Title
Essays on Natural Gas in the Transition to Sustainable Energy

Author
Marks, Levi

Publication Date
2019

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Essays on Natural Gas in the Transition to Sustainable Energy

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

Doctor of Philosophy
in
Economics

by

Levi Marks

Committee in charge:

Professor Olivier Deschenes, Chair
Professor H.E. Frech III
Professor Kyle Meng
Professor Matthew Zaragoza-Watkins, Vanderbilt University

June 2019

The Dissertation of Levi Marks is approved.

Professor H.E. Frech III

Professor Kyle Meng

Professor Matthew Zaragoza-Watkins, Vanderbilt University

Professor Olivier Deschenes, Committee Chair

June 2019

Essays on Natural Gas in the Transition to Sustainable Energy

Acknowledgements

I am deeply thankful to my advisor Olivier Deschenes for the teaching, guidance, and support he has provided over the past six years. I am also very grateful to Matthew Zaragoza-Watkins for his invaluable mentorship in economics, professional development, and other areas of life. I greatly appreciate Ted Frech and Kyle Meng for their support throughout my studies and their insightful feedback that helped shape this dissertation.

I also extend my thanks to my cohort members and especially to Andrew Ayres, Sarah Bana, Chang Lee, and Michael Cooper for their consistent willingness to offer advice on my dissertation research and for their friendship over the years. I am highly appreciative of Kristina Mohlin and others at Environmental Defense Fund for assistance acquiring data used in this dissertation and for the opportunities they provided to engage with environmental policy during my program. I am grateful to my undergraduate professors Brian Treanor, Dorothea Herreiner, and Andrew Healy for inspiring what has become an immensely fulfilling life path. Finally, I am ever thankful to my parents for their unconditional generosity, love, and encouragement.

Curriculum Vitæ
Levi Marks

Education

2019	Ph.D. in Economics, University of California, Santa Barbara
2014	M.A. in Economics, University of California, Santa Barbara
2011	B.A. in Economics and Philosophy, Loyola Marymount University

Research Fields

Environmental Economics, Energy Economics

Working Papers

Vertical Market Power in Interconnected Natural Gas and Electricity Markets
with Charles F. Mason, Kristina Mohlin, and Matthew Zaragoza-Watkins

The Abatement Cost of Methane Emissions from Natural Gas Production

An Estimate-Based Approach to Emissions Pricing

How did the BP oil spill affect environmental attitudes?

Abstract

Essays on Natural Gas in the Transition to Sustainable Energy

by

Levi Marks

As the world transitions toward a low-carbon energy economy, two features of natural gas set it apart from other fossil fuels. First, the combustion of natural gas generates about half as much carbon dioxide as coal. Second, gas-fired electricity generation can be ramped up and down quickly and efficiently, making it well-suited for balancing out intermittently available wind and solar energy. These two factors lead many to believe natural gas will play an extended role in the electricity generation mix. However, there are many ways in which current markets for natural gas not well-adapted for its role in the energy transition. This dissertation explores two areas where natural gas markets can be improved for its efficient utilization in the transition to low-carbon energy. The first two chapters develop policy tools to help reduce methane emissions from the natural gas supply chain, and the third chapter investigates a previously unknown market failure that can arise in interconnected natural gas and electricity markets.

The first chapter of this dissertation empirically estimates the cost of reducing methane emissions from the extraction segment of the natural gas industry. Although natural gas has important climate benefits, it is composed of about 90 percent methane, which is itself a greenhouse gas that is far more potent than carbon dioxide. A small fraction of emitted gas (in the form of equipment leaks and intentional venting) can therefore have severe warming effects. This chapter estimates the cost of reducing these emissions by examining how production facilities' emission rates respond to changes in natural gas prices. Because firms mitigate emissions up to the point at which their marginal cost

of mitigation equals their marginal private benefit of being able to sell captured gas, an estimated relationship between emission rates and prices can be used to determine mitigation costs. Results indicate that methane emissions from natural gas production can be reduced at very low cost relative to other sources of greenhouse gas emissions. For example, an emissions price equivalent to the social cost of methane is predicted to decrease emissions by about 76 percent while increasing the net cost of natural gas extraction by less than one percent.

Building on this result, the second chapter explores how emissions pricing can be used to regulate methane emissions in practice. Previously, emissions pricing programs have been implemented based on the carbon content of fossil fuels or by using continuous emissions monitoring sensors placed in smokestacks. However, because methane emissions from the natural gas industry are released from many different sources in a variety of different ways, comprehensively monitoring them is prohibitively costly at this time. This chapter outlines a novel estimate-based approach for implementing emissions pricing in this setting. Rather than monitoring emissions at all facilities continuously, the regulator randomly selects a subset of each firm's facilities to perform measurements at. The regulator then uses these measurements to develop a firm-level estimate of emissions, which can then be used to apply an emissions tax or account for the use of permits. A theoretical model demonstrates that this approach preserves the efficiency benefits of emissions pricing with comprehensive measurement. Furthermore, a simulation calibrated to be representative of the U.S. natural gas industry predicts that this approach can achieve climate mitigation benefits roughly two orders of magnitude greater than the cost of measurement.

The third chapter, which is coauthored with Charles F. Mason, Kristina Mohlin, and Matthew Zaragoza-Watkins, explores a market failure that can arise from the increasing interdependence of natural gas and electricity markets. It develops a theoretical

model that illustrates conditions under which a firm that owns both electricity generation plants and contracts for natural gas pipeline capacity may find it optimal to withhold those contracts from secondary markets. By artificially limiting the available supply of pipeline capacity on constrained days, this behavior increases electricity prices in the downstream electricity generation market, which benefits non-gas generators owned by the withholding firms. We document pipeline scheduling patterns exhibited by two firms in New England that are consistent with this behavior. We then estimate the impacts of this behavior, finding that it increased wholesale natural gas and electricity prices by 35 percent and 18 percent, respectively. We estimate that substitution from natural gas generation to coal and oil generation due to these artificial supply constraints resulted in economic losses of $1.5 billion over a three-year period. While this behavior may have been within the firms' contractual rights, these findings underscore a need to improve regulation and coordination of these increasingly linked energy markets.

Contents

Chapter 1

The Abatement Cost of Methane Emissions from Natural Gas Production

1.1 Introduction

Consumption of natural gas has increased considerably over the last decade. In the U.S. electricity sector, natural gas is now the predominant generation resource and its share of the generation mix is expected to continue to increase in the future.[1] While this trend has been driven primarily by low extraction costs following the shale revolution, it has also been influenced by the fact that natural gas produces far less carbon dioxide than other fossil fuels. Looking forward, gas-fired generation's ability to quickly and efficiently ramp up and down is likely to become increasingly important as investment in intermittent wind and solar generation increases. These two environmental features may enable natural gas to play a useful role in the transition to low-carbon energy. At present, however, the potential climate benefits of natural gas are being largely undermined by methane emissions from the gas supply chain.

Methane (CH_4), the principal component of uncombusted natural gas, is itself a greenhouse gas that is shorter-lived than carbon dioxide (CO_2) but vastly more potent. A small fraction of gas escaping anywhere along the supply chain, either

[1] The Energy Information Administration predicts increasing investment in gas-fired electricity generation both with and without fulfillment of the Clean Power Plan (EIA, 2018).

through equipment leaks or intentional venting, can have severe climate impacts. Currently, between 2-2.7 percent of total gas production is emitted in the United States, resulting in warming effects similar in magnitude to the warming caused by CO_2 emissions from the combustion of natural gas (Alvarez *et al.*, 2018). To what extent natural gas may be useful for addressing climate change in the future will depend on the cost of reducing these emissions.

This paper investigates these costs for the extraction segment of the natural gas supply chain, where the majority of methane emissions from the industry are generated. My empirical strategy consists of two parts. In the first part, I spatially link natural gas production facilities to geographically dispersed trading hubs to examine how methane emission rates respond to changes in wholesale gas prices. Because in this setting the pollutant is also a priced commodity, the estimated relationship between emissions and price can be directly mapped to a relationship between emissions and cost.[2] In the second part, I use this estimated relationship to simulate how production facilities' methane emissions would change following the implementation of methane pricing. I then aggregate these results to construct a sector-wide marginal abatement cost curve (MACC).

My results imply that methane emissions from natural gas production are an area of substantial low-cost opportunities for greenhouse gas abatement relative to other sectors. In particular, I estimate that imposing an emissions tax or permit price on (leaked) methane emissions equivalent to a $5 per ton carbon price would reduce methane emissions by 56 percent.[3] This represents a decrease of about 52 million

[2] In other words, if there are no market failures, profit-maximizing firms will choose an emission rate that sets the marginal cost of capturing one unit of gas equal to the marginal private benefit of being able to sell that unit of gas, i.e. the gas price.

[3] Note that accurately monitoring CH_4 emissions from the gas supply chain presents a significant challenge to successfully implementing methane pricing at this time. Unlike smokestack CO_2 emissions, fugitive CH_4 emissions are inherently difficult to measure. However, technological advancements are rapidly lowering monitoring costs and market-based instruments may still be effectively deployed under conditions of imperfect measurement (Stranlund *et al.*, 2009; Cremer & Gahvari, 2002).

tons of CO_2-equivalent emissions per year at an annual net cost of \$70 million, which is only about 0.1 percent of the wholesale value of all gas produced in the United States.[4] I further estimate that a methane price designed to fully internalize its social cost would reduce emissions by about 76 percent at an annual net cost of only \$261 million.[5] Under such a policy, the average cost per ton of CO_2-equivalent emissions abated would be about \$3.70, which is substantially lower than empirical estimates of abatement costs for many proposed and existing climate policies.

Previously, abatement costs for methane emissions have been primarily estimated using bottom-up engineering approaches.[6] While these engineering cost studies are useful, they are limited in their ability to account for opportunity costs, learning, heterogeneity in real-world conditions, and various other factors. This is well-documented for GHG abatement through investments in energy efficiency (Fowlie *et al.*, 2018; Gillingham & Palmer, 2014; Allcott & Greenstone, 2012) and carbon sequestration (Lubowski *et al.*, 2006; Stavins, 1999).

Instead of relying on engineering cost estimates, this paper relies on the condition that profit-maximizing firms equate marginal private benefits with marginal costs. This condition enables the use of spatial and temporal variation in natural gas prices to identify how much firms expend to reduce emissions.[7] By implicitly capturing the firm's decision-making process to employ the most efficient abatement measures first, this approach is able to account for all factors that are known to the firm

[4] This figure represents only physical abatement costs and sets aside questions of how tax or permit revenue might be distributed. It also accounts for firms being able to sell captured gas.

[5] Note that this is an out-of-sample prediction, as average annual gas prices range from about \$2-\$6 per thousand cubic feet (Mcf) over the study period while the social cost of methane is about \$27 per Mcf leaked. This figure is for emissions generated in 2020 assuming a 3 percent discount rate and normalized to 2018 dollars (EPA, 2016).

[6] See, for example, ICF (2016), EPA (2015), or Delhotal *et al.* (2006).

[7] I collect data on prices from S&P Global and data on methane emissions from the EPA's Greenhouse Gas Reporting Program (GHGRP). As is discussed in detail in Section 1.4, while the GHGRP is the most comprehensive dataset on methane emissions currently available, it does not provide a direct measurement of emissions, but rather an estimate based on equipment characteristics, emission factors, records of firm activity, and many other inputs. The empirical strategy used in this paper is designed to address noise and potential biases in this measure in order to make use of the signal that is available.

but not directly observed by the econometrician. This makes it particularly useful for predicting the effect of regulating methane using an emissions tax or trading program, which would similarly incentivize firms to exploit the least costly abatement opportunities first.

This work falls under a broad strand of literature in economics that uses empirical methods to estimate abatement costs.[8] Previous studies have estimated abatement costs from various existing or proposed environmental policies (Fowlie *et al.*, 2018; Meng, 2017; Anderson & Sallee, 2011) and from the deployment of specific abatement technologies (Callaway *et al.*, 2018). One particularly related example is Cullen & Mansur (2017), who use variation in natural gas prices following the shale revolution to recover a short-run CO_2 abatement cost curve for the U.S. electricity sector. This paper also contributes to an emerging economics literature on methane leakage. Focusing on the distribution sector, Hausman & Muehlenbachs (2016) quantify regulatory distortions that allow gas utilities to pass the cost of leaked gas through to their ratepayers, resulting in inefficient levels of abatement. In the production sector, Lade & Rudik (2017) study the effects of a 2015 mandate limiting at flaring at oil and gas wells in North Dakota and estimate potential efficiency gains under a counterfactual market-based regulation.

This paper is the first to empirically estimate a marginal abatement cost curve for methane emissions from natural gas production, which accounts for about 60 percent of methane emissions from the U.S. gas industry (Alvarez *et al.*, 2018).[9] I introduce a novel identification strategy that exploits the fact that the pollutant is also a priced commodity to make detailed predictions about the potential impacts of methane policy.[10] While I have applied this strategy to production, it may be similarly

[8] A separate approach, which is typically employed to estimate abatement costs of global climate policies, is the use of computational general equilibrium modeling (e.g. Morris *et al.* 2012; Klepper & Peterson 2006).

[9] Lade & Rudik (2017) construct MACCs for avoided flaring as part of their analysis. However, their approach relies on engineering cost estimates and considers just one specific abatement technology.

[10] While this is the first paper to directly leverage this feature of methane to construct a MACC, the

employed to estimate methane abatement costs for the natural gas processing and storage sectors.[11]

I proceed by providing further background on methane leakage in the next section. Section 1.3 presents a model of firms' extraction and emission decisions that provides intuition for the empirics. Section 1.4 describes data sources for emissions, production, prices, and other variables used in the analysis. Section 1.5 presents the empirical strategy used to recover the relationship between emission rates and prices. Section 1.6 presents the simulation model of methane pricing and compares the estimates in this paper to other estimates of abatement costs. Section 1.7 concludes.

1.2 Background

Methane (CH_4) accounts for about 16 percent of greenhouse gas emissions worldwide in terms of warming, making it the second most important greenhouse gas following carbon dioxide (CO_2) (IPCC, 2014). The three primary anthropogenic sources are agriculture, landfills, and the energy sector, where natural gas used for heating and electricity generation is composed of about 90 percent methane. Although there is still considerable uncertainty as to how much gas is being emitted by each sector, the EPA's Greenhouse Gas Inventory (GHGI) estimates that the energy sector is presently the largest source in the U.S. and that production is the largest source within the sector (see Figure 1).[12] Natural gas production and consumption is projected to increase substantially both within the U.S. and globally for the foreseeable future, making it important to account for methane emissions in any broad-based climate

approach is similar in spirit to a set of recent studies that have used variation in energy prices as a proxy for carbon pricing (Cullen & Mansur, 2017; Ganapati *et al.*, 2016; Aldy & Pizer, 2015).

[11] This method is not applicable to the transmission or distribution sectors, which (at time of publication) are regulated such that pipeline owners are able to pass cost of lost gas through to their customers.

[12] The GHGI is an EPA emissions monitoring project that is related to, yet distinct from, the GHGRP. While the GHGRP is focused on accurately tracking emissions for high-emitting facilities, the GHGI is focused on creating a comprehensive picture of all U.S. emissions at the industry level.

Figure 1: Methane emissions from the various components of the natural gas supply chain.

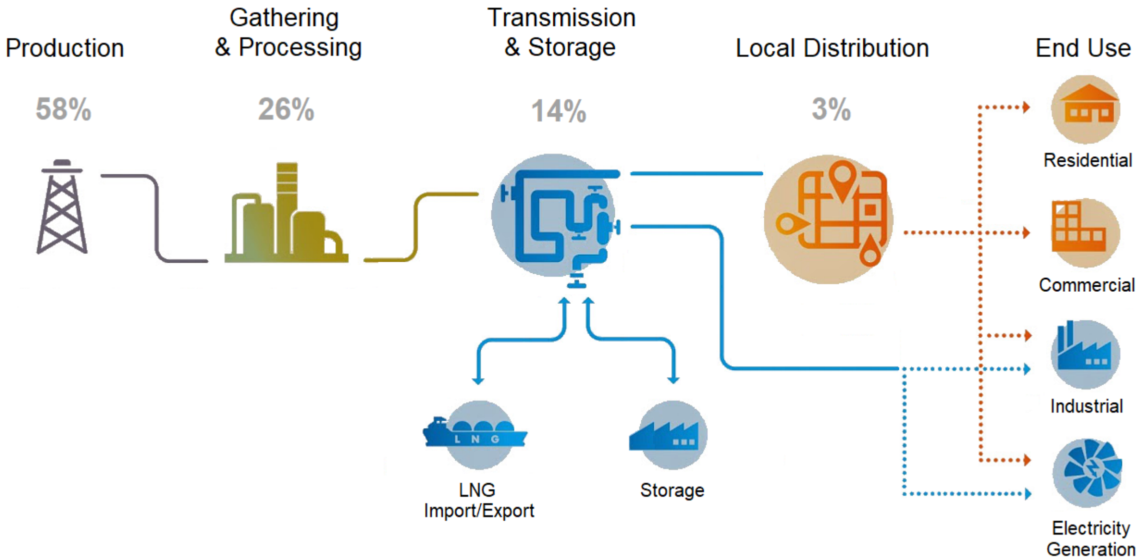

Estimates from Alvarez *et al.* (2018). Graphic adapted with permission from AEMO NGFR.

change mitigation strategy (EIA, 2018; IEA, 2017).[13]

Broadly speaking, CH_4 is unintentionally released into the atmosphere at gas production facilities through leaks in extraction, initial processing, and transmission equipment. It is also intentionally vented during certain procedures in well completions, workovers, and maintenance.[14] There is a high degree of heterogeneity in leakage rates across facilities, which is reflected both in scientific measurement studies (Sanchez & Mays, 2015; Subramanian *et al.*, 2015) and in the GHGRP, where production facilities' emission rates vary from less than .01 percent to over 10 percent. Finally, natural gas is often found alongside petroleum, in which case it may be either vented, flared, or collected and sold (if it is economical to do so) by wells that

[13] The U.S. Energy Information Administration's Energy Outlook 2018 predicts increasing domestic gas production in all seven considered price and technology scenarios, with a 50 percent increase by 2050 in their reference scenario. The International Energy Information Agency's World Energy Outlook 2017 predicts a 20 percent increase in gas production by 2030 in their Sustainable Development scenario and greater increases in other scenarios.

[14] Well completion consists of all activities between actual drilling and extraction of gas for sale, which includes installing equipment and testing, as well as hydraulic fracturing and retrieval of fluids for tight-gas reservoirs. Workovers describe major operations to repair or stimulate gas flow at existing wells.

primarily extract oil.[15]

As of now, regulations on methane emissions from oil and gas production are not well-established in the United States. In late 2016, the EPA introduced performance standards for new wells, processing plants, and compression stations. In 2018, however, the EPA's new administration proposed amendments that would greatly weaken these requirements. Also in 2016, the Bureau of Land Management (BLM) finalized a policy to require wells located on federal and tribal lands to capture high percentages of gas in place of venting and flaring on the basis of conserving federal resources. However, this policy was never implemented and its future remains uncertain. In terms of local regulations, in 2014 Colorado introduced relatively strong performance standards for new and existing wells, including equipment mandates, waste-reducing procedures during well completion, and semi-annual leak detection and repair. In 2015, North Dakota introduced regulations limiting flaring that primarily affected co-produced gas at oil wells.[16]

1.3 Theoretical Framework

This section develops a theoretical model of the production and emission decisions faced by natural gas production firms in order to motivate the empirical analysis of firms' abatement costs. I begin by deriving firms' first order conditions for leakage and abatement and proceed to demonstrate how a relationship between price and abatement costs can be mapped to a relationship between a potential emissions tax and abatement costs.

[15] Because oil and gas are so often co-located, petroleum and natural gas production facilities are not differentiated in the datasets used in this paper.

[16] The EPA regulations came into effect in August 2016. Because this policy affects all production in the United States, its impact should be picked up by time fixed effects. I control for Colorado and North Dakota regulations in the empirical analysis.

1.3.1 The Firm's Problem

Consider the profit function of a gas production firm's operations within a single basin:

$$\pi_t = P_t(Q_t - L_t) - C(Q_t, L_t, X_Q, X_L) \tag{1}$$

P_t is the price of gas in period t, Q_t is the quantity of gas the firm extracts in t, L_t is the quantity of gas it leaks, and $C(\cdot)$ is its total cost. I assume the firm is a price-taker selling into a perfectly competitive wholesale gas market.[17] Because the quantity of gas leaked depends on the amount of gas flowing through the facility's equipment, it is useful to decompose leakage into the product of extraction and a leakage rate $R_t = L_t/Q_t$:

$$\pi_t = P_t(1 - R_t)Q_t - C(Q_t, R_t, X_Q, X_R) \tag{2}$$

In this framework, the firm's problem consists of choosing how much to extract alongside choosing how careful to be to avoid leaks. This characterization makes it possible to separate $C(\cdot)$ into costs of extraction that are unrelated to the facility's leakage rate (i.e., costs of obtaining leases, capital costs for equipment gas does not pass through) and costs that determine the leakage rate (i.e., the additional up-front capital costs for equipment that emits less, labor costs for leak detection and repair). If we assume leakage-related costs are separable for each unit of extraction, we can write the firm's optimization problem as the following:

$$\pi_t = \max_{Q_t, R_t} \; P_t(1 - R_t)Q_t - C_1(Q_t) - Q_t c_2(R_t) \tag{3}$$

[17] The U.S. gas production sector has been generally viewed as competitive following deregulation in the 1980s and 1990s (Gabriel *et al.*, 2005).

Here, $C_1(\cdot)$ is the *total* cost of extraction not related to the leakage rate and $c_2(\cdot)$ is the *per-unit* cost of having leakage rate R_t. This decomposition allows costs not associated with leakage to be nonlinear in production. For example, one might imagine that the cost of acquiring new leases in a given basin increases as the firm increases production because the total number of leases is finite. On the other hand, costs associated the leakage rate are assumed to be the same regardless of the firm's level of production. For example, paying a worker to inspect one well site for leaks is assumed to cost the same amount whether the firm operates 50 wells or 5,000 wells. However, $c_2(\cdot)$ is nonlinear in R_t—in particular, it is decreasing and convex such that it approaches infinity as the leakage rate approaches zero. The convexity captures the intuition that, due to diminishing returns, bringing the leakage rate down from 5 percent to 4.5 percent will be significantly cheaper than bringing it down from 1 percent to 0.5 percent.

The firm's first-order condition for Q_t sets the marginal revenue generated by extracting one unit of gas equal to the marginal cost of extracting it:

$$P_t(1 - R_t) = \frac{\partial C_1(Q_t)}{\partial Q_t} + c_2(R_t) \tag{4}$$

Note that the firm's marginal revenue for one unit of extraction is lower than just the gas price, as only the portion that is not leaked may be sold. In the firm's first-order condition for R_t, however, the firm's marginal revenue of avoiding one unit of leakage is simply the gas price, since the whole unit may be sold:

$$P_t = -\frac{\partial c_2(R_t)}{\partial R_t} \tag{5}$$

Equation (5) forms the basis for the empirics: When maximizing profits, the firm chooses a leakage rate that sets the price equal to their marginal cost of leakage abatement.[18,19] Intuitively, if one unit of gas can be sold for P_t, the firm will be willing to expend up to P_t to prevent it from being lost.

1.3.2 Addition of an Emissions Tax

The implementation of a tax on methane emissions adds another term to the firm's profit function as follows:[20]

$$\pi_t = \max_{Q_t, R_t} \; P_t(1 - R_t)Q_t - C_1(Q_t) - Q_t c_2(R_t) - Q_t R_t T \tag{6}$$

Here, in addition to costs associated with extraction and costs associated with preventing leakage, the firm must pay $\$T$ for each unit of methane emitted.[21] The first-order conditions for the optimal emissions rate now simplifies to:

$$P_t + T = -\frac{\partial c_2(R_t)}{\partial R_t} \tag{7}$$

Equation (7) illustrates that the firm now chooses a leakage rate that sets its marginal cost of preventing one unit of gas from escaping equal to the commodity value of that unit of gas plus the avoided emissions tax. This implies that an emissions tax on CH_4 would have the same effect on fugitive emissions as a change in the price of gas of the same amount, which makes it possible to use an estimated relationship between leakage and prices to predict how leakage would respond to the

[18] Note that $-\frac{\partial c_2(R_t)}{\partial R_t}$ is positive because c_2 is decreasing in R_t.

[19] The one-period framework presented here is useful for setting up a tractable empirical model, but it oversimplifies some important temporal aspects of the firm's true decision making process. In Appendix A.1.1, I extend this framework to a dynamic model and discuss empirical applications that may become possible with more detailed emissions data, should it become available.

[20] In this section as well as in the empirical analysis I consider a hypothetical emissions tax; however, results are also applicable to permit prices under an emissions trading approach. Discussion of whether one instrument may be more appropriate than the other for regulating methane is beyond the scope of this paper.

[21] For simplicity, the theoretical model assumes extracted gas is 100 percent methane. I account for the methane content of extracted gas when simulating the effect of a methane tax in Section 1.6.

implementation of an emissions tax.[22]

1.4 Data

The EPA's Greenhouse Gas Reporting Program provides an annual measure of fugitive methane emissions for nearly 700 onshore gas production facilities during the period from 2011 through 2016. Facilities are delineated at the firm-basin level, meanings most facilities include hundreds or thousands of wells. Emissions from all equipment at all wells operated by a firm within a single basin along with all of the firm's completion and well maintenance activity of the firm within that basin are aggregated into a facility-level estimate. Most of the variables used to construct the facility-level emissions estimate are also reported (at the facility level), including specific emissions from various types of equipment and procedures, equipment counts, and levels of extraction.[23] These data are collected through a comprehensive survey that is mandatory for all U.S. facilities producing at least 25,000 tons of CO_2-equivalent GHG emissions (tCO_2e) annually.[24]

In contrast to emissions from fuel combustion, which firms are generally required to report to the GHGRP using continuous emissions monitoring sensors placed in smokestacks, fugitive methane is not measured directly. Instead, the GHGRP provides firms with a framework for calculating these emissions using equipment characteristics and emissions factors (either type-specific or estimated averages) in combi-

[22] The emissions tax will have some impact on the firm's production decision as well, but to a much lesser degree. The firm's first order condition for Q_t with an emissions tax is $P_t(1 - R_t) - R_tT = \frac{\partial C_1(Q_t)}{\partial Q_t} + c_2(R_t)$. R_t is generally very low (the average emission rate for the quality-trimmed sample is just 1.1 percent) and in fact will decrease further as the firm decreases leakage in response to the emissions tax, so the impact of an emissions tax on production will be far smaller than the impact of a price increase of the same amount.

[23] Unfortunately, many useful supplementary variables (e.g. gas production, oil production, well IDs) are only available for 2015 and 2016.

[24] This includes a large number of gas processing plants, compression stations, and storage sites, as well as power plants, factories, refineries, landfills, and other types of facilities that are not part of this analysis.

nation with records of throughput, maintenance, installation, etc. For some devices, firms are also instructed to test for leaks around individual pieces of equipment. The firm is also required to report venting and flaring activity associated with well completions and workovers.[25]

I use the GHGRP because it is the most comprehensive and consistent source of panel data on methane emissions from natural gas currently available for this analysis. However, it should be noted that a number of scientific studies have shown it to be relatively noisy and subject to some biases.[26] The source of bias that has the biggest implications for this analysis is that GHGRP methodology fails to effectively capture all sources of methane emissions, meaning that both total emissions and total abatement will be understated in this paper. One approach would be to scale emissions up using the best available scientific estimates for actual emissions. However, rather than introducing the additional assumptions necessary to do so, I elect to exclude underrepresented emission sources and present results in terms of emissions as reported to the GHGRP.[27]

Although this overall downward bias in facility-level emissions estimates is carried through to the empirical analysis, particular biases relevant to the responsiveness of firms' emitting behaviors to prices are not problematic as long as they are not systematically correlated with unobserved determinants of natural gas prices. Examples

[25] See Table A8 in the Appendix for a partial list of factors that enter GHGRP CH_4 emission calculations that are directly dependent on firm abatement decisions.

[26] For example, aerial surveys of the Denver-Julesburg basin and Barnett shale regions have detected CH_4 emissions from production sites about three times greater than those reported to the GHGRP (Lyon *et al.*, 2015; Petron *et al.*, 2014). Subramanian *et al.* (2015) perform a bottom-up study of compression stations that estimates similar levels of underreporting due to "super emitting" facilities with severe leaks, the usage of incorrect emissions factors, and the failure to account for some sporadic emissions sources. Lastly, a recent meta-analysis of many bottom-up and top-down studies concluded that actual CH_4 emissions from the U.S. gas industry are about 60 percent greater than those estimated by the EPA's Greenhouse Gas Inventory (Alvarez *et al.*, 2018).

[27] The scientific literature shows that emissions related to firm decisions about maintenance (i.e. equipment leaks) are particularly poorly captured by the GHGRP. In order to linearly scale estimated abatement up to incorporate these emissions as well, it would be necessary to assume that emissions related to maintenance respond in the same way to prices as emissions related to the other two categories.

Figure 2: GHGRP onshore gas production facilities and natural gas trading hubs.

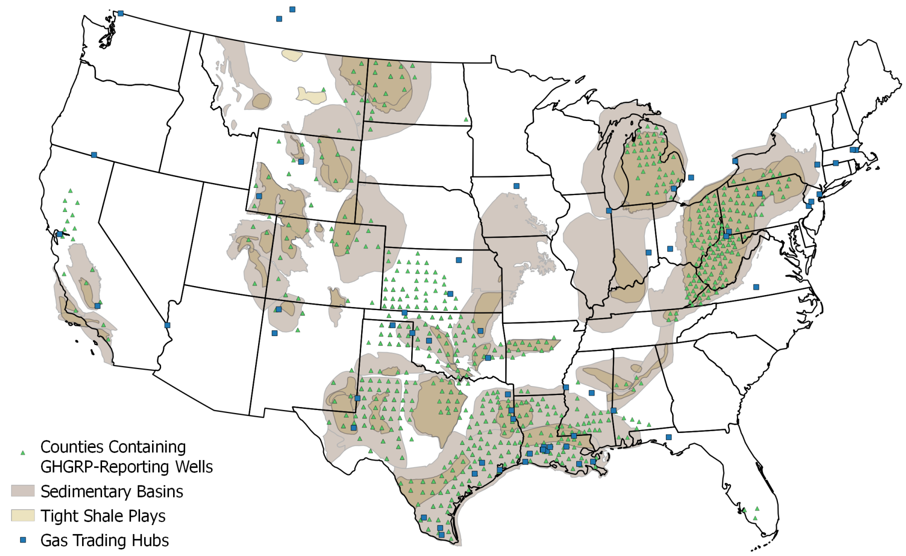

Production facilities are delineated at the firm-basin level in the GHGRP. Each green triangles marks a county containing at least one well that is part of a GHGRP facility, which in most cases means many wells associated with many different facilities. Basin boundaries are sourced from the Energy Information Administration (these boundaries are for illustrative purposes only and are not used in the analysis).

of these biases include differences across firms in levels of effort put toward accurate reporting, changes to the GHGRP methodology over time, and differences in firms' beliefs about the effectiveness of various abatement activities.[28] Such potential biases are netted out by facility and fixed effects, making it possible to accurately recover the abatement behaviors that are effectively captured by the GHGRP.

I collect data on facility-level gas and oil production through DrillingInfo (DI), an industry data provider that collects and digitizes government records of well and permit filings in near real-time. Through DrillingInfo, I am able to observe extraction activity at a daily level for the vast majority of wells in the United States. Because

[28] To elaborate, when a firm purchases higher quality equipment in response to a price signal, its decision is based on its belief about how much additional gas the new unit will recover. This may differ from the equipment's actual abatement potential or from GHGRP emission factors. If, for example, the firm believes a purchased device's actual emission factor is lower than the factor used in the GHGRP, and the firm's beliefs are closer to reality, the sensitivity of that firm's emission rates to prices will be understated. However, if these beliefs are consistent within firms over time or are updated for all firms in a region in ways that are not correlated with prices, they will be picked up by fixed effects.

Table 1: Summary statistics for the full GHGRP sample and the quality-trimmed sample.

		Full Sample		Trimmed Sample	
	Source	**Mean**	**SD**	**Mean**	**SD**
CH_4 Emissions Rate	GHGRP & DI	0.3894	4.0953	0.0108	0.0152
CH_4 Emitted (MMcf)	GHGRP	217	518	266	389
From Completions	GHGRP	29	169	34	134
From Equipment	GHGRP	117	276	143	222
From Maintenance	GHGRP	49	110	58	116
Gas Production (MMcf)	DrillingInfo	57,729	164,731	63,436	98,459
Oil Production (Mbbl)	DrillingInfo	4,199	10,854	4,523	10,992
Wells Per Facility	DrillingInfo	797	1,409	879	1,489
Completions	DrillingInfo	35	73	47	90
Gas Price ($/Mcf)	S&P	3.23	0.83	3.20	0.85
Number of Facilities		683		222	
Total Observations		2,980		1,150	

Mcf ≡ Thousand cubic feet; MMcf ≡ Million cubic feet; Mbbl ≡ Thousand barrels

production facilities in the GHGRP are delineated at the firm-basin level (i.e. all of the drilling, extracting, and initial processing equipment used by one firm within one basin is considered to be a single facility), I link the two datasets by aggregating wells in DrillingInfo to the firm-basin level (see Figure 2). Firm names are not always consistent across the two datasets and asset sales are common in the oil and gas industry, both of which present potential sources of error in manually matching the two datasets. I ensure the quality of matches by removing facilities that differ in production in excess of 25 percent between the GHGRP and DrillingInfo, using the variable for production that is reported to the GHGRP in 2015 and 2016 but not in earlier years.[29]

Facility i's methane emission rate R_{it} is constructed by dividing i's total methane emissions in year t by its total gas production in t. To reduce the potential influence

[29] This implies that facilities that stopped reporting before 2015 are excluded from the analysis. Figure A3 shows the distribution of how well facilities match on production and the 25 percent cutoff—which is admittedly somewhat arbitrary, though results are robust to using other thresholds.

Figure 3: Variation in natural gas spot prices (each line represents one facility).

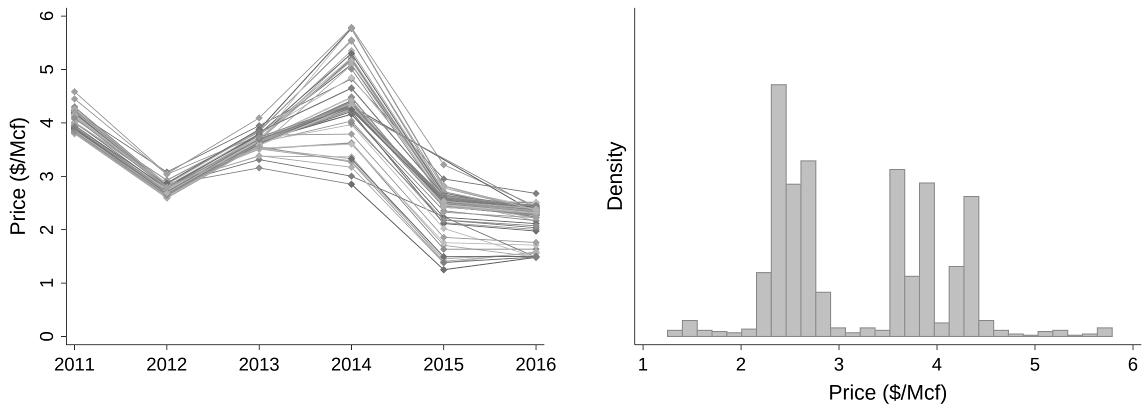

of inaccurate reporting, I further trim the 5 percent of outliers in leakage rates on either end (10 percent total). The 223 facilities that remain in the trimmed sample tend to be slightly larger and perform slightly more completions on average, but are otherwise representative of the full sample (see Table 1). On average, each facility has about 900 wells, though there is a substantial degree of variation in facility size, ranging from only a handful of wells to over 10,000. Gas production, oil production, and CH_4 emissions are similarly highly heterogeneous across facilities. The average emission rate for the trimmed sample is 1.08 percent.

I collect spot natural gas prices from S&P Global Market Intelligence. Natural gas is traded at "hubs" that are geographically dispersed across the United States, which sometimes correspond to specific points where many interstate pipelines intersect, but more often actually represent an aggregation of all transactions along certain sections of one or more pipelines. For simplicity, I use the centerpoint of hubs that consist of stretches of pipelines, which are geocoded from S&P's energy mapping interface. Spot prices are available at 96 hubs for the six-year period for which GHGRP data are available. I link GHGRP facilities with hubs by taking a weighted average of the prices at hubs closest to the centroids of the counties that the facility operates in (see Figure 2). For example, if a facility operates in 3 counties closest to hub A and 2 counties closest to hub B, the price for that facility would be $\frac{3}{5}P_A + \frac{2}{5}P_B$.

Gas prices are spatially correlated, as gas moves continuously through a nationwide network of pipelines, but this correlation diminishes with distance due to transportation costs and transmission constraints. Accordingly, prices at two hubs close to one another will usually be highly correlated, while prices at hubs located across the country from one another will be much more divergent. As shown in Figure 3, there is considerably more variation in prices in the last three years of the study period than there is in the first three years, which may be in part due to binding transmission constraints during the particularly cold winter of 2014-15.

1.5 Empirical Framework and Results

In this section, I estimate the relationship between price and emission rates at gas production facilities. I exploit temporal and spatial variation in gas prices, control for a wide array of potential sources of endogeneity with facility and region-by-year fixed effects, and employ a second-order fractional polynomial (FP) model to capture nonlinearities. However, I also demonstrate that my results are robust to a variety of more restrictive and more flexible models. I additionally explore the mechanisms by which firms reduce emissions in response to higher gas prices, including equipment upgrades, avoiding waste during completions, and leak detection and repair.

1.5.1 Fractional Polynomial Regression

To account for potential nonlinearity, I estimate the relationship between firms' emission rates and gas prices as a second-order fractional polynomial model:[30]

$$R_{it} = \beta_0 + \beta_1 P_{it}^{A} + \beta_2 P_{it}^{B} + \mathbf{X}_{it}\psi + \gamma_i + \lambda_{rt} + \varepsilon_{it} \tag{8}$$

R_{it} is the facility i's emission rate in year t and P_{it} is the price of gas it faces. The powers A and B are determined by the data from a set of predefined possibilities as the parameters that provide the best fit under maximum likelihood estimation.[31] Time-varying controls $\mathbf{X}_{it}$ include oil extraction, completions, number of wells, and indicators for whether the majority of the facility's wells were located in Colorado after 2014 or North Dakota after 2015.[32] Facility fixed effects γ_i net out facility-specific determinants of emissions and potential biases in GHGRP reporting that are persistent within a facility over time. Region-by-year fixed effects λ_{rt} net out potential biases that are consistent across facilities within a particular region and year, such as regional economic shocks that could affect both prices and behaviors associated with emissions.[33] Region-by-year effects also control for changes to the GHGRP methodology over time that affect all facilities. I weight observations by

[30] Fractional polynomial models are an extension of traditional polynomial models that allow a more diverse set of transformations of the independent variable of interest. They overcome a number of limitations, such as oversensitivity to tails of the data, while still maintaining the desirable characteristics of linear regression, such as ease of incorporating fixed effects (Royston & Altman, 1994).

[31] The fractional polynomial methodology requires separately estimating specifications for all possible combinations of A and B to determine the best fit. Following standard practice in the literature, I use -2, -1, -.5, .5, 1, 2, and 3 as the possible values of A and B, as well as the natural log (i.e. $\log(P_{it})$ in place of P_{it}^{A}) (Sauerbrei *et al.*, 2006). In the second-order model—meaning two transformations of P—this implies 44 potential models. Each of these 44 models is separately estimated using maximum likelihood estimation, and only results from the model that provides the best fit in terms of having the highest likelihood are reported.

[32] These two fixed effects control for the impact of methane regulations introduced in those states. A robustness check excluding the Mountain region altogether produces results that are highly similar in character but less precisely estimated (see Appendix A.1.4).

[33] Regions follow the U.S. Energy Information Agency's five natural gas storage regions (Pacific, Mountain, Midwest, South Central, and East).

Figure 4: Relationship between emission rates and prices estimated using a second-order fractional polynomial (left) and comparison with alternative specifications (right).

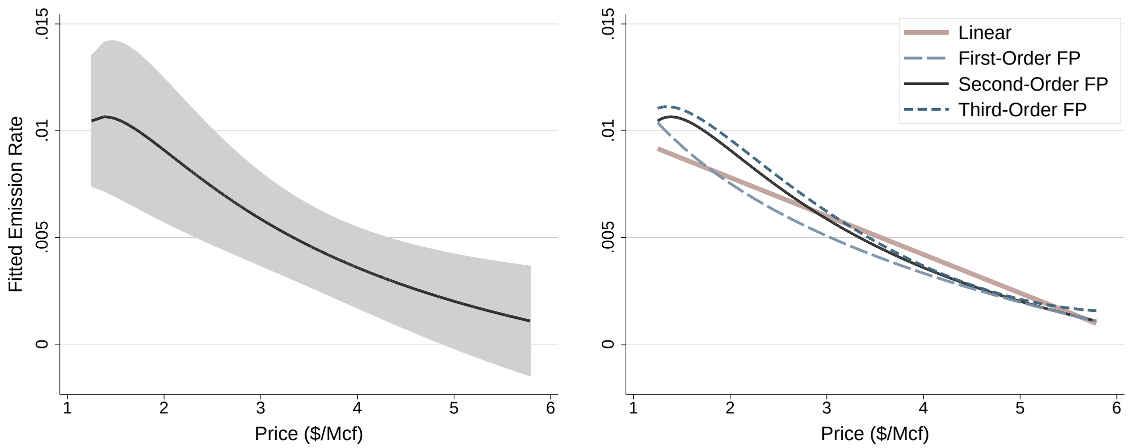

facilities' average gas production over the study period.[34]

The predicted relationship between price and emission rates from the second-order fractional polynomial model is presented in the left panel of Figure 4. It is evident that at nearly all gas prices there exists a downward-sloping relationship between price and emission rates. Conditional on the included fixed effects, production at low gas prices is predicted to leak at about 1 percent and production at the highest average annual gas prices observed during the study period is predicted leak at about .15 percent. The apparent convexity of this relationship is consistent with diminishing returns to abatement activities—i.e., that those facilities facing generally higher prices have already exploited the cheapest abatement opportunities.[35]

Results are similar across a range of more restrictive and more flexible models. A comparison with first- and third-order fractional polynomials is presented in the left

[34] This makes the estimated curve representative of the effect of price on the emission rate of an average unit of gas production, rather than on the emission rate of an average facility, which is preferable for constructing results for the sector in aggregate.

[35] Below about $2/Mcf, the curve becomes concave and even changes sign at the very lowest prices (though this change in slope is not significant). Though it is possible there exists some unknown phenomenon that generates a positive relationship at exceptionally low gas prices, limited support over this range and the fact that this upward-sloping segment disappears in many alternative specifications suggests it is likely to be spurious.

Table 2: Relationship between natural gas spot price and CH_4 emission rate estimated using linear and fractional polynomial models.

Model	(1) Linear	(2) 1st-Order FP	**(3) 2nd-Order FP**	(4) 3rd-Order FP
P_{it}	-0.0018*** (0.0006)			
$\log(P_{it})$		-0.0061*** (0.0017)		
$P_{it}^{-0.5}$				0.0493*** (0.0168)
P_{it}^{-1}			0.0460*** (0.0154)	
P_{it}^{-2}			-0.0319*** (0.0123)	-0.0202** (0.0085)
P_{it}^{3}				0.00001 (0.00001)
Facility FE	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes
N	1,150	1,150	1,150	1,150
adj. R^2	0.632	0.633	0.633	0.632

Standard errors in parentheses (clustered at the parent firm level with 146 firms)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

panel of Figure 4. All three demonstrate a downward-sloping, convex relationship between price and emissions. The third-order FP, which tests 164 possible functional forms, is nearly identical to the second except with slightly greater convexity. Coefficient estimates for all three models and a linear specification are reported in Table 2. For the fractional polynomial specifications, the transformations of P_{it} that best fit the data are identified by the presence of coefficient estimates—for example, the best fit for the second-order model is $R_{it} = \beta_0 + \beta_1 P^{-1} + \beta_2 P^{-2}$. The best fit for the first-order model is simply a log transformation of price. In this model, a 1 percent increase in price is associated with a 0.006 percentage point decrease in emission rates. Scaling this result up, a 30 percent increase in price—about \$1 for an average facility—would be associated with a 0.18 percentage point decrease in emission rates (i.e. from 1 percent to .82 percent).

