NBER WORKING PAPER SERIES

SECRECY RULES AND EXPLORATORY INVESTMENT:
THEORY AND EVIDENCE FROM THE SHALE BOOM

Thomas R. Covert
Richard L. Sweeney

Working Paper 30548
http://www.nber.org/papers/w30548

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022

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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Secrecy Rules and Exploratory Investment: Theory and Evidence from the Shale Boom
Thomas R. Covert and Richard L. Sweeney
NBER Working Paper No. 30548
October 2022
JEL No. D82,L51,L71

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.

Thomas R. Covert
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
thomas.covert@chicagobooth.edu

Richard L. Sweeney
Economics Department
Boston College
Chestnut Hill, MA 02467
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. In the first part of the paper, we theoretically analyze investment behavior under two extreme information disclosure policies: 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 level and timing of investment across these two regimes, and allows us to characterize the conditions under which each regime generates more efficient investment than the other. We then generalize our theoretical analysis to consider intermediate information disclosure policies. 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]

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 a full disclosure policy.[2] 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 drilled first. HK show that 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.

We modify the HK model to understand strategic responses to a policy of complete secrecy (CS). In this environment, a firm that waits learns no information about the outcome of investment if its rival invests in period one, but it does learn that investment happened.[3] This induces some free riding, as firms who wait learn something about their rival's *signal*. We show that this game also has a symmetric equilibrium in cutoff strategies, in which firms with high signals drill early, while firms with lower signals wait. Consistent with the fact that the information revealed to those who wait is less valuable in CS, we show that there is less free riding, compared to FD. Whether this reduction in free riding generates a welfare gain depends on the value of the informational spillovers that do not occur in CS but will in FD. While the net effect depends on the primitives of the problem, we provide a sufficient

[1]In this paper, we only consider the private benefits and costs associated with investment. In the case of 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]The economic theory on bandit games has also explored the equilibrium effects of investment disclosure in games with information spillovers. Rosenberg et al. (2013) studies an irreversible exit version of the exponential bandit setting pioneered by Keller et al. (2005), finding that games with publicly disclosed outcomes deliver higher equilibrium payoffs than similar games with no outcome disclosure. Heidhues et al. (2015) studies a traditional discrete time multi-armed bandit game, allowing for communication between players, and comes to the opposite conclusion, that equilibrium payoffs can actually be higher in a game without publicly disclosed outcomes. Both papers assume that players face an infinite investment horizon. In contrast, our motivating policy example and empirical application is oil and gas exploration, which is inherently finite horizon. Furthermore, both papers envision an intensive margin of investment, whereas mineral exploration (and clinical trials) are inherently binary decisions.

[3]In the motivating examples of mineral exploration, clinical trials and real estate development, the act of investment is essentially unhideable. However, in other settings this may not be the case. Bonatti and Hörner (2017) studies the role of disclosure of investment *actions* in games with information spillovers, assuming outcomes are always observable. Using a Keller et al. (2005) style model they conclude that equilibrium payoffs 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.

condition which guarantees that FD generates more welfare than CS. The key force behind this condition is the extent of firm patience, and sufficiently impatient firms will always prefer full disclosure to complete secrecy.

Although full disclosure and complete secrecy are common real world information policies, there is no reason why a regulator could not choose something in between these two extremes. We make this point formally by studying a class of partial disclosure (PD) policies, in which a binary characterization of the return to investment (is it above or below a pre-specified threshold) is revealed in the second period if one player invests in the first period. We show that all partial disclosure policies have a similar symmetric cutoff equilibrium structure, with a first period cutoff signal which is non-increasing in the disclosure threshold. This generalizes our result above, in that more disclosure (a lower disclosure threshold) increases free-riding.

Our first result from studying partial disclosure is that some amount of disclosure is "free." Specifically, there is a disclosure threshold, which we call maximum nondistortionary disclosure (MND), which generates the same amount of first period investment as CS (it has the same first period equilibrium cutoff signal), but generates strictly more welfare in the second period for players who wait. Thus, in this class of partial disclosure policies, complete secrecy (a disclosure threshold equal to infinity) can never be the optimal disclosure policy, regardless of the primitives of the problem. Our second theoretical result is that more disclosure (a lower disclosure threshold) delivers *higher* expected player welfare, up to the point at which it induces a phenomenon we call "no news" drilling. Depending on the primitives of the environment, it is possible for firms to have signals that are lower than the first period cutoff of a partial disclosure game, but are high enough that if they learn their rival also waited, they can still profitably drill. When this happens, we say a firm has drilled after learning "no news." If the primitives of the environment are such that no news drilling is impossible, even for the FD game (a partial disclosure game with a threshold equal to the point of profitability), then full disclosure is the optimal policy.

Our inspiration for this 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, regulators in 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 free riding 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 model of complete secrecy 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.

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 about the value of X. Conditional on the true state of the world ($X = x$), the signals are independently drawn from some probability distribution $\Pr(S \leq s \mid X = x) = F(s \mid x)$, and the distribution $F(\cdot \mid \cdot)$ 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 $h(x \mid s)$. 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 complete secrecy 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 *does not* learn the true value of X. This setup 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. Our first task will be to establish the existence of a symmetric cutoff equilibrium of the complete secrecy game, and to compare it to the behavior in the full disclosure game.

To begin, we first introduce some notation and basic results that we'll use throughout this section. First, we'll write $\mathbb{E}\left[g(X) \mid s\right] = \int_0^\infty g(x)h(x \mid s)dx$ for any function $g(\cdot)$ and any single signal realization s. Next, we derive the expected value of functions of X under knowledge of your own signal and the knowledge that your rival's signal is bounded above (or below) by a known constant.

Lemma 1 *Let p and q be any signal realizations. The expected value of a function $g(X)$ when your signal, S, is equal to p, and your rival's signal, T is less than or equal to q, is given by:*

$$\mathbb{E}\left[g(X) \mid S = q, T \leq p\right] = \frac{\mathbb{E}\left[g(X)F(q \mid X) \mid p\right]}{\mathbb{E}\left[F(q \mid X) \mid p\right]}$$

Analogously, the expected value when your rival's signal is larger than q is:

$$\mathbb{E}\left[g(X) \mid S = q, T \geq p\right] = \frac{\mathbb{E}\left[g(X)(1 - F(q \mid X)) \mid p\right]}{\mathbb{E}\left[1 - F(q \mid X) \mid p\right]}$$

For a proof, see appendix A.

We also re-state an existing result about the expected value of monotonic functions of X, conditional on ordered signals:

Lemma 2 *Adapted from Milgrom (1981), Proposition 4. For any signals t, u and v, and any monotonic function $g(x)$:*

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

As applied to the setting here, a player with signal t who knows that its rival's signal is greater than u 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 v. This is true for any values of u and v. 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.

2.1 Equilibrium under full disclosure

With this notation and one application of Lemma 1, we can write the equilibrium condition for the first period cutoff signal t_1 in the full disclosure game as:

$$\begin{aligned}\delta^{-1}\mathbb{E}\left[\pi(X) \mid t_1\right] &= \mathbb{E}\left[1-F(t_1 \mid X) \mid t_1\right]\mathbb{E}\left[\max(\pi(x),0) \mid S=t_1, T \geq t_1\right] \\ &\quad + \mathbb{E}\left[F(t_1 \mid X) \mid t_1\right]\max(0, \mathbb{E}\left[\pi(x) \mid S=t_1, T<t_1\right]) \\ &= \mathbb{E}\left[\max(\pi(x),0)(1-F(t_1 \mid x)) \mid t_1\right] + \max\left(0, \mathbb{E}\left[\pi(x)F(t_1 \mid x) \mid t_1\right]\right)\end{aligned}$$

In words, a firm with the cutoff signal t_1 is indifferent between drilling today, and waiting until the next period to either drill under full disclosure (if its rival drilled in the first period), or optionally drill in the second period after learning that its rival did not drill. Firms with signals larger than t_1 strictly prefer drilling in the first period to waiting, and firms with lower signals prefer to wait.

Whether a firm who waits ($s < t_1$) is willing to drill after learning that its rival's signal was lower than t_1 depends on whether $\mathbb{E}\left[\pi(X)F(t_1 \mid X) \mid s\right]$ is non-negative. When this object is strictly positive for $s = t_1$, there is another cutoff, t_2, that determines whether players who wait in the first period will drill after "no news" in the second period, given by:

$$0 = \mathbb{E}\left[\pi(X)F(t_1 \mid X) \mid t_2\right]$$

A firm with signal t_2 is indifferent to drilling when its rival has a signal lower than the first period cutoff, t_1, and so all firms with signals in the interval $[t_2, t_1)$ will wait, and drill after no news. It is possible for $t_2 > t_1$, so that any player who would drill after no news would have already drilled in the first period.[4] However, as we describe later, this situation is economically interesting, as it guarantees that a full disclosure policy delivers higher welfare than a complete secrecy policy.

2.2 Equilibrium under complete secrecy

We can derive a symmetric cutoff equilibrium of the complete secrecy game in a similar fashion, summarized in the following proposition:

Proposition 1

1. *Firms with signals at or above u_1 drill in the first period, where u_1 is the unique solution to*
$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid u_1\right] = \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_1\right] \tag{1}$$

2. *Firms with lower signals wait, and if a firm who waits learns in the second period that its rival drilled, then it will drill if its signal is above u_2, where u_2 is the unique solution to*
$$\mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_2\right] = 0. \tag{2}$$

3. $u_2 < u_1$

For a proof, see appendix B.

Firms with the first period marginal signal u_1 are indifferent between drilling in the first period, and waiting until the second period to see whether their rival drilled. If they wait and learn that their rival drilled, they will drill too, but if their rival does not drill, they won't drill after waiting. Unlike in the FD game, drilling after no news is impossible in the CS game.

Another important difference between the cutoffs in the FD game and the CS game is that the second period cutoff, u_2, determines whether a firm who waits will drill after observing "some news," namely that its rival drilled, whereas t_2, the second period cutoff in the FD game, determines whether a firm who waits drills after no news.

2.3 Comparison of the equilibria in the two games

The qualitative structure of the equilibria of these two games are different, at least in the second period. In the full disclosure game, late drilling can occur even if there was no early

[4] Hendricks and Kovenock (1989) make an additional assumption (Assumption 3, on page 172) in their analysis to rule out this possibility.

drilling, while a player drills late in the complete secrecy game only when its rival drilled early. It is also possible to ordinally rank the two sets of cutoffs.

