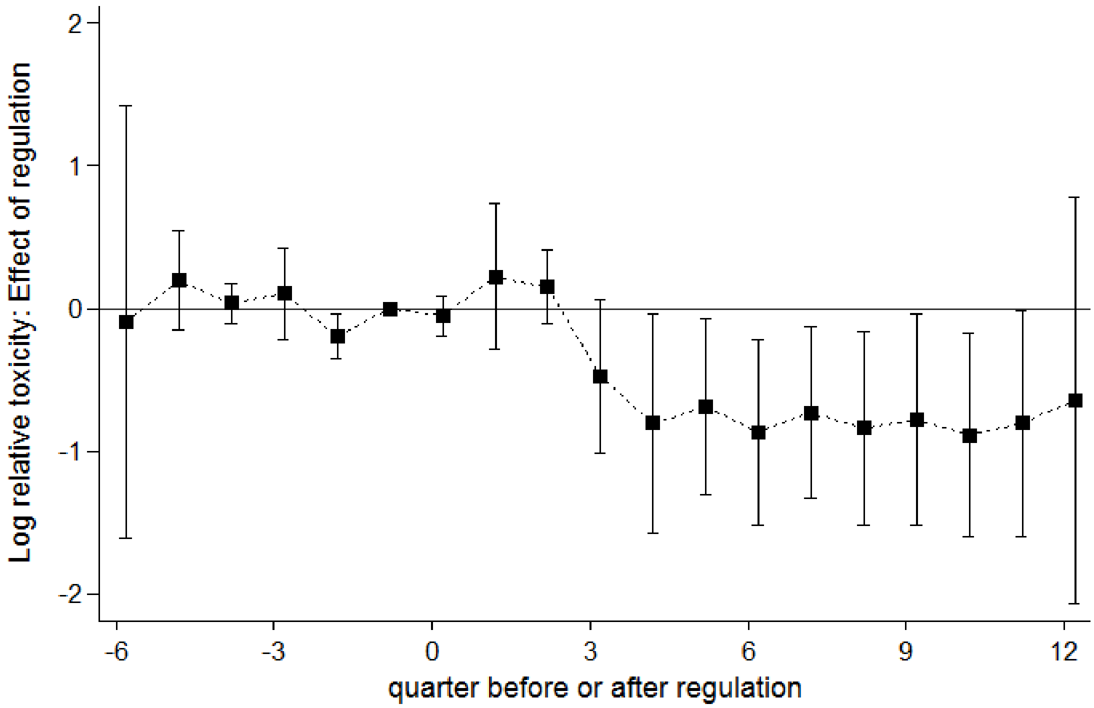

FIGURE 2.2: Difference in differences for log relative toxicity score, VR75 operators

N=33,914 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

(2) suggest that the combination of the full reporting effect and disclosure pressure effect is slightly negative, though not significantly different from zero. The results for VR75 operators in columns (4) and (5) are also negative, with the point estimate implying the disclosure regulation caused a reduction in the concentration of priority toxic and regulated chemicals of about 15-22%; however, this result is not statistically different from zero. Figures 2.3 and 2.4 show the results of the regulations over time. Both figures support the validity of the parallel pre-trend assumption. The results for all operators (Figure 2.3) show little change over time; the point estimate for the effect of the regulation is negative for most of the period after regulation, but generally not statistically significant. Focusing on VR75 operators, which more cleanly isolates the disclosure-pressure effect, the regulation appears to have had a consistently negative

impact on the use of toxics over time, with the largest effect starting about one year after the regulation. The magnitude of the coefficient estimate is on the order of -1.5 for the period starting about three quarters after the regulation (Table 2.9, column 6), implying a decrease in regulated states of average concentrations of priority toxic and regulated chemicals by about 78% ($e^{-1.5} - 1$).

Table 2.9: Regression results for log ppm of PTRCs

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
Depth (10^4 ft)	0.70***	0.77***	0.78***	0.14	0.27	0.27
	(0.14)	(0.16)	(0.16)	(0.37)	(0.59)	(0.57)
Fluid volume	-0.03***	-0.04**	-0.04**	-0.08**	-0.09***	-0.09***
(10^6 gal)	(0.01)	(0.01)	(0.01)	(0.03)	(0.02)	(0.02)
Oil well	0.01	-0.09	-0.07	0.15	-0.11	-0.08
	(0.15)	(0.06)	(0.05)	(0.41)	(0.17)	(0.19)
Vertical	-0.00	0.04	0.04	-0.09	0.15	0.15
	(0.16)	(0.06)	(0.06)	(0.16)	(0.13)	(0.13)
Mandatory	-0.19	-0.07		-0.26	-0.16	
disclosure	(0.22)	(0.11)		(0.20)	(0.12)	
6 qtrs before reg.			-0.04			-0.30
			(0.92)			(1.50)
5 qtrs before reg.			0.10			0.01
			(0.24)			(0.22)
4 qtrs before reg.			0.18**			0.07
			(0.07)			(0.11)
3 qtrs before reg.			0.21			0.10
			(0.17)			(0.12)
2 qtrs before reg.			0.02			-0.25**
			(0.05)			(0.08)
0 qtrs after reg.			-0.06			-0.14
			(0.11)			(0.14)
1 qtr after reg.			-0.05			-0.02
			(0.19)			(0.25)
2 qtrs after reg.			-0.06			-0.31**
			(0.14)			(0.12)
3 qtrs after reg.			-0.21			-1.15*
			(0.28)			(0.55)

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
4 qtrs after reg.			-0.48			-1.72**
			(0.27)			(0.64)
5 qtrs after reg.			-0.52*			-1.53**
			(0.22)			(0.57)
6 qtrs after reg.			-0.45			-1.85**
			(0.31)			(0.52)
7 qtrs after reg.			-0.19			-1.38*
			(0.35)			(0.59)
8 qtrs after reg.			-0.26			-1.58*
			(0.40)			(0.65)
9 qtrs after reg.			-0.10			-1.41*
			(0.38)			(0.66)
10 qtrs after reg.			-0.08			-1.50*
			(0.35)			(0.67)
11 qtrs after reg.			0.24			-1.15
			(0.37)			(0.70)
12 qtrs after reg.			0.44			-1.38
			(0.28)			(1.02)
FEs: Oper., state, year, play	✓	✓	✓	✓	✓	✓
FEs: Play × year	✓	✓	✓	✓	✓	✓
FEs: Oper. × year		✓	✓		✓	✓
State-year trends		✓	✓		✓	✓
R^2	0.30	0.38	0.38	0.29	0.36	0.37
N	71,989	71,989	71,989	33,914	33,914	33,914

Standard errors are clustered by state, and adjusted for "few clusters" by using $T(G-1)$ critical values (Cameron and Miller, 2015).
* p<0.10, ** p<0.05, *** p<0.01.

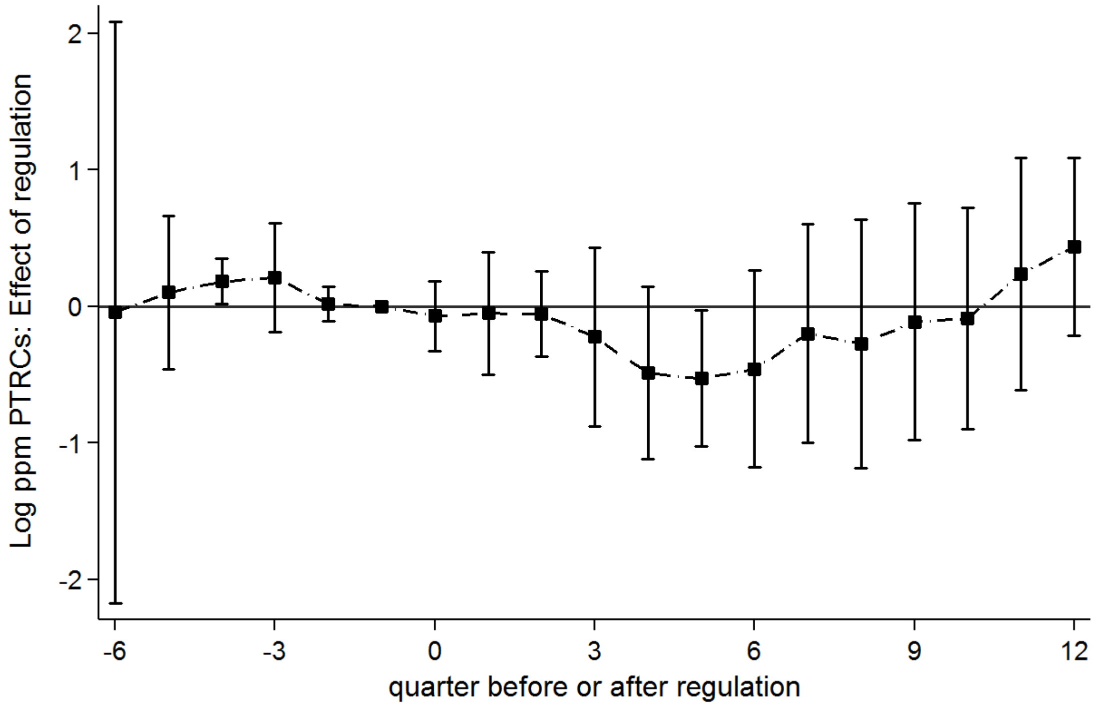

FIGURE 2.3: Difference in differences for log concentration PTRCs, all operators

N=71,989 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

2.5.3 High-media-profile chemicals

Results for high-media-profile chemicals are shown in Table 2.10. As with the previous sets of results, columns (1) and (2) show the effect for all operators, and suggest that the combined disclosure pressure effect and full reporting effect resulted in greater use of these chemicals. Furthermore, the results for the VR75 subset in columns (4) and (5) suggest the counterintuitive result that the disclosure pressure effect was positive; that is, the disclosure regulations increased operators' use of chemicals identified by media reports as potentially dangerous. Further, the increase is of relatively high magnitude, at least in the point estimate: about 63-86% for the VR75 subset, and about 43-58% for all operators.

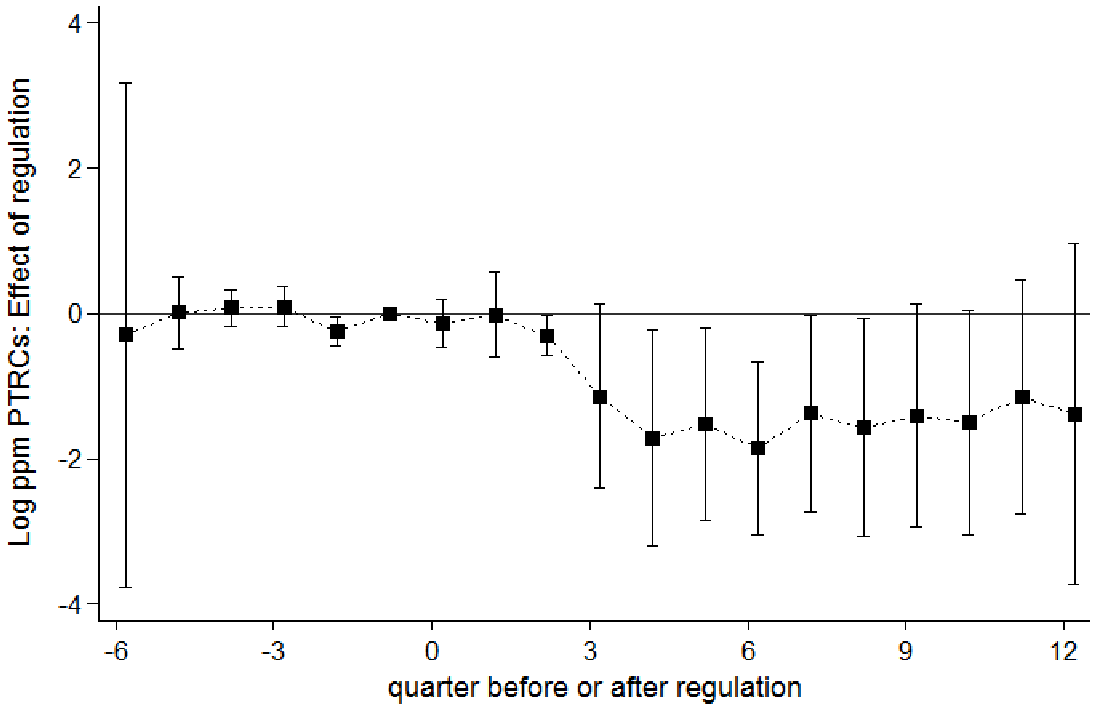

Figure 2.4: Difference in differences for log concentration PTRCs, VR75 operators

N=33,914 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

Figures 2.5 and 2.6 show the effects of the disclosure laws over time. Considering the pre-law trends first, it appears that the use of high-media-profile chemicals was decreasing for treatment states (compared to control states) from about 18 to 12 months prior to the effective date of the disclosure laws, but was about the same for the year preceding the effective date. This could reflect an attempt to deflect disclosure regulations (or other, more costly, regulations) by reducing the use of high-media-profile chemicals during a time of increased regulatory attention, although without further quantitative or qualitative evidence this is uncertain. Turning to the trends after the laws came into effect, firms' immediate response, even in Figure 2.6 where the disclosure-pressure effect should dominate, appears to be to increase the use of high-profile chemicals. This increase is persistent for at least two years after the effective date.

Table 2.10: Regression results for log ppm high-media-profile chemicals

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
Depth (10^4 ft)	-0.50	-0.50*	-0.50*	-0.07	0.03	-0.01
	(0.27)	(0.25)	(0.25)	(0.63)	(0.42)	(0.44)
Fluid volume	-0.07***	-0.07***	-0.06***	-0.09**	-0.09**	-0.08**
(10^6 gal)	(0.02)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)
Oil well	0.29**	0.19	0.17	0.46**	0.26	0.23
	(0.12)	(0.17)	(0.17)	(0.14)	(0.22)	(0.20)
Vertical	-0.00	0.07	0.07	-0.32***	-0.15	-0.13
	(0.12)	(0.07)	(0.07)	(0.06)	(0.12)	(0.13)
Mandatory	0.36	0.47		0.50	0.62	
disclosure	(0.37)	(0.43)		(0.47)	(0.49)	
6 qtrs before reg.			0.53			0.62
			(0.27)			(0.32)
5 qtrs before reg.			0.33**			0.45**
			(0.10)			(0.15)
4 qtrs before reg.			0.14			0.29
			(0.12)			(0.25)
3 qtrs before reg.			0.11			0.32**
			(0.07)			(0.12)
2 qtrs before reg.			-0.11			-0.02
			(0.13)			(0.11)
0 qtrs after reg.			0.46			0.60
			(0.44)			(0.44)
1 qtr after reg.			0.48			0.54
			(0.39)			(0.47)
2 qtrs after reg.			0.39			0.77
			(0.50)			(0.68)
3 qtrs after reg.			0.49			0.55
			(0.44)			(0.51)
4 qtrs after reg.			0.26			0.04
			(0.37)			(0.38)
5 qtrs after reg.			0.21			0.08
			(0.38)			(0.42)
6 qtrs after reg.			0.26			0.07
			(0.44)			(0.54)
7 qtrs after reg.			0.19			0.06
			(0.42)			(0.53)

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
8 qtrs after reg.			-0.13			-0.59
			(0.20)			(0.33)
9 qtrs after reg.			-0.13			-0.77*
			(0.14)			(0.33)
10 qtrs after reg.			-0.42**			-1.44***
			(0.14)			(0.34)
11 qtrs after reg.			-0.19			-1.27***
			(0.15)			(0.27)
12 qtrs after reg.			0.03			-1.32
			(0.25)			(0.70)
FEs: Oper., state, year, play	✓	✓	✓	✓	✓	✓
FEs: Play × year	✓	✓	✓	✓	✓	✓
FEs: Oper. × year		✓	✓		✓	✓
State-year trends		✓	✓		✓	✓
R^2	0.36	0.44	0.44	0.33	0.38	0.39
N	71,989	71,989	71,989	33,914	33,914	33,914

Standard errors are clustered by state, and adjusted for "few clusters" by using $T(G-1)$ critical values (Cameron and Miller, 2015).
* p<0.10, ** p<0.05, *** p<0.01.

These apparently counterintuitive results are consistent with the idea that operators respond to disclosure laws by reducing the use of toxic chemicals regardless of their perceived danger, at least to the degree that perception is measured by media reports. While some companies may have responded to disclosure laws by focusing on reducing the use of chemicals mentioned frequently in media reports, others (including one company who shared with me a summary of their new policy) focus on reducing the use of toxics, regardless of media coverage. To the extent that media focus on more recognizable chemicals regardless of their actual toxicity, the apparent use of greater quantities of high-media-profile chemicals, if they substitute for more toxic but lower-media-profile chemicals, may be desirable from a public health and environmental standpoint. In addition, the observation that reductions in PTRCs occurred sooner, were more consistent over time, and were generally larger in mag-

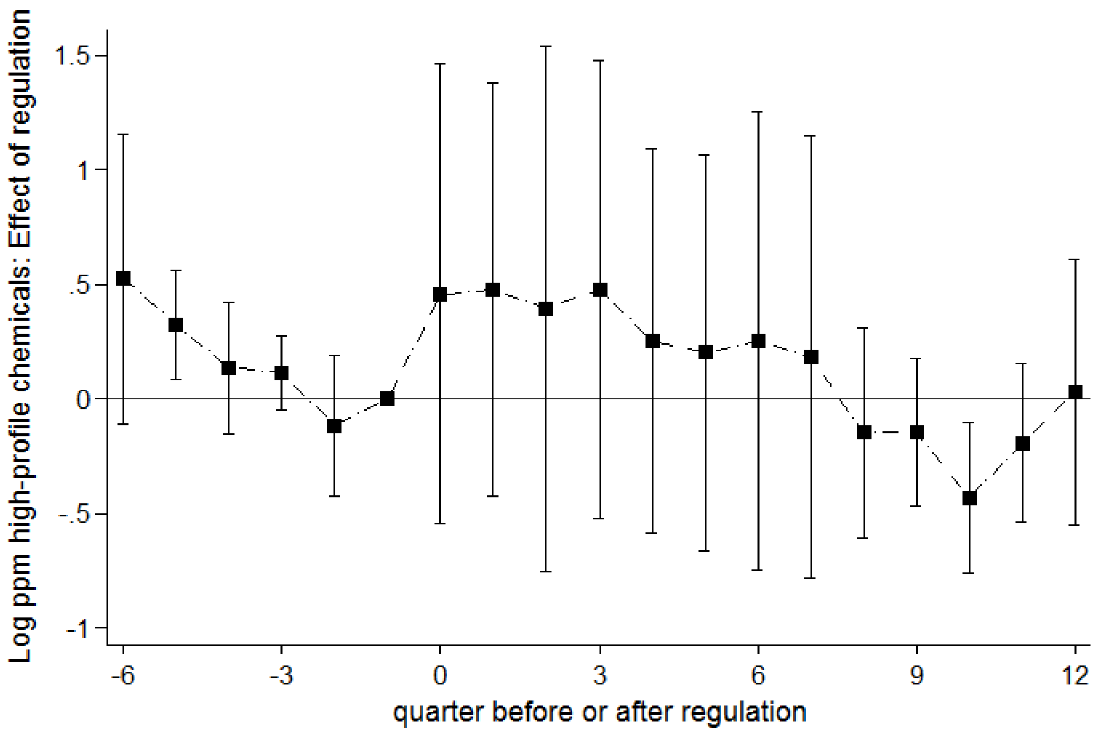

FIGURE 2.5: Difference in differences for log concentration high-media-profile chemicals, all operators

N=71,989 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

nitude also assuages a concern that might otherwise arise, that firms took advantage of disclosure policies that rely on self-reported data and misrepresented their use of toxic chemicals. If misreporting were widespread, we would expect reductions in observed use of high-profile chemicals that were at least as large, if not larger, than that of PTRCs.

