Secrecy Rules and Exploratory Investment: Theory and Evidence from the Shale Boom *

Thomas R. Covert †
Richard L. Sweeney ‡

November 2022

Abstract

We analyze how information disclosure policy affects investment efficiency in non-cooperative settings with information externalities. In a two-firm, two-period model, we characterize equilibrium behavior under policies which disclose whether investment returns exceed a predefined level. These policies include complete secrecy, in which players only observe rival actions, as well as full disclosure, in which players also perfectly observe rival returns. With less disclosure (higher disclosure thresholds), there is less free riding, but additional losses from incomplete information aggregation. We characterize the surplus maximizing disclosure threshold in this environment, and show how it depends on firms' patience. We then apply the model to the early years of the shale boom in Pennsylvania and West Virginia, which at the time were governed by complete secrecy and full disclosure, respectively. We find that full disclosure would have maximized surplus in both states, generating 49% and 160% more value than complete secrecy.

JEL Codes:

Keywords:

*Funding provided in part by the Initiative on Global Markets at the University of Chicago Booth School of Business. Both authors declare they have no interests, financial or otherwise, that relate to the research described in this paper, nor do they have any current ties, directly or indirectly to the energy industry. We thank Mehmet Ekmekci, Alex Frankel, Michael Grubb, Ryan Kellogg, Jacob Leshno, and Pete Maniloff for comments. Yixin Sun, Eric Karsten, and Christoph Schlom provided excellent research assistance.

†University of Chicago Booth School of Business and NBER, Thomas.Covert@chicagobooth.edu.

‡Boston College, sweeneri@bc.edu.

1 Introduction

In non-cooperative environments, correlated outcomes from risky investment opportunities generate information externalities: one player's investment choice can change another player's beliefs about the returns to investment. Oil and gas exploration with decentralized landownership is a classic example of this phenomenon. Once one firm drills its land, neighboring firms can learn that this investment happened, possibly learn the returns to it, and, armed with this knowledge, make a less risky investment decision. The potential for firms to "free ride" on other's exploratory effort has been shown to generate costly delay, suboptimal information acquisition and inefficient sequencing of investment choices (Hendricks and Kovenock, 1989). Although this paper is about oil and gas exploration, the underlying economic forces we study generalize to other innovative settings, like uncertain demand (Chamley and Gale, 1994; Rob, 1991), real estate investment (Grenadier, 1999), and pharmaceutical development (Krieger, 2021).

A common response to the prospect of free riding on innovative efforts is to allow firms to keep secrets (Friedman et al., 1991). For example, in oil and gas exploration, many governments keep the production information of newly drilled wells that they collect for tax purposes confidential. By delaying or eliminating the possibility of observing rival outcomes, strict secrecy reduces free riding. However, secrecy can exacerbate losses from incomplete information aggregation. During a secrecy period, some wells which shouldn't be drilled (based on all available information) will be, while other opportunities which are profitable will go unexploited.

In this paper, we study the net effect of these forces, and ask whether confidentiality laws improve investment efficiency. We begin with a theoretical analysis of an investment game where the regulator sets the amount of information that is disclosed to rivals. The two extremes of this policy decision are "complete secrecy," under which firms who wait learn only whether or not their rival invested, and "full disclosure," under which firms whose rivals invest also learn the outcome of that investment. Our analysis generates predictions about the relationship between the amount of disclosure and the level, timing, and efficiency of investment. In the second part of the paper, we apply this model to rich oil and gas exploration data in the Appalachian shale basin where, at the start of the recent shale boom, Pennsylvania had a policy of complete secrecy, while neighboring West Virginia had a policy of full disclosure. We fit our theoretical model of each regime to the data, estimate the underlying primitives for this setting, and simulate outcomes under counterfactual disclosure policies to identity the optimum.[1]

[1]In this paper, we only consider the private benefits and costs associated with investment. In the case of

To model strategic responses to secrecy policy, we build on the social learning model of Hendricks and Kovenock (1989) (henceforth HK), who study the behavior of competing oil firms under full disclosure. The HK model features two firms, each endowed with a noisy signal of the common return to drilling, as well as a two period investment opportunity. The environment they study has a full disclosure (FD) information policy because a firm who waits will learn whether its rival waited, as well as the true return to drilling if its rival drills first. The unique symmetric Bayes-Nash equilibrium of this game involves cutoff strategies: firms with high signals drill in the first period, while firms with lower signals wait. If one firm drills early, and the other waits, the waiting firm learns the returns to drilling and makes an efficient investment decision. In this model, HK establish that equilibrium investment behavior is less efficient than what a planner endowed with both firms' signals can accomplish by coordinating their efforts. A key source of this inefficiency is the additional waiting that strategic firms do that a planner does not.

Absent centralization, one way to reduce the amount of strategic waiting firms do is to reduce the amount of information firms can learn from their rivals. We demonstrate this formally by extending the HK model to environments with less information disclosure. Specifically, we model a policy where, if one firm drills in the first period and its rival waits, the rival firm observes only whether the returns to drilling exceed a predefined threshold. A threshold of zero is equivalent to full disclosure, and a threshold of infinity is equivalent to complete secrecy.[2] Higher disclosure thresholds are less informative, because learning that returns are lower than a high threshold does not necessarily mean that the returns are negative. We show that the amount of waiting is decreasing in the disclosure threshold, up to a point. However, at very high disclosure thresholds, further increases in the threshold have no impact on waiting. We call the disclosure threshold at which this phenomenon begins to occur maximum non-distortionary disclosure (MND).

Having established that regulators can reduce waiting by coarsening information in this fashion, we then ask whether doing so increases surplus. In HK, waiting is "bad" because of discounting, but also because the more firms wait, the less information they generate for their rivals to learn from. Higher (less informative) disclosure thresholds do cause firms to drill earlier, and thus increase the amount of information generated. However, at high thresholds, this information is less valuable to rivals, making the net effect ambiguous. Our main result

oil and gas extraction however, investment also generates significant external costs. Truly "optimal," social welfare maximizing, disclosure policy in this context would incorporate these social costs. Nevertheless, for ease of exposition, throughout the text we refer the policy that maximizes private payoffs as "optimal," and private surplus as "welfare."

[2]Our analysis also allows for negative disclosure thresholds, but for brevity, our discussion here focuses on positive disclosure thresholds.

(Theorem 1) shows that for disclosure thresholds above a critical point, raising the threshold necessarily reduces welfare. This critical threshold is always lower than MND, so complete secrecy generically discloses too little information, and is never the optimal disclosure policy. In contrast, it is possible for the critical threshold to coincide with full disclosure. When the critical threshold is higher than FD, the optimal disclosure policy lies somewhere between FD and the critical disclosure threshold.

In real world settings, including our empirical application of mineral exploration, we only observe complete secrecy or full disclosure. Motivated by this, we also examine the optimal policy choice constrained to these two extremes. Although our partial disclosure results establish that CS is never the unconstrained optimal disclosure policy, while FD sometimes is, that does not necessarily imply that FD always delivers higher surplus than CS. Indeed, in section 2 we provide numerical examples in which firms in a CS environment do better, on average, than they do under FD. We then provide a sufficient condition which guarantees that FD does perform better, and show that this condition is equivalent to a patience requirement: sufficiently impatient firms always prefer full disclosure to complete secrecy.

Our inspiration for this theoretical analysis comes from differences in oil and gas disclosure policy across US states. Every oil and gas regulator requires firms to report aspects of the drilling and production process, for public safety and taxation purposes. However, different regulators have different rules for publicizing these reports. We focus on disclosure policies among oil and gas regulators in the Appalachian basin. Prior to 2011, Pennsylvania allowed firms to request confidential treatment of all of their reports (production and engineering) for 5 years. As this is the typical length of a mineral lease, this policy was effectively one of complete secrecy. In neighboring West Virginia, confidential treatment expired at after 1 year, which effectively amounted to a policy of full disclosure.

Our empirical analysis employs data on mineral leases and drilling outcomes in these two states from the start of the shale boom until the end of 2010, when Pennsylvania abruptly revoked its confidentiality law and adopted a full disclosure policy. In a descriptive analysis of investment behavior and outcomes, we focus on a narrow corridor around the border between the two states, within which the underlying mineral resources are essentially the same. We show that firms in Pennsylvania invest earlier and more often than in West Virginia, consistent with there being less waiting under complete secrecy. However, we also show that firms in West Virginia get substantially more output per well, consistent with the drilling program under full disclosure being more efficient.

While these cross-sectional differences are consistent with our theoretical predictions comparing FD and CS, they could also be driven by other factors that change at the border.

In order to hold these factors fixed and analyze counterfactual policy within each state, we estimate a structural econometric version of the theoretical model. We divide each state into square-mile drilling opportunities representing half of a latent two-player waiting game. This allows us to write the likelihood of observed drilling decisions in each grid as a function of unknown cost and signal distributions in each state, which we can maximize to estimate those primitives. We recover investment costs of approximately \$4-\$5 million per well, which is in line with estimates produced by the Energy Information Administration during this time. In each state, we find that the model corresponding to the disclosure policy on the books rationalizes the data better than a model imposing the alternative, although this test is imprecise in West Virginia where our sample is small.

Having estimated these primitives, we compute expected welfare over the full range of partial disclosure thresholds, holding other confounding factors fixed. In both states, we find that full disclosure is optimal. Compared to complete secrecy, full disclosure generates 49% and 160% gains relative to complete secrecy in Pennsylvania and West Virginia, respectively. We also identify the maximum nondistortionary disclosure level in each state, and find that the welfare of maximum nondistortionary disclosure is 85% and 63%, respectively, of full disclosure.

This paper contributes to the literature on the role of information in strategic decision-making in the oil and gas industry. As mentioned above, Hendricks and Kovenock (1989) provide a theoretical model of social learning. While that paper shows that full disclosure generates inefficiencies relative to planner behavior, we show that complete secrecy could substantially exacerbate those losses. Hendricks and Porter (1996) study investment on offshore wildcat tracts in the Gulf of Mexico, and find the patterns match a non-cooperative "war of attrition" game. Lin (2013) revisits that setting, and incorporates extraction externalities.

In a related paper, Hodgson (2021) also studies optimal disclosure policy in non-cooperative mineral exploration. Whereas we are interested in the question of *what* information the regulator chooses to disclose, Hodgson (2021) asks *when* a regulator should release information in a full disclosure regime. In an empirical model of offshore drilling in the United Kingdom, which has a five year secrecy period, he finds that revenue would be maximized under a confidentiality period half that long. An important assumption in that paper is that, during the secrecy period, firms not only fail to observe the returns to rival investment, but even the fact that any investment occurred. In contrast, the partial disclosure model we develop takes as its starting point the observation that, in settings like mineral exploration, real estate development, and pharmaceutical trials, the act of investment is practically un-hideable.

Outside of mineral exploration, the economic theory literature has also explored the equilibrium effects of investment disclosure in games with information spillovers.[3] This literature mostly follows the assumptions laid out in Keller et al. (2005), with a long or infinite horizon game, an intensive margin of investment, and initially common beliefs about a binary state of the world. Whether disclosure of investment outcomes helps our hurts equilibrium payoffs seems to depend on details specific to the problem. For example, Rosenberg et al. (2013) finds that games with publicly disclosed outcomes deliver higher equilibrium payoffs than similar games with no outcome disclosure, while Heidhues et al. (2015), which allows for communication between players, comes to the opposite conclusion. In a contest setting, Halac et al. (2017) find that no outcome disclosure is best whenever the social value of an innovation is sufficiently high, similar to the team production results in Campbell et al. (2014). In contrast, our motivating policy example and empirical application is oil and gas exploration, which features binary investment choices, over an inherently short finite horizon, and, crucially, with outcomes that occur over a continuum. Moreover, our analysis envisions a key role for asymmetric beliefs before the game even starts, for which the existing literature on oil and gas exploration Hendricks and Porter (1996) has found convincing evidence.

Finally, this paper also contributes to a literature on the learning during the shale boom. Covert (2015) estimates the extent to which firms learned about fracking input choices. Steck (2022) estimates whether the possibility of this learning lead to free riding. Agerton (2020) considers the way in which selection into drilling resources of heterogeneous quality biases estimates of productivity gains. Here we document considerable uncertainty about the location of shale resources at the start of the boom, and highlight the role that government disclosure policy played in identifying efficient investment opportunities in this uncertain environment.

2 A model of incomplete social learning

Our starting point is the model studied in Hendricks and Kovenock (1989), in which two firms each have a two period mineral lease. Each firm can choose to drill in the first period, drill in the second period, or not drill at all and allow the lease to expire. The actual returns to drilling are $\pi(X)$, where $\pi(\cdot)$ is a monotonically increasing function, and $X \in (0, \infty)$ is a common, but unknown, resource quantity. Drilling is profitable whenever X is larger than a known value x^*, and is unprofitable otherwise.

To decide whether or not to drill, each firm observes a private signal $S \in (-\infty, +\infty)$ about the value of X. Conditional on the true state of the world ($X = x$), the signals are

[3] Instead of outcome disclosure, Bonatti and Hörner (2017) studies the role of action disclosure, assuming outcomes are always disclosed if investment occurs. In this environment, welfare can increase or decrease in response to making actions observable, depending on whether the information that actions convey reflects "good" or "bad" news about the unobserved state of the world.

independently drawn from some probability distribution $\Pr(S \leq s \mid X = x) = F(s \mid x)$, which satisfies the monotone likelihood ratio property (MLRP). Thus, the probability that X is high when a firm observes a high signal is larger than the probability that X is low when it observes that signal, and *vice versa* (Milgrom, 1981). The analogous distribution of X, conditional on a firm's observation of a signal s, $\Pr(X \leq x \mid S = s) = H(x \mid s)$, can be obtained using Bayes' rule, and its density is $h(x \mid s)$. We assume that $h(x \mid s) > 0$ and $0 < F(s \mid x) < 1$ for all s and x. Though the value of X is common to both leases, there is no common pool problem, so if one firm drills, it does not affect the quantity of resources available for the other firm to recover. Both firms have a common discount factor $0 < \delta < 1$.