Proposition 2 *Drilling happens earlier under complete secrecy than it does under full disclosure ($u_1 < t_1$).*

We prove this by contradiction. Suppose that $u_1 \geq t_1$. Then the following algebra must hold:

$$\begin{aligned}
\mathbb{E}\left[\pi(X) \mid u_1\right] &> \delta\mathbb{E}\left[\max(\pi(X),0)(1-F(t_1 \mid X)) \mid u_1\right] + \max(0, \delta\mathbb{E}\left[\pi(X)F(t_1 \mid X) \mid u_1\right]) \\
&\geq \delta\mathbb{E}\left[\max(\pi(X),0)(1-F(t_1 \mid X)) \mid u_1\right] \\
&\geq \delta\mathbb{E}\left[\max(\pi(X),0)(1-F(u_1 \mid X)) \mid u_1\right] \\
&> \delta\mathbb{E}\left[\pi(X)(1-F(u_1 \mid X)) \mid u_1\right] \\
&= \mathbb{E}\left[\pi(X) \mid u_1\right], \text{ a contradiction.}
\end{aligned}$$

The first inequality comes from the fact if $u_1 \geq t_1$, then a player with signal u_1 who is in the first period of the full disclosure game strictly prefers drilling to waiting. The second comes from the fact that the payoff to waiting in FD at signal u_1 when it is revealed that the other player has a signal below t_1 is non-negative, so removing it can only reduce the total payoff to waiting. The third inequality comes from the fact that we have assumed $u_1 \geq t_1$, so that for any X, $1 - F(t_1 \mid X) \geq 1 - F(u_1 \mid X)$, and that $\max(\pi(X), 0)$ is non-negative. The fourth inequality comes from "adding back" the negative payoffs that can occur when drilling, which players can sometimes avoid in the full disclosure game. The last (contradicting) equality is the definition of the u_1 cutoff. Thus, our assumption that $u_1 \geq t_1$ must be wrong, and we can conclude that $u_1 < t_1$.

Proposition 2 tells us that there is less free-riding in CS than FD, in the sense that there are signals $s \in [u_1, t_1)$ which will drill in the first period of a CS game, but will wait in an FD game.

Proposition 3 *The second period cutoff is also lower ($u_2 < t_2$).*

When no news drilling is impossible ($t_2 > t_1$), this is trivial, since $u_2 < u_1 < t_1$. When no news drilling is possible, we again prove by contradiction. Suppose $u_2 \geq t_2$. Then the following algebra holds:

$$\begin{aligned}
0 &= \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_2\right] \\
&> \frac{\mathbb{E}\left[(1 - F(u_1 \mid X)) \mid u_2\right]}{\mathbb{E}\left[F(t_1 \mid X) \mid u_2\right]}\mathbb{E}\left[\pi(X)F(t_1 \mid X) \mid u_2\right] \\
&\geq 0, \text{ a contradiction.}
\end{aligned}$$

The first line is simply the definition of the equilibrium value for u_2. The next line, a strict inequality, comes from a re-arrangement of Lemma 2: at any signal s, value of drilling when you know your rival's signal is *greater than or equal to* u_1 is better news than the value of drilling when you know your rival's signal is less than or equal to t_1. The final line comes from the fact that the definition of t_2 tells us that for any signal $s \geq t_2$, the value of drilling after observing that your rival has a signal lower than t_1 is non-negative, and we have assumed that $u_2 \geq t_2$. However, this implies that the value of drilling when you have signal u_2 and you know your rival has a signal larger than u_1 (the first fine) is strictly positive, which violates the definition of u_2. Thus, $u_2 < t_2$.

With these results, we can partially order the equilibrium cutoff values that occur in these two games. When no news drilling is possible, we'll either have $u_2 < u_1 < t_2 < t_1$ or $u_2 < t_2 < u_1 < t_1$. When no news drilling is impossible, we'll have $u_2 < u_1 < t_1 < t_2$.

2.4 Which game generates more value?

Let $\sigma^g(s) \in \{0, 1\}$ denote a firm's decision to drill in period 1 as a function of its signal in the game $g \in \{\text{Full Disclosure}, \text{Complete Secrecy}\}$. The value of playing g given signal s is:

$$V^g(s) = \sigma^g(s)\mathbb{E}\left[\pi(X) \mid s\right] + \delta(1 - \sigma^g(s))W^g(s)$$

where W^g is value of waiting until period 2 to make a decision. While $\mathbb{E}\left[\pi(X) \mid s\right]$ is the same for both games, the value of waiting differs:

$$\begin{aligned}
W^{\text{Full Disclosure}}(s) &= \mathbb{E}\left[\max(0, \pi(X))(1 - F(t_1 \mid X)) \mid s\right] + \max(0, \mathbb{E}\left[\pi(X)F(t_1 \mid X) \mid s\right]) \\
W^{\text{Complete Secrecy}}(s) &= \max(0, \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid s\right]) + \max(0, \mathbb{E}\left[\pi(X)F(u_1 \mid X) \mid s\right])
\end{aligned}$$

With these objects, we can compute the difference in expected outcomes between playing the two games for a player with a fixed signal s. We are able to sign this difference, in favor full disclosure, under an important special case.

Proposition 4 *When $t_2 > u_1$, the expected value of playing the full disclosure game is larger* *than the expected value of playing the complete secrecy game, for all signals s:*

$$V^{Full\ Disclosure}(s) \geq V^{Complete\ Secrecy}(s)$$

We provide a full proof of this in Appendix C, and provide a sketch here, by verifying the assertion for each relevant range of signal realizations. Firms with especially high signals will drill in either game, so their payoffs are identical. Firms with especially low signals will *only* drill in the FD game, and when they do they earn positive profits, so they must strictly prefer the FD game. Firms with signals that are high enough to drill early in the CS game but wait in the FD game must be better off in the FD game by revealed preference. That leaves firms with signals that are too low to drill early in either game, but high enough to drill after waiting in both games. For these signals, which lie in $[u_2, u_1]$, under our condition that $u_1 < t_2$, the probability that a player drills after waiting in the complete secrecy game, conditional on X, is *always* higher than the probability that the same player drills after waiting in the full disclosure game. This is because $u_1 < t_1$ and because signals in this range can't do no news drilling in the FD game. When drilling is worthwhile ($X \geq x^*$), this additional chance of drilling in CS is good, and when its not, it is bad. However, even at the most positive signal realization in this region ($s = u_1$), we know that there is enough probability mass on X to the left of x^* to make this difference negative in expectation (otherwise, a firm with signal u_1 wouldn't strictly prefer FD). Thus, all signals weakly prefer the FD game to the CS game.[5]

When the above condition on u_1 and t_2 fails, it is still possible that that equilibrium payoffs are higher under FD information policy than CS policy. When that occurs, it will not necessarily be the case that all signals weakly prefer FD to CS. Instead, as we describe in the proof in the appendix, there will be a range of intermediate signals that prefer CS to FD, and as long as this range is sufficiently small, average welfare will still be in favor of FD. However, it is possible to construct environments in which $u_1 \geq t_2$ **and** average equilibrium payoffs are higher under CS than FD.[6]

2.4.1 When is $u_1 < t_2$?

Given that it is sometimes possible to easily rank the two games in expected value terms, it is worth investigating the conditions that could lead to this special case. Intuitively, this

[5] We can use similar logic to show that equilibrium payoffs in the CS game is larger than single agent payoffs. See Online Appendix H.1

[6] 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. Then for $\delta \leq 0.80$, we will have $u_1 < t_2$. However, we'll still have $\mathbb{E}_s\left[V^{\text{Full Disclosure}}(s)\right] > \mathbb{E}_s\left[V^{\text{Complete Secrecy}}(s)\right]$ for values of δ up to 0.97. Once $\delta > 0.97$, we'll have the opposite.

must occur when firms are not especially patient, as u_1 and t_1 are both increasing in δ, while t_2 is *decreasing* in t_1. When δ is small enough, t_1 will be small, which requires that t_2 be large. However, it is always the case that $u_1 < t_1$, so a small enough δ will guarantee $u_1 < t_2$. We formalize this here:

Proposition 5 *Fix a distribution of X, a conditional distribution of signals $F(\cdot \mid \cdot)$ satisfying the MLRP, and a monotonically increasing profit function $\pi(\cdot)$ that crosses zero once. Then there is always a $0 < \overline{\delta} < 1$ such that for all $\delta < \overline{\delta}$, the equilibrium values of u_1 and t_2 when the discount factor is δ satisfy $u_1 < t_2$.*

See Appendix D for a proof.

3 Partial Disclosure Policies

We define a partial disclosure (PD) game, with the threshold Z, as a game in which a firm who waits learns additional information in the second period with positive probability. This information takes a simple form: if the rival firm drills in the first period, the laggard firm learns whether the true value of X is at least Z (good news), or below Z (bad news). If neither firm drills (no news), there is no additional information in the second period. Partial disclosure games nest both the full disclosure game, where $Z = x^*$, and the complete secrecy game, where $Z = \infty$.

3.1 Maximum nondistortionary disclosure

We provide an abbreviated analysis of the symmetric equilibrium of PD games by starting with interesting special case. To do this, we define the value of waiting in a partial disclosure game with threshold Z, when the equilibrium first period cutoff is $v_1(Z)$, as:

$$W(v_1(Z), Z) = W_{\text{good}}(v_1(Z), Z) + W_{\text{bad}}(v_1(Z), Z) + W_{\text{none}}(v_1(Z))$$

where

$$\begin{aligned} W_{\text{good}}(v_1(Z), Z) &= \mathbb{E}\left[\mathbb{I}(X \geq Z)\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right] \\ W_{\text{bad}}(v_1(Z), Z) &= \max(0, \mathbb{E}\left[\mathbb{I}(X < Z)\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right]) \\ W_{\text{none}}(v_1(Z)) &= \max(0, \mathbb{E}\left[\pi(X)F(v_1(Z) \mid X) \mid v_1(Z)\right]) \end{aligned}$$

The equilibrium first period cutoff signal must satisfy the usual condition:

$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid v_1(Z)\right] = W(v_1(Z), Z)$$

Though we can directly solve for the equilibrium v under any specific choice of Z, we can also ask whether there is a Z that satisfies this equilibrium condition at a specific signal $v = v_1(Z)$. Here, we will focus on finding a Z such that $v_1(Z) = u_1$, the first period cutoff under CS. This particular disclosure threshold is interesting because we know average welfare under a partial disclosure game with a first period cutoff equal to u_1 must be no lower than average welfare in a complete secrecy game. The two games have the same first period cutoffs, so any difference in welfare comes from firms who wait, and this is the same set of firms (signals) in both games. However, firms who wait in a partial disclosure game with a first period cutoff u_1 receive no less information than they do in a complete secrecy game, and unless $Z = \infty$ or $Z = 0$, they receive strictly more. Thus the value of waiting in a partial disclosure game with a first period cutoff equal to u_1 can't be lower than the value of waiting in the complete secrecy game, for firms who strictly prefer waiting.

As we show below, this Z exists, and u_1 is the first period cutoff for all values of Z at or above it.