As shown in Figure A4 in the Appendix, the existence of a downward-sloping relationship over the range for which there is substantial variation in price (about $2-$4.50) persists across many alternative specifications. These include removing weights, trimming emission rates at the 1 percent level instead of at the 5 percent level, and using basin-by-year fixed effects in place of region-by-year fixed effects. As shown in Table A3, wide confidence intervals for the specification with basin-by-year fixed effects are driven by the constant term rather than the coefficients on price, which are precisely estimated. However, a model using only year fixed effects in place of region-by-year effects does not generate meaningful results, indicating the existence of important regional trends that obscure the effect of price on emissions. I additionally find that the existence of a negative relationship between emissions and price is robust to the application of a negative binomial model, which specifically addresses the potential failure of the assumption of normally-distributed errors that may arise when OLS is used in a setting where the dependent variables is a rate. The methodology and results for this model are presented in Appendix A.1.2.

I investigate two other potential threats to identification using an instrumental variables (IV) approach. Although the included fixed effects control for possible omitted variables that are constant within facilities over time or across facilities within a particular region and year, a valid instrument for price would eliminate the impact of possible omitted variables that vary at the facility-year level. Furthermore, isolating variation from demand-side price shocks would ensure there is no chance of reverse causality.[36] I therefore explore using various weather variables to instrument for price. Results, presented in Appendix A.1.3, are broadly similar to the results from

[36] In other words, lower emission rates caused by some other exogenous force could decrease prices by increasing the amount of natural gas available for sales. This effect would attenuate my results, as it would imply a positive correlation between emission rates and prices.

non-instrumented specifications, but not statistically significant.[37]

1.5.2 Abatement Mechanisms

To assess whether the aggregate results presented in the previous section are indeed driven by firms adjusting their abatement behaviors in response to changes in prices, I examine a subset of variables that compose the GHGRP's facility-level emissions estimate. In particular, I test whether price changes predict the installation or removal of four types of equipment that are straightforward to measure and known to have high abatement potential, as well as two measures of gas conservation during hydraulic fracturing completions.[38]

The first type of equipment considered is pneumatic pumps, which are used at some wells for injecting chemicals that encourage the flow of natural gas or oil. "Pneumatic" in this context means the pumps rely solely on pressure from gas exiting the well for power, and they are designed continuously emit or "bleed" some fraction of this power gas. Pneumatic pumps can be replaced by electric units that have a higher up-front capital cost but near-zero emissions, so ex-ante one would expect higher prices to predict fewer pneumatic-type pumps. Next are pneumatic controllers, which regulate the flow of gas through equipment or connections. The GHGRP classifies high-bleed and low-bleed controllers based on whether they emit more or less than 6 Scf per day. While some purposes require high-bleed devices, the majority of high-bleed devices can potentially be replaced with more costly low-bleed devices or zero-bleed devices that are powered by electricity rather than gas (McCabe *et al.*,

[37] Inconclusive results from the IV model are at least partly due to limited statistical power. The limiting dataset is the GHGRP, which reports emissions at a facility-year level (production and price variables vary at a daily level). This approach may become viable in the future if satellite methane emissions data with sufficient temporal and spatial resolution becomes available.

[38] Although data on about a dozen equipment types are available in the GHGRP, pneumatic controllers and pumps are well-suited for regression analysis because their contributions to aggregate emissions (as reported in the GHGRP) are based solely on equipment counts and operating times, with a uniform emissions factor.

Table 3: The relationship between price and equipment counts for four types of emitting devices (Columns 1-4) and two measures of firms' activities to avoid leakage during well completions and workovers (Columns 5 and 6).

	(1)	(2)	(3)	(4)	(5)	(6)
	Pneumatic Pumps	Intermittent Pneumatic Controllers	High-Bleed Pneumatic Controllers	Low-Bleed Pneumatic Controllers	Venting Days	Mcf of Gas Recovered For Sales
P_{it}	-212.5**	-692.9***	-9.8	3.9	-6.7	67,064,000
	(93.9)	(260.3)	(284.8)	(20.5)	(5.9)	(71,686,000)
Facility FE	Yes	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	737	1,055	1,055	1,055	716	716

Standard errors in parentheses (clustered at the parent firm level with 146 firms)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

2015).[39] Additionally, another class of controllers only releases emissions intermittently. Intermittent-bleed controllers are much more heterogeneous in emission rates, but they can also often be replaced by low-bleed or zero-bleed devices.

Another substantial source of methane emissions is flaring and venting gas into the atmosphere during well completions and workovers. Up to 2014, the GHGRP required firms to report the number of days gas was vented into the atmosphere for each completion or workover, as well as the quantity of gas (if any) that was captured for sales.[40] Although changes in gas prices should not affect firms' decision between venting or flaring gas, higher gas prices will incentivize firms to capture gas for sales rather than either flare or vent it.

I separately estimate linear regressions of each of these variables using the same set of independent variables as before.[41] Results, presented in table Table 3, are consis-

[39] There is no clear ex-ante prediction for low-bleed devices, as higher prices may cause firms to switch from high-bleed to low-bleed devices and/or cause firms to switch from low-bleed to zero-bleed devices.

[40] In 2015, the GHGRP changed its methodology to allow two potential equations for firms to calculate emissions from completions and workovers. The new methodologies likely improved the quality of measurement, but are considerably less straightforward to analyze.

[41] For consistency, these regressions use the same trimmed sample as above. However, because these variables are in levels rather than rates and are thus not reliant on matching with DrillingInfo

tent with firms adjusting their emitting behaviors in response to price in most cases. In particular, I find that higher prices predict fewer intermittent-bleed controllers, fewer gas-driven pneumatic pumps, fewer days on which gas from completions or workovers was vented, and more gas from these operations being recovered for sales (though only the former two are statistically significant).[42] There is no evidence that counts of high-bleed devices or low-bleed devices are affected by price.[43] Although results for these GHGRP microdata variables are mixed, they are consistent with identifying a stronger result for specifications using the facility-level emissions estimate. The aggregation of many inputs captures more signal than can be recovered from any individual component while weakening the influence random noise caused by reporting errors.

1.5.3 Emissions by Source

Moving up one level in the GHGRP microdata to CH_4 emissions from various source categories enables further exploration of which behaviors drive the curve estimated above. Rather than estimate 15 separate regressions for each of the 15 separately-reported sources, I group sources into three broad categories: Emissions resulting from equipment purchase decisions, emissions from completions and workovers, and emissions associated with leak detection and repair. For example, in addition to pneumatic controllers and pumps, the equipment category includes emissions from dehydrators (which vary in components, dimensions, and input chemicals), and storage tanks (which may or may not use vapor recovery apparatus).[44] Emissions from

data, it is possible to estimate them using entire GHGRP sample. The results of this robustness check are presented in Table A4 in the Appendix.

[42] Smaller sample sizes for pneumatic devices and pumps are due to missing data and smaller sample sizes for completion variables are due to missing data and the fact that the GHGRP stopped recording these variables in 2016.

[43] The expected impact of a higher opportunity cost for lost gas on low-bleed pneumatic devices is ambiguous ex-ante, as they may either be used to replace high-bleed or intermittent-bleed devices or themselves replaced with zero-bleed devices.

[44] A full account of which variables compose each category is provided in Table A8 in the Appendix.

Figure 5: Relationship between emission rates and price by emission source.

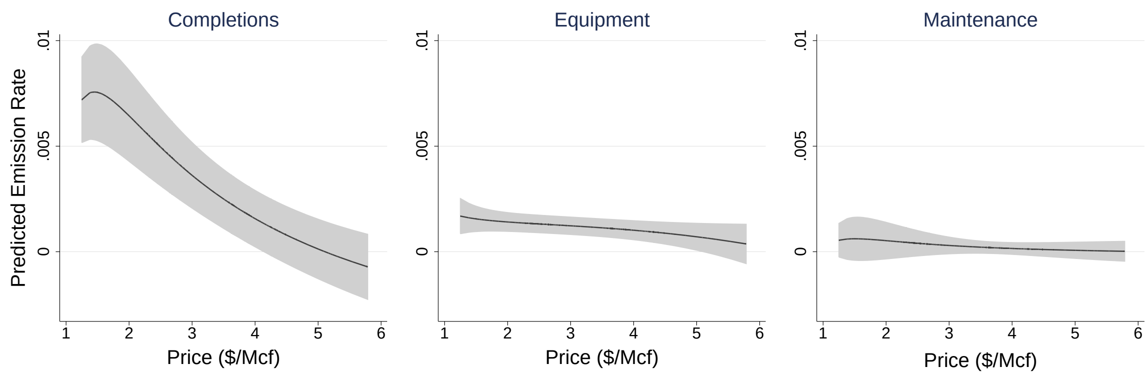

sources that do not directly involve any firm decisions about emitting behavior, such as combustion CH_4 emissions, are excluded from this portion of the analysis.

I separately estimate the relationship between price and emission rates for each of the three source categories using the same second-order fractional polynomial model as before. Results, presented in Figure 5, show that the responsiveness of emissions rates to price detected above is driven primarily by emissions from well completion and to a lesser extent by emissions associated with equipment purchase decisions. Although there may be many reasons for this result, it is likely that timing plays a large role. In a given year, a facility's emissions from completions derive from decisions about how careful to be to avoid wasting gas when completing wells that year. In contrast, emissions related to the type of equipment installed at a facility derive from decisions made in previous years as well as decisions made the same year. Furthermore, past and present equipment choice decisions are made considering expectations of future prices as well as the current spot price.[45] Although there is insufficient power to separately identify the effect of lagged and forward prices on a facility's equipment emissions, it is plausible they are decreasing in these prices as well, making the estimate for the sensitivity of overall emissions to price a conservative

[45] I formalize these conditions in Appendix A.1.1, which extends the theory section of this paper to a dynamic framework.

one.[46]

Emissions from leak detection and maintenance do not appear to be responsive to changes in the natural gas price.[47] However, it is important to note that leaks from equipment failure are most difficult to measure, making it likely this result is driven by the GHGRP methodology being less effective in detecting emissions reductions through improved maintenance. If this is the case, it would be another avenue by which my estimates of the sensitivity of overall leakage rates to price are conservative.

1.6 Predicting the Effect of an Emissions Tax

This section builds upon the results of Section 1.5 by using a straightforward simulation model to predict the effect of a tax on methane. Starting facilities at their average emission rates and prices faced over the study period, I incrementally increase prices and adjust facilities' emission rates following the slope of the estimated curve. I calculate emissions reductions and costs as prices increase, then aggregate these values to construct a marginal abatement cost curve for the sector. I examine abatement costs and benefits at a subset of policy-relevant methane prices and demonstrate that these results are robust to a variety of alternative model selection choices. Finally, I conclude the section by comparing these predictions to engineering estimates of abatement costs for methane emissions and to estimates of abatement costs in GHG-emitting sectors.

[46] By excluding lagged prices and forward prices, the reduced-form framework used in this paper is an oversimplification of the firm's true decision-making process. As the only measure of price included on the right-hand side, the spot price serves as a proxy for past prices, expectations of future prices, and past expectations of future prices (as well as current prices). A dynamic model would be necessary to separately identify the effects of these different price measures. Unfortunately, the facility-year delineation of the GHGRP affords limited statistical power for including additional measures of price as explanatory variables.

[47] I also investigate 1-year lagged maintenance emissions under the hypothesis that leaks may not be detected and reported until the following year, which also produces a null result.

1.6.1 Simulation Model

The core of the simulation model consists of increasing the effective prices faced by facilities and decreasing their emission rates based on the slope of the curve estimated in the previous section.[48] Section 1.3 illustrates that the effect of higher prices on facilities' emission rates directly maps to the effect of a tax, as both increase the opportunity cost of lost gas in the same way. For simplicity, the theoretical framework presents this mapping as 1-to-1. In practice, however, a properly-implemented tax would only affect the methane content of the emitted gas. In this simulation, I assume all facilities extracted gas is 83 percent methane (the average for facilities in the GHGRP sample), such that a \$1 tax will decrease facilities emission rates by the same amount as would a \$0.83 price increase.[49]

As a reasonable baseline, the simulation starts facilities at their average values for emission rates and prices over the study period. A tax is then applied and increased incrementally in discrete steps of Δ_{T} up to \$32/Mcf, which roughly corresponds to a \$50/ton tax on CO_2.[50] Each step k increases facilities' opportunity cost of lost gas (denoted ρ below) by Δ_{T} times the methane content of the extracted gas (denoted μ). With $\bar{P}_i$ as facility i's baseline price, the opportunity cost of lost gas facility i faces in step k is then $\rho_{ik} = \bar{P}_i + \mu\Delta_{\mathrm{T}}k$. Facility i's emission rate in step k evolves according to the first derivative of the estimated second-order fractional polynomial

[48] I choose the second-order FP as my preferred specification primarily because it produces the most reasonable curve for out-of-sample predictions. Although the second- and third-order fractional polynomials produce highly similar curves over the data's support for gas prices, in the third-order model the cubic term dominates at higher prices, leading to an upward-sloping segment that is implausible in reality. As a robustness check, I run the model using the first-order fractional polynomial curve, which predicts slightly greater abatement but at a slightly higher cost than the second-order model.

[49] In the context of describing methane emissions across the gas supply chain, Section 1.2 states that natural gas is composed of about 90 percent methane. That figure refers to "pipeline quality" gas, which has been processed to remove impurities and heavier gaseous hydrocarbons.

[50] For computational tractability, I use $\Delta_{\mathrm{T}} = \$0.05$. Results are not sensitive to choice of step size below \$1.

Figure 6: Predicted change in facilities' emission rates as an emissions tax is implemented.

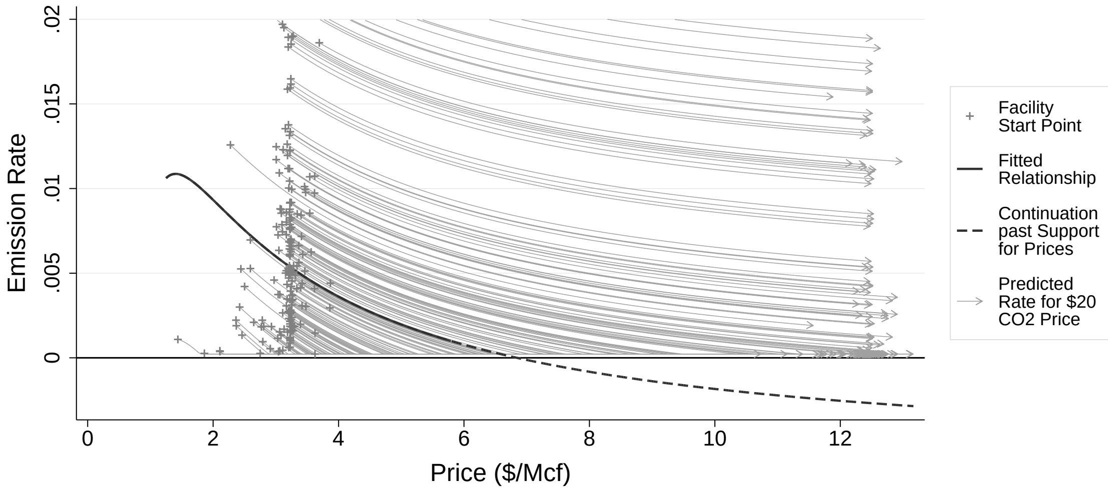

Each facility is assumed to start at its average emission rate and average price faced over the study period, indicated by a +. As a tax on CH_4 is applied and increased, facilities decrease emission rates following the slope of the estimated relationship between emission rates and prices. The dotted line shows the continuation of the curve past the support of variation in prices (i.e. the range past which estimates become out-of-sample predictions). A tax on methane corresponding to a $20/ton tax on CO_2 is illustrated here. Note that emission rates are censored above 2 percent for readability.

fit:

$$R_{ik} \;=\; R_{ik-1} + \mu\Delta_{\mathrm{T}} R'(\rho_{ik}) \;=\; R_{ik-1} + \mu\Delta_{\mathrm{T}}(-\beta_1\rho_{ik}^{-2} - \beta_2\rho_{ik}^{-3}) \tag{9}$$

β_1 and β_2 are the estimated regression coefficients from Column 3 of Table 2. Rather than assume facilities can achieve zero emissions, I lower-bound emission rates at the lowest observed average emission rate among facilities in the trimmed sample (0.0223 percent).[51] Figure 6 illustrates this process.

Emissions reductions are recovered for each facility at each step as change in the facility's emission rate times the its initial level of gas production. Because the quality-trimmed sample of GHGRP facilities accounts for only about 40 percent of total gas production in the United States, I scale production up in order to make

[51] Results are robust to lower-bounding facilities' emission rates at 0.1 percent. About one-tenth of facilities in the trimmed sample have average emission rates below 0.1 percent.

Figure 7: Marginal abatement cost curve for methane emissions from natural gas production.

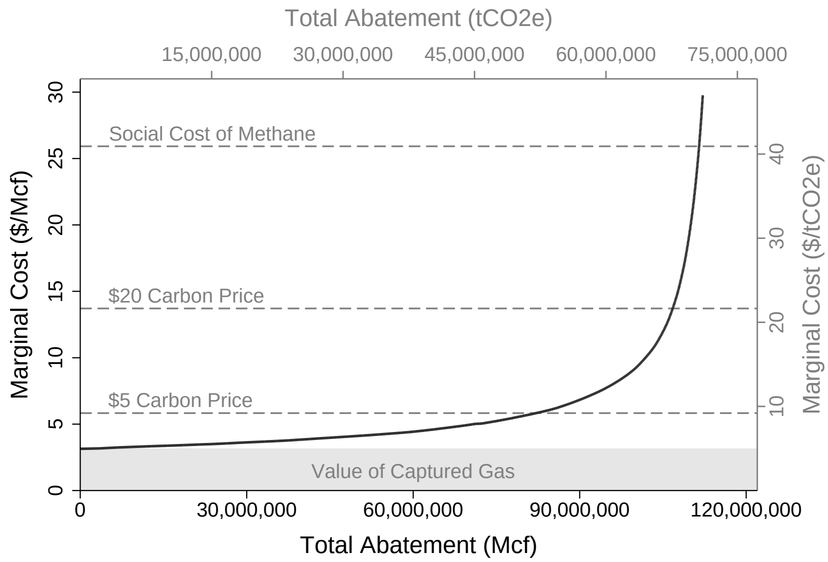

Note that carbon price policies do not directly correspond to marginal abatement costs because a.) firms expend about \$5/tCO_2e (just over \$3/Mcf) to capture gas in the absence of policy and b.) carbon price policies only affect the methane content of extracted gas.

the estimated abatement cost curve reflective of a sector-wide emissions tax. To appropriately capture heterogeneity across facilities in leakage rates and prices (which are correlated with facility size), I proportionally increase facilities' production before running the simulation.[52] With $\bar{Q}_i$ denoting facility i's scaled baseline production, total abatement A at step K is calculated as:

$$A_K = \sum_{k=1}^{K} \sum_{i} \bar{Q}_i (R_{ik} - R_{ik-1}) \tag{10}$$

Here, A_K is equivalent to predicted abatement under a methane tax of $\Delta_{\mathrm{T}} k$/Mcf. Marginal abatement cost at each step is the abatement-weighted average of ρ_{ik}.[53] Plotting total abatement against marginal abatement costs produces the marginal

[52] Specifically, I multiply each facility's production by the ratio of the EIA estimate for gross gas production in the United States in 2016 (32,635 Bcf) to total gas production from the trimmed sample (13,012 Bcf).

[53] It is also possible to recover marginal cost at each step as the change in total cost divided by the change in total abatement.

abatement cost curve shown in Figure 7. To facilitate comparison with other polluting sectors, I convert these variables to tons of CO_2-equivalent emissions on the alternate axes.[54] In general, the curve demonstrates that methane emissions from natural gas production are an area with substantial low-cost opportunities for greenhouse gas mitigation. Total CH_4 emissions from the natural gas production (as estimated by the GHGRP methodology and scaled up to include all U.S. production) are about 147,000,000 Mcf, meaning the majority of emissions from the sector can be abated. While the cost of realizing these reductions depends on the target level of abatement, it evident that a large portion of these reductions can be achieved at very low cost.

Point estimates and bootstrapped standard errors for three selected policy-relevant tax levels are presented in Table 4.[55] I estimate that a \$5 carbon price (corresponding to a \$3.17/Mcf tax on methane) would decrease emissions from the sector by 56 percent. This corresponds to a decrease of about 82 billion cubic feet of fugitive methane emissions annually, which is about 52 million tons of CO_2-equivalent emissions. At this tax level, the marginal unit of abatement would cost firms about \$5.83/Mcf (\$9.20/tCO_2e) and the total cost to the sector would be \$334 million. However, the total wholesale value of the captured gas (calculated at the facility level using average gas prices faced over the study period) would be \$264 million, implying an overall net cost increase of \$70 million, which is only about 0.24 cents per Mcf of gas sold.[56]

The convexity of the MACC demonstrates diminishing returns to increasing taxes as the cheapest abatement opportunities are exploited. I estimate that a \$20 carbon

[54] I use the 100-year warming potential of 34 from the IPCC's Fifth Assessment Report (i.e. one ton of emitted methane results in warming equivalent to 34 tons of CO_2). One ton of methane at standard pressure contains 53.68 Mcf of gas, so 1 Mcf of methane = 34/53.68 tons of CO_2-equivalent emissions.

[55] For the bootstrap, I impose the functional form $R_{it} = \beta_0 + \beta_1 P^{-1} + \beta_2 P^{-2}$, which is the best second-order FP fit for the full sample, rather than allowing the bootstrapped sample to fit the fractional polynomial in each iteration. This prevents cases where bootstrap samples may generate a functional form that becomes upward sloping at higher prices. Observations are clustered at the facility level for resampling. For each of 100 iterations, a random sample of 222 facilities is drawn with replacement, then used to estimate the β_1 and β_2 used in that iteration. However, the original sample is used for the baseline prices and emission rates of facilities.

[56] This calculation uses the EIA's 2016 estimate for marketed U.S. gas production (28,479 Bcf).

Table 4: Simulation results for a subset of potential methane prices.

Methane Tax ($/Mcf)	Equiv. Carbon Price ($/t$CO_2$e)	Total Abatement (tCO_2e)	Total Abatement (Percent)	Total Cost ($)	Value of Captured Gas ($)	Net Cost ($/Mcf)
3.17	5.00	51,974,000	56.1	333,574,000	264,142,000	0.0024
		(20,432,000)	(22.1)	(130,763,000)	(103,386,000)	(0.0010)
12.67	20.00	67,480,000	72.9	534,022,000	342,871,000	0.0067
		(28,060,000)	(30.3)	(237,556,000)	(141,536,000)	(0.0035)
27.37	43.21	70,632,000	75.9	624,502,000	359,045,000	0.0093
		(30,378,000)	(32.8)	(309,830,000)	(153,109,000)	(0.0057)

$27.37/Mcf is the social cost of methane under a 3% discount rate following (EPA, 2016)
Variables in Mcf and dollars rounded to nearest 1,000
Bootstrapped standard errors in parentheses

price (corresponding to an $12.67/Mcf tax on methane) would decrease emissions by 73 percent (about 106 Bcf or 67 million tCO_2e).[57] The total cost would be $534 million and the value of conserved gas would total $343 million, implying a (still relatively modest) net cost increase of 0.67 cents per Mcf of gas sold. A tax designed to fully internalize the social cost of methane would reduce emissions by 76 percent at a net cost of 0.93 cents per Mcf sold.[58] As this is less than 1 percent of the wellhead price of gas anywhere in the country, this result indicates that natural gas is likely to remain competitive in a world where fugitive methane emissions are incorporated in

[57] Note that this level of tax is an out-of-sample sample prediction, as support in gas prices only ranges from about $1.50-$6. Bias could go in either direction. For example, many low-cost abatement opportunities possible at lower gas prices (i.e. those detected by this analysis) may not applicable at higher prices, creating an upward bias. However, it is also likely that many powerful abatement technologies only become cost effective at prices greater that $6, and thus are not at all reflected in this MAC, meaning actual reductions at higher taxes would be greater than those predicted here.

[58] I use $27.37/Mcf for the social cost of methane, which reflects emissions generated in 2020 assuming a 3 percent discount rate and normalized to 2018 dollars. This figure is drawn from 2016 EPA recommendations based on research by Marten *et al.* (2015). Marten et al.'s estimate has the advantage of directly estimating damages from methane instead of converting them from CO_2. However, it is important to note that their estimate is based on a warming potential for methane from the IPCC's Fourth Assessment Report, and has not been updated to account for the higher warming potential recommended by the IPCC's Fifth Assessment Report. I use the Fifth Assessment Report's recommended warming factor of 34 elsewhere in this paper.

climate legislation.

As with any simulation model, these results are dependent to some extent on model selection choices. I find that they are robust to three intuitive modifications: Increasing the lower-bound for facilities' emission rates, starting facilities at 2016 prices and emission rates, and using the estimated relationship between emissions and price from the first-order fractional polynomial model. Results, presented in Tables A5-A7 in the Appendix, generally indicate that choices that decrease total abatement correspondingly decrease costs, and vice-versa. The modification that raises costs the most is using the first-order FP curve, which is steeper at higher gas prices than the second-order fit. However, even in this specification, fully internalizing the social cost of methane reduces the net cost of gas extraction by only about half a percent.

1.6.2 Comparison with Other Abatement Cost Estimates

The MACC estimated above suggests substantially lower abatement costs than most engineering studies of methane leakage. For example, a 2015 EPA cost-benefit analysis of a proposed set of regulations that would affect the entire gas supply chain estimated that they would reduce emissions by only 3.8-4.2 million tCO_2e annually at a net cost of $150-210 million (EPA, 2015).[59] In contrast, this paper estimates these initial reductions to be near costless under the implementation of methane pricing. Although substantial methodological differences undoubtedly contribute to this disparity, the regulatory instrument considered also has an impact. The proposed EPA regulations mandate certain types of equipment and practices for new wells, which will be more or less cost-effective at different well sites, and which are also unlikely to be the most cost-effective measures on average due to the regulator having imperfect information. However, a methane tax or permit trading system characteristically results in the

[59] These figures are for emissions generated in 2020 with a social cost of methane based on a 3 percent discount rate.

most cost-effectiveness abatement measures being undertaken first.[60]

Another reference for methane abatement costs is a 2016 technical report by ICF, which constructs a MACC for the entire natural gas industry using engineering cost estimates (ICF, 2016).[61] That study aligns somewhat more closely with the findings in this paper, identifying abatement opportunities covering 88 Bcf per year that could be achieved at a net cost of $296 million. The abatement cost curve estimated in this study predicts that a reduction of 88 Bcf per year would cost roughly $87 million. One other notable difference is that the ICF MACC predicts negative abatement costs for about 17 Bcf of this abatement. As with the McKinsey curve (Enkvist *et al.*, 2007), the existence of GHG abatement opportunities that have positive private benefits indicates either a failure of their methodology to fully capture some nuanced costs or the presence of a market failure that prevents firms from realizing these potential savings.

One study with predictions that align quite closely with those made in this paper is Mayfield *et al.* (2017), who estimate that the the optimal level of abatement would reduce emissions from the transmission segment of the gas industry by 80 percent.[62] Mayfield et al. use engineering cost estimates as an input for a Monte Carlo simulation model in which a social planner employs lowest-cost abatement technologies first, which is broadly analogous to the implementation of methane pricing. The fact that the MACC for the production sector estimated here—which uses an entirely separate methodology and does not use any data on costs—predicts abatement costs generally in line with or below previous engineering estimates greatly strengthens the conclusion that methane emissions emissions from the natural gas industry can be reduced at

[60] For example, the $5 carbon tax scenario considered in this paper, which roughly doubles the opportunity cost for firms to emit gas, would make a large number of equipment upgrades that were not quite cost effective before worthwhile. However, the very same equipment upgrades might be much less cost-effective at other wells due to heterogeneous real-world conditions, and these upgrades would be passed over.

[61] The 2016 ICF study is an update to a highly similar analysis conducted in 2014 (ICF, 2014).

[62] This estimate is based on a slightly lower social cost of methane equivalent to $24.77/Mcf.

relatively low cost.

This implication is especially clear when comparing the estimates in this paper to abatement costs for greenhouse gas emissions from other sectors. At time of publication, the EU ETS permit price was roughly \$25 and the California permit price was \$15, implying that any additional abatement in sectors covered by their respective permit trading programs would cost at least as much. In contrast, my results imply that cutting methane emissions from natural gas production in half could be achieved at a carbon price below \$5.

Considering average abatement costs, the methane policy equivalent to a \$5 carbon tax, which would reduce sector emissions by 56 percent, is predicted to cost only \$1.34 on average per ton of CO_2-equivalent emissions captured. The policy equivalent to fully internalizing the social cost of methane is predicted to have an average cost of only \$3.76/tCO_2e. Meng (2017) estimates that industry believed an emissions trading scheme proposed in the U.S. in 2009 would have cost \$5-19/tCO_2. Callaway *et al.* (2018) estimates that the abatement cost of installing new renewable energy generation to be at least \$25/tCO_2e for wind and \$43/tCO_2e for solar. Finally, Fowlie *et al.* (2018) estimates CO_2 abatement costs from household weatherization to be \$201/tCO_2. This disparity in abatement costs indicates that methane emissions from natural gas production are an efficient area to prioritize to mitigate greenhouse gas emissions in the short run.

1.7 Conclusion

This paper estimates the marginal abatement cost curve for methane emissions from the natural gas production industry. Because identification is derived from actual firm behavior, results implicitly capture firms' decision-making process to engage in cost-effective abatement. This methodology is therefore well-suited for predicting the effects of regulating methane using market-based instruments, which generate the same incentives.

I find evidence that market-based regulation of methane emissions would achieve substantial greenhouse gas abatement at very low cost. The equivalent of a $5 carbon tax applied to methane could reduce emissions from the sector by 56 percent. This corresponds to roughly 46 million tons of CO_2-equivalent emissions per year, which is close to 1 percent of total U.S. greenhouse gas emissions. Such a policy would imply a net cost of $73 million annually (not including administrative costs) while reducing future climate damages on the order of $1.7 billion. Fully internalizing the social cost of methane would reduce emissions from the sector by roughly 75 percent while increasing the net cost of gas production by less than $0.01/Mcf, indicating that methane regulation could be established with minimal competitiveness impacts.

A number of important caveats to these results have been raised throughout the paper, and two in particular merit further discussion here. First, estimated CH_4 emission reductions are only representative of emissions as they are reported to the GHGRP. While the GHGRP is the most comprehensive record of methane emissions from the natural gas industry currently available, it does not effectively capture many ways in which facility operators mitigate leakage, and these are therefore not picked up in this analysis. For example, the role of leak detection and repair in reducing emissions is only minimally captured by the GHGRP. Fortunately, recent advancements in satellite CH_4 monitoring may soon enable more accurate estimation of abatement costs and open the door to many other avenues for empirically investigating methane emissions from all parts of the gas supply chain (Jacob *et al.*, 2016).

Second, realizing abatement at the costs estimated in this paper requires the successful implementation of a methane tax or trading program.[63] Designing such a

[63] Abatement costs under the implementation of conventional regulation (such as equipment mandates) would be higher than those predicted here, as regulators have imperfect information as to the lowest cost abatement technologies. However, given the current challenges in monitoring CH_4, the advantage of many types of conventional regulation that they are straightforward to enforce

program in a setting where an accurate, low-cost monitoring technology is not readily available presents a formidable challenge. One approach would be to use an inventory calculation such as the GHGRP. Indeed, since the results of this paper are based upon emissions as estimated by the GHGRP, it is reasonable to believe enforcing a market-based instrument based on a reporting survey would be effective in reducing emissions at low cost. Although not all emissions would be captured, this approach is advantageous in being readily practicable. However, the introduction of real penalties would incentivize firms to abate based on emissions as detected by the GHGRP, rather than based on their own knowledge about which abatement technologies are most efficient for their specific facilities, which would further increase the divergence from the theoretical optimum level of abatement. Another approach would be to use direct measurements. Although continuous monitoring of all production sites promises to be cost-prohibitive for many years, intermittent sampling by sensors mounted on aircraft or ground vehicles may be practically feasible in the very near future or even today (Emran *et al.*, 2017; Fredenslund *et al.*, 2017; van den Bossche *et al.*, 2017). Such a program might be particularly cost-effective if sampling were randomly structured to develop firm-level estimates rather than to estimate emissions for individual wells. Beyond applied questions regarding which technologies to use and how frequently to sample emissions, this approach necessitates deeper consideration into how to handle measurement error in a way that is fair to firms and preserves incentives.

As natural gas continues to expand its role in the transition to sustainable energy, it is critical that its particular externalities be effectively managed to optimally balance its utilization with other energy sources. So long as methane emissions are minimally regulated, CO_2-focused regulations that shift usage of other fossil fuels to gas will be severely attenuated in their intended climate impacts. Moreover, comparatively low abatement costs establish a case for prioritizing methane regulations from

merits considerable weight in the short run.

the gas supply chain. Although our knowledge of the causes and scale of methane pollution from the natural gas sector has expanded enormously over the last decade, there are still many unanswered questions surrounding the design of policy to reduce it to efficient levels. Estimating the costs and benefits of various regulatory approaches, exploring of the equilibrium effects of climate policy that does not address methane emissions, and developing the theory of regulation under conditions of imperfect measurement are key areas where further research might inform such policy.

References

Aldy, Joseph E, & Pizer, William A. 2015. The competitiveness impacts of climate change mitigation policies. *Journal of the Association of Environmental and Resource Economists*, **2**(4), 565–595.

Allcott, Hunt, & Greenstone, Michael. 2012. Is there an energy efficiency gap? *The Journal of Economic Perspectives*, **26**(1), 3–28.

Alvarez, Ramón A, Zavala-Araiza, Daniel, Lyon, David R, Allen, David T, Barkley, Zachary R, Brandt, Adam R, Davis, Kenneth J, Herndon, Scott C, Jacob, Daniel J, Karion, Anna, *et al.* 2018. Assessment of methane emissions from the US oil and gas supply chain. *Science*, eaar7204.

Anderson, Soren T, & Sallee, James M. 2011. Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards. *American Economic Review*, **101**(4), 1375–1409.

Callaway, Duncan S, Fowlie, Meredith, & McCormick, Gavin. 2018. Location, location, location: The variable value of renewable energy and demand-side efficiency resources. *Journal of the Association of Environmental and Resource Economists*, **5**(1), 39–75.

Cremer, Helmuth, & Gahvari, Firouz. 2002. Imperfect observability of emissions and second-best emission and output taxes. *Journal of Public Economics*, **85**(3), 385–407.

Cullen, Joseph A, & Mansur, Erin T. 2017. Inferring carbon abatement costs in electricity markets: A revealed preference approach using the shale revolution. *American Economic Journal: Economic Policy*, **9**(3), 106–33.

Davis, Lucas W, & Muehlegger, Erich. 2010. Do Americans consume too little natural gas? An empirical test of marginal cost pricing. *The RAND Journal of Economics*, **41**(4), 791–810.

Delhotal, K Casey, De la Chesnaye, Francisco C, Gardiner, Ann, Bates, Judith, & Sankovski, Alexei. 2006. Mitigation of methane and nitrous oxide emissions from waste, energy and industry. *The Energy Journal*, 45–62.

EIA. 2018. *Annual Energy Outlook 2018 with Projections to 2050*. US Energy Information Administration.

Emran, Bara J, Tannant, Dwayne D, & Najjaran, Homayoun. 2017. Low-altitude aerial methane concentration mapping. *Remote Sensing*, **9**(8), 823.

Enkvist, P, Nauclér, Tomas, & Rosander, Jerker. 2007. A cost curve for greenhouse gas reduction. *McKinsey Quarterly*, **1**, 34.

EPA. 2015 (August). *Regulatory Impact Analysis of the Proposed Emission Standards for New and Modified Sources in the Oil and Natural Gas Sector*. Office of Air and Radiation, U.S. Environmental Protection Agency.

EPA. 2016 (August). *Addendum to Technical Support Document on Social Cost of Carbon for Regulatory Impact Analysis under Executive Order 12866: Application of the Methodology to Estimate the Social Cost of Methane and the Social Cost of Nitrous Oxide*. Interagency Working Group on Social Cost of Greenhouse Gases, United States Government.

Fowlie, Meredith, Greenstone, Michael, & Wolfram, Catherine. 2018. Do energy efficiency investments deliver? Evidence from the weatherization assistance program. *The Quarterly Journal of Economics*, **133**(3), 1597–1644.

Fredenslund, Anders Michael, Mønster, J, Kjeldsen, Peter, & Scheutz, Charlotte. 2017. Screening method to assess the greenhouse gas mitigation potential of old landfills, based on downwind methane concentration measurements. *In: Sardinia 2017, 16th International Waste Management and Landfill Symposium*. CISA Publisher.

Gabriel, Steven A, Zhuang, Jifang, & Kiet, Supat. 2005. A large-scale linear complementarity model of the North American natural gas market. *Energy economics*, **27**(4), 639–665.

Ganapati, Sharat, Shapiro, Joseph S, & Walker, Reed. 2016. *The Incidence of Carbon Taxes in US Manufacturing: Lessons from Energy Cost Pass-Through*. Tech. rept. National Bureau of Economic Research.

Gillingham, Kenneth, & Palmer, Karen. 2014. Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence. *Review of Environmental Economics and Policy*, **8**(1), 18–38.

Greene, William. 2008. Functional forms for the negative binomial model for count data. *Economics Letters*, **99**(3), 585–590.

Hausman, Catherine, & Kellogg, Ryan. 2015. *Welfare and distributional implications of shale gas*. Tech. rept. National Bureau of Economic Research.

Hausman, Catherine, & Muehlenbachs, Lucija. 2016. Price Regulation and Environmental Externalities: Evidence from Methane Leaks.

ICF. 2014 (March). *Economic Analysis of Methane Emission Reduction Opportunities in the U.S. Onshore Oil and Natural Gas Industries*. Tech. rept. ICF International. Prepared for Environmental Defense Fund.

ICF. 2016 (May). *Economic Analysis of Methane Emission Reduction Potential from Natural Gas Systems*. Tech. rept. ICF International. Prepared for ONE Future Inc.

IEA. 2017. *World Energy Outlook 2017*. International Energy Agency.

IPCC. 2014. *Climate Change 2014: Mitigation of Climate Change*. Working Group III Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.

Jacob, Daniel J, Turner, Alexander J, Maasakkers, Joannes D, Sheng, Jianxiong, Sun, Kang, Liu, Xiong, Chance, Kelly, Aben, Ilse, McKeever, Jason, & Frankenberg, Christian. 2016. Satellite observations of atmospheric methane and their value for quantifying methane emissions. *Atmospheric Chemistry and Physics*, **16**(22), 14371–14396.

Klepper, Gernot, & Peterson, Sonja. 2006. Marginal abatement cost curves in general equilibrium: The influence of world energy prices. *Resource and Energy Economics*, **28**(1), 1–23.

Lade, Gabriel E, & Rudik, Ivan. 2017. *Costs of inefficient regulation: Evidence from the Bakken*. Tech. rept. National Bureau of Economic Research.

Lubowski, Ruben N, Plantinga, Andrew J, & Stavins, Robert N. 2006. Land-use change and carbon sinks: econometric estimation of the carbon sequestration supply function. *Journal of Environmental Economics and Management*, **51**(2), 135–152.

Lyon, David R, Zavala-Araiza, Daniel, Alvarez, Ramon, Harriss, Robert, Palacios, Virginia, Lan, Xin, Talbot, Robert, Lavoie, Tegan, Shepson, Paul, Yacovitch, Tara I, *et al.* 2015. Constructing a spatially resolved methane emission inventory for the Barnett Shale region. *Environmental science & technology*, **49**(13), 8147–8157.

Marten, Alex L, Kopits, Elizabeth A, Griffiths, Charles W, Newbold, Stephen C, & Wolverton, Ann. 2015. Incremental CH4 and N2O mitigation benefits consistent with the US Government's SC-CO2 estimates. *Climate Policy*, **15**(2), 272–298.

Mayfield, Erin N, Robinson, Allen L, & Cohon, Jared L. 2017. System-wide and superemitter policy options for the abatement of methane emissions from the US natural gas system. *Environmental science & technology*, **51**(9), 4772–4780.

McCabe, David, Greertsma, Meleah, Matthews, Nathan, Fleischman, Lesley, & Schroeder, Darin. 2015 (January). *Waste Not: Common Sense Ways to Reduce Methane Pollution from the Oil and Natural Gas Industry*. Tech. rept. Clean Air Task Force, Natural Resources Defense Council, and Sierra Club. Technical Report.

Meng, Kyle C. 2017. Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy. *American Economic Review*, **107**(3), 748–84.

Morris, Jennifer, Paltsev, Sergey, & Reilly, John. 2012. Marginal abatement costs and marginal welfare costs for greenhouse gas emissions reductions: results from the EPPA model. *Environmental Modeling & Assessment*, **17**(4), 325–336.

Petron, Gabrielle, Karion, Anna, Sweeney, Colm, Miller, Benjamin R, Montzka, Stephen A, Frost, Gregory J, Trainer, Michael, Tans, Pieter, Andrews, Arlyn, Kofler, Jonathan, *et al.* 2014. A new look at methane and nonmethane hydrocarbon emissions from oil and natural gas operations in the Colorado Denver-Julesburg Basin. *Journal of Geophysical Research: Atmospheres*, **119**(11), 6836–6852.

Royston, Patrick, & Altman, Douglas G. 1994. Regression using fractional polynomials of continuous covariates: parsimonious parametric modelling. *Applied statistics*, 429–467.

Sanchez, Nicolas, & Mays, David C. 2015. Effect of methane leakage on the greenhouse gas footprint of electricity generation. *Climatic change*, **133**(2), 169–178.

Sauerbrei, Willi, Meier-Hirmer, Carolina, Benner, A, & Royston, Patrick. 2006. Multivariable regression model building by using fractional polynomials: description of SAS, STATA and R programs. *Computational Statistics & Data Analysis*, **50**(12), 3464–3485.

Stavins, Robert N. 1999. The costs of carbon sequestration: a revealed-preference approach. *The American Economic Review*, **89**(4), 994–1009.

Stranlund, John K, Chavez, Carlos A, & Villena, Mauricio G. 2009. The optimal pricing of pollution when enforcement is costly. *Journal of Environmental Economics and Management*, **58**(2), 183–191.

Subramanian, R, Williams, Laurie L, Vaughn, Timothy L, Zimmerle, Daniel, Roscioli, Joseph R, Herndon, Scott C, Yacovitch, Tara I, Floerchinger, Cody, Tkacik, Daniel S, Mitchell, Austin L, *et al.* 2015. Methane emissions from natural gas compressor stations in the transmission and storage sector: Measurements and comparisons with the EPA greenhouse gas reporting program protocol. *Environmental science & technology*, **49**(5), 3252–3261.

van den Bossche, Michael, Rose, Nathan Tyler, & De Wekker, Stephan Franz Joseph. 2017. Potential of a low-cost gas sensor for atmospheric methane monitoring. *Sensors and Actuators B: Chemical*, **238**, 501–509.

