The Costs of Inefficient Regulation: Evidence from the Bakken

Gabriel E. Lade and Ivan Rudik*

(Preliminary and Incomplete: Do not cite)

May 12, 2017

Abstract

Oil wells often produce large volumes of lighter hydrocarbons such as natural gas. In regions that are primarily valued for their oil reserves, well operators often resort to flaring these gases. In 2015, the state of North Dakota implemented a regulation requiring operators to capture a minimum fraction of all gas produced across their wells. The regulation is enforced uniformly and does not allow for inter-firm trading. In this paper, we estimate the effectiveness of this regulation and study its relative efficiency compared to a market-based approach. We find that the regulation reduced flaring rates by 2 to 7 percentage points and that firms primarily complied by connected wells to gas capture infrastructure more quickly. We then exploit information on natural gas collection infrastructure costs to construct firm-specific marginal compliance cost curves, and construct counter-factual compliance scenarios that achieve the same flaring reductions but reallocate abatement from high- to low-cost firms. We find that allowing greater flexibility in the regulation would reduce aggregate compliance costs by tens of millions of dollars.

JEL Codes: L71, Q3, Q4

Keywords: North Dakota Bakken, hydraulic fracturing, flaring, efficient regulation

*Lade: Iowa State University and the Center for Agricultural and Rural Development. Email: gelade@iastate.edu. Rudik: Iowa State University and the Center for Agricultural and Rural Development. Email: irudik@iastate.edu. We gratefully acknowledge financial support for this research from the National Bureau of Economic Research Economics of Energy Markets Grant. We thank Dave Keiser, Quinn Weninger, John Rupp, and seminar participants at Indiana University for helpful comments. Erix Ruiz-Mondaca and T.J. Rakitan provided excellent research assistance.

1 Introduction

A necessary condition for cost-effective regulation is that marginal compliance costs are equal across firms. Policies that achieve this condition include a pollution tax or capping aggregate emissions and allowing firms to trade pollution allowances. Despite their straightforward nature, real-world policy often deviates from first-best. This is particularly true in the energy sector, home to many second-best policies such as renewable energy mandates in the transportation fuel and electricity sectors (Fischer and Newell, 2008; Holland et al., 2009; Lapan and Moschini, 2012; Lemoine, 2016; Hollingsworth and Rudik, 2016). Inefficient environmental policies are often implemented because they are politically expedient, but the size of the efficiency costs of second-best policies is often unknown.

This paper studies the impacts and efficiency of a command-and-control regulation aimed at reducing natural gas flaring in North Dakota's oil-producing region: the Bakken. When firms drill for unconventional oil, their wells also produce natural gas and natural gas liquids (NGLs) like butane, methane, and ethane. In regions like the Bakken that are valued primarily for their oil production, these gases are often flared, i.e. burned at the well site (Swanson, 2014). Flaring has become an acute problem in North Dakota as infrastructure to capture and process natural gas has lagged behind the explosion of drilling activity due to the fracking revolution.

Regulators have several incentives to regulate flaring. First, flaring is associated with a number of environmental externalities. Worldwide, flaring results in 300 million tons of CO_2 emissions each year, equivalent to the emissions of 50 million cars, and contributes to local air pollution that may have substantive health impacts on oil field workers and nearby populations (Tedesco and Hiller, 2014; World Bank, 2015). Second, flaring results in economic losses to lease-owners, firms, and the government. Policymakers have enacted a range of policies to reduce the practice. The World Bank has a Zero Routine Flaring initiative that seeks to eliminate routine flaring by 2030. In the U.S., both federal and state agencies have passed or considered a number of regulations to reduce flaring. For example, the U.S. Bureau of Land Management released proposed rules to reduce flaring and venting from oil and gas production on public and Indian lands in February 2016 (Bureau of Land Management, 2016). The EPA has considered a number of rules to regulate methane emissions from oil and gas operations under the Clean Air Act, and the Fish and Wildlife Service considered regulating hydraulically fractured wells drilled on and near protected habitat areas.

In July 2014 the North Dakota Industrial Commission (NDIC) enacted Commission Order 24665. The order established some of the most aggressive flaring standards in the U.S., and its implementation is followed by other states' regulatory agencies (Storrow, 2015). The order requires each operator in the Bakken to capture a minimum percentage of gas produced by all their wells, with an ultimate objective of capturing 90% of produced gas by 2020. Unlike a first-best policy prescription, the NDIC order is firm-specific. If operators face differential compliance costs, as we argue they do, the policy may allocate abatement (gas flaring reductions) inefficiently across firms, and the same state-wide flaring targets would be achievable at a lower cost. Understanding the efficiency of the regulation is important since gas capture regulations are among to most difficult to comply with, and therefore regulatory inefficiencies may come at substantial dollar costs (Zirogiannis et al., 2016).

In this paper, we study the impacts of the NDIC order on oil operators in North Dakota. We first estimate the impact of the regulation on firms' operations, including their flaring rates, oil and gas production, well completion time, and the time it takes to connect wells to gas capture infrastructure. Second, we estimate firm-specific marginal compliance cost curves (i.e. marginal flaring abatement cost curves) that we use to construct counterfactual flaring scenarios and estimate gains from reallocating flaring from high-cost to low-cost firms. We find that, on average, the order decreased flaring at new wells by 2-7 percentage points. In addition, we find that firms subject to the regulation took longer to complete their wells and connected their wells more quickly to gas capture infrastructure. Consistent with previous literature, we do not find that firms responded to the regulation by curtailing oil or gas production (Kellogg, 2011; Anderson et al., 2016). We also document significant heterogeneity in compliance costs across firms and find that firms with historical flaring rates near the flaring restriction and those with high historical flaring rates were most responsive the regulation. Using our estimates of firms' marginal compliance costs, we find that reallocating which wells were connected to gas capture infrastructure could achieve the same flaring reduction at a cost savings of tens of millions of dollars.

Our work contributes to a growing literature studying the economic impacts of the fracking revolution. Previous work has studied the health and pollution impacts of fracking (Hill, 2012, 2015; Olmstead et al., 2013); how nearby drilling is capitalized into housing values (Gopalakrishnan and Klaiber, 2014; Muehlenbachs et al., 2015; Bartik et al., 2016); the efficiency of landowner-firm leases (Vissing, 2016); the supply elasticity of fracked versus conventional wells (Newell et al., 2016); and the economic and welfare impacts of these newly reachable resources (Hausman and Kellogg, 2015; Feyrer et al., 2017). Only recently have others begun to analyze firm decision making, specifically learning, in this setting (Covert, 2015). To date, little work has studied the effects of government regulations and policies on oil and natural gas firms' decisions. One significant contribution of our work is to take advantage of a rich dataset to develop two novel identification strategies to study the impact of policies on well operations.

The paper contributes more generally to an extensive literature studying efficient pollution regulation and the efficiency costs of second-best policies. A number of authors have studied the economic and efficiency implications of second-best policies (Baumol and Oates, 1988; Helfand, 1991). Building on the theoretical models, others have used computational and empirical approaches to estimate the efficiency gains of counterfactual first-best policies. For example, Carlson et al. (2000) find billions of dollars of gains from trade under Title IV of the Clean Air Act Amendments of 1990. Fowlie et al. (2012) estimate that the gains from equalizing nitrogen oxide marginal abatement costs across mobile and non-mobile sources would be larger than a billion dollars. Holland et al. (2015) find that renewable energy mandates for transportation fuels are two to four times more costly than a cap and trade program to achieve the same carbon emission reductions.

The paper proceeds as follows. Section 2 provides a background on oil production in the Bakken, as well as the institutional and regulatory setting in the state. We also describe the new flaring restrictions. In Section 3, we develop a model of firms' production and gas connection decisions to clarify the margins through which firms may respond to the regulation and motivate our subsequent estimates of firms' marginal flaring abatement cost curves. Section 4 describes our data and provides summary statistics, and Section 5 discusses our empirical strategy and contains our results of the effects of the regulation on firms' flaring and production decisions. In Section 6 we study heterogeneity in the impacts and costs of the regulation. We also estimate firm-specific marginal abatement cost curves and construct counterfactual flaring scenarios. Last, Section 7 discusses potential extensions on the current work and concludes.

2 Background

2.1 The Bakken Shale Formation

Unlike the conventional reservoirs from which oil has been traditionally extracted, much of North Dakota's geology is characterized by "tight" oil formations where oil is locked into the structure of shale rock. Two advances since the 1990s have drastically improved the economic viability of shale oil extraction. First, drilling operations have become more efficient at drilling wells horizontally. Since shale is found in a horizontal layer in the earth, drilling horizontally exposes the well to much more shale, and hence trapped oil, than vertically drilled wells. Second, firms have become more effective at fracturing shale rock. The fracturing process involves injecting fluids into wells at extremely high pressures to fracture the surrounding rock to release oil.

These innovations have transformed the oil and gas industry. Oil production from fracked wells now accounts for nearly half of U.S. production (Energy Information Administration, 2015), and in North Dakota production has increased from 90,000 barrels of oil per day (bpd) in 2005 to over 1.2 million bpd in 2015 (North Dakota Industrial Commission, 2016). Firms have also dramatically increased the efficiency with which they extract oil from shale formations, such that break-even oil prices have been estimated to be as low as $35/bbl (Bailey, 2015). North Dakota is likely to continue producing substantial quantities of oil into the future. The U.S. Geological Survey estimates that the Bakken and Three Forks shale formations contain roughly 7.4 billion barrels (bbls) of oil, accounting for nearly 20% of proven recoverable reserves in the United States (Gaswirth et al., 2013; Energy Information Administration, 2016a).

In addition to oil, the U.S. Geological Survey estimates that the Bakken formation contains 6.7 trillion cubic feet of associated natural gas and 530 million barrels of NGLs (Gaswirth et al., 2013). When oil is extracted from a fracked well, these gas co-products come along with it. Historically, much of this gas has been flared, at a significant loss to royalty owners and the state. The commercial value of NGLs flared by North Dakota well operators in May 2013 alone was estimated to be $3.6 million. Compared to an estimated $73 million value of oil produced in the same month, the lost private value of the gas is non-negligible (Salmon and Logan, 2013).[1]

2.2 The North Dakota Flaring Regulation

In 2014, the NDIC passed Order 24665 to reduce flaring in the state (North Dakota Industrial Commission, 2015).[2] Prior to passage of the regulation, the only major regulation that discouraged flaring was a requirement that well operators pay taxes and royalties on any gas flared after the first year of production (Energy Information Administration, 2016b). This existing requirement was not particularly burdensome since was not strictly enforced and wells produce the vast majority of their total oil and gas within the first year.

Order 24665 instituted ambitious gas capture goals for oil producers in the state. The order requires that each producer captures 77% of the gas produced by all the wells they own from January 2015 to March 2016; 80% from April 2016 through October 2018; 85%

[1]Flaring is much preferred to venting, or releasing gases directly into the atmosphere. Vented gases contain compounds like hydrogen sulfide that are hazardous to human health. Flaring converts methane and other pollutants to CO_2, which is a less potent greenhouse gas and reduces the release of other harmful by-products. Venting is prohibited in North Dakota.

[2]A task force was first organized to develop a plan to reduce flaring in North Dakota in September 2013. In March 2014 the task force released its report and the ruling was subsequently adopted.

from November 2016 through October 2018; 88% from November 2018 through October 2020; and 91% after November 2020.[3] The gas capture requirements are applied uniformly across firms.[4] Thus, the policy is akin to a firm-specific cap and trade program, where firms may over-comply at low-cost wells and flare more than required at other wells. However, operators are unable to trade flaring rights across firms.[5] Violating the standard can result in two types of penalties. Firms that are out of compliance can have wells ordered to curtail production to as low as 100 barrels per day until the firm can meet its flaring target. In addition, if a firm is out of compliance for three months or longer, it can incur civil penalties of up to $12,500 per day for each well that is below the mandated firm-wide capture rate (North Dakota Industrial Commission, 2015).