A key informational assumption in the HK model is that the true value of X is uncertain until one (or both) firms drill, and if one firm drills its lease in the first period, its rival can observe X perfectly in the second period. This is the full disclosure information policy we described above. HK establish that the unique symmetric Bayes-Nash equilibrium of this game satisfies a cutoff property: if a firm receives a signal at or above a cutoff level, it drills in the first period, and otherwise it waits. If a firm waits, but its rival drills in the first period, this laggard firm can "free-ride" and make an efficient drilling choice in the second period, because X is revealed by its rival's behavior. If both firms wait, then they make single-agent optimal drilling decisions, updating their beliefs about X with knowledge that their rival also had a signal lower than the cutoff level.

We study the effects of a broader class of information disclosure policy in this environment: if one firm drills in the first period and its rival waits, the rival learns that investment occurred, but only learns whether the true value of X is above ("good" news) or below ("bad" news) a pre-defined threshold Z. If the rival also waits, we call this "no" news. If $Z = x^*$, this is equivalent to the HK model, a full disclosure environment. If $Z = 0$ or $Z = \infty$, this is a complete secrecy environment. Intermediate values of Z correspond to some disclosure, but not as much as in a full disclosure environment.

Complete secrecy is a reasonable approximation to the information environment in which oil and gas exploration firms operate, in jurisdictions with strong confidentiality policies. In this industry, it is impossible for a firm to hide the fact that it decided to drill, because drilling requires a visibly large piece of capital (a drilling rig) which will sit on the firm's lease for weeks or even months. In other contexts, this informational assumption is meant to capture the fact that real investment (construction in real estate or clinical trials in pharmaceuticals) is either physically hard to hide, or regulators mandate its public disclosure. Thus, to the extent that cutoff equilibria exist in this modified game, a firm who decides to invest indirectly informs its rival that its signal is above some threshold.

In the analysis that follows, we'll write $\mathbb{E}\left[g(X) \mid s\right] = \int_0^\infty g(x)h(x \mid s)dx$ for any function

$g(\cdot)$ of X and any single signal realization s. We also re-state two existing results which we make extensive use of in the analysis that follows.

Lemma 1. *Adapted from Milgrom (1981), Proposition 4. For any signals s, t and u, and any non-decreasing function $g(x)$:*

$$\frac{\mathbb{E}\left[g(X)(1-F(t\mid X))\mid s\right]}{\mathbb{E}\left[(1-F(t\mid X))\mid s\right]} > \frac{\mathbb{E}\left[g(X)F(u\mid X)\mid s\right]}{\mathbb{E}\left[F(u\mid X)\mid s\right]}$$

As applied to the setting here, a player with signal s who knows that its rival's signal is greater than t has a larger expected return to drilling than a player with the same signal who knows that its rival's signal is less than or equal to u. This is true for any values of t and u. That is, finding out that your rival's signal is below some threshold is worse news than finding our that your rival's signal is above some other threshold.

Lemma 2. *Adapted from Hendricks and Kovenock (1989), Lemma 2 and Karlin and Rubin (1956), Lemma 1. Suppose $\phi(x)$ is a single crossing function, crossing zero at $\widehat{x}$. Then $\mathbb{E}\left[\phi(x)\mid s\right]$ is a single crossing function in the signal s, so that there is a signal s^* at which $\mathbb{E}\left[\phi(x)\mid s^*\right]=0$. If $\phi(x)$ is negative below $\widehat{x}$ and positive above $\widehat{x}$, and if $\mathbb{E}\left[\phi(x)\mid s\right]\geq 0$ then for all $s'>s$ then $\mathbb{E}\left[\phi(x)\mid s'\right]>0$. Alternatively, if $\phi(x)$ is positive below $\widehat{x}$ and negative above $\widehat{x}$, and if $\mathbb{E}\left[\phi(x)\mid s\right]\geq 0$ then for all $s''<s$, $\mathbb{E}\left[\phi(x)\mid s''\right]>0$.*

We will use this result to analyze how the net returns to drilling in the first period, relative to waiting, vary with the strength of a firm's signal and the information that may come from waiting.

2.1 Properties of symmetric equilibrium play

In this section we characterize symmetric equilibrium play in cutoff strategies. We first establish existence and uniqueness of a symmetric Bayes Nash Equilibrium in cutoff strategies for each Z. Then, we show how the equilibrium cutoff signal depends on Z, establishing that more disclosure (values of Z closer to x^*) lead to more waiting (higher cutoff signals). In the next section we study how welfare varies with the choice of Z.

In an equilibrium in cutoff strategies, firms with high signals, those at or above a cutoff value, drill in the first period, and firms with lower signals wait. After waiting, firms make single agent optimal decisions after observing their rival's choice (drill or wait) and possibly information about X. In what follows, we call $W(s,c,Z)$ the value of waiting for a firm with signal s. In this expression, c denotes the cutoff signal that the firm's rival is using and Z is the disclosure threshold. Mathematically, $W(s, c, Z)$ satisfies:

$$\begin{aligned} W(s, c, Z) &= W_{\text{good}}(s, c, Z) + W_{\text{bad}}(s, c, Z) + W_{\text{none}}(s, c), \text{ where} \\ W_{\text{good}}(s, c, Z) &= \max(0, \mathbb{E}[\pi(X)(1 - F(c \mid X))(X > Z) \mid s]) \\ W_{\text{bad}}(s, c, Z) &= \max(0, \mathbb{E}[\pi(X)(1 - F(c \mid X))(X \leq Z) \mid s]) \\ W_{\text{none}}(s, c) &= \max(0, \mathbb{E}[\pi(X)F(c \mid X) \mid s]) \end{aligned}$$

The value of waiting is the value of optimally acting (drilling or not) after learning what the rival firm did and whatever information is available if the rival drilled. If the rival drilled, which happens with probability $1 - F(c \mid X)$, the news may be "good" $(X > Z)$ or "bad" $(X \leq Z)$. If the rival also waited, which happens with probability $F(c \mid X)$, there is no news. We define the value of drilling under signal s as $D(s) = \mathbb{E}[\pi(X) \mid s]$.

We say s is a best response to c if a firm with signal s is indifferent between drilling and getting the discounted expected value of waiting, so that $D(s) = \delta W(s, c, Z)$. The signal v is a symmetric equilibrium if $D(v) = \delta W(v, v, Z)$, and in what follows, we always refer to v as a symmetric equilibrium first period cutoff signal. When comparing equilibrium play for different values of Z, we index the cutoff signals by Z to indicate that $v(Z)$ is an equilibrium cutoff for a partial disclosure game with disclosure threshold Z.

Proposition 1. *There is a symmetric equilibrium in cutoff strategies for any Z.*

Proof. By our assumption that $X > 0$ it is without loss of generality to assume $Z \geq 0$. First, observe that no firm with $D(s) < 0$ can drill in the first period, and $W(s, c, Z) \geq 0$ for all s, so if s is a best response to c, then $D(s) > 0$. By Lemma 1, we must therefore conclude that $\mathbb{E}\left[\pi(X)(1 - F(c \mid X)) \mid s\right] > 0$ at any best response s. This, in turn, implies that $W_{\text{good}}(s, c, Z) > 0$ at any best response s, so good news drilling is always possible for a best response signal. Second, observe that it cannot be possible for a best response signal to find it profitable to drill at both no news and bad news, for if it was, then $W(s, c, Z) = D(s)$, but our best response condition requires $D(s) = \delta W(s, c, z)$ for $\delta < 1$.

To find the best response to c, we first locate the unique and finite signal s_g which satisfies $0 = D(s_g) - \delta W_{\text{good}}(s_g, c, Z)$. This is possible because we can write this indifference condition as $0 = \mathbb{E}\left[\phi_1(X) \mid s_g\right]$, where $\phi_1(x) = \pi(x)\left(1 - \delta\left(1 - F(c \mid x)\right)\mathbb{I}(x > Z)\right)$ is a single crossing function which is negative for $x < x^*$ and positive for $x > x^*$, so Lemma 2 applies. If both $W_{\text{bad}}(s_g, c, Z) = W_{\text{none}}(s_g, c) = 0$, then s_g is the unique best response to c, because $W(s_g, c, Z) = W_{\text{good}}(s_g, c, Z)$.

If $W_{\text{bad}}(s_g, c, Z) > 0$, so that bad news drilling is possible for s_g, then $D(s_g) < W(s_g, c, Z)$ and there must be a unique $s_b > s_g$ at which $D(s_b) = \delta W(s_b, c, Z)$. This follows from the fact that we can write the indifference condition as $0 = \mathbb{E}\left[\phi_2(X) \mid s_g\right]$, where $\phi_2(x) = \pi(x)\left(1 - \delta\left(1 - F(c \mid x)\right)\right)$ is another single crossing function, so again Lemma 2 applies. In this case s_b is the unique best response to c.

Alternatively, if $W_{\text{none}}(s_g, c) > 0$, so that no news drilling is possible for s_g, then there is a unique $s_n > s_g$ at which $D(s_n) = \delta W(s_n, c, Z)$. Again, this is because we can write the indifference condition as $0 = \mathbb{E}\left[\phi_3(X) \mid s_g\right]$, where we define $\phi_3(x) = \pi(x)\left(1 - \delta\left(F(c \mid x)\mathbb{I}(x \leq Z) + \mathbb{I}(x > Z)\right)\right)$, which is also a single crossing function. In this case s_n is the unique best response.

Thus, there is a unique best response to c and it is finite for all c. Moreover, both $D(s)$ and $W(s, c, Z)$ are continuous functions in all of their arguments, so the best response must be continuous in c. Because the best response function is continuous and finite, it must cross the 45 degree line at some point. When it does, we have found signal v satisfying $D(v) = \delta W(v, v, Z)$. □

Next, we establish that there is a single symmetric equilibrium for each disclosure threshold and that the first period cutoff is increasing in the disclosure threshold for $Z < x^*$, non-decreasing for $Z \geq x^*$, and is eventually flat for larger values of Z. This formally establishes that more disclosure (Z closer to x^*, the point of profitability) always results in more waiting (a higher first period cutoff signal).

Proposition 2.

1. *The symmetric equilibrium is unique.*

2. *There is a disclosure threshold $\widetilde{Z}$ at which bad news drilling is possible in equilibrium, and bad news is possible for all $Z > \widetilde{Z}$. $\widetilde{Z}$ is defined by the value of Z which satisfies* $0 = \mathbb{E}\left[\pi(X)\left(1 - F(v(Z) \mid X)\right)\mathbb{E}(X \leq Z) \mid v(Z)\right]$.

3. *The symmetric equilibrium cutoff is increasing in Z for $Z < x^*$, decreasing in Z for $x^* < Z < \widetilde{Z}$, and is constant in Z for $Z \geq \widetilde{Z}$.*

Proof. For item 1, suppose that there are multiple symmetric equilibria. In that case, there must be at least 3 intersections of the best response function with the 45 degree line. Denote these intersections as $v_1 < v_2 < v_3$, with v_1 intersecting from above, v_2 intersecting from below, and u_3 again intersecting from above. With these equilibrium points, the best response function must be downward sloping (in c) at v_1, upward sloping at v_2, and again downward sloping at v_3.

To assess the feasibility of this assumption, we compute the slope of the best response function with respect to c by totally differentiating the indifference condition, rearranging, and evaluating at a symmetric equilibrium, obtaining $\frac{dv}{dc} = \frac{\delta W_c(v,v,Z)}{D'(v) - \delta W_s(v,v,Z)}$. The denominator of this expression must be positive if v is the indifference signal for a cutoff strategy, so the sign of this expression depends entirely on the sign of $W_c(v, v, Z)$. Under our assumption that there are multiple equilibrium points, we must have $W_c(v_1, v_1, Z) < 0$, $W_c(v_2, v_2, Z) > 0$, and $W_c(v_3, v_3, Z) < 0$. To check the feasibility of these multiple equilibrium points, it is necessary consider the location of Z relative to x^*, whether no-news drilling is possible for the cutoff signal, and if no-news drilling is impossible, whether bad news drilling is possible.

First, assume that bad news drilling is impossible, and begin with the case where $Z > x^*$. If no news drilling is impossible for the cutoff signal in the v_2 equilibrium, then $W_c(v_2, v_2, Z) = -\mathbb{E}\left[\pi(X)f(v_2 \mid X)\mathbb{I}(X > Z) \mid v_2\right]$. This expression is negative for any value of v_2, so it cannot be possible that $W_c(v_2, v_2, Z) > 0$, and we reach a contradiction. If no news drilling is possible in the v_2 equilibrium, then it must also be possible in the v_3 equilibrium, by Lemma 1. Then for both $v = v_2$ and $v = v_3$ we have $W_c(v, v, Z) = \mathbb{E}\left[\pi(X)f(v \mid X)\mathbb{I}(X \leq Z) \mid v\right]$. If $W_c(v_2, v_2, Z) > 0$, then so is $W_c(v_3, v_2, Z)$, by Lemma 2. Note that $W_c(v_3, v_2, Z) = W_c(v_2, v_3, Z)$. If $W_c(v_2, v_3, Z) > 0$, then another application of Lemma 2 implies $W_c(v_3, v_3, Z) > 0$. This violates our assumption that $W_c(v_3, v_3, Z) < 0$, so we again reach a contradiction.

Next, continue to assume bad news drilling is impossible, and now examine the case where $Z < x^*$. If no news drilling is impossible for v_2, then by Lemma 1 it is also impossible for v_1. So, for both equilibrium points we have $W_c(v, v, Z) = -\mathbb{E}\left[\pi(X)f(v \mid X)\mathbb{I}(X > Z) \mid v\right]$. We can write the condition that $W_c(v_1, v_1, Z) < 0 < W_c(v_2, v_2, Z)$ as $\mathbb{E}\left[\pi(X)F(v_2 \mid X)\mathbb{I}(X > Z) \mid v_2\right] < 0 < \mathbb{E}\left[\pi(X)F(v_1 \mid X)\mathbb{I}(X > Z) \mid v_1\right]$. If the left half of this expression is negative, then by the same logic employed in the previous paragraph, the right half must be negative as well, a contradiction. If no news drilling is possible for v_2 then it must also be possible for v_3 by Lemma 1. In that case, for both $v = v_2$ and $v = v_3$ we have $W_c(v, v, Z) = \mathbb{E}\left[\pi(X)f(v \mid X)\mathbb{I}(X \leq Z) \mid v\right]$ and this expression is negative for any v, so its not possible for $W_c(v_3, v_3, Z) > 0$, delivering another contradiction.