Chapter 2

An Estimate-Based Approach to Emissions Pricing

2.1 Introduction

One essential component of any emissions pricing program is a sufficiently accurate measure of emissions upon which to apply a tax or account for the use of permits. For the 44 countries and many more sub-national jurisdictions with active carbon prices, emissions are either directly measured using continuous emissions monitoring sensors placed in smokestacks or calculated based on the carbon content of fossil fuels. While measurement is straightforward for most sources of CO_2, however, it presents a significant challenge for many non-point source pollutants, including methane emissions from natural gas infrastructure.

Methane emissions from the oil and gas industry currently generate about 40 percent as much warming as all gasoline vehicles in the United States (Alvarez *et al.*, 2018).[1] While methane emissions from the natural gas supply chain are not the largest slice of the greenhouse gas emissions pie, recent research in economics and engineering has shown that they can be abated at very low cost relative to other sources throughout the economy.[2] Various abatement strategies may be effective,

[1] This figure uses the warming potential of methane over a 100-year time horizon. Because methane is a short-lived greenhouse gas, warming impacts are much higher on a shorter time horizon. Over a 20-year time horizon, these emissions are responsible for about 120 percent as much warming as U.S. gasoline vehicles.

[2] e.g. Marks (2019); Tyner & Johnson (2018); Mayfield *et al.* (2017); ICF (2016)

but a large degree of heterogeneity in abatement costs across facilities and mitigation technologies suggests high potential efficiency gains from regulating these emissions using market-based instruments (Newell & Stavins, 2003).

Fundamentally, the successful implementation of an emissions tax or permit trading program requires a sufficiently accurate measure of emissions. Because methane emissions from the natural gas supply chain are generated from a variety of different sources in many different ways, comprehensively monitoring them at the facility level is prohibitively costly at this time. However, it would be feasible to develop a robust firm-level estimate of emissions by conducting measurements at a randomly selected subset of facilities operated by each firm.[3] This paper presents a theoretical framework that demonstrates that this "estimate-based" approach to emissions pricing under imperfect information preserves the efficiency benefits of efficiency pricing.

In particular, the model illustrates that because firms' expected benefit of abatement under an estimate-based emissions pricing program is equivalent to their expected benefit under a program with comprehensive monitoring, they are incentivized to invest in cost-effective abatement technologies. The model also determines the optimal number of facilities to sample as a function of the firm's total number of facilities, the per-site measurement cost, and various other parameters. Furthermore, I develop an extension in which the regulator uses the lower bound of a confidence interval for each firm's estimated emissions instead of the point estimate to allay concerns that many firms would be overcharged for emissions they are not actually generating.

The purpose of the simulation model is to predict the impacts of implementing estimate-based emissions pricing for methane emissions from the U.S. oil and gas industry. It uses data on industry composition from government, industry, and academic sources, data on methane emissions from recent scientific measurement studies, and estimated marginal abatement cost curves from recent economic studies. The

[3] "Facilities" in this paper refer to individual wells, gathering stations, processing plants, compression stations, or storage sites.

model predicts that at current measurement costs, it would be optimal to sample all facilities in the gathering, processing, transmission, and storage industry segments (about 7,000 facilities total) and about 8 percent of the approximately 500,000 oil and gas wells that make up the production segment. The total estimated measurement cost is roughly $160 million per year, while the estimated annual abatement benefit (net of physical abatement costs) is roughly $13 billion.

This work builds upon an expansive literature exploring innovative strategies to regulate pollution under imperfect information. For example, Blundell *et al.* (2018) demonstrate that the dynamic tier-based structure of the EPA's approach to enforcing the Clean Air Act incentivizes firms to undertake mitigation investments that are many times more costly than the potential fines they face. Duflo *et al.* (2013) perform a large-scale field experiment that introduces a novel approach to auditing for polluting plants in India, significantly improving the accuracy of reporting and reducing emissions among treated facilities. Even more relevant to this paper is the subset of this literature that focuses on non-point source pollution. One particularly relevant example is Segerson (1988), which proposes a novel regulatory approach for controlling water pollution given conditions of imperfect monitoring. Segerson develops a theoretical model that illustrates how subsidizing and penalizing firms based on ambient pollution levels can induce efficient abatement. Variations of her approach have since been explored theoretically (Xepapadeas, 2011; Hansen, 2002), tested in laboratory and field experiments (Suter & Vossler, 2013; Cochard *et al.*, 2005; Poe *et al.*, 2004) and implemented in at least one jurisdiction (Dowd *et al.*, 2008).

Methane emissions from the oil and gas industry are similarly a non-point source of pollution, and it's possible that previously developed approaches would be useful for controlling them. However, unlike other proposed strategies for non-point source pollution control, the approach explored in this proposal is intended to realize the efficiency benefits of emissions pricing in a setting where pollution is costly—but not impossible—to measure.

This paper begins by providing further background on methane emissions and exploring alternative approaches to addressing them. Section 2 presents a theoretical model of estimate-based emissions pricing. Section 3 develops an extension in which the regulator taxes firms based on the lower bound of a confidence interval for their estimated emissions. Section 4 provides an overview of the data used in the simulation model. Section 6 presents the methodology and results of the simulation, and Section 6 concludes.

2.2 Background

2.2.1 Methane Emissions

Consumption of natural gas has increased dramatically over the last decade, primarily due to the introduction of hydraulic fracturing and horizontal drilling technologies in the late 2000s. In the U.S. electricity sector, it recently displaced coal to become the predominant generation resource, and its share of the generation mix is expected to continue increasing both with and without climate policy (EIA, 2018). This trend has some positive implications for mitigating climate change, as natural gas produces about half as much CO_2 as coal when burned to produce electricity. However, natural gas is composed of about 90% methane, which is itself a greenhouse gas that is about 34 times more potent than CO_2 on a 100-year time horizon (IPCC, 2014).[4] This high warming potential implies that a small fraction of gas escaping anywhere along the gas supply chain can have severe warming impacts.

A recently published metanalysis of about two dozen scientific measurement studies estimates that 2.3 percent of total natural gas production is currently being emit-

[4] Throughout this paper, a warming factor of 34 (the most recent estimate from the IPCC's Fifth Assessment Report) is used to convert tons of methane emissions to tons of CO_2-equivalent emissions.

Figure 1: Methane emissions from the natural gas supply chain

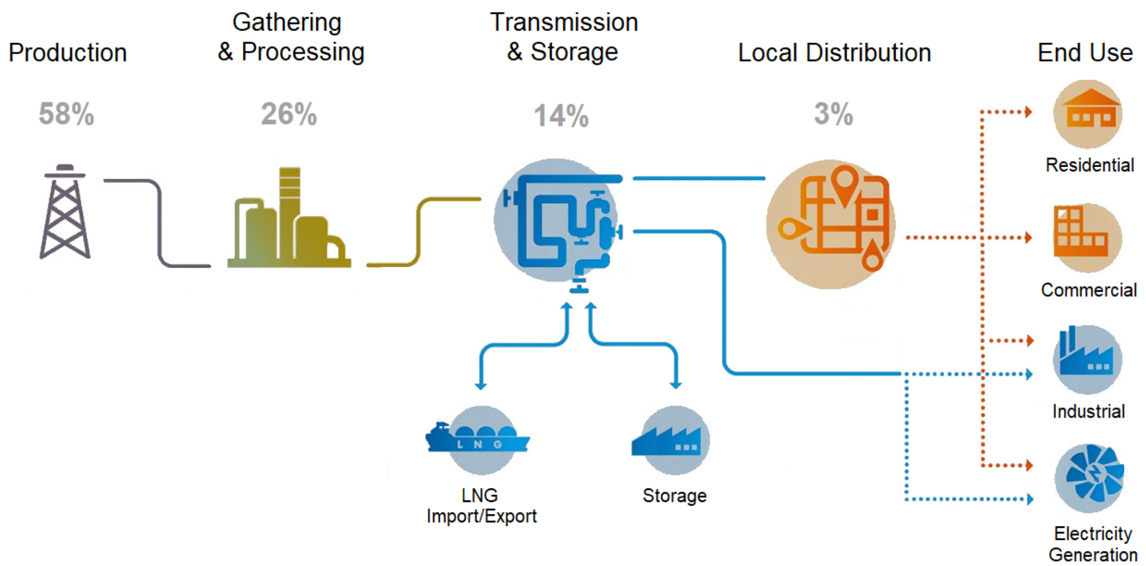

Estimates from Alvarez *et al.* (2018). Graphic adapted with permission from AEMO NGFR.

ted, resulting in warming effects equivalent to about 40 percent of all gasoline vehicle emissions in the United States on a 100-year time horizon. About 60 percent of these emissions are generated in the onshore production segment of the supply chain, which consists of over 500,000 individual oil and gas wells (see Figure 1).[5]

Across the supply chain, methane emissions are caused by unintentional leaks resulting from equipment failures (or operational mistakes) and by intentional venting as part of the regular operation of certain types of equipment or during maintenance. Additionally, in the production sector, gas is intentionally vented during the well completion process. Completion emissions from sites employing hydraulic fracturing are significantly higher, as a large amount of gas is brought to the surface alongside fracking liquid that is often uneconomical to capture in the absence of regulation (Howarth *et al.*, 2011). Lastly, natural gas and oil are very often co-located in reservoirs, and in cases where a primarily oil-producing well produces a limited quantity

[5] Methane emissions from offshore facilities, which represent about 16 percent of U.S. oil production and about 5 percent of U.S. gas production, have not yet been studied in the scientific literature. Because emission rates and abatement costs are likely to be substantially different from onshore production, these emissions are not considered in this paper.

of gas that does not justify installing equipment to capture it, this co-produced gas is either vented or flared.[6]

Corresponding to this diverse set of emission sources, there exists a wide array of potential abatement measures firms can undertake to reduce their emissions. Broadly, methane abatement measures can be categorized into equipment upgrades and improved operational practices. Equipment upgrades include replacing high-bleed pneumatic devices with low-bleed and zero-bleed devices, replacing compressor wet seals with dry seals, installing vapor recovery units for equipment that would otherwise vent gas, and utilizing flares or installing gathering pipelines for well completions and co-produced gas (Tyner & Johnson, 2018; ICF, 2016). Improved operational practices include leak detection and repair (LDAR), reducing waste during blowdowns for maintenance and testing, and improved training to reduce the frequency of operational mistakes (Munnings & Krupnick, 2017; Lowell & Russell, 2016).[7]

2.2.2 Potential Abatement Policies

2.2.2.1 Conventional Regulations

One viable approach to mitigating methane emissions from the gas supply chain is to employ conventional regulations, such as equipment and operational mandates. Colorado became the first state to enact significant regulations on methane in 2014, requiring firms to replace many types of equipment with lower-emitting models and mandating leak detection and repair at certain frequencies depending on the facility type. California adopted similar regulations in 2017, North Dakota introduced flaring and venting standards for its primarily oil-producing production infrastructure in 2015, and Canada is currently developing regulations designed to reduce methane

[6] Because methane is substantially more potent than CO_2, flaring natural gas instead of venting it reduces its warming impacts by about an order of magnitude. Flaring also reduces emission of local co-pollutants such as volatile organic compounds.

[7] For example, improperly sealed hatches on storage tanks at production sites are a primary cause of "superemitting" natural gas facilities (Zavala-Araiza *et al.*, 2017).

emissions from its oil and gas industry by 40-45% by 2025.[8]

Equipment and operational mandates circumvent the measurement challenge by focusing on specific interventions rather than outcomes.[9] Engineering cost studies provide regulators with rough estimates of the expected benefits and costs of various abatement measures, and it is straightforward for auditors to verify that certain types of equipment are installed where they are required. However, notwithstanding measurement costs, a large body of theoretical and empirical research has demonstrated conventional regulations are significantly less efficient in achieving a target level of abatement versus market-based approaches.[10]

2.2.2.2 Inventory-Based Emissions Pricing

Another option would be to apply an emissions tax or permit trading program based on facility-level estimates of emissions constructed using equipment counts, equipment characteristics, equipment emission factors, and records of firm activities in well completion, processing, and maintenance.[11] The EPA's Greenhouse Gas Reporting Program (GHGRP) currently tracks methane emissions at U.S. facilities in this manner. This approach would be relatively straightforward to implement given the existence of an established inventory framework, and would incur lower abatement costs relative to conventional regulations by allowing firms greater flexibility in achieving mitigation.

However, an inventory-based approach has two significant disadvantages. First, direct measurement studies have shown that current inventories fail to capture a number of important emission sources, resulting in substantial downward bias.[12] Sec-

[8] Methane emissions are currently almost entirely unregulated in most other jurisdictions around the world.

[9] Site-level performance standards do not have this benefit, as enforcement also requires measurement.

[10] e.g. Burtraw & Szambelan (2010); Carlson *et al.* (2000); Baumol *et al.* (1988); Montgomery (1972)

[11] A variation of this approach is proposed alongside other potential methane abatement policies in Munnings & Krupnick (2017).

[12] Alvarez *et al.* (2018) estimate that U.S. emissions across the gas supply chain are 60% higher than

ond, the effectiveness of an inventory-based approach in incentivizing cost-effective abatement would directly depend on the accuracy with which regulators are able to identify and quantify available abatement opportunities.[13] For abatement through investments in lower-emitting equipment, emission factors would need to be accurate under real operating conditions for as many different equipment types as possible. For abatement through improved operational practices, some activities (such as increased frequency of LDAR) can be roughly captured by an inventory framework, but other activities (such as procedures for double-checking hatches and seals) are inherently very difficult for inventory methodologies to capture.[14]

2.2.2.3 Estimate-Based Emissions Pricing

An estimate-based approach would price emissions based on firm-level estimates constructed by performing direct measurements at a randomly chosen subset of each firm's facilities. So long as the facilities are selected randomly, these measurements can be used to develop a robust firm-level estimate of emissions with a known level of statistical uncertainty. By only measuring a fraction of facilities, this approach would cost only a fraction of the price of comprehensive monitoring.

The pivotal advantage of this approach is that it covers all potential emissions sources, including sources that are difficult to capture using inventory methodologies. Because this mechanism is completely ambivalent about how emissions are generated, the incentive to reduce emissions applies equally to all potential sources, including both intentional venting and unintentional leaks. Correspondingly, this approach of-

---

what is reported by the GHGRP, and Atherton *et al.* (2017) estimate that actual emissions from British Columbia's natural gas production sector are about twice as high as reported emissions.

[13] This is also true of conventional regulations.

[14] One other approach to pricing methane would be an output-based tax or permit trading program. In other words, the emissions price would be applied based on production at a well or throughput at a midstream facility multiplied by a common assumed emission factor for each facility type. This approach would effectively incentivize firms to produce less natural gas without incentivizing them to engage in cost-effective methane abatement activities, which runs counter to the intended goal of emissions pricing.

fers firms full flexibility in how they choose to reduce their emissions. Rather than being constrained by established abatement technologies and emissions factors used by an inventory, firms would be incentivized to use any technologies and practices they determine to be cost-effective (and there would be no risk of firms being incentivized to use inefficient technologies that might have been misspecified by an inventory framework). Furthermore, firms would have a direct incentive to develop novel technologies and strategies for reducing emissions, which could have positive externality benefits of reducing methane abatement costs in unregulated jurisdictions.

Section 2 of this paper demonstrates that an estimate-based approach to emissions pricing fully preserves firms' incentives to engage in cost-effective abatement in a stylized setting where firms operate an arbitrarily large number of facilities. However, it is important to recognize that practical implementation of this policy may face a number of challenges, centering primarily around the introduction of stochasticity in firms' costs of paying taxes or acquiring permits. For firms with many facilities, divergence between estimated and actual emissions will be small proportional to overall revenues. However, for firms with only a few facilities, an unlucky measurement immediately following a major equipment failure could severely disrupt their profitability.[15] Additionally, because the statistical nature of this approach implies that some firms would be taxed for more emissions than they are actually generating, it may be subject to substantial legal challenges. Section 3 of this paper explores a possible solution to this issue, where instead of charging firms based on the point estimate of their emissions, the regulator charges firms based on the lower bound of a confidence interval for their estimated emissions.

[15] This issue might be overcome by including special considerations for firms below a certain threshold size, such as an effective ceiling for measured emission rates at individual facilities or an allowance for a certain number of re-measurement opportunities. For firms that operate a few larger facilities (such as gas processing plants or compression stations), the simulation model predicts that it would be cost-effective to monitor all facilities (where measurements would still be performed randomly over time).

2.2.3 Measurement Technologies

A handful of technologies that can accurately quantify methane emissions at the site level are currently available.[16] Two of the most accurate are "downwind tracer flux" and "other test method 33-A" (OTM-33A). Downwind tracer flux (hereafter referred to simply as tracer flux) operates by placing one or two sources of an inert tracer gas at the emitting site, then measuring air composition at a location downwind of the site and using the proportion of the inert gas to identify dispersion. OTM-33A is a method developed by the EPA that does not require on-site access, but instead is conducted by driving a vehicle equipped with sensors nearby the emitting site and using an inverse Gaussian air dispersion model to calculate methane emissions.

Both methods produce estimates of emissions with known levels uncertainty. Under good conditions, tracer flux is able to generate an unbiased site-level estimate of methane emissions with a standard deviation of approximately 14.5 percent (Omara *et al.*, 2016). OTM-33A is able to generate a site-level estimate of methane emissions with a downward bias of approximately 10 percent (which can be corrected) and a standard deviation of approximately 28 percent Robertson *et al.* (2017). The per-site cost of measurement is roughly \$3,000 for tracer flux and \$500 for OTM-33A.[17]

Both measurement technologies estimate a snapshot of emissions typically covering a few hours, and there are no examples of continuous monitoring of methane emissions from natural gas facilities in the scientific literature. While continuous monitoring would be feasible using either of these technologies, a substantial part of the cost savings from employing an estimate-based approach to emissions pricing derives from the ability to move costly equipment from site to site to perform measurements.

[16] This is evidenced by their use in at least a dozen scientific measurement studies (Alvarez *et al.*, 2018; Brandt *et al.*, 2014).

[17] Cost estimates are based on a conversation with one of the coauthors of Alvarez *et al.* (2018).

2.3 Model

This section presents a static theoretical framework in which a profit-maximizing representative natural gas firm interacts with a welfare-maximizing regulator. The model is framed in terms of an emissions tax; however, without introducing uncertainty in emission damage or abatement cost functions, optimal abatement under an emissions trading program will be equivalent (Baumol *et al.*, 1988).

2.3.1 Firm's Problem

Consider a natural gas firm that operates a fixed stock N individual facilities, indexed by i. Each facility produces $q_i - e_i$ units of output, where q_i is exogenously determined extraction (or throughput) at facility i and e_i is that facility's emissions.[18] In the absence of an emissions pricing program, the firm's profit function is:

$$\Pi = \max_{\mathbf{a}_i} \sum_{i=1}^{N} [P(q_i - e_i) - c(a_i, q_i, \theta_i)] \quad \text{S.T.} \quad e_i = e_i(a_i, q_i, \theta_i) \tag{1}$$

Total revenue is the price of gas P times output, and total cost c is a function of facility size (indicated by q_i), abatement effort a_i, and a vector of site-specific characteristics θ_i.[19] So long as emissions are monotonically decreasing in abatement effort, total cost is implicitly a function of emissions, such that the firm's problem

[18] The assumption that q_i is exogenous is supported in the context of oil and gas production by Anderson *et al.* (2018), who show that once drilled, a well's production is determined by reservoir pressure as opposed to market forces. Making N a choice variable would be a more appropriate way to incorporate the firm's production decision; for simplicity, however, N is assumed to be constant.

[19] θ_i does not impact the results developed from this model, but appropriately captures the real-world feature that abatement costs will be heterogeneous across facilities for reasons other than facility size, such as gas composition, depth of a reservoir, distance to the pipeline network, etc.

can be simplified to choosing the optimal level of emissions:[20]

$$\Pi = \max_{\mathbf{e}_i} \sum_{i=1}^{N} [P(q_i - e_i) - c(e_i, q_i, \theta_i)] \tag{2}$$

Where before c was increasing and convex in a_i, it is now decreasing and convex in e_i. The firm chooses a vector of emissions across all of its facilities $\mathbf{e}_i$, and its first-order condition for emissions at any particular facility i is

$$\frac{-\partial c}{\partial e_i} = P \tag{3}$$

This implies that the firm sets their marginal cost of capturing one unit of gas equal to their marginal private benefit of being able to sell that unit of gas.[21]

2.3.2 Introduction of Estimate-Based Emissions Pricing

Now suppose a regulator taxes the firm based on an estimate of its total emissions constructed by measuring emissions at M of the firm's N facilities, which are randomly selected. The regulator's estimate for the firms' total emissions $\hat{E}$ is the sample average of those measurements times the number of facilities:

$$\hat{E} = N \tfrac{1}{M} \sum_{j=1}^{M} e_j \tag{4}$$

The firm's profit function becomes

$$\mathbb{E}[\Pi] = \max_{\mathbf{e}_i} \sum_{i=1}^{N} [P(q_i - e_i) - c(e_i, q_i, \theta_i)] - T\mathbb{E}[\hat{E}] \tag{5}$$

[20] In practice, although firms may not be able to choose emissions at a particular site with certainty due to the stochastic nature of leaks, their level of abatement effort directly determines expected emissions.

[21] Note that $\frac{-\partial c}{\partial e_i}$ is actually positive, since c is decreasing in e_i.

which simplifies to[22]

$$\mathbb{E}[\Pi] = \max_{\mathbf{e}_i} \sum_{i=1}^{N} [P(q_i - e_i) - c(e_i, q_i, \theta_i) - Te_i] \tag{6}$$

The firm's first-order condition for emissions is now

$$\frac{-\partial c}{\partial e_i} = P + T \tag{7}$$

If the tax is set at the social cost of methane, this first-order condition implies the firm will undertake the socially optimal level of abatement. In other words, the firm will capture all potential emissions that can be abated at a cost less than the sum of the private and external social cost of allowing that gas to escape into the atmosphere. This result allows for heterogeneity in abatement costs across facilities, such that e_i^* is unique for each facility i. The basic intuition underlying this result is that although any individual facility i only has an $\frac{M}{N}$ probability of being sampled, its contribution to the firm's total taxes owed if it is sampled is scaled up by $\frac{N}{M}$, such that in expectation the incentive to reduce emissions at i is the same as it would be if all facilities were sampled with certainty.

2.3.3 Regulator's Problem

The regulator is concerned with maximizing social welfare W by choosing the optimal number of facilities to sample. In addition to the the firm's profits, the regulator values tax revenues generated $T\hat{E}$, the social cost of emission damages TE (assuming the tax is optimally chosen to reflect the social cost of methane), measurement costs, and

[22] A full derivation is presented in Appendix Section A.2.1.

a penalty for greater stochasticity in the firm's tax payment:

$$\mathbb{E}[W] = \mathbb{E}\Big[\max_{M} \sum_{i=1}^{N} [P(q_i - e_i^*) - c(e_i^*, q_i, \theta_i)] - T\hat{E}^* + T\hat{E}^* - TE^* - c_m M - \alpha T|\hat{E}^* - E^*|\Big] \tag{8}$$

Here, e_i^* is the firm's optimal emissions at i (implicitly defined by Equation 7) and E^* is the firm's optimal actual total emissions. For simplicity, the marginal cost of measuring emissions at each facility c_m is assumed to be constant. The term $\alpha T|\hat{E}^* - E^*|$ captures the regulator's incentive to construct an accurate estimate of emissions.[23] There are a number of reasons the regulator would prefer a more accurate estimate. Fundamentally, greater randomness in the tax burden faced by firms distorts firms' costs, causing a divergence from a competitive equilibrium outcome that would result in deadweight loss.[24] Alternatively, one could frame firms as having risk-averse preferences, or reason that increasing the randomness of the tax burdens faced by firms would decrease the political viability of the policy. Without building one of these potential framings into the model, $\alpha T|\hat{E} - E|$ incorporates the regulator's valuation of accurate estimation by penalizing incorrect measurement at a constant ratio determined by the scaling parameter $\alpha \in (0, 1]$. For example, if $\alpha = 1$, then if a firm is either overcharged or undercharged by \$1,000, social welfare is \$1,000 lower.[25]

The only stochastic component from the regulator's perspective is $\mathbb{E}[|\hat{E}^* - E^*|]$,

[23] Without this term, the regulator's first-order condition for M is $\frac{\partial W}{\partial M} = -c_m$, implying the optimal number of facilities to measure is 1 (since at least one measurement is required for the tax to exist).

[24] In other words, if a firm that would otherwise be a competitive producer happens to be charged for more emissions than it is actually generating, it may be forced to exit the market, causing firms with less efficient production to increase their market share, ultimately decreasing the overall supply of natural gas. Although this might increase overall social welfare if CO_2 emissions from natural gas production are not otherwise regulated, it would be an inefficient policy instrument compared to pricing CO_2 emissions separately. Optimal climate policy would price both CO_2 emissions and methane emissions to incentivize both efficient allocation across energy sources and efficient abatement of methane emissions within the oil and gas sector.

[25] $\alpha = 0.5$ would correspond to a scenario where the regulator internalizes overcharging firms at a one-to-one ratio, but does not care about undercharging firms.

which is distributed folded normal with mean 0 and variance $\frac{\sigma^2}{M}$, where σ is the population standard deviation of measured emissions.[26] Assuming tax revenues are used productively (i.e. the firm's taxes paid cancel out with the regulator's taxes received), the regulators problem reduces to:

$$\mathbb{E}[W] = \max_{M} \sum_{i=1}^{N} [P(q_i - e_i^*) - c(e_i^*, q_i, \theta_i)] - TE^* - c_m M - \alpha T \mathbb{E}[|\hat{E}^* - E^*|] \qquad (9)$$

The first-order condition for M is[27]

$$M^* = \left(\frac{\alpha T N \sigma}{c_m \sqrt{2\pi}} \right)^{\frac{2}{3}} \qquad (10)$$

Intuitively, M^* is increasing in the level of the tax, the number of facilities operated by the firm, and the variance of emissions across the firm's facilities, and decreasing in the cost of measurement (π is the mathematical constant). While this model is simplified by using a representative firm, in practice N and σ will vary across firms. Section 2.5 calculates M^* for U.S. natural gas firms using actual facility counts and estimated distributions of emissions from the scientific literature on methane emissions.

[26] i.e. σ incorporates both the variance of actual emissions across sites and additional variance introduced by random noise from the measurement technology (see Section 2.2.3). Assuming actual emissions are distributed across sites with mean μ_e and standard deviation σ_e, and the measurement technology introduces classical measurement error with mean $\mu_m = 0$ and standard deviation σ_e (where σ_e is proportional rather than in levels), then the population variance of measured emissions $\sigma^2 = \sigma_e^2 + \sigma_m^2(\sigma_e^2 + \mu_e^2)$ (see Appendix Section A.2.2.1 for a full derivation).

[27] A full derivation is presented in Appendix Section A.2.2.

2.4 Using the Lower Bound of Estimated Emissions

The previous section demonstrates that applying the tax based on the point estimate of the firm's emissions leads to the socially optimal level of abatement. However, this policy would also imply that roughly half of emitting firms would be taxed (or would be required to acquire permits) for more emissions than they actually generate, which could present a significant barrier to real-world implementation. This section explores a modified approach, in which the regulator taxes firms based on the lower bound of a confidence interval for their estimated emissions. This policy may achieve nearly the same level of abatement while substantially increasing the political feasibility of implementing estimate-based emissions pricing.

Let $\hat{E}_L$ be the lower bound of a confidence interval for the firm's estimated emissions. Assuming that emissions across a particular firm's facilities are identically distributed and that sampled facilities are independently drawn,[28] and that a sufficient number of facilities are measured (i.e. $M > 30$), the Central Limit Theorem establishes that the distribution of the firm's total estimated emissions will approximate a normal distribution (see Figure 2). The probability that the firm is charged for more emissions than it is actually generating $\mathbb{P}(\hat{E} > E)$ then corresponds to the normal CDF.[29]

[28] The regulator can ensure this assumption holds by determining the sampling methodology.

[29] This section supposes that the probability the policy passes and remains in place is increasing in $\mathbb{P}(\hat{E} > E)$. While developing a political economy model of the decision-making process surrounding the introduction of new environmental policy is beyond the scope of this paper, it may be useful to explicate the intuition underlying this proposition. First, using a lower bound reduces the costs of compliance, which diminishes the natural gas industry's incentive to oppose the policy. Second, throughout the legislative process and potential legal challenges, the argument that individual firms may be unfairly taxed is less likely to be impactful if the probability of over taxation for any particular firm is low. Finally, firms are less likely to devote resources to independently measuring their emissions as part of a legal challenge if there is a low probability that their independently measured emissions will be below $\hat{E}_L$.

Figure 2: With a sufficiently large sample size, the distribution of estimates of the firm's total emissions approaches a normal distribution.

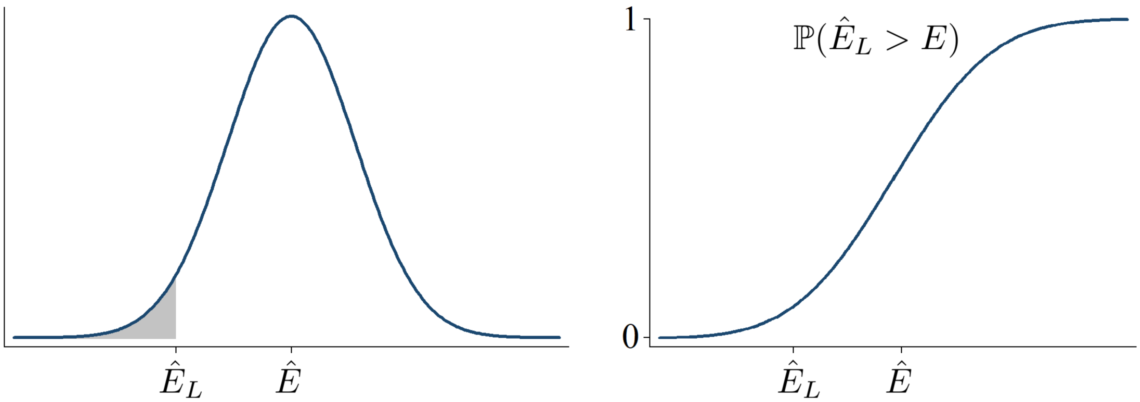

The modified estimate of the firm's total emissions becomes:

$$\hat{E}_L = N\Big(\tfrac{1}{M}\sum_{j=1}^{M} e_j - Z\frac{S}{\sqrt{M}}\Big) \tag{11}$$

where S is the sample standard deviation among measured facilities and Z is the Z-statistic for the desired level of confidence.[30] Entering $\hat{E}_L$ into the firm's expected profit function, the component of $\hat{E}_L$ representing the sample average simplifies as in Equation 6, and the firm's optimization problem can be written as:

$$\mathbb{E}[\pi] = \max_{\mathbf{e}_i} \sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i) - Te_i] + \tfrac{TNZ}{\sqrt{M-1}\sqrt{M}}\mathbb{E}\Big[\sqrt{\textstyle\sum_{j=1}^{M}(e_j - \frac{1}{M}\sum_{j=1}^{M} e_j)^2}\Big] \tag{12}$$

Because optimal emissions at facility i now depend on emissions at all of the firm's other facilities, the first-order condition for e_i becomes highly complex. However, by assuming that in expectation, emissions at all facilities other than i are equal to the average emissions across all other facilities (i.e. $e_k = e, \ \forall \ k \neq i$), the first-order

[30] For example, $Z = 1.96$ corresponds to a 2.5% chance that the firm is charged in excess of its actual emissions.

Figure 3: Marginal abatement cost curve for methane emissions the production segment of the natural gas industry, adapted from Marks (2019). If the firm's incentive to reduce emissions at a particular facility is slightly lower than the social cost of methane, it will still be incentivized to undertake almost the same amount of abatement.

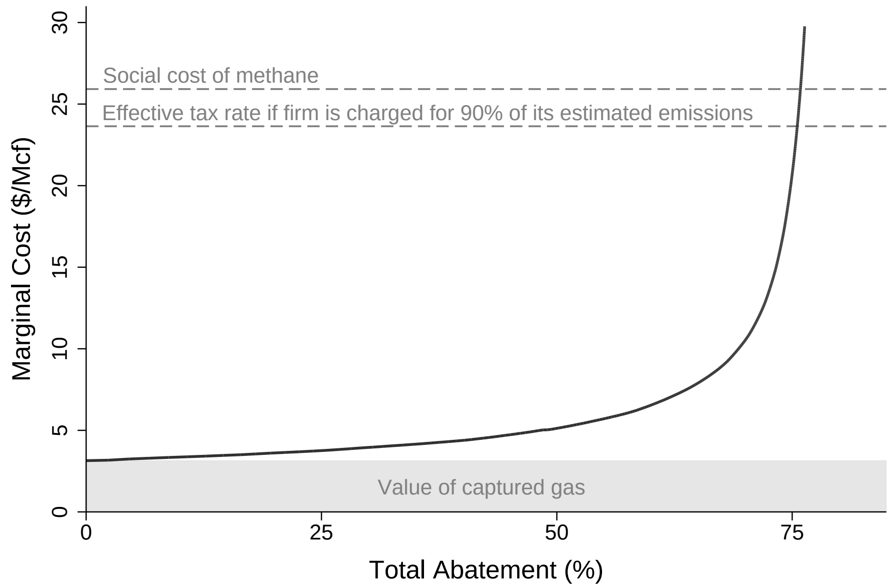

condition for e_i becomes analytically tractable:[31]

$$\frac{-\partial c}{\partial e_i} = P + T - TZ_{\frac{e_i - e}{S\sqrt{M}}} \tag{13}$$

This first-order condition implies that the firm will abate slightly more at facilities where emissions are lower than average and slightly less at facilities where emissions are higher than average. In other words, the firm's incentives are distorted toward increasing the variance of emissions across its facilities, causing a divergence from the socially optimal level of emissions at each facility. However, because $e_i - e$ is of the same magnitude as S, for any reasonably large M, this distortion will be small relative to the impact of the tax in incentivizing the firm to expend more in reducing its emissions (i.e. $T > TZ_{\frac{e_i - e}{S\sqrt{M}}}$). Furthermore, because marginal abatement curves for methane emissions are convex, a small adjustment to firms' incentive to abate at

[31] A full derivation is presented in Appendix Section A.2.3.

a facility i will have an even smaller impact on e_i, as demonstrated in Figure 3.

2.5 Simulation: Data

Without an existent policy or pilot program to assess, the best strategy for predicting the impacts of a potential regulation is to simulate its implementation.[32] Using data from various federal government agencies and the industry data provider DrillingInfo, I reconstruct the distribution of facilities by firm for the production, gathering, processing, transmission, and storage segments of the U.S. natural gas industry. I use data from recent scientific studies that have employed direct measurement technologies to estimate distributions of emissions across these industry segments. Lastly, I draw upon estimated marginal abatement cost curves from recent economics literature to predict emission reductions and abatement costs under an estimate-based emissions pricing program.

2.5.1 Distributions of Facilities by Firm

Table 1 provides an overview of data sources used in this study. For the production segment, I collect data from the industry data provider DrillingInfo, which provides a comprehensive overview of all oil and gas production operations in the United States. I focus on the 585,000 onshore oil and gas wells listed as "active" that had nonzero gas production over the prior 12 months. Because parent firms are not reported in these data, I used the reported operator as a proxy for ownership.[33] While thousands of firms operate only a small handful of wells, the majority of production is concentrated in larger firms: About 86 percent of natural gas is produced by firms that own at

[32] There is a rich literature on non-point source pollution in particular that uses simulations to estimate the effectiveness of potential policy solutions (e.g. Hung & Shaw 2005; Farzin & Kaplan 2004; Kaplan *et al.* 2003).

[33] Because there will be greater aggregation of wells at the ultimate parent level, this will lead estimates for M^* (and correspondingly for the total cost of measurement) will be conservative.

Table 1: Data sources by industry segment

	Facilities	**Emissions**	**Abatement Costs**
Production	DrillingInfo	Alvarez et al. (2018)	Marks (2019)
Gathering	Marchese et al. (2015)	Mitchell et al. (2015)	Marks (2019)
Processing	EPA, EIA	Mitchell et al. (2015)	Marks (2019)
Transmission	EPA, FERC	Subramanian et al. (2015)	Mayfield et al. (2017)
Storage	EPA, EIA	Subramanian et al. (2015)	Mayfield et al. (2017)

least 100 wells, and about 56 percent is produced by firms that own more than 1,000 wells.

For the processing, transmission, and storage segments, I use distributions of facilities by firm from the U.S. Environmental Protection Agency (EPA), Energy Information Administration (EIA), and Federal Energy Regulatory Commission (FERC).[34] Where production firms typically operate hundreds to thousands of smaller sites dispersed across large geographic areas, these three segments are characterized by larger facilities that are effectively more concentrated emission sources. A typical processing plant is owned by a firm that owns 11 plants in total, a typical transmission compression station is owned by a firm that owns about 50 stations total, and a typical storage compression station is operated by a firm that owns about 13 stations total (see Column 3 of Table 2).

The gathering segment consists of localized networks of small-diameter pipelines connecting individual wells to processing plants and interstate pipelines. These networks are interspersed with compression stations to pressurize the gas for transport, and they typically also include smaller processing sites and short-term storage facilities. Likely due to this amorphous physical structure, there is no comprehensive dataset of gathering facilities available at this time. Instead, I rely on a partial dataset

[34] I access these data through the Homeland Infrastructure Foundation-Level Database (HIFLD). For a handful of natural gas storage sites, HIFLD uses on data collected from corporate websites in place of government sources.

Table 2: Distributions of facilities by firm and emissions by industry segment, with a comparison to estimates from Alvarez *et al.* (2018)

	Total Facilities	**Total Firms**	**Facilities per Firm**	**Facilities per Firm** *(Weighted*)*	**Total Emissions** *(Modeled)* (10^3tCO$_2$e/yr)	**Total Emissions** *(Alvarez et al.)* (10^3tCO$_2$e/yr)
Production	585,514	7,655	76 (560)	4,180 (5,411)	268,200 (34,000)	244,800 (+54,400/-54,400)
Gathering	4,548	477	10 (28)	89 (123)	85,800 (9,600)	88,400 (+20,400/-6,800)
Processing	537	189	3 (5)	11 (12)	25,100 (3,300)	24,500 (+6,800/-2,400)
Transmission	1,360	71	19 (24)	48 (30)	36,200 (4,300)	52,400 (+14,600/-11,200)
Storage	454	105	4 (6)	14 (11)	15,500 (2,400)	

*Column 4 shows a weighted average, indicating the average number facilities owned by the operator of a typical facility. Values in parentheses are standard deviations in Columns 2-4, standard error in Column 5, and 95% CI in Column 6. Note that Alvarez et al. combine the transmission and storage segments.

constructed by Marchese *et al.* (2015), who use a combination of methods to separate out gathering facilities from larger state-level datasets of oil and gas infrastructure.[35] Marchese et al. identify 2,519 gathering facilities (with ownership) in eight states, and extrapolate these results to estimate that there are 4,549 (+921/-703) total gathering facilities in the U.S.[36] Each gathering facility as delineated by Marchese et al. is an individual compression, processing, and/or storage site (implying direct measurement is feasible),[37] and a typical facility is owned by a firm that owns about 90 such

[35] In particular, Marchese et al. make an initial determination of each listed facility's industry segment based on the primary segment in which the facility's owner operates, then verify these initial assignments using expert opinion and satellite imagery. The authors are explicit in acknowledging and quantifying sources of uncertainty in their methodology; however, for simplicity, the simulation in this paper only uses their point estimates.

[36] To approximately capture the distribution of ownership of gathering facilities that are not directly identified by Marchese et al., I randomly duplicate 2,030 facilities, clustered by firm. To the extend that ownership of gathering facilities crosses state boundaries, the actual distribution of ownership will be more concentrated, leading my estimates of M^* and the cost of measurement to be conservative (i.e. overstated).

[37] While some emissions are produced by leaks from gathering lines outside of these facilities, Marchese et al. estimate that they account for less than 10% of total gathering segment emissions.

facilities overall.

2.5.2 Distributions of Emissions

To estimate methane emissions for the facilities in my dataset, I rely on direct measurement data from a number of recent scientific measurement studies. I replicate methodologies used by these studies (and subsequent scientific papers that applied improved methods to the same data) to construct segment-specific distributions of emissions to sample from in the simulation model. While these replications simplify some elements of the more complex methodologies they are based on, they successfully capture the approximate scale and distribution of methane emissions across the natural gas supply chain.

For the production segment, I construct distributions of emissions conditional on well-level output using emission factors estimated by Alvarez *et al.* (2018). Alvarez et al. synthesize data from four recent scientific studies that used direct site-level measurement technologies to quantify emissions (i.e. the same technologies that could be used by a regulator to implement estimate-based emissions pricing).[38] Using measurements taken at 433 individual well sites across six major gas producing basins, they fit lognormal distributions of emissions conditional on production for each basin. The authors cite Zavala-Araiza *et al.* (2015)'s observation that a lognormal distribution of emissions is consistent with a system where many independent stochastic events multiplicatively determine the occurrence and magnitude of emissions, and they empirically verify that methane emissions are lognormally distributed in the datasets they use. To save computing time, the authors use these estimated distributions to predict emission factors with uncertainty intervals for six sets of 600 production cohorts, which they then randomly sample in a simulation model to predict aggregate

[38] Alvarez et al. also validate their results using data from eight recent top-down measurement studies that used aircraft, satellite, or tower measurements to quantify emissions.

emissions across all U.S. wells.[39]

In evaluating methane emissions from the remaining sectors, I follow previous scientific literature in grouping together the gathering and processing sectors and the transmission and storage sectors. Gathering and processing facilities are spatially linked, and there is often overlap in their roles in the natural gas supply chain. Accordingly, their equipment and emission profiles are similar. I utilize direct measurements of 114 gathering facilities and 16 processing plants from Mitchell *et al.* (2015) to estimate distributions of emissions for all sites in my dataset.[40] For the gathering segment, aggregate emissions from this approach closely align with the estimates in Alvarez *et al.* (2018) without adjustment.[41] For the processing segment, however, this approach generates an estimate of aggregate emissions that is about 1.5 times higher than the more reliable estimate from Alvarez et al., which is outside of its 95% confidence interval.[42] I correct for this disparity by rescaling the distribution of processing plant emissions by the inverse ratio of the two estimates.

For the transmission and storage segments, I use direct measurements of methane emissions at 45 compression stations from Subramanian *et al.* (2015). I again estimate parameters of a lognormal distribution to draw from in the simulation, then scale up storage segment emissions by the ratio of average storage facility emissions to average

[39] Note that in addition to methane emissions from completed, onshore gas-producing wells, Alvarez et al. also estimate emissions from a number of other production sources, including offshore oil and gas platforms, abandoned wells, and well completions and workovers. These other sources collectively account for about 10% of total estimated production emissions. While it would probably be desirable to cover these sources under an emissions pricing program, they are estimated in ways that do not rely on direct site-level measurements and therefore do not fit easily into the modeling framework used in this paper. These sources are excluded from this analysis, leading estimates of total abatement potential to be conservative.

[40] Following Alvarez *et al.* (2018), I assume emissions are lognormally distributed and recover the mean and variance of the logged distribution of measured facilities (separately for gathering and processing). I then use these parameters to randomly draw emissions for each facility in each iteration of the simulation.

[41] The estimates for gathering and processing emissions in Alvarez *et al.* (2018) are drawn from two more complex modeling studies that use the same data, Marchese *et al.* (2015) and Zavala-Araiza *et al.* (2015).

[42] This is likely due primarily to the facilities sampled by Subramanian et al. not being representative of the overall distribution of U.S. processing facilities.

processing plant emissions as estimated by the more complex modeling framework used by Zimmerle *et al.* (2015). As shown in Columns 5 and 6 of Table 2, without further adjustment this approach generates aggregate emission estimates for these sectors that are consistent with the estimate in Alvarez et al.[43]

2.5.3 Abatement and Abatement Costs

The simulation relies on marginal abatement cost curves (MACCs) drawn from two recent economics papers to predict abatement and abatement costs. For the production sector, I use the MACC from Marks (2019), which is econometrically estimated by examining how production firms' abatement behavior responds to natural gas price shocks. While this approach is generally well-suited for predicting the effects of emissions pricing (since a price increase has the same impact on firms' opportunity cost of allowing gas to be emitted as would an emissions tax of the same amount), one limitation is that it is estimated using data from the EPA Greenhouse Gas Reporting Program, an inventory that captures less than half of total U.S. production emissions (Alvarez *et al.*, 2018).