Firms must comply with the NDIC regulation every month. Each month, the NDIC calculates each firm's capture rate as:

$$(\%\ \text{Capture})_i = \frac{\sum_j \left(\mathrm{g}_{s,j} + \mathrm{g}_{u,j} + \mathrm{g}_{p,j}\right)}{\sum_j \mathrm{g}_j}.$$

where j indexes the wells owned by firm i, $\mathrm{g}_{s,j}$ are gas sales from well j, $\mathrm{g}_{u,j}$ is gas used by the firm on the well site, $\mathrm{g}_{p,j}$ is the total gas processed in an approved manner, and g_j is total gas produced by well j.[6,7] Firm's primary compliance mechanism is to connect wells to pipeline infrastructure that captures the gas coming out of the well. This involves installing small pipelines, called gathering lines, which connect the well site to larger product pipelines that transport gases to processing plants. However, connecting a well to a gathering line does not reduce all flaring on a well pad, and wells that are connected to gas capture infrastructure still routinely flare. Flaring at connected wells is typically due to insufficient capacity of downstream gathering lines or product pipelines, as well as "pigging" (clogging) of pipelines. Thus firms can also reduce flaring by changing practices on the well site. For example, a firm

[3]The original version of the regulation had a more ambitious timescale that was subsequently delayed due to industry pressure. However, the ultimate target has remained unchanged.

[4]The NDIC was cognizant of efficiency concerns to some degree. Order 24665 explicitly states that it is firm-specific instead of well-specific in order to give firms "maximum flexibility" (North Dakota Industrial Commission, 2015).

[5]This type of regulation is sometimes referred to as averaging, or taking an average effluent level across all of a firm's facilities rather than instituting a facility-specific standard. While more efficient than a facility-specific requirement, this type of enforcement may nonetheless lead to different abatement costs across firms.

[6]Gas may be used on site to power an electric generator or processed using a natural gas stripping unit that removes heavier NGLs that can then be sent by truck to a gas processing plant. Both practices are limited in the amount of gas used and/or economically recovered.

[7]The regulation allows for limited forms of banking. Firms may accumulate and bank gas capture credits for up to three months if they are above the mandated gas capture level. NDIC staff indicated that there have been only four instances where a firm has banked credits since the beginning of the flaring regulation so we ignore banking in this paper.

could temporarily curtail oil and gas production or use gas for other purposes on-site during events where gas cannot be sent through the gathering line. Alternatively, when installing gathering line infrastructure at new wells, firms could build "looping" lines to circulate and store gas if line pressure builds up.

The NDIC began enforcement in January 2015, and all active wells in the state were included in firms' gas capture calculations at that time. However, the order includes important exemptions for new wells. A new well is not included in a firm's gas capture calculation for the well's first 90 production days. As a result, firms have substantial flexibility with regards to their flaring rates at new wells until the fourth month of production.

While our paper focuses on installation of gathering lines by oil operators, meeting the NDIC's flaring reduction targets relies on a large expansion in the state's natural gas processing capacity and product pipelines. Leading up to the regulation there was a large increase in natural gas processing capacity and a roll out of larger gas product pipelines. Natural gas processing capacity in North Dakota increased from just over 200,000 mcf/day in 2006 to 1,600,000 mcf/day in 2015, with more capacity anticipated to come online in coming years (North Dakota Pipeline Authority, 2015). Despite this rapid growth, the increase in the number of wells and gains in well productivity have pushed the limits of how much natural gas can be processed.

2.3 Oil Production in the Bakken

Understanding the effects of Order 24665 on firm behavior requires an understanding of firms' decision-making and well production functions. After a suitable location has been determined and mineral rights are obtained, firms drill or "spud" a well. Most producers hire independent drilling companies to drill and frack the well. Drilling is completed in multiple stages, including: (i) drilling the vertical segment of the well; (ii) drilling one or more "laterals" or horizontal segments; and (iii) inserting and securing production casing to protect surface water and ensure the structural integrity of the well. After drilling, firms hydraulically fracture the well. Fracking involves perforating the well casing and then injecting large amounts of water, sand, and other additives at extremely high pressure to create fissures in the surrounding rock. After the well has been fractured it is "completed" and ready to produce oil and natural gas. At this stage, firms install a permanent wellhead and other onsite equipment to capture and store oil and gas produced by the well. Oil, gas, and water flow from the wellhead through several flow lines and tanks that separate the products. After separation, oil is stored in large containers until it is picked up by trucks to be delivered to the nearest oil pipeline or refinery. If the well has gas gathering infrastructure, the separated gas is transported to nearby gas plants through pipelines. If the well does not have gathering lines installed, separated gas is flared.

The amount of oil and gas that a well produces is primarily determined by two factors: (i) the amount of hydrocarbons in the underlying shale; and (ii) the length of the well and intensity with which firms frack the well. Firms can affect the former by drilling in more productive areas. However, firms are not perfectly informed, and they do not always drill into the most productive shale (Covert, 2015). The latter is determined in the fracking phase. After a well is producing, the amount of oil and gas that comes out of the well is largely determined by the pressure within it. While operators can curtail production or shut in a well, they are unable to make the well more productive unless they re-fracture the well.[8]

3 A Model of Gas Capture

We develop a simple model to understand better the economic incentives faced by firms that are subject to the gas capture requirement, as well as to identify factors that drive potential gains to using market-based mechanisms. We model a single firm facing the flaring regulation in a two-stage, static setting, and focus primarily on the second stage. In the first stage, the firm selects the number of wells to drill, J, the location of these wells, the length of the horizontal segment of the well, and how much of each input (water, sand, etc.) to use when fracking the well. Between the first and second stages, the well is fracked and completed. At the beginning of the second stage, the firm realizes its oil and gas production and decides whether to connect each well to gas capture infrastructure. At the end of the second stage, oil is sold at price P^o and gas is sold at price P^g if the well is connected to gas collection infrastructure.

We make two additional assumptions. First, the firm's connection decision is independent of its oil production, i.e., connecting a well has a negligible effect on oil-related profits.[9] This allows us to abstract from wells' oil production when considering the firm's connection decisions. Second, we assume the firm knows the total amount of gas a well will produce when it makes the connection decision. This allows us to abstract from uncertainty that generates an option value to delaying investments.[10]

[8] Kellogg (2011) and Anderson et al. (2016) study conventional oil wells in Texas and argue that well drilling is driven by market oil prices, not oil production of a given well as is commonly assumed. They show that along an equilibrium path, although firms are able to curtail production, firms always keep wells producing at their maximum possible level regardless of the prevailing oil price. This result has one caveat in unconventional oil setting: that firms may re-pressurize unconventional wells by re-fracking.

[9] Connecting a well may require production to be temporarily curtailed; however, we assume that the well will maintain its pressure and production will be shifted into the future.

[10] There is evidence that unconventional wells are lower risk than conventional wells regarding realized

We argue that neither assumption is overly restrictive in our setting. We are unaware of any literature citing production losses from installing gas capture infrastructure. In addition, after completion the oil and gas production function of a well follows a relatively stable decline curve. A common characterization is the 'Arps' model (Fetkovich, 1980):

$$\begin{aligned} \mathrm{o}_{it} &= \mathrm{O}_{i0} t^{\beta_{\mathrm{o}}} \exp(\epsilon_{it}) \\ \mathrm{g}_{it} &= \mathrm{G}_{i0} t^{\beta_{\mathrm{g}}} \exp(e_{it}) \end{aligned} \tag{1}$$

where o_{it} and g_{it} are the oil and gas production from well i in period t, O_{i0} and G_{i0} are the initial levels of oil and gas production from the well, β_{o} and β_{g} are the oil and gas decline rates, and ϵ_{it} and e_{it} are noise terms capturing potential variability in the production function .[11] Thus, to the extent that ϵ_{it} and e_{it} are small, conditional on observing O_{i0} and G_{i0}, firms have a good sense of the total oil and gas that a well will produce.

The firm's profit-maximization problem can be solved with backwards induction. Here we consider only the second stage. Wells are heterogeneous in both the amount of gas produced and in connection costs. Well j produces g_j units of gas over its lifetime, which can be calculated by taking the integral of equation (1) over t. We denote the connection costs for well j as $C_j(h_j)$, where $h_j \in \{0, 1\}$ where 1 indicates that the well is connected to a gathering line and 0 indicates that it was left unconnected. We assume that $C_j(0) = 0,\ C_j(1) > 0$.[12] Last, we assume the regulator sets a minimum fraction of gas that must be captured by the firm across all its wells, $\bar{F} \in (\alpha, 1]$ where $\alpha > 0$ is sufficiently high so that the flaring constraint binds. The firm's problem is then:

$$\max_{h_1,\ldots,h_J} \sum_{j=1}^{J} P^g\, g_j\, h_j - C_j(h_j)$$

$$\text{subject to:} \quad \frac{\sum_{j=1}^{J} g_j\, h_j}{\sum_{j=1}^{J} g_j} \geq \bar{F} \quad \text{and} \quad h_j \in \{0,1\}\ \forall j = 1, ..., J$$

Let λ denote the Lagrange multiplier on the firm's flaring constraint. The firm connects well

production. Newell et al. (2016) describe how conventional wells in Texas may have low productivity, but every so often a firm will drill a massively productive well. Unconventional wells, however, were found to be more consistent.

[11] Covert (2015) estimates a non-parametric well production function over fracking inputs that determine O_{i0} and G_{i0}.

[12] Gathering line costs vary along two important dimensions: (i) distance to the nearest product pipeline; and (ii) the diameter of the line. Estimates from a recent report by ICF international put gathering line costs at \$25,000-\$34,000 per inch-mile for two- to eight-inch gathering lines (International, 2014).

j if:

$$P^g + \lambda \geq \frac{C_j(1)}{g_j}, \quad j = 1, ..., J. \tag{2}$$

The firm connects well j if the marginal benefits of a unit of gas production from a connected well, the market price of gas and the firm's shadow price of the constraint, are greater than the per-unit costs of capturing gas (abating flaring) from that well, the cost of connecting a well spread out over the well's total gas production.

The first-order condition yields a key insight with regards to efficient regulation. Suppose that flaring imposes a fixed per-unit external cost due to the environmental externalities, lost royalties, and lost tax payments and this external cost enters the regulator's objective. The regulator's first-best policy would equalize shadow prices across all firms. In this case, the left-hand side of equation (2) would be identical across firms, and the wells with the cheapest connection cost per unit of gas production across all firms would be connected. Thus, any differences in $C_j(1)/g_j$ across firms would indicate that the flaring regulation inefficiently allocates abatement. Under the uniform flaring regulation modeled above, λ will differ across firms because firms own portfolios of wells with different levels of productivity and connection costs.