Proposition 6

1. *The solution, in Z, to $0 = \mathbb{E}\left[\mathbb{I}(X \leq Z)\pi(x)(1 - F(u_1 \mid x)) \mid u_1\right]$ exists, is finite, and is greater than x^*. Call this solution $\overline{Z}$.*

2. *For $Z = \overline{Z}$, there is a symmetric equilibrium in cutoff strategies in which the first period cutoff, $v_1(Z)$, is equal to u_1.*

3. *For $Z > \overline{Z}$, the first period cutoff is also equal to u_1.*

To prove items 1 and 2, recall that in the complete secrecy game, a firm with signal u_1 can profitably drill after waiting if its rival drills first. That is, drilling under the knowledge of (a) your own signal is u_1 and (b) your rival's signal is at least u_1, is profitable. We can partition this information set, in which drilling is profitable on average, into two pieces: when $X \geq Z$, and when $X < Z$, for any choice of Z. Drilling will always be profitable under the $X \geq Z$ partition, and for high enough Z, drilling can be profitable in the $X < Z$ partition, for example, when $Z = \infty$. However, drilling may be unprofitable in the $X < Z$ partition if Z is low enough (e.g., it must be unprofitable when $Z = x^*$). This means that there must be some intermediate level of Z, which we define as $\overline{Z}$, at which drilling has exactly zero profits in the $X < Z$ partition. Because there are no profits to drilling in the $X < \overline{Z}$ partition, all of the profits to drilling after waiting in a complete secrecy game must be coming from the $X \geq \overline{Z}$ partition. However, these partitions are exactly the information structure available after waiting in a partial disclosure game with threshold $\overline{Z}$, so u_1 must be the first period cutoff for such a game.

To prove item 3, note that since a firm with signal u_1 is indifferent between drilling or not when it learns that its rival's signal is at least u_1 and $X < \overline{Z}$, the same firm will strictly prefer to drill when it learns that $X < Z$ for values of $Z > \overline{Z}$. Thus, a firm with signal u_1 will drill after waiting if it sees its rival drill first, regardless of whether X is above or below a threshold $Z > \overline{Z}$. We also know that a player with signal u_1 won't drill after observing its rival's signal is smaller than u_1. As a result, the value of waiting in a partial disclosure game with $Z > \overline{Z}$ for a player with signal u_1 is identical to the value of waiting in a complete secrecy game, and we already know that is equal to $\delta^{-1}\mathbb{E}\left[\pi(X) \mid u_1\right]$. This means that for $Z > \overline{Z}$, we must have $v_1(Z) = u_1$.

We call $\overline{Z}$ the maximum nondistortionary disclosure (MND) threshold because disclosure above this point has no consequence on the equilibrium first period cutoff signal. $\overline{Z}$ is the most that can be disclosed about X without increasing free-riding from the level attained in CS.

Despite the fact that the first period cutoff remains u_1 for any $Z \geq \overline{Z}$, it is nevertheless the case that behavior in the second period does depend on the specific value of Z.

Proposition 7

1. *The solution, in v, to $0 = \mathbb{E}\left[\mathbb{I}(X \leq Z)\pi(x)(1 - F(u_1 \mid x)) \mid v\right]$ exists, and is the minimum signal at which players who wait will drill after learning $X < Z$ (bad news). Call this cut-off signal $v_3(Z)$.*[7]

2. *For $Z \geq \overline{Z}$, $v_3(Z) \leq u_1$, so some signals that wait will drill after bad news.*

3. *For $Z \geq \overline{Z}$, all signals $s < u_1$ will drill after learning $X > Z$ (good news).*

4. *For $Z \geq \overline{Z}$, no signals $s < u_1$ will drill after observing their rival did not drill (no news).*

5. *For $Z \geq \overline{Z}$, average player welfare is strictly decreasing in Z.*

We provide proofs in appendix E.1.

Maximum nondistortionary disclosure ($\overline{Z}$) thus provides valuable information in environments that are currently operating under complete secrecy, with no equilibrium "cost." That is, a regulator currently collecting information on outcomes but not disclosing any information could instead disclose whether $X \geq \overline{Z}$ or not. The above results show that this would generate strictly positive welfare gains to players with pessimistic signals, while maintaining investment incentives for players with optimistic signals, and thus leaving free-riding

[7] We choose to call this $v_3(\cdot)$ in order to reserve $v_2(\cdot)$ as the label for a no-news drilling cutoff signal in the next section.

leaving unchanged. In section 5, we compute the MND thresholds in Pennsylvania and West Virginia and measure the counterfactual gains to using these policies in place of complete secrecy.

3.2 Partial disclosure for $Z < \overline{Z}$

At lower levels of Z, equilibrium play no longer involves a first period cutoff equal to u_1. Instead, the entire equilibrium structure depends on the level of disclosure.

Proposition 8

1. *For $x^* \leq Z < \infty$, there is a symmetric equilibrium in cutoff strategies in which the first period cutoff, $v_1(Z)$, is non-increasing in Z. For $Z \geq \overline{Z}$, $v_1(Z) = u_1$. For $Z < x^*$, $v_1(Z)$ is non-decreasing in Z.*

2. *For $x^* \leq Z < \overline{Z}$, the "bad news" cutoff $v_3(Z) \geq v_1(Z)$, so there is never any bad news drilling in equilibrium.*

3. *If there is no news drilling in the full disclosure game ($t_1 > t_2$), then there is a threshold $\widetilde{Z}$, with $x^* \leq \widetilde{Z} < \overline{Z}$, such that when $x^* \leq Z \leq \widetilde{Z}$, there is a "no news" second period cutoff $v_2(Z) < v_1(Z)$, and players with signals between $v_2(Z)$ and $v_1(Z)$ will drill after waiting and learning that their rival waited as well.*

4. *At $Z = x^*$, $v_1(Z) = t_1$ and $v_2(Z) = t_2$.*

Proofs for these results are provided in appendix E.

Proposition 8 shows how more disclosure (lower Z) implies more free-riding (higher $v_1(Z)$), a generalization of the result that $u_1 < t_1$. The fact that there is no bad news drilling in partial disclosure games with thresholds below $\overline{Z}$ means that the structure of partial disclosure equilibrium behavior is similar to full disclosure behavior, even though "bad news" may still reflect a profitable state of the world (e.g., when $x^* < X < Z$). Similarly, like some full disclosure games, partial disclosure games can exhibit no news drilling.

As in our comparison between full disclosure and complete secrecy, we can provide a partial characterization of the welfare differences between partial disclosure games with different disclosure thresholds.

Proposition 9

1. *Average player welfare is increasing in Z for $Z \in [0, x^*]$.*

2. *If there is some no-news drilling when $Z = x^*$, then average player welfare is decreasing in Z for $Z \in [\widetilde{Z}, \infty]$.*

3. *If there is no no-news drilling at $Z = x^*$, then average player welfare is maximized at $Z = x^*$.*

The mechanisms behind these results are similar to the logic in the proof for Proposition 4, and we provide proofs in the appendix. Proposition 9 shows that optimal disclosure is either at full disclosure (when there is no no-news drilling in FD) or somewhere between full disclosure and $\widetilde{Z}$.

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 most active gas plays in the world over the past two decades, the Marcellus and the Utica, underlie two states with starkly different disclosure policies during this time period: Pennsylvania and West Virginia (Figure 1).[8] 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 often quite disaggregated, with many firms owning mineral leases in the same small geographic area. Once a firm acquires a mineral leases from a landowner, it has a finite

[8] 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

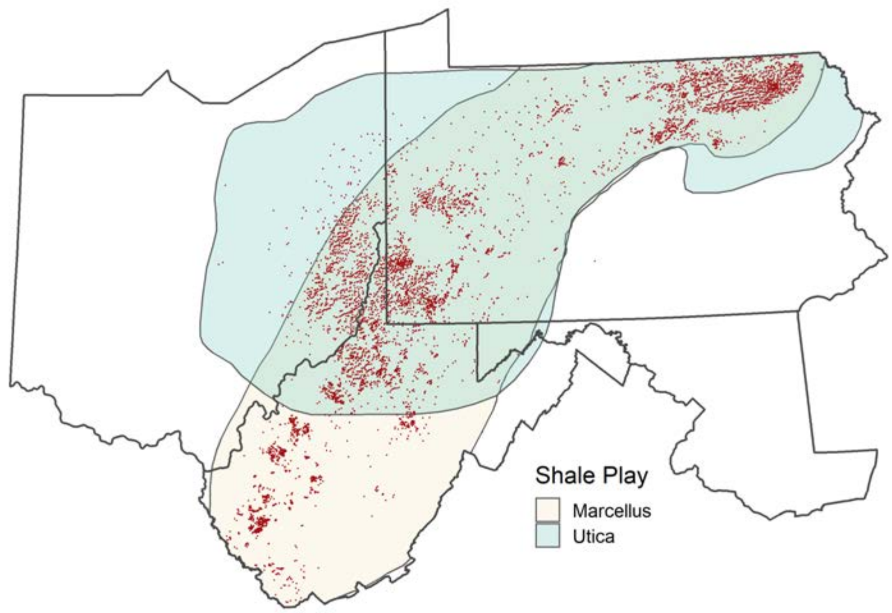

Shale well drilling (red dots) in Ohio, West Virginia, and Pennsylvania, between 2004 and 2020, relative to the extents of the Marcellus and Utica shale plays.

amount of time to 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 make 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 May 2010, the Pennsylvania State Legislature passed Senate Bill 297, a set of amendments to the original Pennsylvania Oil and Gas Act. 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 periods when the secrecy policies differ prior to 2011.[9] 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.[10]

4.2 Data

We acquired data on leasing, drilling, and production from DrillingInfo (DI), a major commercial provider of information to the oil and gas industry.[11] For each lease, DI 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. DI collects this information from county courthouses, and DI's 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 DI.

We combine this lease data with DI's comprehensive drilling and production data. DrillingInfo collects this information from state oil and gas regulators and organizes it into a standard format. The DI 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 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

[9] Brown et al. (2020) estimate the impact of state hydrocarbon taxes on extraction.

[10] Black et al. (2018) study the impact of introduction of the well fee in Pennsylvania.

[11] Recently DrillingInfo has re-branded itself as Enverus.

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.[12] 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.

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 policy. As a result, drilling behavior in these grids is unlikely to be accurately explained by a single disclosure model. To avoid 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 that firms likely had a common set of pre-existing knowledge about. Prior to the start of

[12] 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).

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 Figure 2. We focus our analyses on the parts of Pennsylvania and West Virginia covered by two of these regions, "Gas Plays" 12 and 13. The two regions were estimated to have similar average gas production potential, and they are the only regions which span both states.

Table 1 presents grid-level summary statistics. 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.

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 2. 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.