Finally, assume bad news drilling is possible, which requires both that $Z > x^*$ and that no news drilling is impossible. Then for any equilibrium v we have $W_c(v, v, Z) = -\mathbb{E}\left[\pi(x)f(v \mid X) \mid v\right]$. If we have multiple equilibria with $v_1 < v_2 < v_3$ then we must have both $\mathbb{E}\left[\pi(X)f(v_2 \mid X) \mid v_2\right] < 0 < \mathbb{E}\left[\pi(X)f(v_1 \mid X) \mid v_1\right]$ and $\mathbb{E}\left[\pi(X)f(v_2 \mid X) \mid v_2\right] < 0 < \mathbb{E}\left[\pi(X)f(v_3 \mid X) \mid v_3\right]$. However, $\mathbb{E}\left[\pi(X)f(v \mid X) \mid v\right]$ is proportional to the expected profits from drilling when both firms' signals are equal to v. Since $v_1 < v_2 < v_3$, it is not possible for drilling to be expected value negative when both firms' signals are v_2 but expected value positive at both v_1 and v_3, delivering a contradiction. Thus, multiple symmetric equilibria are not possible, and at the unique symmetric equilibrium we have $W_c(v, v, Z) < 0$.

For item 2, note that if there is no Z at which bad news drilling is possible for the equilibrium cutoff, then for $Z = \infty$ the waiting value is given by $W(v, v, \infty) = \max(0, \mathbb{E}[\pi(X)F(v \mid X) \mid v])$. This expression can never be larger than $D(v)$ for any v, so if there is no Z at which bad news drilling is possible, then there is no v at which the discounted value of waiting is equal to the value of drilling under $Z = \infty$, which is impossible if a symmetric equilibrium exists. Thus, there must be some $\widetilde{Z}$ at which bad news drilling delivers exactly zero expected profits for the equilibrium cutoff signal, and if $Z > \widetilde{Z}$, then $\mathbb{E}[\pi(X)(1 - F(v(Z) \mid X))\mathbb{I}(X \leq Z) \mid v(Z)] > 0$.

For item 3, differentiate the indifference condition with respect to Z and evaluate at an equilibrium, to obtain $\frac{dv}{dZ} = \frac{\delta W_Z(v,v,Z)}{D'(v) - \delta W_s(v,v,Z) - \delta W_c(v,v,Z)}$. As we saw above, the first two terms in the denominator sum to a positive value for any equilibrium v, and that $W_c(v, v, Z) < 0$. Thus, the denominator is always positive, and the sign of $\frac{dv}{dZ}$ is the same as the sign of $W_Z(v, v, Z)$. For a Z at which bad news drilling is impossible, this is equal to $-\pi(Z)(1-F(v \mid Z))h(Z \mid v)$, so that $W_Z(v, v, Z)$ has the opposite sign of $\pi(Z)$, positive for $Z < x^*$, and negative for $Z > x^*$. When bad news drilling is possible, $W_Z(v, v, Z) = 0$. □

Firms with signals lower than v will wait, and after observing what their rival did, make single agent optimal choices. Because the profit function is monotonic and signals satisfy the monotone likelihood ratio property, these second period choices also have cutoff structures, summarized below.

Proposition 3.

1. *After observing good news (the rival firm had a signal at or above v, and $X > Z$), all firms with signals at or above v^+ will drill, with v^+ defined by $0 = W_{good}(v^+, v, Z)$. If $Z \geq x^*$, then $v^+ = -\infty$, otherwise v^+ is finite, but smaller than v.*

2. *After observing bad news (the rival firm had a signal at or above v, and $X \leq Z$), all firms with signals at or above v^- will drill, with v^- defined by $0 = W_{bad}(v^-, v, Z)$. If $Z \leq x^*$ then $v^- = +\infty$, otherwise it is finite. v^- is smaller than v when $Z > \widetilde{Z}$.*

3. *After observing no news (the rival firm had a signal below v), all players with signals at or above v^0 will drill, with v^0 defined by $0 = W_{none}(v^0, v)$. v^0 can be larger or smaller than v.*

Proof. For item 1, if $Z \geq x^*$ then there is no uncertainty about whether drilling is profitable, so all signals that wait can profitably drill. If $Z < x^*$, it is possible for some signals to find drilling after good news to nevertheless be expected value negative. However, $W_{\text{good}}(s, v, Z)$ is the expected value of a single crossing function with respect to an MLRP density, so there must be a signal $s = v^+$ at which it equals zero (Lemma 2). Since we know that the cutoff signal v must be able to profitably drill after good news, v^+ cannot be larger than v.

For item 2, if $Z \leq x^*$, then there is also no uncertainty about whether drilling is profitable, and no signals can profitably drill after learning that $X < Z$. If $Z > x^*$, the logic used in the previous paragraph shows that v^- must exist. When the firm with the cutoff signal can drill after bad news without incurring losses (e.g., when $Z > \widetilde{Z}$), we know that $W_{\text{bad}}(v, v, Z) > 0$, so there must be a lower signal at which that expression reaches 0, and we know $v^- < v$. However, if the firm with the cutoff signal can't profitably drill after observing bad news, then no firms with lower signals can.

For item 3, note that v^0 must be *decreasing* in v, as the news that the rival firm did not drill is increasingly bad as v gets smaller. If the firm with the cutoff signal cannot profitably drill after no news, then no firm with a lower signal can, and it must be the case that $v^0 > v$. However, if the firm with the cutoff signal can profitably drill after no news, then some firms with lower signals can as well, and in that case $v^0 < v$. □

2.2 Welfare consequences of different disclosure thresholds

Although more disclosure (Z closer to x^*) results in more waiting, it also can deliver better information to firms who wait. Thus, the welfare consequences of more vs. less disclosure reflect a tradeoff between limiting free-riding by firms with higher signals and delivering useful information to firms with lower signals. Here we measure this tradeoff, and characterize what can be said about ranking different disclosure policies in expected value terms.

Theorem 1.

1. *There is a finite $\overline{Z}^0 < \widetilde{Z}$ such that for all $Z \geq \overline{Z}^0$, average firm welfare is decreasing in Z. If no news drilling is possible for v when $Z = x^*$, then $\overline{Z}^0$ is the unique $Z > x^*$ which satisfies $0 = \mathbb{E}\left[\pi(X)F(v(Z) \mid X) \mid v(Z)\right]$. If no news drilling is impossible for v at $Z = x^*$, then $\overline{Z}^0 = x^*$.*

2. *For every $Z < x^*$, there is an associated $Z' > x^*$ that delivers higher average player welfare.*

Proof. For item 1, first observe that because bad news drilling and no news drilling cannot simultaneously be profitable for an equilibrium cutoff signal, it must be the case that $\overline{Z}^0 < \widetilde{Z}$. Next, we establish that average player welfare is decreasing in Z for $Z \geq \widetilde{Z} > \overline{Z}^0$. Let $\widetilde{Z} \leq Z_1 < Z_2$. The first period cutoff is the same at Z_1 as at Z_2 (which we simply refer to as v), so the same firms drill early in both games. The difference in waiting values for firms with signals high enough to drill after bad news in both games will also be identical, as in either case those firms drill after "some" news but not after no news. For firms with signals too low to drill after bad news in Z_1 but high enough to drill after bad news in Z_2, the difference in waiting values can be written as $\mathbb{E}[\phi_4(X) \mid s]$, with $\phi_4(x) = -\pi(X)(1-F(v \mid X))\mathbb{I}(X \leq Z_1)$. The function $\phi_4(x)$ is a single crossing function which is positive for $x < x^*$, negative for $x^* \leq x < Z_1$, and zero otherwise. We know that $\mathbb{E}[\phi_4(X) \mid v] > 0$ by revealed preference, so by Lemma 2 for signals $v^-(Z_2) \leq s < v$ we must have $\mathbb{E}[\phi_4(X) \mid s] > 0$. Firms with signals s lower than $v^-(Z_2)$ won't do bad news drilling in either game. However, in the Z_1 game they'll sometimes drill when $Z_1 \leq x < Z_2$, whereas this drilling never happens in the Z_2 game. Because this is profitable drilling, waiting in the Z_1 game must be more profitable than waiting is in the Z_2 game for signals below $v^-(Z_2)$. Thus, all signals weakly prefer Z_1 to Z_2.

Next, we show that all players weakly prefer Z_1 to Z_2 for $\overline{Z}^0 \leq Z_1 < Z_2 \leq \widetilde{Z}$. By Proposition 2, we know $v(Z_1) > v(Z_2)$, so all signals at or above $v(Z_2)$ weakly prefer Z_1. Because bad news drilling is impossible for the equilibrium cutoff signal at both Z_1 and Z_2, the difference in waiting values for signals s below $v(Z_2)$ can be written as $\mathbb{E}[\phi_5(X) \mid s]$, with $\phi_5(x) = \pi(X)\left((1 - F(v(Z_1) \mid X))\mathbb{I}(X > Z_1) - (1 - F(v(Z_2) \mid X))\mathbb{I}(X > Z_2)\right)$. The function $\phi_5(x)$ is a single crossing function which is 0 for $x < Z_1$, positive for $Z_1 \leq x < Z_2$, and negative for $x \geq Z_2$. We know that $\mathbb{E}[\phi_5(X) \mid v(Z_1)] > 0$ by revealed preference, so by Lemma 2, for signals $s < v(Z_1)$ we must also have $\mathbb{E}[\phi_5(X) \mid s] > 0$. Thus, all signals weakly prefer Z_1 to Z_2, for $\overline{Z}^0 \leq Z1 < Z2$.

For item 2, because $v(Z)$ is continuous, increasing in Z for $Z < x^*$, non-increasing in Z for $Z > x^*$, and because $v(0) = v(\infty)$, then for any $Z < x^*$ there is a corresponding $Z' > x^*$ such that $v(Z) = v(Z')$. Z and Z' have the same first period cutoffs, so all signals at or above that cutoff get the same welfare under Z as Z', and any differences in welfare must be derived from signals s smaller than $v = v(Z) = v(Z')$. If $Z > 0$, then $Z' < \widetilde{Z}$ and bad news drilling is impossible for the cutoff signal in both Z and Z'. If $Z = 0$, given the above results, it suffices to compare it against $Z' = \widetilde{Z}$, at which bad news drilling earns zero profits for the cutoff signal. We can write the difference between waiting at Z' and waiting at Z as $\mathbb{E}[\phi_6(X) \mid s]$, with $\phi_6(x) = -\pi(X)(1 - F(v \mid X))\mathbb{I}(Z < X < Z')$. As before, the function $\phi_6(x)$ is a single crossing function that is 0 for x larger than Z' or smaller than Z, positive for $Z \leq x \leq x^*$ and negative for $x^* \leq x \leq Z'$. We know that $\mathbb{E}[\phi_6(X) \mid u] \geq 0$, since the discounted value of waiting equals the value of drilling for these signals, and this will be the same under Z and Z'. Thus, by Lemma 2, we must conclude that $\mathbb{E}[\phi_6(X) \mid s] > 0$ for $s < v$, and thus all signals weakly prefer Z' to Z. □

The first statement establishes that average player welfare is decreasing in Z for the range of disclosure thresholds in which no news drilling is impossible. The optimal disclosure threshold can be no larger than the critical value $\overline{Z}^0$, and if no news drilling is impossible at $Z = x^*$, so that $\overline{Z}^0 = x^*$, then its impossible everywhere, and full disclosure is necessarily optimal. The second statement establishes that disclosure thresholds *below* the point of profitability can never be optimal, as there are always equivalent disclosure thresholds *above* x^* with identical first period drilling incentives but strictly higher waiting values for the signals that wait. Thus, Theorem 1 partially characterizes the optimal choice of Z.

A useful corollary of Proposition 2 (part 3) and Theorem 1 (part 1) is that complete secrecy $Z = \infty$ can never be an optimal disclosure policy among the space of partial disclosure policies. Indeed, there is some of amount of disclosure ($Z = \widetilde{Z}$) which is effectively "free" in the sense that it maintains the low-free riding properties of complete secrecy, but allows some useful information to be disclosed. In our empirical application, we refer to $\widetilde{Z}$ as the maximum nondistortionary disclosure (MND) threshold, because it is the most that can be disclosed about X without increasing free-riding from the level attained in CS.

3 Full disclosure and complete secrecy

The previous section lays out a general theory of symmetric equilibrium play in environments with partial disclosure of outcomes. However, most real world settings that this framework could plausibly describe have one of just two disclosure policies, either exactly disclosing outcomes ($Z = x^*$), which we call "full disclosure" or disclosing nothing about outcomes ($Z = 0$ or, equivalently, $Z = \infty$), which we call "complete secrecy." Here we explore what our model has to say about the differences between these two commonly used disclosure policies.

First, a direct application of Proposition 2 tells us that there should be more early drilling, and thus less free-riding, in environments with complete secrecy (e.g., $Z = \infty$) than similar environments with full disclosure ($Z = x^*$). This matches the natural intuition that in a world with less to learn from free-riding, there are fewer incentives to do it in the first place.

Second, Theorem 1 shows that full disclosure is the optimal disclosure policy among all disclosure thresholds whenever no news drilling is impossible in a full disclosure environment. When this condition holds, it of course implies that full disclosure is better than complete secrecy. However, this condition is a strong one. As we demonstrate below, full disclosure has higher average welfare than complete secrecy under a weaker sufficient condition, that there isn't "too much" no news drilling.

Proposition 4. *If $v^0(x^*) \geq v(\widetilde{Z}) = v(\infty)$, then average player welfare is higher under $Z = x^*$ than any $Z \geq \widetilde{Z}$.*

Proof. When $Z \geq \widetilde{Z}$, we know all signals at or above $v(Z)$ weakly prefer x^* by revealed preference. Lower signals will wait in both games, and for those below $v(Z)$ but above $v^-(Z)$ the difference in waiting values is is $\mathbb{E}\left[\phi_7(X) \mid s\right]$, with the function $\phi_7(x) = \pi(x)\left((1 - F(v(x^*) \mid x))\mathbb{I}(x > x^*) -\right.$

This is a single crossing function which is positive for $X < x^*$ and negative for $X > X^*$. We know that $\mathbb{E}[\phi_7(X) \mid v(Z)] > 0$ by revealed preference, so by Lemma 2 we know all signals $s \in [v^-(Z), v(Z)]$ weakly prefer x^* to $\widetilde{Z}$.