4 Data Description and Summary Statistics

Our dataset consists of monthly, well-level production, flaring, and sales data from over 9,250 horizontal wells owned and operated by 55 firms in North Dakota between 2007 and 2016. For most of our analysis, we focus on around 6,700 wells completed since 2012. We limited the production data from the NDIC in a few ways. First, we focus on oil and gas wells that produced in the Bakken or Three Forks shale formation since the NDIC order applies only to these wells. Second, we drop wells where we observe the maximum level of oil production occurring more than five months after first production. These wells have likely been re-fracked and are therefore not comparable to other wells.[13]

We have a rich set of well-level characteristics including: the year and month a well was spudded (drilled) and completed; the latitude/longitude of the well; the depth and horizontal length of the well bore; and the current and original owner of the well. We also calculate a number of spatial characteristics of each well. First, we proxy for the distance between a well and the nearest gas pipeline infrastructure as the distance to the nearest well that is

[13]Just over 1,000 wells are dropped from the data as a result of this restriction.

connected to gas gathering infrastructure.[14] We also calculate the distance to the nearest natural gas processing plant and the monthly capacity utilization at the plant. We include weather data from the nearest weather monitoring station provided by the North Dakota Agricultural Weather Network, and snowfall data from the NOAA National Operational Hydrologic Remote Sensing Center. Last, we control for historical oil and gas price data using the Henry Hub (HH) natural gas and West Texas Intermediate (WTI) oil prices from the EIA.[15]

Figure 1 graphs the monthly oil and gas production, gas flaring, the number of completed and connected wells, and the number of unconnected wells from January 2007 to February 2016 for all wells in our sample. Oil and gas production, as well as the number of wells completed, grew dramatically until mid-2014 when oil prices began to fall precipitously. The number of unconnected wells peaked in early 2013, and began to decline rapidly in January 2015. Total gas flaring in the state was relatively constant through 2015 and has since declined steadily.

Figure 2 graphs the average oil and gas production, flaring rates, and the fraction of wells connected to gas capture infrastructure in well production time. Production time is in terms of months since the well began producing, with month one corresponding to the first production month, regardless of the calendar month that a well was completed. We differentiate wells by their completion year, and include wells completed in 2012, 2013, 2014, and 2015-2016. The figures document the tremendous gains in the average productivity of wells over this period. Initial oil and gas production averaged 600 barrels per day (bbls/day) and 600 thousand cubic feet per day (mcf/day) in 2012. In 2015-2016, the average well initially produced 750 bbls/day and nearly 900 mcf/day. The figures also illustrate the typical declining production profile of wells where the majority of a well's lifetime oil and gas production occurs in the first year.

The bottom panels of Figure 2 compare average flaring rates at the same sets of wells. Flaring rates declined slowly over production time, but remained around 40% in the fourth month of production for wells completed in 2012 and 2013. Flaring decreased more rapidly in 2013 after the fourth month, but average flaring rates remained above 20% at the end of the first year of production and remained around 15% two years after the well was completed. Wells completed in 2014 and 2015-2016 display nearly identical flaring rates over the first two

[14] Data on gathering line locations in North Dakota is confidential. We determine the month that a well is connected to gas capture infrastructure by flagging the first month that the well has positive gas sales.

[15] Oil and gas prices for North Dakota crude oil are not publicly available at the frequency we require over the full sample period. In a monthly online webinar, the director of the NDIC stated that while there is no traded Bakken oil price, it is typically paid as a basis off of the WTI, and that a good estimate of the price received by Bakken producers is 85% of the WTI price. We are unaware of any posted prices for natural gas in the state, and therefore rely on similar dynamics with regards to gas price determination.

production months. However beginning in month three, wells completed in 2015-2016 show a rapid decline in flaring until around the eighth production month where the gap starts to close. Notably, in the fourth month of production, the month where flaring at new wells starts to count towards firms' aggregate flaring rates, average flaring at wells completed in 2015-2016 is nearly exactly equal to 23% - firms' mandated flaring rates set by the NDIC in 2015. Figure 2d graphs the fraction of wells that connected in a given production month. In 2012 and 2013, just around 40% of wells were connected to gas gathering infrastructure in their first month of production, and some wells remained unconnected after the first full year of production. Wells completed in 2014, and 2015-2016 have higher initial connection rates, and nearly all wells are connected after the first year. However, wells completed after 2015 have a sharp increase in their connection rates in the third, fourth, and fifth production months around when a well becomes subject to the regulation.

Table 1 presents other relevant summary statistics. We divide the statistics by wells completed in 2012-2014 and those completed in 2015-2016. Gas flaring rates in the first year of production fell from about 34% in 2012-2014 to 25% in 2015-2016. Flaring rates are lower at connected wells. However, they are not zero and are roughly the same across wells completed before and after the regulation.[16] The statistics also reflect the large increase in gas capture infrastructure in the Bakken. In 2012-2014, new wells were on average 1 mile from the nearest well connected to gas gathering infrastructure and 14 miles from the nearest gas processing plant.[17] In 2015-2016, those distances decreased to 0.3 miles and 13 miles, respectively. This has been associated with a dramatic decrease in time to gas connection, from 3.5 months to 1.5 months on average. Oil and gas prices have also varied substantially over our sample. Average oil and gas prices were $94/bbl and $3.80/mcf from 2012-2014. Both prices have fallen since the summer of 2014, averaging just $46/bbl for gas and $2.55/mcf for oil in 2015-2016.

[16]Flaring at connected wells is typically the result of issues with or excess pressure in pipelines, or to natural gas plants operating at or near capacity. While average capacity utilization rates at the nearest gas plants have been around 50% since 2011, we observe many months in our sample with capacity utilization at or near 100%.

[17]Data on the location of gas gathering lines in North Dakota is confidential. As such, we proxy for the distance from gas gathering infrastructure in this paper using the distance between a well and the nearest well that has sold a positive amount of gas - our proxy for connected to gathering lines. Conversations with regulators at the NDIC and ND Pipeline Authority confirm that this is a good proxy for distance from pipeline infrastructure.

5 Effects of the NDIC Flaring Order

In this section, we describe our empirical strategy to estimate the impact of the regulation on average flaring rates in North Dakota. We then describe our methods for disentangling the mechanisms along which firms respond to the order. We focus on: (i) time to complete wells; (ii) time to connect wells to gas capture infrastructure; and (iii) oil and gas production.[18] Last, we present our results.

5.1 Empirical Strategy: Flaring

We begin with a reduced form description of the effects of the regulation. Our first empirical strategy uses a differences-in-differences estimation framework that takes advantage of the richness of our data. Because the majority of production, and therefore flaring, occurs in the first few months after a well is completed, we focus on the effects of the policy over the first year of production. We naturally define the treatment group as wells that were completed in January 2015 or later.

Ideally, we would observe wells drilled in similar locations over the same period that are exempt from the regulation. While some unconventional wells are drilled in the Montana and Saskatchewan portions of the Bakken shale formation, the number of wells outside North Dakota are limited, and data is not available at the same scale or temporal resolution. We instead take advantage of the fact that wells drilled before the regulation in North Dakota faced very similar production profiles and define our control wells as those that were completed in 2014.[19] We define time in our differences-in-differences model as production time and include a number of covariates and fixed effects to control for productivity, economic, infrastructure, and weather differences across treatment and control wells. Specifically, we estimate a regression of the following form:

$$Y_{ift\tau} = \rho \mathbf{1}[\text{Treated}] + g(t; \Theta) + \mathbf{X}'_{ift\tau}\beta + \alpha_i + \mathbf{\Gamma}_\tau + \varepsilon_{ift\tau}, \tag{3}$$

[18] We do not consider other dimensions along which firms may respond to the regulation including well location, well length, or fracking input choice. Conversations with regulators and operators in North Dakota suggest that drilling and location decisions are driven by economic conditions, which is primarily determined by oil sales profitability.

[19] Wells completed in 2014 were eventually subject to the regulation. Flaring from a well completed in July 2014 would be included in firm's flaring calculations beginning in January 2015. To address this, we drop any control well observations in 2015-2016. In the appendix, we test the sensitivity of our specification of our control group by including wells completed in earlier years. In addition, we run a number of placebo tests defining, for example, wells completed in 2014 as our placebo treatment group and wells completed in 2013 as our control group. Results are not sensitive to the specification, and our placebo tests support the validity of our difference-in-differences design.

where $Y_{ift\tau}$ is the flaring rate at well i owned by firm f in production month t and calendar month τ.[20] We include a number of covariates $\mathbf{X}_{if\tau}$ to control for important differences between wells completed in 2014 and those completed after 2015. To control for differences in the economic environments, we include oil and gas production and prices. We proxy for distance to gas capture infrastructure using the distance between well i and the closest well that is already connected to gas capture infrastructure.[21] We also include a flexible function of production time, $g(t;\theta)$, to control for practices common across wells in each production month. In our main specification, we specify $g(t;\Theta)$ as production time fixed effects. Last, we include well fixed effects (α_i) and both calendar year and month-of-year fixed effects ($\mathbf{\Gamma}_\tau$) to control for time-invariant features of wells (e.g. initial productivity), common trends in flaring rates over time, and common seasonal effects.

Our main coefficient of interest is ρ, the coefficient on the indicator variable $\mathbf{1}[\text{Treated}]$. The indicator is equal to one after the fourth production month, when wells are included in firms' gas capture calculation, for all wells completed after 2015 and zero otherwise. Our identifying assumption is that absent the NDIC regulation and conditional on our full set of controls, flaring at wells completed in 2015 would be similar over the first year of the production as at wells completed in 2014. For our identification to be invalid, well flaring rates would need to be correlated with an unobserved variable that is varying in production time, but is not due to seasonal factors, annual factors, weather, oil and gas prices, nor distance to gas gathering line infrastructure.

Our second strategy uses a matching estimator that compares average flaring rates after the fourth production month at wells completed in 2015 versus those completed in 2014. Given our large dataset, we match wells on a number of characteristics. In our main specification, we match wells based on their initial oil and gas production, depth of the well, the distance from a connected well in the first production month, and the number of months that we observe the well.[22]

Both strategies above identify changes in flaring at wells after they are subject to the regulation. Neither captures changes in flaring before the fourth production month. To explore firms' responses to the flaring regulation more flexibly, we also estimate regressions

[20]For example, $f_{if,1,\tau}$ is the percent of produced gas that is flared at well i in its first month of production, and $f_{if,12,\tau}$ is the percent of produced gas flared in the twelfth production month.

[21]We have also explored specifications using distance to the nearest gas plant as well as whether the gas plant was near its capacity in a given month. Neither variable appeared to affect well flaring or our results systematically, and were therefore omitted here.

[22]Depth of the well includes both the vertical and lateral length. We match wells exactly on the number of months that we observe them after the fourth production month to ensure we are comparing flaring over the same production time. We include up to five matches per well, and define the distance between well characteristics as the inverse sample covariate covariance.

of the form:

$$Y_{ift\tau} = \sum_{s=1}^{T} \rho_s \mathbf{1}[\text{Treated, t=s}] + g(t;\Theta) + \mathbf{X}'_{ift\tau}\beta + \alpha_i + \mathbf{\Gamma}_\tau + \varepsilon_{ift}. \tag{4}$$

Equation (3) allows for separate coefficients ρ_t for each production month. The coefficients estimate the average difference in flaring rates within each production month between wells completed in 2014 and 2015, conditional on our control variables. The estimates would capture, for example, increasing rates of connecting to gas capture infrastructure in the first, second, or third production month.[23]

5.2 Empirical Strategy: Mechanisms

We use a similar empirical strategy as above to disentangle how firms comply with the regulation. We consider three margins of behavior. First, we test whether firms take systematically longer to complete wells after spudding. Delays between spudding and completion may indicate that firms are installing more on-site infrastructure, including gas capture infrastructure. Second, we test whether, conditional on completion, firms connect to gas capture infrastructure more quickly. Given that gas output is highest in the first months of production, reducing time to connection can significantly increase the total amount of gas captured. Last, we test whether firms curtail oil and gas production at wells subject to the regulation.