For the transmission and storage sector, I use the MACC estimated in Mayfield *et al.* (2017). This MACC is constructed by using engineering cost estimates of various potential abatement activities as inputs to a simulation model where a social planner implements the lowest-cost abatement opportunities first. Their model focuses specifically on the transmission and storage sectors and incorporates uncertainty in emission factors, operating hours, and abatement costs. For the gathering and processing segments, which have not yet been studied in the economics literature on methane emissions, I apply the MACC from Marks (2019), since the distributions of emitting sources for gathering the processing segments are more similar to the

[43] Note that emissions from LNG import/export terminals and from leaks along transmission pipelines are excluded (see Footnote 39). Combined, these two sources generate about 10,000 tCO_2e of additional methane emissions from the transmission sector annually.

production segment than to the transmission and storage segments.

One inconsistency in using these MACCs for this application is that they are estimated for emissions reductions across entire sectors, but the simulation framework applies them at the site level. However, the increasing, convex structure of these sector-level MACCs inherently also applies to site-level emissions (which similarly face diminishing returns to abatement activities), and all of the outcomes of interest from the simulation model are aggregated to the sector level.

2.6 Simulation: Methodology and Results

I perform the simulation model for 100 iterations to incorporate uncertainties in emissions and abatement. Within each iteration, the basic structure of the simulation model consists of randomly assigning facility-level emissions from the distributions described above, randomly assigning facility-level abatement under an estimate-based emissions pricing program using the MACCs above, calculating the optimal number of sampled facilities by firm, and recovering abatement, measurement costs, and various other outputs.

I set the emissions tax at \$42/tCO_2e, which corresponds to the EPA's estimate for the social cost of emissions generated in 2020 assuming a 3% discount rate. I assume the regulator uses tracer flux as the measurement technology, which has a cost of \$3,000 per site and implies site-level measurement error of $\sigma = 0.145$. Finally, I assume $\alpha = 0.5$, which corresponds to a scenario where the regulator internalizes overcharging firms at a one-to-one ratio, but does not internalize undercharging firms.

For the production segment, distributions of emissions are constructed within each iteration using the conditional emission factors estimated by Alvarez *et al.* (2018). Following their methodology, I randomly assign each of the 55 basins in my dataset to one of the six basins from their study and separate wells by production levels into 600 cohorts. I then randomly sample emissions for each well from a lognormal distribution parameterized by the mean and standard deviation of the corresponding basin-cohort emission factor reported in their supplemental data. For the other segments, site-level emissions are randomly sampled from unconditional lognormal distributions parameterized as described above.[44] While this methodology incorporates the correct level of uncertainty at the site level, it does not account for uncertainty in aggregate methane emissions. To incorporate uncertainty in aggregate emissions estimated by the most recent scientific literature, I adjust emissions at all facilities within each segment by a factor drawn from a normal distribution with a standard deviation calculated from the confidence intervals reported in Alvarez *et al.* (2018).[45]

To incorporate uncertainty in abatement activity and costs, within each iteration I randomly assign predicted abatement at the segment level using bootstrapped MACCs from Marks (2019) for the production, gathering, and processing segments and MACCs drawn from iterations of the simulation implemented in Mayfield *et al.* (2017) for the transmission and storage segments.[46] On average, abatement cost curves from both studies (independently) predict about 80 percent abatement under a methane price equivalent to a $42 carbon price, which carries through to the final results.

Because the regulator would not know the actual population standard deviation of emissions for individual firms, I suppose the regulator assumes that firm-level emis-

[44] Although neglecting to condition on facility size in the gathering, processing, transmission, and storage segments implies that within any given firm, the distribution of facility-level emissions will be inaccurate, the aggregate distribution of facility-level emissions across firms will be accurate.

[45] For example, for the gathering segment, I use Alvarez et al.'s point estimate (88,400 tCO_2e) and 95% confidence interval (+20,400/-6,800) to construct an average proportional standard deviation of 7.8 percent. I then multiply emissions at all gathering facilities within one iteration by $1 + \epsilon$, where $\epsilon \sim N(0, 0.078)$. This adjustment does not capture the skewed distributions of uncertainty in most industry segments, but appropriately captures the magnitude of the uncertainty.

[46] Although the theoretical model appropriately incorporates heterogeneous abatement costs at the facility level, randomly assigning MACCs at the facility level from the same distribution would artificially smooth out uncertainty from those estimates. Assigning MACCs at the segment level within each iteration appropriately incorporates the uncertainty from Marks and Mayfield et al.'s estimates.

Figure 4: Left: Optimal number of facilities to sample by firm size. Right: Actual vs. estimated emissions by firm.

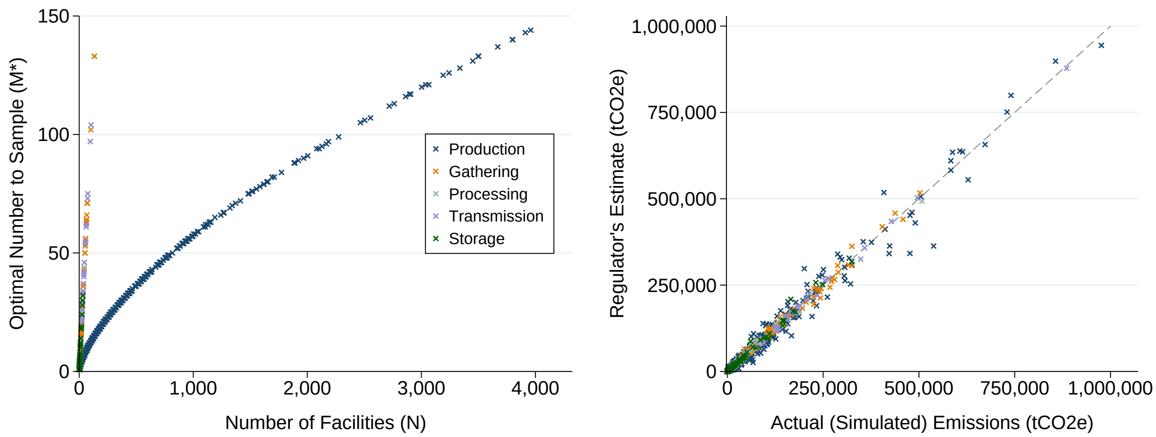

Note: For readability, results are shown for just one iteration, the left graph is truncated for M^* above 150 (23 firms), and the right graph is truncated for emissions above 1,000,000 tCO_2e (9 firms).

sions exhibit the population standard deviation of emissions within a given segment across all firms when calculating M^*. In cases where the optimal number of sampled facilities exceeds the number of facilities operated by the firm, I assume the regulator samples all facilities.[47] Interestingly, this is the case for all facilities in the processing, transmission, and storage industry segments in all iterations, and almost all facilities in the gathering segment. This is likely due to high standard deviations of emissions relative to the number of facilities operated by each firm. For the production segment, however, the simulation predicts that the regulator will optimally sample about only about 8% of facilities on average. Because the regulator assumes each firms' facilities exhibit the population standard deviation of emissions, the number of sampled facilities actually reduces to a direct function of N, as illustrated in Figure 4.[48] Given an estimate of M^* for each firm, I then randomly assign M sampled facilities for

[47] This is not equivalent to comprehensive monitoring, since the regulator is only taking a single snapshot measurement of emissions at a particular (randomly determined) time of the year. At this time, the theoretical model does not explicitly incorporate the timing of measurements. However, to the extent that emissions and abatement behaviors are dispersed over time, resampling the same facility at random intervals would generate the same incentive for firms to engage in cost effective abatement.

[48] In particular, the optimal number of facilities to sample can be reduced to $M^* = 0.57N^{\frac{2}{3}}$.

Figure 5: Simulated emissions under an estimate-based emissions pricing program

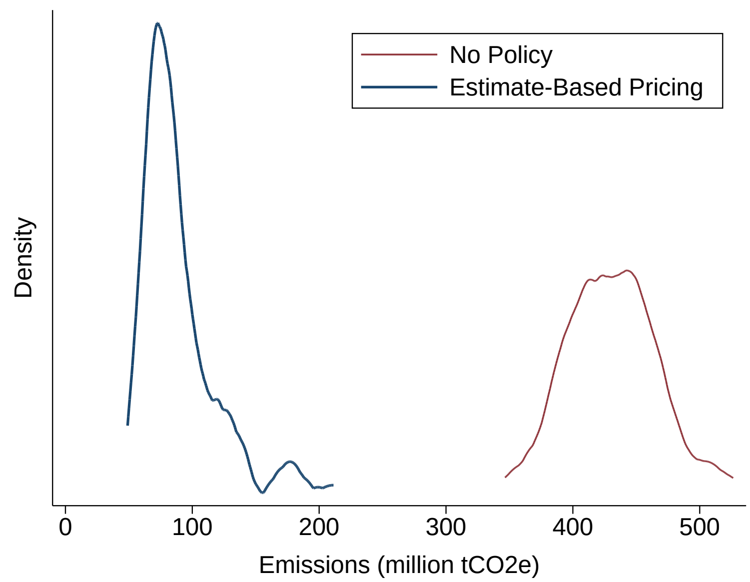

each firm and predict regulator emission estimates using the proportional uncertainty parameter σ_m. As shown in the right panel of Figure 4, there is some variation between actual emissions and the regulator's estimate, which is proportionally greater for smaller facilities.

Figure Figure 5 illustrates simulated total emissions under an estimate-based emissions pricing policy versus the status quo, and detailed results by sector are presented in Table 3.[49] The total amount of abatement predicted by the model is roughly 340 million tons of CO_2-equivalent emissions per year, which corresponds to avoided climate damages of roughly \$14 billion. The physical cost necessary to achieve that abatement (net of the value of recovered gas) is approximately \$1 billion dollars, and the annual cost of measurement using the tracer flux technology is approximately \$160 million. Measurement costs are particularly low in the processing, transmission, and storage sectors, where emissions are concentrated at relatively few individual facilities. Tax revenues for the roughly 90 million tCO_2 that would continue being emitted

[49] Total abatement cost is determined by the sum of the area under the MACC for all facilities (accounting for the wholesale value of conserved gas). Total climate benefit of avoided carbon emissions is simply calculated as total abatement times the level of the tax, and the total cost of measurement is calculated as c_m times M^*.

Table 3: Simulation results by industry segment

	Total Abatement ($10^3 tCO_2e$)	**Abatement Benefit** ($ millions)	**Abatement Cost** ($ millions)	**Total Cost of Measurement** ($ millions)	**Total M***
Production	211,254	8,872	671	123	40,910
	(39,986)	(1,679)	(105)	(36)	(11,933)
Gathering	68,474	2,875	213	13	4,502
	(10,794)	(453)	(40)	(0.3)	(88)
Processing	19,851	833	63	2	537
	(3,221)	(135)	(11)	-	-
Transmission	29,158	1,224	125	4	1,360
	(4,570)	(191)	(19)	-	-
Storage	12,542	526	53	1	454
	(2,313)	(97)	(9)	-	-
Total	341,000	14,334	1,127	143	47,763
	(39,000)	(1,643)	(121)	(36)	(11,941)

under optimal methane policy are approximately $3 billion, implying that this program would generate net revenues for the government. The overall benefit-to-cost ratio predicted by the simulation model is approximately 11:1.

2.7 Conclusion

This paper has outlined a novel strategy for implementing emissions pricing in a setting where emissions are costly to measure. A theoretical framework demonstrates that in a stylized setting where firms have an arbitrarily large number of facilities, this approach preserves the efficiency benefits of emissions pricing. Furthermore, a modified policy in which firms would be charged based on the lower bound of a confidence interval for their estimated emissions implies only a modest divergence from the optimal level of abatement.

The simulation model presented in this paper, while coarse, characterizes the approximate magnitude of the expected costs and benefits of using estimate-based emissions pricing to address methane emissions from the U.S. natural gas industry. In general, the benefits far exceed the costs, which is consistent with previous literature in economics and engineering that finds that methane emissions from the natural gas supply chain are a highly efficient area to prioritize for greenhouse gas mitigation. The additional cost of measurement is about an order of magnitude lower than the physical cost of abatement, and about two orders of magnitude lower than the total abatement benefit. Furthermore, the main parameter determining the cost of these estimates is based on small-scale scientific measurement studies. It is likely that the per-site cost of measurement would be much lower for a large-scale program due to economies of scale, and measurement costs will likely decline over time as methodologies are improved and new measurement technologies are developed.

While an estimate-based approach to emissions pricing is promising, further research is needed to establish its feasibility and effectiveness in practice. There are many possible extensions and improvements that could be incorporated into the basic theoretical framework presented here. For example, it would be highly useful to adapt the model to finite sample theory to more accurately capture firms with only a few facilities, which make up a large portion of the overall distribution. Additionally, future research could explore improved sampling methodologies beyond random selection, such as using satellite measurements to identify high-emitting facilities for sampling.[50] Furthermore, as with previously introduced strategies for addressing other non-point source pollution problems, laboratory experiments could be conducted to verify that participants appropriately internalize expected emissions taxes across facilities. Finally, field experiments would be highly valuable both for verifying behavioral responses at the firm level and for developing better estimates of

[50] Note that any modification that changes the distribution of sampled facilities would need to explicitly correct for bias introduced by using a non-random sample.

the costs and benefits of this policy.

References

Alvarez, Ramón A, Zavala-Araiza, Daniel, Lyon, David R, Allen, David T, Barkley, Zachary R, Brandt, Adam R, Davis, Kenneth J, Herndon, Scott C, Jacob, Daniel J, Karion, Anna, *et al.* 2018. Assessment of methane emissions from the US oil and gas supply chain. *Science*, eaar7204.

Anderson, Soren T, Kellogg, Ryan, & Salant, Stephen W. 2018. Hotelling under pressure. *Journal of Political Economy*, **126**(3), 984–1026.

Atherton, Emmaline, Risk, David, Fougère, Chelsea, Lavoie, Martin, Marshall, Alex, Werring, John, Williams, James P, & Minions, Christina. 2017. Mobile measurement of methane emissions from natural gas developments in northeastern British Columbia, Canada. *Atmospheric Chemistry & Physics*, **17**(20).

Baumol, William J, Baumol, William Jack, Oates, Wallace E, Bawa, Vijay S, Bawa, WS, & Bradford, David F. 1988. *The theory of environmental policy.* Cambridge university press.

Blundell, Wesley, Gowrisankaran, Gautam, & Langer, Ashley. 2018. *Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations.* Tech. rept. National Bureau of Economic Research.

Brandt, Adam R, Heath, GA, Kort, EA, O'sullivan, F, Pétron, G, Jordaan, SM, Tans, P, Wilcox, J, Gopstein, AM, Arent, D, *et al.* 2014. Methane leaks from North American natural gas systems. *Science*, **343**(6172), 733–735.

Burtraw, Dallas, & Szambelan, Sarah Jo. 2010. *US emissions trading markets for SO2 and NOx.* New York: Routledge.

Carlson, Curtis, Burtraw, Dallas, Cropper, Maureen, & Palmer, Karen L. 2000. Sulfur dioxide control by electric utilities: What are the gains from trade? *Journal of political Economy*, **108**(6), 1292–1326.

Cochard, François, Willinger, Marc, & Xepapadeas, Anastasios. 2005. Efficiency of nonpoint source pollution instruments: an experimental study. *Environmental and Resource Economics*, **30**(4), 393–422.

Dowd, Brian M, Press, Daniel, & Los Huertos, Marc. 2008. Agricultural nonpoint source water pollution policy: The case of California's Central Coast. *Agriculture, ecosystems & environment*, **128**(3), 151–161.

Duflo, Esther, Greenstone, Michael, Pande, Rohini, & Ryan, Nicholas. 2013. Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India. *The Quarterly Journal of Economics*, **128**(4), 1499–1545.

EIA. 2018. *Annual Energy Outlook 2018 with Projections to 2050.* US Energy Information Administration.

Farzin, Y Hossein, & Kaplan, Jonathan D. 2004. Nonpoint source pollution control under incomplete and costly information. *Environmental and Resource Economics*, **28**(4), 489–506.

Hansen, Lars Gårn. 2002. Regulation of non-point emissions–a variance based mechanism. *Environmental and Resource Economics*, **21**(4), 303–316.

Howarth, Robert W, Santoro, Renee, & Ingraffea, Anthony. 2011. Methane and the greenhouse-gas footprint of natural gas from shale formations. *Climatic Change*, **106**(4), 679.

Hung, Ming-Feng, & Shaw, Daigee. 2005. A trading-ratio system for trading water pollution discharge permits. *Journal of Environmental Economics and Management*, **49**(1), 83–102.

ICF. 2016 (May). *Economic Analysis of Methane Emission Reduction Potential from Natural Gas Systems*. Tech. rept. ICF International. Prepared for ONE Future Inc.

IPCC. 2014. *Climate Change 2014: Mitigation of Climate Change*. Working Group III Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.

Kaplan, Jonathan D, Howitt, Richard E, & Farzin, Y Hossein. 2003. An information-theoretical analysis of budget-constrained nonpoint source pollution control. *Journal of Environmental Economics and Management*, **46**(1), 106–130.

Lowell, Dana, Jones Brian Seamonds David, & Russell, Pye. 2016. *Analysis of Pipeline and Hazardous Materials Safety Administration Proposed New Safety Rules: Pipeline Blowdown Emissions and Mitigation Options*. Tech. rept. M.J. Bradley and Associates LLC. Prepared for Environmental Defense Fund and Pipeline Safety Trust.

Marchese, Anthony J, Vaughn, Timothy L, Zimmerle, Daniel J, Martinez, David M, Williams, Laurie L, Robinson, Allen L, Mitchell, Austin L, Subramanian, R, Tkacik, Daniel S, Roscioli, Joseph R, *et al.* 2015. Methane emissions from United States natural gas gathering and processing. *Environmental science & technology*, **49**(17), 10718–10727.

Marks, Levi. 2019 (April). *The Abatement Cost of Methane Emissions from Natural Gas Production*. Working Paper.

Mayfield, Erin N, Robinson, Allen L, & Cohon, Jared L. 2017. System-wide and superemitter policy options for the abatement of methane emissions from the US natural gas system. *Environmental science & technology*, **51**(9), 4772–4780.

Mitchell, Austin L, Tkacik, Daniel S, Roscioli, Joseph R, Herndon, Scott C, Yacovitch, Tara I, Martinez, David M, Vaughn, Timothy L, Williams, Laurie L, Sullivan, Melissa R, Floerchinger, Cody, *et al.* 2015. Measurements of methane emissions from natural gas gathering facilities and processing plants: Measurement results. *Environmental science & technology*, **49**(5), 3219–3227.

Montgomery, W David. 1972. Markets in licenses and efficient pollution control programs. *Journal of economic theory*, **5**(3), 395–418.

Munnings, Clayton, & Krupnick, Alan. 2017. *Comparing Policies to Reduce Methane Emissions in the Natural Gas Sector*. Tech. rept. Resources for the Future. Washington, DC.

Newell, Richard G, & Stavins, Robert N. 2003. Cost heterogeneity and the potential savings from market-based policies. *Journal of Regulatory Economics*, **23**(1), 43–59.

Omara, Mark, Sullivan, Melissa R, Li, Xiang, Subramanian, R, Robinson, Allen L, & Presto, Albert A. 2016. Methane emissions from conventional and unconventional natural gas production sites in the Marcellus Shale Basin. *Environmental science & technology*, **50**(4), 2099–2107.

Poe, Gregory L, Schulze, William D, Segerson, Kathleen, Suter, Jordan F, & Vossler, Christian A. 2004. Exploring the performance of ambient-based policy instruments when nonpoint source polluters can cooperate. *American Journal of Agricultural Economics*, **86**(5), 1203–1210.

Robertson, Anna M, Edie, Rachel, Snare, Dustin, Soltis, Jeffrey, Field, Robert A, Burkhart, Matthew D, Bell, Clay S, Zimmerle, Daniel, & Murphy, Shane M. 2017. Variation in methane emission rates from well pads in four oil and gas basins with contrasting production volumes and compositions. *Environmental science & technology*, **51**(15), 8832–8840.

Segerson, Kathleen. 1988. Uncertainty and incentives for nonpoint pollution control. *Journal of environmental economics and management*, **15**(1), 87–98.

Subramanian, R, Williams, Laurie L, Vaughn, Timothy L, Zimmerle, Daniel, Roscioli, Joseph R, Herndon, Scott C, Yacovitch, Tara I, Floerchinger, Cody, Tkacik, Daniel S, Mitchell, Austin L, *et al.* 2015. Methane emissions from natural gas compressor stations in the transmission and storage sector: Measurements and comparisons with the EPA greenhouse gas reporting program protocol. *Environmental science & technology*, **49**(5), 3252–3261.

Suter, Jordan F, & Vossler, Christian A. 2013. Towards an understanding of the performance of ambient tax mechanisms in the field: evidence from upstate New York dairy farmers. *American Journal of Agricultural Economics*, **96**(1), 92–107.

Tyner, David R, & Johnson, Matthew R. 2018. A Techno-Economic Analysis of Methane Mitigation Potential from Reported Venting at Oil Production Sites in Alberta. *Environmental Science & Technology*.

Xepapadeas, Anastasios. 2011. The economics of non-point-source pollution. *Annu. Rev. Resour. Econ.*, **3**(1), 355–373.

Zavala-Araiza, Daniel, Lyon, David R, Alvarez, Ramon A, Davis, Kenneth J, Harriss, Robert, Herndon, Scott C, Karion, Anna, Kort, Eric Adam, Lamb, Brian K, Lan, Xin, *et al.* 2015. Reconciling divergent estimates of oil and gas methane emissions. *Proceedings of the National Academy of Sciences*, **112**(51), 15597–15602.

Zavala-Araiza, Daniel, Alvarez, Ramón A, Lyon, David R, Allen, David T, Marchese, Anthony J, Zimmerle, Daniel J, & Hamburg, Steven P. 2017. Super-emitters in natural gas infrastructure are caused by abnormal process conditions. *Nature communications*, **8**, 14012.

Zimmerle, Daniel J, Williams, Laurie L, Vaughn, Timothy L, Quinn, Casey, Subramanian, R, Duggan, Gerald P, Willson, Bryan, Opsomer, Jean D, Marchese, Anthony J, Martinez, David M, *et al.* 2015. Methane emissions from the natural gas transmission and storage system in the United States. *Environmental science & technology*, **49**(15), 9374–9383.

Chapter 3

Vertical Market Power in Interconnected Natural Gas and Electricity Markets

with Charles F. Mason, Kristina Mohlin, & Matthew Zaragoza-Watkins

3.1 Introduction

Natural gas has replaced coal as the predominant electricity generation resource in the United States and provided more than a third of the country's utility-scale electricity generation in 2016. With half the burn-point emissions of other fossil fuels and physical properties that make it ideal for balancing out the intermittency of renewables, its share of the generation mix is only expected to increase in the immediate future.[1] With over 50% of its total generation already coming from gas, New England is at the leading edge of this transition and an ideal environment in which to explore issues that may arise from the growing interdependencies between gas and electricity markets.

In recent years, New England's wholesale natural gas and electricity markets have experienced severe, concurrent price spikes. During the months of extreme cold that marked the winter of 2013-14 (i.e., the "Polar Vortex"), for example, New England gas prices averaged \$17.86 per MMBtu (million British Thermal Units) —almost four times the Henry Hub price—and reached a record high of \$78/MMBtu on January 22,

[1] See EIA (2017).

2014.[2] These extreme price spikes have been commonly attributed to limited pipeline capacity serving New England, and this "scarce capacity" narrative has been used in recent proposals to expand natural gas pipeline capacity serving the region.[3]

Limited pipeline capacity is indeed partly responsible for these extreme prices. But we also find strong evidence that two firms that held significant shares of the contracts to flow gas on the Algonquin Gas Transmission Pipeline—one of the two major pipelines serving New England—regularly restricted capacity to the region by scheduling deliveries without actually flowing gas. These unusual scheduling practices tied up capacity that, in a well-functioning market, should have been released, or would have otherwise made available, to other shippers. Instead, significant quantities of pipeline capacity went unutilized on many of the coldest days of the year, pushing up the price of gas.

While most shippers had little incentive to sacrifice revenue from gas sales by withholding capacity, the two firms observed to withhold capacity also own large portfolios of electric generation units located in the region, giving them an incentive to increase gas prices in order to *raise rivals' costs* (Salop & Scheffman, 1983). That is, by restricting sales of a necessary input to production for their downstream competitors in the wholesale electricity market, the capacity-withholding firms increased the quantity of electricity their largely non-gas units were called upon to generate and the price those units earned.

In this paper, we analyze three recent years of scheduling data on Algonquin for evidence of firms withholding pipeline capacity in this manner. We find clear patterns of withholding at a subset of delivery nodes operated by Avangrid and Eversource (henceforth referred to as Firm A and Firm B), the only two firms operating on the pipeline with substantial assets and operations in both the gas distribution market

[2] The natural gas price at Henry Hub in Lousiana is the benchmark price for natural gas traded in the United States and is considered reflective of the commodity value without transportation costs.
[3] See, for example, Rose *et al.* (2014) and ICF (2015).

and the electricity generation market. Using a panel data regression model, we empirically demonstrate these nodes were disproportionately served by specialized types of contracts that allow firms to call for gas on demand and to make large adjustments, without notice, in the last few hours of the day, two necessary conditions for executing a penalty-free withholding strategy. Over our three year study period, aggregate withholding at these few nodes reduced the pipeline's effective capacity by approximately 50,000 MMBtu per day, on average. On 37 days, over 100,000 MMBtu of capacity—about 7% of the pipeline's total daily capacity and about 28% of the daily capacity that is typically used to supply gas-fired generators—was withheld at these nodes.

This behavior significantly impacted both natural gas and electricity prices in New England. We employ an instrumental variables model to estimate a counterfactual gas price series, finding that gas prices were \$1.68/MMBtu (39%) higher on average during our entire study period and \$3.82/MMBtu (68%) higher during the winters. We proceed to construct a simulation model of New England's wholesale electricity market and use our estimated counterfactual gas price series as an input to estimate the effect on the region's electricity market. We find that electricity prices were about \$10/MWh (19%) higher on average over our study period due to capacity withholding. Our simulation predicts that underutilized pipeline capacity ultimately resulted in a transfer from New England electricity ratepayers to generators (and their fuel suppliers) of about \$3.4 billion over the course of our study period, about half of which occurred during the particularly cold winter of 2013-14. While the studied behavior may have been within the two firms' contractual rights, the significant impacts in both the gas and electricity markets show the need to consider improvements to market design and regulation as these two energy markets become increasingly interlinked.

The remainder of this paper is structured as follows: Section 2 reviews the relevant literature on raising rivals' costs and similar market power scenarios in electricity markets. In section 3 we provide background on the three markets that comprise our institutional setting. Section 4 presents the theoretical framework that forms the basis for our empirical strategy. Our empirical analysis is broken into two parts: Section 5 presents our investigation into firms' patterns of capacity withholding and section 6 presents our estimations of gas and electricity market impacts. Section 7 concludes.

3.2 Literature Review

Market power refers to the ability of a firm (or group of firms) to raise and maintain price above the level that would prevail under competition. Like all network utilities, energy transportation infrastructure is characterized by large initial capital investments and spatial differences in supply and demand that create an environment susceptible to the exercise of market power. When transmission constraints bind, they effectively segment the network into a set of smaller markets wherein firms that don't own or control a significant share of total assets across the network may have significant local abilty to set prices (Borenstein *et al.*, 1995). So far, this situation has mostly been studied in the context of electricity markets and much less so for gas markets.

Network congestion fluctuates with demand, meaning markets may be highly concentrated at some times and highly competitive at others. Consequently, Borenstein *et al.* (1999) discourage applying traditional measures of market concentration such as the Herfindahl-Hirsch Index (HHI) to electricity markets. Instead, they suggest modeling energy markets to investigate whether firms employ strategic behavior in their production decisions. Borenstein & Bushnell (1999) employ this method to predict significant potential for market power at the outset of the deregulation of the California electricity market and again to empirically verify ex-post the majority contribution of market power to California's extremely costly 2000 energy crisis (Borenstein *et al.*, 2002).

The instance of market power discussed in this paper stems from the contracts that serve as property rights to natural gas transportation capacity, which are in many ways analogous to transmission rights in electricity markets. Joskow & Tirole (2000) analyze the interaction between transmission rights and market power in electricity markets. They present a model of a two-node grid, where an upstream node with many competitive, low-cost generators is separated by a single transmission line from a downstream node where a single firm controls more expensive generation resources. Different marginal costs lead the independent system operator (ISO) to pay different prices at each node, which enhances efficiency in ideal conditions but also introduces the possibility of gaming the system. In this setting, if the downstream generator obtains physical transmission rights (which allocate capacity for generators to use to transmit electricity at no additional cost), inefficiency may arise. Under some realistic conditions, the downstream generator finds it more profitable to use physical rights to withhold transmission capacity to increase the downstream node's price, leading to welfare losses due to productive inefficiency. Further, the downstream generator is incentivized to acquire all the physical rights so they can simultaneously decide transmission capacity to and production at the downstream node.

Joskow and Tirole's analysis provides an interesting parallel to our setting, where firms that own downstream electric generation and pipeline capacity rights are, under some conditions, incentivized to use those rights to tie up capacity. Interestingly, Joskow & Tirole (2000) advocate for adapting the capacity release regulations of the gas pipeline industry to the electricity market to mitigate the potential abuse of physical rights in this manner. However, our study clearly shows that capacity release rules as they stand are insufficient to overcome the incentives toward inefficiency that are created by physical transportation constraints.

Cremer & Laffont (2002) adapt Joskow and Tirole's two-node, two-producer electricity model to natural gas to show similar results, although their model is limited in the depth to which it incorporates the institutional differences of the gas market. A much more heavily studied area of market power in natural gas is the supply-side market concentration in the European gas market, which imports a majority of its gas from only three countries – Russia, Norway, and Algeria (see e.g., Lise & Hobbs 2009, Boots *et al.* 2003, Holz *et al.* 2008).

While the ability to influence prices emerges from the physical capacity constraint in our setting, the firm's primary incentive to withhold capacity comes from vertical integration across the gas and electricity markets. One commonly-studied concern in the literature on vertical market power is *foreclosure* (sometimes also termed *raising rivals costs*), wherein a vertically-integrated firm instructs its upstream entity to restrict sales of a necessary production input to its downstream entity's competitors to increase the prices and market share enjoyed by that arm of the firm (e.g., Hart *et al.* 1990, Ordover *et al.* 1990).

Adapting the concept of raising rivals costs specifically to energy markets, Hunger (2003) raises the concern that a merger between a gas company and an electricity generation firm may incentivize it to withhold gas from the generation market to raise the wholesale electricity price received by its generators. Withholding is profitable if its impact on the firm's revenues in the electricity market, determined by the level of generation capacity and the elasticity of the generation supply curve, exceeds the opportunity cost of not selling the gas to other generators. Vazquez *et al.* (2006) expands on this opportunity to exert market power in the context of examining a real-world merger in Spain between a dominant natural gas firm and an electricity firm with a large quantity of gas-fired generation resources. In their model, a monopolistic gas producer restricts output beyond the level required to capture monopolistic rents in the power market in order to increase the wholesale electricity price and the revenues of their generators in that market. In this paper, we expand the theory developed by Vazquez *et al.* (2006) and Hunger (2003) by integrating a careful consideration of the role of transmission constraints and rights to capacity, adapted from the literature on market power in electricity markets, and empirically identify a real-world example of this scheme at play in New England.

3.3 Background: Three Interconnected Markets

3.3.1 The market for natural gas transportation

The modern US market for natural gas transportation was established through a series of reforms implemented by the Federal Energy Regulatory Commission (FERC) in the 1980s and 90s. These reforms effectively separated the gas transportation market from the physical commodity market by requiring interstate pipeline companies to sell their transportation services through long-term contracts for pipeline capacity.[4] Under this regime, capacity purchasers enter into multi-year contracts with a pipeline at FERC-regulated rates, which are designed to allow pipeline companies to earn a "just and reasonable" rate of return on their investment (FERC, 2017). Local gas distribution companies (gas utilities or LDCs which in turn provide gas to retail residential, commercial and industrial customers) have tended to be the largest subscribers of pipeline capacity, procuring sufficient contracts to meet retail customer demand.[5] As regulated utilities that are able to pass procurement costs through to their ratepayers, LDCs assume little commercial risk associated with entering into these long-term contracts. Other purchasers of long term capacity contracts include industrial facilities

[4] Previously, interstate pipeline companies would buy gas from producers at the wellhead, transport it to centers of demand, and sell it for a single price incorporating both their cost of the gas itself and their cost plus allowed profit from transportation. Following the reforms, the pipeline companies do not take ownership of the gas at any point. For a detailed history on the restructuring of the natural gas transportation market, see Oliver & Mason (2018).

[5] In addition, when the pipelines were converted from merchant to transportation-only entities, their firm sales contracts, which were almost exclusively with LDCs, were converted to firm transportation contracts. Outside of rare situations like those in Arizona, Florida, and Louisiana, in particular, there were almost no end-users or electricity generators with preexisiting firm sales contracts with interstate pipelines to be converted to firm transportation as part of industry restructuring.

and gas marketers (taking speculative positions). Gas-fired electric generators, which represent an increasing fraction of wholesale gas demand, have tended not to purchase long term contracts , because the cyclicality of their demand—both daily and seasonal depending on conditions in the wholesale electricity market—has made procuring long-term pipeline capacity contracts cost-prohibitive.

In 2008, FERC amended the rules governing long-term contracts to allow contract holders to sell temporary use of their pipeline capacity at unregulated prices on a secondary "capacity release market."[6] Additionally, pipeline operators may use the capacity release market to allocate unreserved capacity or sell unreserved capacity on an interruptible basis at an associated volumetric rate up to a maximum FERC-set rate. These policies are designed to promote a more liquid market for gas transport and to allow pipeline operators to efficiently allocate scarce capacity. Secondary capacity release sales can last anywhere from several hours to a year. Most on Algonquin fall in the range of a few days to a few weeks.

Firms that hold capacity rights on a pipeline are known as "shippers."[7] They exercise their capacity rights by electronically submitting "nominations" to the pipeline company on a daily basis. Nominations consist of an intake "receipt" point, an outflow "delivery" point, and a scheduled daily quantity of gas to flow.[8] This quantity must be flowed at a roughly even rate over the course of the 24-hour gas day, which runs from 9 a.m. til 9 a.m. Central Time the following day.[9] To induce shippers to judiciously manage nominations and flows, differences between scheduled nominations and actual

[6] While the capacity release market had been in existence since 1997, it was not until 2008 that capacity with a duration of a year or less could be sold into the market at an uncapped price.

[7] In many cases, the "shipper" that manages gas and capacity purchases is a separate subsidiary arm of their parent company. For simplicity, we refer to an LDC itself, its parent energy firm, and its shipping arm interchangeably throughout the paper.

[8] Net a small percentage that is skimmed by the pipeline operator to power compression stations.

[9] The precise rule for most contracts is that the gas must be flowed over a period lasting from 16 to 24 hours. For example, a capacity owner on Algonquin holding 24,000 MMBtu would generally be entitled to flow between 1,000 MMBtus and 1,500 MMBtus per hour as specifically set forth in their service agreement. Note that currently, the only service agreements providing for an associated 6% hour (like the 1,500 MMBtus per hour in the example) were those that were converted from sales agreements during restructuring.

flows incur imbalance penalties upon the shipper, which can be more or less severe depending on the size and nature of the infraction. Imbalance penalties are generally more severe when there is less slack available in the system to compensate, as in the winter.[10]

Shippers are able to make adjustments to their nominations during the gas day. FERC requires pipelines to offer a minimum of three "intraday" scheduling cycles, though some pipelines (including Algonquin) offer more frequent scheduling opportunities. On Algonquin and a few other lines, (like Transcontinental Gas Pipeline serving the east coast and New York), the last intraday cycle generally occurs a few hours before the end of the gas day and is commonly know as the "clean up" cycle. During this cycle, shippers match their scheduled nominations to their actual flows.

In addition to firm contracts, some pipelines (including Algonquin) offer "no-notice" contracts, which are a form of legacy contract generally only available to LDCs. On Algonquin, a no-notice contract is tied to storage capacity and service on Texas Eastern Transmission Company (which is owned and operated by same parent company as Algonquin). Together the storage service and no-notice contract allow an LDC to adjust its scheduled flows without prior notice, and to flow their total scheduled quantity of gas on an uneven hourly basis and over a period of less than a full 24 hours. FERC allows these contracts on the basis that they are necessary in order for LDCs to reliably serve their retail customers.

3.3.2 The wholesale natural gas market

In New England, gas is typically traded without the benefit of an exchange where bids and asks can be matched and settled for market-wide price determination. Instead, buyers and sellers must search for willing and able counter-parties and negotiate prices. This is partly because sales on New England's wholesale spot market frequently involve

[10] See Appendix A.3.1 for further detail on imbalance penalties.

delivery to a specific pipeline node, effectively re-bundling the physical commodity with the transportation service.[11]

Gas-fired generators purchase the vast majority of their gas on the wholesale spot market, on a "delivered to their location" basis, because their energy needs are typically much more variable and less predictable than those of LDCs and therefore not well served by long-term contracts. Most LDCs are both consumers and marketers of gas. Because they must hold sufficient long-term contracts to reliably supply their gas heating rate-paying customers, they find themselves, on all but the coldest winter days, with excess capacity rights. That excess capacity can be used either to ship gas to the region and sell it on the spot market or can be sold directly on the capacity release market.[12] Independent marketers do not themselves use gas, but instead hold long-term contracts in anticipation of profiting from short-run sales to the other firm types (primarily generators).[13]

Gas utilities (LDCs) are typically regulated monopolies. By regulatory design, they are allowed to make a fixed rate of return on their shareholder's capital investments. Ratepayers finance the LDCs' purchases of capacity rights which protect against gas price shocks. Hence, LDCs are subject to rules that limit their ability to profit from their excess contracts. These "revenue-sharing" rules are set by public utility

[11] Spot market prices therefore incorporate the wellhead price of the gas, the cost of the pipeline capacity needed to transport it, and the shadow price of the pipeline capacity constraint, which captures the difference in prices between the receipt and delivery regions due to differences in available supply when the pipeline is at full capacity (Cremer *et al.*, 2003). Because the prices of primary contracts for capacity are regulated but capacity-release and spot-market prices are not, the owners of capacity are able to extract congestion rents when capacity is scarce (Oliver *et al.*, 2014).

[12] Indeed, FERC regulations intended to ensure pipelines are fully utilized require them to sell any capacity (that is not scheduled by firm shippers) to interruptible contract shippers requesting access to that capacity by means of nominations to use it. In Section 3.5, we show how LDCs circumvent the intent of this rule in New England by initially nominating more capacity than they intend to use and then adjusting their scheduled quantity downward at the end of the gas day.

[13] Another set of participants in the spot market for gas are asset managers who act as third party agents and/or principles (depending on the Asset Management Agreement's terms) for for contract holders. These independent marketers hold long-term contracts either as principal directly with the pipeline; or as replacement shipper under a long-term capacity release transaction.

commissions and vary across states.[14] In general, they require LDCs to return a certain percentage of revenues from capacity release and spot market sales (sometimes referred to as "non-firm margin" sales) to their ratepayers.

For the regulator, choosing an appropriate revenue-sharing rule is a balancing act between protecting ratepayers and allowing the LDC to keep enough profit such that they are incentivized to transact their excess capacity efficiently. It is generally held that LDCs require little incentive to efficiently market their excess capacity. If the incremental cost of marketing an additional unit of excess capacity is small, then the LDC's share of the profit from marketing that additional unit of excess capacity can be similarly small without significantly distorting the LDCs behavior. However, this reasoning fails to consider the incentives of firms that earn profit in other interconnected markets. In particular, for firms that operate in natural gas supply and delivery markets as well as electricity markets, shrinking the incentive to efficiently market excess capacity could have the unintended consequence of increasing the relative weight those firms assign to profits earned in the market where their incentives are not diminished. Indeed, it is likely that revenue-sharing rules in New England have contributed to the extent, and location, of capacity withholding. In Connecticut, for example, the state where most of the capacity-withholding behavior we observe takes place, LDCs must return 99% of non-firm margin sales to ratepayers, while in Massachusetts the revenue-sharing rule distributes 90% to ratepayers, and in Rhode Island the rule is slightly more complicated but works out to about 83%.[15]

[14] Revenue-sharing rules sometimes also vary across firms within states, though this is not the case for the three New England states that are the focus of this analysis.

[15] Specifically, Rhode Island's Public Utilities Commission requires the single LDC that operates in the state to return 100% of the first $1 million of all non-firm margin sales to ratepayers and 80% of all additional revenues.

3.3.3 The wholesale electricity market

In 1999, New England became one of the first regions to implement a competitive wholesale electricity market. Under the current structure, energy (i.e., the electricity commodity) is traded in the *day-ahead* market, which is operated by the Independent System Operator of New England (ISO-NE).[16] The market takes the form of a first-price auction, in which generators bid in quantities of energy they will supply at given prices for each hour of the day. In an idealized setting, competition incentivizes generators to bid in their true marginal costs of production, allowing ISO-NE to utilize the lowest-cost combination of generation resources to use to meet demand.[17] As illustrated in Figure 1, the most expensive unit of generation required to meet demand (which is almost perfectly inelastic in the short run) sets the wholesale electricity price paid to all generators called upon operate.[18]

While virtually all generators competitively bid into ISO-NE's day-ahead market,[19] the way in which they earn profits depends on whether they are considered to be *regulated* or *merchant unregulated* assets. Merchant unregulated generators are typically

[16] ISO-NE additionally operates a *real-time* electricity market that balances short-term fluctuations in supply and demand, a forward capacity market, markets for transmission rights, and markets for system services such as regulation and reserves. Although many of these markets are affected by gas pipeline underutilization, we focus our analysis on the day-ahead energy market because it is the most significant in terms of trading volumes and generally sets price expectations for other markets.

[17] In practice, the first-price auction structure and local market power due to transmission constraints give generators some ability to mark up their bids above marginal cost (Kim, 2016). We do not address markup in this paper as it is not the focus of our analysis.

[18] This market-clearing price corresponds to the cost of the energy itself, which ISO-NE terms the "energy price." Line losses, transmission constraints, and spatially heterogeneous demand imply different values of energy in different areas of the grid, which ISO-NE accounts for by adjusting the energy price up or down at various nodes in the network to construct "locational marginal prices" (LMPs). LMPs are typically within a few cents to a few dollars of the energy price. Reserve requirements, heterogeneous ramping rates for various generation technologies, and other physical properties of the system further complicate the true dispatch procedure employed by ISO-NE. Throughout this paper, we model the day-ahead market in a simplified setting without these considerations, which successfully captures the relevant dynamics between the wholesale gas and electricity markets without getting weighed down in detail.

[19] A few hundred small-scale solar, wind, run-of-river hydro, and landfill gas facilities representing less than 1% of New England's total capacity do not participate in the day-ahead market and instead sell energy directly to utilities or commercial customers through bilateral contracts.

Figure 1: An illustration of New England's day-ahead electricity market, which is constructed and cleared for each hour of each operating day. Generators bid in their marginal costs of generation (plus markup when they have market power) which can be used to construct a market "bid supply" curve ranking generation resources from lowest to highest cost. Electricity distribution utilities bid in their predicted levels of demand to construct a market demand curve (shown here as perfectly inelastic as there is very little operational demand response in New England). The most expensive generation resource required to meet demand for a given hour sets the wholesale price received by all generators called upon to operate.