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 2.
Drilling outcomes restricted to wells spud by the end of 2010. "Drilled Early" refers to wells spud by the end of 2008.

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 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 (Figure 3). 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.[13] 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

[13] 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.

as the outcome. Pennsylvania grids drill 11% fewer wells, but obtain 80% less output per grid.

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.

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 exploration 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

Figure 2: Border radius sensitivity

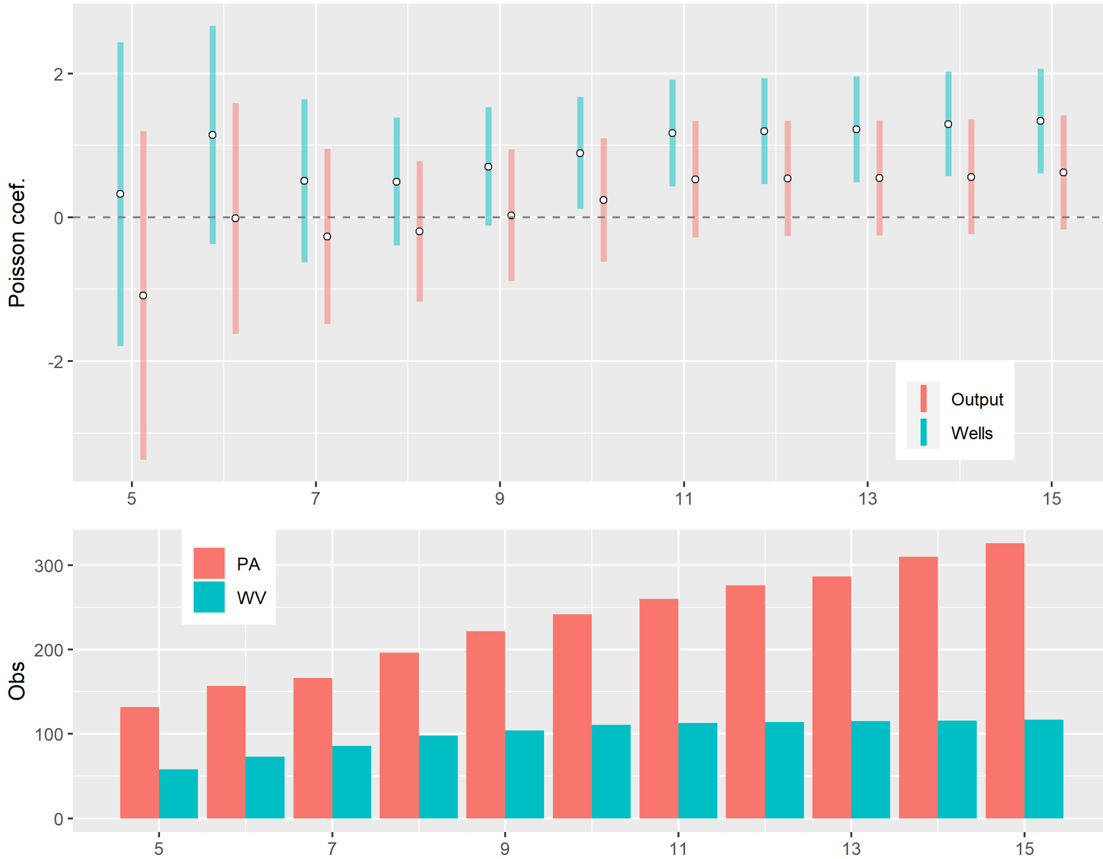

2011. Finally, in column 6 we match on both border index and lease year on this restricted sample. Across both of these matched samples, the incidence of early drilling and the number of wells per square mile is much higher. However, unlike the raw comparisons in Table 1, Pennsylvania wells are much 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 that 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 recreated in Figure 2. 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.[14] 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 hor-

[14]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 Appendix I.1 for details.

Figure 3: Sample grids and drilling

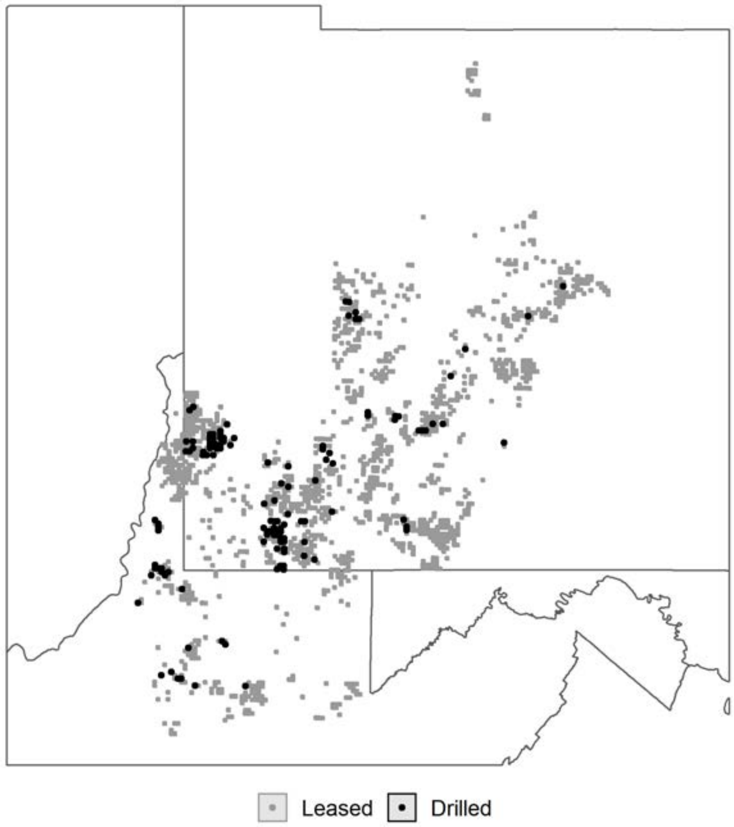

Leased grids are 33% leased by the end of 2005. Drilling indicates grids which have a shale well by the end of 2010.

izontal length of the well that was drilled, and fixed effects for the year in which drilling takes place.[15] We then compute the predicted output of that well, assuming it was drilled in 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.[16] At the end of this process, we have a prediction of average

[15] 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 data.

[16] This procedure is similar to the approaches used in Covert (2015), Agerton (2020), Hodgson (2021), and

resource quality, X, for each grid in our sample, including grids that have not yet been drilled by the end of 2010. Figure 4 presents a map of these predictions, and Figure A.1 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.[17] We set $P = \$3.1/\text{mcf}$, which is the median realized price for wells drilled during 2009. Royalty 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.[18] 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

Herrnstadt et al. (2020).

[17] For details, see Appendix I.1.

[18] 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.

$s = \log X + \epsilon$, with ϵ independent of X, and distributed normally, with zero mean and 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.[19] 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. 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.

Although we are specifically interested in estimating the primitives (ν, K), these 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 state.[20] 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 ex-

[19] 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.

[20] See appendix I for details.

ample, 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.[21] 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% production tax liability.[22]

To assess the suitability of these models for explaining drilling behavior in their respective

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

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

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.

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 2. 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.

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 Virginia, 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 value of playing an FD game, represented as a partial disclosure game with disclosure at that state's estimated x^*. The right-most part similarly represents the value 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 $v_1(Z)$ always lies above the no-news cutoff $v_2(Z)$ for all Z, a condition which guarantees that full disclosure is the optimal disclosure policy. The decrease in $v_1(Z)$ 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 $v_2(Z) < v_1(Z)$, 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.[23] 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.

As shown in Figure 4, our parameter estimates for both states generate equilibrium cutoff signals under both games that satisfy the $u_1 < t_2$ 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 $t_2 > t_1$

[23]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

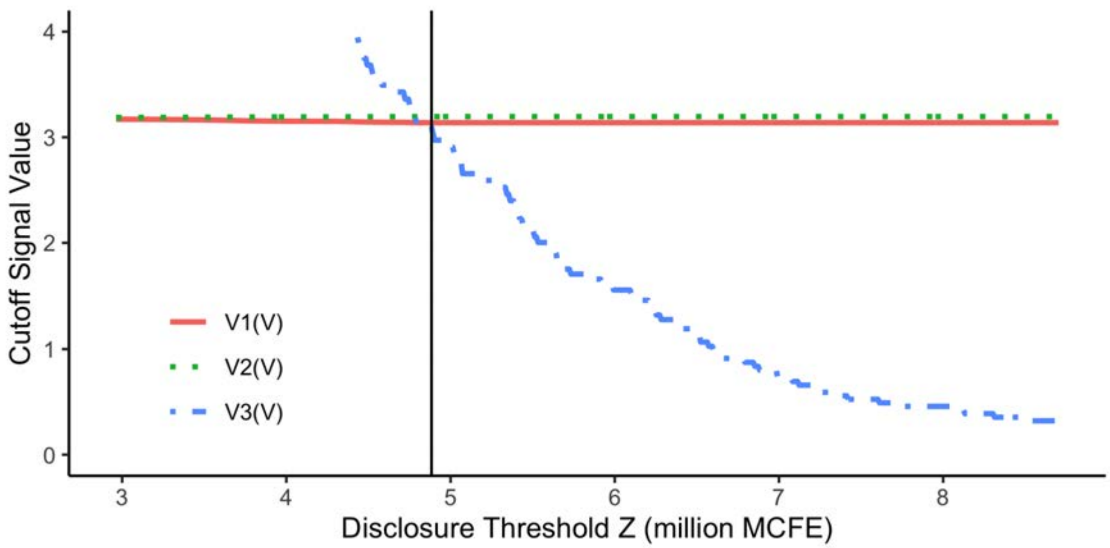

(b) West Virginia

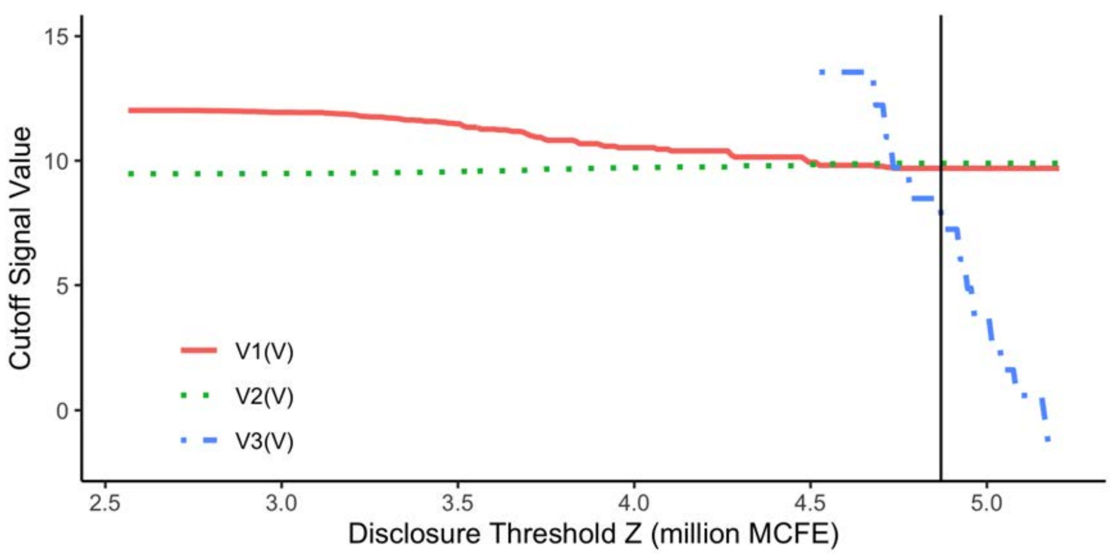

The solid red line is the first period cutoff, the dotted green line is the second period "no news" cutoff, and the dash-dotted blue line is the second period "bad news" cutoff. The bad news cutoff is especially high for low disclosure thresholds so it is not plotted for all values of Z in the figure.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

(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 Proposition 9.