We estimate survival (hazard) models for the spud-to-completion time and first production-to-connection time. In the former, wells "survive" if they are still not completed (i.e. not producing) t months after spudding, and "die" if they are completed. In the latter, firms "survive" if they remain unconnected to gas capture infrastructure t months after initial production and "die" if they connect. We use the same control and treatment group as before, comparing the completion and connection time of wells completed in 2014 to wells completed after 2015. As before, we consider only up to the first twelve months of production and throw out data for wells completed in 2014 during the 2015-2016 calendar years. The data are therefore right censored, with some wells remaining incomplete or unconnected in our last observation.

We first estimate a non-parametric Kaplan-Meier (KM) survivor function for each outcome. Let $\bar{t}_j$ denote the month a well is completed or connected to gas capture infrastructure, i_j be the number of wells not completed or connected before month $\bar{t}_j$, and c_j be the number

[23] Increasing rates of connecting to gas capture infrastructure may be expected in the month or two before fourth production month. It takes time to connect a well and firms may want to ensure the well is connected when it becomes subject to the regulation by connecting it earlier.

of wells that are completed or connected in month $\bar{t}_j$. The KM function is given by:

$$\hat{S}(t) = \prod\nolimits_{j|\bar{t}_j \le t} \left(\frac{i_j - c_j}{i_j} \right). \tag{5}$$

We estimate equation (5) separately for wells completed in 2014 and those completed after 2015.

Equation (5) does not control for any differences in the economic environment, gas capture infrastructure, or weather differences between the treatment and control groups. We therefore also estimate a parametric survival model with time-varying controls. Specifically, we estimate a hazard function for wells that are either completed or connected in period t as:

$$h(t, 1[\text{Treated}], \mathbf{X}_{it}; \Theta) = \frac{f(t, 1[\text{Treated}], \mathbf{X}_{it}; \Theta)}{1 - F(t, 1[\text{Treated}], \mathbf{X}_{it}; \Theta)}, \tag{6}$$

where $f(\cdot)$ and $F(\cdot)$ are the density and cumulative density functions, which we specify as a Weibull distribution in our main specifications, of the spud-to-completion time or first production-to-connection time.[24] The covariates $\mathbf{X}_{it}$ include a similar set of well characteristics as before: completion month, oil and gas production, natural gas and oil prices, and weather conditions. Our coefficient of interest is on the indicator function, 1[Treated], for whether the well was completed in 2015 or later.

Coefficients of a survival function can be unintuitive and not particularly policy relevant. As a result, we also estimate regression-adjusted average treatment effects (ATE) of the regulation on the spud-to-completion time and first production-to-connection time. To do this, we first estimate separate Weibull survival models for wells completed in 2014 and those completed after 2015. To ensure we have one predicted survival time for each well, we estimate a time-invariant version of equation (6).[25] We then predict and compare the average survival times for each group to estimate an ATE of the regulation on the time to completion and time to connection.[26]

[24] Results are similar using an exponential and Gompertz survival distribution. Newell et al. (2016) use a generalized gamma distribution to estimate spud-to-completion times for conventional and unconventional oil wells in Texas. Results using a generalized gamma model are also similar to our Weibull results when we do not include covariates. However, including controls in the model leads to convergence issues with the generalized gamma model.

[25] Coefficients for the time-invariant survival function are largely similar to the time-varying parameter survival model.

[26] The ATEs are estimated using the `STATA` command `stteffects ra`. In the spud-to-completion regressions, we include as controls: the well's initial oil and gas production; oil and natural gas prices in the completion month; the distance to a completed well in the completion month; well depth; and local weather conditions in the completion month. In the connection duration regressions, we include controls for: the well's initial oil and gas production; oil and natural gas prices in the connection month; the distance from a completed well in the completion month; and local weather conditions in the connection month.

Last, we test whether the regulation affected wells' oil and gas production. We use the same identification strategies as in our flaring regressions. We first estimate regression equation (3). However, we replace $Y_{ift\tau}$ with the natural logarithm of oil or gas production from a well. In these regressions $g(t;\Theta)$ explicitly controls for the expected oil and gas decline curve. Similar to Newell et al. (2016), we explore three forms of $g(\cdot)$, allowing for increasing flexibility in wells' average decline rates: (i) an Arps model where $g(\cdot)$ is the natural logarithm of production time;[27] (ii) a cubic spline in production time;[28] and (iii) production time fixed effects. Controls in the regressions include oil and natural gas prices and local weather conditions. Our identifying assumption is that, conditional on our set of controls and fixed effects, deviations from the average well production path after the fourth month of production for wells subject to the regulation represent responses to the regulation. As before, we also estimate an ATE using a matching estimator comparing wells' oil and gas production after the fourth month for wells completed in 2014 versus those completed after 2015.

5.3 Results: Flaring Treatment Effects

Table 2 presents our estimates of the effect of the regulation on flaring. Columns (1)-(3) estimate the treatment effects using all wells. Columns (4)-(6) estimate the same regressions but only for wells that connected to gas capture infrastructure in their first two production months. Columns (1)-(2) and (4)-(5) present our difference-in-difference estimates and columns (3) and (5) present our nearest-neighbor matching estimates.

When we include only well fixed effects, we find that the regulation decreased flaring by 23 percentage points at all wells and 11 percentage points at connected wells. This is consistent with the comparison of average flaring rates between 2014 and 2015-16 wells in Figure 2c. Once we control for differences in well productivity, production, and our other covariates, the estimated treatment effect decreases to a 2 percentage point decrease in flaring rates for all wells and nearly zero for connected wells. Nearest-neighbor estimates comparing flaring in months four to twelve suggest larger treatment effects of 7 percentage point declines in flaring rates for all wells and 2 percentage point declines for connected wells. The point estimates of all other covariates have intuitive signs and magnitudes. Flaring rates increase in gas production and distance to the nearest connection, and is largely unaffected by oil production and oil and gas prices.

Figure 3 graphs the estimates and 95% confidence intervals from equation (4), our more

[27] Taking the log of equation (1) leads to a linearly estimable Arps model.

[28] Specifically, we estimate a four-knot restricted cubic spline with knots at 1.1, 1.4, 2, and 2.4 months. Knots are focused early in the production lifetime since this is where the most curvature is in the production path.

flexible treatment effect specification. The estimates are from a regression with the same controls as in column 2 of Table 2. The plotted effects are relative to the regulation's effect in the first production month. When we include all wells in the regression, we find little differential response in the second production month. The largest treatment effects are estimated in the third to seventh production months, where we estimate decreases of 6 to 7 percentage points. We continue to find negative treatment effects after the eighth production month but they are statistically indistinguishable from zero. At connected wells, we estimate small average reductions of up to 2 percentage points in months three to six, but all point estimates are statistically insignificant.

The findings suggest that, on average, the regulation had a moderate but significantly negative effect on flaring, reducing flaring rates between 2-7 percentage points on average. Wells already connected to gas capture infrastructure do not appear to respond to the regulation, suggesting that the order had little effect on well operations after firms connect to gas capture infrastructure.

5.4 Results: Mechanisms

Figure 4 graphs the KM survival probabilities and corresponding 95% confidence intervals for wells' spud-to-completion time and production-to-connection time. We graph separate survival functions for wells completed in 2014 and those completed after 2015. Figure 4a graphs the survival probabilities for each month since initial spudding. In all months, the estimated survival probability (non-completion probability) is higher for wells spudded after 2015 than wells spudded in 2014. For example, six months after spudding, only 42% of 2014 wells remained incomplete, while over 55% of 2015-2016 wells remained incomplete after six months. This is despite the substantial decrease in drilling activity in the region and is consistent with the hypothesis that it is taking longer to install gas capture infrastructure on site.

Figure 4b graphs survival probabilities for the time-to-connection duration models. Wells completed after 2015 have lower survival rates in all months. In the first production month, 45% of wells completed in 2014 remain unconnected while only 39% of wells completed after 2015 are unconnected in the first production month. Survival probabilities are similar in the second and third production month; however, in month four the number of unconnected wells falls steeply for wells completed after 2015, and the survival function for those wells remains statistically significantly lower than wells completed in 2014 until the ninth production month. The findings are consistent with firms responding to the regulation by connecting more wells in the first production month and connecting wells more quickly once the regulation becomes effective after the fourth production month.

Tables 3 and 4 present the estimates from our structural survival models. Coefficients from the survival model in columns (1) to (3) are specified in accelerated failure-time so that a one unit change in explanatory variable x_j increases the failure time by $\exp(\beta_j)$. Columns (4) and (5) present the regression-adjusted mean completion time for 2014 wells and the difference in completion time (measured in months) between wells completed in 2014 and those completed after 2015.

Consistent with our KM estimates, we find that wells spudded after 2015 have longer spud-to-completion times and quicker connection times than those in 2014. When we control for well productivity, fuel prices, well length, and distance to nearest connection, we find that post-2015 wells have over 100% larger non-completion rates, and take just over 0.5 months longer to complete on average. Once production has begun, wells completed post-2015 have 30% shorter non-connection rates and connect to gas capture infrastructure 2.5 to 3.4 months quicker on average than wells completed in 2014.

Table 5 presents our results for the effects of the regulation on firms' oil and gas production. When we control for expected production using the Arps model, the most restrictive production decline curve, we find that treated wells produce more oil and gas on average after the fourth production month. This is inconsistent with the intuition that firms have an incentive to curtail production after the fourth production month under the regulation. However, when we control more flexibly for wells' expected production using either the cubic spline or production time fixed effects, we find much smaller positive effects that are statistically indistinguishable from zero. Except for gas production, we have similar findings with our nearest-neighbor matching estimates. Consistent with previous research, we find that conditional on producing, oil and natural gas prices do not have a statistically significant effect on production after controlling for wells' natural declining curve.

Our results have several takeaways. First, the regulation induced firms to reduce flaring by connecting wells to gas capture infrastructure more quickly. Second, while firms took longer to complete wells, once oil and gas production began, we do not find evidence that firms curtailed oil and gas production to comply with the regulation.

6 Heterogeneous Costs and Gains from Trade

In this section, we explore heterogeneity in firm characteristics as well as in the estimated impacts of the NDIC order across firms. We first present summary statistics of firms' characteristics that are likely correlated with their compliance costs. We then explore heterogeneous impacts of the flaring regulations using similar methods to Section 5. Last, we make use of the theory developed in Section 3 to estimate the net costs of connect wells to gas capture infrastructure and study the efficiency of the regulation.

6.1 Results: Heterogeneous Firm Characteristics

Table 6 summarizes select firm characteristics for January 2015, the month the regulation went into effect. In January 2015, 52 firms owned and operated 7,944 active wells in our sample. The average firm owned and operated 150 wells. However, well ownership is highly skewed, with a few firms owning and operating the majority of wells: the median firm owns only 40 wells, one firm owns a single well, and one firm owns over 1,000 wells.

Total oil and gas production across firms is also heterogeneous. The average (median) firm produced 570,000 (130,000) barrels of oil and 700 (130) million cubic feet of gas, while the largest firm produced nearly 4 million barrels of oil and over 540 billion cubic feet of gas. There is also heterogeneity in firms' flaring rates. On average firms flared 20% of their gas. However, a few firms flared no gas while some flared all of their gas. Firms also own well portfolios in regions with very different installed gas capture infrastructure. The average (maximum) distance firms' wells are from other connected wells is 1 mile (4 miles), but some firms have all their wells extremely close to infrastructure (0.01 miles) while other firms' wells are 10 miles away on average.