Signals smaller than $v^-(Z)$ will not do bad news drilling, so for these values of s the difference in waiting values is $\mathbb{E}[\phi_8(X) \mid s]$, with the single crossing function $\phi_8(x) = \pi(x)((1 - F(v(x^*) \mid x))\mathbb{I}(x >$ This function is 0 for $X < x^*$, positive for $x^* \leq X < Z$, and negative for $X > Z$. Again, we know that $\mathbb{E}[\phi_8(X) \mid v^-(Z)] > 0$ by revealed preference, so by Lemma 2 all lower signals must also prefer x^* to Z. Thus, all signals weakly prefer x^* to $\widetilde{Z}$. □

As motivated above, full disclosure and complete secrecy are common information disclosure policies, so being able to rank them in welfare terms is useful. As we show next, the condition on signal values that allow us to rank the two policies is equivalent to a requirement that firms are sufficiently impatient.

Proposition 5. *There is always a discount factor $0 < \overline{\delta} < 1$ such that for all $\delta \leq \overline{\delta}$, the equilibrium cutoffs satisfy $v^0(x^*) \geq v(\widetilde{Z}) = v(\infty)$.*

Proof. In the limit as $\delta \to 0$, both $v(x^*)$ and $v(\widetilde{Z})$ converge to the single agent cutoff $\underline{s}$, defined by $\mathbb{E}[\pi(X) \mid \underline{s}] = 0$. No news drilling is clearly not possible for this cutoff signal, by Lemma 1, and so as $\delta \to 0$, $v^0(x^*) > v(\widetilde{Z})$. As $\delta \to 1$, $v(x^*)$ tends toward $+\infty$, while $v(\widetilde{Z})$ tends to the finite $\overline{s}$ which solves $\mathbb{E}[\pi(X) \mid \overline{s}] = \mathbb{E}[\pi(X)(1 - F(\overline{s} \mid X)) \mid \overline{s}]$. With an infinite first period cutoff, $v^0(x^*) \to \underline{s} < \overline{s}$, so at the other extreme, we know the condition is not satisfied. Next, observe that v is monotonically increasing in δ for any Z. To see this, differentiate the indifference condition and rearrange terms to obtain $\frac{dv}{d\delta} = \frac{W(v,v,Z)}{D'(v) - \delta W_s(v,v,Z) - \delta W_c(v,v,Z)}$. The numerator of this expression is positive as it must be equal to the value of drilling at the cutoff signal, and we've previously established that the denominator is positive for all equilibrium v. Finally, recall that $v^0(Z)$ is decreasing in $v(Z)$, which implies that $v^0(Z)$ is decreasing in δ. This, combined with the initial observation, shows that there must be some $\overline{\delta}$ at which $v^0(x^*) = v(\infty)$, and for all $\delta \leq \overline{Z}$ we have $v^0(x^*) > v(\infty)$. □

Proposition 4 provides a "signal-by-signal" result, in that every signal at least weakly prefers full disclosure to complete secrecy, and Proposition 5 shows that the condition which delivers this result is always attainable with a low enough discount factor. When firms are patient enough, and these conditions fail, there will be a range of signals which prefer *less* disclosure, not more. If this range is sufficiently small, average welfare will still be in favor of FD. However, it is possible to construct environments in which *ex ante* average equilibrium payoffs are higher under CS than FD. For example, suppose $\log X \sim N(0,1)$, signals are given by $s_i = \log X + \epsilon_i$, with $\epsilon \sim N(0, 4^2)$ and independent of X. Additionally, assume that the value of output is 1 and the fixed cost of drilling is equal to exp(0.5), so that the unconditional average well has zero profits. When $\delta \leq 0.80$, we have $v^0(x^*) \geq v(\widetilde{Z})$, so Proposition 4 applies. For all higher values of δ, $v^0(x^*) < v(\widetilde{Z})$. However, letting $V^g(s)$ denote the expected value of playing game g with signal s, it is still the case that $\mathbb{E}_s\left[V^{\text{Full Disclosure}}(s)\right] > \mathbb{E}_s\left[V^{\text{Complete Secrecy}}(s)\right]$ for values of δ up to 0.97. Once $\delta > 0.97$, complete secrecy deliver higher average welfare.

4 Empirical application

The early years of the shale boom provide a good setting to estimate the impact of disclosure policy in a high stakes environment with uncertain investment outcomes. While the existence of shale formations had been known since at least the 1980's, they were regarded as prohibitively expensive to exploit, and thus the exact levels of economically recoverable hydrocarbons at each location remained unknown. The rapid maturation of hydraulic fracturing at the turn of the century, combined with the emergence of horizontal drilling, abruptly changed this predicament, unleashing a wave of shale exploration and development.

We focus on the Appalachian shale basin, where two of the world's most active gas plays, the Marcellus and the Utica, underlie two states that until recently had starkly different disclosure policies: Pennsylvania and West Virginia (Figure 1).[4] In this section, we first review the institutional and policy background in these states. We then describe the data and relate it to the theoretical model presented above, and present some initial comparisons of investment across these states.

4.1 Background

Oil and gas investment involves two primary phases: *leasing*, in which firms acquire the right to explore from landowners overlying potential petroleum reservoirs, and *drilling*, during which firms decide whether to drill into land they have leased.

In the Appalachian shale basin, almost all mineral rights are privately held. Unlike the American west, these rights are delineated by irregular shapes, and their ownership is relatively dispersed. As described in Covert and Sweeney (2019), the private market for mineral leases is often informal and decentralized. As a result, the ownership of mineral leases is quite disaggregated, with many firms owning mineral leases in the same small geographic area. This creates a natural mechanism for relevant information to spill-over between firms. Once a firm acquires a mineral leases from a landowner, it has a finite amount of time to

[4] These shale formations also underlie Ohio. Ohio has a 6 month secrecy period, which is considerably shorter than the typical length of a mineral lease, so this also amounts to a full disclosure policy. However, there is essentially no shale exploration in Ohio prior to 2010, when our sample ends. We therefore exclude Ohio from all of our analysis.

Figure 1: Appalachian Shale Plays

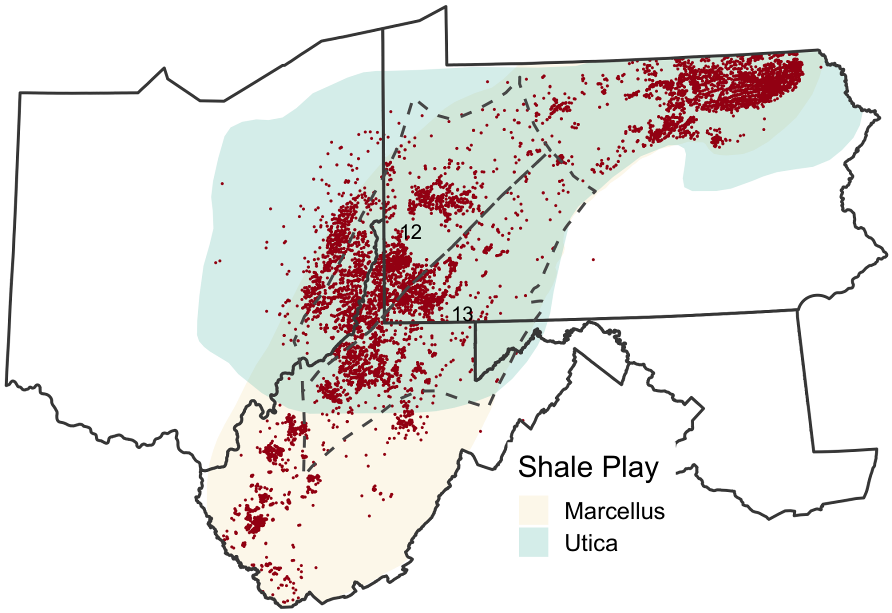

Shale well drilling (red dots) in Ohio, West Virginia, and Pennsylvania, between 2004 and 2020, relative to the extents of the Marcellus and Utica shale plays. The dashed lines give the boundaries of the USGS Gas play regions 12 and 13. For further details, see Figure 6 in the online appendix.

drill a well and establish whether it is productive. This "primary term" is generally three to five years.

In order to drill, firms must acquire a drilling permit from a state regulator, and the existence of these permits is generally public information. After receiving permission to drill, a firm will hire a drilling rig and hydraulic fracturing crew to visit its lease and develop the well. The physical scale of drilling and hydraulic fracturing is large enough that the decision to drill is also effectively public information. Finally, after a well is drilled and fracked, the firm will start producing from it, and regulators will collect records from the firm about what occurred during drilling and fracking (i.e., how long did it take, at what depths did the drilling rig encounter each formation, etc.). At this point firms are also required to report production outcomes to regulators.

The Marcellus and Utica shale basins span parts of both Pennsylvania and West Virginia, and in both of these states, firms are required to submit these drilling and fracking reports to regulators within 90 days. In West Virginia, regulators publicize these reports available after 1 year. Prior to 2011, Pennsylvania's regulators allowed for substantially more secrecy. All well-level records, including drilling reports and production outcomes, were kept secret for five years. In March 2010, the Pennsylvania State Legislature passed Senate Bill 297, a set of amendments to the original Pennsylvania Oil and Gas Act dating back to 1984. The new law provided for semi-annual, instead of annual reporting, and, crucially, repealed Pennsylvania's secrecy statute entirely. All existing production and drilling information held under the old secrecy rule was made public shortly after the law went into affect, in November 2010.

Apart from disclosure policy, these two states differ in another important dimension. West Virginia imposes a 5 percent tax on oil and gas revenue. Initially, Pennsylvania did not tax oil and gas production, and this difference confounds a simple comparison of activity across the two states when the secrecy policies differ prior to 2011.[5] In February 2012, it implemented an annual per-well fee, which is thought to increase overall costs by about 5 percent. The fact that this change came so quickly after the repeal of its secrecy policy also confounds a difference in difference estimation strategy.[6]

4.2 Data

We acquired data on leasing, drilling, and production from Enverus, a major commercial provider of information to the oil and gas industry. For each lease, Enverus provides a spatial representation of the covered minerals, the date it was signed, the firm it was signed with, and the length of the lease's primary term. Enverus collects this information from county courthouses, and their data covers most (but not all) counties overlying active portions of the Marcellus and Utica shale plays. In all the following analysis, we exclude counties not covered by Enverus.

We combine this lease data with Enverus' comprehensive drilling and production data. Enverus collects this information from state oil and gas regulators and organizes it into a standard format. The Enverus drilling data data covers about 90,000 wells drilled between 2000-2016. For each well, we observe its GPS coordinates, the date it was drilled, the firm that drilled it, the well's target formation (ie "Marcellus"), whether the well was a horizontal well (and thus likely to have received a hydraulic fracturing treatment) or a vertical well. For every well, we also observe a time series of all subsequent oil and gas production.

4.3 Sample construction

The model presented in section 2 envisions two identically sized leases, physically remote from other sources of informative drilling activity, each of which is just large enough to support a single well. While this theoretical setup provides a useful setting to study the role of incentives and information in mineral exploration, it is unfortunately far from reality. Landownership is quite dispersed in the Appalachian basin, and this fact is reflected in the

[5] Brown et al. (2020) estimate the impact of state hydrocarbon taxes on extraction.
[6] Black et al. (2018) study the impact of introduction of the well fee in Pennsylvania.

size distribution of our lease data. The median lease is just over 3 acres, well below the approximately 80 acres needed to drill a single horizontal well at the conventional length of 5,000 feet.

Because the majority of leases are far too small to be drilled on their own, oil and gas companies frequently agglomerate adjacent leases into a single drilling permit (called a "unit"), and it is these pools of leases which, in theory would match the unit of observation in the models of the previous section. However, unlike lease, drilling, or production data, there are no administrative pooling records in any state in the Appalachian basin, so it is not possible to measure outcomes at the unit level. In order to measure outcomes that likely arise from this pooling process without having data on it, we construct an artificial grid of 1-mile by 1-mile squares, each of which could potentially contain a drilling unit.[7] In all the analysis that follow, we measure outcomes at the grid level, treating each grid as a single player's manifestation from one version of the game described in section 2.

Having constructed potential drilling opportunities, we next identify grids that experience enough leasing activity to be drilled. In our main specification, we assume that a grid becomes "active," and thus can be drilled, once 33% of its area has been leased. Although 33% of a square mile is a bit more than twice as much land that is needed to drill a single well, as we mentioned above, lots of leasing in this context is dispersed, and so grids with exactly 80 acres leased will not necessarily have 80 contiguous acres leased. We view this kind of threshold assumption as a reasonable compromise between excluding grids that may indeed have enough contiguous acreage to be drilled but not enough to satisfy this screen, and including others that pass the screen but nevertheless don't have 80 contiguous acres. The results we report are very similar to those we obtain if we define a grid as active once it instead reaches 25% or 50% leased.

Most leases in this setting have 5 year primary terms, and thus will be active for 5 years after they reach the 33% leasing threshold. In grids that reach this threshold after 2005, firms will be able to drill during primary term after Pennsylvania's change in its information disclosure policy. As a result, drilling behavior in these grids is unlikely to be accurately explained by a single disclosure model. To avoid the complications that arise from this, in our structural analysis, we will focus exclusively on grids that reach the 33% threshold by the end of 2005, and measure drilling outcomes on each grid until the end of 2010. In early 2011, Pennsylvania posted the production information for all of its previously drilled wells online.

Finally, we will ultimately focus our attention on areas of Pennsylvania and West Virginia

[7] 1 square mile corresponds to a "section," or 640 acres. In many other states, sections serve as predefined drilling opportunities; for example, Louisiana (Herrnstadt et al., 2020).

that firms likely had a common set of pre-existing knowledge about. Prior to the start of the shale boom, the only publicly disclosed geological data characterizing the Marcellus and Utica shale basins was summarized in Charpentier et al. (1993). This U.S. Geological Survey report describes the results of the U.S. Department of Energy funded work done by the Eastern Gas Shales Program in the late 1970s and early 1980s, which was tasked with measuring the geological conditions and underlying resource quality of various American shale basins. The USGS report includes a map which divides the entire Appalachian basin into several regions within which geological conditions and resource quality were thought to be similar, which we recreate in Online Appendix Figure 6. We focus our analyses on the parts of Pennsylvania and West Virginia covered by two of these regions, "Gas Plays" 12 and 13, outlined in Figure 1. The two regions were estimated by Charpentier et al. (1993) to have similar average gas production potential, and they are the only regions which span both states.