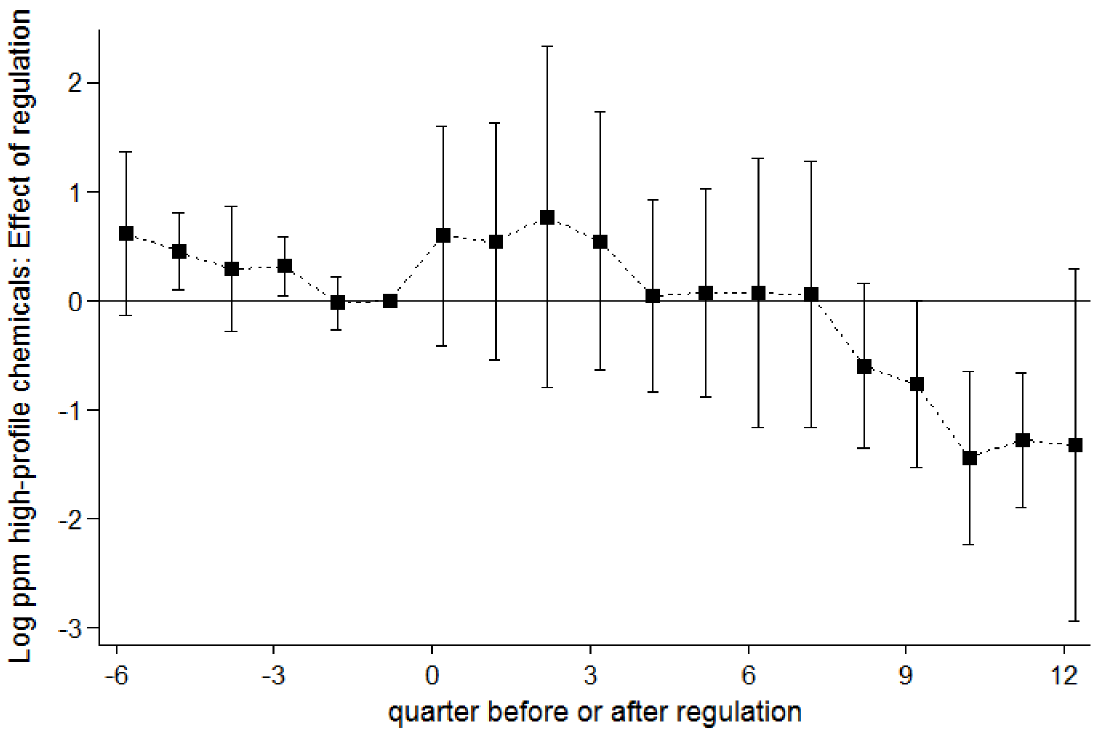

Figure 2.6: Difference in differences for log concentration high-media-profile chemicals, VR75 operators

N=33,914 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

2.5.4 Proprietary chemicals

Table 2.11 and Figures 2.7 and 2.8 show the impact of the disclosure regulations on concentration of chemicals declared as proprietary information. The pre-regulation trend evident in the two figures suggests the parallel trends assumption does not hold: that is, operators used higher concentrations of proprietary chemicals in treatment states, compared to control states, leading up to the passage of the mandatory disclosure regulation. Thus, while the concentration of proprietary chemicals increased in treatment states relative to control states after the mandatory disclosure regulations, it is not clear that this increase is caused by the regulations; it appears that the increase would have happened regardless.

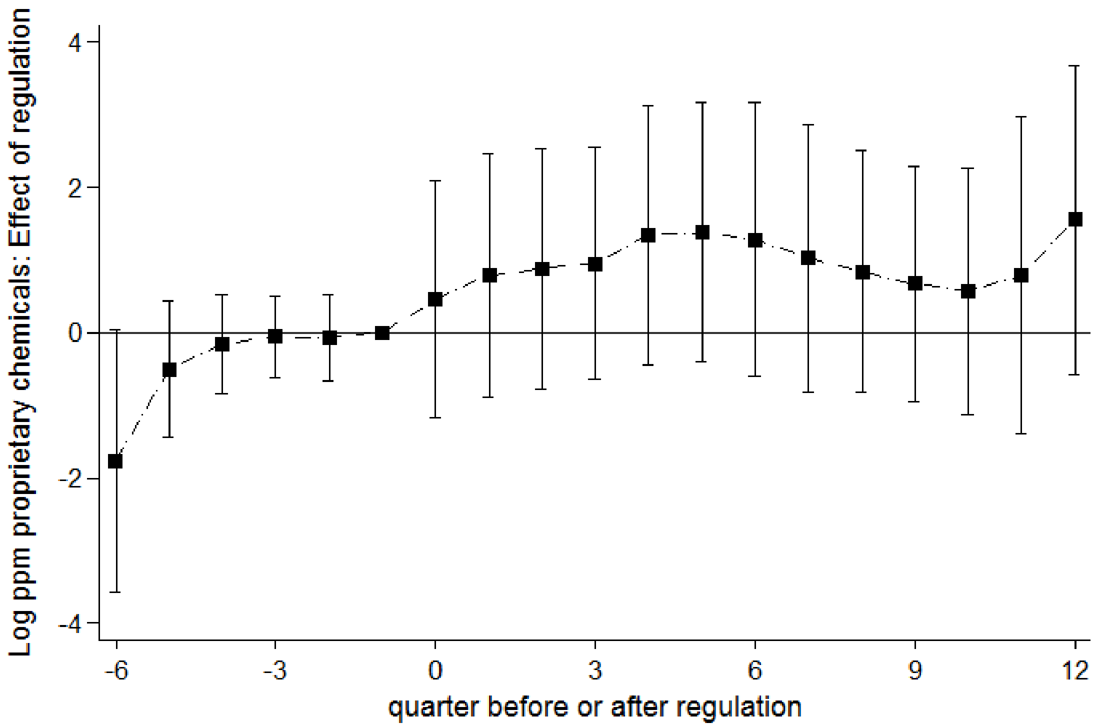

FIGURE 2.7: Difference in differences for log concentration proprietary chemicals, all operators

N=71,989 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

Table 2.11: Regression results for log ppm proprietary chemicals

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
Depth (10^4 ft)	-0.12	-0.05	-0.07	-0.18	0.08	0.07
	(0.53)	(0.56)	(0.56)	(0.58)	(0.62)	(0.62)
Fluid volume (10^6 gal)	-0.03	-0.02	-0.02	-0.09	-0.08	-0.08
	(0.04)	(0.03)	(0.03)	(0.05)	(0.04)	(0.04)
Oil well	0.00	0.20	0.16	0.21	0.48*	0.46
	(0.17)	(0.12)	(0.12)	(0.49)	(0.24)	(0.24)
Vertical	-0.01	-0.13	-0.12	0.07	-0.08	-0.03
	(0.19)	(0.26)	(0.25)	(0.36)	(0.38)	(0.36)
Mandatory disclosure	0.69	0.65		1.04	1.05	
	(0.60)	(0.69)		(0.81)	(0.87)	

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
6 qtrs before reg.			-1.76*			-2.10***
			(0.79)			(0.47)
5 qtrs before reg.			-0.51			-1.11*
			(0.40)			(0.50)
4 qtrs before reg.			-0.16			-0.67
			(0.30)			(0.54)
3 qtrs before reg.			-0.05			-0.38
			(0.24)			(0.37)
2 qtrs before reg.			-0.08			-0.21
			(0.26)			(0.35)
0 qtrs after reg.			0.46			0.73
			(0.71)			(0.85)
1 qtr after reg.			0.79			1.26
			(0.72)			(0.96)
2 qtrs after reg.			0.88			1.35
			(0.72)			(0.94)
3 qtrs after reg.			0.95			1.28
			(0.69)			(0.89)
4 qtrs after reg.			1.34			1.75
			(0.78)			(1.01)
5 qtrs after reg.			1.38			1.91
			(0.77)			(1.08)
6 qtrs after reg.			1.28			1.68
			(0.82)			(1.20)
7 qtrs after reg.			1.03			1.48
			(0.80)			(1.15)
8 qtrs after reg.			0.84			1.11
			(0.72)			(0.99)
9 qtrs after reg.			0.67			0.90
			(0.70)			(0.98)
10 qtrs after reg.			0.57			0.72
			(0.74)			(1.02)
11 qtrs after reg.			0.79			1.09
			(0.94)			(1.58)
12 qtrs after reg.			1.55			1.35
			(0.93)			(1.71)

	All operators			VR75 operators		
	(1)	(2)	(3)	(4)	(5)	(6)
FEs: Oper., state, year, play	✓	✓	✓	✓	✓	✓
FEs: Play × year	✓	✓	✓	✓	✓	✓
FEs: Operator × year		✓	✓		✓	✓
State-year trends		✓	✓		✓	✓
R^2	0.36	0.46	0.46	0.30	0.40	0.40
N	71,989	71,989	71,989	33,914	33,914	33,914

Standard errors are clustered by state, and adjusted for "few clusters" by using $T(G-1)$ critical values (Cameron and Miller, 2015).
* p<0.10, ** p<0.05, *** p<0.01.

Nevertheless, it is worth considering whether aspects of mandatory disclosure regulations would have inspired operators to increase their declaration of proprietary chemicals. When reporting is voluntary, some operators may choose not to list or not to quantify proprietary chemicals, instead listing only those chemicals whose identity is declared. Thus, the observed increase in the quantity of chemicals declared proprietary after disclosure laws may result from the fact that the disclosure laws formalized the necessity to report proprietary chemicals when they are used in fracturing fluid.

Other explanations may also be at work. One possibility is that operators were concerned about revealing trade secrets in the process of complying with disclosure laws, and used the proprietary declaration to avoid revealing strategically valuable information. Another possibility, considering the results from this section together with Sections 2.5.1 and 2.5.2, is that operators sought to avert public or stakeholder pressure by using the proprietary declaration to cover the ongoing use of toxic or regulated chemicals.

2.5.5 Robustness checks

One critical identifying assumption of the difference-in-differences approach is that in the absence of the regulatory treatment, the ex post trends in chemical use would

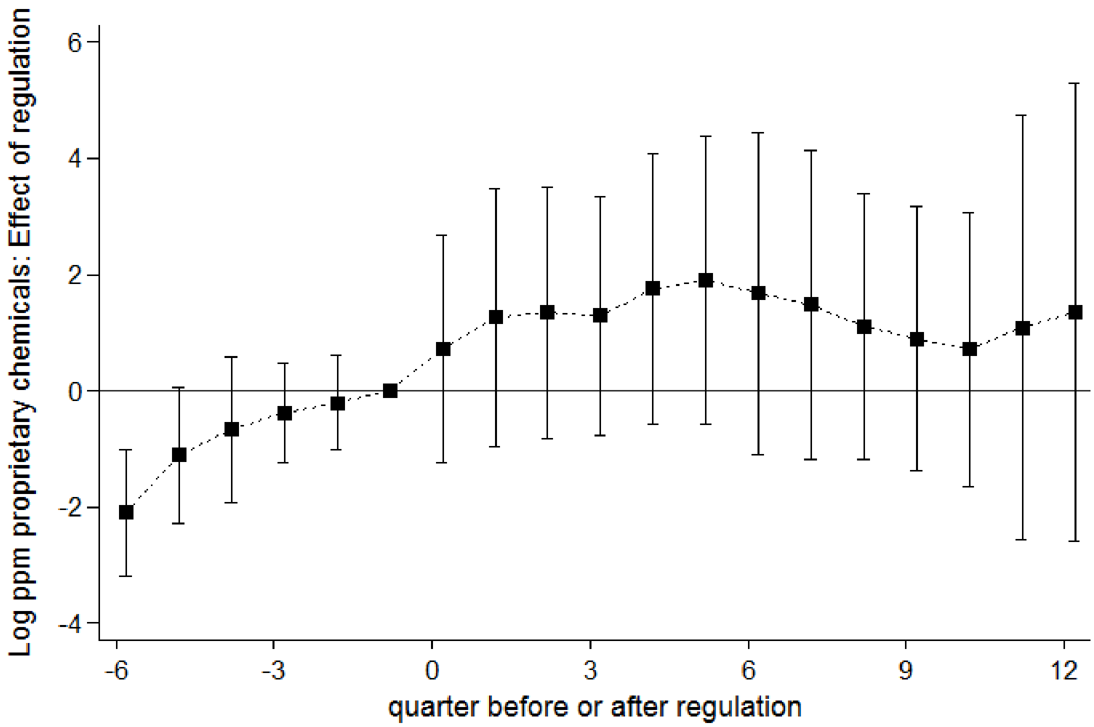

FIGURE 2.8: Difference in differences for log concentration proprietary chemicals, VR75 operators

N=33,914 wells. Error bars show 95% confidence interval. Standard errors, clustered by state, use $T(G-1)$ critical values (Cameron and Miller, 2015).

have been parallel for states with and without the mandatory disclosure regulation. One way to verify the validity of this assumption is to implement a "placebo" regulation that occurs prior to the actual regulation, then run the same regressions as in the main specification and look for significant effects of the placebo regulation. If there are, then the main estimates could reflect some other process (e.g., a secular time trend) rather than a causal effect of the regulation.

To implement this check, I drop all treated observations (i.e., wells that were subject to mandatory disclosure) and then create a placebo policy variable that turns on one year before disclosure becomes mandatory in the corresponding state. I run a difference-in-differences analysis for each dependent variable, using the same specifications described above, to see if the analysis of the placebo policy suggests a significant change due to this (placebo) treatment. In each specification (including the time-varying approach as in (2.3)), the coefficient on the placebo treatment indicator is not statistically significant. This result increases my confidence in the validity of the main conclusion.

As a second robustness check, I consider the possibility that operators decreased the use of all chemical additives after the mandatory disclosure regulations. If this were the case, then this would provide an alternative explanation for the observed decrease in total concentration of toxic and regulated chemicals in fracturing fluid. To test this alternative explanation I run two additional specifications. First, I employ an alternative dependent variable equal to the total concentration of all chemicals (other than water and sand). Second, I choose a random subset of 45 chemicals (i.e., about the same number of elements as the subset of priority toxic and regulated chemicals) and calculate a dependent variable equal to the log of total concentration of that random set. In both cases, the analysis does not indicate the mandatory disclosure regulations had a significant effect.

Finally, I recalculate t-statistics for the main regressions using a wild cluster bootstrap. As noted in Section 2.4, hypothesis tests based on White (2000) cluster-robust standard errors tend to over-reject the null hypothesis. Several alternatives exist to correct for this (Cameron et al., 2008; Cameron and Miller, 2015; Mackinnon and Webb, 2016; Webb, 2014). The results reported above reflect the correction that is most straightforward to implement, which is to use $G - 1$ degrees of freedom for t-tests. I also run an analysis using a wild cluster bootstrap with the six-point Webb distribution (Mackinnon and Webb, 2016; Webb, 2014) to recalculate p-values. The bootstrapped standard errors are somewhat greater than the errors using the $T(G - 1)$ distribution, but coefficients are still statistically significant in general; for instance, for the results corresponding to column 6 of Table 2.9 (that is, log

PTRCs for voluntary reporters) the resulting p-values over the range from four to ten quarters after the regulation range from 0.068 to 0.086. This reinforces the validity of the main conclusion.

2.6 Conclusion and policy implications

Taken together, the difference-in-difference analysis and robustness checks suggest that the requirement for oil and gas operators to disclose the chemical additives used in their hydraulic fracturing operations caused a decrease in both the overall concentration of toxic and regulated chemicals, and the relative toxicity of chemicals used in fracturing fluid. The resulting decrease is not immediate, and its magnitude is difficult to quantify precisely, because the observed effect of switching from a voluntary to a mandatory reporting regime combines a "full reporting effect" and a "disclosure pressure effect." I show that the full reporting effect is large and positive, and that the disclosure pressure effect is large and negative—to the best of my ability to estimate, it is on the order of 78% for average concentrations of priority toxic and regulated chemicals, for the period starting about one year after the regulation. This effect persists for several years after the regulation (and may be permanent).

These results offer encouraging evidence for the hypothesis that mandatory information disclosure regulations can influence companies to change their behavior in ways that decrease potential threats to external stakeholders—in this case, potential harms to human health and the environment that may arise from the use of toxic chemicals. Furthermore, this analysis demonstrates that even companies in non-consumer-facing settings may change their behavior in this way. This supplements analyses in other, more consumer-facing domains such as restaurant hygiene (Jin and Leslie, 2003), drinking water provision (Bennear and Olmstead, 2008), and electricity generation (Delmas et al., 2010). Thus, my analysis provides useful insight for researchers and practitioners regarding how firms respond to mandatory disclosure regulations in settings where consumers have little or no direct influence on firms' activities. In addition, the conclusions of this paper are relevant for policymakers who must choose between alternative regulatory instruments, or specific design elements of instruments, for promoting public welfare.

3

Learning by Viewing? Social Learning, Regulatory Disclosure, and Firm Productivity in Shale Gas (with A. Steck, C. Timmins, and D. Wrenn)

3.1 Introduction

Disclosure laws are used by regulators to disseminate information about potentially damaging activities by firms, with the hope of providing firms with incentives to reduce those activities. While disclosure policies may lead to an environmental benefit, they can also have a cost if the lack of secrecy limits firms' ability to realize the benefits of innovation. The existence of this tradeoff depends on the answers to two questions. First, does the disclosed information have value to firms? Second, does disclosure effectively enable social learning, and thereby limit firms' ability to capitalize on innovation? We study both questions in the context of hydraulic fracturing chemicals, and find affirmative evidence for each. Thus, while the use of toxic substances in hydraulic fracturing fluids makes a case for a public benefit from disclosure, policy makers should weigh that benefit against disclosure's costs on industry.

Hydraulic fracturing is a growing technology that has transformed the nature of energy production in the U.S. and the world, and contributed to substantial economic growth in many areas (Hausman and Kellogg, 2015). It has also raised a number of concerns regarding local economic impacts and averting behavior (e.g. Muehlenbachs et al. (2015); Wrenn et al. (2016)) and local environmental impacts, including the use of toxic chemical additives in the fracturing process (Elgin et al., 2012; Stringfellow et al., 2014). Those concerns have led to policies since 2010 requiring public disclosure of information about the chemicals used, now instituted in 18 states.

Similar disclosure policies designed to disseminate information about pollution releases (or the potential for such releases) have become increasingly popular in many industries. See, for example, policies regulating electricity generation (Delmas et al., 2010; Kim and Lyon, 2011), and drinking water provision (Bennear and Olmstead, 2008). By harnessing the collective power of stakeholders internal and external to the disclosing entity, information-based policies may improve public welfare via market, political, or managerial mechanisms (Fung et al., 2007; Bennear and Olmstead, 2008), as well as provide stakeholders with information that supports their "right to know" about potential environmental or health risks. For example, the regulations may reduce the use of toxic chemicals and potential threats to environmental quality or human health. Analyses within industries that are less consumer-facing, such as manufacturing facilities that report to the Toxics Release Inventory (TRI), find that reported releases of toxic chemicals have declined over the course of the program by as much as 50 percent (Bennear and Coglianese, 2005). In non-environmental contexts, information-based policies have been used to regulate financial disclosure, nutritional labeling, restaurant hygiene, and other matters (Fung et al., 2007; Jin and Leslie, 2003). Information-based policies seem to be especially popular in settings where conventional regulatory approaches are ill-suited, such as when risks are not well understood, but are not anticipated to be extraordinary. They often represent a "pragmatic compromise" (Fung et al., 2007) that is a politically viable response to emerging risks.

In the case of hydraulic fracturing, operators are likely to have better information than regulators regarding which chemicals are most effective for releasing trapped hydrocarbons in a particular shale formation. A consequence of this asymmetric information is that direct regulations (i.e., command-and-control policies) are unlikely to be efficient. In contrast, disclosure regulations, by opening firms to the threat of public pressure or legal action, can incentivize responses that take into account the tradeoffs between productivity impacts and potential sanctions. For instance, firms may respond by using less toxic formulations, or take further actions to prevent leaks and spills.

At the same time, mandatory information disclosure may impose unintended costs on firms. In particular, these policies may limit operators' ability to monetize the value of innovations. Firms in the oil and gas industry, as in many industries, rely on secrecy as a mechanism to capture value from their investments in innovation and maintain their competitive advantage (Wang and Krupnick, 2013; Cohen et al., 2000). Indeed, the industry has made the argument that fluid formulas represent "trade secrets" in response to calls for mandatory disclosure regulations. In this paper, we shed light on whether innovations in fracturing fluids do in fact constitute a competitive advantage. In particular, we compare detailed, well-level information on inputs and outputs to test whether the chemicals used become more similar with public disclosure. We observe evidence that chemical mixes do converge for wells across different firms. We then test whether using chemicals more similar to high-performing wells appears to improve productivity for poorer-performing firms, and find an affirmative answer. These findings suggest policy-makers need to weigh the tradeoff between right-to-know disclosure rules and industry secrecy. A further question is: to what extent does forced disclosure undermine the incentives for firms in this industry to invest in research and development for future innovations? We leave this latter question to on-going work.