Figure 5: Average Player Welfare Under Counterfactual Disclosure Policies

(a) Pennsylvania

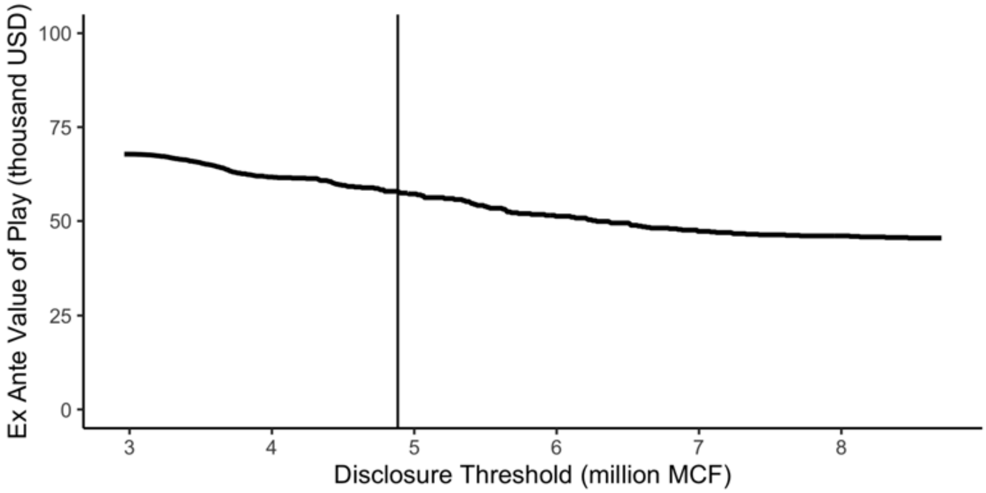

(b) West Virginia

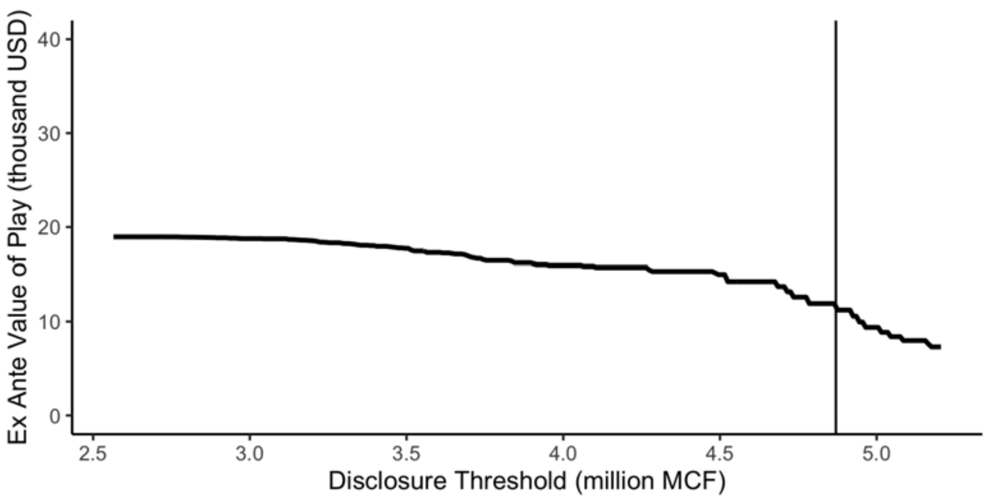

The black vertical line represents the location of minimal disclosure.

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.[24] 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% 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.

References

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

Black, Katie Jo, Shawn J McCoy, and Jeremy G Weber, "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*, 2018, *5* (1), 107–153.

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

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

Chamley, Christophe and Douglas Gale, "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica: journal of the Econometric Society*, 1994, *62* (5), 1065–1085.

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

Charpentier, Ronald R, Wallace de Witt, GE Claypool, LD Harris, RF Mast, JD Megeath, JB Roen, JW Schmoker, and RC Kepferle, "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)*, 1993.

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

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

Friedman, David D, William M Landes, and Richard A Posner, "Some Economics of Trade Secret Law," *The journal of economic perspectives: a journal of the American Economic Association*, 1991, *5* (1), 61–72.

Grenadier, Steven R, "Information Revelation Through Option Exercise," *The review of financial studies*, January 1999, *12* (1), 95–129.

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

Hendricks, Kenneth and Dan Kovenock, "Asymmetric Information, Information Externalities, and Efficiency: The Case of Oil Exploration," *The Rand Journal of Economics*, 22 1989, *20* (2), 164.

_ **and Robert H. Porter**, "The Timing and Incidence of Exploratory Drilling on Offshore Wildcat Tracts," *The American Economic Review*, June 1996, *86* (3), 388–407.

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

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

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

Keller, Godfrey, Sven Rady, and Martin Cripps, "Strategic experimentation with exponential bandits," *Econometrica*, 2005, *73* (1), 39–68.

Krieger, Joshua L, "Trials and terminations: Learning from competitors' R&D failures," *Management Science*, 2021, *67* (9), 5525–5548.

Lin, C-Y Cynthia, "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*, 2013, *95* (5), 1601–1621.

Milgrom, Paul R, "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380–391.

Rob, Rafael, "Learning and Capacity Expansion under Demand Uncertainty," *The Review of Economic Studies*, June 1991, *58* (4), 655–675.

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

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

U.S. Energy Information Administration, "Trends in U.S. Oil and Natural Gas Upstream Costs," Technical Report March 2016.

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

Wang, Zhongmin and Alan Krupnick, "A retrospective review of shale gas development in the United States: What led to the boom?," *Economics of Energy & Environmental Policy*, 2015, *4* (1), 5–18.

Appendix A Proof of Lemma 1

$$
\begin{aligned}
\mathbb{E}\left[\pi(X) \mid Q=q, Q' \leq p\right] &= \int_0^{\infty} \pi(x) \Pr(X=x \mid Q=q, Q' \leq p) dx \\
&= \frac{\int_0^{\infty} \pi(x) \Pr(Q=q, Q' \leq p \mid X=x) \Pr(X=x) dx}{\int_0^{\infty} \Pr(Q=q, Q' \leq p \mid X=z) \Pr(X=z) dz} \\
&= \frac{\int_0^{\infty} \pi(x) \Pr(Q=q \mid X=x) \Pr(Q' \leq p \mid X=x) \Pr(X=x) dx}{\int_0^{\infty} \Pr(Q=q \mid X=z) \Pr(Q' \leq p \mid X=z) \Pr(X=z) dz} \\
&= \frac{\int_0^{\infty} \pi(x) f(q \mid x) F(p \mid x) \Pr(X=x) dx}{\int_0^{\infty} f(q \mid z) F(p \mid z) \Pr(X=z) dz} \\
&= \frac{\int_0^{\infty} \pi(x) h(x \mid q) F(p \mid x) dx}{\int_0^{\infty} h(z \mid q) F(p \mid z) dz} \\
&= \frac{\mathbb{E}\left[\pi(X) F(p \mid X) \mid q\right]}{\mathbb{E}\left[F(p \mid X) \mid q\right]}
\end{aligned}
$$

The second line comes from the first as an application of Bayes' rule. The third line comes from the fact that signals are *iid*, conditional on X. The fourth line is cosmetic, changing "probabilities" to densities and distribution functions. The fifth line is an additional application of Bayes' rule, and the final line is a simplification of that result.

Appendix B Proofs of the CS equilibrium

A player with this signal must be indifferent between drilling in the first period and waiting until the next period to decide whether or not to drill:

$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid u_1\right] = \mathbb{E}\left[F(u_1 \mid X) \mid u_1\right] \max\left(0, \frac{\mathbb{E}\left[\pi(X)F(u_1 \mid X) \mid u_1\right]}{\mathbb{E}\left[F(u_1 \mid X) \mid u_1\right]}\right) + \\ \mathbb{E}\left[1 - F(u_1 \mid X) \mid u_1\right] \max\left(0, \frac{\mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_1\right]}{\mathbb{E}\left[1 - F(u_1 \mid X) \mid u_1\right]}\right)$$

The first object on the right-hand side of this expression is the probability that your opponent's signal is less than u_1 times the payoff to optimal behavior when it turns out that your opponent's signal is less than u_1. Similarly, the second object is the probability your opponent's signal is greater than u_1 times the payoff to optimal behavior when it turns out that your opponent's signal is greater than u_1.