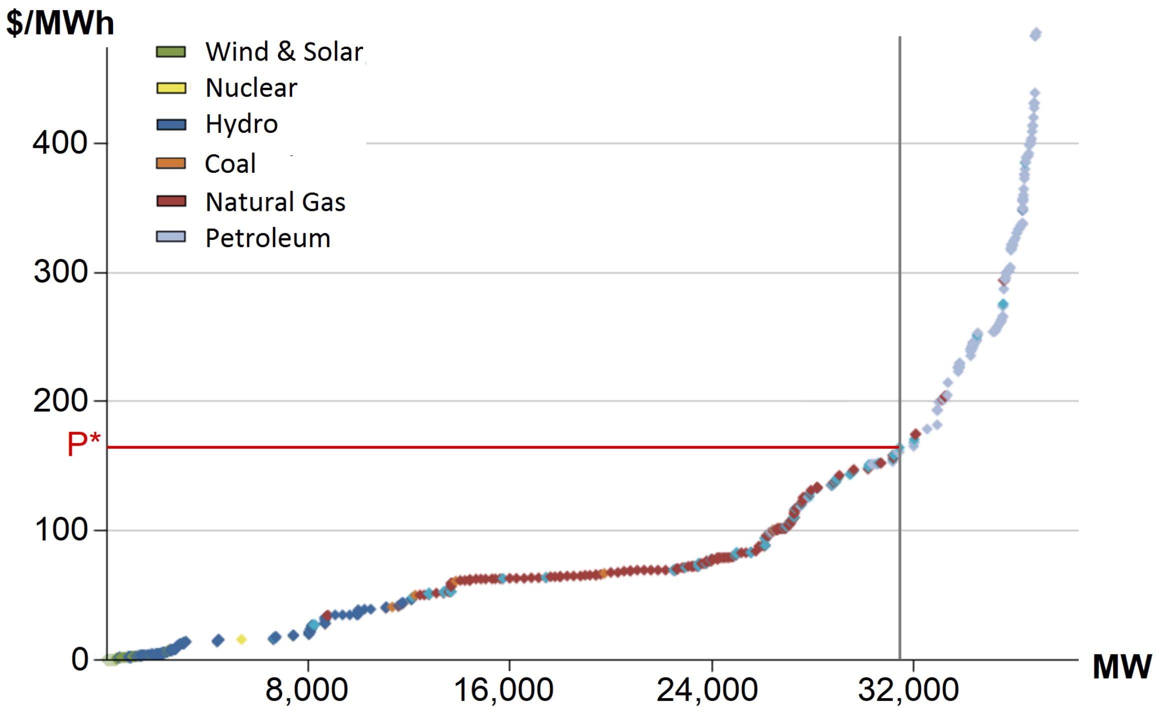

Note: The underlying data corresponds to marginal cost of generation and capacity rather than bid supply offers. We are not able to use the latter here in a straightforward manner because bid data is anonymized; however, the curve here roughly matches the distribution of generators' price and quantity supply offers to the day-ahead market by fuel type and serves the purpose of illustrating the bid supply curve.

owned by independent power generation firms that do not operate any transmission lines or distribution services.[20] These generators, which represent about 85% of New England's total capacity, pay all of their own fuel and capital costs and retain all of their own revenues.[21]

[20] One notable exception is 76.5 MW of unregulated wind generation owned by one of the two firms observed to withhold gas pipeline capacity on Algonquin.

[21] This accounting is admittedly more complex for generators that enter into bilateral contracts with utilities and commercial customers; however, because the prices in these contracts are set by expectations of wholesale market prices, we are able to focus on day-ahead market outcomes

The regulated generation assets that make up the other 15% are typically owned by electricity distribution utilities. In some cases, regulated plants exist because they provide reliability services that are not economical under the current market structure but which are necessary to provide reliable service (i.e. peaking plants in urban load pockets). In other cases, they are holdovers from the regulated environment two decades prior. Instead of profiting based on the outcomes of the wholesale electricity market, the electric utilities that own these generation resources are entitled to make a fixed rate-of-return on the capital investment that goes into them, much in the same way they profit from their distribution assets.

New England is heavily reliant on gas-fired electric generation, with gas supplying about half of all electricity generated in the region. As shown in Figure 1, natural gas occupies the middle portion of the bid supply curve, and consequently a gas-fired plant is the marginal generator about three-fourths of the time. Accordingly, higher wholesale gas prices usually imply higher electricity prices, and this effect is amplified at higher prices due to the convexity of the bid-supply curve.

As will be discussed in detail in Section 5, the gas LDCs engaged in capacity withholding are owned by parent energy firms that also own hundreds of megawatts of regulated generation capacity in New England, and one of the two additionally owns about 75 megawatts of unregulated capacity. These parent firms are primarily gas and electricity distribution utilities, but each have meaningful margins in the wholesale electricity market. As the markets for gas and electricity become increasingly interdependent, new opportunities emerge for firms with operating arms in both markets to exert market power. In the next section, we formalize the relationships between the three markets described here in order to demonstrate how a vertically-integrated firm may benefit by using its contracts for capacity to reduce availability on the pipeline rather than to actually transport gas.

without loss of generality.

3.4 Theoretical framework

In this section we present a graphical discussion of the incentives to withhold pipeline capacity.[22] The central point developed in this section is that firms that operate in both the pipeline transportation market and the electricity market have a clear incentive to restrict gas deliveries during periods of scarcity (raising the wholesale gas price), in order to capture rents (and raise rivals costs) in the electricity market. The incentive for these firms to withhold gas, rather than sell it in the wholesale market, is amplified for gas LDCs by revenue-sharing rules that require gas utilities to dividend most of the profit from non-firm-margin sales back to their ratepayers. We consider how state-level variation in these rules impacts which nodes firms use to withhold supply. Finally, we consider how the spatial nature of the pipeline network contributes to where in the system firms are likely to withhold capacity, noting that firms will have a stronger incentive to withhold capacity at points where it is likely to have the greatest impact on the wholesale electricity price.

Based on the institutional features discussed in Section 3, suppose firms may operate in three vertically related markets: Furthest upstream is the market for natural gas pipeline transportation, which is used to deliver gas to a wholesale gas market, in which LDCs sell gas to electric generators. Second is the wholesale gas market (in which LDCs sell gas after serving retail gas demand). Third is the wholesale electricity market in which some generating units are gas-fired and others are not.

Our firm of interest operates as a seller in both the wholesale electricity market as well as in the wholesale and retail gas markets. It owns electric generating units (for simplicity we can assume non-gas fired). Through its LDC operating arm, the firm also holds pipeline capacity – which, as with the other LDCs – is a source of market power in the wholesale gas market during periods of scarcity. Similar to the other LDCs, this firm is required to first serve demand in the retail gas market, and

[22] An algebraic presentation is relegated to the Appendix.

may after that sell any excess pipeline capacity in the wholesale gas market under the same regulated rates of return and profit sharing rules, respectively. Unlike the other LDCs on the system, this firm's incentives derive from the interaction of its positions in the electricity and gas markets.

We start by considering the impact of a reduction in total deliveries in the gas transportation market upon the wholesale gas market. We interpret this reduction in flows as arising because one firm chooses to overschedule, *i.e.* it schedules for larger deliveries in the day ahead timely cycle than are ultimately executed the next day. From the firm's perspective, the important consideration for assessing the impact on the wholesale gas price is the degree of total excess pipeline capacity prior to overscheduling.

Figure 2 illustrates the interaction between the pipeline transport and wholesale gas markets. We model the supply curve of gas transport as Leontief in nature: marginal costs are constant (reflecting a constant per-unit commodity price and a constant cost of transporting gas along the pipeline) so long as scheduled deliveries fall below total pipeline capacity. The wholesale price of natural gas will then equal marginal cost. However, when scheduled deliveries rise to the level of maximum available capacity there is no possibility to increase transportation capacity at any cost in the short-run, as illustrated by the vertical turn in the supply curve. In this situation, the wholesale price of gas will be determined by the level of demand (rather than the marginal cost), i.e., by how much potential buyers are willing to pay for those last units of capacity, and thus make it possible for sellers of pipeline capacity to extract scarcity rents.

When the quantity demanded at the (constant) marginal cost is less than maximum available capacity so that there is a lot of excess pipeline capacity, as in the left panel of Figure 2, overscheduling will have little impact upon the wholesale price of natural gas. But when the demand curve shifts so that the quantity demanded at marginal cost is close to the maximum available capacity and the pipeline comes closer to its

Figure 2: Impact of overscheduling capacity on the wholesale gas market on warmer days when the pipeline is uncongested (left panel) and on colder days when it is fully scheduled (right panel)

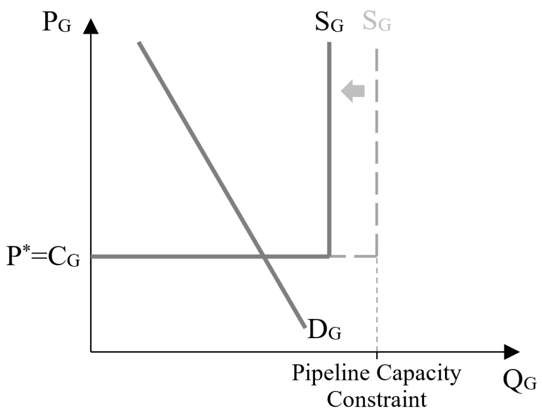

Figure 3: Impact of overscheduling capacity on the wholesale electricity market on warmer days when the pipeline is uncongested (left panel) and on colder days when it is fully scheduled (right panel)

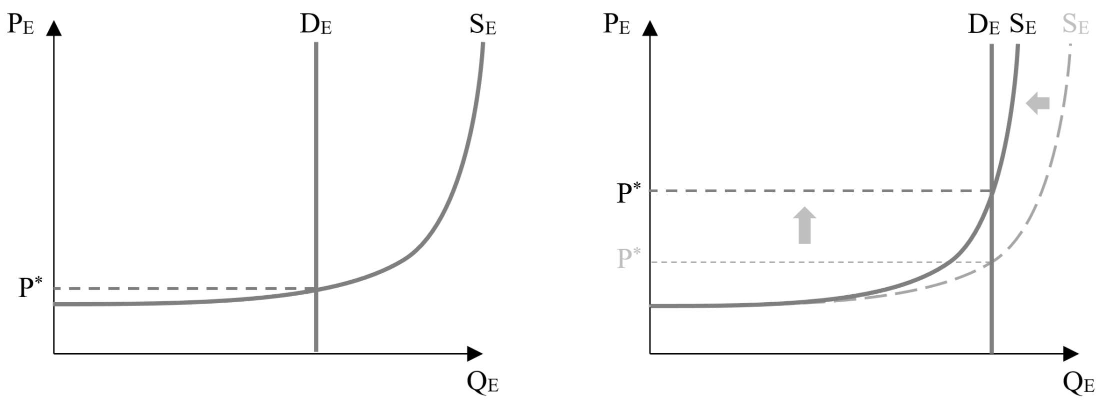

capacity constraint, as in the right panel of Figure 2, overscheduling can induce an increase in the wholesale price of natural gas; the magnitude of this increase in price depends on the level of demand compared to available capacity (i.e., excess capacity) and the magnitude of overscheduling. The smaller is the initial level of excess capacity (*e.g.,* as a result of large retail demand for gas – perhaps as a result of colder than expected temperatures) the easier it is to force wholesale gas prices up.

Any increases in natural gas wholesale prices will have a derivative effect upon

Figure 4: Impact of overscheduling capacity on the wholesale gas market (left panel) and on the electricity market (right panel). The dark shaded region in the left panel represents gas market revenues the withholding LDC sacrifices by letting their capacity go unused, which are restricted by revenue-sharing rules. The dark shaded area in the right panel corresponds to the additional revenues earned by the LDC's generation capacity from a higher wholesale electricity price.

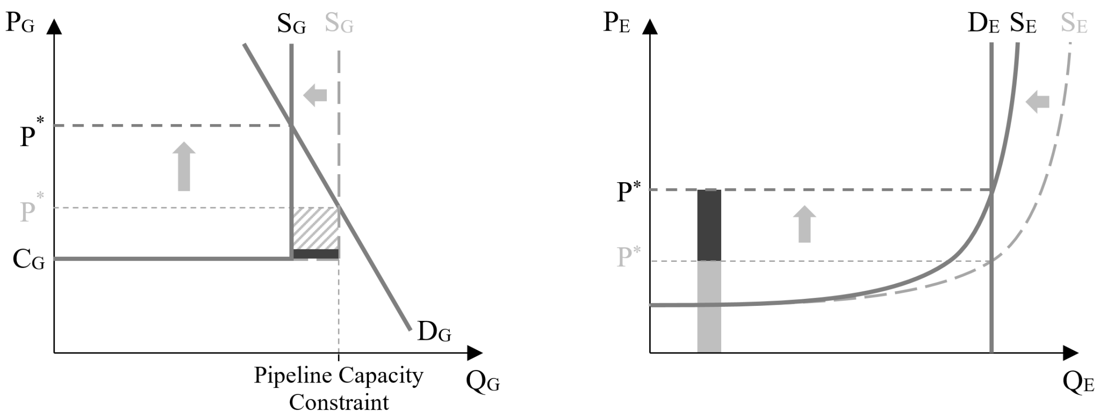

electricity markets, as illustrated in Figure 3. Two factors are in play here: the initial interaction of supply and demand in the electricity market, and the degree to which increased natural gas wholesale prices shift the supply curve for electricity. We model the demand for electricity as perfectly inelastic, with wholesale electricity prices reflecting the marginal cost of supplying electricity at the market-clearing quantity (as depicted in the left panel). The industry aggregate marginal cost, in turn, reflects the incremental cost of the marginal producer. This marginal cost curve rises slowly over a significant range of quantities. As electricity sales rise, ever-less efficient (and therefore more expensive) sources of electricity supply are called on, and so the marginal cost curve becomes steeper and steeper.

Importantly, many of these more costly sources of electric generation are gas-fired. Increases in gas wholesale prices are likely to cause an upward shift in the electricity supply curve, particularly at higher quantities. If demand intersects supply in the elastic region, or if the overscheduling in the gas transport market does not induce much of an increase in wholesale natural gas prices, then the impact on electricity prices is inconsequential (left panel). But if demand intersects supply in the less elastic portion and wholesale gas prices rise significantly as a result of overscheduling in the gas transport market, then there will be a noticeable increase in electricity prices (right panel).

Of course, any overscheduling that leads to unused pipeline capacity means that there are foregone profits from not supplying the wholesale gas market. This effect is illustrated as the light rectangle in the left panel of Figure 4. These profits are only partly foregone by the overscheduling LDC, who is subject to a revenue sharing rule: most of these foregone earnings would have had to be returned to the LDC gas customers. As such, the ultimate amount of foregone profits for the firm's shareholders is substantially smaller, as illustrated by the darker (small) rectangle in the left panel. By contrast, any increases in profits in a separate market need not be subject to this regulatory effect. In particular, if the firm has holdings in the wholesale electricity market then any extra profits that arise therein as a result of the higher wholesale gas prices are retained; this effect is illustrated in the right panel of Figure 4.[23]

The model sketched above articulates the differentially large incentive a firm that holds positions in both the electricity and gas markets can have to influence the price of wholesale natural gas, especially in the presence of a revenue sharing rule. We also wish to explore the spatial nature of incentives when these integrated firms are located at different points along the pipeline. To this end, we imagine two delivery points (nodes), one upstream and one downstream. There are retail gas customers located at both nodes; we presume the market downstream is larger than the market upstream

[23] The scenario illustrated in this Figure implicitly assumes the firm is not subject to cost of service regulation in the electricity market. In the markets we analyze empirically in the next section, one firm is subject to rate-of-return regulation also on its electricity generation holdings. But then there is a separate, and subtle, motive for pushing up electricity prices. If this rate-of-return regulation requires a minimum level of operation of that generation capacity and the firm in question holds high-cost electricity generation capacity there is a risk that it will not be able to meet this stipulation. By forcing up electricity prices, the firm may be able to convert its high-cost units into economically viable units, thereby accessing a revenue stream that might otherwise be unavailable. This problem will be particularly acute if the firm has established costly capital, such as expensive scrubbers, only to find these units priced out of the market at most points in time.

(*e.g.,* on the Algonquin pipeline, which flows from South to North, the upstream node might be Hartford and the downstream node Boston). As above, the LDCs are obliged to meet all retail gas demand at that node. One important distinction to the model above is that a firm located at the downstream node has the right to sell gas in either the upstream or the downstream wholesale gas market, whereas a firm located at the upstream node can only sell gas in the upstream market.[24]

With minor adaptation, the raising rivals' costs arguments described above can be applied here. The key point is that the incentive to raise costs by influencing the price of delivered gas is larger for a firm located upstream than downstream, for two reasons: first, because there is necessarily less spare pipeline capacity upstream it takes less withholding to engender any particular level of increase in delivered price. Second, because of the additional pipeline tariff a firm must pay to utilize the segment between the upstream and downstream points, the upstream firm has a natural cost advantage over the downstream firm. As we noted above, any such cost advantage is the root source of motives to influence markets by raising input prices (and thereby raising rivals' costs).

The final point we wish to make relates to demand shocks, as might arise in inclement weather conditions. In weather conditions that raise demand, we expect to see a larger impact in the downstream market than in the (comparatively smaller) upstream market. If these shocks are anticipated when flows are scheduled, as seems likely, then the disproportionate increase in downstream demand will have spillover effects in the upstream market, because an increase in scheduled deliveries downstream also raises the amount of gas shipped to or through the upstream point they must

[24] One can think of these firms holding contracts for delivery, with one firm holding a contract guaranteeing delivery to the upstream node, and another firm holding a contract allowing delivery downstream or any other points further upstream. With this interpretation, the latter could choose to withdraw gas at either point, while the former firm would be obliged to remove gas only upstream. This case is most operative when the system approaches scheduling near to full contractual entitlements. At these times, restrictions to receipts and deliveries along primary paths cause the delivery rights of the contracted capacity to the upstream location to be the terminus of their firm rights.

reduce spare capacity upstream. In essence, the increased downstream demand creates conditions where it is easier to hold upstream markets hostage – and where it takes less intervention to force wholesale gas prices up through overscheduling.

3.5 Detecting Capacity Withholding

In this section, we show empirically that capacity withholding occurred on the Algonquin pipeline, and that it resulted in significant levels of pipeline capacity going unused in aggregate. Moreover, withholding occurred at only a small subset of LDC-designated delivery nodes operated by just two parent energy companies that also own significant generation capacity in the region. These nodes are primarily located in Connecticut, where revenue-sharing rules are least generous to LDCs, and they are disproportionately served by no-notice contracts that allow shippers to make large schedule adjustments without incurring imbalance penalties.

We detect capacity withholding by analyzing hourly-level scheduling data for all 117 delivery nodes on the Algonquin pipeline for every day in our three-year study period.[25] This reveals a unique pattern exhibited by a handful of nodes wherein shippers consistently reserve more capacity than they use to actually flow gas. These shippers avoid incurring imbalance penalties designed to prevent this type of overscheduling by reducing their scheduled quantity at the last moment, such that their final scheduled quantity matches what they actually flowed.

It is striking to observe these patterns at all, but we note that withholding increases in both magnitude and variance in the winter months, when greater demand for gas for both heating and generation increases the value of the fixed stock of pipeline capacity. As illustrated in the previous section, revenue-sharing regulations limit

[25] Hourly scheduled quantities for all nodes going back three years are publicly available through the Algonquin pipeline's FERC-mandated electronic bulletin board, generating about seven million node-hour observations. Note that we only observe scheduled quantities; actual flows are known only to the pipeline company and individual nodal operators.

LDCs' ability to extract these congestion rents through the gas transportation market or the wholesale gas market, thereby making the less-efficient extraction pathway—raising the wholesale electricity price—comparatively more attractive. While the available data are insufficient to conclusively determine whether this is indeed the motivation of the withholding firms, we demonstrate that the observed patterns of withholding are consistent with this explanation and inconsistent with several other possibilities.

3.5.1 Analysis of Scheduling Patterns

Each contract for capacity gives the shipper holding it the right to use a certain amount of space along the pipeline between one or more specifically listed *receipt* (intake) nodes and one or more specifically listed *delivery* (outflow) nodes. To actually exercise this right, the shipper must electronically submit a *nomination* to the pipeline company, which states the quantity of gas they intend to move, where it will enter the pipeline, and where it will exit. This capacity scheduling process is carried out on a daily basis for each gas day, which runs from 9 a.m. til 9 a.m. the following day.[26] Importantly, capacity is nominated not as a rate of flow, but rather as a total quantity to be transported over the course of the gas day at a roughly constant rate.

Shippers must submit their initial nominations by the close of the *timely cycle*, which occurs at 1 p.m. the day before the gas day, in order to be guaranteed the capacity provided by their contracts.[27, 28] In contrast to the majority of interstate

[26] The gas day and all associated scheduling times are in Central Clock Time for all interstate pipelines to facilitate harmonization of the gas transportation industry across the US.

[27] If shippers neglect to nominate by the timely cycle, the pipeline may use their contracted capacity to allow other shippers to move gas to other points on the pipeline that are not their specifically listed nodes. These nominations, which are called *secondary in-path*, *secondary out-of-path*, and *interruptible* nominations, are used by LDCs and independent marketers to sell gas on the wholesale market to generators located at other parts of the pipeline. On a day when the pipeline is fully scheduled, the shipper that did not nominate on time will only be able to utilize pipeline capacity if another shipper will adjust its nomination downward later in the scheduling period to free up some capacity.

[28] *No notice* contracts are exempt from this requirement: For the most part, provided a non-zero

pipelines, Algonquin allows shippers to adjust their nominations on an hourly basis over the 44-hour scheduling period, which begins with the timely cycle at 1 p.m. the day before and concludes at 9 a.m. at the end of the gas day.[29] While schedule changes are observed at all hours, the vast majority of adjustments are made within three specific windows: In the eight hours following the initial nomination at the timely cycle; between 6 a.m. and 6 p.m. encompassing the start the gas day; and between 6 a.m. and 8 a.m., just before the end of the gas day.

Figure 5 illustrates the scheduling pattern at a typical LDC-designated node, which is characterized by frequent, relatively large adjustments in either direction in either of the earlier two scheduling windows and less frequent, relatively small adjustments in the final window. Each delivery node's pattern is unique, but the vast majority of LDC-designated nodes can be broadly characterized by this description.

The substantial, bidirectional schedule changes made in the first two common time windows are consistent with LDCs getting better information about their expected retail demand, generators getting better information about electricity market conditions, and gas being traded on the wholesale market, either directly or through independent marketers, to efficiently allocate capacity between these two firms types.[30] On a cold day when pipeline capacity is fully scheduled, a downward adjustment at one node in either of these first two windows is therefore accompanied by an increase at another node, and the pipeline remains fully scheduled. In contrast, the schedule adjustments made in the window between 6 a.m. and 8 a.m. at the end of the gas day represent shippers matching the node's final scheduled daily nomination to the quantity of gas that was actually delivered to it, in what the industry refers to as the *clean up* or *true*

nomination is submitted in the Timely cycle, these contracts enable scheduled quantities to be adjusted at any time during the scheduling period with guaranteed approval.

[29] Most other pipelines allow schedule adjustments at just five specific times following the initial timely cycle nomination: A *late cycle* in the evening the day before the gas day, three *intraday cycles* during the gas day, and a final *cleanup cycle* near the end of the gas day.

[30] For simplicity, we disregard the two industrial end users directly connected to Algonquin here as they account for less than 1% of the market.

Figure 5: The scheduling pattern of a typical LDC delivery node.

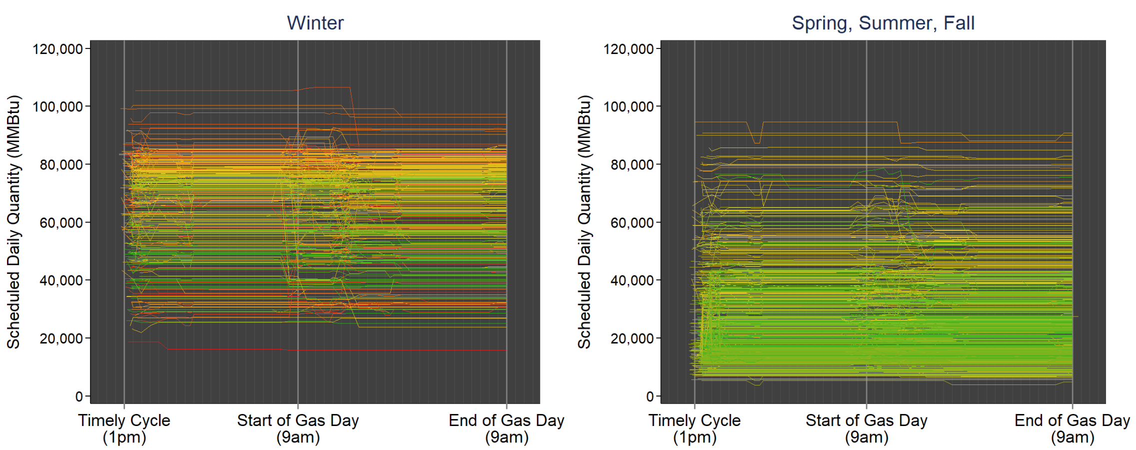

Each line represents one gas day in our three year study period. Winter is defined as December 1 through March 31 following the delineation used by Algonquin Gas Transmission. The X-axis covers the 44-hour scheduling period and the Y-axis is the total daily quantity of gas scheduled for delivery to that node at a given time. Line color represents the Algonquin Citygate price, wherein redder lines indicate higher-priced days. The scheduling pattern at this node and at most other LDC delivery nodes on Algonquin is characterized by most adjustments being made shortly after the timely cycle or around the start of the gas day with some slight balancing either direction in the final hours on some days. We constructed equivalent graphs for all 117 delivery nodes on the Algonquin pipeline.

up cycle. Beyond accurate bookkeeping, this adjustment is necessary for shippers to avoid the accounting imbalance penalties assessed for monthly deviations between scheduled and actual flows in excess of 5%.[31] The relatively small, less frequent, bi-directional end-of-day adjustments observed in Figure 5 are consistent with an LDC shipper that nominates capacity with the intent to use their entire nomination to transport gas to customers. In this case, the adjustments reflect only differences between their prediction of retail customer demand and realized demand, which will be minimal given previous opportunities for adjustment if predictions are accurate in expectation.

We focus our attention on ten LDC-designated delivery nodes that do not exhibit the same pattern, and instead are observed to make large, consistently negative schedule adjustments in the final three hours of the gas day. Figure 6 illustrates an example of this distinct scheduling pattern, which is exhibited very prominently at six

[31] See Appendix A.3.1 for further detail on imbalance penalties.

Figure 6: The scheduling pattern of an LDC delivery node that consistently down-schedules its nomination in the final three hours of the gas day.

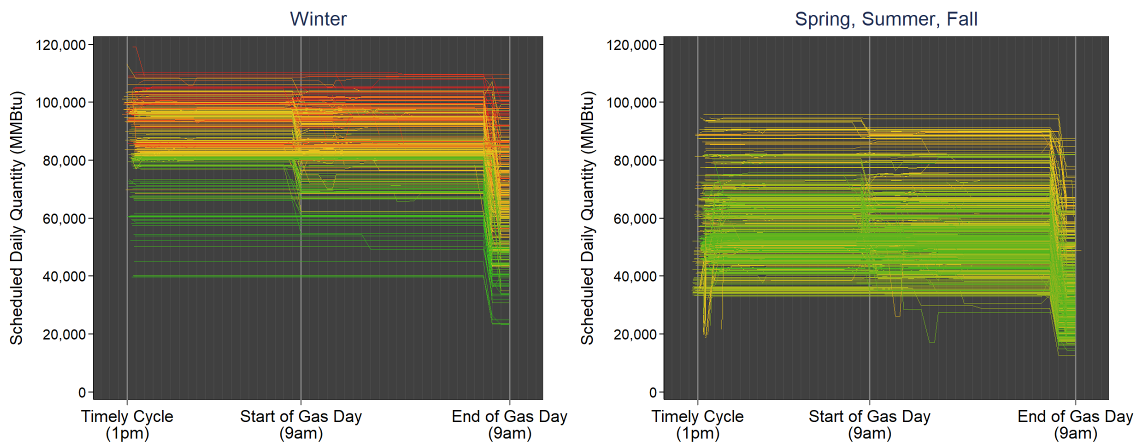

The y-axis corresponds to a total daily quantity rather than a flow rate, meaning these large negative schedule adjustments correspond to unused pipeline capacity. For example, if this node schedules 72,000 MMBtu at the beginning of the scheduling period, it is indicating to the pipeline company that it will be flowing gas at a rate of 72,000/24=3,000 MMBtu per hour over the course of the gas day and the pipeline company then reserves that capacity for them. When it reduces its scheduled quantity to 48,000 MMBtu three hours before the end of the gas day, it is not reducing its rate of flow at that time, but rather indicating to the pipeline company that it had been flowing gas at a rate of 2,000 MMbtu per hour over the gas day.

nodes that are clear outliers, and to a lesser extent at four additional nodes.

These large, consistently negative adjustments just before the end of the gas day are consistent with an LDC shipper that intentionally nominates capacity in excess of its predictions of its customers' total daily demand. In this case, the adjustments incorporate the differences between predicted and realized demand *plus* the offloading of overscheduled capacity to avoid imbalance penalties. Capacity cannot be double booked, and the pipeline company manages nominations such that the scheduled flows through any point on the pipeline do not exceed safe operating limits. Thus, a negative adjustment in the clean up cycle indicates capacity that was scheduled but not utilized to ship gas to that node. When aggregate nominations reach the pipeline's capacity constraint and the negative schedule adjustment is not accompanied by a positive adjustment at another node, the negative adjustment in the cleanup cycle corresponds to capacity that went unused across the entire system for that gas day.

Table 1: Average schedule change in the last three hours of the gas day by node in MMBtu. Six nodes are clear outliers from the rest of the distribution, and the next four nodes that downschedule the most are also operated by either Firm A or Firm B.

Rank	**Schedule Change in Last 3 Hours**	**Schedule Change (Winter Only)**	**Node Operator**	**Node Type**
1	-18,444	-17,865	Firm A	LDC
2	-10,576	-9,281	Firm A	LDC
3	-7,116	-7,766	Firm A	LDC
4	-3,889	-2,529	Firm A	LDC
5	-3,808	-8,426	Firm B	LDC
6	-2,401	-4,963	Firm B	LDC
7	-861	-286	Firm A	LDC
8	-711	-1,645	Firm B	LDC
9	-563	-59	Firm A	LDC
10	-479	-975	Firm B	LDC
11	-348	18	Firm L	Generator
12	-250	-594	Firm M	Generator
13	-229	-261	Firm C	LDC
...	...	...	...	...
116	395	449	Firm N	Generator
117	653	1,017	Firm K	LDC

The six nodes we clearly observe making consistent, substantial negative adjustments in the final hours of the gas day (hereafter referred to as *downscheduling*) are all operated by shippers owned by just 2 out of the 27 parent energy firms that operate delivery locations, which we will refer to as Firm A and Firm B. To establish that the scheduling patterns at these nodes systematically differ from the rest of the distribution, we plot each node's average schedule changes over each hour of the 44-hour scheduling period in Figure 7. We note also that following these six nodes, the four nodes with the next largest negative schedule adjustments in the final hours of the gas day are also operated by Firms A and B, suggesting they may be involved as well (see Table 1). Across the entire system, aggregate schedule adjustments in the final three hours of the gas day averaged -48,493 MMBtu over our three-year study period and -51,152 MMBtu in the winters. On 37 days in the study period, the aggregate adjustment exceeded -100,000 MMBtu, which is roughly 7% of the pipeline's total

Figure 7: Schedule change from the previous hour over the scheduling period. Each line represents the average behavior of one the 117 delivery nodes on Algonquin. Six nodes operated by either Firm A or Firm B are clear outliers from the rest of the distribution in consistently making large negative schedule adjustments in the final hours of the gas day. Nodes operated by Firm A engage in this practice year round, while nodes operated by Firm B primarily perform these schedule adjustments primarily in the winter.

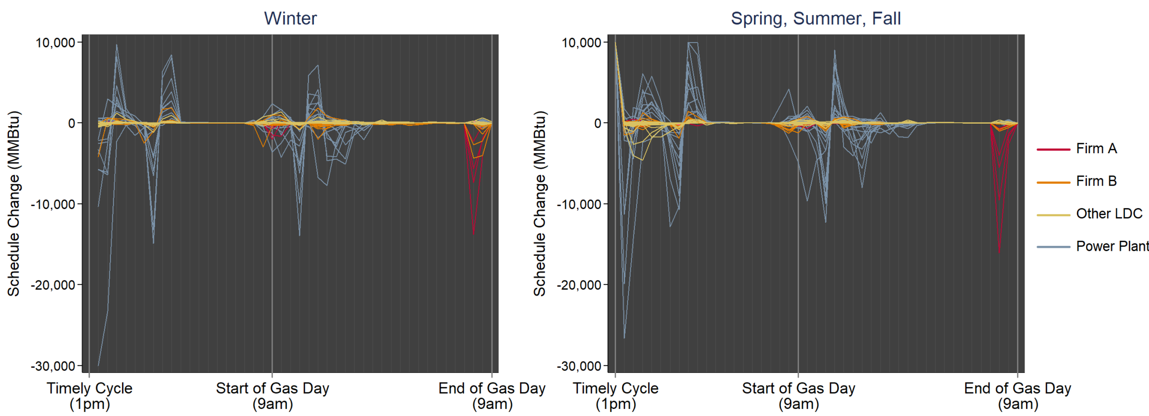

capacity[32] and roughly 28% of the total supply to electricity generators connected to the Algonquin pipeline.

Examining scheduling patterns separately for the winter season and the rest of the year reveals a behavioral difference between these two firms: Firm A consistently downschedules in both the summer and winter seasons, while Firm B engages in downscheduling primarily during the winter season.[33] In general, the suspect Firm A nodes appear to be exhibiting a blanket policy of always withholding excess capacity to the extent their contract holdings enable them to do so, whereas the suspect Firm B nodes appear to "turn on" this policy at some point toward the beginning of the winter season and then turn it off again sometime in the spring. This difference could be interpreted as suggestive evidence that the incentives for Firm A to increase electricity prices are stronger than those for Firm B.[34] However, we refrain from placing

[32] Measured at the Stony Point compression station, which is the most frequent bottleneck for deliveries to New England.

[33] We define winter as between December 1 and March 31 following the delineation used by Algonquin Gas Transmission (Spectra, 2016).

[34] We explore this possibility further in Section 5.4.

much weight on this interpretation, because the decision to withhold or release excess capacity is relatively inconsequential in the summer. Blocking capacity will not affect gas and electricity prices when the pipeline is not fully scheduled, and capacity rights cannot be used to extract scarcity rents through the capacity release or wholesale gas spot market on warm days when the pipeline is uncongested. In contrast, as explained in the theory section, this behavior will raise gas and electricity prices when the pipeline is constrained, and when scarcity rents are available, revenue-sharing mechanisms will diminish capacity rights holders' ability to extract them through capacity release or wholesale gas market sales. We proceed to explore how these mechanisms relate to observed downscheduling in the following sections.

3.5.2 Geography of Suspect Nodes

Spatially, we observe that eight of the ten most frequently downscheduling nodes, including five of the six clear outliers, are located in close proximity to one another in Connecticut (see Figure 8). Importantly, this section of the pipeline is downstream of its primary bottleneck at the Stony Point compression station, meaning nominations to these locations can exclude others from delivering gas to New England when the pipeline is congested.[35] However, all nodes downstream of Stony Point, including those in Rhode Island and Massachusetts, have the same capability to influence the pipeline's effective capacity constraint. While there are some potential alternative explanations,[36] we believe downscheduling behavior occurs primarily at nodes in

[35] In order to keep gas flowing at a high rate across long distances, interstate pipelines have compression stations every 50 to 100 miles that effectively break the pipeline into a series of segments. On the Algonquin pipeline, Stony Point is the compression station that is most frequently scheduled up to its operating capacity first, and it is located downstream of all of Algonquin's major Western receipt points.

[36] One potential alternative explanation is that there is significantly less generation capacity located in Connecticut, meaning there is less demand for natural gas in that state. This may combine with frictions for using contracts to deliver gas further downstream than its listed delivery nodes to make excess contracts delivering gas to Connecticut less valuable for capacity release or spot market sales and comparatively more valuable for capacity overscheduling (see Appendix A.3.2 for a more detailed explanation).

Figure 8: The locations of the 10 nodes that downschedule the most on average. Eight are located in Connecticut, where revenue sharing rules are strongest.

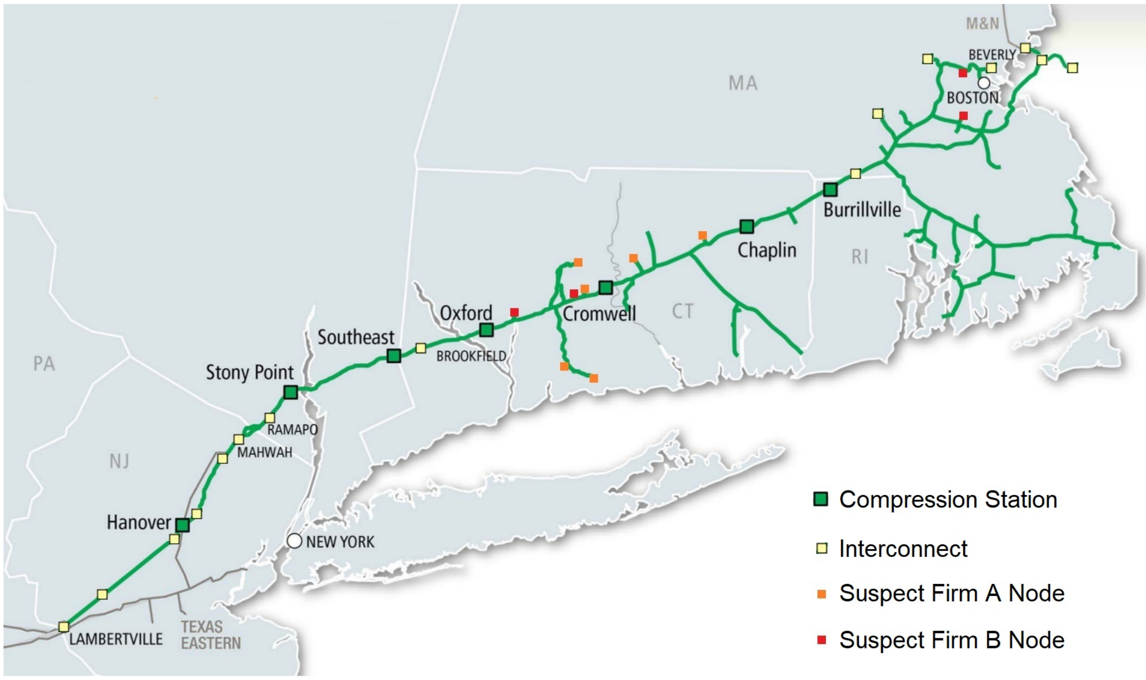

Connecticut because of the strength of the state's extra-marginal revenue sharing rules for LDCs.[37]

While it is relatively straightforward to see in the data that downscheduling primarily occurs at LDC-designated nodes in Connecticut, we formally test this hypothesis using the following regression model:

$$D_{it} = \alpha_0 + \beta_1 LDC_{it} + \beta_2 CT_{it} + \beta_3 (LDC_{it} * CT_{it})$$
$$+ \alpha_1 HDD_{it} + \alpha_2 HDD_{it}^2 + \alpha_3 W_t + \lambda_t + \varepsilon_{it}$$

Downscheduling D_{it} is defined as node i's scheduled daily quantity at 6 a.m. on gas day t (three hours before the end of gas day) minus its scheduled daily quantity at 9 a.m. at end of that gas day. LDC_{it} and CT_{it} are binary indicators for whether node i

[37] Recall from Section 3.3.2 that LDCs in Connecticut are required to return 99% of non-firm margin (i.e. capacity release and gas spot market) sales to ratepayers, versus 90% in Massachusetts and a slightly more complicated rule that works out to about 83% in Rhode Island.

is an LDC or located in Connecticut, respectively.[38] While neither of these variables change over the study period, we use node-days as the unit of analysis to allow for the inclusion of time-varying controls, and we adjust standard errors for clustering at the node level.[39] These controls include temperature in the form of heating degree days (HDD_{it} and HDD_{it}^2),[40] an indicator for whether the day is a weekend (W_t), and quarter fixed effects (λ_t) to capture seasonal variation common to all nodes.

Results are presented in Columns (1) and (2) of Table 2. We find that downscheduling is concentrated at LDC-designated nodes with significance at the 5% level. Column (2) demonstrates that downscheduling occurs primarily at such nodes located in Connecticut, as the coefficient on the interaction is an order of magnitude greater than the coefficient on either indicator variable. This is consistent with the hypothesis that revenue sharing rules being less generous to LDCs in Connecticut distort their incentives to use their excess pipeline capacity efficiently in the upstream market.[41]

3.5.3 Ability to Avoid Imbalance Penalties

As described in Section 3.3.1, shippers must pay accounting imbalance penalties to the pipeline company if their scheduled flows deviate from their actual flows in excess

[38] We elect to use a binary indicator for Connecticut for ease of interpretation, especially given our frequent use of this variable for interactions in subsequent regressions. Our results are robust to using a continuous variable for the revenue-sharing mechanism (see Table A1 in the Appendix).

[39] The variable we use for contracts in subsequent regressions does vary across nodes over time; however, given our sample size, the inclusion of node fixed effects diminishes our statistical power to the extent that we are unable to draw any meaningful conclusions. In essence, some nodes engage in downscheduling behavior and some do not, and the objective of this section is to empirically explore the factors that contribute to this behavior to the extent that the available data enable us to do so. We acknowledge that this is primarily a correlational analysis, and that while our findings provide supporting evidence for our hypotheses about the institutional conditions and firm incentives that drive capacity withholding, we do not causally identify these mechanisms.

[40] *Heating degree days* refers to the number of degrees below 65 Fahrenheit and is commonly used as a better proxy for heating demand than temperature. We average this variable across Connecticut, Rhode Island, and Massachusetts, using data from the National Climatic Data Center.

[41] We acknowledge, of course, that any other characteristic of Connecticut could be driving this result. In particular, we believe one other significant contributing factor may be lower demand for natural gas for generation in Connecticut (discussed in detail in Appendix A.3.2). We have focused on differential revenue sharing rules for the moment because it is more straightforward, but the results of this subsection also pick up the effect of alternative mechanisms such as reduced generation capacity.

Table 2: Relationships between downscheduling in the final 3 hours of the gas day and LDC status, whether the node is located in Connecticut (as a proxy for strength of LDC extra-marginal revenue sharing mechanisms), "no notice" contracts delivering gas from the Texas Eastern pipeline, and interactions. Coefficients on Contracts and CT×Contracts are omitted in columns (3) and (5) because only LDCs hold "no notice" contracts on Algonquin.

	(1)	(2)	(3)	(4)	(5)
LDC	585.0**	114.2**	315.1		99.53**
	(240.6)	(48.20)	(207.9)		(45.61)
CT		61.95		236.9	61.95
		(89.38)		(375.3)	(89.38)
Contracts			0	0.00319	0
			(.)	(0.00344)	(.)
LDC×CT		1325.8**			212.1
		(647.5)			(473.8)
LDC×Contracts			0.0502		0.00214
			(0.0521)		(0.00344)
CT×Contracts				0.419***	0
				(0.119)	(.)
LDC×CT×Contracts					0.417***
					(0.122)
HDD	7.630	7.630	7.630	7.630	7.630
	(4.660)	(4.660)	(4.660)	(4.660)	(4.660)
HDD^2	-0.266**	-0.266**	-0.266**	-0.266**	-0.266**
	(0.129)	(0.129)	(0.129)	(0.129)	(0.129)
Weekend	83.48***	83.48***	83.48***	83.48***	83.48***
	(28.87)	(28.87)	(28.87)	(28.87)	(28.87)
Quarter FE	Yes	Yes	Yes	Yes	Yes
N	133,029	133,029	133,029	133,029	133,029

Standard errors in parentheses (clustered at the node level)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

of 5% on a monthly basis.[42] Additionally, shippers must pay *Operational Flow Order* (OFO) imbalance penalties if they cause a physical imbalance in the system on days

[42] This threshold gives shippers some room to block capacity on occasion without later making adjustments in the clean up cycle, but frequently withholding large shares of excess capacity would require the downscheduling behavior we observe.

where the pipeline issues an OFO warning.[43] Overscheduling could potentially be executed under regular firm-service contracts. But to avoid imbalance penalties, the shipper would need to source gas from a supplier that is complicit in injecting an amount of gas into the pipeline that differs from what has been scheduled. Any storage provider or producer could potentially play this role,[44] but the on-demand storage service offered by the Texas Eastern pipeline provides a convenient tool for shippers using Algonquin to accomplish this without involving a third party.[45] This on-demand storage service does not automatically inject gas into the pipeline when nominations are made, but instead is used by the pipeline company to balance pressure across the entire system.