Our analysis is primarily empirical, and is set in Pennsylvania. This state provides a unique data environment in which to analyze our questions of interest. Pennsylvania was one of the earliest states to experience a dramatic increase in exploration and production of unconventional shale gas, and continues to experience extensive unconventional development. Furthermore, Pennsylvania features an unusual regulatory episode in which operators were required to disclose information about chemical additives in fluid to the regulator, but not in a format that was easily accessible for the general public (or to one another), for a 14-month period in the height of the boom. We have recovered this information through a combination of "Right-to-Know" law requests and other methods. Observing this information helps us to distinguish the effect of the public disclosure law from other simultaneous phenomena, including a general improvement in technology.

We perform a series of empirical analyses designed to test whether regulations requiring the mandatory disclosure of chemicals created opportunities for social learning—even when those regulations included provisions that permitted operators to declare some chemicals "proprietary." As a first step, we look for evidence that the unexplained variation in production declined over time; that is, if there is a tighter distribution of residuals from a regression of output on observable well characteristics. We then test for convergence in inputs, as well as test for convergence in productivity across operators; in the latter test, we focus on evidence that low-performing firms used information revealed by the disclosure law to catch up with more successful firms. There could, naturally, be other reasons for this convergence. We therefore focus on the specific mechanism of interest, and examine the role of input similarity in the design of fracturing fluid. We find that convergence is indeed explained specifically by weaker firms adopting the chemical mixtures used by more successful wells in previously drilled wells, after the disclosure rules have gone into place.

This paper makes three primary contributions. First, we provide evidence that secrecy is valuable to hydraulic fracturing firms. This fact raises the tradeoff between the public's right to know about fracturing chemicals used and firms' right to secrecy. Second, we study the interplay between information disclosure regulations and social learning. Third, we examine the role that chemicals play in the technology of hydraulic fracturing.

Our first contribution is to the literature on the relationship between secrecy and innovation. We provide some insight into the importance of the social tradeoffs policy makers face when considering disclosure laws. There may be a compelling social benefit from disclosing specific information, such as when production processes involve toxic chemicals in residential areas. This social benefit must be weighed against the ability of firms to realize economic returns from their innovations. Theoretically, patents solve the problem of imperfect appropriability, but in practice, evidence in many industries suggests firms realize a competitive advantage from innovation through a combination of secrecy, lead time, and investment in complementary capabilities (Cohen et al., 2000). We provide evidence that innovations in hydraulic fracturing chemicals are valuable, suggesting that this tradeoff may be important in this case.

Second, this paper expands the literature on the empirical effects of transparency or disclosure regulations in a new direction. Our analysis is the first, to our knowledge, to combine a study of the effects of information disclosure regulations with social learning. We find strong evidence that information disclosure regulations about an emerging technology can enable social learning.[1] Prior papers have documented the effects of disclosure laws on environmental, health, or other outcomes (e.g., Jin and Leslie, 2003; Fung et al., 2007; Bennear and Olmstead, 2008; Delmas et al.,

[1] This could also be interpreted as a form of technology diffusion facilitated by regulation.

2010). In addition, some work has analyzed the effect of environmental disclosures on investor behavior (e.g. Hamilton, 1995). Separately, other authors have studied the phenomenon of social learning (Conley and Udry, 2010; Covert, 2015), without considering the role of disclosure regulations specifically.

Third, we study the role that chemical additives have played in the development of hydraulic fracturing technology. A number of recent papers have studied the rise of hydraulic fracturing, but none have addressed the role played by chemicals. We take the role of chemicals seriously, and thereby contribute to the literature on firm learning about hydraulic fracturing specifically, and emerging technologies more generally.

The paper proceeds as follows. In Section 3.2, we provide context and background for our study, including a review of the emergence of the technologies that have made development of shale gas economically feasible, the environmental concerns that led to the requirement for public disclosure, and the literature on the role of alternative strategies—including secrecy—that firms use to preserve competitive advantage. Section 3.3 describes our data. In Section 3.4 we present our analysis of convergence and catch-up in operators' output, as well as evidence on how the disclosure rules affected chemical input choices. Section 3.5 offers a discussion and concluding remarks.

3.2 Background

3.2.1 Engineering and geology of hydraulic fracturing

The rise of shale gas to prominence in the US energy landscape has been well-documented.[2] The ability to profitably recover hydrocarbons from shale has been

[2] Shale gas grew from 5% of total US dry gas supply in 2004 to 56% in 2015 (http://www.eia.gov/conference/2015/pdf/presentations/staub.pdf). Spurred on by recent developments in hydraulic fracturing and horizontal drilling technologies, natural gas has largely replaced coal in the production of electricity (http://www.cnbc.com/2015/07/14/natural-gas-tops-coal-as-

largely based on advances in four key areas of technology: horizontal drilling, three-dimensional seismic imaging, micro-seismic fracture mapping, and massive hydraulic fracturing (Wang and Krupnick, 2013). Elements of these technologies have been in development for several decades, spurred by both private and public investments in research and development. Technological advances since the 1970s have ranged from changes in the major compounds comprising fracturing fluid to greater control over directional drilling of wellbores. Foam was replaced by gels in the formulation of fracturing fluid, and order-of-magnitude changes were made in the quantity of proppant used. In the 1990's, there were important advances in the role of directional drilling; in combination with massive hydraulic fracturing, the ability to drill horizontally through a shale formation made shale gas development economical. At the same time, "slick water" fracturing fluid replaced gels. Most recently, fracturing fluid has been refined in multiple dimensions for maximizing output and minimizing costs.

A wide array of chemicals are used in hydraulic fracturing fluid to enhance the productivity of the primary inputs—water and sand. In particular, firms use chemical additives to help open fractures in the rock, transport the proppant along the length of the fracture, lower viscosity in order to allow faster pumping and higher pressures, minimize fluid loss into the face of the formation, reduce scaling on the formation, reduce chemical corrosion or bacterial growth that might threaten the integrity of metal casings, facilitate breakup of other chemicals post-fracture, and serve other purposes (Stringfellow et al., 2014; Montgomery, 2013; Gulbis and Hodge, 2000). In short, fracturing fluid is a complex mixture in which an additive that improves performance in one dimension may reduce performance in another. Although the

top-source-of-electric-power-generation-in-us.html). The largest contribution from any one shale play to the growth described above has come from the Marcellus Shale, located in Pennsylvania and West Virginia. Due in part to the availability of pre-disclosure data, Pennsylvania will be our area of study (see Section 3.3).

cost of the chemicals themselves is usually small in comparison to the overall cost of the well stimulation operation, the proper choice of chemicals may have dramatic effects on the overall cost and productivity of a well.[3]

As of 2014, the total estimated recovery for shale oil wells was on the order of 5%, compared to 50% for conventional oil wells.[4] As engineers seek to improve recovery from shale wells, innovation continues on several elements of the technology, including the use of longer fractures, greater use (per foot) of water and proppant, shorter stages and "micro-perforations", and improved identification of naturally existing fractures through higher-resolution micro-seismic mapping. Designing fracturing fluid for optimal performance is complementary to several of these elements, and represents a significant area of focus for oil and gas engineers, for advancement of shale production technology (Robart et al., 2013; Montgomery, 2013; Gulbis and Hodge, 2000).

In some cases, the quest for superior fracturing fluids has led engineers to consider the use of highly toxic chemicals (Stringfellow et al., 2014). Indeed, many of the chemicals used in fracturing fluid are known to be toxic to human health, or may cause damage to the ecosystem. These risks have raised concerns among environmental groups (Elgin et al., 2012; Haas et al., 2012). Early concerns about shale gas development were driven by the possibility that toxic fracturing fluid might migrate to or be accidentally released into ground water or surface water. The industry responded to these concerns by pointing to the small percentage of fracturing fluid that is actually comprised of substances other than water and sand.[5] Public concerns about risks to water sources were accentuated by the often close proximity of

[3] Mark Boling, Southwestern Energy, personal communication.

[4] R. Kleinberg, Schlumberger, April 2014: "Shale Gas & Tight Oil Technology: Evolution & Revolution", presentation to US Association for Energy Economics.

[5] The typical proportion of chemicals in slickwater fracturing fluid, other than water and sand, is 2 to 3%. Nonetheless, for a typical operation that uses on the order of five million gallons of fluid, even 1% of the fracturing fluid would represent 50,000 gallons.

well-pads to residential and other non-industrial land uses, and by a few high-profile incidents of water pollution.[6] Media coverage of fracturing chemicals has highlighted the toxicity of some chemicals along with the industrys desire to maintain secrecy over the specific chemicals involved (Elgin et al., 2012; Haas et al., 2012).

3.2.2 Policy tradeoffs of information disclosure

With respect to potential external impacts associated with the use of toxic chemicals, economic theory suggests a number of alternative approaches that public policymakers might use to regulate such externalities. Command-and-control regulations require firms to undertake particular technologies or practices. When regulators have less information about production processes than do firms, this "one size fits all approach may be suboptimal. Market-based regulations (e.g., severance taxes and impact fees) modify firms incentives via price effects, but may affect firms only on the extensive margin, and may be difficult to implement if chemical releases are difficult to monitor. In contrast, information-based regulations require regulated entities to disclose elements of their production process that may have external impacts, but which would be difficult for outside stakeholders to ascertain without the disclosure requirement. Disclosure regulations are motivated both by the notion that the public has a "right to know the details of firms (and governments) production decisions, especially when those elements conceivably affect public welfare. In addition, disclosure regulations may also motivate regulated entities to change their behavior, either through pressure applied in the marketplace or the courts. The former means that information-based regulations are most commonly used in consumer-facing industries (Fung et al., 2007).

Scholars have posited alternative mechanisms to explain why mandatory disclosure laws may positively influence regulated actors even without a consumer channel.

[6] See, for example, http://www.vanityfair.com/news/2010/06/fracking-in-pennsylvania-201006.

Bennear and Olmstead (2008) identify a political mechanism: information may increase the ability of a concerned public to lobby for stronger regulation. These authors also note that disclosure may affect an organizations internal decision making, as individuals within the firm change their behavior as a result of measuring and reporting data. Another interpretation focuses on the role of liability, which is facilitated by information disclosure. Olmstead and Richardson (2014) list factors that work in favor of liability as a regulatory approach, including asymmetric information (i.e., regulated firms have better information about technologies) and limited ability to avoid payment (e.g., by spinning off risky activities into operations with little exposed assets.[7] If it is unlikely that a suit is brought, either because litigation or information costs are high, the disciplining force of liability will also be muted. Because shale gas development is a highly technical process that often occurs in areas with restricted public access, it may be difficult for the general public, and even regulators, to monitor operator activities. Providing information to nearby landowners and other interested parties reduces this cost.

Konschnik (2014) summarizes other reasons why disclosure is valuable, many of which are more direct. For instance, in the event of an accident, disclosure to emergency medical personnel and medical staff may improve treatment and protect the staff members themselves. In addition, disclosure provides information to nearby landowners and local government authorities so that they can test their water supplies, increasing their ability to bargain with operators. Disclosure can also help establish liability for contamination, making liability a more effective regulatory tool. If the information made available extends over a period of time, then disclosure can facilitate the monitoring of environmental releases, exposures, or health impacts over time. Finally, disclosure satisfies the public's "right-to-know" about possible release

[7] Insurance and bonding requirements in Pennsylvania are small (either $2,500 per well or $25,000 for all wells in the state).

of, or exposure to, hazardous materials, in the spirit of the Emergency Planning and Community Right-to-Know Act (EPCRA) of 1986 (although the shale gas industry is currently exempt from some provisions of this Act).

At the same time, to the extent that industry actors rely on secrecy to maintain competitive advantage from investments in research and development, disclosure potentially undermines incentives to invest in innovation. It is well understood that to the extent that new knowledge generated by investments in research and development can be copied or imitated by other firms that do not pay the full cost of the investment, the social returns to investments in R&D may exceed private returns. In theory, patents solve the problem of imperfect appropriability by granting inventors the right of exclusive use for a period of time (Cohen, 2010). In practice, however, the protection afforded by patents varies substantially across industries, and would-be innovators in many industries rely on secrecy, lead time, and investments in complementary assets in order to maximize the returns to innovative activity (Cohen et al., 2000; Cohen, 2010; Teece, 1986). A survey administered in 1994 to R&D laboratories in the US manufacturing sector indicated that on average, lab managers in the petroleum and chemicals industries considered secrecy more effective than any other mechanism, including patents, for protecting both product and process innovations (Cohen et al., 2000).[8] Thus, in this context, regulations that require companies to disclose information about their use of specific inputs may result in decreased private returns to innovation.[9]

[8] This does not necessarily imply patents are unimportant, as their use may be complementary in an overall strategy to maintain competitive advantage.

[9] Note that much of the technological development that facilitated the "shale revolution" was based on R&D that was directly sponsored by the government, or by government-subsidized R&D in the form of tax credits for development of unconventional shale and tight gas. In response to persistent supply shortages in the 1970s, the US government began to directly and indirectly fund R&D investment in the natural gas industry. The Natural Gas Policy Act of 1978 removed wellhead price controls and provided tax incentives for developing new natural gas resources. The Crude Oil Windfall Profit Tax Act of 1980 provided tax credits for developing unconventional fuels, which increased their financial return and reduced their risk. The Department of Energy (DOE)

3.2.3 Disclosure rules

There are currently eighteen states with significant hydraulic fracturing activity and chemical disclosure laws. There is general uniformity across these states in terms of the required information that operators must disclose, including ingredient name, chemical abstract service (CAS) number, concentration in the fracturing fluid (typically the maximum concentration used in any fracturing stage), the name of the supplier, and the trade name if applicable.

There is less uniformity in terms of where the information must be registered. Five of the eighteen states, including several of the states that passed the earliest disclosure rules, require operators to report to a state regulatory agency or commission. Six require that operators report to FracFocus, an online database created by a multi-state commission in partnership with a non-profit organization (GWPC and IOGCC, 2015). When uploading information to FracFocus, operators are also asked to provide information about well location and characteristics including vertical depth, volume of water used, latitude and longitude, and well name. The seven remaining states allow operators to choose their reporting location (i.e., to FracFocus or the state).[10]

States have adopted similar approaches to accommodating the need for trade secrets, partly due to the Uniform Trade Secrets Act.[11] In particular, all states allow exemptions for the disclosure of additives considered to be confidential business information that firms believe gives them a competitive advantage. Operators

initiated a number of Unconventional Natural Gas Research Programs, including the Eastern Gas Shales Program, the Western Gas Sands Program, the Methane Recovery from Coalbeds Program, the Seismic Technology Program, and the Drilling, Completion and Stimulation Program. While these programs did not have a role in the development of horizontal drilling or 3D seismic imaging technologies, they did play an important role in high-volume fracturing and micro-seismic fracture mapping (Shellenberger et al., 2012; Wang and Krupnick, 2013).

[10] Although Oklahoma notes that the state regulator will upload to FracFocus any information it receives.

[11] The Uniform Trade Secrets Act, which seeks to harmonize standards for trade secret protection, was promulgated by the Uniform Law Commission in 1979 and passed by 46 states.

must declare an exemption for individual chemical ingredients for which they claim trade secret status. This is accommodated by FracFocus, which allows for uploaded information to include the concentration of the chemical used but not its name or chemical identification number. Some states also require operators to report the chemical family to which the proprietary substance belongs.

3.3 Data

We obtained production data for oil and gas wells in Pennsylvania from DrillingInfo, a national provider of information on the oil and gas industry. We limit our analysis to wells in unconventional reservoirs, as identified by the Pennsylvania Department of Environmental Protection (PADEP). This is consistent with our focus on innovation within unconventional shale development: according to industry engineers and geologists we consulted, the areas of current, active technological innovation that is relevant to fracturing operations are largely distinct between unconventional and conventional production. In other words, learning about production in conventional reservoirs does not transfer readily to provide insights into production in unconventional reservoirs. These experts did advise us that learning about fracturing vertical wells is transferable to fracturing in horizontal wells (and vice versa). Thus, we include both vertical and horizontal wells.

We identified 7,028 unconventional wells in Pennsylvania that reported initial production between January 2007 and December 2015. We collect data on inputs and operating parameters from two sources: Well Completion Reports and Stimulation Fluid Additive reports from PADEP, and the FracFocus database. Well Completion Reports, which operators must submit within 30 calendar days following completion, contain firm identifying information, well location, and information about the perforation and stimulation process. The requirement to submit Well Completion

Reports dates back to 1989.[12] Effective in February 2011, operators were also required to submit information on chemicals used in the stimulation process, including the name and concentration of each chemical additive in the fracturing fluid. Operators were instructed to submit information about chemicals either along with the Well Completion Report or on a separate DEP form, the Stimulation Fluid Additive report.

Some Pennsylvania operators elected to submit chemical additive information to the national FracFocus registry. Operators from other states (not just Pennsylvania) were uploading chemical additive information to FracFocus at the same time, typically on a voluntary basis.[13] FracFocus permitted operators to upload data in a standardized format, and the template contained the same information that operators had to report under Pennsylvania's 2011 disclosure law. Thus the operators who uploaded chemical additive information to FracFocus were in compliance with the reporting law.[14] In April 2012, Pennsylvania amended its reporting regulation to require operators to upload chemical additive information to FracFocus, in place of the requirement to submit information to the DEP on Well Completion Reports.

Thus, we obtained chemical additive information from both Well Completion Reports and associated Stimulation Fluid Additive reports, or from the FracFocus database. Notably, information that operators submitted to FracFocus was more readily observable by competitors and the public.[15] By contrast, Well Completion

[12] http://www.pacode.com/secure/data/025/chapter78/s78.122.html

[13] Operators in Wyoming and Arkansas were required to disclose chemical additive information as of September 2010 and January 2011, respectively; however, both states provided their own reporting websites for this purpose. Later in 2011 and 2012 several other states passed laws requiring chemical disclosure. Some of those laws required operators to upload reports to FracFocus, and others offered operators the option to use FracFocus or a state registry.

[14] There may have been confusion on this point at the time; some operators, evidently uncertain how DEP would enforce the requirement, provided printouts of their FracFocus disclosure forms along with their Well Completion Reports. Other operators submitted disclosures to FracFocus but not in the Well Completion Reports.

[15] When FracFocus was launched, GWPC and IOGCC provided fracturing fluid chemical reports

Reports and Stimulation Fluid Additive reports submitted to the DEP were available for review by a subscription service or by in-person review at regional DEP offices. Subscribers could view Well Completion Reports through the Exploration and Development Well Information Network (EDWIN), in which some chemical disclosure forms were available as scanned PDF documents.[16] However, there was a long wait time for reports to be scanned and uploaded to this system, especially in the height of the fracturing boom.[17] Members of the public could also review reports in person, but would have had to identify the permit number of a specific well, contact the appropriate regional DEP office, file a request, schedule an appointment to visit in person (typically three to four weeks in advance), and would then be allowed to review a limited number of hard copy documents onsite (on the order of 25 per day). Therefore we do not believe operators would have expected others to observe their fluid contents, prior to the April 2012 rule that required public disclosure of chemical information on FracFocus.

To capture productivity, we use the standard industry metric of initial gas output per foot of wellbore. Total output is highly correlated with initial output, and dividing by the length of the perforated interval normalizes output by well size to facilitate comparison across wells. In the process of creating this metric, we find that operators failed to provide the length of the perforated interval for 1,158 wells,

as individual PDF files that were easy to download individually but challenging to compile en masse. At the time, GWPC and IOGCC specifically stated their intent to provide a forum where the public could view individual reports but not look at many reports at once. At least two entities had successfully scraped the entire FracFocus database by late 2012. One was a for-profit consulting firm motivated by commercial interest in the database; the other was an environmental NGO that claimed, on right-to-know grounds, that the public should be able to review and compare information across wells (Skytruth, 2013).

[16] Until 2015, EDWIN was known as the Integrated Records and Information System (IRIS).

[17] One of the authors accessed the system several times in 2012-13 and found that the wait time for stimulation fluid additive information to be available in the system was highly variable, and the delay for its addition could be up to 18 months or more after well completion.

so we cannot use these in the analysis.[18] We drop an additional 94 wells that have nonsensical completion dates (the recorded completion date is after the date of initial production). Finally, we drop 21 wells with zero recorded gas production. This leaves us with a sample of 5,755 wells.[19]

For each of these wells we observe identifying information (operator, location, and completion date) and output. We collected data on inputs, including volume of water, quantity of proppant, and chemical additives to the stimulation fluid, from DEP reports and FracFocus. The DEP reports proved unamenable to optical character recognition scraping: they featured to a wide variety of formats (over ten different formats with different page headers), the use of numbers that were sometimes handwritten or crossed out and overwritten, overlaid date stamps, and raster images. Thus the relevant information was digitized manually, by a team of data entry contractors. This effort took about 1,800 person-hours over 4 months, and involved the entry of about 200 data items per report. We verified the quality of data entry by using standard procedures including systematic checks for consistency and reasonableness, spot checks comparing hand entries to the original reports, and comparison of duplicate entries by different contractors.