To simplify the right-hand side expression, note that the first maximum operator is binding. That is, $\frac{\mathbb{E}[\pi(X)F(u_1|X)|u_1]}{\mathbb{E}[F(u_1|X)|u_1]} < 0$. Why? Suppose not. This is the expectation of $\pi(X)$, conditional on your signal being exactly u_1 and the knowledge that your rival's signal is less than or equal to u_1. Because the distribution of X conditional your and your rival's signals has the MLRP in both signals, we can apply Lemma 2. The "news" that your opponent's signal is less than or equal to u_1 is worse than the news that your opponent's signal is greater than u_1, so the expected value of $\pi(X)$ conditional on your opponent's signal being below u_1 is less than or equal to the expected value when your opponent's signal is above t_1. Thus if we've assumed that $\frac{\mathbb{E}[\pi(X)F(u_1|X)|u_1]}{\mathbb{E}[F(u_1|X)|u_1]} \geq 0$, then $\frac{\mathbb{E}[\pi(X)(1-F(u_1|X))|u_1]}{\mathbb{E}[1-F(u_1|X)|u_1]} \geq 0$, too. In that case, the right maximum of the right-hand side expression in the definition of t_1 is also non-binding. Then, the entire right-hand side of the t_1 condition is simply $\mathbb{E}\left[\pi(X) \mid u_1\right]$, the *undiscounted* version of the left hand side. But the only way for that to hold is for $\delta = 1$. Thus, $\frac{\mathbb{E}[\pi(X)F(u_1|X)|u_1]}{\mathbb{E}[F(u_1|X)|u_1]} < 0$, and we can simplify our expression for $_1$ to:

$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid u_1\right] = \mathbb{E}\left[1 - F(u_1 \mid X) \mid u_1\right] \max\left(0, \frac{\mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_1\right]}{\mathbb{E}\left[1 - F(u_1 \mid X) \mid u_1\right]}\right)$$

We can do one more step of this process and simplify this even further. Suppose it was the case that at the equilibrium value for u_1, the remaining maximum was binding: the expected value of $\pi(X)$ conditional on your signal being u_1 and the knowledge that your opponent's signal was bigger than u_1 was negative. This would imply that the *worse* news, that your opponent's signal was only bigger than or equal to $-\infty$ (what we effectively have on the left hand side), must be at least as negative. However, the right-hand side is wrapped inside a maximum operator, which does not fall below 0, so we have a contradiction. It must be the case that $\frac{\mathbb{E}[\pi(X)(1-F(u_1|X))|u_1]}{\mathbb{E}[1-F(u_1|X)|u_1]} \geq 0$, and we can simply to:

$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid u_1\right] = \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_1\right] \tag{3}$$

A firm with the equilibrium cutoff signal u_1 is indifferent between drilling today and waiting until tomorrow to drill. If that firm waits, and does learn its rival drilled first, then it drills as well. If it waits and learns its rival didn't drill, it also won't drill. Thus, if neither firm drills in the first period, there is no drilling in the second period.

A firm with signal $s < u_1$ will wait to see what its rival does. As above, it won't drill after learning that its rival didn't drill. But, if it does see its rival drill, and as such updates its beliefs about X upwards, it is possible that it will drill in the second period. In order for this to be profitable, the signal s must be high enough. In particular, we need $s \geq u_2$, given by:

$$\mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_2\right] = 0. \tag{4}$$

It is clear from this definition that $u_2 \leq u_1$. If not, then the expected payoff to drilling with signal u_2 upon observation that a rival already drilled, and thus had a signal at or above u_1, would be larger than the same payoff with signal u_1. However, u_2 is selected to make this payoff equal to zero, while we already saw that at u_1, this payoff is positive. Thus, $u_2 \leq u_1$.

Appendix C Proof of Proposition 4

To see why $V^{\text{Full Disclosure}}(s) \geq V^{\text{Complete Secrecy}}(s)$, we will examine the payoffs to signals in each of the intervals contained in the partition $-\infty < u_2 \leq u_1 \leq t_1 < \infty$. First, note that firms with especially high signals $s \geq t_1$ drill early in both games, so for them, $V^{\text{Full Disclosure}}(s) = V^{\text{Complete Secrecy}}(s)$. Firms with moderately high signals, those with $u_1 \leq s < t_1$, will drill early in the complete secrecy game, but will wait in the full disclosure game. For this range of signals, we know that the discounted value of waiting under full disclosure game is more valuable than drilling immediately, so we can conclude that $V^{\text{Full Disclosure}}(s) > V^{\text{Complete Secrecy}}(s)$. Players with especially low signals, $s \leq u_2$, will wait in either game. However, a player with $s < u_2$ in the complete secrecy game will never drill after waiting, while a player with the same signal might *efficiently* drill after waiting in the full disclosure game. Thus, for signals in this range, $V^{\text{Full Disclosure}}(s) > V^{\text{Complete Secrecy}}(s)$.

Finally, consider the case where $u_2 < s \leq u_1$. In this range, firms wait in both games, and firms may drill in the second period in both games, so we must directly examine the expected difference in the value of waiting between the games. Because we have assumed that $u_1 < t_2$, a firm with a signal in this range will not do no-news drilling in the full disclosure game, which allows us to simplify its value of waiting to:

$$W^{\text{Full Disclosure}}(s) = \mathbb{E}\left[\max(0, \pi(X))(1 - F(s_1 \mid X)) \mid s\right]$$

With this simplification, we can write the difference between the expected value of waiting in the two games as:

$$\begin{aligned}
\Delta_1(s) &= W^{\text{Complete Secrecy}}(s) - W^{\text{Full Disclosure}}(s) \\
&= \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid s\right] - \mathbb{E}\left[\max(0, \pi(X))(1 - F(t_1 \mid X)) \mid s\right] \\
&= \mathbb{E}\left[\pi(X)\left(\mathbb{I}(X < x^*)(1 - F(u_1 \mid X) + (F(t_1 \mid X) - F(u_1 \mid X))\right)\right) \mid s] \\
&= \mathbb{E}\left[\phi_1(X) \mid s\right]
\end{aligned}$$

where x^* is the value of output x for which $\pi(x) = 0$. For values of $x < x^*$, the function $\phi_1(x) < 0$, and for values of $x \geq x^*$, $\phi_1(x) \geq 0$, so $\phi_1(x)$ has a single crossing of 0 at the point $x = x^*$. We already know that at $s = u_1$, the value of drilling is equal to the value of waiting in the complete secrecy game, but less than the value of waiting in full disclosure game, so $\Delta_1(u_1) \le 0$. Firms with *worse* signals $s < u_1$ put more probability mass on *lower* values of X than firms with $s = u_1$, so $\Delta_1(s) \le 0$ as well.[25]. Thus, for $u_2 < s \le u_1$ we have $V^{\text{Full Disclosure}}(s) > V^{\text{Complete Secrecy}}(s)$. This completes the proof.

N.B. When $u_1 \ge t_2$, the above logic no longer holds. To see why, consider the difference in the value of waiting for firms who can do no news drilling in the full disclosure game ($t_2 \le s < u_1$):

$$\begin{aligned}
\Delta_2(s) &= W^{\text{Complete Secrecy}}(s) - W^{\text{Full Disclosure}}(s) \\
&= \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid s\right] - \mathbb{E}\left[\max(0, \pi(X))(1 - F(t_1 \mid X)) \mid s\right] - \mathbb{E}\left[\pi(X)F(u_1 \mid X) \mid s\right] \\
&= \mathbb{E}\left[\pi(X)(\mathbb{I}(X < x^*)(1 - F(t_1 \mid X)) - F(u_1 \mid X)) \mid s\right] \\
&= \mathbb{E}\left[\phi_2(X) \mid s\right]
\end{aligned}$$

To characterize the function $\phi_2(x)$, let $\widetilde{x}$ be the value of x at which the functions $1 - F(t_1 \mid x)$ (which is increasing) and $F(u_1 \mid x)$ (which is decreasing) cross, such that $1 - F(t_1 \mid \widetilde{x}) = F(u_1 \mid \widetilde{x})$. Then $\phi_2(x) \ge 0$ for $x \le \min(x^*, \widetilde{x})$, and $\phi_2(x) \le 0$ for $x > \min(x^*, \widetilde{x})$. Thus, $\phi_2(x)$ is positive for low realizations of X and negative for high realizations. We already know that $\Delta_2(u_1) = \mathbb{E}\left[\phi_2(X) \mid u_1\right] \le 0$ by revealed preference: players with signal u_1 would rather wait than drill in the full disclosure game. However, unlike the case above, the single crossing function $\phi_2(x)$ is negative for good x's and positive for bad x's, so we can't rule out the possibility that there are signals $s \in [t_2, u_1)$ for which $\Delta_2(s) \ge 0$, who *prefer* a complete secrecy information policy to a full disclosure policy.

Appendix D Proof of Proposition 5

We first prove that first period cutoffs are monotonically increasing in δ. Then, we show the key impatience result.

D.1 FD and CS cutoffs are monotonically increasing in δ

To show that the both FD and CS cutoffs are monotonically increasing in δ, first note that its impossible for either first period cutoff to be non-monotonic in δ. It it was, then there would have to be two distinct values of δ, e.g., δ^1 and δ^2, with equal first period cutoffs in, for example, the CS game, such that $u_1^1 = u_1^2 = u$. This would imply $\delta^1\mathbb{E}\left[\pi(X)(1 - F(u \mid X)) \mid u\right] = \delta^2\mathbb{E}\left[\pi(X)(1 - F(u \mid X)) \mid u\right]$, and this isn't possible unless $\mathbb{E}\left[\pi(X)(1 - F(u \mid X)) \mid u\right] = 0$, which cannot occur if there is to be any drilling at all, so u_1 must be strictly monotonic in δ. A similar argument applies for the cutoffs in the FD game.

Next, note that in the limit as $\delta \to 0$, we must have $u_1 = s_0$, where $\mathbb{E}\left[\pi(X) \mid s_0\right] = 0$, while in the limit as $\delta \to 1$, we must have $u_1 = s_1$, where $\mathbb{E}\left[\pi(X) \mid s_1\right] = \left[\pi(X)(1 - F(s_1 \mid X)) \mid s_1\right]$. Because $0 = \mathbb{E}\left[\pi(X) \mid s_0\right] < \mathbb{E}\left[\pi(X)(1 - F(s_1 \mid X)) \mid s_1\right] = \mathbb{E}\left[\pi(X) \mid s_1\right]$, we must have $s_0 < s_1$. This means u_1 must be strictly monotonically increasing in δ.

In the FD game, in the limit as $\delta \to 0$, we also must have $t_1 = s_0$, where $\mathbb{E}\left[\pi(X) \mid s_0\right] = 0$. However, in the limit as $\delta \to 1$, t_1 becomes unbounded and tends to ∞, as there becomes no delay cost to waiting. By the same logic as we saw in the CS case, this implies that t_1 must be strictly monotonically increasing in δ.

[25]This essentially an application of Lemma 2 of Hendricks and Kovenock (1989)

D.2 Proof of the impatience result

As we saw above, when firms are completely impatient, as $\delta \to 0$, $t_1 \to s_0$. At this point, any signal at which drilling after no news is profitable would have to be larger than s_0, so we know $t_2 > u_1$. At the other extreme, as $\delta \to 1$ and firms are perfectly patient, the first period cutoff for FD tends to ∞, as there is no delay cost for any signal to waiting. This means t_2 satisfies $0 = \mathbb{E}\left[\pi(X)F(\infty \mid X) \mid t_2\right] = \mathbb{E}\left[\pi(X) \mid t_2\right]$, so $t_2 \to s_0$. However, the first period cutoff for CS remains finite, at the u_1 which satisfies $\mathbb{E}\left[\pi(X) \mid u_1\right] = \mathbb{E}\left[\pi(X)(1 - F(u_1 \mid X)) \mid u_1\right]$. We can rearrange this expression to obtain $0 = \mathbb{E}\left[\pi(X)F(u_1 \mid X) \mid u_1\right]$ so that at $\delta = 1$ (and only there), the cutoff signal is indifferent to drilling after no news in a CS game, and this signal must be higher than the zero profit signal s_0.