Table 1 presents grid-level summary statistics of this data. In the first two columns, we include all grids which ever become active (33% leased) by 2020. Below the grid counts, we report summary statistics on drilling activity that occurs prior to 2011, when Pennsylvania changes its disclosure policy. We see that only 6% of Pennsylvania and West Virginia grids have any drilling by this point, and in both states most of this drilling is "late," occurring after 2008. Wells in Pennsylvania produce nearly three times as much output as wells in West Virginia and earn 2.5 times as much revenue.[8]

In the rightmost columns of the table, we restrict the sample grids to those that become active by the end of 2005, and which lie in the relevant, shared, gas regions presented in Figure 1. On this sample, which is approximately 10 percent of the larger population in the left-most columns, slightly more than 8 percent of grids drill by the end of 2010. However, in Pennsylvania, the majority of these grids drill "early," before 2009, while opposite is true in West Virginia. This relationship matches one of the core predictions of section 2: drilling happens earlier under full disclosure. Unlike the full population, in this sample, where underlying geological quality is likely similar across the two states, Pennsylvania wells produce less output and earn less revenue.

4.4 Border comparison

The summary statistics in the two rightmost columns of Table 1 suggest that Pennsylvania, which was governed by complete secrecy prior to 2011, drilled earlier but less efficiently than West Virginia, which was governed by full disclosure. While this is consistent with the

[8]We describe our process for cleaning the production data in Section 5, and our process for constructing prices in Online Appendix B.

Table 1: Grid summary stats

	All Grids		Main Sample	
	PA	WV	PA	WV
Number of square mile grids				
N	10322	2844	1227	284
Mean (sd)				
Drilled	0.063 (0.24)	0.058 (0.23)	0.083 (0.28)	0.081 (0.27)
Drilled Early	0.016 (0.12)	0.02 (0.14)	0.051 (0.22)	0.025 (0.16)
Wells	0.16 (0.85)	0.12 (0.63)	0.24 (1.1)	0.23 (0.98)
Output/ Well (bcf)	3 (3.2)	1.3 (1.5)	1.2 (1.2)	1.7 (1.6)
Price ($/mcf)	3.4 (0.7)	3.6 (1.1)	3.8 (0.86)	3.3 (0.48)
Revenue / Well (millon $)	9.3 (9.8)	4 (4.5)	4 (3.8)	5.4 (5)

States are divided into square mile grids. In the first two columns (labeled "All Grids"), we include all grids we ever observed at least 33% leased by the end of 2020. In the last two columns, the sample is restricted to grids that are 33% leased by 2005, overlying Gas Plays 12 and 13 in Figure 1. Drilling outcomes restricted to wells spud by the end of 2010. "Drilled Early" refers to wells spud by the end of 2008.

model's core prediction, the comparison is complicated by the fact that, while these states share the same underlying shale formations, *ex post* resource quality within these formations is not constant, and as a result may be different, on average, between the two states. In this section, we make comparisons between grids in Pennsylvania and West Virginia that are restricted to lie within a narrow bandwidth (10 miles) around the Pennsylvania border (see Online Appendix Figure 7). This restriction ensures that the underlying rock quality is similar across the two states.

Table 2 presents the results. In column 1, we include all grids within 10 miles of the Pennsylvania border that become "active" (33% leased) by the end of 2020.[9] Row 1 contains a simple projection of an indicator for whether the grid was drilled by the end of 2010 onto an indicator for the state of Pennsylvania. On this sample, Pennsylvania grids are 2 percentage points less likely to drill, but are 0.7 percentage points more likely to drill "early" (prior to 2009). In row three we perform a Poisson regression with the number of wells per square-mile as the outcome. Pennsylvania grids drill 11% fewer wells, but obtain 80% less output per grid.

In column 2, we restrict the sample to grids that became active by 2010. This cuts more than half the sample. However, within the set of grids sufficiently leased for mineral explo-

[9]Grids less than one mile away are excluded because the relevant information environment is not obvious; the outcomes in grids less than 1 mile south of the border may be knowable by firms less than 1 mile north of the border.

Table 2: Border Regression Results, 10 Mile Radius

		Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Linear	Drilled	-0.020	-0.012	0.015	0.016	0.079	0.086
		(0.008)	(0.013)	(0.013)	(0.012)	(0.030)	(0.033)
	Drilled Early	0.007	0.025	0.035	0.031	0.084	0.092
		(0.004)	(0.008)	(0.008)	(0.007)	(0.022)	(0.025)
Poisson	Wells	-0.111	0.234	0.664	0.741	0.892	1.175
		(0.243)	(0.246)	(0.244)	(0.253)	(0.397)	(0.401)
	Output	-0.797	-0.499	-0.020	0.047	0.238	0.414
		(0.295)	(0.296)	(0.295)	(0.311)	(0.438)	(0.429)
	Grids: PA	2232	993	708	708	242	242
	Grids: WV	1016	625	625	625	111	111
	Leased By	none	2010	2010	2010	2005	2005
	Matching Vars	None	None	Idx	Idx, Year	Idx	Idx, Year

Sample restricted to grids between 1 and 10 miles of the border. All outcomes discounted to the start of the lease sample, January 2003. Wells are restricted to shale wells drilled by the end of 2010. Drilled "early" refers to wells drilled by the end of 2008. Output refers to total discounted output within the grid. In models 3 through 6, the sample is first balanced across Pennsylvania and West Virginia using coarsened exact matching. The Matching Vars row indicates whether observations were matched spatially using their position along the Pennsylvania border ("Idx"), and the year the grid became active.

ration leading into the shale boom, the estimated difference in early drilling in Pennsylvania more than doubles. In this specification, Pennsylvania drills 23% more wells but obtains 50% *less* output. In column 3 we use coarsened exact matching (CEM, Iacus et al. (2012)) to match grids spatially based on their position along the Pennsylvania border (denoted "Idx" in the table). The motivation here is that, even in this narrow bandwidth, the underlying rock quality is (ex post) quite different at different points along the border, and we want to ensure balance across the two states in latent quality. Within this sample, Pennsylvania is more likely to drill across both periods, drilling more than 60% more wells than West Virginia, but obtaining weakly less output. In column 4 we match on both the border index and the year the grid reaches 33% leasing.

In column 5 of Table 2 we restrict the sample to grids that become active by the end of 2005. The motivation for this restriction is that, given typical lease term lengths, it is likely that we observe both early and late drilling before the Pennsylvania policy changes in 2011. Finally, in column 6 we match on both border index and lease year on this restricted sample. Across both of these matched samples, the incidence of early drilling and the number of wells per square mile is much higher. However, unlike the raw comparisons in Table 1,

Figure 2: Border radius sensitivity

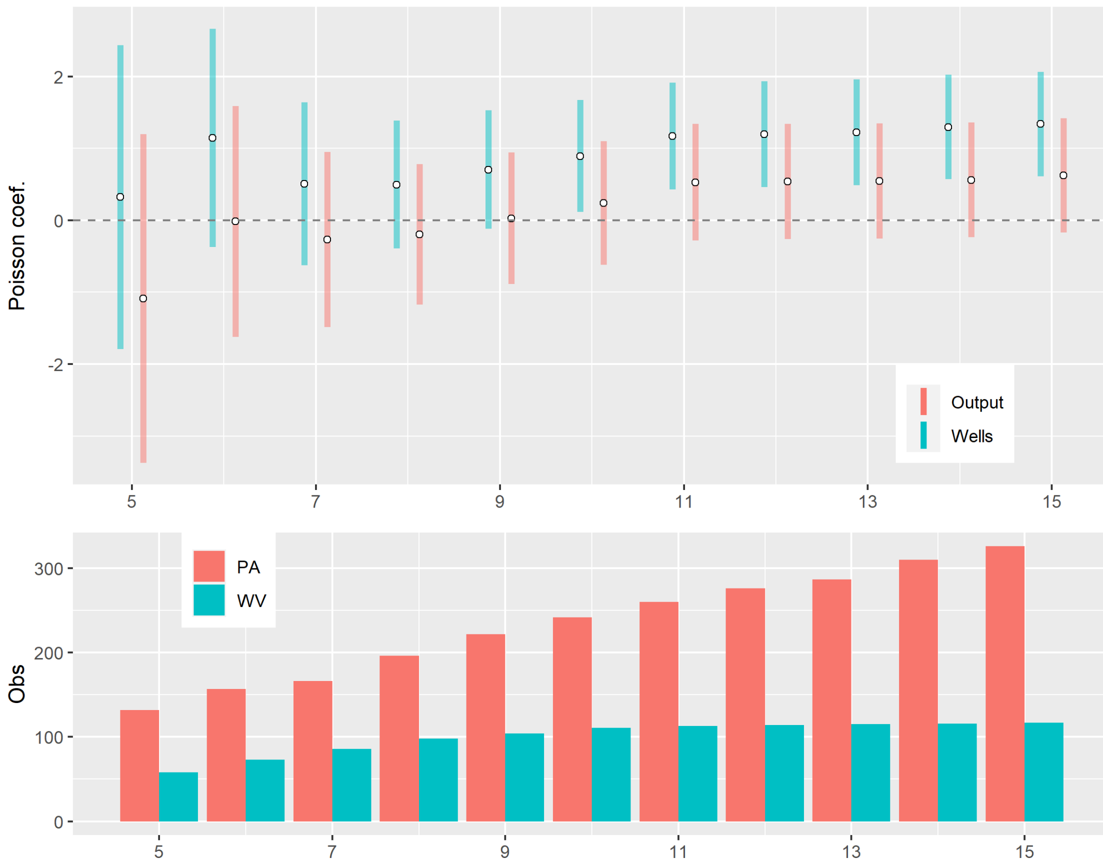

Pennsylvania wells are much less productive, consistent with the model's predictions.

In Figure 2 we explore the sensitivity of the Poisson results to the border radius imposed. The sample is restricted to grids leased by 2005, matched on the border index (corresponding to column 5 in Table 2). The bottom panel demonstrates the rapidly diminishing sample as we narrow the bandwidth. The top panel presents point estimates and confidence intervals for the well per square mile and output Poisson regressions. While the confidence intervals are wide, it is clear that Pennsylvania drills many more wells, but generates proportionally much less output than is implied by its additional drilling.

While these comparisons of drilling behavior and subsequent production provide suggestive evidence that information disclosure policy can affect real outcomes in a manner consistent with the model in section 2, they lack a traditional causal interpretation because of other differences in economic incentives between Pennsylvania and West Virginia. Although underlying resource quality is unlikely to meaningfully change at the border, other policies, aside from information disclosure, do. The most salient observable difference is taxes: West Virginia has a 5% tax on oil and gas production, as well as partially offsetting investment tax credits for drilling new wells. In contrast, during this time there were no taxes or tax credits associated with oil and gas production in Pennsylvania. These tax differences are likely to be economically meaningful, especially in settings like this one where drilling is relatively rare. It is also possible that key unobserved factors, like drilling costs and the precision of pre-drilling signals, may differ between the states.

5 Structural estimation

In this section, we estimate models of drilling under the information policy in each state, full disclosure in West Virginia and complete secrecy in Pennsylvania. We incorporate the important differences in production tax treatment into the underlying profit functions, and estimate the drilling costs and signal variances that would rationalize the outcomes we observe in each state. With estimates of these primitives in hand, we compute the average welfare under alternative disclosure policies, holding other factors fixed, to characterize the optimal disclosure policy in each state and quantify its benefits relative to other disclosure policies.

The data we use in estimation consists of the square mile grids in Pennsylvania and West Virginia which were active by the end of 2005, and which lie in either Gas Play 12 or Gas Play 13 from the Charpentier et al. (1993) map (see Online Appendix Figure 6). This sample is described in the right columns of Table 1. We assume that each square-mile grid constitutes a random half of a waiting game, so that the outcome in a grid represents one firm's drilling choice. Figure 3 presents a map of sample grids and whether drilling is observed by the end of 2010.

As we describe below, our model delivers probabilities for each of these outcomes, *conditional on the true resource quality* X in a grid. For grids that are drilled during our time period, X is observed.[10] However, for grids that aren't drilled until later, or not at all, this information is missing. To overcome this, we'll do two things. First, for all drilled grids, including those drilled after our time period ends, we'll adjust the observed X values for the progress in technology that occurred. We do this by regressing each well's realized output onto a dummy variable for horizontal vs. vertical drilling, a logarithmic term for the horizontal length of the well that was drilled, and fixed effects for the year in which drilling takes place.[11] We then compute the predicted output of that well, assuming it was drilled in

[10] Technically, we observe a fraction of eventual production, as most drilled wells in this sample are still producing to this day. We convert this incomplete history of production into a discounted forecasted value using an engineering model. See Online Appendix B for details.