Furthermore, if the shipper is transporting gas under a no notice contract, the pipeline guarantees the shipper's ability to make changes to their schedule at any point during the scheduling period. Such an arrangement facilitates overscheduling by ensuring that any downscheduling adjustments in the final few hours will be automatically approved by the pipeline company. [46] As shown in Figures 9 and 10, both Firm A and Firm B hold sufficient *no notice* contracts and sufficient contracts originally sourcing gas from on-demand storage locations on Texas Eastern to engage in the levels of downscheduling we observe without incurring imbalance penalties.

[43] OFO warnings are extremely frequent on the Algonquin pipeline during the winter.

[44] It would be circuitous for a marketer to be involved here, as they would need their supplier to be complicit in injecting less gas than was scheduled.

[45] The second party is the pipeline operator: Both Texas Eastern and its storage service are operated by the same parent energy firm that operates the Algonquin pipeline. It is a near certainty the pipeline operator is aware of the scheduling practices on its pipelines that result in underutilized capacity. However, as pipeline companies make 99% of their revenues from fixed-charge payments associated with the quantity of reserved firm-service and only 1% on the use of these capacity contracts, their incentives lead them to favor constructing new pipeline capacity – to sell more contracts – rather than ensuring existing capacity is fully utilized. In New England, this incentive took the concrete form of the proposed Access Northeast pipeline expansion project for which the pipeline operator, Firm B, and another New England gas utility are co-developers.

[46] A requirement that interstate pipeline companies offer no notice contracts was included in FERC Order 636, the policy that mandated the unbundling of gas transportation service from the physical commodity, at the request of LDCs, who argued that continuation of their no notice service would also be needed in the new market structure to ensure they could reliably serve unexpected fluctuations in demand.

Figure 9: Average holdings of all contracts for capacity on the Algonquin pipeline from 2012 through 2016. Each bar represents a shipper; shippers are grouped by parent company for Firms A, B, and C, all of whom transport a majority of their gas through Algonquin's interconnects with the Texas Eastern Pipeline. Many of these contracts are no-notice.

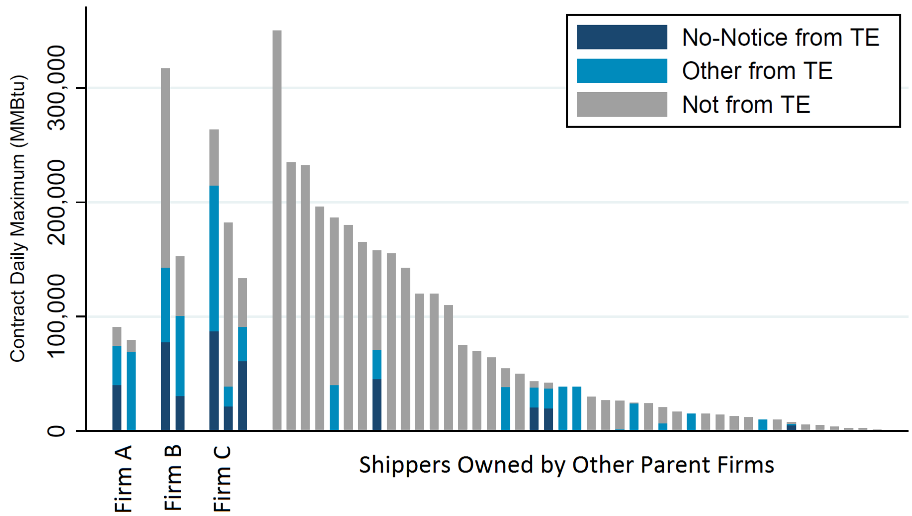

Figure 10: Average holdings of all contracts for capacity on the Texas Eastern pipeline for gas delivered to its two interconnects with the Algonquin pipeline. Of the gas sourced from Texas Eastern by Firms A, B, and C, much of it comes from storage. The upper limits of the aggregate downscheduling behavior observed (around 100,00- MMBtu) roughly match the sum of Firm A's and Firm B's no-notice contracts sourcing gas from storage on Texas Eastern.

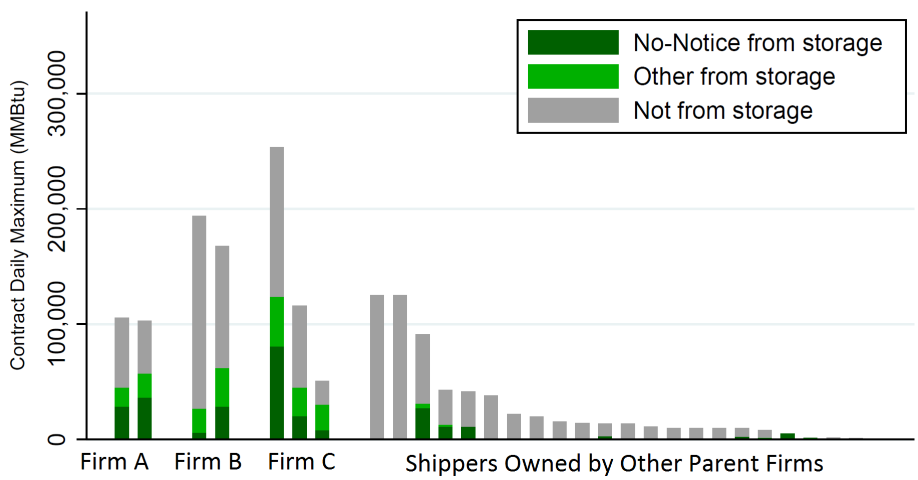

Contract rights vary on a roughly quarterly basis on the Algonquin pipeline, allowing us to exploit temporal as well as spatial variation. Graphically, Figures 11-13 show a relationship between these two firms' holdings of *no notice* contracts sourcing gas from the Texas Eastern pipeline (NN from TE contracts) and downscheduling over time at the segment level. For most segments of the pipeline, there appears to be no relationship between downscheduling behavior and NN from TE contracts. However, for the segment between Cromwell and Chaplin, where two suspect nodes are located, aggregate downscheduling is of roughly the same order of magnitude as Firm A's and Firm B's holdings of NN from TE contracts delivering to that segment. Between Oxford and Cromwell, where six suspect nodes are located, their NN from TE contract holdings appear to be an approximate upper bound on the level of downscheduling that occurs. We examine these apparent trends in greater detail using a set of specifications that introduce NN from TE contracts as an independent variable, the most inclusive of which is the following:

$$
\begin{aligned}
D_{it} = \alpha_0 &+ \beta_1 LDC_{it} + \beta_2 CT_{it} + \beta_3 C_{it} \\
&+ \beta_4 (LDC_{it} * CT_{it}) + \beta_5 (LDC_{it} * C_{it}) + \beta_6 (CT_{it} * C_{it}) + \beta_7 (LDC_{it} \times CT_{it} \times C_{it}) \\
&+ \alpha_1 HDD_{it} + \alpha_2 HDD_{it}^2 + \alpha_3 W_t + \lambda_t + \varepsilon_{it}
\end{aligned}
$$

Here, C_{it} is the total quantity of NN from TE contracts delivering gas to node i that are in effect on day t.[47] Column (3) of Table 2 indicates that NN from TE contract holdings are not necessarily a predictor of downscheduling across all LDC nodes, but Columns (4) and (5) demonstrate that they are an extremely strong predictor of downscheduling behavior at nodes located in Connecticut. In particular, for nodes in Connecticut, every additional MMBtu of NN from TE contracts is associated with an

[47] These data are available through "Index of Customers" reports, which interstate pipelines are required to make publicly accessible on their reporting web sites for the previous three years. As both contracts are reported separately for each the two pipelines, it impossible to confirm that a particular nomination or contract on Algonquin is used in combination with a particular contract on Texas Eastern. We therefore consider the "from TE" component of "NN from TE" to be proxy for a delivery right sourcing from storage from Texas Eastern, noting that both Firm A and Firm B hold sufficient *no notice* contracts on Texas Eastern sourcing gas from on-demand storage sites to engage in the levels of downscheduling we observe, as shown in Figure 10.

Figure 11: Aggregate downscheduling and contract positions of Firms A and B over time for the segment between the Oxford and Cromwell compression stations, in which six of the ten nodes that downschedule the most are located. These two firms' holdings of NN from TE contracts roughly correspond to an upper bound on the amount of downscheduling that occurs in this segment. Note: "All Contracts" is truncated at 350,000 in this series of charts for ease of presentation.

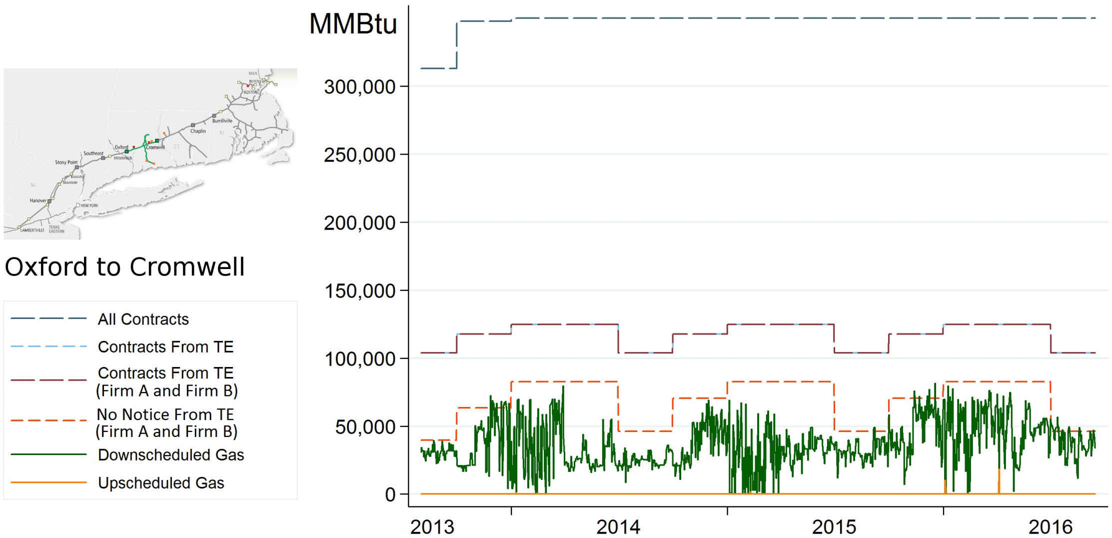

Figure 12: Aggregate downscheduling and contract positions of Firms A and B over time for the segment between the Cromwell and Chaplin compression stations, in which two of the ten nodes that downschedule the most are located. The level of downscheduling behavior is of roughly the same order of magnitude as these two firms' holdings of NN from TE contracts delivering gas to this segment.

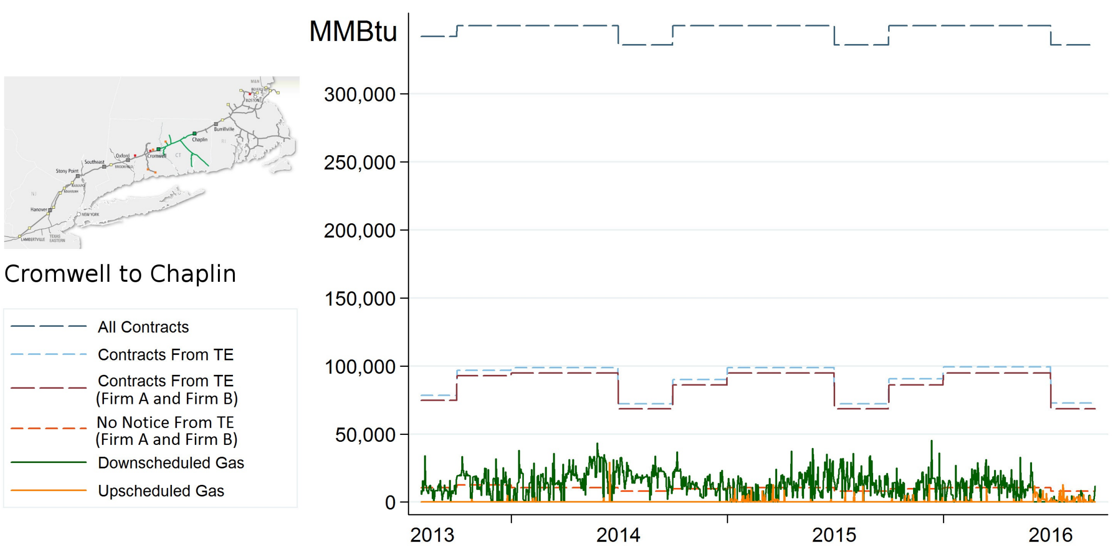

Figure 13: Aggregate downscheduling and contract positions of Firms A and B over time for the pipeline's "J System," which includes the Boston metropolitan area and serves many large electricity generators in addition to heating demand, and in which two of the ten nodes that downschedule the most are located. There is no obvious relationship here between downscheduling and NN from TE contracts held by Firms A and B.

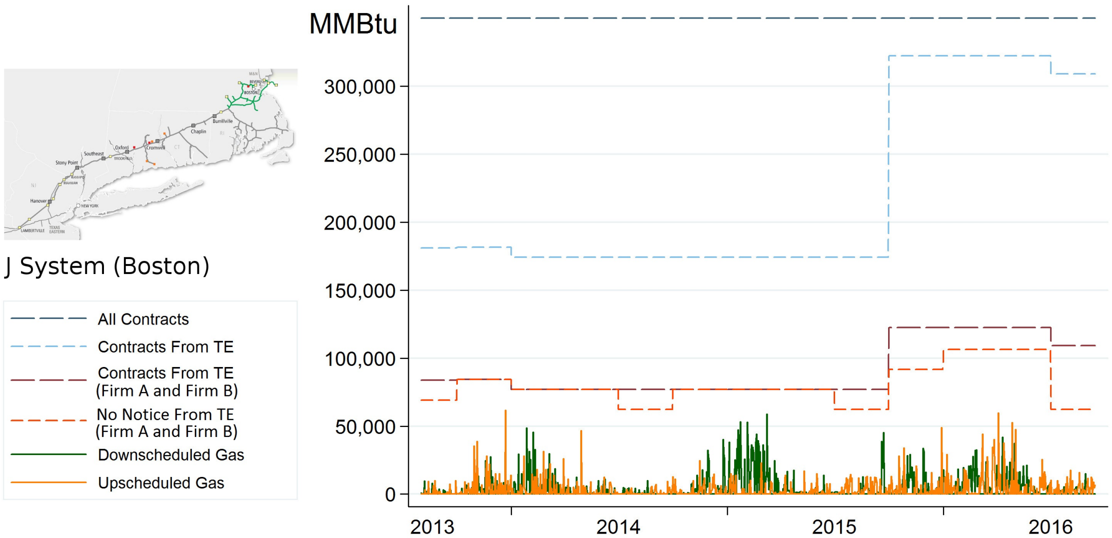

average increase in downscheduling of 0.42 MMBtu.[48]

3.5.4 Electricity Market Incentives

In Section 3.4, we outlined the incentives by which a firm owning generation capacity may benefit from higher electricity prices, and the connection between electricity prices and downscheduling in the gas transportation market. Here, we explore how generation capacity ownership relates to downscheduling behavior seen in the data.

Table 3 demonstrates that the only two firms that consistently engage in withholding behavior are also the two firms that hold the most and third-most generation capacity in New England among the LDCs served by the Algonquin pipeline. We note also that the firm holding the second most generation capacity holds no *no*

[48] Both *no notice* contracts and contracts sourcing gas from Texas Eastern are independently strong predictors of downscheduling behavior at nodes in Connecticut. Because these two contract characteristics overlap so heavily for Firms A and B, we are unable to empirically separate whether one or the other is critical to downscheduling.

notice contracts on Algonquin and an order of magnitude fewer contracts of any kind (including regular firm service contracts) than either Firm A or Firm B, meaning that while this firm benefits from higher electricity prices, its ability to affect the Algonquin pipeline's effective capacity constraint is very limited.

To bring electricity generation ownership into our empirical model, it is necessary to restrict our sample to LDCs only, as power plant nodes connected to Algonquin are of course operated by firms holding large quantities of generation capacity in the region. We therefore remove the LDC indicator from our previous specifications and replace it with a measure of generation capacity ownership. We run five separate specifications that mirror those of Table 2, the most inclusive of which is the following:

$$\begin{aligned} D_{it} = \alpha_0 &+ \beta_1 MW_{it} + \beta_2 CT_{it} + \beta_3 C_{it} \\ &+ \beta_4(MW_{it} \times CT_{it}) + \beta_5(MW_{it} \times C_{it}) + \beta_6(CT_{it} \times C_{it}) + \beta 7(MW_{it} \times CT_{it} \times C_{it}) \\ &+ \alpha_1 HDD_{it} + \alpha_2 HDD_{it}^2 + \alpha_3 W_t + \lambda_t + \varepsilon_{it} \end{aligned}$$

Here, MW_{it} is an indicator for whether the parent firm operating node i owns at least 100 megawatt (MW) of any type of generation capacity in the region.[49] Columns (1) and (2) of Table 4 demonstrate a significant correlation between owned generation

[49] We elect to use an indicator here to capture incentives created by ownership of both regulated and unregulated generation capacity. Our choice of 100 MW as the threshold is admittedly somewhat arbitrary, but we believe the reader will agree that it effectively separates firms with significant generation capacity in the region from those without (see Table 3). Beyond making the interpretation of interactions more difficult, using a continuous variable for the parent firm's generation capacity ownership is problematic for this set of regressions because the incentive pathways are quite different for different types of capacity. These pathways depend especially on whether the capacity is merchant unregulated or regulated, and additionally on many other characteristics, such as the fuel type, whether it is baseload or peaking, and the age of the facility. We demonstrate this with a robustness check wherein we perform the same set of regressions, but using the continuous variable for MW of generation owned rather than the binary one, producing generally erratic results (see Table A2 in the Appendix). We understand these counterintuitive results to be driven by the fact that Firm A downschedules more than Firm B, yet Firm B owns more generation capacity in New England, and we take this as suggestive evidence that, by merit of its directness, the incentive pathway for merchant unregulated generation is much stronger than the derivative pathways for regulated generation. Indeed, when we use a continuous measure of merchant unregulated MW owned as our independent variable of interest, our results are once again intuitively signed and highly significant (see Table A3 in the Appendix).

Table 3: The two firms observed to consistently engaged in downscheduling are also the two that hold the most and third-most generation capacity in New England out of the 11 parent energy companies that operate LDC-designated nodes on Algonquin.

LDC	Schedule Change in Last 3 Hours (MMBtu)	Schedule Change (Winter Only) (MMBtu)	Generation Capacity (MW)	Unregulated Capacity (MW)	All Contracts (MMBtu)	NN from TE Contracts (MMBtu)
Firm A	-41,506	-37,903	309	101	185,300	39,200
Firm B	-8,832	-18,320	1,177	0	632,400	105,100
Firm C	-225	-262	25	25	166,500	0
Firm D	-4	-10	0	0	32,000	0
Firm E	5	8	0	0	1,300	800
Firm F	7	36	0	0	24,500	0
Firm G	17	38	81	81	42,200	19,700
Firm H	21	56	23	0	7,600	5,000
Firm I	108	146	1,008	1,008	12,500	0
Firm J	177	227	9	0	768,900	167,100
Firm K	298	3	0	0	156,700	44,300

Not shown: 14 electricity generation firms and 2 industrial end users that operate nodes on Algonquin. We include firms' contract holdings on Algonquin as a proxy for the relative size of their natural gas operations on the pipeline, which reflects their ability to affect prices by withholding capacity. Firm I holds a large quantity of unregulated generation capacity, but its LDC operations on Algonquin are extremely limited, which limits ability to constrict pipeline capacity. Firm G appears to have both some ability and some incentive to withhold pipeline, though significantly less than the two firms we observe downscheduling. We note also that Firm A appears to withhold slightly more on average that it holds NN from TE contracts (with the "no notice" aspect the binding constraint). This suggests that "no notice" contracts may facilitate downscheduling, but are not absolutely required for it.

capacity and downscheduling behavior. In particular, nodes operated by firms that own at least 100 MW of generation capacity in New England downschedule over 1000 MMBtu more than nodes that do not. We find that this relationship is again driven primarily by nodes in Connecticut.

Interacting contract holdings with generation capacity ownership and state allows us to investigate nodes that have both the ability and the incentive to withhold capacity. Column (5) of Table 4 demonstrates that such nodes are clear outliers from the rest of the distribution in terms of their schedule adjustments in the clean-up cycle. With a high degree of statistical confidence, we find that nodes operated by firms owning at least 100 MW of generation capacity downschedule about 40% on average of their total NN from TE contract holdings if they are located in Connecticut, but only about 2% of their NN contract holdings if they are located outside of Connecticut.[50,51]

The available data enable us to detect downscheduling behavior and explore correlations with various mechanisms and incentives that may determine it. Our results are consistent with the hypothesis that these scheduling practices are representative of intentional capacity withholding intended to raise gas and electricity prices. However, we are unable to prove this conclusively without a source of plausibly exogenous variation to electricity market incentives.

One potential alternative explanation, for example, is that these two firms are exercising risk aversion by reserving an upper bound on the capacity they think they might need to ensure they will have access to it if demand turns out to be higher than expected. Although we cannot concisely reject this hypothesis given our data, we find it unlikely for two reasons: First, The two firms engaging in downscheduling are the only two firms that have a significant LDC presence on Algonquin and also have strong incentives to have higher electricity prices. Nodes operated by other firms appear to be able to consistently do a very good job of predicting the next day's demand. Second, by their very nature, the *no notice* contracts held by Firms A and B guarantee their ability to ramp up capacity usage at any point during the gas day,

[50] To challenge our understanding that *no notice* contracts sourcing gas from Texas Eastern in particular are requisite for systematic withholding, we re-run this set of regressions using contracts of any type and point of origin in their place. Table A4 in the Appendix presents these results, which are characteristically similar but generally smaller in magnitude. For example, the coefficient on the triple interaction of contracts, capacity, and Connecticut is still significant at the 1% level, but about 2.5 times smaller than when using *no notice* contracts from Texas Eastern, suggesting that these contracts are indeed particularly useful for downscheduling. However, the available data are insufficient to eliminate the alternative possibility that by happenstance some other unobserved factor causes these two LDCs that operate in Connecticut and also own significant generation capacity to hold NN from TE contracts in large quantities.

[51] While the ten most-downscheduling nodes are all operated by Firm A or Firm B, several independent marketers also manage contracts delivering gas to these locations. (The nodal operator manages actual flows and is typically responsible for the majority of the node's deliveries, but it is not necessary to operate a node to make deliveries to it. Independent marketers in particular use contracts to deliver gas to the region but do not operate nodes on the Algonquin pipeline.) To confirm that Firms A and B are indeed responsible for the withholding we observe, we run this set of regressions first considering only contracts held by them and then considering only contracts held by other firms. The results, presented in Table A5, clearly demonstrate that contracts held by Firms A and B drive the downscheduling we observe.

Table 4: Relationship between downscheduling and generation capacity ownership for LDC-designated nodes on Algonquin. "MW" here is a binary indicator for whether the node operator's parent firm owns at least 100 MW of generation capacity in New England.

	(1)	(2)	(3)	(4)	(5)
MW	1012.1**	101.5**	322.9		-18.28
	(445.2)	(43.46)	(386.9)		(21.29)
CT		2.959		274.4	-18.29
		(20.81)		(465.6)	(12.45)
Contracts			-0.00259	0.00225	-0.00262
			(0.00271)	(0.00346)	(0.00269)
MW×CT		1623.8**			501.9
		(768.1)			(551.2)
MW×Contracts			0.156		0.0190***
			(0.116)		(0.00574)
CT×Contracts				0.417***	0.00435
				(0.122)	(0.00447)
MW×CT×Contracts					0.395***
					(0.120)
HDD	9.814*	9.814*	9.814*	9.814*	9.814*
	(5.537)	(5.537)	(5.537)	(5.537)	(5.537)
HDD^2	-0.328**	-0.328**	-0.328**	-0.328**	-0.328**
	(0.155)	(0.155)	(0.155)	(0.155)	(0.155)
Weekend	83.48**	83.48**	83.48**	83.48**	83.48**
	(33.89)	(33.89)	(33.89)	(33.89)	(33.89)
N	109,152	109,152	109,152	109,152	109,152

Standard errors in parentheses (clustered at the node level)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

making it unnecessary for them to overbook capacity in this manner.

We therefore proceed to explore these firms' incentives further by modeling the effect of downscheduling behavior on gas and electricity prices, enabling us to compare their lost revenues in the gas market to their increased revenues in the electricity market, and estimate costs and distributional impacts in the process. Our methodology and results are presented in the next section.

3.6 Distributional and Welfare Effects

Next, we investigate the impact of downscheduling on wholesale gas prices and simulate the welfare effects on the wholesale electricity market, which market we believe bore the majority of the incidence of the downscheduling behavior. First, we use an instrumental variables approach to estimate the elasticity of demand for natural gas in the New England wholesale spot market, and use our estimated demand elasticity to construct a counterfactual Algonquin City Gate (ACG) price series.[52] Then, we use the observed and counterfactual gas prices to simulate the effects on economic dispatch in the New England wholesale electricity market, which we use to calculate the welfare and distributional impacts on generators and consumers. Lastly, we use our simulation results to compare LDCs' electricity market incentives to withhold to the opportunity cost of sacrificing excess capacity in the gas market.

3.6.1 Effect on Natural Gas Prices

Ideally, we would directly measure the effect of downscheduling on the the Algonquin City Gate (ACG) price, the main price index for wholesale gas transactions in New England.[53] Unfortunately, downscheduling is correlated with demand, through temperature, on a seasonal and daily basis. That is, these firms primarily engage in downscheduling on colder days and during the winter months when capacity is more likely to be constrained. Moreover, we believe day-to-day variation in downscheduling during the winter is partly driven by the quantity of excess contracts each firm has available (i.e., after supplying demand from residential and commercial heating customers), which is largely driven by temperature. We therefore use an instrumental variables approach to estimate the elasticity of demand for natural gas in the New

[52] Here, we suppose that all other marketers would not change their quantity supplied in response to decreases in downscheduling by other firms. That is, residual supply is inelastic.

[53] During our study period, the ACG was constructed by Platts, which collected price data by surveying market participants about their recent transcations.

England wholesale spot market and then use our estimated demand elasticity to construct a counterfactual Algonquin City Gate (ACG) price series.[54]

3.6.1.1 Elasticity of Demand for Gas

Typically, one would use the price of the same good in another market as an instrument for the price of a good in the market of interest (see for example, Hausman (1996) or Nevo (2001)). Our instrument for the ACG price is the Henry Hub (HH) price. The HH, which is located in Erath, LA, is a major distribution node and the primary pricing point for natural gas futures traded on the New York Mercantile Exchange. The HH price is determined by macroeconomic price shifters and unlikely to be affected by local supply and demand shifters affecting the ACG price, satisfying the exclusion restriction. Our instrumental variables specification is as follows:

$$D_t = \alpha_0 + \beta_1 P_t^G + \alpha_1 HDD_t + \alpha_2 HDD_t^2 + \alpha_3 W_t + \lambda_t + \epsilon_t$$

D_t is the natural log of quantity demanded (by electric generators) in the ACG wholesale gas market on day t and P_t^G is the natural log of the ACG price instrumented by the natural log of the HH price. We include controls for temperature using heating degree days (HDD_t and HDD_t^2), an indicator for whether the day is a weekend (W_t), and month-of-year fixed effects λ_t. The parameter β_1, captures the instrumented price elasticity of demand for ACG gas. To account for heteroskedasticity and autocorrelation, we use Newey and West's optimal lag selection criteria to specify the covariance matrix. Because we are primarily interested in this relationship on cold days, when downscheduling is most likely to have impacted the ACG price, we restrict the estimation to days with positive heating degrees (*i.e.* days where the maximum temperature is below sixty-five degrees farenheit).

[54] Here, we suppose that all other marketers would not change their quantity supplied in response to decreases in downscheduling by other firms. That is, residual supply is inelastic.

Table 5: Estimating the elasticity of demand for pipeline natural gas using OLS and Instrumental Variables, where the Algonquin City Gate (ACG) price is instrumented with the Henry Hub (HH) price.

	(1)	(2)	(3)	(4)
	$\log(q^G)$	$\log(p^G)$	$\log(q^G)$	$\log(p^{HH})$
$\log(p^G)$	-0.239***		-0.266***	
	(0.0197)		(0.0210)	
HDD	-0.00178	0.0172***	-0.00139	0.00768*
	(0.00314)	(0.00325)	(0.00313)	(0.00426)
HDD^2	-2.42e-05	0.000403***	-8.00e-06	5.94e-05
	(6.93e-05)	(6.92e-05)	(6.77e-05)	(7.68e-05)
Weekend	-0.0932***	-0.0477	-0.0945***	-0.00230
	(0.0182)	(0.0309)	(0.0181)	(0.0266)
$\log(p^{HH})$		1.129***		
		(0.0403)		
Constant	13.17***	-0.388***	13.19***	0.771***
	(0.0550)	(0.0612)	(0.0558)	(0.0752)
Observations	795	795	795	795
R-squared	0.660	0.788	0.659	0.089

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 5 summarizes the results of our instrumental variables approach. As a benchmark, Column (1) reports estimates from an OLS regression of the natural log of the ACG spot market quantity on natural log of the ACG price and covariates. The ACG price coefficient is -0.24 with a standard error of 0.02. Column (2) reports the results for a first-stage regression of the natural log of the ACG price on the natural log of the HH price and the same covariates described above. The t-statistic on the natural log of the HH price is 27.19 and the joint F-statistic is 209.27, suggesting the HH price is a strong instrument. Column (3) reports results for the main IV regression. The coefficient on the ACG price is -0.27 with a standard error of 0.02. Column (4) tests for endogeneity in the relationship between the HH price and the New England HDD terms. Only the coefficient on HDD is significant at the 10 percent level, though the magnitude (0.008) is quite small relative to the HH price coefficient in Column 3 (1.13), which is significant at the 1 percent level.

Our OLS and instrumented estimates of the elasticty of demand fall squarely within the range of estimates found in (Davis & Kilian, 2011) (*i.e.* , -0.1 to -0.34). For the calculations and simulations that follow, we use the instrumented coefficient and standard error.

3.6.1.2 Parameterizing the Demand Function

Next, we plug in our estimated price elasticity of demand into a constant-elasticity demand function: $D(p) = kp^{-0.27}$. Here, quantity demanded D is a function of a multiplier k, the ACG price p, and the elasticity of demand, which we estimated in the previous section to be -0.27. Using the daily quantity of gas demanded by electric generators (the primary source of spot-market demand) and the ACG price, we can solve for the vector of daily ks. Substituting this vector of ks back into the constant-elasticity demand function fixes the relationship between quantity and price, which allows us to calculate a counterfactual vector of prices from a counterfactual vector of quantities.

3.6.1.3 Constructing Counterfactual Quantities and Prices

To construct a vector of counterfactual daily quantities, we begin by summing the downscheduled quantities at the ten nodes operated by Firms A and B where downscheduling occured most frequently and intensively. As above, downscheduled quantities are measured as the change in scheduled quantities during the last three hours of the gas day, less the average fraction, across all other nodes, of capacity downscheduled in the last three hours of the same gas day. To account for downstream capacity constraints, the daily downscheduled quantity is bounded from above by unused capacity at the

Figure 14: The observed Algonquin Citygate gas price and the counterfactual gas price in a scenario without pipeline capacity withholding estimated by our IV model (weekly averages).

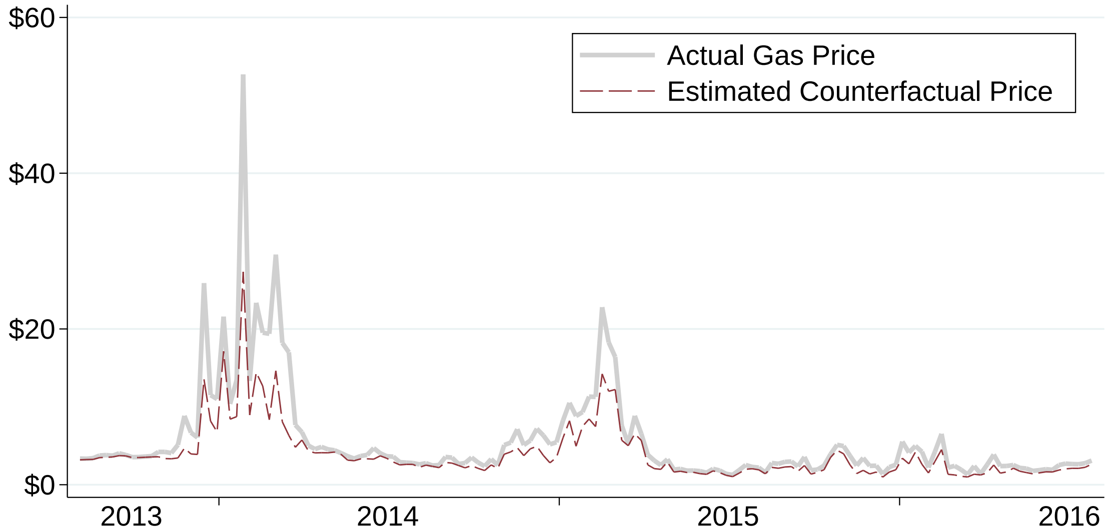

compression station in Burillville, RI.[55] Next, we add this to daily deliveries to nodes serving electric generators, our measure of observed quantity demanded, to arrive at a counterfactual vector of quantities. Substituting the counterfactual vector of quantities into the parameterized demand function yields a counterefactual vector of prices. We bound our counterfactual price vector from below using the vector of Texas Eastern Zone M3 (TEM3) prices. The price of gas at TEM3, which sits at the junction of the Texas Eastern and Algonquin pipelines, captures the cost of delivering gas to – or the price paid by buyers procuring gas for takeaway at – New England's doorstep.[56]

Differences between the series of observed and counterfactual Algonquin Citygate gas prices, shown in Figure 14, measure the impact of withholding. The average

[55] The Burrillville compression station has a very high average rate of capacity utilization and is the last potential bottleneck before the Algonquin pipeline branches into a nodal network. Unused capacity at Burillville is measured as the difference between end-of-day scheduled quantity and daily operational capacity.

[56] That is, using the price at TEM3 as a lower bound allows us to identify and correct for instances when our counterfactual quantity implies a slack capacity constraint. Failing to account for this would overstate the effect of capacity withholding.

price difference in our three-year sample is \$1.68 (39 percent) and the price difference during only the winter months of each year is \$3.82 (68 percent). We note again that this is a spot market price and, as such, is not directly representative of prices paid by residential and commercial heating customers, who are supplied by LDCs that acquire the vast majority of their gas through long-term contracts. Because electricity generators are the principal buyers of spot market gas, these price differences have distributional consequences primarily in the wholesale electricity market, where a convex generation supply curve amplifies their effects. We proceed in the next section to estimate electricity market impacts using a simulation model with our estimated counterfactual gas price series as an input.

3.6.2 Electricity Market Impacts

Measuring the effect of capacity withholding on the wholesale electricity price is crucial both to understanding firms' incentives to withhold and to determining the resultant costs borne by ratepayers. We accomplish this by simulating New England's wholesale electricity market under both observed and counterfactual gas price series and comparing the two. Broadly, our estimation procedure consists of a.) re-constructing ISO-NE's wholesale electricity auction market, b.) non-parametrically estimating a distribution of counterfactual generator supply bids using our counterfactual gas price series, and c.) clearing the market a large number of times using draws from the counterfactual bid distribution.

3.6.2.1 New England's Wholesale Electricity Market

In New England's deregulated wholesale electricity market, prices and allocations are determined by a day-ahead auction operated by the nonprofit market administrator ISO-NE. At 10 a.m. the day before each operating day,[57] generators submit supply

[57] Unlike the gas day, ISO-NE's electricity market operating day simply runs from 12 a.m. to 11:59 p.m. Eastern Time.

bids stating how much generation they will produce at given prices for each hour of the day and distribution utilities submit demand bids stating how much energy they will require at each hour of the day. The market-clearing price (expressed in \$/MWh) is the price bid in by the generator providing the most expensive unit of generation required to meet demand (the "marginal generator"). All generators that bid in generation at prices lower than the clearing price are paid the clearing price for each MWh of energy they produce.[58]

Our re-construction and clearing of the day-ahead market is a departure from ISO-NE's actual dispatch procedure in three ways. First, we exclude exports to and imports from the neighboring New York and Canadian energy markets. Second, we do not attempt to incorporate transmission constraints or lines losses.[59] Finally, our model does not incorporate no-load fees[60] or startup costs.[61] We believe assuming zero no-load fees and startup costs is reasonable for our application because a.) although they compose a substantial portion of total power costs paid by consumers, they do not enter the wholesale market price, as it is determined by the variable cost component of generators' bids, and b.) the key input we are adjusting—the gas price—is fundamentally a component of generators' variable costs.

[58] In addition to the day-ahead market, ISO-NE also operates a real-time auction market that adjusts for unforeseen fluctuations in supply and demand throughout the operating day. We focus our analysis on the day-ahead market as 95 percent of ultimately consumed energy is traded there and real-time prices generally closely track those of the day-ahead market (Kim, 2016).

[59] We perform a bias correction targeted at adjusting for transmission constraints on high-price days, which is discussed further in Section 3.6.2.3 and Appendix A.3.3.2.

[60] In addition to the principle variable cost component of their bids, generators can also optionally submit a no-load fee. No-load fees are a fixed sum to be paid to the generator for each hour it is called upon to operate that are designed to capture its fixed operating costs (such as labor).

[61] Similarly, generators can choose to include cold-, intermediate-, and/or hot-start fees. These startup fees are paid to generators each time they switch from being offline or in standby mode to operating. While we observe the various startup costs as submitted in generators' bids, the associated times for each type of start vary by generator and are not made publicly available.

Figure 15: Left: A sample of three actual bid supply offers submitted by a typical fuel-switching gas and oil generator. Each corresponds to one hour of one day for three separate days in the study period. The steep increase at 36 MW indicates switching to oil. Right: The market bid supply curves for those three hours, constructed by aggregating all generators' supply bids.

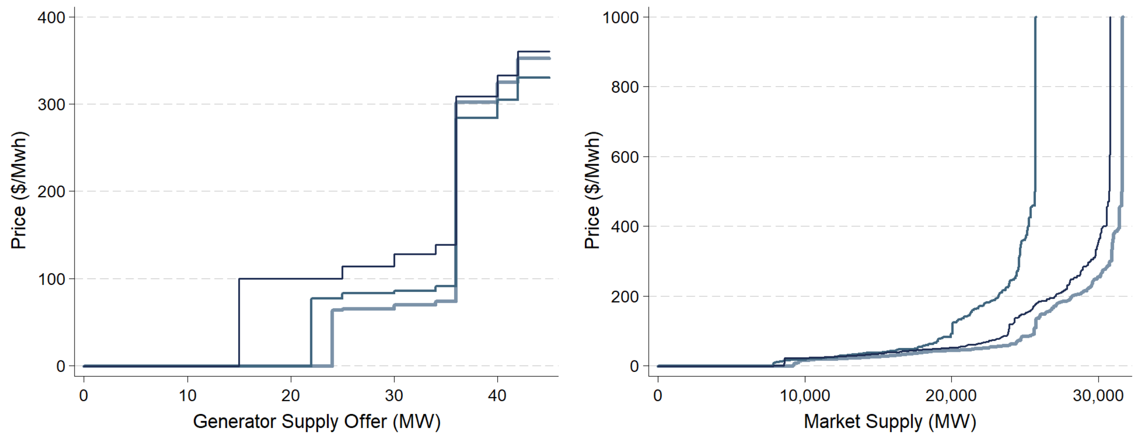

3.6.2.2 Estimating Generator Bid Functions

A single generator is allowed to offer different quantities of generation at different prices, and accordingly their supply bids take the form of step functions of up to ten price-quantity pairs for each hour of each day (see Figure 15 for an illustration). Because the size and number of steps often varies across days even within an individual plant, it is infeasible to parametrically estimate a relationship between the gas price and generators' bids. We therefore estimate generator bids using a non-parametric approach that employs a variant of nearest neighbor matching combined with a resampling procedure. Our approach is a novel extension of Reguant (2014)'s technique for estimating generators' expectations of competitors' bid functions.

Our simulation model is built upon the complete set of actual generator bids for the study period, which is publicly available in an anonymized format through ISO-NE's website.[62] We define the bid function generator i actually submitted on day t as $\bar{b}_{it}$ and the full set of actual generator bids as $\bar{B}$. Each bid function consists

[62] Although we de-anonymize these bid data in later sections to explore distributional impacts of withholding, the core of our simulation model applies the same methodology to all generators regardless of fuel type and thus does not require de-anonymization.

of 24 sets of up to 10 quantity-price pairs. Instead of of imposing a structure on these price-quantity pairs, we use matching to estimate the entire bid function b_{it} conditional on electricity market demand (D_t^E) and temperature (HDD_t) using a two-step matching procedure.[63] We employ this procedure to estimate two bid functions for each generator for each day: one corresponding to their expected bid given the actual observed gas price (b_{it}^1) and the other corresponding to their expected bid given the counterfactual gas price we estimate in the preceding section (b_{it}^2).[64]

$$b_{it}^1 = \mathbb{E}_i[b_{it}(p_t^G)|D_t^E, HDD_t]$$
$$b_{it}^2 = \mathbb{E}_i[b_{it}(p_t^{G,cf})|D_t^E, HDD_t]$$

In the first stage, we construct a sub-sample of potential match days that are highly similar to the target day in relevant electricity market conditions that determine generator bids. In particular, we isolate the 5 percent of days in the three-year study period that are most similar to day t in temperature and electricity market demand by Mahalanobis distance. For generators that were online during the entire three-year period, this generates a subsample of 56 days that are highly similar to the target day in these two key determinants of generators' bids.

In the second stage, we select from the first stage subsample the 3 days that experienced gas prices closest to either the actual gas price on t (for estimating b_{it}^1) or

[63] For market demand D_t^E, we use the predicted peak demand for the day-ahead market, which is published at the same time bids must be submitted for the day-ahead market and therefore represents market expectations of demand at that time.

[64] Although we observe generator i's actual bid function on day t, which is intrinsically the best measure of their bid function given the actual gas price, we estimate (b_{it}^1) in order to employ the same simulation methodology in both the actual and counterfactual scenarios to remove any bias introduced by the methodology that is common to both.

to the estimated counterfactual gas price for t (for estimating b_{it}^2).[65,66] We use 3 days because it is again roughly 5 percent of the first-stage subsample for generators that were online during the entire period.[67] Indexing the scenario (actual or counterfactual gas price) as s, each predicted bid function $\widehat{b}_{it}^s$ then consists of three of generator i's actual bids $\{\bar{b}_{i\tau_1^s}, \bar{b}_{i\tau_2^s}, \bar{b}_{i\tau_3^s}\}$ from three days (τ_1^s, τ_2^s, and τ_3^s) that are highly similar to t in temperature, electricity demand, and either gas price (for $\widehat{b}_{it}^1$) or the estimated counterfactual gas price (for $\widehat{b}_{it}^2$):

$$\widehat{b}_{it}^1 = \{\bar{b}_{i\tau_1^1}, \bar{b}_{i\tau_2^1}, \bar{b}_{i\tau_3^1}\} = f(b_{it}(p_t^G)|D_t^E, HDD_t)$$
$$\widehat{b}_{it}^2 = \{\bar{b}_{i\tau_1^2}, \bar{b}_{i\tau_2^2}, \bar{b}_{i\tau_3^2}\} = f(b_{it}(\widehat{p_t^{G,cf}})|D_t^E, HDD_t)$$

In order to incorporate the statistical uncertainty of the IV-estimated counterfactual gas price series into our final results, for each iteration of the simulation we randomly draw $\widehat{p_t^{G,cf}}$ from a distribution of estimated counterfactual prices $\widehat{P_t^{G,cf}}$.[68] We draw 100 complete counterfactual price series and perform the second stage matching procedure for each, generating 100 sets of 3 match days for each generator for each day.