Despite our careful collection and entry of input data, not all variables are available for each well. Fluid volume is available for 4,353 wells (76%); proppant volume is available for 4,316 wells (75%); and both fluid and proppant quantity are available for 3,184 wells (55%). We have detailed chemical information for 4,290 wells, including 3,982 wells with chemical information disclosed to FracFocus and 308 wells (from the period February 2011 - April 2012) with chemical information disclosed only to PADEP.[20] Of the 1,465 wells with no information on chemicals, 1,250 were

[18] This omission of perforated interval appears to be out of line with basic reporting requirements.

[19] To ensure the analysis is not driven by outliers, we winsorize per-foot initial gas production at the 99th percentile.

[20] Mandatory disclosure started in Pennsylvania in 2011, but prior to this date some operators

not subject to disclosure; 215 wells appear to be out of compliance, as they were fractured after February 2011 but chemical additive information was not released to PADEP nor FracFocus.[21] We were able to obtain other information on other well parameters for a larger set of wells; we have information on the number of stages, for instance, for 4,699 wells (82%).

Table 3.1 provides a summary of information about the wells for which we observe chemical input information. Although we have other information for wells beginning in 2007, chemicals data are available only for wells starting in 2010. The table shows that the median and mean number of chemicals per well has increased slightly over time, fluid volume has increased somewhat more substantially, and gas production per foot has generally increased.

Table 3.1: Summary information for wells in sample with chemicals data

	Fluid Volume		# Chemicals / Well		First 6 Month Gas per foot	
Year	Median	Mean	Median	Mean	Mean	Median
2010	3.8	4.1	11.5	10.7	160.2	141.2
2011	4.4	4.5	9	10	154.7	122.8
2012	4.2	4.4	11	12.5	186.5	131.5
2013	5.6	5.9	13	15.4	198.1	153.9
2014	7.8	8.3	13	15.6	200.6	159.9
2015	8.9	9.1	13	14.3	151.3	128.1

Data sources are described in Section 3.3.
Chemicals are listed, non-proprietary chemicals with legitimate CAS numbers.
Fluid volumes are expressed in millions of gallons.
Gas volumes are expressed in thousands of cubic feet (MCF).

Table 3.2 provides additional information for wells in the sample for which we

voluntarily submitted chemical information for some wells.

[21] Another possibility is that operators provided chemical information for these wells but the reports had not been digitized and uploaded to EDWIN or its predecessor system by January 2017 when we searched the database for Well Completion Reports and Stimulation Fluid Additive reports. Assuming the reports were provided to DEP within 30 days of well completion, this would imply a delay of 4 to 5 years from when DEP received the reports to entry of the report into the system. While such a delay may seem unlikely, we did note that approximately 800 Well Completion Reports, some pertaining to wells completed in 2011 or early 2012, were uploaded into the system between December 2015 and January 2017.

observe chemical use. This table demonstrates that operators continued to innovate by introducing new chemicals each year, with the largest innovation (in terms of number of new chemicals) occurring in 2011-2012—right around the time that the state was issuing laws requiring mandatory disclosure. Companies also retired some chemicals from use (at least within our sample's timespan).

Table 3.2: Additional information for wells in sample with chemicals data

			Chemicals		
Year	# Wells	# Operators	# New	# Retiring	# Total
2010	38	8	51	3	51
2011	881	24	87	17	132
2012	1055	27	92	20	199
2013	961	31	62	66	249
2014	897	27	43	121	223
2015	148	17	4	–	112

Data sources are described in Section 3.3.
Chemicals are listed, non-proprietary chemicals with legitimate CAS numbers.
New chemicals refers to the count of unique chemicals appearing in our dataset for the first time in a given year.
Retiring chemicals refers to the count of unique chemicals appearing in our dataset for the last time in a given year.
Total chemicals refers to the count of unique chemicals used in a given year.

3.4 Analysis

In this section, we perform a series of empirical analyses designed to answer two questions. First, did mandatory disclosure of chemicals, even with provisions allowing for some chemicals to be declared proprietary, enable social learning. Second, did the learning enabled by disclosures have value, helping less productive operators to "catch up" to more productive operators. In future work, we plan to also analyze the extent to which this may have reduced operators' incentives to innovate.

For all of our analysis, we define two periods: pre- and post-disclosure. Thanks to Pennsylvania's regulatory history, the pre-disclosure period covers the 14 months from February 2011 to April 2012. During this period, we as econometricians are able to observe the chemicals used, but it is implausible that operators had access to this information (see Section 3.3). The post-disclosure period covers all of the wells from April 2012 through the end of our sample.

If the answer to both questions is affirmative, we would expect to see convergence in productivity across operators, with less able firms using information revealed by the disclosure laws to catch up with more successful firms. After confirming that this is the case, we examine each of the two questions more closely to rule out alternative explanations. To answer the first question, we examine how the design of fracturing fluids changed after disclosure. After finding a convergence in the chemicals used consistent with copying from disclosure, we turn to the second question. To answer it, we construct a two-stage estimator that combines information on chemical formulas and well productivity. We find evidence that convergence is indeed explained specifically by weaker firms adopting the chemical mixtures used by more successful firms in previously drilled wells. This finding is consistent with firms learning through chemical disclosure. Finally, we conduct several robustness and placebo tests to rule out alternative explanations.

3.4.1 Change in well productivity

As a starting point in our analysis, we consider the hypothesis that disclosure led to convergence in well productivity. To investigate this, we run a simple regression at the well-level:

$$g_w = \beta_0 + \beta_1 s_w + \beta_2 f_w + \epsilon_w, \tag{3.1}$$

where the w subscript indicates a well, g denotes per-foot first-six-month gas production, and s and f are fracturing proppant and fluid volumes. We then take the estimated residuals $\hat{\epsilon}_w$ and calculate a Gaussian kernel-weighted standard deviation over time. The results are shown in Figure 3.1. The curves plot the estimated standard deviation, for various bandwidth choices; the vertical black line marks April

2012, the date that disclosure came into effect. All three bandwidths show a marked fall in the standard deviation of the estimated residuals after disclosure. This fact motivates the rest of our analysis that follows.

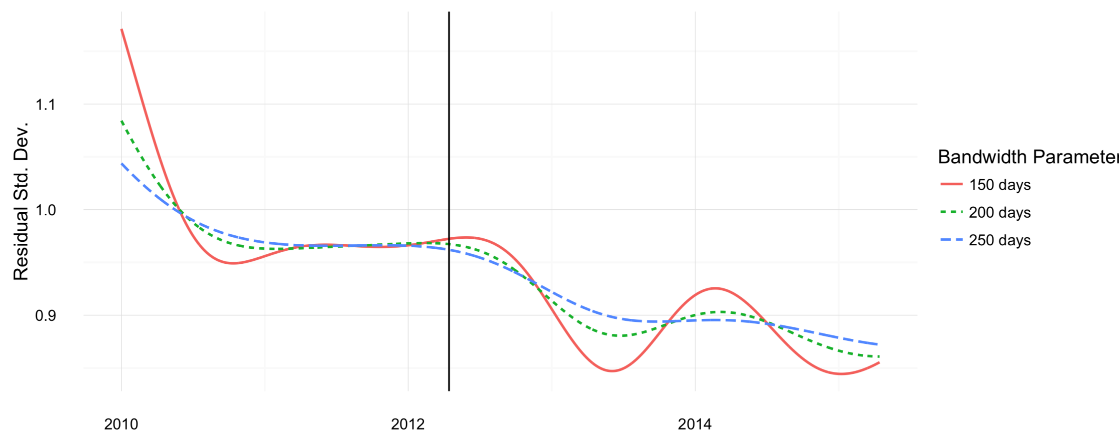

FIGURE 3.1: Kernel-weighted standard deviation of productivity residuals over time

The curves indicate the kernel-weighted estimated residual standard deviation over time, with different bandwidths. The vertical line indicates the date that public disclosure came into effect.

3.4.2 Well-to-well similarity measures

It is possible that the drop in residual standard deviation described in Figure 3.1 is simply the result of the shale gas industry maturing, with individual operators all converging towards a production frontier. One way to test the learning hypothesis would be to look for evidence that poor performers improve their productivity specifically when they copy the chemical mixtures of more successful firms.

We first consider the effect of Pennsylvania's disclosure rules on operators' chemical choices. If chemical disclosure compelled the release of valuable (formerly private) information, we might expect to find that operators' chemical recipes exhibit more similarity after disclosure than beforehand. Our detailed chemical input data (discussed in Section 3.3) allows us to explore this hypothesis by constructing measures of similarity for each well-to-well pair we observe.

We begin by introducing some notation, defining our similarity metric of choice, and providing a summary of the metric in our data.

Define $s_{ij} \in [0,1]$ as the pairwise similarity between wells i and j.[22] A value $s_{ij} = 1$ implies that the hydraulic fracturing fluids used in wells i and j are indistinguishable according to the chosen metric; $i = j \Rightarrow s_{ij} = 1$.

We now introduce some notation to define the similarity metrics. Let $\mathcal{C}$ denote the set of all possible chemicals, and let $x_{ic} \in [0,1]$ be the concentration of chemical $c \in \mathcal{C}$ in well i. So the sum of all such contributions $\sum_{\mathcal{C}} x_{ic} = 1$. We also define the binary variable $y_{ic} = \mathbb{1}\{x_{ic} > 0\}$. Then we can define the following quantities for a given (i,j) pair:

- $a_{ij} = \sum_{\mathcal{C}} y_{ic}$ is the number of chemicals found in well i
- $b_{ij} = \sum_{\mathcal{C}} y_{jc}$ is the number of chemicals found in well j
- $c_{ij} = \sum_{\mathcal{C}} y_{ic} y_{jc}$ is the number of chemicals found in both wells i and j

In principle we can calculate "abundance" or "correlation" similarity metrics that take into account the more detailed information in x_{ic}. Two features of our setting argue against taking this approach in practice. First, chemical quantities are reported as maxima rather than actual concentrations; second, FracFocus and the PADEP reports differ in terms of recording concentrations in terms of mass and volume. Due to these limitations, we focus on "binary" metrics that consider only the absence or presence y_{ic} of chemical c from well i. We proceed with the Sorensen binary metric, but find similar results with the alternative Jaccard binary metric.[23,24]

[22] Most metrics, including those we use, feature the property $s_{ij} \in [0,1]$, although other measures like correlation, with support $s_{ij} \in [-1,1]$, are possible.

[23] See chapter 12 of Krebs (2014).

[24] The Jaccard binary metric is defined as: $s_{ij} \equiv c_{ij}/(a_{ij} + b_{ij} - c_{ij})$.

The Sorensen binary metric is defined as:

$$s_{ij} \equiv \frac{2c_{ij}}{a_{ij} + b_{ij}}.$$

3.4.3 Disclosure and similarities

We calculate the similarity metric for each i, j well-pair in our data. Our chemical data on 4,015 wells gives us 8 million such pairs. Figure 3.2 shows the distribution of these measures for two sub-samples of pairs: those where both wells are drilled by the same operator, and those where the wells are drilled by different operators. The figure also plots the medians of the two distributions. A few observations are noteworthy: the different-operator distribution has a mass point at $s_{ij} = 0$, with most of the mass roughly centered around $s_{ij} = 0.25$. While the same-operator distribution also has most of its mass below $s_{ij} = 0.5$, it is noticeably more uniform and has a mass point at $s_{ij} = 1$. This confirms what we might have expected: intra-operator well-pairs tend to use more similar chemical mixes than inter-operator well-pairs.

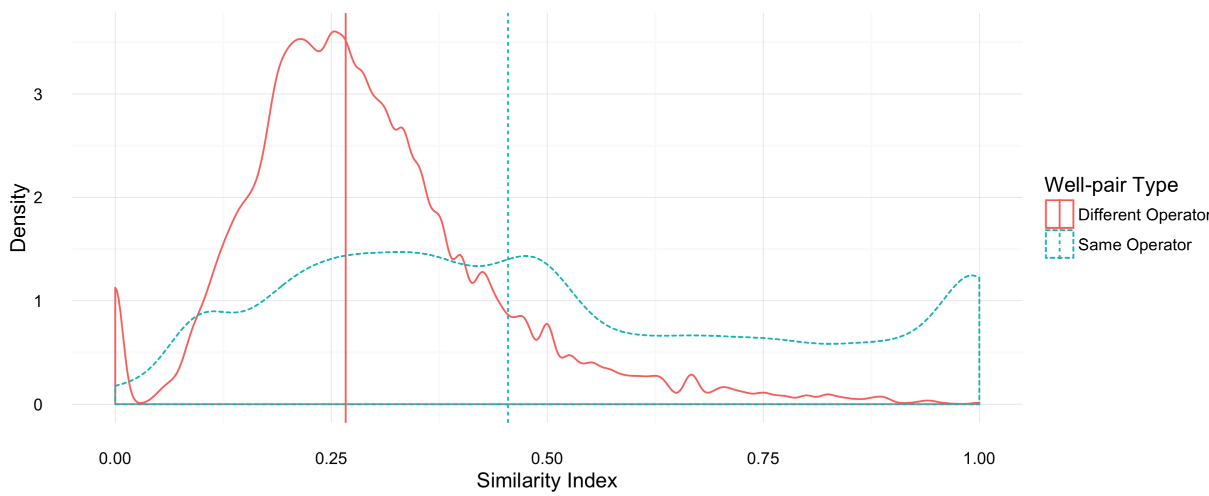

FIGURE 3.2: Sorensen similarities by same / different operator status

The curves plot estimated densities for well-pair similarities s_{ij}, conditional on whether the wells in the pair had the same or different operators. The vertical lines indicate the distribution medians.

We next consider how well-pair similarities have changed with the advent of disclosure. We do this by running a series of regressions of the following form:

$$s_{ij} = \beta_0 + \beta_1 oper_{ij} + \beta_2 post_{ij} + \beta_3 dist25_{ij} + \\ \beta_4 oper_{ij} * post_{ij} + \beta_5 dist25_{ij} * post_{ij} + \alpha_i + \alpha_j + \epsilon_{ij}, \quad (3.2)$$

where s_{ij} is the Sorensen similarity metric defined above, $oper_{ij}$ is a binary variable equal to 1 if wells i and j were drilled by different operators, $post_{ij}$ is a binary variable equal to 1 if either well i or j was fractured in the post-disclosure period, and $dist25_{ij}$ is a binary variable indicating whether the wells are located more than 25km apart. Operator fixed effects are included with α_i and α_j. The results of regressions featuring subsets of these variables are shown in Table 3.3.

A few things are worth noting from Table 3.3. First, all columns show a significantly negative estimate of the coefficient on $oper_{ij}$: this indicates that well-pairs with different operators tend to use less-similar chemical mixes, confirming the graph in Figure 3.2. Second, the coefficient of $post_{ij}$ is also always estimated to be significantly negative, indicating that well-pairs use less-similar chemical mixes on average if at least one of the wells is drilled in the post-disclosure period. Third, the coefficient on $dist25_{ij}$ is estimated to be negative, indicating that more geographically distant well-pairs use less-similar chemical mixes on average.[25]

[25] Similar results hold for other cut-off distances, or for letting distance enter the regression directly.

Table 3.3: Disclosure and well-pair similarities

	Similarity Index				
	(1)	(2)	(3)	(4)	(5)
Different Operator	−0.227***	−0.220***	−0.322***	−0.314***	−0.302***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Post	−0.024***	−0.023***	−0.117***	−0.115***	−0.123***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
Different Operator			0.101***	0.099***	0.086***
x Post			(0.001)	(0.001)	(0.001)
Distance > 25km		−0.014***		−0.014***	−0.035***
		(0.000)		(0.000)	(0.001)
Distance > 25km					0.022***
x Post					(0.001)
Operator FEs	✓	✓	✓	✓	✓
Observations	8,058,105	8,058,105	8,058,105	8,058,105	8,058,105
R^2	0.183	0.183	0.184	0.184	0.185

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.
The similarity index is the Sorensen index, described in detail in Section 3.4.2. Post is a binary variable, indicating whether *at least* one of the wells in the well-pair was fractured after disclosure.

The interaction terms are more interesting. The interaction between $dist25_{ij}$ and $post_{ij}$ is estimated to be positive and significant, indicating that after disclosure, distant wells use more similar chemical mixes than before disclosure. Similarly, the interaction between $oper_{ij}$ and $post_{ij}$ is estimated to be positive and statistically significant. This indicates that different-operator well-pairs use more similar chemical mixes post-disclosure than pre-disclosure. This is consistent with disclosure facilitating the transfer of knowledge about chemical mixes, and motivates our further analysis in the next section.

3.4.4 Convergence in output

First stage regressions

To evaluate the effect of changes in chemicals and other inputs on productivity, we set up a two-stage regression framework. In the first stage, we regress gas production on a set of observables and fixed effects. The fixed effects include operator by period fixed effects, which represent the relative productivity of each operator in each period, conditional on observables. The difference between the first-period and second-period fixed effects, in turn, represents an operator-level measure of change in average productivity from the first to the second period. In the second stage, we calculate the difference between first- and second-period fixed effects and regress this difference on the first-period fixed effect, along with a "quality-similarity measure that is explained in section 3.4.4 below. The second-stage regression is similar in spirit to simple tests of convergence in the economic growth literature (e.g., Bernard and Durlauf (1996)). As we shall see, our results suggest there is convergence among operators from the first to the second period, but this convergence is driven strongly by the fact that after disclosure, ex-ante lower-performing operators use chemical formulas that are more similar to those used by ex-ante higher-performing wells.

In our first stage regression, we regress per-foot initial gas production on a set of observables and fixed effects, using the form

$$y = F \times g(X)$$

where F is an operator fixed effect and $g(X)$ is a function of variables X that affect the well in question. Suppose $g(X) = e^{X'\beta}$. Taking logs on both sides, this implies that

$$\log y = \log F + X'\beta$$

In our preferred specification, the function $g(X)$ includes a township fixed effect, a year fixed effect, the density of unconventional wells previously completed in the township (as a quadratic), and the current output price.[26] The year fixed effect should absorb secular technological change, and the township fixed effect helps to control for spatial differences in resource quality.[27] We explored several alternative specifications, including specifications without price, without prior well density, with alternative measures of prior well density, and with well density (or alternative measures) in linear form rather than quadratic; all of these alternatives produced results similar to those reported here.

We estimate the first stage regression using the log of first 18 month gas output per foot for the dependent variable, and then exponentiate the fixed effects (to transform them from lnF back into F). We interpret these exponentiated effects as the

[26] For output prices, we use prices for the Henry Hub exchange for the month in which well i was completed. City gate prices for Marcellus production have often been lower than Henry Hub prices (http://www.eia.gov/todayinenergy/detail.php?id=24712), but we have not been able to obtain a continuous price series for Dominion South or other exchanges that may be more appropriate to the Marcellus. In any case the Henry Hub price is likely highly correlated with these other exchanges; thus, to the extent that the main point of our first stage analysis is to recover operator-period fixed effects, our use of the Henry Hub prices likely does not introduce significant error.

[27] Alternatively, we could perform a semiparametric regression that conditions flexibly over space. However, Covert (2015) implemented such a procedure in a related setting—shale oil production in North Dakota—and found the results were not substantially different from those using fixed effects for administrative jurisdictions.

contribution of the firm—a scale parameter, or multiplicative effect—that is positive by construction.

Table 3.4 shows the results of the first stage regression, with initial gas per foot as the dependent variable. Gas price appears to be positively correlated with initial productivity, consistent with other research showing a positive elasticity of production with respect to price (Newell et al., 2016).

Table 3.4: Results of first stage regression

	Log First 18 Months Gas Production
Gas Price	0.078***
	(0.017)
Prior Well Density	0.000
	(0.001)
Prior Well Density (squared)	−0.000
	(0.000)
Township FE	✓
Year FE	✓
Operator-period FE	✓
Observations	5,755
R^2	0.988

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.
This regression omits the constant in order to calculate a fixed effect for each operator-period; therefore, the R^2 value should not be given the usual interpretation.