[11] We use Poisson pseudo maximum likelihood regression, as opposed to least squares, in order to capture the proportional nature of technical change, and account for the fact that there are some dry holes in our

Figure 3: Sample grids and drilling

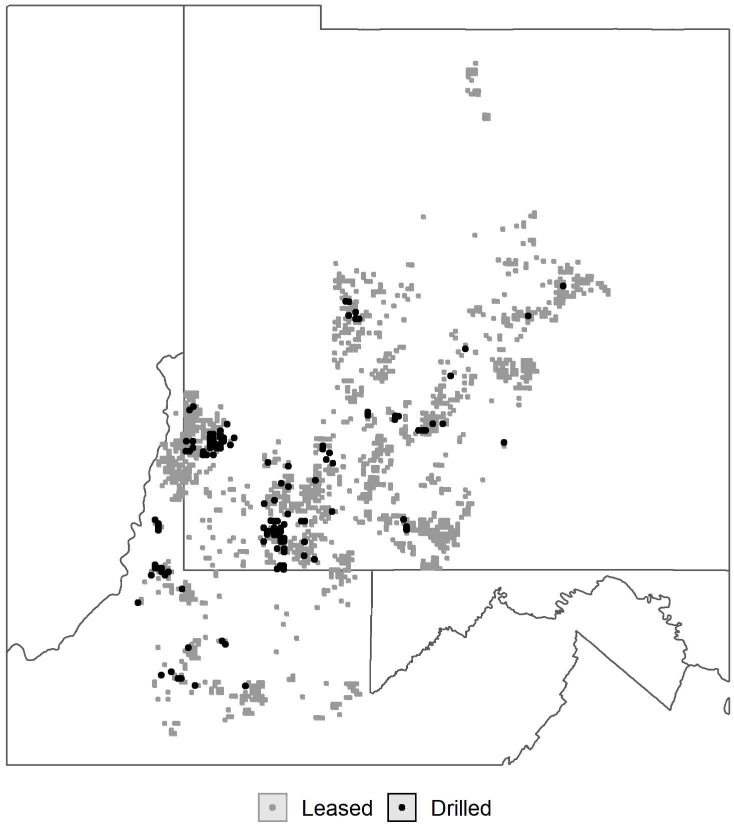

Leased grids are 33% leased by the end of 2005. Drilling indicates grids which have a shale well by the end of 2010.

a fixed year (we use 2009), with a fixed horizontal length (we use the median value in 2009), using our regression coefficients and that well's regression error. This procedure delivers X values for each *well* that are normalized to the technology available during our time period.

Next, we use this well-level normalized production data to estimate the average resource quality in each square mile grid with a geospatial krigging procedure. In the first step of this process, we compute the empirical spatial covariance between wells using the location of each well and our estimate of its normalized realized production. In the second step, we use this covariance function to compute predicted average resource quality at the centroid of each of our square mile grids.[12] At the end of this process, we have a prediction of average resource quality, X, for each grid in our sample, including grids that have not yet been drilled by the end of 2010. Online Appendix Figure 8 presents a map of these data.

[12]This procedure is similar to the approaches used in Covert (2015), Agerton (2020), Hodgson (2021), and Herrnstadt et al. (2020).

predictions, and Online Appendix Figure 9 presents the distribution of predicted output for the estimation sample. Across states, the two distributions share a lot of overlap. However the Pennsylvania distribution has a much longer right tail than West Virginia.

5.1 Empirical model

For each grid, we observe $X \in (0, \infty)$, the true resource quality at the potential drilling location, and the observed drilling outcome D. Drilling outcomes take on three possible values: $D =$ "early" means that the location was drilled in the first period of a game, $D =$ "late" means the location was drilled in the second period of a game, and $D =$ "never" means the location was not drilled before the end of 2010.

To translate this data into our model, we make two functional form assumptions. First, we assume that profits are linear in X:

$$\pi(X) = X(P(1 - \text{royalty} - \text{tax}) - \text{o\&m}) - K$$

We compute the output price P as the average realized spot natural gas price over the life of a typical well in our sample.[13] We set $P = \$3.1$/mcf, which is the median realized price for wells drilled during 2009. The royalty rate reflects the share of revenue that must be paid to landowners. In the lease data for this sample, the average royalty rate is 14%, and this is consistent across states, so we impose that value in estimation. As described above, production taxes in PA are 0%, while they are 5% in WV. We subtract operating and maintenance costs of $1/mcf based on estimates from publicly traded oil and gas company financial reports.[14] This leaves the drilling cost K as the only parameter in the profit function that we need to estimate.

Our second assumption is that the distribution of signals, conditional on the true resource quality, takes the "signal plus noise" form. We assume that a signal s can be written as $s = \log X + \epsilon$, with ϵ independent of X, and distributed normally, with zero mean and

[13] For details, see Online Appendix B.

[14] The largest publicly traded firms that were active in the Appalachian basin during our time period were Chesapeake, Range Resources, EQT, and Cabot Oil & Gas. While none of these companies reported lease operating costs specific to the Appalachian basin, together their reports are supportive of $1 per mcfe as a sensible operating cost. For example, in 2010, Chesapeake reported a total of $1.30/mcfe in "production expense per mcfe" and "General and administrative expense per mcfe." In the same year, Range Resources reported a total of $0.72/mcfe in "lease operating" and "workover" expense. EQT's lease operating expenses are considerably lower in that year, at $0.24/mcfe, but their operations include many conventional gas plays with lower operating costs. However, at the time EQT also operated a gas gathering and processing business, serving other producers in the Appalachian basin and other places. That business unit simultaneously reported $1.11/mcfe in "average gathering" fees. Finally, Cabot Oil & Gas reported $0.45/mcfe for their "North" operations, which include both Appalchain basin assets, as well as assets in the Rockies.

standard deviation ν. Because the normal distribution has a log-concave density, the joint distribution of s and X satisfies the monotone likelihood ratio property, a key assumption in the model we've developed above.[15] Thus, a second primitive we must estimate is ν.

Based on these assumptions, we can write the probability that we observe early drilling, conditional on X, in each game, as:

$$\Pr(D = \text{early} \mid X)_{CS} = 1 - \Phi\left(\frac{u_1 - \log X}{\nu}\right)$$

$$\Pr(D = \text{early} \mid X)_{FD} = 1 - \Phi\left(\frac{t_1 - \log X}{\nu}\right)$$

where $\Phi(\cdot)$ is the CDF of a standard normal random variable, and for convenience we have written $u_1 = v(\infty)$ and $t_1 = v(x^*)$. Similarly, we can write the probability of observing late drilling, again conditional on X, as:

$$\Pr(D = \text{late} \mid X)_{CS} = \underbrace{\left(1 - \Phi\left(\frac{u_1 - \log X}{\nu}\right)\right)}_{\text{Rival drills early}} \times \underbrace{\left(\Phi\left(\frac{u_1 - \log X}{\nu}\right) - \Phi\left(\frac{u_2 - \log X}{\nu}\right)\right)}_{\text{Own signal in late drilling range}}$$

$$\begin{aligned}\Pr(D = \text{late} \mid X)_{FD} = & \underbrace{\Phi\left(\frac{t_1 - \log X}{\nu}\right)}_{\text{Own signal too low to drill early}} \times \underbrace{\left(1 - \Phi\left(\frac{t_1 - \log X}{\nu}\right)\right) \mathbb{I}(X > x^*)}_{\text{Rival drills early and outcome is good}} \\ & + \underbrace{\left(\Phi\left(\frac{t_1 - \log X}{\nu}\right) - \Phi\left(\frac{t_2 - \log X}{\nu}\right)\right)}_{\text{Own signal in no news drilling range}} \times \underbrace{\Phi\left(\frac{t_1 - \log X}{\nu}\right)}_{\text{Rival waits}} \\ & \times \underbrace{\mathbb{I}(t_1 \geq t_2)}_{\text{No news drilling possible}}\end{aligned}$$

where $x^* = \frac{K}{P(1-\text{royalty}-\text{tax})-\text{o\&m}}$ is the value of X at which drilling has exactly 0 profits, and we write $u_2 = v^-(\infty)$ and $t_2 = v^0(x^*)$.

Although we are specifically interested in estimating the primitives (ν, K), the above probabilities are a function of the equilibrium cutoff signal values. To estimate the value of the primitives which best fit our data, we must compute the cutoff signals as a function of the primitives. We do this using the *empirical* distribution of the predicted X's in each

[15] See, for example, https://sites.stat.washington.edu/jaw/RESEARCH/TALKS/Toulouse1-Mar-p1-small.pdf. By the same logic, our assumption that ϵ is normally distributed can be relaxed, at some computational cost, to an assumption that ϵ has some unspecified log-concave density.

state.[16] For any conditional distribution of signals and any monotone profit function, we can solve for the implied equilibrium cutoffs consistent with the empirical distribution of X. For example, after substituting in our functional form assumptions for the conditional distribution of signals and the profit function, we obtain the u_1 condition for equilibrium first period drilling in the complete secrecy game as:

$$0 = \frac{1}{N}\sum_i (\tilde{P}X_i - K)\left(1 - \delta\left(1 - \Phi\left(\frac{u_1 - \log X_i}{\nu}\right)\right)\right)\phi\left(\frac{u_1 - \log X_i}{\nu}\right)$$

where $\tilde{P}$ is the output price net of taxes, royalty payments and operating expenses, X_i is the i-th observed value of X in our sample and $\phi(\cdot)$ is the density function for a standard normal random variable. We construct similar expressions for u_2, and for the cutoffs t_1 and t_2 in the full disclosure game.

We maximize the likelihood of the observed distribution of (X, D) over different values of (ν, K), with an inner "fixed point" step in which we solve for the equilibrium values of the game-specific signal cutoffs that are consistent with those parameters, using a numerical root-finding routine.

5.2 Results

The first row of Table 3 reports state-specific estimates of the standard deviation of the noise in firms' signals, ν. In both states, our estimates suggest that signals are quite noisy, though less so in Pennsylvania. The underlying distribution of $\log(X)$ in Pennsylvania has a standard deviation of 0.54, so this estimate of ν implies that about 87% of the variation in signals that Pennsylvania firms receive is noise.[17] In West Virginia, where the standard deviation of log output is smaller, but the ν estimate is higher, 99% of the variation in signals that firms receive is noise. These high values for ν underscore the importance of accounting for the information firms may receive about X after choosing to wait.

The second row of Table 3 reports estimates of the fixed cost of drilling, K. These point estimates are comparable to estimates of drilling costs for Marcellus shale wells reported in the U.S. Energy Information Administration (2016) drilling cost report. For the time period we study, EIA's estimate of drilling and completion costs ranges from about \$3 to \$6 million per well. Our estimates suggest drilling costs of about \$5 million in Pennsylvania, and a bit under \$4 million in West Virginia. One possible explanation for the lower drilling costs in West Virginia is the presence of a Manufacturing Investment Tax Credit that is available to natural gas producers in the state, specifically intended to offset up to 60% of a firm's 5%

[16] See Online Appendix C for details.

[17] $\frac{1.407^2}{1.407^2+0.54^2} \approx 0.87$.

Table 3: Structural Model Parameter Estimates and Fit

	Pennsylvania		West Virginia	
Parameters (Point Estimates, Standard Errors)				
ν	1.407	0.208	5.914	2.419
K	4.937	0.159	3.876	0.059
Fit Statistics (Empirical, Fitted)				
early	0.051	0.051	0.025	0.029
late	0.032	0.032	0.056	0.053
Sample Size				
N	1,227		284	
Negative Log-Likelihood	386.883		91.663	

Pennsylvania sample estimated assuming a complete secrecy information policy. West Virginia sample estimated assuming a full disclosure information policy. Both samples restricted to grids that are 33% or more leased up by the end of 2005, and overlie Gas Play 12 or Gas Play 13, from Figure 6. The parameter ν has the same units as log output, while the parameter K is in millions of dollars. Early drilling refers to the share of grids first drilled by the end of 2008, while late drilling refers to the share of grids first drilled in 2009 or 2010.

production tax liability.[18]

To assess the suitability of these models for explaining drilling behavior in their respective states, we compare observed patterns of drilling to the fitted predictions under our parameter estimates in the middle panel of Table 3. In both states, we are able to match the quantitative features of drilling probabilities — that late drilling is less likely in Pennsylvania and more likely in West Virginia — and our fitted probabilities are additionally close to their empirical counterparts.

Our parameter estimates are predicated on our assumption that drilling behavior in Pennsylvania is governed by CS information disclosure policies, while behavior in West Virginia is governed by FD policies. We can test these assumptions by estimating ν and K in each state under the opposite information policy, and use the method in Vuong (1989) to test the hypothesis that our chosen policy fits the data better. In both cases, the log-likelihoods of the data under our assumed information structure are higher than under the opposite information structures, which means that our Vuong test statistics have "correct" sign. In Pennsylvania, where we have a fairly large sample, we can reject both a a one-sided null hypothesis that FD policy fits the data better than CS as well as the traditional two-sided hypothesis that the two models fit the data equally well ($p = 0.01$). In contrast, in West Vir-

[18] https://tax.wv.gov/Documents/TSD/tsd110.pdf

ginia, with about one fifth as much data, we fail to reject the null that the two informational assumptions are the same.

With estimates of ν and K, we can evaluate the equilibrium consequences of counterfactual information disclosure policies. In Figure 4 we plot out the equilibrium cutoffs of partial disclosure games, across an evenly spaced grid of disclosure thresholds, in each state. The left-most part of each figure represents the equilibrium structure of an FD game, represented as a partial disclosure game with disclosure at that state's estimated x^*. The right-most part similarly represents the equilibrium structure of playing a CS game, represented as a partial disclosure game with disclosure at the maximum value of X observed for that state. The vertical black line corresponds to the estimated maximum nondistortionary disclosure (MND) threshold for the state. In Pennsylvania, the MND is 4.88 million mcf, which is the 95th percentile of the unconditional output distribution, while in West Virginia, it is 4.87 million mcf, the 97th percentile.

The top panel shows these cutoffs for Pennsylvania, where the first period cutoff $v(Z)$ always lies above the no-news cutoff $v^0(Z)$ for all Z, a condition which guarantees that full disclosure is the optimal disclosure policy. The decrease in $v(Z)$ as Z gets large is limited in Pennsylvania because our estimate ν is small. As a result, when a player receives a high signal in Pennsylvania, there is a good chance that the signal is high *because* X is high. Similarly, when a player waits and observes that its rival waited too, there is a good chance that X must be low, and thus drilling after no news is unprofitable. This highlights that, all else equal, a reduction in ν reduces the value of the information that can arrive if a player waits. In contrast, in West Virginia, which we plot in the bottom panel, cutoff signals are relatively high as a result of our much higher estimate of ν. For low disclosure thresholds, we see that $v^0(Z) < v(Z)$, implying that there is no-news drilling in West Virginia, which would be absent in Pennsylvania. Higher values of ν increase the likelihood that both high and low signals could be driven by noise, so a player who waits and learns that its rival has waited as well places far less weight on the event that X is unprofitable than it would in Pennsylvania.

In Figure 5 we plot out the *ex ante* average value of playing these partial disclosure games.[19] In both states, we estimate that the gains to moving from CS to FD are substantial. In Pennsylvania, the ex ante value of play under FD is \$67,806, while its \$45,471 under CS. This means CS information policies only generate 67% of the welfare that FD would. In West Virginia the differences are even more stark, with CS only capturing 39% of FD welfare. However, in both states, a maximum non-disortionary disclosure policy substantially improves upon CS, capturing 85% of FD in Pennsylvania, and 63% in West Virginia.