At the (limit of the) lowest possible value for δ, we know $u_1 < t_2$, and at the (limit of the) highest possible value for δ, we know $u_1 > t_2$. Since u_1 is increasing in δ and t_2 is decreasing in δ, and since everything in this setup is continuous, there must be a $\overline{\delta} \in (0, 1)$ at which $u_1 = t_2$, and for $\delta < \overline{\delta}$, we must have $u_1 < t_2$. This completes the proof.

Appendix E Proofs of partial disclosure propositions

E.1 Proof of Proposition 7

When $Z > \overline{Z}$ we know that a player with signal u_1 can profitably drill under bad news, that $X < Z$, because that player gets zero profits from drilling under bad news in the $Z = \overline{Z}$ case. By continuity, this implies the existence of a signal smaller than u_1 that earns zero profits from drilling after waiting and learning that its rival's signal is at or above u_1, and that $X < Z$, for $Z > \overline{Z}$. This proves items 1 and 2.

Because good news involves X values strictly larger than $\overline{Z}$, and because we know $\overline{Z} > x^*$, any signal can profitably drill after good news. This proves item 3.

From our analysis of the complete secrecy game, we know a player with signal u_1 can't earn strictly positive profits from drilling after no news when the cutoff signal is u_1, so no smaller signals can either. Since the cutoff signal remains at u_1 for $Z > \overline{Z}$, we know that all signals who wait in such partial disclosure games will not drill after no news. This proves item 4.

To prove item 5, note that because $v_1(Z) = u_1$ for $Z \geq \overline{Z}$, variation in $Z \geq \overline{Z}$ only affects the value of waiting, so it is sufficient to show that the value of waiting is decreasing in Z. To do this, write out the value of waiting for a player in a partial disclosure game with signal s and threshold Z as:

$$\begin{aligned} W^{\text{Partial Disclosure}}(s, Z) = {} & \mathbb{E}\left[\pi(X)\mathbb{I}(X \geq Z)(1 - F(u_1 \mid X)) \mid s\right] \\ & + \max\left(0, \mathbb{E}\left[\pi(X)\mathbb{I}(x < Z)(1 - F(u_1 \mid X)) \mid s\right]\right) \end{aligned}$$

For values of s where the maximum operator in the second expression is binding ($s < v_3(Z)$), the value of waiting is decreasing in Z because $\pi(x)$ is positive for all $x \geq Z \geq \overline{Z} > x^*$. For higher values of s, when the maximum operator is non-binding, the expression reduces to:

$$W^{\text{Partial Disclosure}}(s, Z) = \mathbb{E}\left[\pi(x)(1 - F(t_1 \mid x)) \mid s\right]$$

and this is independent of Z. Thus for any signal that waits, the value of waiting is non-decreasing in Z.

E.2 First period cutoffs when $x^* \leq Z < \overline{Z}$

When Z is smaller than $\overline{Z}$, it is no longer the case that the first period cutoff is equal to u_1, so the structure of the equilibrium must change. To analyze this, let us return to the equilibrium definition:

$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid v_1(Z)\right] = W(v_1(Z), Z)$$

where

$$W(v_1(Z), Z) = W_{\text{good}}(v_1(Z), Z) + W_{\text{bad}}(v_1(Z), Z) + W_{\text{none}}(v_1(Z))$$

and the individual waiting values are given by:

$$\begin{aligned}
W_{\text{good}}(v_1(Z), Z) &= \mathbb{E}\left[\mathbb{I}(X \geq Z)\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right] \\
W_{\text{bad}}(v_1(Z), Z) &= \max(0, \mathbb{E}\left[\mathbb{I}(X < Z)\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right]) \\
W_{\text{none}}(v_1(Z)) &= \max(0, \mathbb{E}\left[\pi(X)F(v_1(Z) \mid X) \mid v_1(Z)\right])
\end{aligned}$$

We first consider whether it is possible for the firm with the cutoff signal $v_1(Z)$ to drill *unconditionally* in the second period if it observes that its rival drilled in the first period. If that were true, then $W_{\text{bad}}(v_1(Z), Z) > 0$, and the definition for $v_1(Z)$ would collapse to the definition we've already seen for u_1. The solution to that definition is, of course, u_1, and at u_1, a firm who learns its rival has a signal higher than u_1, but that $X \leq Z < \overline{Z}$, earns *negative* profits from drilling. This contradicts our assumption that the firm would drill unconditionally in the second period if its rival drills first. Thus, in partial disclosure games where $x^* \leq Z < \overline{Z}$, the firm with the cutoff type *does not* drill if it waits, learns that its rival drilled, and that $X < Z$. This lets us simplify the equilibrium condition one step further to:

$$\delta^{-1}\mathbb{E}\left[\pi(X) \mid v_1(Z)\right] = \mathbb{E}\left[\mathbb{I}(X \geq Z)\pi(x)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right] + W_{\text{none}}(v_1(Z))$$

With this simplification, we can show that for $x^* \leq Z < Z' \leq \overline{Z}$, $v_1(Z) > v_1(Z')$, or more simply, that $v_1(Z)$ is decreasing over that interval. Again, assume the opposite, that $v_1(Z) \leq v_1(Z')$. If this were true, then a player with signal $v_1(Z)$ would wait in a game with threshold Z', while the same player would drill in a game with threshold Z. Since the value of drilling is the same in either scenario, this would imply that the value of waiting with signal $v_1(Z)$ in the partial disclosure game with threshold Z is smaller than the value of waiting with signal $v_1(Z)$ in the game with threshold Z':

$$\begin{aligned}
0 > \mathbb{E}\left[\mathbb{I}(X \geq Z)\pi(X)(1 - F(v_1(Z))) \mid v_1(Z)\right] - \left[\mathbb{I}(X \geq Z')\pi(X)(1 - F(v_1(Z'))) \mid v_1(Z)\right] \\
+ \left(\max\left(0, \mathbb{E}\left[\pi(X)F(v_1(Z) \mid X) \mid v_1(Z)\right]\right) - \max\left(0, \mathbb{E}\left[\pi(X)F(v_1(Z') \mid X) \mid v_1(Z)\right]\right)\right)
\end{aligned}$$

To see if this assumption is feasible, we must examine three separate cases regarding which of the maximums are binding. We will see in each of these cases that the assumption we've made, that $v_1(Z) \leq v_1(Z')$, leads to a contradiction.

E.2.1 First (and second) maximum non-binding

If the first maximum is non-binding, then because we are assuming that $v_1(Z) \leq v_1(Z')$, the second is non-binding as well, leaving us with:

$$\begin{aligned}0 > &\mathbb{E}\left[\mathbb{I}(Z \leq X \leq Z')\pi(X)(1 - F(v_1(Z') \mid X)) \mid v_1(Z)\right] \\ &- \mathbb{E}\left[\mathbb{I}(X < Z)\pi(X)(F(v_1(Z') \mid X) - F(z_1(Z) \mid X)) \mid v_1(Z)\right]\end{aligned}$$

The first term in the right-hand side expression is positive, since $Z \geq x^*$. By Lemma 2, the second term is bounded above by $\mathbb{E}\left[\mathbb{I}(X < Z)\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right]$, and we know from the previous discussion that a firm with signal $v_1(Z)$ does not drill when it observes its rival did drill, but that $X < Z$. Thus, the second term is negative. As a result, the entire right-hand side is positive. But that can't be smaller than zero, so we have a contradiction, and must conclude that $v_1(Z) > v_1(Z')$.

E.2.2 Both maximums binding

When both $\mathbb{E}\left[\pi(X)F(v_1(Z) \mid X) \mid v_1(Z)\right] < 0$ and $\mathbb{E}\left[\pi(X)F(v_1(Z') \mid X) \mid v_1(Z)\right] < 0$, the expression simplifies to:

$$\begin{aligned}0 > &\mathbb{E}\left[\mathbb{I}(X \geq Z')\pi(X)(F(v_1(Z') \mid X) - F(v_1(Z) \mid Z)) \mid v_1(Z)\right] \\ &+ \mathbb{E}\left[\mathbb{I}(Z \leq X < Z')\pi(X)(1 - F(v_1(Z) \mid Z)) \mid v_1(Z)\right]\end{aligned}$$

Both terms in the right-hand side of this expression are positive under the assumption that $v_1(Z') \geq v_1(Z)$, so we again have a contradiction, and must conclude that $v_1(Z) > v_1(Z')$.

E.2.3 Only second maximum binding

Finally, when $\mathbb{E}\left[\pi(X)F(v_1(Z) \mid X) \mid v_1(Z)\right] < 0$ but $\mathbb{E}\left[\pi(X)F(v_1(Z') \mid X) \mid v_1(Z)\right] \geq 0$, the (simplified) expression would be

$$\begin{aligned}0 > &\mathbb{E}\left[\mathbb{I}(Z \leq X \leq Z')\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right] \\ &- \mathbb{E}\left[\pi(X)F(v_1(Z) \mid X) \mid v_1(Z)\right] \\ &- \mathbb{E}\left[\mathbb{I}(X \leq Z')\pi(X)(F(v_1(Z') \mid X) - F(v_1(Z) \mid X)) \mid v_1(Z)\right]\end{aligned}$$

The first term on the right-hand side is positive because $Z \geq x^*$ and we are assuming that the second term is negative, so the net effect of the first two terms is positive. To see why the third term is also negative, start from the payoff to drilling for a player with signal $v_1(Z')$ who learns that its rival's signal is above $v_1(Z')$ but that $X < Z'$. We know this payoff is negative, and can show that its larger than the third term we are trying to sign:

$$\begin{aligned}0 &> \mathbb{E}\left[\mathbb{I}(X \leq Z')\pi(X)(1 - F(v_1(Z') \mid X)) \mid v_1(Z')\right] \\ &> \mathbb{E}\left[\mathbb{I}(X \leq Z')\pi(X)(1 - F(v_1(Z') \mid X)) \mid v_1(Z)\right] \\ &> \mathbb{E}\left[\mathbb{I}(X \leq Z')\pi(X)(1 - F(v_1(Z) \mid X)) \mid v_1(Z)\right] \\ &> \mathbb{E}\left[\mathbb{I}(X \leq Z')\pi(X)(F(v_1(Z') \mid X) - F(v_1(Z) \mid X)) \mid v_1(Z)\right]\end{aligned}$$

The first inequality holds because we already saw that a firm with the marginal signal does not drill on bad news when $Z < \overline{Z}$. The second inequality holds because drilling with signal

$v_1(Z)$, under any news about a rival's signal, is worse than drilling with signal $v_1(Z')$, under our assumption that $v_1(Z') > v_1(Z)$. The third inequality holds because news that the rival signal is better than $v_1(Z)$ is worse news than when the rival signal is better than $v_1(Z')$. The final inequality holds because the news that the rival's signal is between $v_1(Z')$ and $v_1(Z)$ is worse than the news that it is better than $v_1(Z)$. We must conclude that our assumption that $v_1(Z') \geq v_1(Z)$ is wrong, and that, instead, we have $v_1(Z) > v_1(Z')$.