Quality-Similarity metric

To investigate the relationship between changes in productivity and the convergence in chemical mixes across different-operator well-pairs shown in Table 3.3, we define a quality-similarity index, QS_f for each firm f. It is designed to capture the quality of an operator's chemical matching, where quality is defined by estimated pre-disclosure fixed effects. It is constructed as follows:

$$QS_f \equiv \frac{1}{|M_f|} \sum_{m \in M_f} s_m(\theta_m^{PRE} + \hat{\epsilon}_m)$$

$$M_f \equiv \{i, j \mid i \text{ is } f \text{ POST well, } j \text{ is non-}f \text{ POST well, } j \text{ fractured before } i\}$$

$$s_m \equiv \text{similarity index for well-pair } m$$

$$\theta_m^{PRE} = \text{estimated PRE period fixed effect for well } j\text{'s firm.}$$

$$\hat{\epsilon}_m = \text{(exponentiated) first-stage residual for well } j.$$

QS_f is thus a quality-weighted measure of the post-period similarity between f's wells and other firms' wells. We require that well j was stimulated before well i to ensure the possibility that firm f would have had the opportunity to view the chemicals used in well j and adjust its mix for well i if it chose.

Quality-Similarity difference-in-differences test

With calculated values of QS, we proceed to the second stage regression:

$$\theta_f^{POST} - \theta_f^{PRE} = \beta_0 + \beta_1\theta_f^{PRE} + \beta_2 QS_f + \beta_3 QS_f\theta_f^{PRE} + \epsilon_f. \tag{3.3}$$

Because the "data" for this regression features estimation error from the first stage, we calculate standard errors via a bootstrap. We use a modified bootstrap to ensure that we do not lose power in the second stage, which takes place at the operator-period level. To get our bootstrap sample, we first calculate the number of wells fractured within each operator-period. Then we draw, with replacement, that number of observations from that operator-period. This method, which is in the spirit of a block or panel bootstrap (Cameron and Trivedi, 2005), ensures that we do not lose any second stage observations due to the random draws inherent in the bootstrap routine. Unless noted otherwise, all second-stage results presented below use this bootstrap routine with 200 bootstrap replications.

The results can be seen in Table 3.5. The first column shows the regression when only θ_f^{PRE} is included; the coefficient is found to be significantly negative, which suggests those with lower values of θ_f in the first period experience greater growth in their θ_f with the advent of disclosure – i.e., there is convergence in the values of θ_f over time. The second column shows estimates from Equation (3.3), with the addition of QS_f and the interaction term. With these additions, the point estimate of β_1 has flipped to be positive, suggesting divergence in θ_f over time. The point estimate of β_2 is also positive and statistically insignificant, suggesting that a higher QS_f is associated with higher growth in θ_f. Finally, β_3 is estimated to be significantly negative. This suggests that when well similarities are accounted for, the convergence of θ_f over time is due not to θ_f^{PRE}, but rather to the interaction term.

Table 3.5: Second stage results

	$\theta^{POST} - \theta^{PRE}$	
	(1)	(2)
θ^{PRE}	-0.289	2.582**
	(0.243)	(1.018)
QS		2.335
		(2.597)
θ^{PRE} x QS		-4.501**
		(2.177)
Constant	0.881***	-1.473
	(0.237)	(1.431)
Observations	28	24
R^2	0.002	0.346

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.
Standard errors are calculated using 200 bootstrapped samples, collected with binning at the operator-period level. See text for details.

Selection and exiting firms

The results reported above support the notion that convergence occurred over the period we study, but that convergence is driven primarily by the interaction of θ_f^{PRE} with QS_f. However, a substantial number of operators exited Pennsylvania in the POST period (i.e., did not drill any new wellbores). We must therefore consider the extent to which selection affects the results. Would we find a different relationship if these firms had continued to operate in the Pennsylvania shale fields?

To address this concern, we assign values of θ_f^{POST} and QS_f to the exiting firms, re-run the second-stage regression in Equation (3.3), and check for stability of results. We use several alternative assumptions in an effort to test for robustness under a range of plausible but conservative scenarios.[28]

We consider the following four alternative scenarios:

1. Assume exiting firms would have performed at the 25th percentile of the firms that did continue to operate.

2. Assume exiting firms would have performed at the median of the firms that did continue to operate.

3. Assume exiting firms in the bottom quartile of the distribution in the pre period would have performed at the bottom of the distribution in the post period; assume firms in the top three quartiles in the PRE period would have performed at the top of the distribution in the POST period.

4. Assume exiting firms in the bottom half of the distribution in the PRE period would have performed at the bottom of the distribution in the POST period;

[28] At first glance, it might appear that an assumption that all exiting firms would have performed at the top of the distribution in the POST period would provide a conservative test. However, this turns out not to be the case. Instead, such an assumption implies that a number of poor performers (in the PRE period) all improved their performance in the POST period, which would lend support to the convergence hypothesis.

assume firms in the top half of the distribution in the PRE period would have performed at the mean of the distribution in the POST period.

We must also assign values for QS_f for the firms that exited. We make three alternative assumptions:

1. Assume exiting firms would have had a QS_f at the 25th percentile of the firms that remained.
2. Assume exiting firms would have had a QS_f at the median of the firms that remained.
3. Assume exiting firms would have had a QS_f at the 75th percentile of the firms that remained.

It is not clear *a priori* whether a relatively high or low assignment of QS_f should make for a conservative test of its role in achieving convergence, so using a range of alternative assumptions allows us to test across various alternative states of the world. We run distinct regressions that match each of the four assumptions for θ_f^{POST} against each of the three assumptions for QS_f (plus an additional set of regressions that considers a simple convergence without QS_f). Table 3.6 provides a summary of the results for these tests. Note that the table contains the results of 16 regressions; each column and set of rows (between the horizontal lines) represents a distinct regression.

The table shows that this analysis is consistent with the main analysis, implying that selection on surviving firms is not driving our results. This supports our earlier finding that convergence is driven primarily through the interaction of the QS_f term and the θ^{PRE} term, which is consistent with disclosure-enabled "catching up".

Table 3.6: Results of QS-DD tests with alternative assumptions for exiting firms

Assumption for θ^{POST}		No QS (simple convergence)	25th Percentile QS	Median QS	75th Percentile QS
25th Percentile	θ^{PRE}	-0.254^{**}	1.522	3.397^{***}	4.351^{***}
		(0.119)	(1.231)	(1.135)	(0.812)
	QS		2.955	5.23^{**}	6.014^{***}
			(2.188)	(2.014)	(1.460)
	$QS \times \theta^{PRE}$		-3.160	-6.462^{***}	-7.944^{***}
			(2.243)	(2.023)	(1.402)
Median	θ^{PRE}	-0.287^{**}	1.409	3.177^{***}	4.136^{***}
		(0.112)	(1.119)	(1.022)	(0.727)
	QS		1.961	4.584^{**}	6.082^{***}
			(1.989)	(1.814)	(1.306)
	$QS \times \theta^{PRE}$		-3.014	-6.124^{***}	-7.610^{***}
			(2.040)	(1.822)	(1.255)
Bottom Or Top	θ^{PRE}	-0.178	2.950	0.591	-1.613
		(0.340)	(3.367)	(3.609)	(3.213)
	QS		-2.959	-3.700	-2.654
			(5.982)	(6.407)	(5.771)
	$QS \times \theta^{PRE}$		-4.966	-0.621	3.310
			(6.135)	(6.436)	(5.546)
Bottom Or Mean	θ^{PRE}	-0.069	2.014^{*}	3.006^{***}	3.188^{***}
		(0.110)	(1.062)	(1.027)	(0.844)
	QS		4.151^{**}	5.087^{***}	4.562^{***}
			(1.887)	(1.823)	(1.516)
	$QS \times \theta^{PRE}$		-3.704^{*}	-5.407^{***}	-5.577^{***}
			(1.935)	(1.831)	(1.456)
N		44	44	44	44
N (exit)		20	20	20	20

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.
All QS is standard Sorensen. All regressions include a constant term. Growth rate is $\theta^{POST} - \theta^{PRE}$. Significance tests and standard errors do not reflect the bootstrap procedure.

3.4.5 Placebo tests

In addition to our attempt to combat potential sample selection in the prior section, we recreate this analysis with two alternative definitions of QS_f to check the robustness of our findings. In the first alternative, QS_f is generated using all post well-pairs with firm f's wells, regardless of the chronological order of the stimulation. In the second alternative, we create a version that includes a geographic limitation in order to rule out spatial spillovers.

In the first alternative, the estimated coefficient on the interaction of θ^{PRE} and QS_f is not significant, which is in contrast to our finding with the proper definition of QS_f. This contrast supports the hypothesis that disclosure in fact aided learning. The alternative with a geographic restriction has similar results to the original, suggesting that convergence is not due to spatial proximity.

"Time-Fluid QS"

For the first alternative, we define a new quality-similarity index QS_f^{TF}:

$$QS_f^{TF} \equiv \frac{1}{|M_f^{TF}|} \sum_{m \in M_f^{TF}} s_m \theta_m^{PRE}$$

$$M_f^{TF} \equiv \{i, j \mid i \text{ is } f \text{ POST well, } j \text{ is non-}f \text{ POST well}\}$$

Note that s_m and θ_m^{PRE} are defined as before, so the only change is the set of well-pairs used, M_f^{TF}. In the original definition, we imposed the requirement that f's well was completed after the other operator's well, to allow for the possibility that f was able to view the chemical disclosures of the non-f well. In this alternative, we do away with this limitation and construct QS_f^{TF} as though we had no information on the relative timing of the wells. If the learning through disclosure mechanism is operative, we expect that the analysis with the new measure QS_f^{TF} will result in less significantly negative estimate of β_3.

The results of the new second stage regression are shown in column 3 of Table 3.7. While the estimates are significantly different from zero, the coefficient on the interaction term is positive, suggesting that the QS^{TF} is contributing to divergence. This result supports the conclusion that convergence is driven by learning through disclosure: including matches with wells drilled after f's wells in M_f^{TF} weakens the explanatory power of the regression.

Table 3.7: Second stage results with alternative QS definitions

	$\theta^{POST} - \theta^{PRE}$			
	No QS	QS	Time-Fluid QS	Distance QS
	(1)	(2)	(3)	(4)
θ^{PRE}	-0.289	2.582^{**}	-11.137^{***}	2.775^{***}
	(0.243)	(1.018)	(1.013)	(0.934)
QS		2.335	-34.060^{***}	2.495
		(2.597)	(2.483)	(2.381)
θ^{PRE} x QS		-4.501^{**}	23.032^{***}	-4.974^{**}
		(2.177)	(2.280)	(2.100)
Constant	0.881^{***}	-1.473	16.505^{***}	-1.505
	(0.237)	(1.431)	(1.333)	(1.405)
Observations	28	24	28	24
R^2	0.002	0.346	0.318	0.461

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.
Standard errors are calculated using 200 bootstrapped samples, collected with binning at the operator-period level. See text for details.

Distance buffer

The second alternative is designed to rule out an alternative hypothesis: that well-pairs are becoming more similar due to learning spillovers that are limited to, or primarily driven by, spatial proximity (Conley and Udry, 2010). To test this possibility, we construct a final quality-similarity index QS_f^{DIST}:

$$QS_f^{DIST} \equiv \frac{1}{|M_f^{DIST}|} \sum_{m \in M_f^{DIST}} s_m \theta_m^{PRE}$$

$$M_f^{DIST} \equiv \{i, j \mid i \text{ is } f \text{ POST well, } j \text{ is non-} f \text{ POST well, } i \text{ and } j \geqslant 25\text{km apart}\}$$

In this alternative, M_f^{DIST} is defined to include only more distant well-pairs. The second stage results are in the final column of Table 3.7. The point estimates are almost identical to those with the original QS_f measure, and the interaction term is still significantly negative. The similarity of the estimates using QS_f^{DIST} and QS_f rules out spatial spillovers as the source of learning, and thereby lends support to learning through disclosures.

3.4.6 Mechanism: Learning through contractors

One possible channel for the transfer of information about chemical mixtures is the contractors who are hired to hydraulically fracture the wells. Contractors perform a variety of roles; operators hire them to assist with jobs that may include drilling, cementing, well logging operations, and other tasks, as well as designing and conducting the fracturing job. The PADEP requires operators to provide information about contractors on Well Completion Reports. According to these reports, operators hired 8 contractors on the median well in our sample to assist in various roles; for some wells, operators hired up to 40 contractors. However, the roles these contractors play are not always specified. We were able to identify the fracturing contractor for 3,153 wells. This includes 2,558 wells for which the contractor job was specified as involving stimulation or fracturing (distinct from related tasks such as perforation), and 595 wells in which an operator hired a contractor that exclusively performs such services.

To test whether contractors facilitate the transfer of information about chemical mixtures, and whether that role was changed by the institution of disclosure rules, we perform a series of regressions of well-pair similarity indices on dummies for: (i) if the two wells share the same contractor and (ii) if at least one of the wells is in the post-disclosure period. We consider only inter-operator well-pairs.

$$s_{ij} = \beta_0 + \beta_1 contractor_{ij} + \beta_2 post_{ij} + \beta_3 contractor_{ij} * post_{ij} + \epsilon_{ij} \tag{3.4}$$

For all of these regressions, we restrict ourselves to the sample of well-pairs with data on the fracturing contractor *for both wells*(see above). The results of these regressions are shown in Table 3.8.

When it is the only non-constant regressor, the dummy for sharing a contractor is associated with an increase in the similarity index of 0.272. In the second column, the dummy for post-disclosure is included: this dummy leaves the contractor coefficient statistically unchanged, enters negatively and significantly, but is smaller in magnitude. The third column adds an interaction between these two dummies. The coefficient for same contractor increases, the coefficient on POST becomes even smaller, and the interaction term is estimated to be negative and significant. The role of the contractor in facilitating a similar chemical mix is reduced by about a third in the post-disclosure period, suggesting that operators are not as reliant on contractors as a source of information when that information is being published. These results suggest that the contractor channel is associated with more similar wells, but that this channel is less effective in the post-disclosure period.

3.5 Conclusions

Although often motivated primarily by a feeling that the public has the "right to know" about risks that arise from the storage, use, and disposal of toxic chemicals, disclosure laws have been found to affect firms' behavior in other ways. However, prior literature in economics and policy has focused on issues such as how disclosure laws induce voluntary self-regulation, evidently motivated by effects on external

Table 3.8: Determinants of well-pair similarities

	Similarity Index		
	(1)	(2)	(3)
Same Contractor	0.272***	0.271***	0.394***
	(0.005)	(0.005)	(0.020)
Post		−0.077***	−0.057***
		(0.007)	(0.006)
Contractor x Post			−0.126***
			(0.020)
Constant	0.289***	0.366***	0.345***
	(0.002)	(0.007)	(0.006)
Observations	3,544,453	3,544,453	3,544,453
R^2	0.216	0.219	0.220

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.
Results are from the regression described in (3.4); similarities are from the subsample where the fracturing contractor can be identified with confidence. See Section 3.4.6 for details.
Standard errors are clustered at the well-index level, i and j of s_{ij}.

stakeholders such as product consumers, employees, or the general public. We consider a different question—whether these laws can also create pathways for knowledge transmission that were previously inaccessible or overly costly—and assemble detailed data on both the inputs and outputs of shale gas firms to investigate this phenomenon.

We find evidence for convergence in outputs, convergence in inputs, and a link between the two—in sum, suggesting that disclosure laws affecting oil and gas production firms in Pennsylvania created opportunities for social learning that "follower" companies took advantage of, in ways that allowed them to catch up with "leader" companies. This also supports the notion that companies concerned about erosion of competitive advantage due to public disclosure laws may have a valid argument. Nonetheless, whether such disclosure laws reduce public welfare remains an open question. If disclosure undermines secrecy in a way that reduces innovative activity, this could harm welfare, but firms may also experience gains from trade that may outweigh such costs if they exist. In ongoing research, we hope to better elucidate the tradeoffs with respect to overall productivity and investment in innovation.

4

Energy Transitions and Technology Change: "Leapfrogging" Revisited

4.1 Introduction

The recent rapid growth in many emerging economies comes with opportunities for better living conditions, increased investments in human capital, self-sustaining economic growth, and greater global prosperity. Energy consumption is foundational to this economic growth, yet over a billion people still lack access to electricity, and many more are served by underpowered or unreliable energy systems. Over the coming decades, energy systems for both generation and consumption will be built to accommodate the needs of these populations, even as they continue to grow. Indeed, Wolfram et al. (2012) forecast that nearly all of the growth in energy demand in coming decades will arise from the developing world. Meanwhile, technological innovation creates opportunities for "leapfrogging" in energy systems—that is, the use of modern technologies in emerging economies that were not available to today's industrialized countries at the comparable time of their development. Leapfrogging opportunities may exist in interrelated areas that affect the efficiency of energy transformation, the carbon intensity of energy generation, and the energy intensity of economic growth (van Benthem, 2015).

Researchers have pointed out actual or potential technology leapfrogging in various domains (e.g. Goldemberg, 1998; Smil, 2010; Amankwah-Amoah, 2014), with the most notable example being cellular telephone networks that allow developing countries to skip over fixed-line technology. In the energy domain, researchers have documented leapfrogging in the adoption of solar electricity generation technologies in rural areas, ethanol production in Brazil, and biomass cookstoves in China (Goldemberg, 1998), adoption of energy-efficient appliances and fuel-efficient vehicles (van Benthem, 2015), and elsewhere.

At the same time, institutions may hinder the ability of countries to take advantage of these opportunities. Large entrenched industries, often with strong political and economic ties to central governments, may take strategic actions to block or slow the growth of potential competitors (Pearson, 2014). Government planners who wish to utilize the best modern technology may withdraw their support for incremental steps. Unreliable energy grids may create incentives for private firms or capital-rich house-holds to build redundant energy generation systems that, in turn, operate less efficiently due to lower economies of scale. Studies of technological change, including numerous examples from the energy sector, suggest widespread change often takes many decades (Rosenberg, 1972; David, 1990; Grübler et al., 1999; Hall, 2004; Smil, 2010).

This paper contributes to the understanding of technological evolution in the energy sector by exploring the extent to which the energy intensity of economic growth has changed across countries over many decades. I test for evidence of "leapfrogging" in the overall amount of energy used and the intensity of energy used to generate a given change in economic output. This first step builds on recent research by van Benthem (2015); I expand on that analysis by incorporating a longer time series of energy consumption for industrialized countries and by including energy used in energy conversion and distribution, rather than just consumption by end users. Next, I test for evidence of leapfrogging in the carbon intensity of economic output and growth. Finally, I explore the heterogeneity in relationships between economic output, energy use, and pollution, identifying institutional, financial and environmental factors that correlate with different countries' experiences with energy leapfrogging.

4.2 Energy consumption, economic growth, and technology

Research on economic and energy history shows that energy intensity of economic growth varies with the income level (van Benthem, 2015). At the lowest levels of economic development, the energy intensity of economic growth is relatively low, while in the "takeoff" or middle-income period energy intensity increases. In later phases of development, especially as economies transition to more service-based sectors, the energy intensity of economic growth tends to decline once again.

4.2.1 Empirical analysis: Energy and development

A large literature in economics explores cross-sectional differences in energy use and energy intensity of economic growth.[1] Several authors have used panel econometric methods to study these phenomena (e.g., Medlock and Soligo, 2001; Judson et al., 1999; Galli, 1998; van Benthem and Romani, 2009). In general these studies find strong evidence of non-linearity in the relationship between income and energy demand, reflecting the effects of structural economic changes and adoption of new technologies during the course of development. For instance, Galli (1998), studying ten Asian developing countries, finds evidence of declining energy intensity as nations become wealthier. Judson et al. (1999) analyze the relation of growth to energy

[1] Interest is hardly confined to economics; many historians, for instance, have reflected on decadal energy transitions in regions or individual countries.

consumption in different sectors for a panel of 123 countries, and find different patterns for household use (increasing over most of the income range), transportation (increasing over all of the income range), and industrial sectors (an inverted U shape over the income range).

A recent panel analysis (van Benthem, 2015) tests for energy leapfrogging in both levels of energy use for a given level of income, and intensity of economic growth. This analysis takes advantage of a proprietary dataset with information on energy consumption and GDP distinguished into eight sectors, spanning 76 countries with up to 46 years for some countries (1960-2006). The author finds no evidence that today's developing countries experience economic growth with lower energy use or energy intensity than did developing countries in the past. He does, however, find evidence of household-level adoption of energy-efficient appliances and vehicles. He reconciles these observations by noting the combination of industrial outsourcing to less-developed countries, and more energy-intensive consumption by households and public institutions at an earlier stage of development—along the lines of a technology rebound effect. However, the detailed data are available only for a relatively small set of industrialized countries (twelve). This is because IEA data on energy consumption start in 1960, and by this time many industrialized countries had per capita income greater than $10,000, which is too high to be relevant for the tests used (i.e., because the tests compare today's industrialized countries at a time when they had per capita income less than that amount). Thus, one drawback to this analysis is that the conclusions may be due partly to peculiar characteristics of the small set of industrialized countries available for comparison.