[19]By *ex ante* we mean the expected value of playing a game before one's signal is realized.

Figure 4: Equilibrium Cutoff Signals Under Counterfactual Disclosure Policies

(a) Pennsylvania

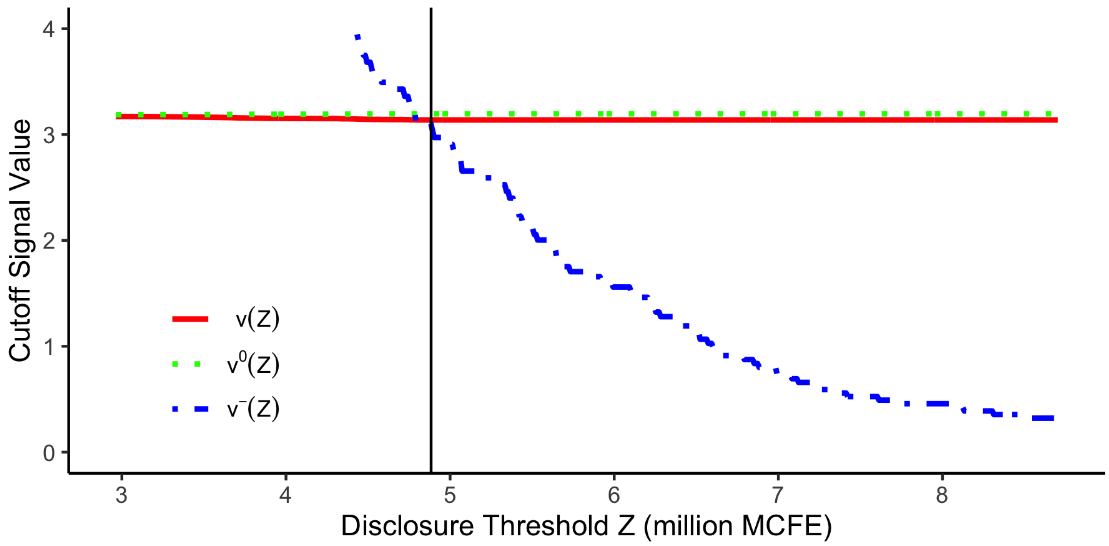

(b) West Virginia

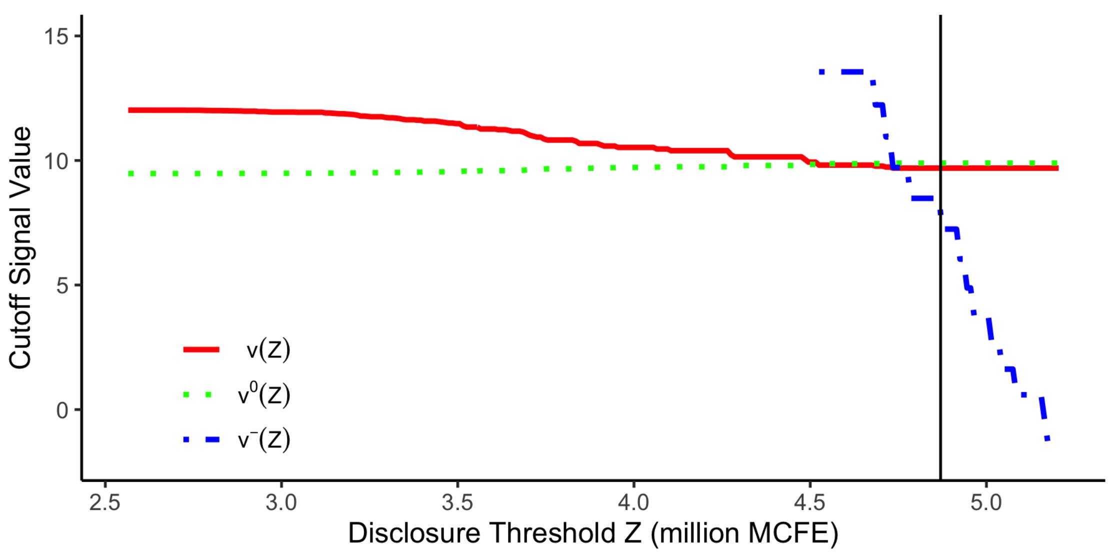

The solid red line is the first period cutoff, the dotted green line is the second period "no news" cutoff, and the dash-dotted blue line is the second period "bad news" cutoff. The bad news cutoff is especially high for low disclosure thresholds so it is not plotted for all values of Z in the figure. The black vertical line represents the location of $\widetilde{Z}$, maximum non-disortionary disclosure.

Though it is hard to see, the equilibrium cutoffs for both states shown in Figure 4 satisfy the $v^0(x^*) \geq v(\infty)$ condition that is at the heart of our proof of Proposition 4. As a result, the plots in Figure 5 necessarily reflect that FD has higher welfare than CS. Additionally, in Pennsylvania (but not West Virginia), it turns out that $v^0(x^*) > v(x^*)$ (there is no no-news drilling), so the reason why the top panel of Figure 5 is monotonically decreasing is precisely the logic employed in the proof of Theorem 1.

6 Conclusion

In non-cooperative settings with costly investment and information externalities, regulators face a tradeoff when it comes to disclosing information on exploratory activities (or requiring that it be disclosed). Disclosing investment outcomes disseminates socially valuable information that can improve the efficiency of subsequent investment, but this prospect also induces free riding. In a two-firm, two-period model, we characterize equilibrium behavior and show how the resulting optimal disclosure policy depends on firms' patience. Moreover, we establish that some amount of disclosure is free, in that it does not distort the timing of early investment, but does increase the welfare from subsequent investment. Thus, if partial disclosure is possible, complete secrecy can never be optimal.

We quantify the gains from optimal disclosure policy in the context of the Appalachian shale boom. The rapid maturation of hydraulic fracking and horizontal drilling at the turn of the century transformed the economics of this shale basin, but, in order to take advantage of this new technology, exploration companies still had to figure out *where* to apply it.[20] During this time, Pennsylvania regulators allowed firms to keep their exploratory efforts secret, while neighboring West Virginia disclosed this information to rivals almost immediately. We show that, had Pennsylvania followed West Virgina's lead, the private value of shale exploration would have been 49% higher during this formative period. Consistent with this, in 2011 Pennsylvania abruptly ended its secrecy policy, and began fully disclosing oil and gas exploration and production outcomes.

The oil and gas industry in noteworthy because disclosure is common. However there is no economic reason why these lessons couldn't be applied to other settings. Governments often collect information closely related to investment outcomes for tax or regulatory purposes. Our analysis shows that some of that disclosure could be free, and the returns to optimal disclosure could be large.

[20]This observation also has also been noted by Agerton (2020) in another gas producing basin.

Figure 5: Average Player Welfare Under Counterfactual Disclosure Policies

(a) Pennsylvania

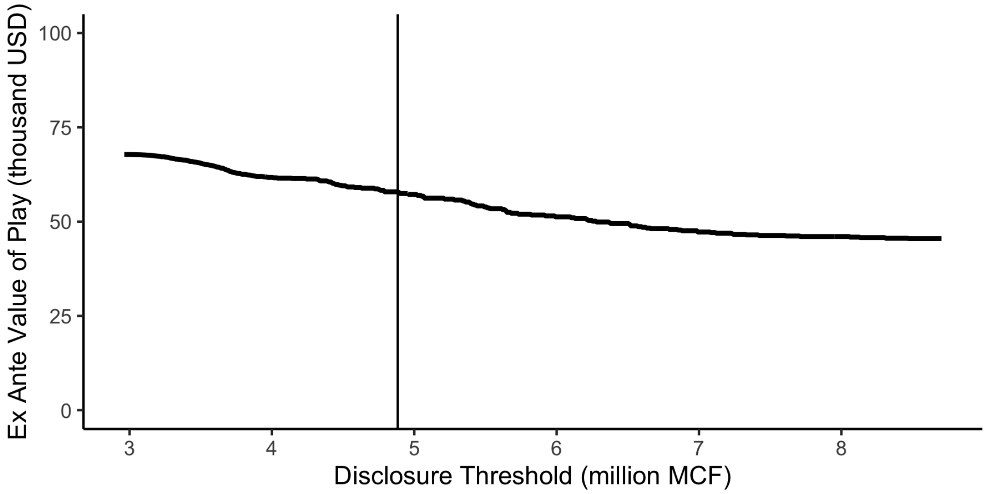

(b) West Virginia

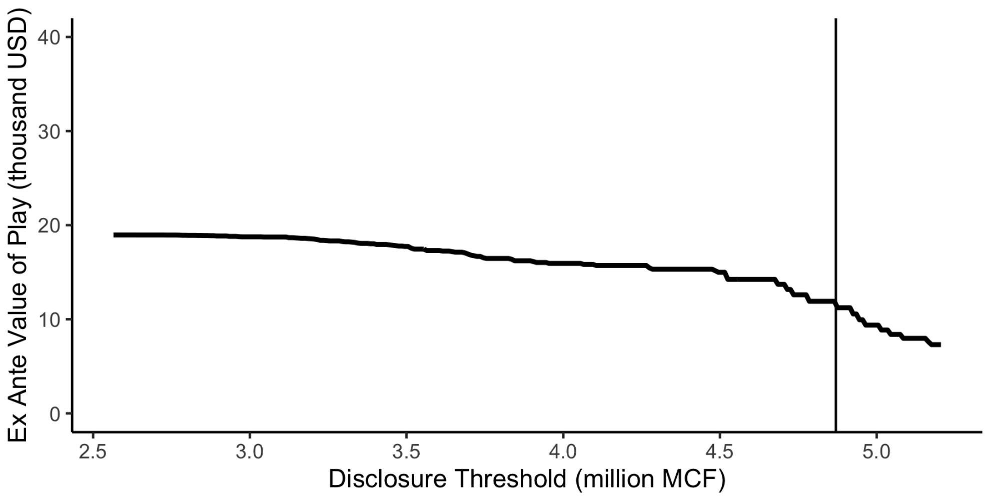

The black vertical line represents the location of maximal non-distortionary disclosure.

Table 4: Lease summary stats

	Pennsylvania	Mean	SD	P01	P25	P50	P75	P99	N
Primary Term (years)	No	4.51	1.68	0.42	3.00	5.00	5.00	10.01	72437
	Yes	4.89	1.64	0.50	5.00	5.00	5.00	10.01	123474
Size (acres)	No	58.96	125.88	0.40	2.18	26.95	83.20	383.43	72437
	Yes	67.09	316.99	0.39	3.90	18.79	69.93	569.14	123474
Royalty Rate (fraction)	No	0.13	0.01	0.12	0.12	0.12	0.12	0.18	52279
	Yes	0.14	0.03	0.12	0.12	0.12	0.15	0.20	60035

Online Appendix

A Additional Tables and Figures

B Construction of average prices

We assume that firms have rational expectations about the future path of natural gas prices. This would imply that the expected discounted weighted average price of natural gas, over the time span a well will produce, can be estimated by its empirical counterpart.

To do this, we compute, for each well i in our sample, a realized gas price as:

$$P_i = \frac{\sum_{t=1}^{T_i} \delta_{\text{monthly}}^t H_{\tau_i+t-1} \gamma_t}{\sum_{t=1}^{T_i} \gamma_t}$$

where $\delta_{\text{monthly}} = 0.992$ is our monthly discount factor, which we set to mimic our annual discount factor of 0.9, H_j is the realize price at Henry Hub in calendar month j, τ_i is the calendar month in which well i was drilled, T_i is the number of months between τ_i and July, 2022, and $\{\gamma_t\}_{t=1}^{240}$ are a set of weights which sum to 1, and match the decline rate structure implied by natural gas production in this setting.

We calibrate the γ_t's to both match the approximate rate of decline which is recorded in our production data, and to fit a traditional Arps decline curve. Specifically, we assume that production in the t-th month of a well's life is given by $y_{it} = y_{i0} t^{\beta}$. Our production data records the cumulative production after 12 and 36 months for each well. We pick β to match the median ratio of 36 month cumulative production to 12 month cumulative production among all wells in our sample. This results in $\beta \approx -0.44$. Finally, using the above model, we compute the average of P_i over all wells drilled during 2009, obtaining $P = \$3.10/mcf$.

Figure 6: Gas Region Map

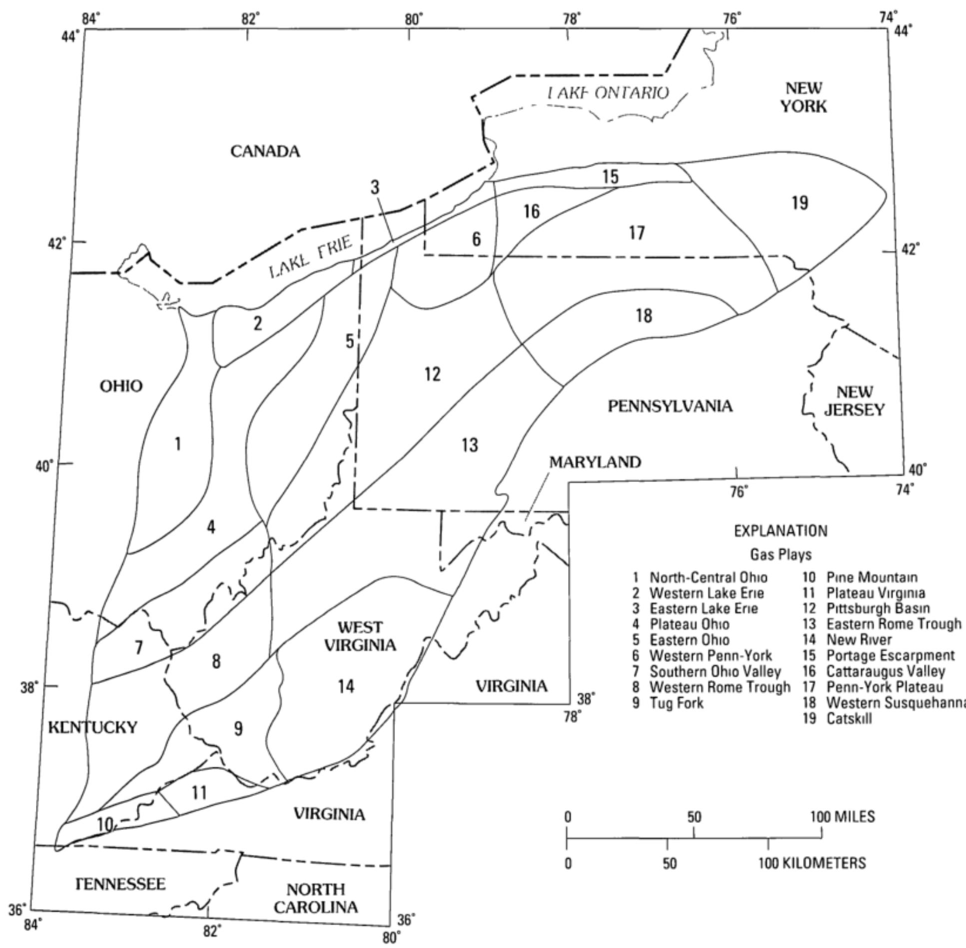

Figure 2. Devonian shale gas plays in the Appalachian basin

From Figure 2 of Charpentier et al. (1993). Each region was studied by U.S. Geological Survey geologists using data generated by the Eastern Gas Shales Project, a publicly funded effort to explore the potential of shale gas resources during the oil crisis of the 1970s (Wang and Krupnick, 2015).