6.2 Results: Treatment Effect Heterogeneity

A number of firms' characteristics may be correlated with the costs of the NDICs order. Here, we focus on one particularly salient characteristic: the baseline flaring rate of firms.[29] Table 7 presents estimates of heterogeneity in the average treatment effect based on firms' average flaring rates before the regulation. To estimate the coefficients, we interact indicators for firms' "baseline" flaring rates with the treatment indicator in equation (3). We consider two baseline periods, 2013 and 2014 calendar years, and create indicators for whether firms' average flaring rates in those years were: (i) less than 20%; (ii) between 20% and 25%; (iii) between 25% and 30%; (iv) between 30% and 40%; and (v) greater than 40%. The estimated coefficients, therefore, represent the average decrease in treated wells after their fourth production month owned by the firm in that bin's baseline flaring range compared to

[29]Since it is firm flaring that is regulated, the treatment effect is likely heterogeneous in current flaring rates. However, there is a simultaneity problem in estimating effects using current firm flaring rates. As a result, we focus on historical flaring rates and use two alternative 'baseline' years to test the sensitivity of our results.

wells owned by firms with similar baseline flaring rates from 2014.[30]

As in Table 2, when we include only well fixed effects, we estimate that the flaring restriction led to large decreases in flaring rates across all wells, with the largest treatment effects estimated for wells owned by firms with baseline flaring rates between 20% and 25% and those with baseline flaring rates above 40%. When we control for production, price, and weather as well as include our calendar and production time fixed effects, all coefficients shrink. When we define our baseline year as 2013, the only remaining flaring bins where we observe significant flaring reductions are those that are just around the flaring limit (20% to 25%) and those with very high baseline flaring rates (>40%). We find similar results when we use 2014 as are baseline year; however, we find significant increases in the 25% to 30% baseline group and decreases in the 30% to 40% baseline group.

The summary statistics suggest that firms face differential marginal compliance costs with the flaring order. Furthermore, our the estimated heterogeneous treatment effects suggest that the regulation has lead to large flaring reductions among only a subset of firms and small reductions among others. Both results suggest that there may be gains from reallocating abatement from high- to low-cost firms.

6.3 Marginal Connection Cost Curves and Gas Capture Efficiency

Section 3 showed how under an efficient regulation, all wells with connection costs below some threshold would be connected and all wells above the threshold would be left unconnected. We use this insight to study the efficiency of the flaring regulation. First, assume that operators use the current gas price as their expectations for future prices and all firms face the same gas price. The regulation is satisfied at minimum total compliance cost if and only if the connection cost per unit of gas produced by a well is lower for all connected wells than all unconnected wells.[31] In the usual theoretical settings, this occurs when the costs for each firm's marginal well are equal. In this setting firms have highly discrete connection decisions so this may not precisely hold. Here we require that no firm can connect an unconnected well at cheaper cost per unit of gas than a well that has already been connected. From here on we call a firm's or the market's marginal connection cost per unit of gas curve the marginal flaring abatement cost curve.

[30] The estimated coefficients are equivalent to estimating separate treatment effects using only wells owned by firms in each baseline flaring bin.

[31] In a dynamic setting, this condition need not hold. For example, a firm may connect a well that statically has too high of connection costs to be efficient because the firm is forward looking and anticipates connected more wells to this newly developed infrastructure in the future. Here we do not study dynamic efficiency so our results can be thought of as an upper bound on the efficiency costs of uniform regulation conditional on firms acting rationally.

Consider a well i, we model it's connection cost per unit of gas as,

$$\frac{(\text{On-site Costs}) + (\text{Line Costs}) \times d_i}{g_i}$$

The first term in the numerator is the cost from the on-site equipment necessary to capture the gas,[32] the second term is the cost of building out the gathering line to well i, which is a function of length of gathering line that needs to be built, d_i. The term in the denominator is the total gas production at the well over its lifetime.

We calculate well i's gas production g_{it} at all future production months t by estimating an Arps production function for wells completed in 2014-2016:

$$log(\text{g}_{it}) = \beta_1 t + \theta_i + \varepsilon_{it}$$

where θ_i is a well fixed effect. Our estimate of the decline rate is $\hat{\beta}_1 = -.045$. To obtain the initial production level Q_0 we do one of two things. First, if a well was connected prior to the current calendar month m, the firm has already observed the initial production level. Here we use the actual initial production level of the well and then calculate all future production levels using the Arps model. For wells that have yet to produce gas, if the firm wants to connect the well in the current month m, it is making this decision before observing production. Here we proxy for expected production in future months using the average of other recently completed nearby wells. Specifically, we find all wells completed in the last 6 months that are within 25 kilometers, and then compute their average gas production in each production month.[33]

We obtain on-site costs and per-mile estimates of gas gathering line costs from the Interstate Natural Gas Association of America (INGAA). INGAA reports the costs for equipment at oil wells is \$210,000 per well. INGAA reports that the median cost per mile of gathering line is approximately \$200,000. To proxy for how long of a gathering line must be built for some well i, we find the nearest well h to well i that was selling gas in our data in the previous month. To be selling gas, well h must be connected to a gathering line. We assume that the distance to well h is the lowest cost way for well i to be connected and sell its gas and use this distance as the length of gathering line that must be laid for well i to be connected.[34]

We can now construct marginal flaring abatement cost curves for the entire oil extraction

[32] This may include dehydrators, compressors, or technology to remove hazardous pollutants like hydrogen sulfide.

[33] This is similar to Newell et al. (2016).

[34] Alternatively one may consider distance to a gas processing facility. However, in our empirics, we find this distance to have a statistically and economically insignificant effect on well flaring which suggests this is not the correct distance to use.

market and for individual firms by sorting wells according to their connection costs per unit of gas. These are displayed in Figure 5 for November 2015 and in Figure 6 February 2016. Since firms comply with the regulation predominately by connecting wells, the marginal flaring abatement cost curve consists entirely of wells that were unconnected before that month. We focus on individual months since firms must comply with the regulation on a monthly basis and the data are also reported monthly.

In the two figures, wells that were connected are displayed with an orange circle, while wells that remained unconnected at the end of the month are displayed with a blue X. The top left panels show the market marginal flaring abatement cost curves. In both months, the wells with the lowest connection cost were connected, and wells with the highest connection costs were left unconnected. However, there are some cases where high-cost wells were connected while low-cost wells were not indicating that flaring reducing are not being achieved at least cost. By looking only at the market marginal connection cost curves, it is unclear whether this is due to issues with our static approach or that it is because firms have different marginal connection cost curves.

The bottom left panels display marginal compliance cost curves for a selected set of individual firms, labeled A-F. Individually, the firms appear to be acting roughly according to our theory and connecting their lowest cost wells. However, the firms have substantially different marginal connection cost curves. For example, firm D has many low-cost wells, while firms B, E and F have much higher cost wells suggesting that there are potential gains from moving to a market-based regulation.

The right two panels graph the connection outcomes if the wells were connected in an efficient manner with the objective of achieving the same flaring reduction. The top right panels show that in November 2015, the costs of connection could have been reduced by \$20 million and in February 2016 the costs could have been reduced by \$30 million. The bottom right panels show that in the efficient outcome in November 2015, only firm D connects wells out of the set of 5 firms displayed. In February 2016, the efficient connection outcome looks similar on a firm-by-firm basis since the efficient wells were largely the ones being connected in reality. The large gains under the efficient outcome are due to a handful of very costly wells being connected.

A final takeaway from the graphs is that the most productive wells, wells that have big horizontal gaps on the graph, also tend to be low connection cost wells, located near the bottom of the graph. This is consistent with firms clustering their drilling activity in productive oil and gas regions and, due to the clustering of the wells in the same area, there is nearby gas capture infrastructure for all of these wells. The unproductive wells high connection cost wells may be exploratory wells far from other wells and existing gas capture infrastructure.

7 Conclusions and Discussion

We use a rich, well-level dataset on oil firms' operations in North Dakota to study the effects and efficiency of a prominent regulation aimed at reducing gas flaring in the state. Our results suggest that the regulation has been effective. Well operators have reduced flaring two to seven percentage points by increasing the speed with which they connect wells to gas capture infrastructure. However, we show that there are substantive efficiency costs due to the presence of heterogeneous compliance costs and the regulation being enforced uniformly across firms. Using a simple counterfactual exercise we show that reallocating abatement from high- to low-cost firms would reduce aggregate compliance costs by tens of millions of dollars.

Our results are subject to a number of caveats. We rely on reduced-form methods to estimate the average treatment effects of the regulation. The methods, therefore, do not allow for strategic decision-making by firms or take full advantage of the feature that connecting to gas capture infrastructure requires potentially large upfront costs or forward looking behavior. In addition, our results are conditional on the existing state of gas capture infrastructure in North Dakota. Thus, our results pertain only to the effects of the regulation on oil operators' gathering line installation decisions and do not allow for strategic investments in other gas capture and processing infrastructure. Discussions with regulators in North Dakota confirm that among the more salient changes since the passage of the regulation is more regular coordination between oil operators and gas processing plants. Future work may more explicitly consider the interactions between these two groups, as well as consider the effects of the NDIC order and other regulations on the development and placement of gas pipelines and processing plants.

In addition, recent work studying the Texas oil and gas industry shows that a failure to internalize environmental risks due to bankruptcy protections shifts industry structure towards smaller firms (Boomhower, 2016). Small firms may also take advantage of the benefits of limited liability in the North Dakota shale fields. The introduction of new, stringent flaring standards may, therefore, act to increase capital costs for drilling new wells. If these new, larger upfront costs affect firms' entry decisions, the new standard may have the effect of pricing smaller, capital constrained firms out of the market. Future research may explore these issues along with a number of other effects of this and similar regulations.

References

Anderson, S., R. Kellogg, and S. Salant (2016). Hotelling Under Pressure. Working Paper.

Bailey, A. (2015). NDPA Director Analysis Reveals Bakken Breakeven Pricing. *The Bakken Magazine*. September 16, 2015. http://thebakken.com/articles/1287/ndpa-director-analysis-reveals-bakken-breakeven-pricing.

Bartik, A., J. Currie, M. Greenstone, and C. Knittel (2016). The Local Economic and Welfare Consequences of Hydraulic Fracturing. Working Paper.

Baumol, W. J. and W. E. Oates (1988). *The theory of environmental policy*. Cambridge university press.

Boomhower, J. (2016). Drilling Like There's No Tomorrow: Bankruptcy, Insurance, and Environmental Risk. Working Paper.

Bureau of Land Management (2016). Proposed Rule: Waste Prevention, Production Subject to Royalties, and Resource Conservation. *Federal Register 81*(25), 6616–6686.

Carlson, C., D. Burtraw, M. Cropper, and K. L. Palmer (2000). Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade? *Journal of Political Economy 108*(6), 1292–1326.

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

Energy Information Administration (2015). Hydraulic Fracturing Accounts for about Half of Current U.S. Crude Oil Production. Today in Energy, March 15, 2016. https://www.eia.gov/todayinenergy/detail.php?id=25372.

Energy Information Administration (2016a). International Energy Statistics.

Energy Information Administration (2016b). Natural gas flaring in North Dakota has declined sharply since 2014. Today in Energy, June 13, 2013. http://www.eia.gov/todayinenergy/detail.cfm?id=11671.

Fetkovich, M. (1980). Decline Curve Analysis using Type Curves. *Journal of Petroleum Technology 32*(6), 1065–1077.

Feyrer, J., E. Mansur, and B. Sacerdote (2017). Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution. *American Economic Review 107*(4), 1313–1334.

Fischer, C. and R. G. Newell (2008, mar). Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management 55*(2), 142–162.

Fowlie, M., C. R. Knittel, and C. Wolfram (2012). Sacred Cars? Cost-effective regulation of stationary and nonstationary pollution sources. *American Economic Journal: Economic Policy 4*(1), 98–126.