4.2.2 Institutions, policy, and technological change

Widespread technological change, especially when it involves large capital investments, can be a slow and gradual process, particularly due to (1) the co-evolution of long-lived technological systems or clusters; (2) dynamic competition between technologies rather than a smooth progression from identifiably "old" to identifiably "new" technologies; and (3) nonlinear patterns of technology adoption with respect to income. These factors suggest reason to be cautious in interpreting the results of panel studies depending on the timeline used for the analysis.

National energy consumption arises from the net effect of countless individual and institutional actors representing households, firms, and governments. Some of the energy-using technologies in which these actors invest have relatively short useful lives, on the order of a few years, while other technologies are useful for several decades. For instance, household appliances may last as little as three to five years; motor vehicles on the order of ten to fifteen years. Industrial capital assets frequently produce value for thirty years or more.[2] Grübler et al. (1999) document diffusion processes for about 25 energy-using or energy generation technologies and note a range of diffusion rates, depending on the relative advantage of the new technology, the degree of interdependence with other technologies, and the degree to which the new technology is complementary within existing technology ecosystems or requires the development of new infrastructure. David (1990) echoes the theme of the slow evolution of technological and economic ecosystems in documenting the slow diffusion of electric dynamo technology. Given that energy infrastructure frequently requires large capital investments, has a long useful life, and is integrally embedded within technological ecosystems, longer-duration studies of the relationship of energy consumption and economic growth may reveal insights not otherwise available.

Furthermore, when technological change involves the displacement of one technology by another, that displacement need not occur quickly or in a uniform direction. For instance, writing about the transition from sail-powered ships to steamships,

[2] Institutions and policies may have much longer-lasting consequences, such as urban development that is oriented around mass transit versus private vehicles.

which involved the use of auxiliary sails on steamships and auxiliary engines on sailing ships, Rosenberg (1972) cautions against the interpretation of technological change as "historical foreshortening" and notes that a careful analysis would describe invention and technological change as "a gradual process of accretion" and "a cumulation of minor improvements." This is especially true when users and manufacturers of older technologies are able to improve the efficiency of the old technology in the face of competition from substitutes—the so-called "last gasp" phenomenon—but can also arise from setbacks in the development of newer technologies, or from political-economic interactions in which bureaucracies or other institutions grow up around certain industries and then act, strategically or otherwise, to slow the innovation or adoption of new technologies (e.g., see Pearson, 2014).

The nonlinear relationship between household income and adoption of energy technologies, well documented in Wolfram et al. (2012) and Gertler et al. (2016), also suggests advantages to longer-run analysis, since relationships based on medium-run trends may not hold at longer time scales. Taken together, these three aspects of technological change imply that long-run analyses of technological change may reveal insights that medium-run analysis may not.

4.3 Data

I assemble a panel dataset of energy consumption, energy-related carbon emissions, population, GDP, and an index of real oil prices for 124 countries, with coverage for as long as 153 years for some countries. The next two subsections describe the sources for this dataset, and Section 4.3.3 provides a summary of the coverage of this panel.

4.3.1 Energy consumption and emissions

I use two measures of energy consumption. The first is total final consumption (sometimes abbreviated TFC), which represents energy consumption by end users, including households and industries. The second is TFC plus energy used in transformation from primary to secondary sources (e.g., the transformation of coal or natural gas into electricity), energy used by energy-producing industries (e.g., to produce coal from source rock), and losses during energy transmission and distribution. I call this total energy use (TEU).[3] In 2013, the components of TEU outside of TFC—energy used to extract, convert, and distribute energy—accounted for one-third of global consumption (IEA, 2015a).

My data on energy consumption come from several sources. The bulk of the data are from the International Energy Agency (IEA), which draws on national statistical organizations, international agencies, and other sources to compile energy consumption (and other) statistics. The IEA Extended Energy Balances series (IEA, 2015a) provides information for countries on the supply and consumption of energy, tabulated over about 60 generation sources ("products" in the IEA jargon) and 90 use categories ("flows").

The earliest energy consumption data from IEA (2015a) are available for countries in the Organization for Economic Cooperation and Development (OECD) from as early as 1960, and non-OECD countries from as early as 1971. However, as of 1960 many industrialized countries had per capita income greater than 10,000 USD. Thus, a data series using only IEA data may not provide an adequate number of points of comparison. To address this concern, I incorporate data from Unger and Thistle (2013) for Canada, and Kander et al. (2014) for eight European countries

[3] My TEU measure differs from the International Energy Agency's "Total Primary Energy Supply" in that I exclude non-energy uses of fossil fuels, such as the production of petrochemicals, lubricants or tar.

(France, Germany, Italy, the Netherlands, Portugal, Spain, Sweden, and the United Kingdom). The start years for these countries vary, but in all cases annual energy consumption is available by 1861. In addition, I incorporate data on historical energy consumption in the USA starting in 1900 (Schurr et al., 1960; EIA, 2016). Appendix A provides additional details about combining the energy consumption data series.

Table 4.1 provides summary statistics for average energy use over time for different groups of countries. As expected, both energy consumption and carbon emissions are higher for countries with higher levels of income.[4] However, Table 4.1 also provides suggestive evidence that other factors besides income affect energy consumption and emissions. For example, comparing both 2013 TEU per capita, and 2013 TFC per capita, between OECD and non-OECD countries at income levels in the $20,000-$30,000 and over $30,000 ranges shows that non-OECD countries have substantially greater total energy use within the same income range. This suggests institutions may moderate the relationship between economic activity and energy use, perhaps by influencing the abilities of governments, firms or households to adopt energy-efficient technologies. A similar relationship holds for carbon emissions between these comparison groups. For lower levels of income (e.g., $10,000-$20,000), OECD countries use greater amounts of energy and emit more carbon from energy uses.

Studies of economic growth and energy use show that predominantly agricultural societies use relatively low levels of energy. With industrialization and increased urban settlement, energy demand increases, then levels off as the economic base becomes more organized around services, manufacturing is often outsourced to other countries, and more efficient energy technologies are deployed, either due to technological advance or economies of scale (Smil, 2010). Figure 4.1 shows, for a few countries, the resulting S-shaped relationship between energy consumption and eco-

[4] The relatively high figure for TEU per capita in 1960 for "OECD countries with income between $10,000 and $20,000" is attributable to the fact that data are available for only two countries in this category in 1960 (Poland and Turkey), so the average is more sensitive to outliers.

Table 4.1: Summary statistics for energy use and carbon emissions

Measure (per capita)	Year	OECD 10-20k	OECD 20-30k	OECD >30k	non-OECD <10k	non-OECD 10-20k	non-OECD 20-30k	non-OECD >30k
TFC	1960	33.3	24.9	88.8				
(GJ)	1971	41	63.9	133.7	14.4	26.4	25.8	76.3
	2013	65.2	84.9	158.7	21.8	51.7	142.1	196.2
TEU	1960	45.7	29.5	116.8				
(GJ)	1971	54	84.8	165.9	17.1	35.4	46.6	152.5
	2013	98.8	119.9	218.5	29.3	69.7	157.2	335.5
CO_2	1960	3.5	1.9	9.7				
(metric	1971	3.9	5.8	11.6	0.4	1.9	2.9	8.3
tons)	2013	6.3	6.6	9.4	1.4	3.9	10.1	19.3

Dataset includes 34 OECD and 90 non-OECD countries.
Income levels are per capita in 2013, measured in 2005 USD at purchasing power parity. Data are available for most OECD countries from 1960, and for most non-OECD countries from 1971.
Source: IEA (2015a, 2015b).

nomic growth. If technology leapfrogging does result in lower energy consumption for a given quantity of economic growth, the S-shaped curves shown in Figure 4.1 should be flatter for countries that develop later in time.[5]

To the extent that energy from low-carbon sources supplies a relatively small part of most countries' energy needs, the relationship between carbon emissions and economic growth follows a similar pattern. This observation runs counter to earlier theoretical predictions and some empirical findings that economic growth eventually leads to lower levels of emissions (e.g., the "Environmental Kuznets Curve" literature). With increased development of renewable energy along with other non-fossil sources such as nuclear fission, some countries may eventually be able to achieve economic growth without increased carbon emissions. Similar to the comment above, to the extent that countries can take advantage of technology leapfrogging opportunities

[5] The relationship between TEU and economic growth, though not shown, looks substantially similar for this subset of countries.

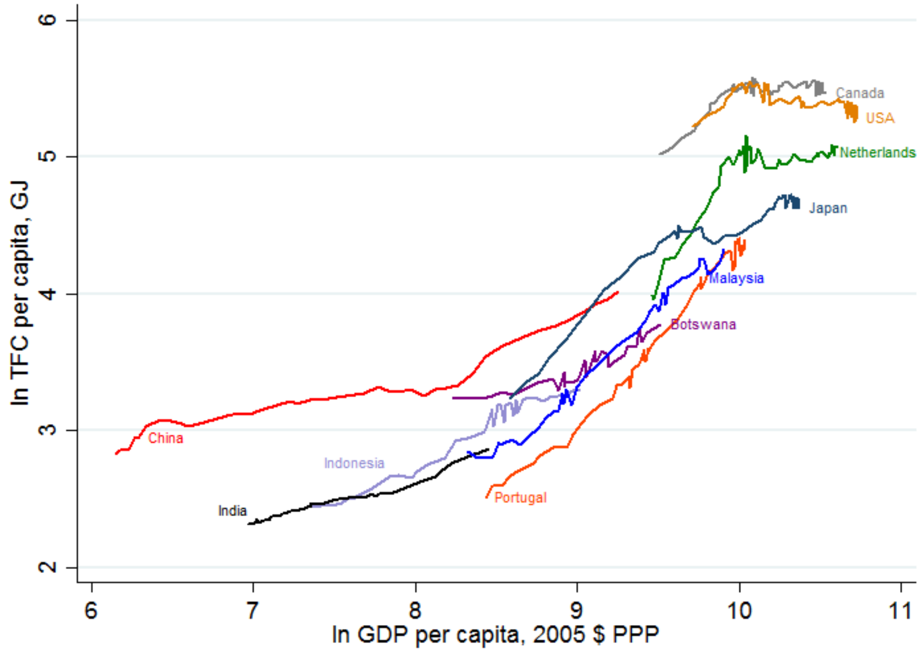

FIGURE 4.1: Energy transitions for select countries, 1960-2013

in energy efficiency or in low-carbon generation, curves illustrating this relationship would be flatter for countries that develop in later periods.

4.3.2 Prices, GDP, and emissions

Data on energy-related carbon dioxide (CO2) emissions come from IEA (2015b), which provides that information for the same years as the energy consumption metrics.[6] Data on population and Gross Domestic Product (GDP) are from the World Bank's World Development Indicators, supplemented by data from the International Monetary Fund's International Financial Statistics and from Maddison (2010). The latter source provides all of the GDP and population data prior to 1960, which is the start year for data from the World Development Indicators. I use the purchasing power parity (PPP) measure of GDP, which facilitates cross-country comparisons by

[6] In a future draft, I plan to calculate carbon emissions for the pre-1960 energy consumption series using the same methodologies as in IEA (2015b).

accounting for differences in relative price levels.

Following van Benthem (2015), I construct a real oil price index using oil price data from the BP Statistical Review of World Energy, converted into country-specific indexes using inflation and exchange rate data from the World Development Indicators (from 1960 onward) and Bordo (2015) (1861-1959 for countries with energy consumption data over that period). The oil price index provides a country-specific measure of price variability over time, compared to the world oil price from the BP Statistical Review.[7]

4.3.3 Summary of panel data

The full panel consists of energy consumption, population, GDP, and real price data for 124 countries, starting as far back as 1861 and extending to 2013. This includes 43 countries classified by the World Bank as high-income (per capita income greater than 12,745 USD in 2013), 36 as upper-middle income (income between 4,126 and 12,745 USD), 31 as lower-middle income (income from 1,046 to 4,125 USD), and 14 as low-income (income below 1,046 USD). The panel begins in 1861 for Canada and eight European countries, in 1900 for the USA, in 1960 for most remaining OECD countries, and in 1971 for most of the non-OECD countries. For a few countries, coverage begins in a different year; the main exceptions are former Soviet republics and certain Eastern European countries, for which coverage begins in 1990.

For the initial exercise of comparing my results to those of van Benthem (2015), I use that paper's definition of developing and industrialized countries. In this framework, "developing" countries are those with per capita income of 10,000 USD or below in 2013 (in 2005 USD PPP), and "industrialized" countries are those with per capita income of 15,000 USD or greater. These classifications are also in rough

[7] I am grateful to Arthur van Benthem for sharing detailed calculations for his construction of a real price index, as well as data on historical exchange rates and consumer price indices.

accordance with the World Bank income groups. Under this classification system, my panel includes 53 developed countries, 52 developing countries, and 19 countries in an intermediate range.

4.4 Empirical methods and results

In documenting the empirical analysis, I begin with an exploration of the relationship between energy consumption and income level for developed and developing countries, comparing the ratio of energy demand to GDP per capita between those two groups. I then analyze the energy intensity and carbon intensity of GDP growth, and look for evidence of technology leapfrogging. I also explore the heterogeneity in energy intensity and carbon intensity of income growth across countries and over time, with an eye toward discerning what characteristics or institutions explain the most variation in the energy intensity of growth in different countries.

4.4.1 Levels

As an initial exploration of the relationship between energy consumption and income, I compare energy consumption per capita with GDP per capita between industrialized countries at the time of their development, and developing countries in 2013. Figure 4.2 provides a graphical overview of this relationship for countries in the range of 2,000 to 10,000 USD GDP per capita. Hollow squares in the figure show less-developed countries as of 2013, and solid squares show the developed nations of 2013 at an earlier stage in their development. This includes the earliest year with IEA data on total energy consumption for which today's developed nations had income per capita below 10,000 USD (1960 for most OECD countries, 1971 for most non-OECD countries, and between 1991 and 1993 for five countries in the former Soviet Union or Yugoslavia). It also includes six additional countries that had income per capita above 10,000 USD by the start of the IEA data series, but for which the other sources noted above provide pre-1960 data on energy consumption. For these countries (Canada, Germany, Netherlands, Sweden, USA, and United Kingdom), Figure 4.2 shows energy use and GDP in 1925.

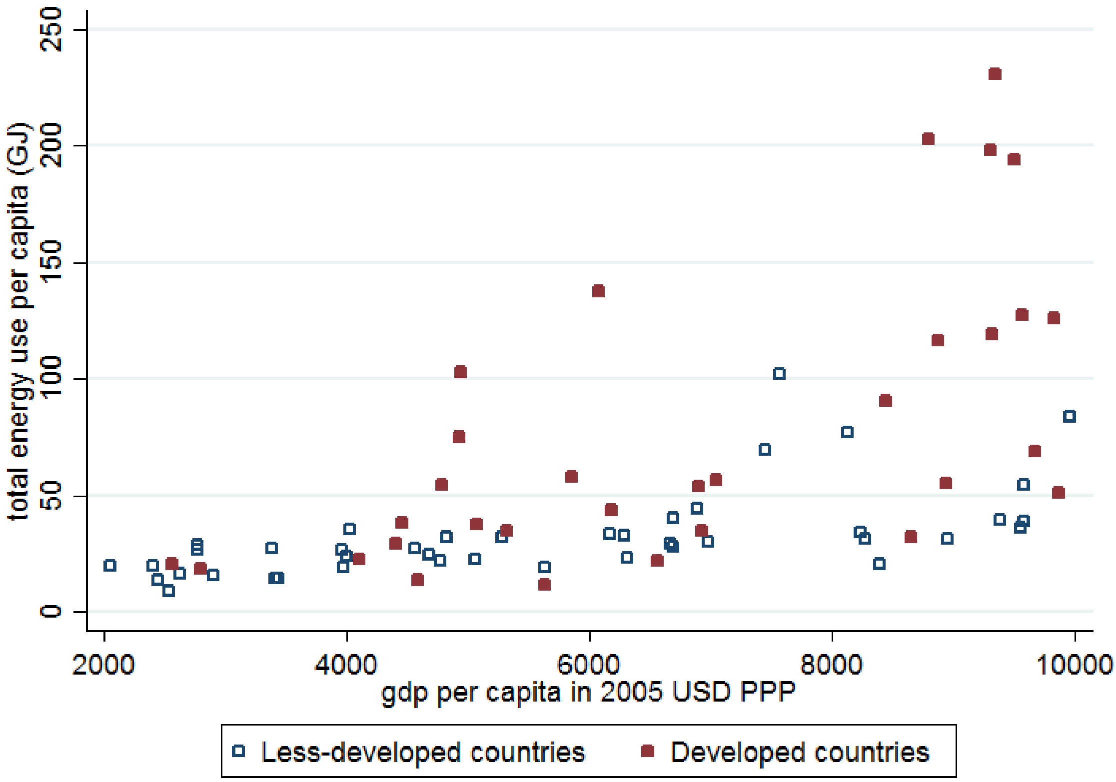

FIGURE 4.2: Total energy use per capita versus GDP for developed countries historically and developing countries in 2013

Figure 4.2 provides suggestive evidence that less-developed countries today may be achieving higher income at a lower level of energy use than industrialized countries did at a comparable time of their development. Many of the solid squares (developed countries) exhibit relatively high levels of energy use for their level of economic development, higher than many of the hollow squares (developing countries). While suggestive, this figure uses only a portion of the data available, since each country is represented only once, and for an arbitrary year.

For a more rigorous analysis, I estimate the quantitative relationship between per capita energy consumption (and, separately, emissions) and GDP, plus a dummy variable for today's less-developed countries. For each dependent variable I use four alternative specifications of the relationship: two in levels and two in logs, and with a linear and a quadratic specification on the income term. Thus, for instance, the specification for energy consumption in logs with a quadratic income term is

$$\ln EC_{it} = \alpha_0 + \alpha_1 \ln GDP_{it} + \alpha_2 (\ln GDP_{it})^2 + 1(LDC_i) + \epsilon_{it}, \tag{4.1}$$

where EC_{it} is energy consumption in country i in year t; GDP_{it} is gross domestic product per capita; and $1(LDC_i)$ is an indicator equal to one if country i is a less-developed country in 2013. I estimate this equation for country-years in which GDP per capita is between 2,000 and 12,000 USD, corresponding to the World Bank definitions of middle-income countries.

In terms of the energy-income relationships illustrated in Figure 4.1, the specification (in logs) forces the paths of less developed countries to be parallel to those of industrialized nations but, through the use of a dummy variable for less developed countries, permits a different intercept. A negative coefficient for the dummy variable would imply that the developing nations of today are experiencing higher income at a lower level of energy consumption compared to previous cohorts of developing countries—that is, to today's industrialized countries during the time of their development.

Table 4.2: Per capita energy use and carbon emissions for developing & developed countries

	DV: Total energy use per capita (GJ)				DV: CO_2 per capita (metric tons)			
Variable	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
$1(LDC_i)$	-21.44***	-21.9***	-0.5***	-0.49***	-1.1	-1.05	-0.37	-0.35
	(7.02)	(7.12)	(0.13)	(0.13)	(0.8)	(0.81)	(0.24)	(0.24)
GDP_{it}	0.007***	0.011**			0.0005***	0.0003		
	(0.001)	(0.004)			(0.0001)	(0.0003)		
$(GDP_{it})^2$		2.9×10^{-7}				1.3×10^{-8}		
		(3.3×10^{-7})				(2.0×10^{-8})		
$\ln(GDP_{it})$			0.78***	0.25			1.43***	0.28
			(0.09)	(2.39)			(0.17)	(4.42)
$(\ln GDP_{it})^2$				0.03				0.07
				(0.14)				(0.26)
Constant	17.52**	8.72	-2.82***	-0.58	0.28	0.66	-11.69***	-6.89
	(7.24)	(10.49)	(0.77)	(10.02)	(0.91)	(0.93)	(1.54)	(18.76)

N=2,563 (total energy use) or 1,809 (energy-related CO_2) country-year observations.
Standard errors, clustered at the country level, in parentheses.
*** $p < .01$, ** $p < .05$, * $p < 0.10$.
Similar to Figure 2, the results focus on country-years where GDP per capita is between 2,000 and 10,000 USD.