Figure 7: Border Grids, Pre-2010 Leasing and Drilling

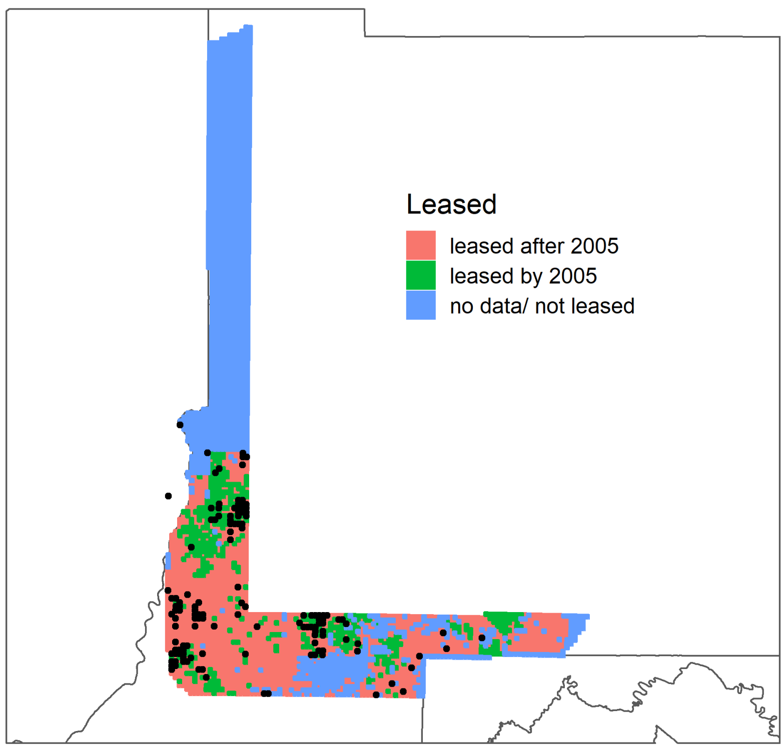

Black circles reflect square-mile grids which were drilled before the end of 2010.

Figure 8: Predicted Output (mmcfe)

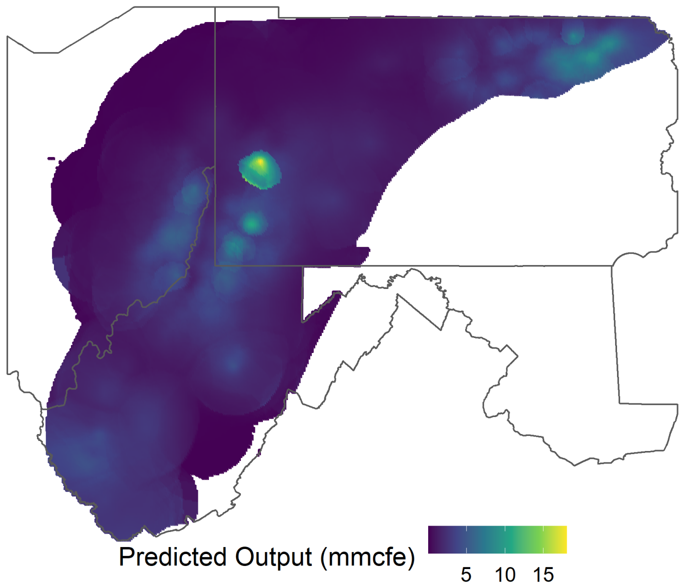

Output predictions generated using the geospatial krigging procedure.

C Integrating over the empirical distribution of X

To estimate the value of the primitives which best fit our data, we must be able to compute the cutoff signals as a function of the primitives. We do this using *empirical* distribution of the X's. To see how, consider the task of computing the expected value of some function $z(X)$ with respect to the distribution of X, conditional on a signal realization u.

$$\begin{aligned}\int z(x)h(x \mid u)dx &= \frac{1}{\int f(u \mid y)q(y)dy}\int z(x)f(u \mid x)q(x)dx \\ &\approx \frac{1}{N}\frac{1}{\int f(u \mid y)q(y)dy}\sum_i z(X_i)f(u \mid X_i)\end{aligned}$$

In the first line above, we invoke Bayes' rule to express the density of X conditional on a signal realization as the unconditional density, $q(x)$, times the density of *signals* conditional on X, divided by a constant. If we have an iid sample of X's, then the approximation in the second line will have a *plim* equal to our initial integral.

Our equilibrium conditions can be written in this fashion, with a suitable choice of the function $z(x)$ that will imply the left-hand side of the above expression is equal to zero. For

Figure 9: Distribution of predicted output

(a) All grids in DeWitt regions 12 and 13

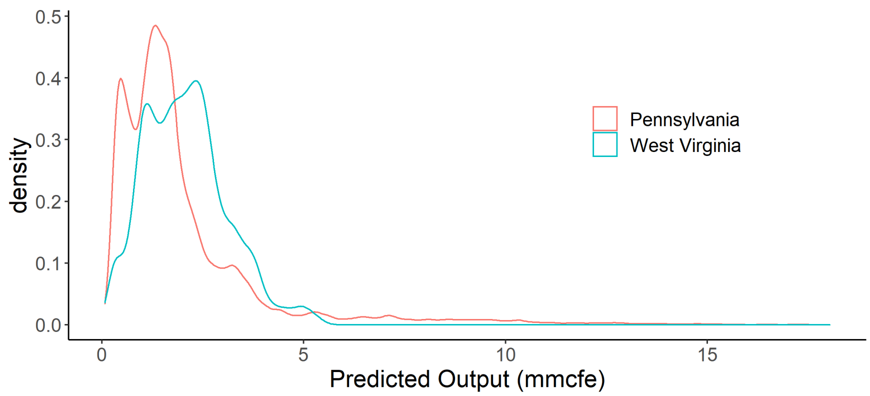

(b) Grids leased by 2005

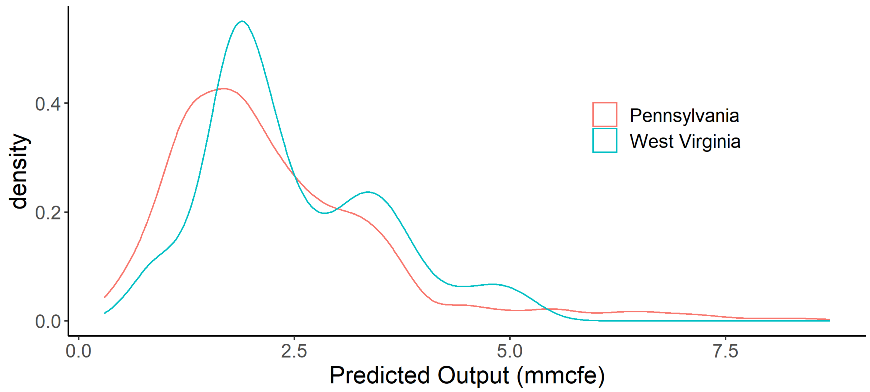

Output predictions generated using the geospatial krigging procedure.

example, we can write the equilibrium condition for u_1 as:

$$\begin{aligned} 0 &= \int \underbrace{\pi(x)\left(1-\delta(1-F(u_1 \mid X))\right)}_{z_{u_1}(x)} h(x \mid u_1)dx \\ &= \int z_{u_1}(x)h(x \mid u_1)dx \\ &\approx \frac{1}{N}\frac{1}{\int f(u_1 \mid y)q(y)dy}\sum_i z_{u_1}(X_i)f(u_1 \mid X_i) \\ &= \frac{1}{N}\sum_i z_{u_1}(X_i)f(u_1 \mid X_i) \end{aligned}$$

In the final line of this expression, we have removed the Bayes' rule integrating constant because it is positive for all values of u_1 and because the left-hand side of the expression is equal to zero. Thus, for a fixed conditional distribution of signals $f(u \mid x)$, and for a fixed profit function $\pi(x)$, we can solve for the implied equilibrium first period cutoff in a complete secrecy game (u_1) consistent with the empirical distribution of X. We can construct similar expressions for u_2, and for the cutoffs in the full disclosure game.

References

Agerton, Mark (2020): "Learning where to drill: Drilling decisions and geological quality in the haynesville shale," .

Black, Katie Jo, Shawn J McCoy, and Jeremy G Weber (2018): "When externalities are taxed: The effects and incidence of Pennsylvania's impact fee on shale gas wells," *Journal of the Association of Environmental and Resource Economists*, 5, 107–153.

Bonatti, Alessandro and Johannes Hörner (2017): "Learning to Disagree in a Game of Experimentation," *Journal of Economic Theory*, 169, 234–269.

Brown, Jason P, Peter Maniloff, and Dale T Manning (2020): "Spatially variable taxation and resource extraction: The impact of state oil taxes on drilling in the US," *Journal of Environmental Economics and Management*, 103, 102354.

Campbell, Arthur, Florian Ederer, and Johannes Spinnewijn (2014): "Delay and deadlines: Freeriding and information revelation in partnerships," *American Economic Journal: Microeconomics*, 6, 163–204.

Chamley, Christophe and Douglas Gale (1994): "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica: journal of the Econometric Society*, 62, 1065–1085.

Charpentier, Ronald R, Wallace de Witt, GE Claypool, LD Harris, RF Mast, JD Megeath, JB Roen, JW Schmoker, and RC Kepferle (1993): "Estimates of unconventional natural-gas resources of the Devonian shale of the Appalachian Basin," *Petroleum Geology of Devonian and Mississippian Black Shale of Eastern North America. Reston, VA: United States Geological Survey (USGS).*

Covert, Thomas (2015): "Experiential and Social Learning in Firms: the Case of Hydraulic Fracturing in the Bakken Shale," Working paper.

Covert, Thomas R and Richard L Sweeney (2019): "Relinquishing Riches: Auctions vs Informal Negotiations in Texas Oil and Gas Leasing," .

Friedman, David D, William M Landes, and Richard A Posner (1991): "Some Economics of Trade Secret Law," *The journal of economic perspectives: a journal of the American Economic Association*, 5, 61–72.

Grenadier, Steven R (1999): "Information Revelation Through Option Exercise," *The review of financial studies*, 12, 95–129.

Halac, Marina, Navin Kartik, and Qingmin Liu (2017): "Contests for experimentation," *Journal of Political Economy*, 125, 1523–1569.

Heidhues, Paul, Sven Rady, and Philipp Strack (2015): "Strategic experimentation with private payoffs," *Journal of Economic Theory*, 159, 531–551.

Hendricks, Kenneth and Dan Kovenock (1989): "Asymmetric Information, Information Externalities, and Efficiency: The Case of Oil Exploration," *The Rand Journal of Economics*, 20, 164.

Hendricks, Kenneth and Robert H. Porter (1996): "The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts," *The American Economic Review*, 86, 388–407.

Herrnstadt, Evan M, Ryan Kellogg, and Eric Lewis (2020): "The Economics of Time-Limited Development Options: The Case of Oil and Gas Leases," Working Paper 27165, National Bureau of Economic Research.

Hodgson, Charles (2021): "Information Externalities, Free Riding, and Optimal Exploration in the UK Oil Industry," .

Iacus, Stefano M, Gary King, and Giuseppe Porro (2012): "Causal inference without balance checking: Coarsened exact matching," *Political analysis*, 20, 1–24.

Karlin, Samuel and Herman Rubin (1956): "The theory of decision procedures for distributions with monotone likelihood ratio," *The Annals of Mathematical Statistics*, 272–299.

Keller, Godfrey, Sven Rady, and Martin Cripps (2005): "Strategic experimentation with exponential bandits," *Econometrica*, 73, 39–68.

Krieger, Joshua L (2021): "Trials and terminations: Learning from competitors' R&D failures," *Management Science*, 67, 5525–5548.

Lin, C-Y Cynthia (2013): "Strategic decision-making with information and extraction externalities: A structural model of the multistage investment timing game in offshore petroleum production," *Review of Economics and Statistics*, 95, 1601–1621.

Milgrom, Paul R (1981): "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 12, 380–391.

Rob, Rafael (1991): "Learning and Capacity Expansion under Demand Uncertainty," *The Review of Economic Studies*, 58, 655–675.

Rosenberg, Dinah, Antoine Salomon, and Nicolas Vieille (2013): "On games of strategic experimentation," *Games and Economic Behavior*, 82, 31–51.

Steck, Andrew L (2022): "Industry dynamics with social learning: Evidence from hydraulic fracturing," .

U.S. Energy Information Administration (2016): "Trends in U.S. Oil and Natural Gas Upstream Costs," Tech. rep.

Vuong, Quang H (1989): "Likelihood ratio tests for model selection and non-nested hypotheses," *Econometrica: Journal of the Econometric Society*, 307–333.

Wang, Zhongmin and Alan Krupnick (2015): "A retrospective review of shale gas development in the United States: What led to the boom?" *Economics of Energy & Environmental Policy*, 4, 5–18.