Gaswirth, S., K. Marra, T. Cook, R. Charpentier, D. Gautier, D. Higley, T. Klett, M. Lewan, P. Lillis, C. Schenk, M. Tennyson, and K. Whidden (2013). Assessment of undiscovered oil resources in the Bakken and Three Forks Formations, Williston Basin Province, Montana, North Dakota, and South Dakota. U.S. Geological Survey Fact Sheet 2013–3013, `http://pubs.usgs.gov/fs/2013/3013/`.

Gopalakrishnan, S. and H. Klaiber (2014). Is the Shale Energy Boom a Bust for Nearby Residents? Evidence from Housing Values in Pennsylvania . *American Journal of Agricultural Economics 96*(1), 43–66.

Hausman, C. and R. Kellogg (2015). Welfare and Distributional Implications of Shale Gas. *Brookings Papers on Economic Activity Spring*, 71–125.

Helfand, G. (1991). Standards versus standards: the effects of different pollution restrictions. *The American Economic Review 81*(3), 622–634.

Hill, E. (2012). Shale Gas Development and Infant Health: Evidence from Pennsylvania. Working Paper.

Hill, E. (2015). Impacts of Shale Gas Development on Ambient Air Pollution and Infant Health in Texas. Working Paper.

Holland, S., J. Hughes, and C. Knittel (2009, feb). Greenhouse Gas Reductions under Low Carbon Fuel Standards? *American Economic Journal: Economic Policy 1*(1), 106–146.

Holland, S., J. E. Hughes, C. Knittel, and N. Parker (2015). Some Inconvenient Truths About Climate Change Policy: The Distributional Impacts of Transportation Policies. *The Review of Economics and Statistics 97*(5).

Hollingsworth, A. and I. Rudik (2016). External Impacts of Local Energy Policy: The Case of Renewable Portfolio Standards.

International, I. (2014). North American Midstream Infrastructure through 2035: Capitalizing on Our Energy Abundance. *INGAA Foundation Report, Prepared by ICF International* March 18, 2014. `http://www.ingaa.org/File.aspx?id=21498`.

Kellogg, R. (2011). Learning by drilling: Interfirm learning and relationship persistence in the texas oilpatch. *The Quarterly Journal of Economics 126*(4), 1961–2004.

Lapan, H. and G. Moschini (2012). Second-Best Biofuels Policies and the Welfare Effects of Quantity Mandates and Subsidies. *Journal of Environmental Economics and Management 63*(2), 224–241.

Lemoine, D. (2016, feb). Escape from Third-Best: Rating Emissions for Intensity Standards. *Environmental and Resource Economics*.

Muehlenbachs, L., E. Spiller, and C. Timmins (2015). The Housing Market Impacts of Shale Gas Development. *American Economic Review 105*(12), 3633–3659.

Newell, R., B. Prest, and A. Vissing (2016). Trophy Hunting vs. Manufacturing Energy: The Price-Responsiveness of Shale Gas. RFF Discussion Paper 16-32.

North Dakota Industrial Commission (2015). Order 24665. Technical report.

North Dakota Industrial Commission (2016). North Dakota Monthly Oil Production Statistics. Accessed March 18, 2016. `https://www.dmr.nd.gov/oilgas/stats/historicaloilprodstats.pdf`.

North Dakota Pipeline Authority (2015). Gas plants. Accessed October 13, 2015. `http://northdakotapipelines.com/gas-plants/`.

Olmstead, S., L. Muehlenbachs, J.-S. Shih, Z. Chu, and A. Krupnick (2013). Shale gas development impacts on surface water quality in Pennsylvania. *Proceedings of the National Academy of Sciences 110*(13), 4962–4967.

Salmon, R. and A. Logan (2013). Flaring Up: North Dakota Natural Gas Flaring More than Doubles in Two Years. Ceres Research Report. July 2013. Accessed October 2015.`https://www.ceres.org/resources/reports/flaring-up-north-dakota-natural-gas-flaring-more-than-doubles-in-two-years/`.

Storrow, B. (2015). Could North Dakota be a Model for How to Reduce Flaring? *The Casper Star Tribune.* March 10, 2015. `http://trib.com/business/energy/could-north-dakota-be-a-model-for-how-to-reduce/article_441dd743-5cb4-5504-8d42-1934f2a852a3.html`.

Swanson, A. (2014). America's Oil Boom is Visible from Space. *The Washington Post.* October 20, 2014. `https://www.washingtonpost.com/news/wonk/wp/2014/10/20/americas-oil-boom-is-visible-from-space/?utm_term=.2f3993e72123`.

Tedesco, J. and J. Hiller (2014). Up in Flames. *San Antonio Express News* (San Antonio, TX). August 22, 2014. Accessed March 17, 2016. http://www.mysanantonio.com/news/local/article/Up-in-Flames-5704385.php.

Vissing, A. (2016). One-to-Many Matching with Complementary Preferences: An Empirical Study of Natural Gas Lease Quality and Market Power. Working Paper.

World Bank (2015). Zero Routine Flaring by 2030. Accessed October 2015. http://www.worldbank.org/en/programs/zero-routine-flaring-by-2030.

Zirogiannis, N., J. Alcorn, J. Rupp, S. Carley, and J. D. Graham (2016). State regulation of unconventional gas development in the U.S.: An empirical evaluation. *Energy Research & Social Science 11*, 142–154.

8 Tables and Figures

Figure 1: Production, Flaring, and Completions in the Bakken

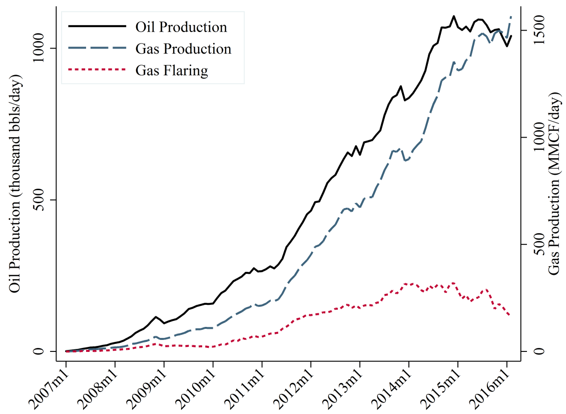

(a) Aggregate Oil & Gas Production and Flaring

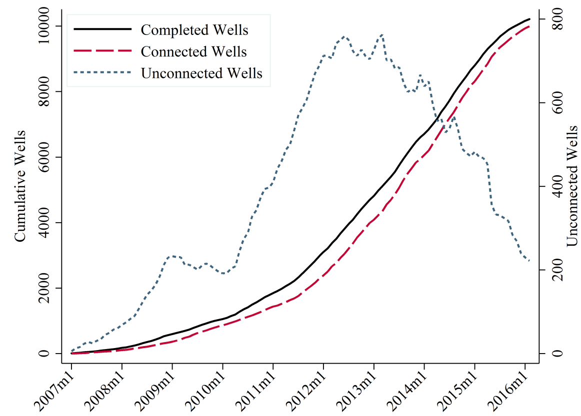

(b) Completed and Connected Wells

Notes: Figure 1a graphs total production and flaring from all horizontal wells in the Bakken/Three Forks formation completed between January 2007 and February 2016. The left axis refers to oil production in thousand barrels per day, and the right axis refers to gas production/flaring in million MCF per day. Figure 1b graphs the cumulative number of completed wells and wells connected to gas gathering infrastructure (left axis), as well as the number of unconnected in each month (right axis).

Figure 2: Well Production, Flaring and Connection Rates by Production Month

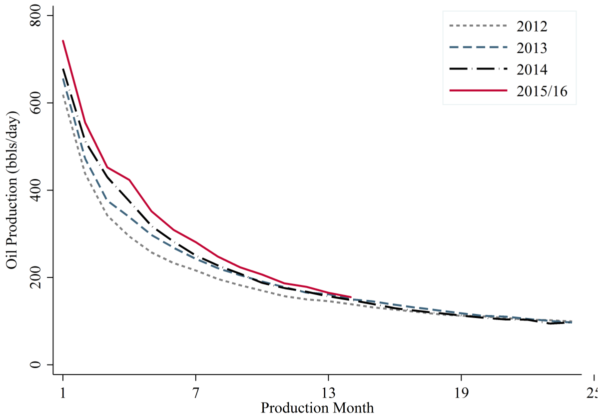

(a) Average Oil Production

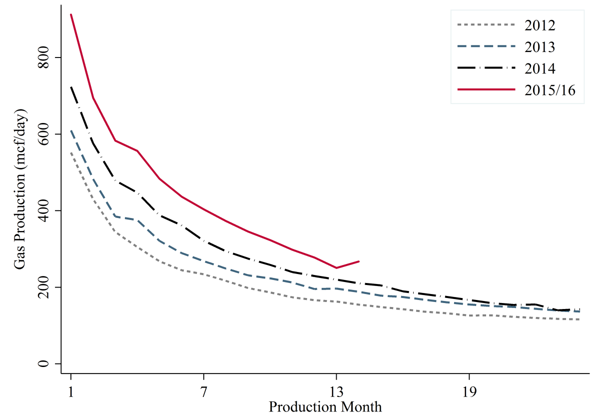

(b) Average Gas Production

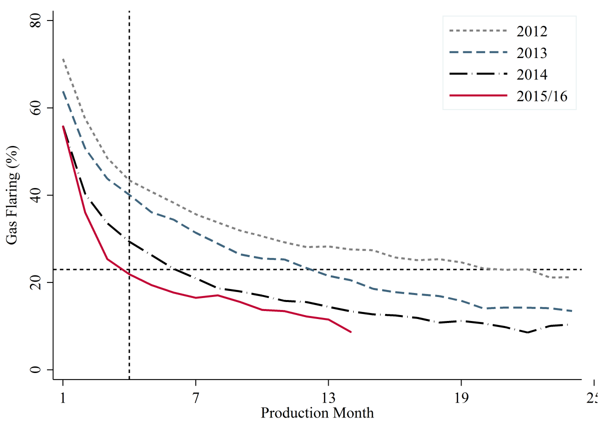

(c) Average Flaring Rate

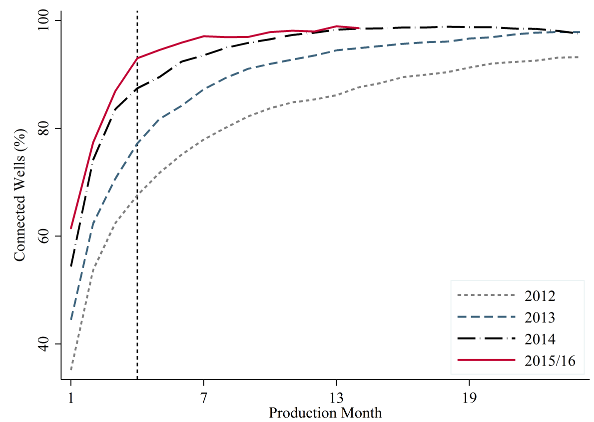

(d) Average Connection Rate

Notes: Figures 2a and 2b graph the average oil and gas production over the first two production years at wells completed in 2007, 2010, 2013, and 2015/16. Figure 2c graphs average flaring rates over the first two production years at wells completed in the same years. The horizontal dashed line is the NDIC flaring targeted flaring rate, and the vertical dashed line represents the production month when wells' flaring rates begin to count towards firms' flaring rate calculations. Figure 2c graphs the fraction of wells connected in each production month for each respective completion year.