Table 4.2 presents the results of this analysis. For both total energy use and energy-related carbon emissions, the coefficient on the LDC dummy is negative in all specifications, supporting the notion that some form of technology leapfrogging is occurring in energy efficiency and perhaps also in the carbon intensity of energy generation. At similar income levels, today's developing countries consume about 49 percent less energy (about 21 GJ less) than did today's industrialized countries during their economic development. Also, at similar income levels, today's developing countries produce about 35 percent less carbon than did today's industrialized countries when they were developing.

With the longer data series, the finding on energy use is opposite to that of van Benthem (2015); in a similar analysis, but with a smaller set of countries and a shorter time period, that study found developing countries using 19-20 percent more energy than developed countries when they had similar income levels. One explanation is that the longer time series allows more time for technological change to take effect, and also allows the use of a larger set of comparison countries—especially industrialized countries.

4.4.2 Intensity of economic growth

The analysis of levels of energy demand for a given income level indicates some support for the occurrence of technology leapfrogging, but a more meaningful analysis would identify the relationship between growth rates of economic activity and energy consumption (or carbon emissions). The time-invariant features of individual countries such as climate or hydroelectric potential affect overall levels of energy demand or carbon emissions, but a key question regarding technology change is whether the next increment of economic growth requires the same proportional increase in energy consumption for today's developing countries as it did for developing countries in the past.

What follows in this section, including the buildup of estimating equations and justification for the approach, hews closely to the analysis of van Benthem (2015). As noted previously, that paper is the most recent and thorough analysis of the possibility of technology leapfrogging in energy systems, and I intentionally set out to replicate the approach with a more expansive dataset and considering a larger range of dependent variables, before analyzing the sources of heterogeneity.

Prior literature (van Benthem, 2015; van Benthem and Romani, 2009; Medlock and Soligo, 2001) considers the relationship between economic growth, energy consumption, and prices using the following specification:

$$\ln EC_{it} = \alpha_0 + \alpha_1 \ln GDP_{it} + \alpha_2 (\ln GDP_{it})^2 + \alpha_3 \ln p_{it} + \theta_i + \lambda_t + \epsilon_{it}, \qquad (4.2)$$

in which EC_{it} and GDP_{it} are defined as in equation (1), p_{it} represents energy prices (the real oil price index discussed in Section 3), and θ and λ represent country and year fixed effects. Country-specific fixed effects allow for time-invariant differences in country-level energy use, such as due to climatic differences. Time fixed effects allow for trends in technological capability or macroeconomic shocks that affect energy consumption over and above the effect of GDP and energy prices.

Equation (4.2) implicitly assumes that energy consumption responds immediately to the effect of prices and income growth, regaining a long-term equilibrium with no lag in response. In reality, when households, firms and governments make decisions about capital stock, those decisions are responsive to prices and income growth over longer periods of time. Incorporating a lagged term of energy consumption in the right-hand side allows the short-run response to differ from that in the long run:

$$\ln EC_{it} = \alpha_0 + \alpha_1 \ln GDP_{it} + \alpha_2 (\ln GDP_{it})^2 + \alpha_3 \ln p_{it} + \gamma \ln EC_{i,t-1} + \theta_i + \lambda_t + \epsilon_{it} \qquad (4.3)$$

In equation (4.3), γ represents the speed of adjustment. The short-run response of energy consumption to prices is still measured by α_3, and response to income growth is still calculated as $\alpha_1 + 2\alpha_2 \ln GDP$. The long-run responses are measured as the short-run responses divided by $(1 - \gamma)$.

Equation (4.3) still has limitations, however. Among the most important is that it restricts the fitted relationship to be quadratic, whereas the true relationship may be more complex. Again following van Benthem (2015), I code a series of dummy variables that split the sample into several income bands based on GDP per capita. I choose these bands to match the 2015 study: per capita income under 3,500 USD; from 3,500 to 10,000 USD; from 10,000 to 20,000 USD, 20,000 to 30,000 USD, and over 30,000 USD (note that in both the 2015 paper and this paper, income is measured in 2005 USD at purchasing power parity). These are in rough accordance with the World Bank classifications of low, middle, and high income countries, although with additional break points within the high income countries. I also code dummy variables that correspond to the classification of countries as being developed or less-developed as of 2013 (i.e., today's developing countries and developing countries of the past).

By interacting these dummy variables with each other and with the explanatory variables, I can distinguish the long-run responsiveness of energy consumption (or carbon emissions) to income growth for the developing countries of today and of the past. In the equation that follows, $1(B_n)$ is an indicator variable with value one if GDP_{it} is within the income band, and zero otherwise. Similarly, $1(LDC_i)$ is an indicator variable with value one if country i is a less-developed country in 2013 (and zero otherwise), and $1(IC_i)$ is defined similarly but for industrialized countries. The equation I estimate is as follows:

$$
\begin{aligned}
\ln EC_{it} = \; & \beta_0 \\
& + 1(LDC_i)[\Sigma_n \beta_{1n,LDC} \ln(GDP_{it})1(B_n) + \beta_{2n,LDC} \ln(P_{it})1(B_n) \\
& + \Sigma_n \gamma_{n,LDC} \ln(EC_{i,t-1})1(B_n)] \\
& + 1(IC_i)[\Sigma_n \beta_{1n,IC} \ln(GDP_{it})1(B_n) + \beta_{2n,IC} \ln(P_{it})1(B_n) \\
& + \Sigma_n \gamma_{n,IC} \ln(EC_{i,t-1})1(B_n)] \\
& + \theta_i + \lambda_t + \epsilon_{it}
\end{aligned}
\tag{4.4}
$$

If technology leapfrogging has occurred, in a way that reduces the energy intensity of growth for today's developing countries, this would correspond to a finding that the long-run responsiveness differs between the LDC parameters and the IC parameters over the same income range. I focus particularly on the income band from 3,500 to 10,000 USD, which encompasses the "takeoff" period of economic development, the period of rising energy intensity according to historical studies, and approximately matches the World Bank classification for middle-income countries. Thus, I test whether $\dfrac{\beta_{1,LDC}}{(1-\gamma_{LDC})}$ is equal to $\dfrac{\beta_{1,IC}}{(1-\gamma_{IC})}$. If the former term is lower, this would suggest that today's developing countries are developing with lower energy intensity.

Table 4.3 presents the results of estimating equation (4.4) and calculating long-run responses. As expected, the long-run response of energy consumption to real energy price is generally negative (or zero). It may seem curious at first that the long-run response of consumption to real energy price is positive and significant for industrialized countries at the highest income level; however, this may be caused by the fact that some industrialized countries are also net energy exporters, so when the real energy price increases their energy consumption also increases because households and other users can afford to use energy less efficiently. Perhaps surprisingly, in most specifications the energy intensity of GDP growth does not decline with higher levels of income (e.g., in column 1, within industrialized countries the magnitudes of the coefficients on $\ln GDP$ are about the same for higher income levels as they are for lower income levels). This runs counter to the historical observation that energy intensity declines for higher income levels, and may be due to the relatively low threshold for the highest category of income; however, this puzzle warrants further investigation.

The main question of interest is whether the long-run response of energy consumption to economic growth is lower for today's less-developed countries (second row of coefficients in Table 4.3) than for today's industrialized countries when they were at similar levels of income (third row of coefficients). The last row in the table provides the results of this test, which suggests that there is some form of technological advance that has allowed less-developed nations to grow with lower energy intensity than in the past. The result is weakly statistically significant but stable across specifications.

This finding is counter to that of van Benthem (2015), who found no significant change using the same empirical test (but a smaller set of countries and years). As with other differences in findings, the difference may arise from several sources. In Table 4.3 the measure of energy use (total final consumption by end users) is the same as in van Benthem (2015), so that is not the source of the different result. One likely cause is that I use a longer time series and a broader set of countries, which may better capture the time scale necessary for widespread technological change. However, where that paper had data on energy consumption and GDP at the country-sector-year level, I have data only at the country-year level. Thus, there may be sectoral shifts in economic activity that I do not capture. In this sense, the 2015 paper suggests an absence of energy leapfrogging within sectors (e.g., leapfrogging is not resulting in more energy-efficient manufacturing or transportation in developing countries), while this paper suggests leapfrogging is occurring within countries (perhaps as a combination of sectoral shifts and leapfrogging over a longer time period).

Table 4.3: Long-run response of total final consumption to income and price, by income group and development status

Variable	Income band (USD $k)	Country group	(1)	(2)	(3)	(4)	(5)
$\ln(GDP)$	0-3.5	LDC	0.49***	0.52***	0.52***	0.25***	0.25***
			(0.07)	(0.08)	(0.09)	(0.09)	(0.09)
	3.5-10	LDC	0.50***	0.52***	0.52***	0.88***	0.94***
			(0.06)	(0.08)	(0.08)	(0.12)	(0.12)
	3.5-10	IC	0.60***	0.64***	0.61***	1.06***	0.94***
			(0.08)	(0.10)	(0.11)	(0.22)	(0.30)
	10-20	IC	0.64***	0.67***	0.64***	0.68***	0.71***
			(0.09)	(0.12)	(0.13)	(0.17)	(0.17)
	20-30	IC	0.62***	0.63***	0.69***	0.23	0.25
			(0.09)	(0.12)	(0.13)	(0.32)	(0.32)
	> 30	IC	0.49***	0.52***	0.50***	0.83***	0.83***
			(0.06)	(0.09)	(0.09)	(0.27)	(0.27)
ln (p)	0-3.5	LDC	-0.01	0.06	0.03	0.03	-0.02
			(0.05)	(0.06)	(0.06)	(0.08)	(0.10)
	3.5-10	LDC	-0.12***	-0.07**	-0.06**	-0.07***	-0.08***
			(0.04)	(0.03)	(0.03)	(0.02)	(0.02)
	3.5-10	IC	-0.12**	-0.04	-0.01	-0.02	-0.02
			(0.05)	(0.04)	(0.04)	(0.03)	(0.03)
	10-20	IC	-0.10*	0.02	0.06	-0.08	-0.08
			(0.05)	(0.07)	(0.08)	(0.22)	(0.23)
	20-30	IC	0.10***	0.13***	0.12***	0.08**	0.09**
			(0.03)	(0.02)	(0.02)	(0.04)	(0.04
	> 30	IC	-0.01	0.06	0.03	0.03	-0.02
			(0.05)	(0.06)	(0.06)	(0.08)	(0.10)
Time fixed effects			no	λ_t	$\lambda_t \times (LDC, IC)$	$\lambda_t \times 1(B_n)$	$\lambda_t \times (LDC, IC) \times 1(B_n)$
Income coefficient (LDC-IC) within 3.5-10k USD income band			-0.09*	-0.12*	-0.09	-0.19	0
			(0.06)	(0.06)	(0.07)	(0.17)	(0.32)

Results of estimating Equation 4.4 with dependent variable log of total final consumption per capita.
Standard errors, clustered at the country level, in parentheses.
Includes country fixed effects. N=3,969 country-year observations.
$^{***}p < .01, ^{**}p < .05, ^{*}p < 0.10$.

In the absence of sector-level data, I cannot distinguish how much of the difference arises from sectoral shifts and how much from the longer time horizon. However, if agents can take advantage of more free trade or other institutions that facilitate sectoral shifts even as they allow income growth, then this may be sufficiently informative for policymakers and others interested in recognizing and facilitating leapfrog opportunities.

I run the same test for total energy use and for energy-related carbon emissions. Table 4.4 presents the results of this test for total energy use, presenting results only for comparisons of developing and industrialized countries in the 3,500-10,000 USD income band (however, the estimation strategy is the same as that shown in equation (4.4) and Table 4.3). The results suggest weak evidence of leapfrogging in total energy use: perhaps surprisingly, weaker than in total final consumption. Since the difference between the two series amounts to energy industry own use and energy losses in transformation and distribution, this suggests that there has been more leapfrogging in end-user consumption than in these segments. Indeed, an analysis of own use and transformation losses separately might suggest the energy intensity of these losses has increased for today's developing countries, which is a surprising result given advances in technology for the production of energy. One possible explanation is that energy sources are becoming harder to access (e.g., oil exploration in more remote areas, coal seams deeper underground).

Table 4.5 provides results of the same test, using per-capita energy-related carbon emissions as the dependent variable. There is weak evidence for "carbon leapfrogging," with the magnitudes of coefficients for developing countries slightly smaller than for developed countries (but not significantly so). The responses shown in Table 4.5 are a composite of the energy intensity of income and the carbon intensity of energy consumption; thus, they partly reflect the same trends as in Table 4.4. The responses are also consistent with the observation that, with a few notable exceptions

Table 4.4: Long-run response of total energy use to income

Variable	Income band (USD $k)	Country group	(1)	(2)	(3)	(4)	(5)
ln (GDP)	3.5-10	LDC	0.59***	0.54***	0.56***	0.83***	0.80***
			(0.07)	(0.08)	(0.08)	(0.11)	(0.11)
	3.5-10	IC	0.62***	0.62***	0.61***	0.87***	0.90***
			(0.06)	(0.07)	(0.08)	(0.11)	(0.17)
Time fixed effects			no	λ_t	$\lambda_t\times$ (LDC, IC)	$\lambda_t\times$ $1(B_n)$	$\lambda_t\times$ (LDC, IC) $\times 1(B_n)$
Income coefficient (LDC-IC) within 3.5-10k USD income band			-0.03	-0.08*	-0.05	-0.04	-0.1
			(0.04)	(0.04)	(0.04)	(0.08)	(0.20)

Results of estimating Equation 4.4 with dependent variable log of total energy use per capita. (A subset of parameter estimates are shown here.)
Standard errors, clustered at the country level, in parentheses.
Includes country fixed effects. N=4,586 country-year observations.
$^{***}p < .01, ^{**} p < .05, ^{*} p < 0.10$.

(like China), most developing countries have invested relatively little in low-carbon generation sources.

4.4.3 Heterogeneity and robustness

Understanding the heterogeneity across countries (and perhaps time periods of development) would contribute to understanding the effects of technology policies and environmental regulations in developing and developed nations. For instance, Popp (2011) finds that policies in developed countries tend to drive innovation in emissions-reducing technologies, and energy-efficient innovations diffuse to low-income countries regardless of domestic environmental policies, but adoption of other technologies (that do not increase firms' profits) does not increase in the absence of environmental policy. Identifying the net effect of current policies in the aggregate furthers our understanding of how policy in low-income and industrialized countries has affected

Table 4.5: Long-run response of CO_2 emissions to income

Variable	Income band (USD $k)	Country group	(1)	(2)	(3)	(4)	(5)
ln (GDP)	3.5-10	LDC	1.11***	1.05***	0.87***	1.07***	0.98***
			(0.18)	(0.19)	(0.18)	(0.22)	(0.24)
	3.5-10	IC	1.06***	1.08***	1.09***	1.21***	1.40***
			(0.14)	(0.17)	(0.20)	(0.23)	(0.32)
Time fixed effects			no	λ_t	$\lambda_t \times (LDC, IC)$	$\lambda_t \times 1(B_n)$	$\lambda_t \times (LDC, IC) \times 1(B_n)$
Income coefficient (LDC-IC) within 3.5-10k USD income band			0.05	-0.04	-0.22	-0.15	-0.42
			(0.20)	(0.20)	(0.22)	(0.22)	(0.40)

Results of estimating Equation 4.4 with dependent variable log of carbon emissions per capita. (A subset of parameter estimates are shown here.)
Standard errors, clustered at the country level, in parentheses.
Includes country fixed effects. N=3,969 country-year observations.
$^{***}p < .01, ^{**}p < .05, ^{*}p < 0.10$.

energy technology innovation and transfer.

A promising research angle would be to pursue further the question of heterogeneous responses, for instance using quantile regression techniques that also account for the panel data structure (Kato et al., 2012; Lamarche, 2010) to analyze energy intensity and carbon intensity of economic growth. In addition to illustrating the heterogeneity across countries, such an approach may also help identify the causes of that heterogeneity. For instance, other researchers have identified various factors that seem to contribute to different responses: the maturity of commercial banking and thus availability of financing for renewable energy technologies (Brunnschweiler, 2010), openness to trade (Lovely and Popp, 2011), natural environmental factors that necessitate more energy use for space conditioning (Smil, 2010), and others. Each of these would suggest a different policy response to maximize potential for technology transfer to reduce the energy intensity of economic growth.

It may also be useful to explore the relationships described in equation (4.4) through the addition of panel-heterogeneous time trends (Bai, 2009), in addition to the time-invariant country fixed effects and uniform-country time fixed effects. It is reasonable to believe that macroeconomic shocks or technological advance may affect different countries differently, and using heterogeneous time trends provides one way to investigate this heterogeneity. This may also provide some insight on countries' heterogeneous relationships between energy consumption and growth.

Additional insight might be gained by separating out OPEC countries and post-Soviet transition economies. For different reasons, these nations may be outliers in terms of having relatively high energy intensity of economic growth. However, they are generally split between industrialized and developing countries, so special treatment of these observations may not affect the core findings as they relate to leapfrogging.

Finally, it seems useful to explore the effects of using different thresholds to distinguish developed and industrialized nations. The present analysis uses 2013 income levels and follows van Benthem (2015) to facilitate comparability with that study, but it may be more instructive to distinguish developed countries based on the era in which their "takeoff" development phase took place. This would further help to characterize the heterogeneity of nations' experiences, and may clarify the relationship of my findings to those of van Benthem (2015), since that study did not find evidence of leapfrogging for countries that developed earlier (roughly in 1970-2000).

4.5 Conclusions

With demand for energy services growing rapidly, and especially across the developing world, planners and policy makers rely on accurate quantitative forecasts of energy demand to balance the need to build out energy infrastructure relative to other pressing needs for public investment. A previous paper (van Benthem, 2015) suggests there are important technology rebound effects that may drive today's developing countries to use more energy per unit of economic growth than did developing countries of the past, which bodes ill for forecasters who are assuming some level of energy savings or "leapfrogging" due to the availability of more efficient technologies.

Using a much longer time series and a broader set of countries in the analysis, I offer additional insight and demonstrate that while technology rebound effect may be of concern for some countries, on average it is not, when considering a broader range of developed countries, a longer timeline for analysis, and a more complete set of energy technologies—especially on the generation and distribution side. That is, the forecasting agencies may be correct after all in assuming some level of energy efficiency available to today's developing countries that did not exist in previous decades. That said, the more valuable policy implications may be those that come after continued exploration, including exploration of heterogeneous effects and interactions with historical or current institutions.

Appendix A

Energy Consumption Data

I use four sources for energy consumption data prior to 1960. The principal source is Kander et al. (2014), which provides data for eight European countries from 1861-2008. I also obtain data for Canada from 1861-2002 from Unger and Thistle (2013). All of these authors worked together to recover and tabulate historical energy consumption data from each respective country using common methods, including conducting an extensive review to verify comparability across countries. For the USA I use two additional sources, compiled by authors working independently from the European and Canadian group (Schurr et al., 1960; EIA, 2016).

In addition to working to verify internal consistency, the authors of Kander et al. (2014) also attempted to make their data series consistent with the IEA Extended Energy Balances series, especially through 1970 (personal communication with Astrid Kander). For most countries, the overlap is reasonably good. Figure A.1, for instance, shows both the series from both IEA (2015a) and Kander et al. (2014) for the Netherlands. The data series match almost identically from 1960 through 1971, then diverge, with the IEA series initially slightly higher and then slightly lower than the series from Kander et al. (2014). Nevertheless, the same trends are apparent in both (albeit with somewhat more scatter in the series from Kander et al. (2014)). The comparability of the data series, especially during the period 1960 to 1971, suggests it is reasonable to use the series from Kander et al. (2014) to extend the IEA series backwards.