3.6.2.3 Simulating the Electricity Market

To test the general viability of our simulation model, we first re-construct and clear the wholesale auction market using actual generator bids $\bar{B}$ without any resampling. This

[65] As with traditional matching estimators, some bias may be introduced due to the fact that matched days will not be exact matches on relevant covariates (Abadie & Imbens, 2006). In particular, within each subset of first-stage match days, correlation between gas price, demand, and weather would cause lower gas price days to systematically be matched with days with slightly lower demand and temperature. To ensure the results are driven by changes in the gas price and not these other variables, we use a regression-based bias correction procedure, which is discussed in detail in Appendix A.3.3.

[66] We exclude the same day for estimation of b_{it}^1.

[67] For generators that were offline for some days in the study period (about 10 percent of observations), we use 5 percent in the first stage and three days in the second stage.

[68] Specifically, we adjust our point estimate for the elasticity of demand by a random shock drawn from a Gaussian distribution of the estimated error, calculate the counterfactual gas price series using that adjusted elasticity, and lower-bound it ex-post with the Texas Eastern M3 price.

Figure 16: Estimated wholesale electricity prices generated by our cur simulation model that reconstructs and clears the day-ahead market using actual bids with no resampling.

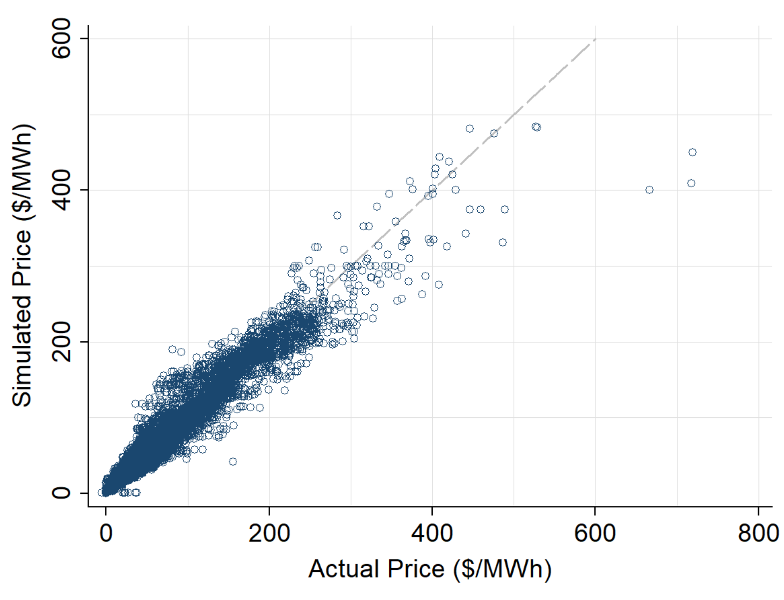

entails aggregating all price-quantity pairs submitted by all generators into a market bid supply function for each hour of each day and determining the most expensive step required to meet that hour's demand, recovering the price from that step as the output.[69]

In comparing the generated price series to the true day-ahead price series, we assess that our model introduces a relatively small degree of noise and bias by assuming away trade, no-load costs, and startup fees. As shown in Figure 16, prices estimated by the simulation using $\bar{B}$ closely track actual day-ahead prices. The coefficient of correlation between the two is .969. Our simulation-generated prices have a slight upward bias, with an average clearing price of $48.77 versus an average actual dayahead price of $47.53. We also calculate total energy costs by independently multiplying price by demand for each hour and then summing across the study period, finding that the

[69] Rather than directly using the day-ahead demand, we adjust demand using a regression-based bias correction procedure that accounts for transmission constraints, which is discussed in detail in Appendix A.3.3.2. This bias correction does not change our market-level results, but greatly improves the model's ability to capture dispatch of peaking generators specifically.

estimated energy cost from the simulation using actual bids is about \$19.184 billion versus \$18.802 billion when using the actual day-ahead clearing price.[70,71]

Next, we re-construct and clear the wholesale auction market using the estimated distribution of generator bids under the actual gas price $\widehat{B_1}$. For each of 100 iterations, we randomly draw one of the three bid functions in $\widehat{b}_{it}^1$ for each generator for each day, use those draws to construct an aggregate bid supply curve, and clear the market for each hour of each day, recovering a distribution of estimated electricity price series.

While this is certainly a roundabout approach to estimating electricity prices given that we observe both actual clearing prices and actual generator bids, it enables us to assess the level of uncertainty and/or bias that is introduced by the matching-resampling procedure and to remove that bias to some extent. At \$45.84, the average estimated clearing price from the simulation that matches based on the actual gas price is slightly lower than the actual clearing price, and the coefficient of correlation between the two is .906. To calculate the estimated distribution of total energy costs $\widehat{C_1^e}$, within each iteration we multiply demand by clearing price for each hour and sum over the three-year study period. The mean of this distribution, \$18.159 billion, is our estimate for total energy cost paid by ratepayers from the resampling simulation model using the actual gas price. This figure is lower than the actual energy cost by \$643 million, which we note is almost an order of magnitude smaller than the final predicted estimate of the cost to ratepayers due to withholding activity.[72]

Next, we re-construct and clear the electricity market using the estimated distribution of counterfactual generator bids under the zero withholding scenario. We

[70] We use actual day-ahead demand instead of bias-adjusted demand for all cost calculations.

[71] Note that while energy costs account for the lions share of ratepayer electricity bills, rates additionally include fixed costs such startup costs, no-load fees, distribution costs, and administrative fees that we do not attempt to estimate.

[72] In other words, while it's clear from the discrepancy that our matching-resampling procedure introduces some bias, we can be reasonably confident that our final estimates are accurately signed and correct in magnitude because this bias is small relative to the final estimated damages from pipeline capacity withholding. Moreover, we correct for bias that's common to both $\widehat{C_1^e}$ and $\widehat{C_2^e}$ by using the difference between the two (rather than the difference between $\widehat{C_2^e}$ and the actual energy cost, which would be higher) for our final damages estimate.

Table 6: Actual electricity price and energy cost (1), estimated prices and costs from simulation models using actual bids without resampling (2), resampling based on actual gas price (3), and resampling based on the estimated counterfactual gas price (4), and our estimate of the excess energy cost due to withholding (5).

	(1) Actual Elec. Price	(2) Actual Bids	(3) Actual Gas Price	(4) Counterfactual Gas Price	(5) Difference Btw. (3) and (4)
Price	47.56	48.77	45.84	37.18	8.66
($)	-	-	(0.11)	(0.62)	(0.62)
Energy Cost	18.802	19.184	18.158	14.746	3.412
($B)	-	-	(.042)	(.255)	(.255)

Bootstrapped standard errors in parentheses (approximate as distributions are not exactly normal)

follow the same methodology as before except now sampling generator bids from $\widehat{B_2}$ to construct a distribution of 100 counterfactual electricity price series. Results are presented in of Table 6. The average clearing price estimated for the counterfactual scenario is $37.18, which is $8.66 (19 percent) lower than our simulation estimate of the actual electricity price using the same matching and resampling methodology. Figure 17 demonstrates that the price differential was largely driven by the "polar vortex" of 2013-14, when electricity prices were $41.46 (56 percent) higher on average due to capacity withholding. Our model estimates that electricity prices were $13.97 (23 percent) higher due to capacity withholding during the winter of 2014-15, $3.54 (11 percent) higher during the winter of 2015-16, and $3.21 (11 percent) higher during the summers in the study period.[73] We attribute the limited impact of withholding on prices during the winter of 2016-17 to generally warmer temperatures (37 °F on average versus 28 °F and 27 °F during the other two winters).

We calculate total energy costs in the counterfactual gas price scenario $\widehat{C_2^e}$ as before and then subtract $\widehat{C_2^e}$ from $\widehat{C_1^e}$ to produce a distribution of estimated excess energy costs due to pipeline capacity withholding. This distribution, shown in Figure 18,

[73] We again define winter in this context as December 1 - March 31 following the Algonquin pipeline's tariff.

Figure 17: Actual and simulated wholesale electricity prices averaged by week. The strongest impacts of pipeline capacity withholding were realized during the winters of 2013-14 and 2014-15.

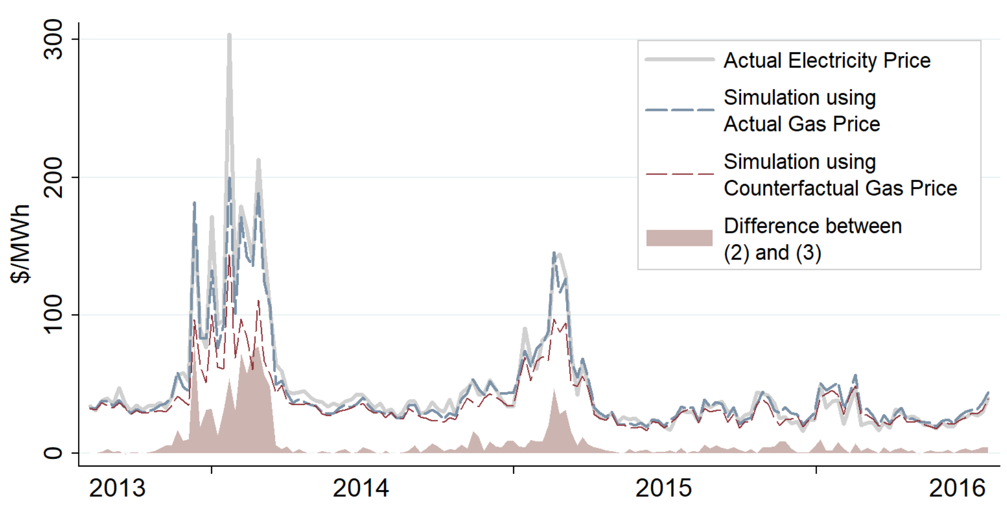

Figure 18: Distribution of estimated excess energy costs due to pipeline capacity withholding over the three-year study period.

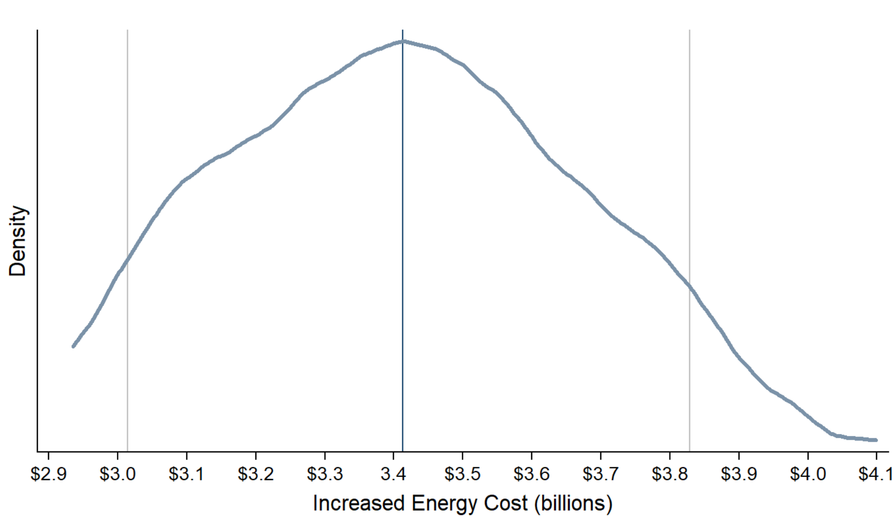

incorporates the uncertainty associated with our wholesale gas price estimates and, to the extent possible, the uncertainty introduced by our resampling procedure. The mean, \$3.412 billion, is the estimated total excess energy cost due to pipeline capacity withholding and the 95 percent confidence interval ranges from \$3.014-\$3.828 billion.

3.7 Conclusion

To date, most growth in gas deliveries for electric generation has utilized legacy pipeline infrastructure. Much of this infrastructure was built to supply heating customers, with institutions designed to manage heating demand. As demand for gas for electric generation grows, gas-fired electric generators will increasingly find themselves in competition with legacy pipeline customers for scarce pipeline capacity. This new competitive environment is likely to affect the incentives of firms, especially those that operate in both markets, in unforeseen ways. Therefore, it will be critical that the institutions governing the trade and transport of gas, both as a heating source and for electric generation, are harmonized and structured efficiently.

In this paper, we identify a major inefficiency spanning the natural gas transportation and wholesale electricity markets. We quantify the extent to which two firms withheld pipeline capacity and detail the institutional arrangements that allowed these firms to execute their withholding strategy. Using an instrumental variables approach, we identify the effect of withholding on the wholesale natural gas price, and we employ an electricity dispatch model to trace that impact through the wholesale electricity market. During our study period, this withholding of pipeline capacity resulted in a transfer from electricity ratepayers to electricity generators (and their fuel suppliers) of $3.6 billion.

While only the firms in question can explain their actions and intentions, we have shown that it is consistent with the exercise of market power and inconsistent with the most plausible alternative explanations, namely risk aversion on the part of these firms. Furthermore, we have demonstrated that the decreased supply of gas to New England generators caused by this behavior increased revenues in the wholesale electricity market for these (and other) firms by margins that likely exceeded the opportunity cost of using the capacity to supply the gas spot market, especially when accounting for state revenue-sharing rules.

Just as transmission constraints create opportunities for market power in the electricity sector, so too can capacity constraints create opportunities for market power in the natural gas transportation industry. As natural gas prices have fallen, demand for gas to power the electric grid has steadily increased, creating new rivals and tying closer the gas transportation and electricity markets. In this context, it has become increasingly important that already existing pipeline capacity is optimally utilized not only to protect the interests of gas and electricity ratepayers, but also to ensure that unbiased price signals lead to an efficient level of new pipeline development. Pipeline market reforms that facilitate more flexible contracting mechanisms, more frequent scheduling cycles, and act to prevent capacity withholding, or impose a cost for capacity withholding and create a publicly-available record of capacity withholding; all of which will serve to better align the gas transport and electricity markets, could help to create more liquid markets in which firms find it more difficult to exert market power.

While the analysis here has focused on identifying the exercise of market power on one particularly congested pipeline serving New England, severe bottlenecks and vertically integrated firms coexist in many other natural gas transportation markets. To what extent capacity withholding has led to pipeline underutilization in other regions is an important area for future study.

References

Abadie, Alberto, & Imbens, Guido W. 2006. Large sample properties of matching estimators for average treatment effects. *econometrica*, **74**(1), 235–267.

Boots, Maroeska G, Rijkers, Fieke AM, & Hobbs, Benjamin F. 2003. GASTALE: An oligopolistic model of production and trade in the European gas market. *Energy Research Centre of the Netherlands (ECN), ECN*, **35**(1).

Borenstein, Severin, & Bushnell, James. 1999. An empirical analysis of the potential for market power in California?s electricity industry. *The Journal of Industrial Economics*, **47**(3), 285–323.

Borenstein, Severin, Bushnell, James, Kahn, Edward, & Stoft, Steven. 1995. Market power in California electricity markets. *Utilities Policy*, **5**(3), 219–236.

Borenstein, Severin, Bushnell, James, & Knittel, Christopher R. 1999. Market power in electricity markets: beyond concentration measures. *The Energy Journal*, 65–88.

Borenstein, Severin, Bushnell, James B, & Wolak, Frank A. 2002. Measuring market inefficiencies in California's restructured wholesale electricity market. *The American Economic Review*, **92**(5), 1376–1405.

Cremer, Helmuth, & Laffont, Jean-Jacques. 2002. Competition in gas markets. *European Economic Review*, **46**(4), 928–935.

Cremer, Helmuth, Gasmi, Farid, & Laffont, Jean-Jacques. 2003. Access to pipelines in competitive gas markets. *Journal of Regulatory Economics*, **24**(1), 5–33.

Davis, Lucas W, & Kilian, Lutz. 2011. The allocative cost of price ceilings in the US residential market for natural gas. *Journal of Political Economy*, **119**(2), 212–241.

EIA. 2017 (January). *Annual Energy Outlook with Projections to 2050*. US Energy Information Agency.

FERC. 2017. *Cost-of-Service Rate Filings*. `https://www.ferc.gov/industries/gas/gen-info/rate-filings.asp`. Accessed: 2017-07-31.

Hart, Oliver, Tirole, Jean, Carlton, Dennis W, & Williamson, Oliver E. 1990. Vertical integration and market foreclosure. *Brookings papers on economic activity. Microeconomics*, **1990**, 205–286.

Hausman, Jerry A. 1996. Valuation of new goods under perfect and imperfect competition. *Pages 207–248 of: The economics of new goods*. University of Chicago Press.

Holz, Franziska, Von Hirschhausen, Christian, & Kemfert, Claudia. 2008. A strategic model of European gas supply (GASMOD). *Energy Economics*, **30**(3), 766–788.

Hunger, David. 2003. Analyzing gas and electric convergence mergers: A supply curve is worth a thousand words. *Journal of Regulatory Economics*, **24**(2), 161–173.

ICF. 2015. New England Energy Market Outlook: Demand for Natural Gas Capacity and Impact of the Northeast Energy Direct Project. *ICF International Inc.*

Joskow, Paul L, & Tirole, Jean. 2000. Transmission rights and market power on electric power networks. *The Rand Journal of Economics*, 450–487.

Kim, Harim. 2016. Heterogeneous Impacts of Cost Shocks, Strategic Bidding and Pass-Through: Evidence from the New England Electricity Market.

Lise, Wietze, & Hobbs, Benjamin F. 2009. A dynamic simulation of market power in the liberalised European natural gas market. *The Energy Journal*, 119–135.

Nevo, Aviv. 2001. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, **69**(2), 307–342.

Oliver, Matthew, & Mason, Charles. 2018. An overview of interstate gas pipeline regulation in the United States. 01.

Oliver, Matthew E, Mason, Charles F, & Finnoff, David. 2014. Pipeline congestion and basis differentials. *Journal of Regulatory Economics*, **46**(3), 261–291.

Ordover, Janusz A, Saloner, Garth, & Salop, Steven C. 1990. Equilibrium vertical foreclosure. *The American Economic Review*, 127–142.

Reguant, Mar. 2014. Complementary bidding mechanisms and startup costs in electricity markets. *The Review of Economic Studies*, **81**(4).

Rose, Judah, Gerhardt, David, Karp, John, Brock, Frank, Sakya, Trisagni, & Shrestha, Mohit. 2014. Of Polar Vortexes and PJM Price Spikes. *ICF International Inc.*

Salop, Steven C., & Scheffman, David T. 1983. Raising Rivals' Costs. *The American Economic Review*, **73**(2).

Spectra. 2016. *FERC Gas Tariff of Algonquin Gas Transmission, LLC*. Sixth revised volume no. 1 edn.

Vazquez, Miguel, Berzosa, Ana, & Vazquez, Carlos. 2006. An analysis of the oligopolistic effects of the integration between electricity and gas suppliers in the Spanish market. *Available at SSRN 1262743.*

Appendices

Appendix for Chapter 1

A.1.1 Theory Extensions

This section extends the one-period theoretical framework in the main text to a dynamic model, then further extends that dynamic model to separately consider emissions associated with completions, equipment, and maintenance. Although the results presented here are not practical for extending to empirical analysis given current data limitations, they provide intuition for many of the intertemporal aspects of firms' production and emission decisions.

A.1.1.1 Dynamic Model

I begin by supposing that instead of choosing a level of production and an emission rate within a single period, the firm owns a stock of wells at the start of each period, chooses how many new wells to drill, and chooses the emission rate of these new wells. For simplicity, I assume wells are homogeneous and that each well generates one unit of production per period. I further suppose the number of wells owned by the firm in a given period W_t can be broken down into new wells drilled that period W_t' and wells leftover from the previous period. Building from Equation 3, the firm's instantaneous profit function is reformulated as:

$$\pi_t = P_t W_t(1 - R_t) - C_1(W_t') - W_t' c_2(R_t') \tag{11}$$

I disregard maintenance costs (for now), so the only costs incurred by the firm are those for new wells. Facility emission rates evolve as the weighted average of the emission rate of existing wells plus the emission rate of new wells, and wells depreciate at a rate of $1-\delta$. Accordingly, the equations of motion for these two variables are as follows:

$$W_t = \delta W_{t-1} + W'_t \tag{12}$$

$$R_t = \frac{\delta W_{t-1}}{W_t} R_{t-1} + \frac{W'_t}{W_t} R'_t \tag{13}$$

Applying a discount factor of β, the firm's intertemporal optimization problem is:

$$\pi = \max_{\{W'_t, R'_t\}_t} \sum_{t=0}^{T} \beta^t \big(\mathbb{E}[P_t] W_t (1-R_t) - C_1(W'_t) - W'_t c_2(R'_t)\big) \tag{14}$$

$$S.T. \quad W_t = \delta W_{t-1} + W'_t \tag{15}$$

$$R_t = \tfrac{\delta W_{t-1}}{W_t} R_{t-1} + \tfrac{W'_t}{W_t} R'_t \tag{16}$$

$$W'_t \geq 0,\ R'_t \in (0,1) \tag{17}$$

Starting from any given period τ, the firm's profit stream is only dependent on their choice of R' and W' in τ and in future periods. Substituting in the equation of motion for R_t facilitates taking the first-order condition for emissions from new wells:

$$\mathbb{E}[\sum_{t=\tau}^{T} \pi_t] = P_\tau W_\tau \big(1 - \tfrac{\delta W_{\tau-1}}{W_\tau} R_{\tau-1} - \tfrac{W'_\tau}{W_\tau} R'_\tau\big) - C_1(W'_\tau) - W'_\tau c_2(R'_\tau) \tag{18}$$

$$\begin{aligned} &+ \beta \mathbb{E}[P_{\tau+1}] W_{\tau+1} \big(1 - \tfrac{\delta W_\tau}{W_{\tau+1}} (\tfrac{\delta W_{\tau-1}}{W_\tau} R_{\tau-1} + \tfrac{W'_\tau}{W_\tau} R'_\tau) - \tfrac{W'_{\tau+1}}{W_{\tau+1}} R'_{\tau+1}\big) - C_1(W'_{\tau+1}) - W'_{\tau+1} c_2(R'_{\tau+1}) \\ &+ \beta^2 \mathbb{E}[P_{\tau+2}] W_{\tau+2} \big(1 - \tfrac{\delta W_{\tau+1}}{W_{\tau+2}} (\tfrac{\delta W_\tau}{W_{\tau+1}} (\tfrac{\delta W_{t-1}}{W_t} R_{t-1} + \tfrac{W'_t}{W_t} R'_t) + \tfrac{W'_{\tau+1}}{W_{\tau+1}} R'_{\tau+1}) - \tfrac{W'_{\tau+2}}{W_{\tau+2}} R'_{\tau+2}\big) - \ \ ... \\ &+ \ \ ... \end{aligned} \tag{19}$$

$$\frac{\partial \pi}{\partial R'_t} = -W'_\tau \tfrac{\partial c_2}{\partial R'_\tau} - W'_\tau P_\tau - \beta \delta W'_\tau \mathbb{E}[P_{\tau+1}] - \beta^2 \delta^2 W'_\tau \mathbb{E}[P_{\tau+2}] + \ \ ... \ \ = 0 \tag{20}$$

$$-\frac{\partial c_2}{\partial R'_\tau} = \sum_{\iota=0}^{T} \beta^\iota \delta^\iota \mathbb{E}[P_{\tau+\iota}] \tag{21}$$

Equation 21 shows that the firm chooses an emissions rate for new wells that sets the marginal cost of having emission rate R'_τ for new wells equal to the present discounted value of expected future prices. Unfortunately, it is not possible within the GHGRP data to separate out emissions from new wells from those of existing wells, as a single emission rate is reported for each facility. With the addition of DrillingInfo variables for the firms' number of existing wells and new wells,[64] it is hypothetically possible to back out the emission rate for new wells. However, this process breaks down in practice, possibly because it introduces additional noise by amplifying any imperfect matching between the two datasets. The advent of satellite methane emission monitoring at the level of spatial resolution requisite for estimating emissions from individual well sites (or well-level reporting within the GHGRP or another survey) would generate a direct measurement of R'_τ that could be used to estimate Equation 21.

A.1.1.2 Incorporating Emission Sources

Finally, I extend the model to separately consider emissions that are associated with capital purchase decisions, emissions associated with maintenance (i.e. leak detection and repair), and emissions associated with completions. This breakdown is relevant specifically within the dynamic model because emissions from each of these sources follow from decisions the firm makes based on very different time frames.[65] For example, when a firm purchases equipment for a new well, their decision on how much to expend to acquire less-emitting equipment is based primarily on expectations of future gas prices over the expected lifetime of the equipment.

[64] Annual production from new wells and existing wells is also necessary here, as the assumption that each well extracts gas at the same rate does not hold in actuality.

[65] This is also true of equipment upgrades for existing wells, where there is an additional cost associated with forgoing the remaining potential lifetime of existing equipment and labor costs for installing the upgrade that would not be incurred in its absence. For simplicity, I omit equipment upgrades from the model.

With these modifications, the firm's optimization problem is now:

$$\pi = \max_{\{W_t', M_t, R_t'^k, R_t'^c\}_t} \sum_{t=0}^{T} \beta^t \big[\mathbb{E}[P_t]\big(W_t^e(1 - R_t^e f(M_t)) + W_t'(1 - R_t'^k - R_t'^c)\big) - C_1(W_t') - W_t^e c_m(M_t) - W_t' c_k(R_t'^k) - W_t' c_c(R_t'^c)\big] \tag{22}$$

$$S.T. \quad W_t^e = \delta(W_{t-1}^e + W_{t-1}') \tag{23}$$

$$R_t^e = \tfrac{\delta W_{t-1}^e}{W_t^e} R_{t-1}^e + \tfrac{\delta W_{t-1}'}{W_t^e} R_{t-1}'^k \tag{24}$$

$$W_t' \geq 0,\ M_t' \geq 0,\ R_t'^k \in (0,1),\ R_t'^c \in (0,1) \tag{25}$$

W_t does not appear in this formulation, as it has been broken down into wells that existed upon entering the period W_t^e and new wells built that period W_t', which is a choice variable as before. Furthermore, now instead of simply choosing how leaky those new wells will be, the firm chooses the emission rate of the equipment at those new wells $R_t'^k$, the leakage rate for completions $R_t'^c$ (i.e. what percent of the wells' first year of production will be allowed to escape during the completion process), and how much effort to devote toward maintenance to repair leaks M_t. Revenues are separated into those generated from existing wells and those generated by new wells. Revenues from existing wells depend on the baseline emissions rate from those wells and a factor $f(M_t)$, which is decreasing and convex in maintenance.[66] Revenues from new wells depend on $R_t'^k$ and $R_t'^c$. Costs resulting from decisions that affect emissions rates are again broken down into equipment, completion, and maintenance categories. The equations of motion for W_t^e and R_t^e are adjusted slightly to reflect the separation of existing wells from new wells.

First order conditions for the decision variables, starting from any time period τ,

[66] $f(\cdot)$ starts at some factor greater than one for zero maintenance effort, as emissions will be greater than the previous years rate due to equipment degradation if no maintenance is performed. $f(\cdot)$ then approaches 1 asymptotically as maintenance effort increases toward infinity. This formulation assumes that firms engage in leak detection and repair on at least an annual basis, such that leaks do not persist across years.

now simplify to the following:

$$-\frac{\partial c_c}{\partial R'^c_\tau} = P_\tau \tag{26}$$

$$-\frac{\partial c_m}{\partial M_\tau} = P_\tau R^e_\tau \tfrac{\partial f}{\partial M_\tau} \tag{27}$$

$$-\frac{\partial c_2}{\partial R'^k_\tau} = \sum_{\iota=0}^{T} \beta^\iota \delta^\iota \mathbb{E}[P_{\tau+\iota}] \prod_{j=0}^{\iota} f(M_{\tau+j}) \tag{28}$$

The first notable observation is that the firm's decision rule for emissions from completion in any given period depends only on the price in the current period. This is likely to be a contributing factor to the empirical result in Section 1.5.2 that emissions from completion are most responsive to current prices. Intuitively, the firm's decision rule for maintenance effort is a function of the price in the current period, the baseline emission rate of wells that exist going into the current period, and the sensitivity of those emission rates to maintenance effort. The first order condition for emissions related to equipment purchase decisions includes the present discounted value of gas as in the previous section, but now also depends on future maintenance decisions. Because the firm's decisions on maintenance effort and emissions for new wells are intertwined in this framework, they are not directly estimable using a reduced-form approach.

A.1.2 Negative Binomial Model

To address the possibility that results could be driven by improper specification in using OLS to estimate the effect of price on CH_4 emission rates, I additionally estimate this relationship using a negative binomial framework. In general, the assumptions of OLS are not precisely satisfied in models that have a rate or proportion as the dependent variable.[67] In this case, models that explicitly treat the dependent variable as a count are likely to provide a better fit. Two of the most frequently used count data models are Poisson regression and negative binomial regression, which is a generalization of Poisson that allows the dependent variable's variance parameter to differ from the mean.[68]

Rather than using emission rates as the dependent variable, this specification uses emissions in levels as the dependent variable and treats the quantity of gas extracted as an exposure variable.[69] This empirical framework operates under the analogy that each unit of extraction may either be leaked or contained, and the coefficient of interest recovers the effect of price on the probability that each unit will be leaked. The probability that in total e units of gas are leaked is then given by:

$$\Pr(E_{it} = e|\mu) = \frac{e^{-\mu}\mu^{e}}{e!}$$

Where E_{it} is the level of emissions at facility i in year t and e is drawn from a negative binomial distribution with parameter μ that takes the form:

$$\mu = exp(\beta_0 + \beta_1 P_{it}^{A} + \beta_2 P_{it}^{B} + \mathbf{X}_{it}\psi + \gamma_i + \lambda_{rt} + \varepsilon_{it}) \tag{29}$$

[67] For example, it is impossible for errors to be normally distributed if the outcome variable is lower-bounded at zero, and rate variables with a large proportion of data clustered near 0 or 1 are less likely to have approximately normally-distributed errors.

[68] Negative binomial models provide a better fit than Poisson when the conditional variance of the dependent variable exceeds it's conditional mean (Greene, 2008). That is the case here, as the conditional mean of CH_4 emissions is about 253 MMcf and the conditional variance is about 89,000 MMcf.

[69] *i.e.* Rate = Count/Exposure

Figure A1: Robustness check employing a negative binomial model instead of OLS. The left panel shows a second-order fractional polynomial fit, and the right panel shows a comparison with higher- and lower-order specifications. For reference, exp(13) = 442,413 Mcf and exp(10.5) = 36,315 Mcf, indicating that emissions are predicted to decrease by about one order of magnitude as prices increase from the lowest to the highest observed in the sample.

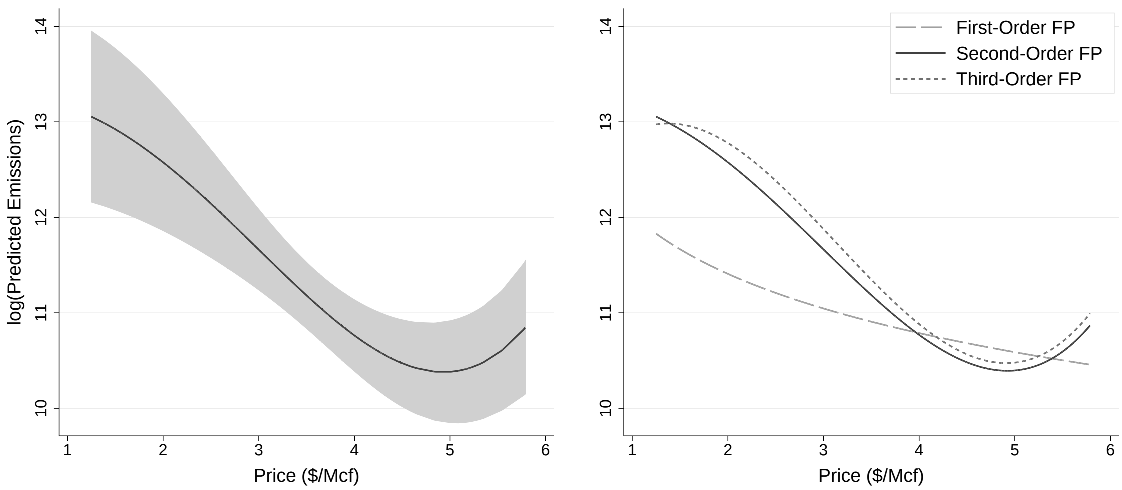

The fractional polynomial methodology is carried through to Equation 29, although it is not possible to include regression weights in the negative binomial models. Controls and fixed effects are also consistent with the specification in the main text, with the exception that extraction Q_{it} is used to determine exposure. Coefficients and model parameters are estimated using maximum likelihood. The best model fit for the second-order fractional polynomial is shown in the left panel of Figure A1. This curve is broadly similar to the result from the OLS second-order fractional polynomial specification and especially similar to the robustness check that omits regression weights (see Figure A4). Coefficient estimates for this specification and for first- and third-order fractional polynomials are reported in Table A1. Although the negative binomial framework may be a more appropriate specification along some dimensions, the OLS framework used in the main text is more directly useful for constructing a sector-wide marginal abatement cost curve.

Table A1: Results from a robustness check using a negative binomial model in place of OLS, including a linear specification (1), and first-, second-, and third-order fractional polynomial fits (2-3).

	(1)	(2)	**(3)**	(4)
Model	Linear	1st-Order FP	**2nd-Order FP**	3rd-Order FP
P_{it}	-0.1905 (0.1799)			
$\log(P_{it})$		-0.8953 (0.5620)		
P_{it}^3			-0.1407*** (0.0308)	
$\log(P_{it}) \times P_{it}^3$			0.0730*** (0.0160)	
P_{it}^2				0.9126*** (0.3272)
$\log(P_{it}) \times P_{it}^2$				-1.3507*** (0.3920)
$\log(P_{it})^2 \times P_{it}^2$				0.4637*** (0.1239)
Facility FE	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes
N	1,114	1,114	1,114	1,114

Standard errors in parentheses (clustered at the parent firm level with 146 firms)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.1.3 Instrumental Variables Regression

To address potential endogeneity from reverse causality or omitted variables that vary over time within regions, I explore instrumenting for price using exogenous weather shocks. In particular, I use average annual heating degree days (HDD), a temperature measure corresponding to degrees below 65 Fahrenheit.[70] To satisfy the exclusion restriction, temperature must not be correlated with emission rates except through its impact on gas prices. Because it is difficult to ensure this holds for a facility's own temperature, I additionally employ a strategy using weather in areas surrounding a facility conditional on weather at that facility, which directly satisfies exclusion restriction (Davis & Muehlegger, 2010; Hausman & Kellogg, 2015).[71] The intuition underlying the second approach is that temperature in surrounding areas will affect demand, but it will be entirely exogenous to production activities in that area conditional on temperature in that area. In both approaches, I also include one-year lagged temperature, as storage volumes from the previous year may also impact demand.

In the first stage, I regress price on a measure of temperature (either own temperature or temperature in nearby areas) and one lag, including the same fixed effects and controls as before. Results are presented in Table A2. Weighting observations based on facilities' mean gas production (as in the main text), the instrument relevance condition fails in the first stage (Columns 1 and 2). Omitting regression weights, a strong relationship of the expected sign is detected in the first stage (Columns 3 and 5). The second stage is estimated as a second-order fractional polynomial, as in the main text. As shown in Figure A2, the relationship between emission rates and prices appears similar to the non-instrumented relationship. However, the second stage results are not statistically significant (Columns 4 and 6).[72]

[70] HDD is recognized throughout the natural gas industry to be a very strong predictor of demand.

[71] I assign facility i's own temperature ($\mathrm{HDD}_{i,t}$) by taking an average of temperature at the hubs used to construct the price for i and create a variable for temperature in areas around facility i ($\mathrm{HDD}_{-i,t}$) by taking an average of temperature at hubs immediately adjacent to i's hubs.

[72] The two transformations of price that provide the best fit in the second-order fractional polynomial

Figure A2: Estimated relationship between emission rates and prices using weather at a facility as an instrument for price at that facility (left) and using weather in regions neighboring a given facility as an instrument for price at that facility (right).

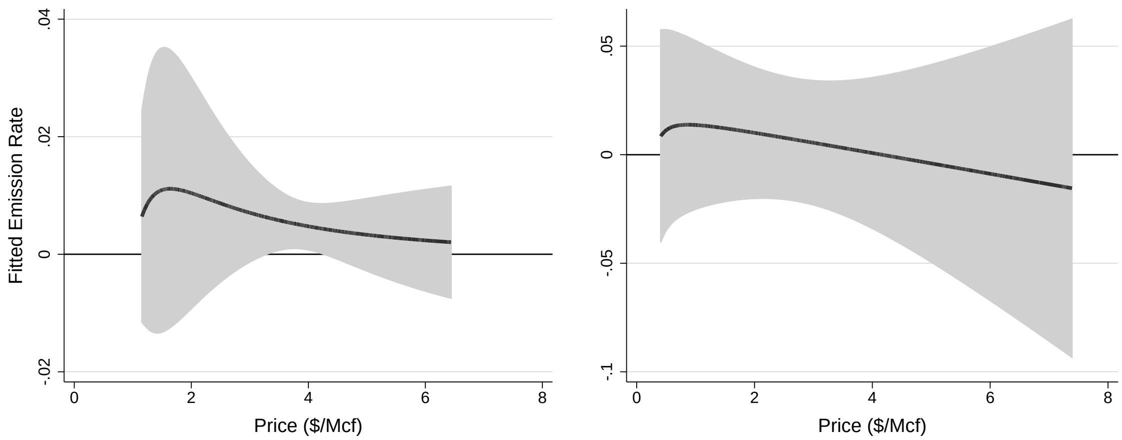

model are the same in both models: $\widehat{P}_{it}^{-2}$ and $log(\widehat{P}_{it}^{-2}) \times \widehat{P}_{it}^{-2}$.

Table A2: Results from robustness checks using weather variables as instruments for price.

	1st Stage with Weights		2SLS Own Weather		2SLS Nearby Weather	
	(1)	(2)	(3)	(4)	(5)	(6)
	P_{it}	$P_{i,t}$	P_{it}	R_{it}	P_{it}	R_{it}
$\text{HDD}_{i,t}$	0.0232*	0.00379	0.0772***		-0.0153	-0.000349
	(0.0135)	(0.0283)	(0.0138)		(0.0138)	(0.000929)
$\text{HDD}_{i,t-1}$	0.0284		0.0429***			
	(0.0191)		(0.00997)			
$\text{HDD}_{-i,t}$		0.0351			0.160***	
		(0.0688)			(0.0240)	
$\text{HDD}_{-i,t-1}$		0.0261			0.0959***	
		(0.0367)			(0.0115)	
$\widehat{P}_{it}^{-2}$				-0.00107		-0.00396
				(0.0103)		(0.00665)
$log(\widehat{P}_{it}^{-2}) \times \widehat{P}_{it}^{-2}$				0.0653		0.0204
				(0.102)		(0.0240)
Weights	Yes	Yes				
Facility FE	Yes	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	1150	1150	1150	1150	1150	1150

Standard errors in parentheses (clustered at the parent firm level with 146 firms)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.1.4 Additional Robustness Checks

Table A3: Results from robustness checks excluding weights (1), using a 1% threshold for Winsorizing facility emission rates (2), using basin-by-year fixed effects in place of region-by-year fixed effects (3), using year fixed effects (4), excluding the Mountain region (5), and excluding 2016 (6). All specifications are second-order fractional polynomials.

	(1)	(2)	(3)	(4)	(5)	(6)
Model	Unweighted Regression	Trimming Leaks at 1%	Basin-Year FE	Year FE	Excluding Mountain	Excluding 2016
P_{it}^2	-0.00291**	-0.0023*				-0.0021
	(0.00118)	(0.0014)				(0.0013)
P_{it}^3	0.0004**	0.0003		0.00006		0.0003
	(0.0002)	(0.0002)		(0.00011)		(0.0002)
$P_{it}^3 \times \log(P_{it})$				-0.00003		
				(0.00005)		
P_{it}^{-1}			0.0536***			
			(0.0169)			
P_{it}^{-2}			-0.0386***		0.0104*	
			(0.0133)		(0.0058)	
$P_{it}^{-2} \times \log(P_{it})$					0.1047	
					(0.0642)	
Facility FE	Yes	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes			Yes	Yes
N	1,150	1,246	1,126	1,156	872	1,036
adj. R^2	0.662	0.340	0.629	0.635	0.278	

Standard errors in parentheses (clustered at the parent firm level with 146 firms)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A4: Robustness check for mechanisms by which firms' emitting behavior responds to price using the full GHGRP sample. Note that using the unrestricted GHGRP sample requires omitting the production control variables retrieved from DrillingInfo.

	(1)	(2)	(3)	(4)	(5)	(6)
	Low-Bleed Pneumatic Devices	High-Bleed Pneumatic Devices	Intermittent Pneumatic Devices	Pneumatic Pumps	Venting Days	Gas Recovered For Sales
P_{it}	-32.38	-9.24	-134.4	-91.94**	-7.816	31,202,000
	(119.6)	(15.99)	(213.5)	(43.82)	(6.083)	(34,662,000)
Colorado_{2014+}	658.3	-205.3	128.2	-142.9	-3.513	9,239,000
	(1183.3)	(136.2)	(563.8)	(96.46)	(3.306)	(10,344,000)
Facility FE	Yes	Yes	Yes	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	2,593	2,593	2,593	1,855	2,017	2,017

Standard errors in parentheses (clustered at the parent firm level with 146 firms)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A5: Robustness check lower-bounding facilities potential emission rates at 0.1 percent instead of at .0223 percent.

Methane Tax	Equiv. Carbon Price	Total Abatement	Total Abatement	Total Cost	Value of Captured Gas	Net Cost
(\$/Mcf)	(\$/tCO$_2$e)	(Mcf)	(Percent)	(\$)	(\$)	(\$/Mcf)
2.79	5.00	65,876,000	45.1	271,485,000	212,562,000	0.0021
		(28,279,000)	(19.3)	(117,476,000)	(91,049,000)	(0.0009)
11.18	20.00	85,582,000	58.5	438,883,000	275,874,000	0.0057
		(40,085,000)	(27.4)	(226,449,000)	(128,785,000)	(0.0035)
27.37	48.97	42,263,000	61.6	530,323,000	290,292,000	0.0084
		(15,388,000)	(30.4)	(321,038,000)	(142,621,000)	(0.0064)
N		1,150	1,150	1,150	1,150	1,150

Variables in Mcf and \$ rounded to nearest 1,000
Bootstrapped standard errors in parentheses

Table A6: Robustness check starting facilities at 2016 values for prices and emission rates rather than average values over the study period.

Methane Tax ($/Mcf)	Equiv. Carbon Price ($/$tCO_2$e)	Total Abatement (Mcf)	Total Abatement (Percent)	Total Cost ($)	Value of Captured Gas ($)	Net Cost ($/Mcf)
2.79	5.00	61,255,000 (25,852,000)	53.7 (22.7)	189,883,000 (82,525,000)	141,938,000 (59,761,000)	0.0017 (0.0008)
11.18	20.00	74,707,000 (35,877,000)	65.4 (31.4)	293,188,000 (166,497,000)	172,902,000 (82,712,000)	0.0042 (0.0030)
27.37	48.97	77,262,000 (43,179,000)	67.7 (34.5)	341,345,000 (239,652,000)	178,772,000 (90,665,000)	0.0057 (0.0054)
N		1,150	1,150	1,150	1,150	1,150

Variables in Mcf and $ rounded to nearest 1,000
Bootstrapped standard errors in parentheses

Table A7: Robustness check using the estimated curve from the first-order fractional polynomial ($R_{it} = \beta_0 + \beta_1 log(P_{it})$) in the simulation model.