Figure 3: Treatment Effects by Production Month

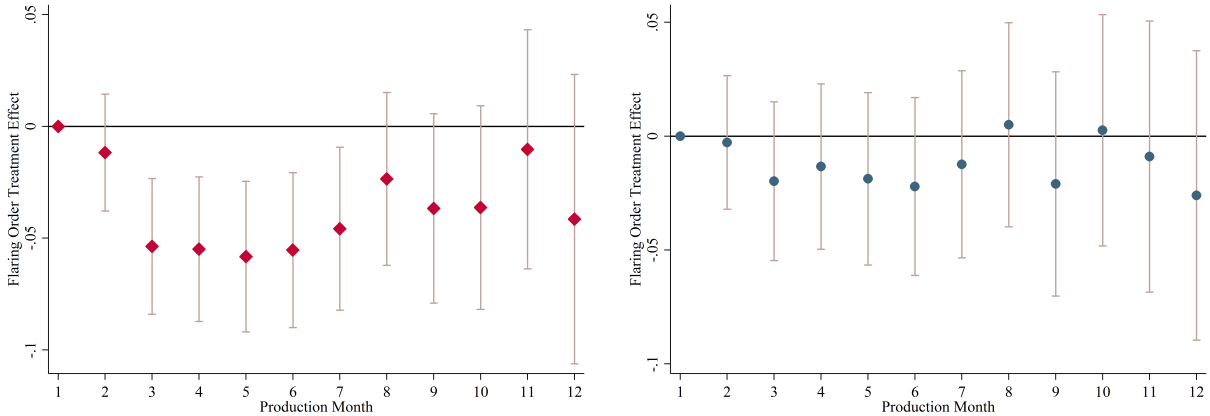

(a) All Wells

(b) Connected Wells

Notes: Figure 3 graphs the point estimates and corresponding 95% confidence intervals from estimating equation (4). Time is specified in production time, with month 1 corresponding to the first production month, and the effects are relative to the regulation's effect in the first production month. Figure 3a includes all wells, and Figure 3b includes only wells that were connected in the first two production months. Both regressions include the same controls as in column 2 of Table 2. Standard errors are clustered at the well level.

Figure 4: Kaplan-Meier Survival Estimates

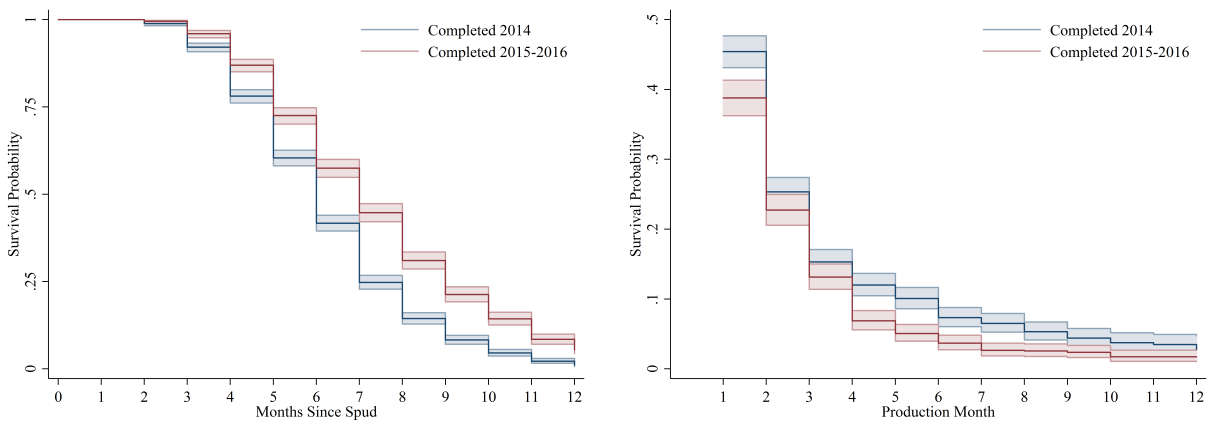

(a) Spud-to-Completion Time

(b) Time to Gas Connection

Notes: Figure 4 graphs separate KM survival probability and 95% confidence intervals for wells completed in 2014 and wells completed after 2015. Figure 4a graphs the KM survival probabilities for spud-to-completion time, where month 1 corresponds to the month a well was spudded. Figure 4b graphs the KM survival probabilities for production to gas connection time, where month 1 corresponds to the first production month of a well.

Figure 5: Market and Firm Level Marginal Connection Cost Curves
(November 2015)

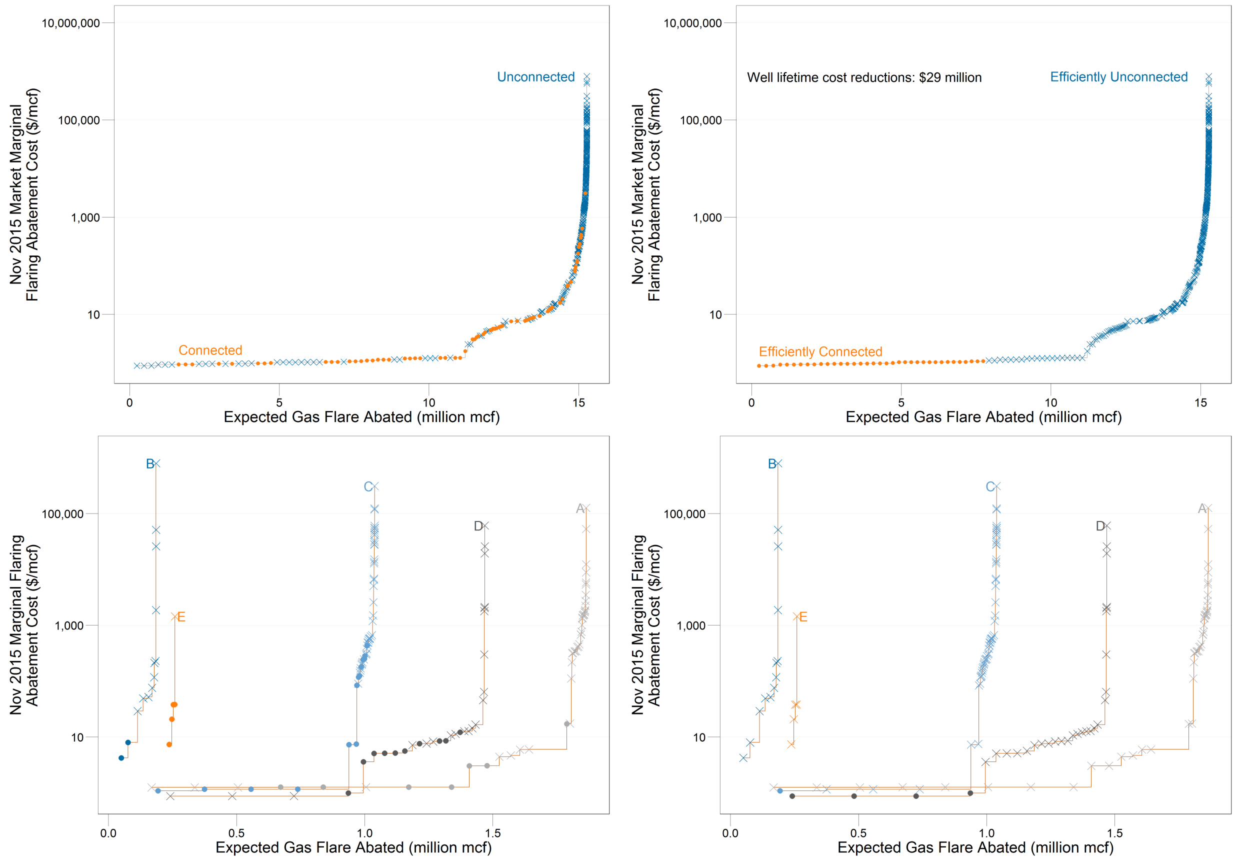

Notes: The top row contains the market marginal connection cost curves, and the bottom row contains marginal connection cost curves for a selected set of firms. The left column shows which wells actually were connected to gas capture infrastructure, while the right column shows which wells should have been connected in order to achieve at least the same flaring reduction but at lowest possible cost.

Figure 6: Market and Firm Level Marginal Connection Cost Curves
(February 2016)

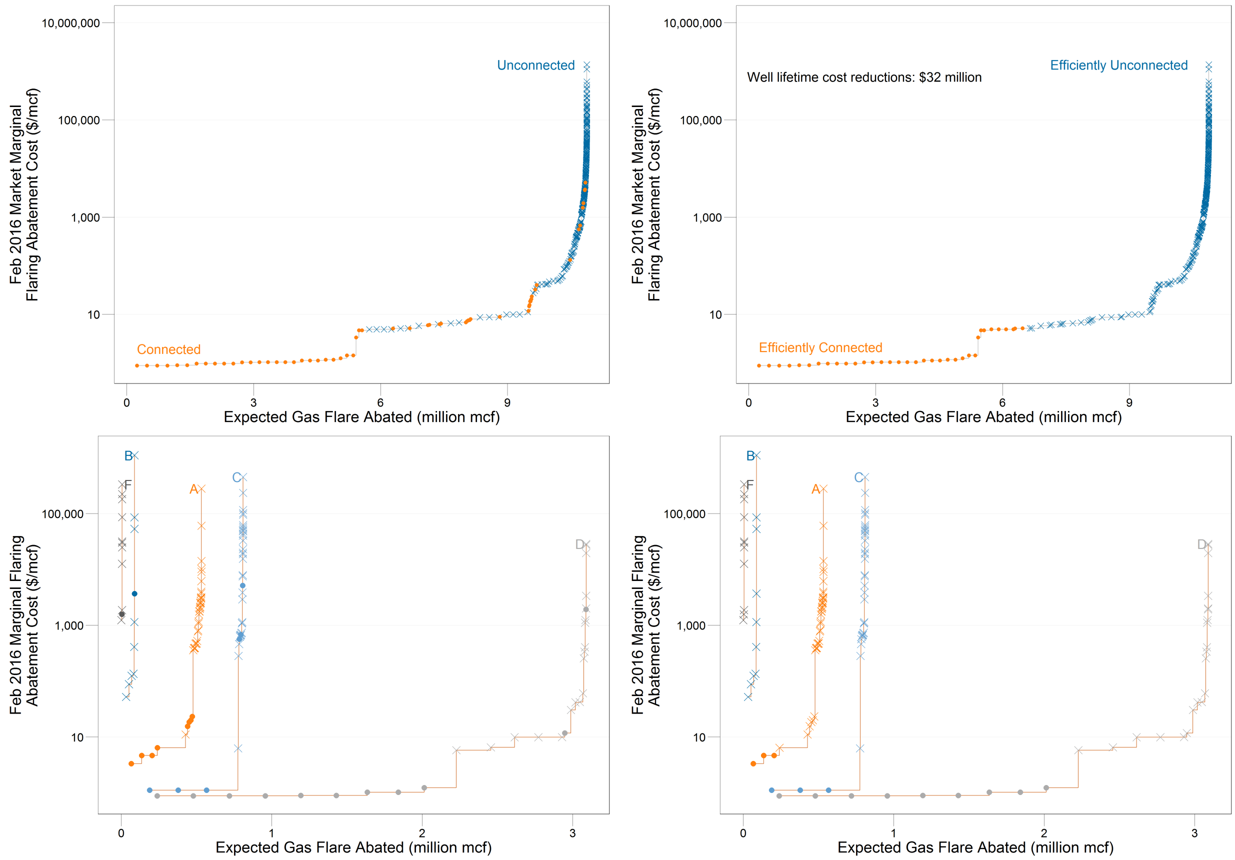

Notes: The top row contains the market marginal connection cost curves, and the bottom row contains marginal connection cost curves for a selected set of firms. The left column shows which wells actually were connected to gas capture infrastructure, while the right column shows which wells should have been connected in order to achieve at least the same flaring reduction but at lowest possible cost.