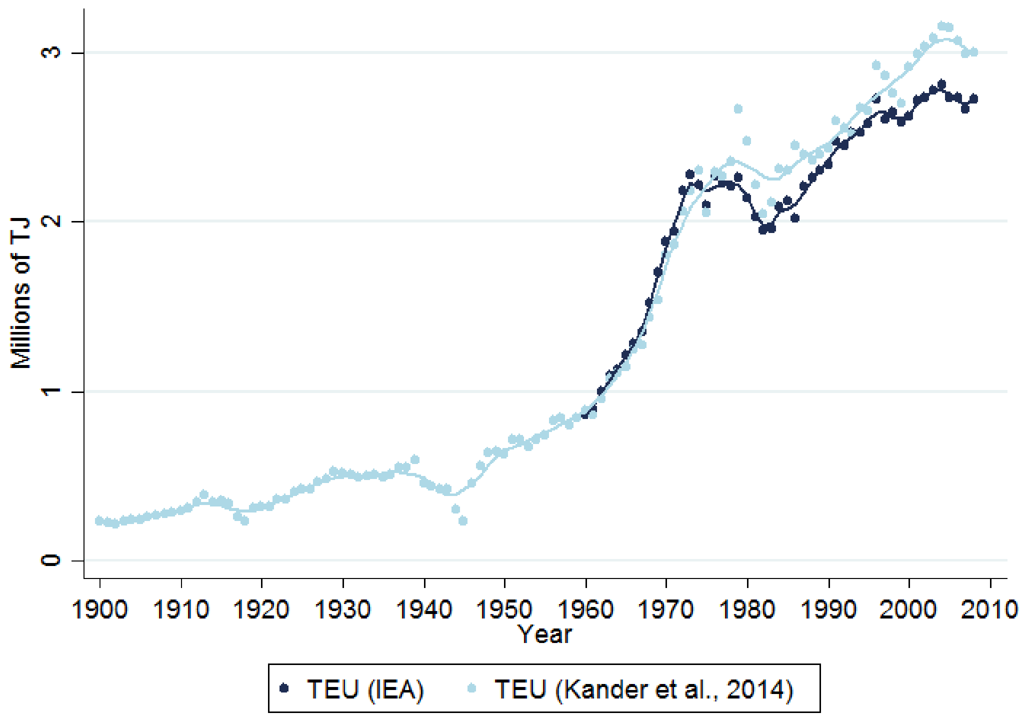

Figure A.1: Total Energy Use in the Netherlands: Comparison of IEA (2015a) and Kander et al. (2014)

The series for Germany (Figure A.2) suggests more reason to be cautious in melding the two series. Here, the series match reasonably well from 1970 through 2008, but from 1960 through 1969 the IEA series indicate substantially lower energy consumption, in a way that is discontinuous with both the preceding observations from Kander et al. (2014) and the subsequent observations from IEA (2015a). In Sweden, on the other hand, the series match reasonably well until the late 1970s, at which point the IEA series indicate substantially higher energy consumption, on the order of 30% (Figure A.3).

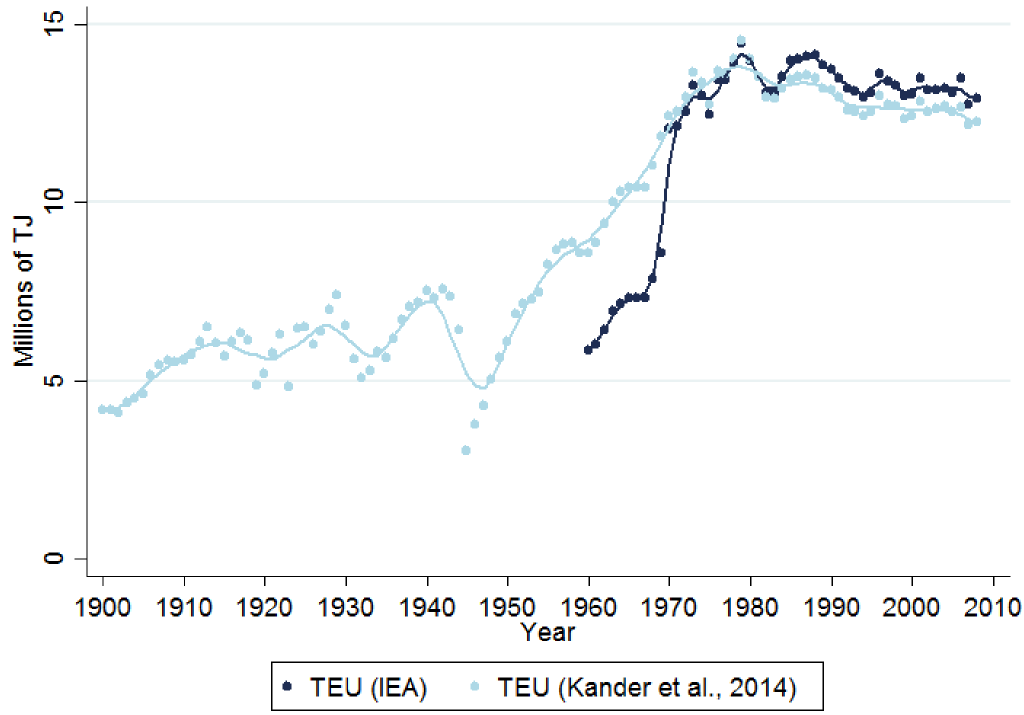

Figure A.2: Total Energy Use in Germany: Comparison of IEA (2015a) and Kander et al. (2014)

One reason for the divergence of the two series may be different assumptions regarding the efficiency of conversion for primary electricity sources, including geothermal and nuclear sources (personal communication with Sofia Henriques). In any case, in the regression analyses I include dummy variables to allow for these structural trend breaks (as well as other breaks within the IEA series, such as methodological changes or switches from the use of one national statistical series to another).

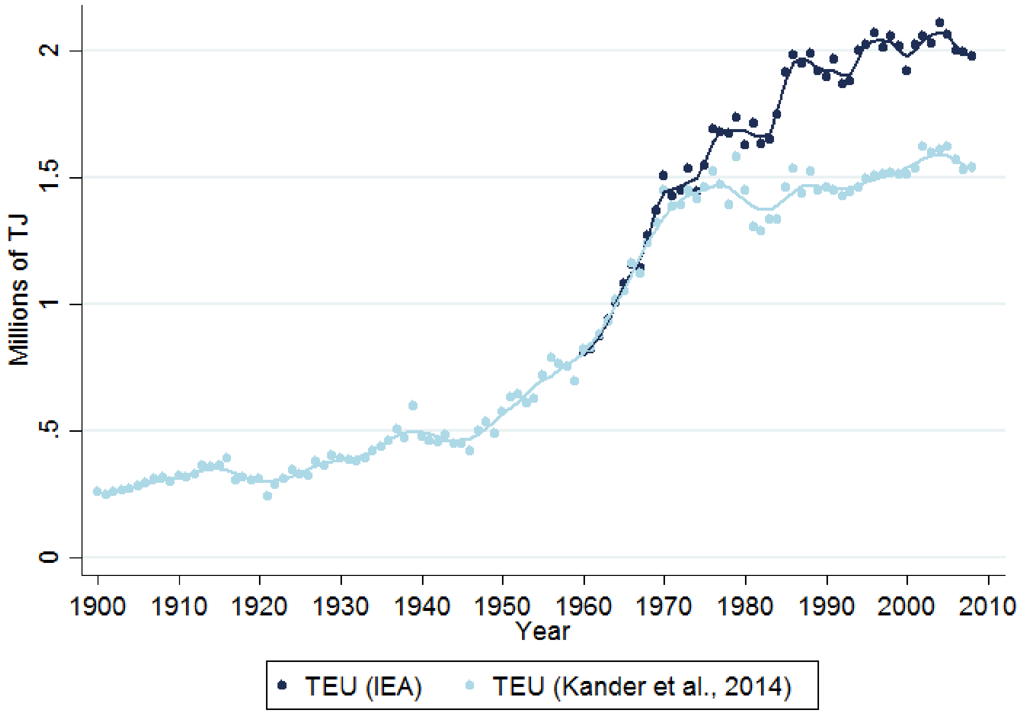

FIGURE A.3: Total Energy Use in Sweden: Comparison of IEA (2015a) and Kander et al. (2014)

Bibliography

Amankwah-Amoah, J. (2014), "Solar Energy in Sub-Saharan Africa: The Challenges and Opportunities of Technological Leapfrogging," *Thunderbird International Business Review*, 57, 15–31.

Angrist, J. D. and Pischke, J.-S. (2008), *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.

Bai, J. (2009), "Panel Data Models With Interactive Fixed Effects," *Econometrica*, 77, 1229–1279.

Bennear, L. S. and Coglianese, C. (2005), "Measuring Progress: Program Evaluation of Environmental Policies," *Environment: Science and Policy for Sustainable Development*, 47, 22–39.

Bennear, L. S. and Olmstead, S. M. (2008), "The impacts of the "right to know": Information disclosure and the violation of drinking water standards," *Journal of Environmental Economics and Management*, 56, 117–130.

Bernard, A. B. and Durlauf, S. N. (1996), "Interpreting tests of the convergence hypothesis," *Journal of Econometrics*, 71, 161–173.

Bertrand, M., Duflo, E., and Mullainathan, S. (2004), "How Much Should We Trust Differences-In-Differences Estimates?" *The Quarterly Journal of Economics*, 119, 249–275.

Bordo, M. D. (2015), Financial Crises Database.

Brunnschweiler, C. N. (2010), "Finance for renewable energy: an empirical analysis of developing and transition economies," *Environment and Development Economics*, 15, 241–274.

Cameron, A. C. and Miller, D. L. (2015), "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 50, 317–372.

Cameron, A. C. and Trivedi, P. K. (2005), *Microeconometrics: Methods and Applications*, Cambridge University Press.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008), "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90, 414–427.

Cohen, W. M. (2010), "Fifty Years of Empirical Studies of Innovative Activity and Performance," in *Handbook of The Economics of Innovation, Vol. 1*, eds. B. H. Hall and N. Rosenberg, pp. 129–213, Elsevier BV, Chapter 4.

Cohen, W. M., Nelson, R. R., and Walsh, J. P. (2000), "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)," Working Paper 7552, National Bureau of Economic Research.

Conley, T. G. and Udry, C. R. (2010), "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, 100, 35–69.

Covert, T. (2015), "Experiential and Social Learning in Firms: The Case of Hydraulic Fracturing in the Bakken Shale," Unpublished manuscript.

Dasgupta, S., Wheeler, D., and Wang, H. (2007), "Disclosure strategies for pollution control," in *The International Yearbook of Environmental And Resource Economics 2006/2007: A Survey of Current Issues (New Horizons in Environmental Economics Series)*, ed. T. Tietenberg, Edward Elgar Publishers.

David, P. (1990), "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review*, 80, 355–61.

Delmas, M., Montes-Sancho, M. J., and Shimshack, J. P. (2010), "Information Disclosure Policies: Evidence From the Electricity Industry," *Economic Inquiry*, 48, 483–498.

Doshi, A. R., Dowell, G. W. S., and Toffel, M. W. (2013), "How firms respond to mandatory information disclosure," *Strategic Management Journal*, 34, 1209–1231.

Elgin, B., Haas, B., and Kuntz, P. (2012), "Fracking Secrets by Thousands Keep U.S. Clueless on Wells," *Bloomberg*.

Fung, A., Graham, M., and Weil, D. (2007), *Full Disclosure: The Perils and Promise of Transparency*, Cambridge University Press, New York.

Galli, R. (1998), "The Relationship Between Energy Intensity and Income Levels: Forecasting Long Term Energy Demand in Asian Emerging Countries," *The Energy Journal*, 19.

Gertler, P. J., Shelef, O., Wolfram, C. D., and Fuchs, A. (2016), "The Demand for Energy-Using Assets among the World's Rising Middle Classes," *American Economic Review*, 106, 1366–1401.

Goldemberg, J. (1998), "Viewpoint: Leapfrog energy technologies," *Energy Policy*, 26, 729–741.

Ground Water Protection Council and Interstate Oil & Gas Compact Commission (GWPC and IOGCC) (2015), FracFocus Chemical Disclosure Registry.

Grübler, A., Nakićenović, N., and Victor, D. G. (1999), "Dynamics of energy technologies and global change," *Energy Policy*, 27, 247–280.

Gulbis, J. and Hodge, R. M. (2000), "Fracturing Fluid Chemistry and Proppants," in *Reservoir Stimulation*, eds. K. Nolte and M. Economides, Chapter 7. Third Edition.

Haas, B., Polson, J., Kuntz, P., and Elgin, B. (2012), "Fracking Hazards Obscured in Failure to Disclose Wells," *Bloomberg*.

Hall, B. H. (2004), "Innovation and Diffusion," Working Paper 10212, National Bureau of Economic Research.

Hamilton, J. T. (1995), "Pollution as News: Media and Stock Market Reactions to the Toxics Release Inventory Data," *Journal of Environmental Economics and Management*, 28, 98–113.

Hausman, C. and Kellogg, R. (2015), "Welfare and Distributional Implications of Shale Gas," .

International Energy Agency (IEA) (2015a), Extended world energy balances.

International Energy Agency (IEA) (2015b), CO2 Emissions from Fuel Combustion.

Jin, G. Z. and Leslie, P. (2003), "The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards," *The Quarterly Journal of Economics*, 118, 409–451.

Judson, R. A., Schmalensee, R., and Stoker, T. M. (1999), "Economic Development and the Structure of the Demand for Commercial Energy," *The Energy Journal*, 20.

Kander, A., Malanima, P., and Warde, P. (2014), *Power to the People: Energy in Europe over the Last Five Centuries (The Princeton Economic History of the Western World)*, Princeton University Press.

Kato, K., Galvao, A. F., and Montes-Rojas, G. V. (2012), "Asymptotics for panel quantile regression models with individual effects," *Journal of Econometrics*, 170, 76–91.

Kim, E.-H. and Lyon, T. P. (2011), "Strategic environmental disclosure: Evidence from the DOE voluntary greenhouse gas registry," *Journal of Environmental Economics and Management*, 61, 311–326.

Konschnik, K. (2014), "Goal-Oriented Disclosure Design for Shale Oil and Gas Development," *Natural Resources Journal*, 55, 319.

Krebs, C. (2014), *Ecological Methodology*, Addison Wesley Longman, Menlo Park, California.

Lamarche, C. (2010), "Robust penalized quantile regression estimation for panel data," *Journal of Econometrics*, 157, 396–408.

Lovely, M. and Popp, D. (2011), "Trade, technology, and the environment: Does access to technology promote environmental regulation?" *Journal of Environmental Economics and Management*, 61, 16–35.

Lyon, T. P. and Maxwell, J. W. (2004), *Corporate Environmentalism and Public Policy*, Cambridge University Press.

Mackinnon, J. G. and Webb, M. D. (2016), "Wild Bootstrap Inference for Wildly Different Cluster Sizes," *Journal of Applied Econometrics*.

Maddison, A. (2010), Historical Statistics of the World Economy: 1-2008 AD.

Maykuth, A. (2012), "More states ordering disclosure of fracking chemicals," Philadelphia Inquirer, January 22.

Medlock, K. B. and Soligo, R. (2001), "Economic Development and End-Use Energy Demand," *The Energy Journal*, 22, 77–105.

Montgomery, C. (2013), "Fracturing Fluid Components," in *Effective and Sustainable Hydraulic Fracturing*, InTech.

Muehlenbachs, L., Spiller, E., and Timmins, C. (2015), "The Housing Market Impacts of Shale Gas Development," *American Economic Review*, 105, 3633–3659.

Newell, R. G., Prest, B. C., and Vissing, A. (2016), "Trophy Hunting vs. Manufacturing Energy: The Price-Responsiveness of Shale Gas," Working Paper 22532, National Bureau of Economic Research.

Olmstead, S. M. and Richardson, N. (2014), "Managing the Risks of Shale Gas Development Using Innovative Legal and Regulatory Approaches," Discussion paper, Resources for the Future.

Pearson, N. O. (2014), "Cheap Electricity for Poor Squeezing Out Solar in India," *Bloomberg News*.

Popp, D. (2011), "International Technology Transfer, Climate Change, and the Clean Development Mechanism," *Review of Environmental Economics and Policy*, 5, 131–152.

Robart, C., Ruegamer, M., and Yang, A. (2013), "Analysis of U.S. Hydraulic Fracturing Fluid System Trends," SPE Paper 163875, Society of Petroleum Engineers.

Rosenberg, N. (1972), "Factors affecting the diffusion of technology," *Explorations in Economic History*, 10, 3–33.

Schurr, S. H., Netschert, B. C., and Resources for the Future (1960), *Energy in the American Economy, 1850-1975*, Johns Hopkins Press.

Shellenberger, M., Nordhaus, T., Trembath, A., and Jenkins, J. (2012), "Where the Shale Gas Revolution Came From: Government's Role in the Development of Hydraulic Fracturing in Shale," .

Sider, A. (2010), "Panel set to expose drill-fluid contents," Arkansas Democrat-Gazette, October 30.

Skytruth (2013), "Fracking Chemical Database," Online database.

Smil, V. (2010), *Energy Transitions: History, Requirements, Prospects*, Praeger.

Stringfellow, W. T., Domen, J. K., Camarillo, M. K., Sandelin, W. L., and Borglin, S. (2014), "Physical, chemical, and biological characteristics of compounds used in hydraulic fracturing," *Journal of Hazardous Materials*, 275, 37–54.

Teece, D. J. (1986), "Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy," *Research Policy*, 15, 285–305.

Unger, R. W. and Thistle, J. (2013), *Energy Consumption in Canada in the 19th and 20th Centuries: A Statistical Outline*, CNR Edizioni.

U.S. Energy Information Administration (EIA) (2016)Annual Energy Review (Table 2.1a, Energy consumption estimates by sector, 1949-2012).

U.S. Environmental Protection Agency (USEPA) (2012a), "Draft Permitting Guidance for Oil and Gas Hydraulic Fracturing Activities Using Diesel Fuels: Underground Injection Control Program Guidance #84," EPA 816-R-12-004.

U.S. Environmental Protection Agency (USEPA) (2012b), "Risk-Screening Environmental Indicators, Technical Appendix A: Listing of All Toxicity Weights for TRI Chemicals and Chemical Categories," EPA 816-R-12-004.

U.S. Environmental Protection Agency (USEPA) (2014), "Final Permitting Guidance for Oil and Gas Hydraulic Fracturing Activities Using Diesel Fuels: Underground Injection Control Program Guidance #84," EPA 816-R-14-001.

van Benthem, A. A. (2015), "Energy Leapfrogging," *Journal of the Association of Environmental and Resource Economists*, 2, 93–132.

van Benthem, A. A. and Romani, M. (2009), "Fuelling Growth: What Drives Energy Demand in Developing Countries?" *The Energy Journal*, 30, 91–114.

Wang, Z. and Krupnick, A. (2013), "A Retrospective Review of Shale Gas Development in the United States: What Led to the Boom?" Discussion paper, Resources for the Future.

Waxman, H. A., Markey, E. J., and DeGette, D. (2011), "Chemicals Used In Hydraulic Fracturing," Tech. rep., United States House of Representatives, Committee on Energy and Commerce.

Webb, M. D. (2014), "Reworking Wild Bootstrap Based Inference for Clustered Errors," Working Paper 1315, Queen's University Department of Economics.

White, H. (2000), *Asymptotic Theory for Econometricians: Revised Edition (Economic Theory, Econometrics, and Mathematical Economics) (Economic Theory, Econometrics, & Mathematical Economics)*, Emerald Group Publishing Limited.

Wolfram, C., Shelef, O., and Gertler, P. (2012), "How Will Energy Demand Develop in the Developing World?" *Journal of Economic Perspectives*, 26, 119–138.

Wrenn, D. H., Klaiber, H. A., and Jaenicke, E. C. (2016), "Unconventional Shale Gas Development, Risk Perceptions, and Averting Behavior: Evidence from Bottled Water Purchases," *Journal of the Association of Environmental and Resource Economists*, 3, 779–817.

Biography

Theodore Robert Fetter was born November 28, 1974, in Saint Paul, Minnesota, USA. He attended two schools as an undergraduate: From 1992-1995 he attended Brown University, with a declared major in Geology-Biology; from 1997-1999 he attended the University of Massachusetts Amherst, and earned a Bachelor of Science in Resource Economics in 1999. He also holds a Master of Science in Resource Economics from the University of Massachusetts Amherst, awarded in 2002; a Master of Science from the Yale School of Forestry and Environmental Studies, awarded in 2015; and a Ph.D. from Duke, awarded in 2017.

He is a recipient of the Joseph L. Fisher Doctoral Dissertation Fellowship from Resources for the Future and the Duke Environmental Economics Doctoral Scholars Fellowship from the Nicholas Institute for Environmental Policy Solutions. He has also received funding from the Duke University Energy Initiative, where he was a Doctoral Student Fellow. In addition, he has received grants from Equitable Origin and the Yale Center for Environmental Law and Policy.

He is the co-author, with Michael A.O. Ash, of "Who Lives on the Wrong Side of the Environmental Tracks? Evidence from the EPA's Risk-Screening Environmental Indicators Model," published in the *Social Science Quarterly* in 2004. He is also the co-author, with Julie A. Caswell, of "Variation in Organic Standards Prior to the National Organic Program," published in the *American Journal of Alternative Agriculture* in 2002.