Methane Tax ($/Mcf)	Equiv. Carbon Price ($/$tCO_2$e)	Total Abatement (Mcf)	Total Abatement (Percent)	Total Cost ($)	Value of Captured Gas ($)	Net Cost ($/Mcf)
2.79	5.00	75,622,000 (27,534,000)	51.7 (18.8)	317,725,000 (114,115,000)	243,513,000 (88,463,000)	0.0026 (0.0009)
11.18	20.00	109,060,000 (40,943,000)	74.6 (28.0)	609,751,000 (245,542,000)	351,036,000 (157,731,000)	0.0091 (0.0042)
27.37	48.97	119,595,000 (49,452,000)	81.8 (33.8)	827,475,000 (455,370,000)	384,703,000 (157,731,000)	0.0155 (0.0108)
N		1,150	1,150	1,150	1,150	1,150

Variables in Mcf and $ rounded to nearest 1,000
Bootstrapped standard errors in parentheses

Table A8: GHGRP variables used to construct emission rates from equipment, completions, and maintenance. Each emissions source is a publically-available variable reported by the GHGRP. The third column consists of components of the equations used to calculate the estimated emissions from various sources (these are not publicly available) that are specifically related to firm decisions about emissions from the various category types. Equation components that are unrelated to firm decisions (such as population emissions factors) are not shown.

Category	Emissions Source	Relevant Decision Component(s)
Equipment	Pneumatic Controllers	Type (High/Low/Intermittent- Bleed)
	Pneumatic Pumps	Number of Devices
	Storage Tanks	Whether has Vapor Recovery
	Associated Gas Venting/Flaring	Whether to Vent, Flare, or Sell
	Centrifugal Compressors	Emissions from Wet Seal Degassing
	Dehydrator Vents	Absorbent Type
		Pump Type
		Use of Stripping Gas
		Use of Flash Tank Separator
		Dimensions of Dehydrator Vessel
		Whether has Vapor Recovery
Completions	Well Testing	Whether Gas is Vented or Flared
	Completion/Workover Venting	Time Gas is Vented
		Whether used Separator
	Liquid Unloading	Time Venting Each Well
		Flow Rate
Maintenance	Centrifugal Compressors	Emissions from Wet Seal Degassing
	Storage Tanks	Direct Emissions Measurement
		Time Dump Valve Not Closed
	Equipment Leak Surveys	Number & Type of Leaking Devices
		Time Assumed to be Leaking

Figure A3: Distribution of the ratio of the production variable from the DrillingInfo dataset to the same variable from the GHGRP for the years 2015 and 2016. Deviations in excess of 25 percent are trimmed from the sample.

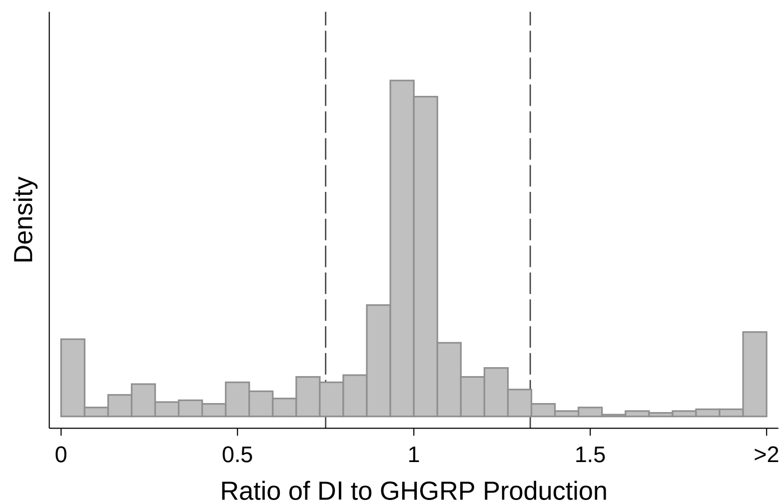

Figure A4: Robustness checks for the second-order fractional polynomial regression of emission rates on price.

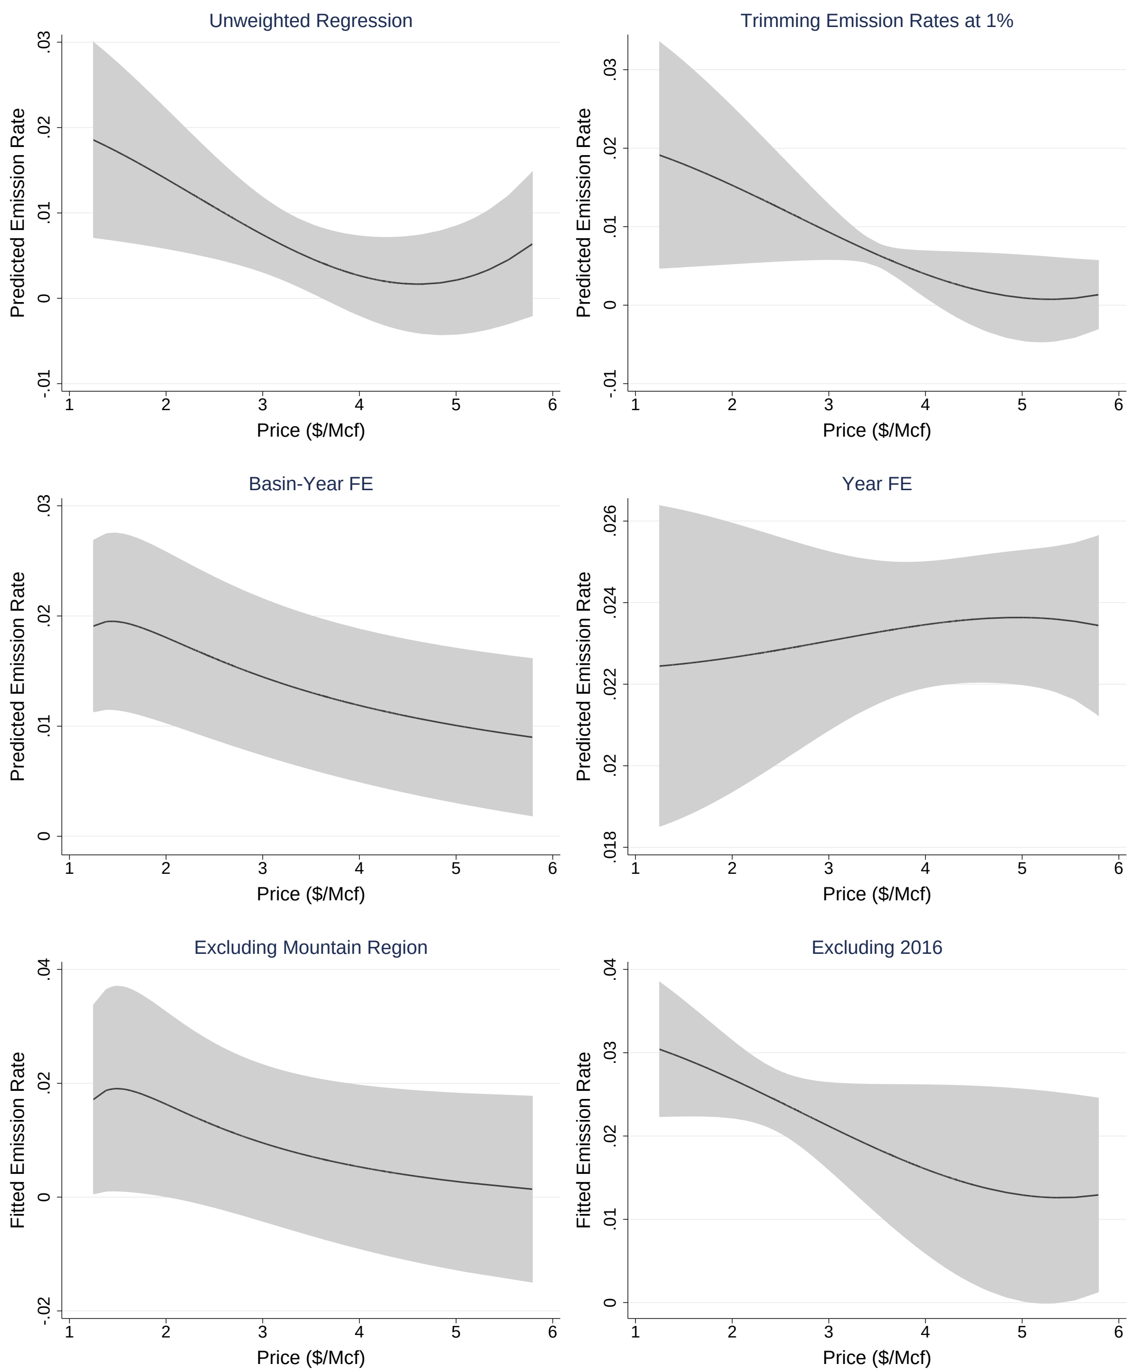

Appendix for Chapter 2

A.2.1 Derivation of FOC for e_i

$$
\begin{aligned}
\hat{E} &= N\tfrac{1}{M}\sum_{j=1}^{M} e_j \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - T\mathbb{E}[\hat{E}] \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - \tfrac{TN}{M}\mathbb{E}[\sum_{j=1}^{M} e_j] \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - \tfrac{TN}{M}\mathbb{E}[\sum_{i=1}^{N} S_i e_i] \qquad \big(S_i \equiv \mathbb{1}(\text{i sampled})\big) \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - \tfrac{TN}{M}\sum_{i=1}^{N}\mathbb{E}[S_i]e_i \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - \tfrac{TN}{M}\sum_{i=1}^{N}\tfrac{M}{N}e_i \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - T\sum_{i=1}^{N} e_i \\
\mathbb{E}[\pi] &= \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i) - Te_i] \\
\frac{\partial \mathbb{E}[\pi]}{\partial e_i} &= 0 \quad \longrightarrow \quad \frac{-\partial c}{\partial e_i} = P + T
\end{aligned}
$$

A.2.2 Derivation of FOC for M

$$
\begin{aligned}
W &= \max_M \sum_{i=1}^{N} [P(q_i - e_i) - c(e_i, q_i, \theta_i)] - TE - c_m M - \alpha T|\hat{E} - E| \\
\mathbb{E}[W] &= \max_M \sum_{i=1}^{N} [\ ...\ ] - TE - c_m M - \alpha T \mathbb{E}[|\hat{E} - E|] \\
\mathbb{E}[W] &= \max_M \sum_{i=1}^{N} [\ ...\ ] - TE - c_m M - \alpha T \mathbb{E}[|\tfrac{N}{M} \sum_{j=1}^{M} e_j - \tfrac{N}{N} \sum_{i=1}^{N} e_i|] \\
\mathbb{E}[W] &= \max_M \sum_{i=1}^{N} [\ ...\ ] - TE - c_m M - \alpha T N \mathbb{E}[|\tfrac{1}{M} \sum_{j=1}^{M} e_j - \tfrac{1}{N} \sum_{i=1}^{N} e_i|] \\
\mathbb{E}[W] &= \max_M \sum_{i=1}^{N} [\ ...\ ] - TE - c_m M - \alpha T N \mathbb{E}[|\bar{e}_m - \mu_e|] \\
\bar{e}_m &\sim N(\mu_e, \tfrac{\sigma^2}{M}) \\
\bar{e}_m - \mu_e &\sim N(0, \tfrac{\sigma^2}{M}) \\
|\bar{e}_m - \mu_e| &\sim \text{“Folded Normal”}(0, \tfrac{\sigma^2}{M}) \\
\mathbb{E}[|\bar{e}_m - \mu_e|] &= \sigma \sqrt{\tfrac{2}{\pi M}} \\
\mathbb{E}[W] &= \max_M \sum_{i=1}^{N} [\ ...\ ] - TE - c_m M - \alpha T N \sigma \sqrt{\tfrac{2}{\pi M}} \\
\frac{\partial W}{\partial M} &= 0 \quad \longrightarrow \quad M^* = \left(\frac{\alpha T N \sigma}{C_m \sqrt{2\pi}} \right)^{\frac{2}{3}}
\end{aligned}
$$

A.2.2.1 Derivation of σ^2

$$\begin{aligned} e_i^m &\equiv \text{measured emissions} \\ e_i &\equiv \text{actual emissions} \\ &\sim (\mu = \mu_e, SD = \sigma) \\ u_i &\equiv \text{measurement error} \\ &\sim (\mu = 0, SD = \sigma_m) \\ u_i &\perp\!\!\!\perp e_i \end{aligned}$$

$$e_i^m = e_i + e_i u_i$$

$$\begin{aligned} Var(e_i u_i) &= \mathbb{E}[(e_i u_i)^2] - (\mathbb{E}[e_i u_i])^2 \\ &= \mathbb{E}[e_i^2]\mathbb{E}[u_i^i] - (\mathbb{E}[e_i]\mathbb{E}[u_i])^2 \\ &= (Var(e_i) + \mathbb{E}[e_i]^2)(Var(u_i) + \mathbb{E}[u_i]^2) \\ &= \sigma_m^2(\sigma_e^2 + \mu_e^2) \end{aligned}$$

$$\begin{aligned} Var(e_i^m) &= Var(e_i) + Var(e_i u_i) \\ \sigma^2 &= \sigma_e^2 + \sigma_m^2(\sigma_e^2 + \mu_e^2) \end{aligned}$$

A.2.3 Derivation of FOC for e_i in lower-bound extension

$$\hat{E_L} = N\Big(\tfrac{1}{M}\sum_{j=1}^{M} e_j - Z\frac{S}{\sqrt{M}}\Big)$$

$$\hat{E_L} = N\Big(\tfrac{1}{M}\sum_{j=1}^{M} e_j - \frac{Z}{\sqrt{M}}\sqrt{\frac{\sum_{j=1}^{M}(e_j - \frac{1}{M}\sum_{j=1}^{M} e_j)^2}{M-1}}\;\Big)$$

$$\mathbb{E}[\pi] = \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i)] - T\mathbb{E}[\hat{E_L}]$$

$$\mathbb{E}[\pi] = \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i - e_i) - c(e_i, q_i, \theta_i) - Te_i] + \tfrac{TNZ}{\sqrt{M-1}\sqrt{M}}\mathbb{E}\Big[\sqrt{\sum\nolimits_{j=1}^{M}(e_j - \tfrac{1}{M}\sum\nolimits_{j=1}^{M} e_j)^2}\Big]$$

$$\mathbb{E}\Big[\sqrt{\sum\nolimits_{j=1}^{M}(e_j - \tfrac{1}{M}\sum\nolimits_{j=1}^{M} e_j)^2}\Big] = \mathbb{E}\Big[\sqrt{\sum\nolimits_{i=1}^{N}(S_ie_i - \tfrac{1}{M}\sum\nolimits_{i=1}^{N} S_ie_i)^2}\;\Big]$$

$$= \mathbb{E}[g(S_1, ..., S_N)]$$

Law of the Unconscious Statistician: $\mathbb{E}[g(XY)] = \sum_x\sum_y f_{XY}(x,y)g(x,y)$

$$\mathbb{E}[g(S_1, ..., S_N)] = \tfrac{1}{\binom{N}{M}}\sum_{s_1}\cdots\sum_{s_N}\underbrace{\sqrt{\sum\nolimits_{i=1}^{N}(S_ie_i - \tfrac{1}{M}\sum\nolimits_{i=1}^{N} S_ie_i)^2}}_{\binom{N}{M}\text{ possible combinations}}$$

$$\begin{aligned}
&= \tfrac{M!(N-M)!}{N!}\Bigg(\sqrt{\big(e_1 - \tfrac{1}{M}(e_1 + ... + e_M)\big)^2 + ... + \big(e_M - \tfrac{1}{M}(e_1 + ... + e_M)\big)^2}\\
&\qquad + \sqrt{\big(e_1 - \tfrac{1}{M}(e_1 + ... + e_{M+1})\big)^2 + ... + \big(e_{M+1} - \tfrac{1}{M}(e_1 + ... + e_{M+1})\big)^2}\\
&\qquad + \quad ...\\
&\qquad + \sqrt{\big(e_{N-M} - \tfrac{1}{M}(e_{N-M} + ... + e_N)\big)^2 + ... + \big(e_N - \tfrac{1}{M}(e_{N-M} + ... + e_N)\big)^2}\Bigg)
\end{aligned}$$

Assume $e_k = e,\ \forall\, k \neq i$

$$\begin{aligned}
&= \tfrac{M!(N-M)!}{N!}\Bigg(\sqrt{\big(e_i - \tfrac{1}{M}(e_i + ... + e)\big)^2 + ... + \big(e - \tfrac{1}{M}(e_i + ... + e)\big)^2}\\
&\qquad + \sqrt{\big(e_i - \tfrac{1}{M}(e_i + ... + e)\big)^2 + ... + \big(e - \tfrac{1}{M}(e_i + ... + e)\big)^2}\\
&\qquad + \quad ...\\
&\qquad + \sqrt{\big(e - \tfrac{1}{M}(e + ... + e)\big)^2 + ... + \big(e - \tfrac{1}{M}(e + ... + e)\big)^2}\Bigg)
\end{aligned}$$

$$= \tfrac{M!(N-M)!}{N!}\Bigg(\sqrt{\big(e_i - \tfrac{1}{M}(e_i+(M-1)e)\big)^2 + (M-1)\big(e - \tfrac{1}{M}(e_i+(M-1)e)\big)^2}$$

$$+ \sqrt{\big(e_i - \tfrac{1}{M}(e_i+(M-1)e)\big)^2 + (M-1)\big(e - \tfrac{1}{M}(e_i+(M-1)e)\big)^2}$$

$$+ \quad ...$$

$(Other\ rows\ don't\ contain\ e_i)$

$$= \tfrac{M!(N-M)!}{N!}\Bigg(\binom{N-1}{M-1}\sqrt{\big(e_i - \tfrac{1}{M}(e_i+(M-1)e)\big)^2 + (M-1)\big(e - \tfrac{1}{M}(e_i+(M-1)e)\big)^2}$$

$$+ \quad ...$$

$$= \tfrac{M}{N}\sqrt{\big(e_i - \tfrac{1}{M}(e_i+(M-1)e)\big)^2 + (M-1)\big(e - \tfrac{1}{M}(e_i+(M-1)e)\big)^2}$$

$$+ \quad ...$$

$$\frac{\partial \mathbb{E}[g(S_1,...,S_N)]}{\partial e_i} = \frac{(M-1)(e_i-e)}{N\sqrt{\Big(e_i-\frac{1}{M}(e_i+(M-1)e)\Big)^2+(M-1)\Big(e-\frac{1}{M}(e_i+(M-1)e)\Big)^2}}$$

$$\mathbb{E}[\pi] = \max_{\mathbf{e}_i}\sum_{i=1}^{N}[P(q_i-e_i) - c(e_i,q_i,\theta_i) - Te_i] + \tfrac{TNZ}{\sqrt{M-1}\sqrt{M}}\mathbb{E}\Big[\sqrt{\sum_{j=1}^{M}(e_j - \tfrac{1}{M}\sum_{j=1}^{M} e_j)^2}\Big]$$

$$\frac{\partial \mathbb{E}[\pi]}{\partial e_i} = 0 \ \rightarrow\ \tfrac{-\partial c}{\partial e_i} = P + T - \frac{TZ(M-1)(e_i-e)}{\sqrt{M-1}\sqrt{M})\sqrt{\Big(e_i-\frac{1}{M}(e_i+(M-1)e)\Big)^2+(M-1)\Big(e-\frac{1}{M}(e_i+(M-1)e)\Big)^2}}$$

$$\frac{-\partial c}{\partial e_i} = P + T - TZ\frac{(M-1)(e_i-e)}{\sqrt{M-1}\sqrt{M}\sqrt{\sum_{j=1}^{M}(e_j-\frac{1}{M}\sum_{j=1}^{M}e_j)^2}}$$

$$\frac{-\partial c}{\partial e_i} = P + T - TZ\frac{e_i-e}{\sqrt{M}\sqrt{\frac{\sum_{j=1}^{M}(e_j-\frac{1}{M}\sum_{j=1}^{M}e_j)^2}{M-1}}}$$

$$\frac{-\partial c}{\partial e_i} = P + T - TZ\tfrac{e_i-e}{S\sqrt{M}}$$

Appendix for Chapter 3

A.3.1 Imbalance penalties

In order to promote the efficient utilization of pipelines, FERC regulations require interstate pipeline companies to charge two types of imbalance penalties. The first is called an OFO imbalance penalty, which is assessed when shippers cause a physical imbalance on the system on a day when there is an "Operational Flow Order" (OFO) in effect. The pipeline company issues an OFO by electronically notifying all of its customers that the pipeline has reached or is about to reach its capacity constraint at at least one bottleneck on the system (usually a compression station). On a congested day, this usually happens well before the gas day starts, but it may be issued or updated during the gas day. Once this notice has been issued, any shipper that causes a physical imbalance in the system by withdrawing in excess of 2% more or less gas at their delivery node than they had injected at their receipt node is assessed a penalty equal to three times the Algonquin Citygate price times the quantity of the infraction.[74] This penalty is so severe because causing a physical imbalance on the system can be extremely harmful on days when the pipeline is near capacity constraint. If a customer withdraws more than they inject, other customers will lose service and be unable to withdraw the quantity they had scheduling using their rights. If a customer injects more than they withdraw, pressure levels could build up at a bottleneck to unsafe levels.

The second is a accounting imbalance penalty, which is assessed symmetrically for deviations in either direction between the quantity of gas the shipper had scheduled to flow through the pipeline and what they actually flowed on a monthly basis. Specifically, shippers are assessed a penalty that ranges from 1.1 to 1.5 times the Algonquin Citygate price times the quantity of the infraction on scale that increases

[74] See Spectra (2016) Section 26

with the size of the infraction.[75] This penalty is much less severe because it does not affect the safety and reliability of the the system, but still substantial—especially for LDCs, who are regulated such that they can pass the cost of gas through to their customers but cannot pass through any imbalance penalties they incur.

A.3.2 Alternative explanation for spatial heterogeneity

Alongside stronger revenue-sharing rules in Connecticut, lower demand for natural gas for electricity generation in the state may interact with capacity scheduling frictions to encourage greater withholding there. Contracts for capacity guarantee the holder's ability to transport gas to the delivery node listed in the contract. However, they may also be used to deliver gas to other nodes if capacity is available. This flexible service, termed "secondary" nominations, can be reliably used to transport gas to nodes in close proximity to the contracted delivery node (*i.e.* in the same segment) or upstream of that node. However, secondary nominations attempting to transport gas further downstream than the contracted delivery node are considerably less reliable. The capacity may be physically unavailable due to a bottleneck further downstream when the pipeline is highly congested during the winter, the pipeline company is not obligated to deliver to secondary locations and has incentive to err toward caution to reliably supply primary nominations, and uncertainty regarding approval of downstream secondary nominations persists until at least the timely cycle and potentially into the scheduling period. Therefore, the further downstream a contract's primary delivery point is, the more valuable it will be for spot market sales of bundled gas and transportation to generators.

The persistent uncertainty is highly relevant to gas-fired generators because the deadline to submit bids to the wholesale electricity market is at 10am the day before,

[75] To be specific, the penalty is $1.1 \times P^G \times Q$ for deviations between 5% and 10%, $1.2 \times P^G \times Q$ for 10-15%, $1.3 \times P^G \times Q$ for 15-20%, $1.4 \times P^G \times Q$ for 20-25%, and $1.5 \times P^G \times Q$ for deviations in excess of 25% (see Spectra (2016) Section 25)

which is three hours before the pipeline's timely cycle deadline.[76] A generator bidding based on an uncertain downstream secondary nomination therefore accepts some risk of being unable to acquire the gas needed to meet its load obligation. Dual-fired generation units can hedge this risk by burning petroleum if necessary, potentially at a net economic loss in the short term, and will weigh the likelihood of this outcome in their bidding strategies. Gas-only units have no such flexibility and face severe penalties from the ISO if they fail to meet their bidden load obligations. Contracts delivering gas upstream of the majority of gas-fired generation are thus less valuable for selling gas to generators, either directly or through the capacity release market, because of both a reduced quantity of gas that is actually delivered and uncertainty regarding when deliveries will be successful.

As shown in Table A6, there is significantly less gas-fired generation capacity connected to fifth and sixth segments of the Algonquin pipeline (where eight of the 10 suspect nodes are located) in comparison to the segments further downstream that serve Massachusetts and Rhode Island. When an LDC-affiliated shipper has excess capacity delivering gas to these segments after supplying their heating demand, if they want to sell gas to generators they will need to rely primarily on secondary nominations and accept some uncertainty of making the sale and a reduced price that incorporates generator uncertainty. However, by scheduling phantom transportation to their own primary delivery locations, they can artificially further constrict an existing bottleneck on the pipeline to reduce overall supply of gas to the region with certainty.

[76] Bids must be submitted to both the day-ahead and real-time electricity markets by 10am the day before the operating day. Bids to the real-time market may subsequently be adjusted; however, 95% of the total energy is traded in the day-ahead market and prices in the real-time market generally closely track those of the day-ahead market.

A.3.3 Simulation Model Bias Corrections

A.3.3.1 Bias from Inexact Matches

As with traditional matching estimators, some bias may be introduced due to the fact that matched observations in a pair will not be exactly identical to their counterparts on observed covariates. If matched days systematically differ from their target day counterparts in demand and/or temperature, our estimated differences in electricity market outcomes could be driven in part by differences in these variables in addition to differences in gas prices.[77] We are concerned that this type of bias may arise in our setting because within the local neighborhood of 56 first-stage match days for each target day there is likely to be a relationship between gas price and other determinants of demand. This means that when sampling three lower gas price days from the counterfactual scenario from among the 56 potential matches, we are more likely to choose days with lower electricity demand and higher temperature as well, and our final estimates would be driven in a small part by differences in these other variables.

We employ a simple regression-based bias correction procedure that adjusts gas prices for potential match days to remove variation due to temperature and electricity demand.[78] Within each neighborhood of 56 potential match days, we estimate the relationship between the gas price, electricity demand, and temperature (HDD):[79]

$$P_t^G = \alpha_0 + \alpha_1 D_t^E + \alpha_2 HDD_t + u_t$$

We then construct a bias-correct gas price for each potential match day $\widetilde{P}_t^G$ by

[77] In this section, we refer to the original day as the "target day." Each target day in the sample is matched with three "match days" that are similar in electricity demand and temperature and either a.) the gas price or b.) our estimated counterfactual gas price, depending on the scenario.

[78] While gas prices are also correlated with oil and coal prices to some extent, we do not include them in our bias correction to keep it consistent across all plants for which we employ matching.

[79] Although the relationships between these variables within the subsample 56 potential match days will be reflective of broader relationships for the entire population, estimating them within first-stage subsamples flexibly allows these relationships to be nonlinear in all first-stage variables.

taking the *predicted* price for the target day (indexed using ξ) and adding in the the residual for each match day:

$$\widetilde{P}_t^G = \widehat{\alpha}_0 + \widehat{\alpha}_1 D_\xi^E + \widehat{\alpha}_2 HDD_\xi + u_t$$

We then select the three second-stage match days using the difference between $\widetilde{P}_t^G$ and either P_t^G or $\widehat{P}_t^{G,cf}$, depending on the scenario. In this framework, $\widetilde{P}_t^G$ incorporates only variation in gas price due to weather and demand of the target day plus only variation in the gas price that is not driven by these determinants in the target day (such as variation that may be driven by capacity withholding).

This bias correction procedure improves the quality of matches on first stage variables. Before employing it, match days for our counterfactual scenario are about .5 °F warmer than target days on average and their peak demand is about 65 MW lower. By systematically matching days that are slightly warmer and have slightly lower demand than the target day, we would be overestimating the impact of withholding. After employing the bias correction, match days are only about .018 °F warmer than target days and their peak demand is only about 13 MW lower.

A.3.3.2 Bias from Out-of-Merit-Order Dispatch during Congestion

The core of our simulation model consists of reconstructing and clearing the day-ahead energy market using generators' actual day-ahead bids and day-ahead demand. This simplification of ISO-NE's actual market clearing process makes it possible to exploit the properties of nearest-neighbor matching estimators to identify how a change in the gas price would affect electricity prices, but necessarily introduces some discrepancy with real world outcomes. In particular, we note that simulated prices will likely deviate from actual day-ahead prices because we do not incorporate imports/exports, startup costs, or transmission constraints in our model. While all three introduce noise, we believe the last in particular may bias our results because transmission constraints are more likely to come into play when the set of available generators is restricted by limited gas supply. Furthermore, this bias will inherently be greater at high electricity prices. Reconstructing and clearing the market without resampling and without applying a bias-correction to demand, we find that simulated prices closely track actual prices when prices are low but are skewed downward when prices are high (see Panel A of Figure 19).

Our bias correction procedure first solves our simulation model with actual bids $\bar{b}_{it}$ backwards to calculate the implied demand necessary to rationalize the observed day-ahead electricity price for each hour in our study period $\widetilde{q}^e_h$. We then estimate the relationship between $\widetilde{q}^e_h$ and actual day-ahead demand q^e_h, and use predicted demand from that model $\widehat{q}^e_h$ in place of actual demand in all simulations in the main text of this paper. Because we are specifically interested in correcting for a nonlinear relationship, we estimate $\widehat{q}^e_h$ using a fractional polynomial specification:

$$\widetilde{q}^e_h = \theta_0 + \theta_1 (q^e_h)^\alpha + \theta_2 (q^e_h)^\beta + \nu_t \tag{1}$$

We estimate this relationship for the entire sample and then separately estimate it for winter and the rest of the year, finding that the latter model is a much closer fit, as shown in Panels C and D of Figure 19. We therefore use the relationships recovered from the seasonal model to predict $\widehat{q}^e_h$, which structurally removes the nonlinearity between actual levels of day-ahead demand and levels required to rationalize observed prices in our simulation model as shown in Panel E. As Panel B illustrates, this adjustment significantly improves the fit of our simulation model price results for prices above \$200/MWh. As a robustness check, we perform a full run of our simulation model without applying this bias correction, finding that it only slightly affects our estimate of energy costs, as shown in Panel F. We proceed to use the $\widehat{q}^e_h$ in place of q^e_h

to determine the market-clearing price in all simulations in the main text.

Figure 19: Bias correction to adjust demand for transmission constraint on high-priced days.

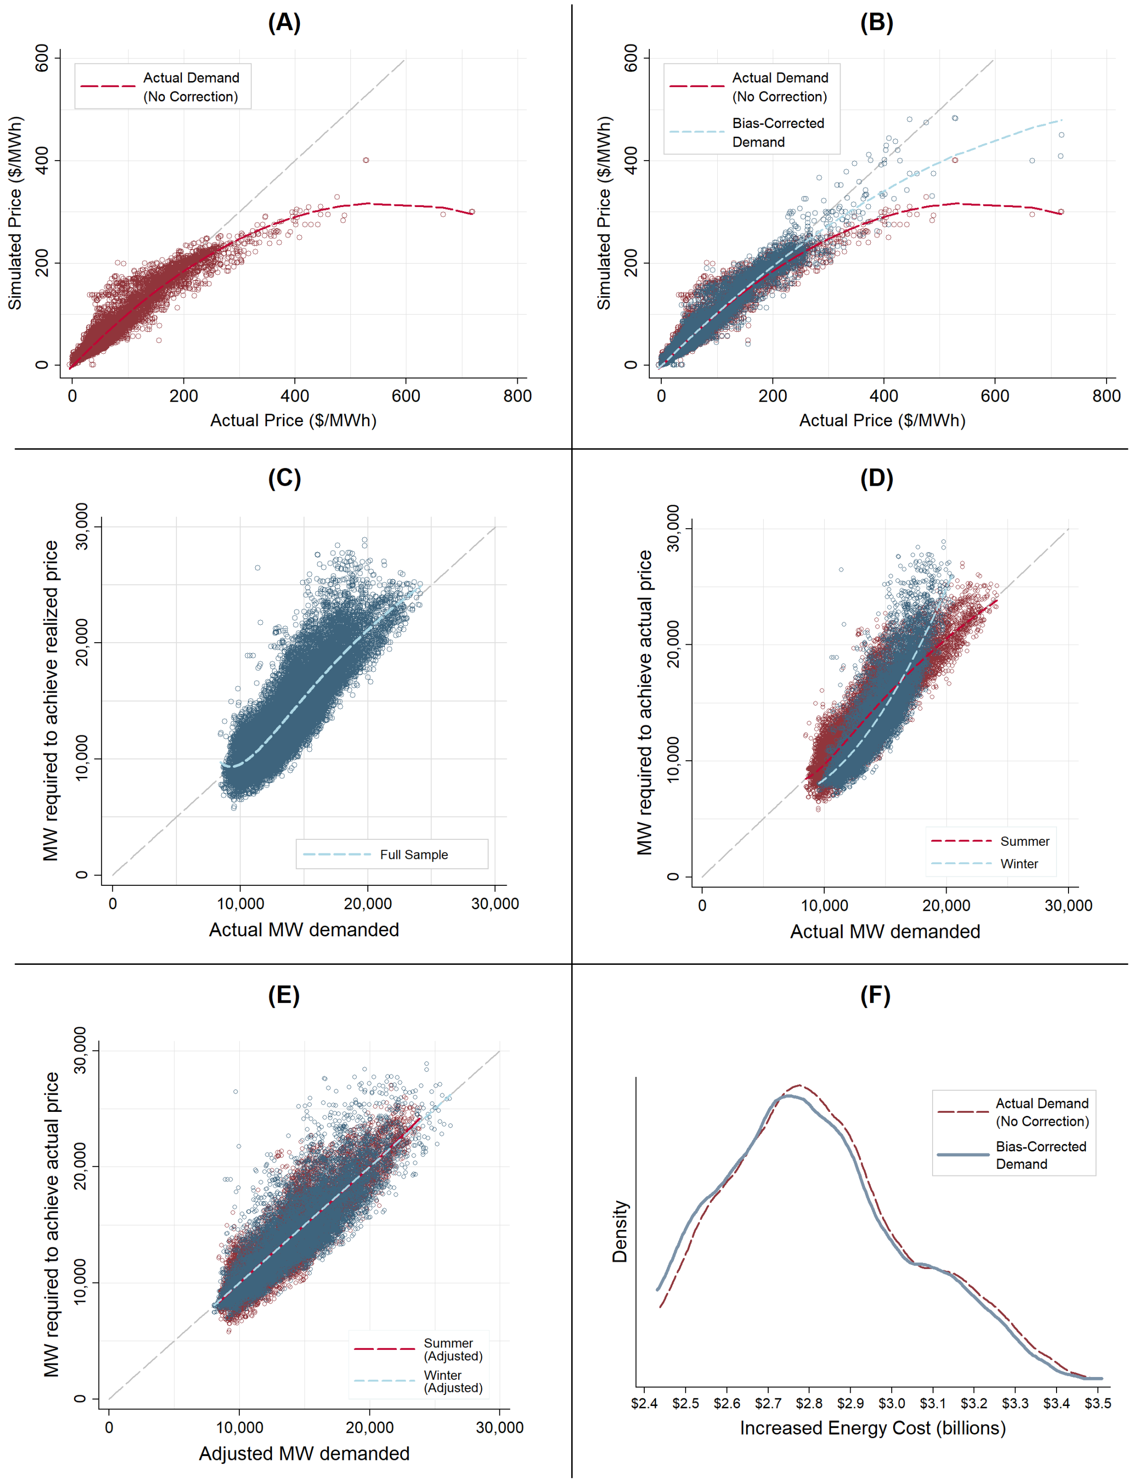

A.3.4 Robustness Checks

Table A1: Robustness check using the actual revenue sharing rule for each state. We consider the revenues retained by the LDC to focus on firm incentives (1% for CT, 10% for MA, 17% for RI) and take the inverse (1 for CT, 0.1 for MA, 0.0588 for RI) to facilitate interpretation of the interactions.

	(1)	(2)	(3)	(4)	(5)
LDC	585.0**	-23.22	315.1		57.00
	(240.6)	(85.03)	(207.9)		(69.45)
Sharing		68.66		278.0	68.66
		(97.85)		(412.6)	(97.85)
Contracts			0	-0.0360***	0
			(.)	(0.0119)	(.)
LDC×Sharing		1462.6**			255.4
		(714.1)			(520.9)
LDC×Contracts			0.0502		-0.0366***
			(0.0521)		(0.0122)
Sharing×Contracts				0.457***	0
				(0.131)	(.)
LDC×Sharing×Contracts					0.454***
					(0.134)
N	133,029	133,029	133,029	133,029	133,029

MW	1012.1**	-79.05	322.9		-72.08
	(445.2)	(97.80)	(386.9)		(65.01)
Sharing		2.759		324.5	5.041
		(21.07)		(512.1)	(34.56)
Contracts			-0.00259	-0.0365***	-0.00136
			(0.00271)	(0.0121)	(0.00182)
MW×Sharing		1804.7**			532.3
		(853.4)			(613.0)
MW×Contracts			0.156		-0.0266*
			(0.116)		(0.0143)
Sharing×Contracts				0.454***	-0.0150
				(0.134)	(0.0258)
MW×Sharing×Contracts					0.459***
					(0.135)
N	109,152	109,152	109,152	109,152	109,152

All regressions include quarter fixed effects and controls for temperature and day of week

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses (clustered at the node level)

Table A2: Robustness check using a continuous variable for MW owned by the node operator's parent firm ("MW_C"). We understand these counterintuitive results to be driven by the fact that Firm A downschedules more than Firm B, yet Firm B owns more generation capacity than Firm A, and we take this as suggestive evidence that the direct incentive pathways for merchant unregulated capacity (which Firm A owns and Firm B does not) are stronger than the indirect ones for regulated capacity.

	(1)	(2)	(3)	(4)	(5)
MW_C	-0.0107	0.126***	-0.138		-0.00800
	(0.194)	(0.0464)	(0.157)		(0.0212)
CT		2503.0*		274.4	867.0
		(1329.7)		(465.6)	(928.3)
Contracts			0.0387	0.00225	-0.00280
			(0.0545)	(0.00346)	(0.00262)
$MW_C \times$CT		-1.695			-0.634
		(1.082)			(0.737)
$MW_C \times$Contracts			0.0000365		0.0000182***
			(0.0000345)		(0.00000482)
CT$\times$Contracts				0.417***	0.612***
				(0.122)	(0.0751)
$MW_C \times$CT$\times$Contracts					-0.000384***
					(0.0000830)
N	109,152	109,152	109,152	109,152	109,152

All regressions include quarter fixed effects and controls for temperature and day of week

Standard errors in parentheses (clustered at the node level)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3: Our results are robust to using MW of merchant unregulated generation owned by the node operator's parent firm ("Merchant") as the dependent variable of interest. Coefficient estimates are characteristically similar in sign and significance but smaller in magnitude as the continuous variable for MW owned has a larger scale than a binary one. For example, the interpretation of the triple interaction in column (5) would be that for each MW of generation owned by the parent firm of a node in Connecticut, that node will downschedule an additional .004 MMBtu on average for each 1 MMBtu of NN from TE contracts they own (additional to the effects of the other six variables and interactions).

	(1)	(2)	(3)	(4)	(5)
Merchant	0.563	-0.0950***	0.0129		-0.0748***
	(0.871)	(0.0290)	(0.301)		(0.0158)
CT		253.4		274.4	-120.2**
		(188.8)		(465.6)	(56.18)
Contracts			0.00882	0.00225	0.00201
			(0.0133)	(0.00346)	(0.00354)
Merchant×CT		45.75**			26.99*
		(22.85)			(15.88)
Merchant×Contracts			0.00385		-0.0000315
			(0.00248)		(0.0000283)
CT×Contracts				0.417***	0.176***
				(0.122)	(0.0462)
Merchant×CT×Contracts					0.00398***
					(0.000952)
N	109,152	109,152	109,152	109,152	109,152

All regressions include quarter fixed effects and controls for temperature and day of week

Standard errors in parentheses (clustered at the node level)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A4: Robustness check using contracts of any type delivering gas from any point of origin ("Contracts$_A$") instead of just "no notice" contracts delivering gas from Texas Eastern. Coefficients on variables involving contracts are characteristically similar in sign and significance, but generally smaller in magnitude, suggesting that "no notice" contracts from Texas Eastern are indeed particularly useful for downscheduling.

	(1)	(2)	(3)	(4)	(5)
MW	1012.1**	101.5**	-155.6		20.78
	(445.2)	(43.46)	(480.7)		(37.57)
CT		2.959		-612.3***	-53.37***
		(20.81)		(203.4)	(19.74)
Contracts$_A$			-0.00219*	-0.00124	-0.00233*
			(0.00122)	(0.00113)	(0.00124)
MW×CT		1623.8**			-656.3**
		(768.1)			(260.6)
MW×Contracts$_A$			0.0695		0.00366
			(0.0459)		(0.00230)
CT×Contracts$_A$				0.153***	0.00229
				(0.0197)	(0.00401)
MW×CT×Contracts$_A$					0.149***
					(0.0168)
N	109,152	109,152	109,152	109,152	109,152

All regressions include quarter fixed effects and controls for temperature and day of week

Standard errors in parentheses (clustered at the node level)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A5: Robustness checks using NN from TE contracts held by Firms A and B only ("Contracts$_{EA}$") and held by other firms only ("Contracts$_{-EA}$").

	(1)	(2)	(3)	(4)	(5)
MW	1012.1**	101.5**	345.0		5.054
	(445.2)	(43.46)	(368.0)		(24.98)
CT		2.959		381.9	1.872
		(20.81)		(442.9)	(20.83)
Contracts$_{EA}$			0.0376	0.0191***	0
			(0.0908)	(0.00464)	(.)
MW×CT		1623.8**			477.4
		(768.1)			(552.1)
MW×Contracts$_{EA}$			0.133		0.0189***
			(0.139)		(0.00480)
CT×Contracts$_{EA}$				0.403***	0.0315
				(0.119)	(0.0194)
MW×CT×Contracts$_{EA}$					0.366***
					(0.111)
N	109,152	109,152	109,152	109,152	109,152

	(1)	(2)	(3)	(4)	(5)
MW	1012.1**	101.5**	1021.6**		86.49**
	(445.2)	(43.46)	(457.0)		(41.72)
CT		2.959		1484.8**	-15.74
		(20.81)		(690.3)	(12.08)
Contracts$_{-EA}$			-0.00313	-0.00471*	-0.00321
			(0.00269)	(0.00273)	(0.00269)
MW×CT		1623.8**			1634.2**
		(768.1)			(768.2)
MW×Contracts$_{-EA}$			-0.0623**		-0.00476
			(0.0306)		(0.00380)
CT×Contracts$_{-EA}$				-0.526*	-0.000313
				(0.288)	(0.00324)
MW×CT×Contracts$_{-EA}$					0
					(.)
N	109,152	109,152	109,152	109,152	109,152

All regressions include quarter fixed effects and controls for temperature and day of week

Standard errors in parentheses (clustered at the node level)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A6: There is limited generation capacity and thus limited demand from generators in the spot market for natural gas within and upstream of segments 5 and 6, where 8 of the 10 most-withholding nodes are located.

Segment	**Segment Name**	**Generation Capacity**	**Upstream Capacity**
1	Lambertville to Hanover	0	0
2	Hanover to Stony Point	0	0
3	Stony Point to Southeast	0	0
4	Southeast to Oxford	6.3	6.3
5	Oxford to Cromwell	441.6	447.9
6	Cromwell to Chaplin	761.2	1209.1
7	Chaplin to Burrillville	861.8	2070.9
8	Burrillville to J and G Systems	447.8	2548.7
9	J System (Boston)	3430.6	5979.3
10	G System (Rhode Island)	2583.4	5132.1

Source: SNL and EIA Energy Mapping System