Table 1: Summary Statistics:
Well Characteristics and Prices by Completion Year

		Mean	Median	Std. Dev.
2012-2014	Oil Production in 1st Year (bbls/day)	297.93	222.89	280.26
	Gas Production in 1st Year (mcf/day)	329.20	219.48	392.16
	Flaring in 1st Year: All Wells (%)	0.34	0.11	0.40
	Flaring in 1st Year: Connected Wells (%)	0.21	0.05	0.30
	Well Depth (ft)	20,082	20496.50	1615.31
	Distance to Connection	0.96	0.40	1.50
	Distance to Gas Plant (miles)	14.23	13.63	7.19
	Capacity Utilization at Closest Gas Plant (%)	0.51	0.58	0.24
	Time to Gas Connection (Months)	3.51	2.00	4.91
	WTI Price ($/bbl)	94.54	95.31	10.48
	Henry Hub Price ($/mcf)	3.80	3.81	0.78
2015-2016	Oil Production in 1st Year (bbls/day)	390.57	313.25	339.13
	Gas Production in 1st Year (mcf/day)	521.53	380.11	519.63
	Flaring in 1st Year: All Wells (%)	0.25	0.08	0.33
	Flaring in 1st Year: Connected Wells (%)	0.20	0.07	0.28
	Well Depth (ft)	20,418	20,715	1,560.39
	Distance to Closest Connected Well (miles)	0.30	0.03	0.61
	Distance from Closest Gas Plant (miles)	13.11	12.91	6.59
	Capacity Utilization at Closest Gas Plant (%)	0.55	0.64	0.24
	Time to Gas Connection (Months)	1.84	1.00	1.44
	WTI Price ($/bbl)	45.91	46.22	8.70
	Henry Hub Price ($/mcf)	2.55	2.66	0.35

Table 2: Average Effect of Regulation on Flaring

	All Wells			Connected Wells		
	(1)	(2)	(3)	(4)	(5)	(6)
	D-in-D	D-in-D	NN Match	D-in-D	D-in-D	NN Match
Flaring Restriction	-0.229***	-0.022**	-0.070***	-0.114***	-0.003	-0.021**
	(0.008)	(0.011)	(0.010)	(0.008)	(0.013)	(0.010)
Log Oil Production		-0.015***			-0.012	
		(0.005)			(0.009)	
Log Gas Production		0.064***			0.030***	
		(0.005)			(0.008)	
Log WTI Price		-0.053			-0.037	
		(0.073)			(0.084)	
Log HH Price		0.042			0.025	
		(0.080)			(0.093)	
Log Dist to Connection		0.056***			0.009**	
		(0.003)			(0.005)	
Observations	23,240	23,211	2,554	13,175	13,166	1,474
Wells	3,045	3,042	2,602	1,751	1,750	1,474
Well FE	Yes	Yes	No	Yes	Yes	No
Year FE	No	Yes	No	No	Yes	No
Month FE	No	Yes	No	No	Yes	No
Production Time FE	No	Yes	No	No	Yes	No
Weather Controls	No	Yes	No	No	Yes	No

Notes: The dependent variable is the well-level flaring rate. Time is specified as production time. Treated wells are those completed after 2015 and have produced for more than four months. Control wells are those completed in 2014. D-in-D denotes our differences-in-differences specification comparing flaring rates at wells completed in 2015 to those completed in 2014. NN Match denotes a nearest-neighbor matching estimate of the effect of the flaring regulation. Wells drilled in 2015 are matched to those drilled in 2014 based on their initial oil and gas production, depth, initial distance to nearest connected well, and the number of observed production months. The NN Match only uses data from treatment and control wells after four months of production. Standard errors are clustered at the well level. *, **, *** denotes significance at the 10%, 5%, and 1% level.

Table 3: Spud-to-Completion Duration:
(Wells Completed in 2014 vs. 2015-2016)

	(1)	(2)	(3)	(4)	(5)
	AFT	AFT	AFT	ATE	ATE
Completed Post-2015	0.180***	0.635***	0.780***	0.625***	0.542*
	(0.015)	(0.027)	(0.032)	(0.234)	(0.304)
Log Oil Production		-0.015	-0.015		
		(0.011)	(0.011)		
Log Gas Production		0.012	0.014*		
		(0.008)	(0.008)		
Log HH Price		-0.025	-0.331***		
		(0.048)	(0.072)		
Log WTI Price		0.731***	1.220***		
		(0.045)	(0.065)		
Log Dist to Connection		-0.044***	-0.045***		
		(0.004)	(0.004)		
Log Total Depth of Well		0.393***	0.359***		
		(0.090)	(0.093)		
Mean Completion Time				6.246***	6.464***
(2014 Wells)				(0.114)	(0.153)
Observations	22,021	21,443	21,443	3,167	3,164
Model	AFT	AFT	AFT	ATE	ATE
Density	Weibull	Weibull	Weibull	Weibull	Weibull
Weather Controls	No	No	Yes	No	Yes

Notes: The dependent variable is the spud-to-completion duration measured in months since a well was first spudded. Columns (1)-(3) present estimates from the parametric survival functions which are specified in accelerated failure time (AFT). Columns (4) and (5) present are estimated average treatment effect of the regulation, which is measured in months. Standard errors in all are clustered at the well level. *, **, *** denotes significance at the 10%, 5%, and 1% level.

Table 4: Connection Duration:
(Wells Completed in 2014 vs. 2015-2016)

	(1) AFT	(2) AFT	(3) AFT	(4) ATE	(5) ATE
Completed Post-2015	-0.109***	-0.317***	-0.328***	-2.698***	-3.397***
	(0.039)	(0.051)	(0.060)	(0.239)	(0.276)
Log Oil Production		-0.296***	-0.294***		
		(0.025)	(0.026)		
Log Gas Production		-0.059***	-0.060***		
		(0.015)	(0.015)		
Log HH Price		-0.666***	-0.297**		
		(0.105)	(0.121)		
Log WTI Price		-0.001	-0.265**		
		(0.072)	(0.103)		
Log Dist to Connection		0.130***	0.130***		
		(0.007)	(0.007)		
Mean Connection Time				4.193***	4.866***
(2014 Wells)				(0.234)	(0.268)
Observations	6,371	6,371	6,351	2,948	2,945
Model	AFT	AFT	AFT	ATE	ATE
Density	Weibull	Weibull	Weibull	Weibull	Weibull
Weather Controls	No	No	Yes	No	Yes

Notes: The dependent variable is the time to connection duration measured in months since a well first began production. Columns (1)-(3) present estimates from the parametric survival functions which are specified in accelerated failure time (AFT). Columns (4) and (5) present are estimated average treatment effect of the regulation, which is measured in months. Standard errors in all are clustered at the well level. *, **, *** denotes significance at the 10%, 5%, and 1% level.

Table 5: Average Effect of Regulation on Oil and Gas Production

Table 5(a)
(Dependent Variable: Oil Production)

	(1) D-in-D	(2) D-in-D	(3) D-in-D	(4) NN Match
Flaring Restriction	0.181***	0.056	0.061	0.036
	(0.041)	(0.040)	(0.042)	(0.038)
Log HH Price	0.064	0.215	0.195	
	(0.316)	(0.314)	(0.313)	
Log WTI Price	0.174	0.453	0.439	
	(0.284)	(0.285)	(0.283)	
Observations	23,211	23,211	23,211	2,602
Production Time Control	ARPS	Cubic Spline	Prod FEs	N/A
Well FE	Yes	Yes	Yes	No
Year FE	Yes	Yes	Yes	No
Month FE	Yes	Yes	Yes	No
Weather Controls	Yes	Yes	Yes	No

Table 5(b)
(Dependent Variable: Gas Production)

	(1) D-in-D	(2) D-in-D	(3) D-in-D	(4) NN Match
Flaring Restriction	0.124***	0.024	0.015	0.095**
	(0.043)	(0.044)	(0.046)	(0.045)
Log HH Price	-0.074	0.046	0.050	
	(0.360)	(0.359)	(0.358)	
Log WTI Price	-0.046	0.170	0.185	
	(0.319)	(0.321)	(0.319)	
Observations	23,211	23,211	23,211	2,602
Production Time Control	ARPS	Cubic Spline	Prod FEs	N/A
Well FE	Yes	Yes	Yes	No
Year FE	Yes	Yes	Yes	No
Month FE	Yes	Yes	Yes	No
Weather Controls	Yes	Yes	Yes	No

Notes: The dependent variable is the log of each well's oil and gas production. D-in-D denotes differences-in-differences specification comparing flaring rates at wells completed in 2015 to those completed in 2014. NN Match denotes a nearest-neighbor matching estimate of the effect of the flaring regulation. Wells drilled in 2015 are matched to those drilled in 2014 based on their initial oil and gas production, depth, initial distance to nearest connected well, and the number of observed production months. The NN Match only uses data from treatment and control wells after four months of production. Standard errors are clustered at the well level. *, **, *** denotes significance at the 10%, 5%, and 1% level.

Table 6: Summary Statistics: Firm Heterogeneity
(January 2015)

	Mean	Median	Std Deviation	Min	Max
No. Wells	152.77	40	239.57	1	1,042
Oil Production (mil bbls)	0.57	0.13	0.89	0.00	3.92
Gas Production (bcf)	0.70	0.13	1.21	0.00	5.45
Avg. Flaring (%)	0.20	0.17	0.23	0.00	1.00
Avg Dist to Connected Well (mi)	1.00	0.47	1.74	0.01	10.14
Max Dist to Connected Well (mi)	3.98	3.26	3.53	0.01	19.78

Notes: The table presents aggregate statistics for all firms with wells actively producing in the Bakken formation in January 2015 to illustrate the heterogeneity in firm characteristics during our sample period. In total, 52 firms owned and operated 7,944 wells in our sample during this month.

Table 7: Average Effect of Regulation on Flaring

	(1)	(2)	(3)	(4)
Flaring Restriction	-0.163***	-0.024	-0.248***	-0.079***
(0%-20% Baseline)	(0.023)	(0.022)	(0.014)	(0.014)
Flaring Restriction	-0.348***	-0.172***	-0.321***	-0.153***
(20%-25% Baseline)	(0.015)	(0.015)	(0.016)	(0.017)
Flaring Restriction	-0.175***	0.014	-0.119***	0.077***
(25%-30% Baseline)	(0.012)	(0.013)	(0.015)	(0.015)
Flaring Restriction	-0.163***	0.006	-0.235***	-0.088***
(30%-40% Baseline)	(0.027)	(0.026)	(0.015)	(0.015)
Flaring Restriction	-0.243***	-0.071***	-0.337***	-0.130***
(>40% Baseline)	(0.020)	(0.018)	(0.046)	(0.039)
Observations	23,240	23,211	23,240	23,211
Base Year	2013	2013	2014	2014
Well FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
Month FE	No	Yes	No	Yes
Production & Price Controls	No	Yes	No	Yes
Weather Controls	No	Yes	No	Yes

Notes: The dependent variable is the well-level flaring rate. In all specifications, time is specified as production time, or time since the well began producing gas and oil. Treated wells are those completed after 2015 and have produced for more than four months. Control wells are those completed in 2014. Separate treatment effects are specified for firms with different baseline flaring rates specified above. Baseline flaring rates are calculated as average firm-wide flaring rates in 2013 for columns (1)-(2) and in 2014 for columns (3)-(4). Standard errors are clustered at the well level. *, **, *** denotes significance at the 10%, 5%, and 1% level.

A Appendix

To Be Completed