This completes the proof of Proposition 8, part 1, for the case where $Z \geq x^*$.

E.3 First period cutoffs when $Z < x^*$

The exact same set of steps above will establish that the equilibrium condition for $v_1(Z)$ when $Z < x^*$ is identical to the equilibrium condition when $x^* \leq Z < \overline{Z}$. Similarly, we can employ the same contradiction approach to show that $v_1(Z)$ is increasing in the interval $[0, x^*]$. This proves Proposition 8, part 1, for the case where $Z < x^*$.

Note that in the limiting case of $Z = 0$, the equilibrium condition for $v_1(Z)$ collapses to the familiar form for the equilibrium condition for u_1, so we additionally know that the lowest value $v_1(Z)$ can take is u_1.

E.4 Equilibrium behavior for players with $s < v_1(Z)$

Firms with signals lower than $v_1(Z)$ wait. Like firms with the cutoff signal, if they observe good news, they drill.[26] Because we saw above that players with the marginal signal, $v_1(Z)$, cannot profitably drill under bad news, firms with lower signals also can't, and as such we know that $v_3(Z) > v_1(Z)$ for $Z < \overline{Z}$ (Proposition 8, part 2).

When a player with a signal below $v_1(Z)$ learns that its rival also has a signal below $v_1(Z)$, e.g. there is no news, that player drills if (a) $v_1(Z)$ isn't too low, so that learning that their rival had a signal below $v_1(Z)$ isn't sufficiently bad, and (b) their own signal is high enough to justify drilling. Formally, there is a second period cutoff $v_2(Z)$, similar to the cutoff t_2 in the full disclosure game, whenever the solution $v_2(Z)$ to the following indifference condition is smaller than $v_1(Z)$:

$$\mathbb{E}\left[\pi(X)F(v_1(z) \mid X) \mid v_2(Z)\right] = 0$$

Like in the full disclosure game, the solution $v_2(Z)$ to the above equation might be larger than $v_1(Z)$, in which case there is no late drilling after "no news." When this can happen for some values of Z but not others, we define $\widetilde{Z}$ as the lowest value of Z at which no news drilling becomes impossible, i.e., the Z which satisfies $v_1(Z) = v_2(Z)$, or, equivalently:

$$\mathbb{E}\left[\pi(X)F(v_1(\widetilde{Z}) \mid X) \mid v_1(\widetilde{Z})\right] = 0$$

For $x^* < Z \leq \widetilde{Z}$, a disclosure threshold of Z will induce some no news drilling, for players with signals smaller than $v_1(Z)$ but larger than $v_2(Z)$.

Finally, given that we know $v_1(Z)$ is decreasing in Z, and that everything in the equilibrium condition for $v_1(Z)$ is continuous, we must have $v_1(x^*) = t_1$ and $v_2(x^*) = t_2$, which

[26] This assumes $Z \geq x^*$. If not, there will be another cutoff signal, similar in spirit to the bad news cutoff $v_3(Z)$, above which players can profitably drill on good news. This will be lower than $v_1(Z)$.

proves Proposition 8, part 4.

E.5 Welfare maximizing partial disclosure

E.5.1 Welfare maximizing Z is no smaller than x^*

We prove the first part of Proposition 9 by taking advantage of the fact that $v_1(Z)$ is increasing on the interval $[0, x^*]$, decreasing on the interval $[x^*, \overline{Z}]$, and that $v_1(0) = v_1(\overline{Z}) = u_1$. Because $v_1(Z)$ is continuous, this means that for any value of $0 \leq Z < x^*$, we can find a Z' that is larger than x^* with the property that $v_1(Z) = v_1(Z')$. This, in turn, means that we can easily compare the expected player welfare of a policy $Z < x^*$ with the corresponding $Z' > x^*$. In both cases, players who drill early experience the same level of welfare, so it suffices to compare the value of waiting for players who wait under Z and Z'. As we'll see, the expected value of waiting under Z' is always greater than the expected value of waiting under Z.

Let $v = v_1(Z) = v_1(Z')$. The values of waiting for a player with signal $s < z$, playing in the Z and Z' games, are:

$$W^Z(s) = \mathbb{E}\left[\mathbb{I}(X \geq Z)\pi(x)(1 - F(z \mid X)) \mid s\right] + \max\left(0, \mathbb{E}\left[\pi(X)F(z \mid X) \mid s\right]\right)$$
$$W^{Z'}(s) = \mathbb{E}\left[\mathbb{I}(X \geq Z')\pi(x)(1 - F(z \mid X)) \mid s\right] + \max\left(0, \mathbb{E}\left[\pi(X)F(z \mid X) \mid s\right]\right)$$

Regardless of whether the player is in the Z or the Z' game, the term inside the maximum will be the same, so we can drop it from our consideration. Thus, the difference in waiting values is

$$\begin{aligned}\Delta(s) &= \mathbb{E}\left[\mathbb{I}(Z \leq X \leq Z')\pi(X)(1 - F(z \mid X)) \mid s\right] \\ &= \mathbb{E}\left[\phi(X) \mid s\right]\end{aligned}$$

The function $\phi(\cdot)$ crosses 0 exactly once at x^*, and is negative below x^*, while its positive above x^*. Additionally, we know a player with signal z must be indifferent between drilling and waiting in either game, and that drilling has the same value in both games, then $\mathbb{E}\left[\phi(X) \mid z\right] = 0$. Players with signals lower than z, who place more weight on lower outcomes than players with siganals equal to z, so we know that for $s < z$, we have $\mathbb{E}\left[\phi(X) \mid s\right] < 0$, which means the value of waiting under Z must be lower than the value of waiting under Z'. This proves Proposition 9, part 1, and we must conclude that the welfare maximizing Z is never lower than x^*.

E.5.2 $\widetilde{Z}$ is the welfare maximizing disclosure threshold among disclosure policies which prevent no news drilling

In comparing full disclosure and complete secrecy, we were able to establish that full disclosure generates higher welfare than complete secrecy by way of a signal-by-signal analysis, under the special case that $t_2 > u_1$. When this is true, for every signal s, players are either indifferent to which game they played, or strictly prefer playing full disclosure. If two disclosure thresholds both prevent no news drilling from happening, it turns out that we can make a similar comparison.

Let $\widetilde{Z}$ be the lowest disclosure threshold that prevents no news drilling from happening, and if there is no no-news drilling at $Z = x^*$, let $\widetilde{Z} = x^*$. Consider two disclosure thresholds,

Z_1 and Z_2, such that $\widetilde{Z} \leq Z_1 < Z_2 \leq \overline{Z}$. We will compare player welfare at a given signal under each of these games.

Because of the monotonicity of $v_1(\cdot)$, we have $v_1(Z_1) > v_1(Z_2)$, so a player with a signal $s > v_1(Z_1)$ drills early in both games and experiences the same welfare. Players with signals between $v_1(Z_2)$ and $v_1(Z_1)$ drill early in the Z_2 game but wait in the Z_1 game, and are strictly better off in the Z_1 game by revealed preference.

Finally, players with signals $s < v_1(Z_2)$ wait in both games. Because we've assumed that both Z's prevent no news drilling, the waiting values are:

$$W^{Z_1}(s) = \mathbb{E}\left[\mathbb{I}(X \geq Z_1)\pi(X)(1 - F(v_1(Z_1) \mid X)) \mid s\right]$$
$$W^{Z_2}(s) = \mathbb{E}\left[\mathbb{I}(X \geq Z_2)\pi(X)(1 - F(v_1(Z_2) \mid X)) \mid s\right]$$

The difference in waiting values is thus:

$$\begin{aligned}\Delta(s) &= W^{Z_1}(s) - W^{Z_2}(s) \\ &= \mathbb{E}\left[\mathbb{I}(Z_1 \leq X \leq Z_2)\pi(X)(1 - F(v_1(Z_1) \mid X)) \mid s\right] \\ &\quad - \mathbb{E}\left[\mathbb{I}(X > Z_2)\pi(X)(F(v_1(Z_1 \mid X)) - F(v_1(Z_2) \mid X)) \mid s\right] \\ &= \mathbb{E}\left[\phi(X) \mid s\right]\end{aligned}$$

The first term in this expression is profits that a player who waits in Z_1 earns, but would not earn in Z_2, due to the higher disclosure threshold. Because $Z_1 > x^*$, this is positive. The second term represents profits that a player who waits in Z_2 earns, but would not earn in Z_1, due to its rival's lower cutoff strategy. This must be a loss, relative to Z_2. Like in our exploration of the full disclosure and complete secrecy waiting values under the $u_1 > t_2$ condition, this difference in waiting values can be expressed as the expected value of a function $\phi(X)$ which has a single crossing, and so we can (possibly) sign it with knowledge of its value at $s = v_1(Z_2)$.

To see this, first note that a player with signal $s = v_1(Z_2)$ is indifferent between drilling and waiting when the threshold is Z_2. However, the same player strictly prefers to wait when the threshold is Z_1. Thus, we know that $\Delta(v_1(Z_2)) > 0$. Next, note that for x in the range $[Z_1, Z_2)$, $\phi(x) > 0$, and $\phi(x) < 0$ when $x \geq Z_2$, so $\phi(\cdot)$ has a single crossing at $x = K_2$. Since we know $\Delta(v_1(Z_2)) > 0$, and since $\phi(x)$ is positive for lower values of x and negative for higher values of x, then the Hendricks & Kovenock MLRP/single crossing lemma again implies that $\Delta(s) > 0$ for $s < v_1(Z_2)$.

Thus, players with signals below $v_1(Z_2)$ strictly prefer playing a partial disclosure game with threshold Z_1 to playing a partial disclosure game with threshold Z_2, and we can conclude that for all signals, players do no worse playing the Z_1 game than the Z_2 game. Since our choice of $Z_1 < Z_2$ was generic aside from the requirement that $\widetilde{K} \leq Z_1 < Z_2 \leq \overline{K}$, we can conclude that $\widetilde{K}$ offers the highest player welfare among all partial disclosure thresholds in the range $\left[\widetilde{K}, \infty\right)$. This proves Proposition 9, part 2.

As a corollary, note that if there is no no-news drilling at $Z = x^*$, then there is no no-news drilling at any partial disclosure level, and we can conclude that the optimal partial disclosure level is $Z = x^*$.

Appendix F Additional Tables and Figures

Figure A.1: Distribution of predicted output

(a) All grids in DeWitt regions 12 and 13

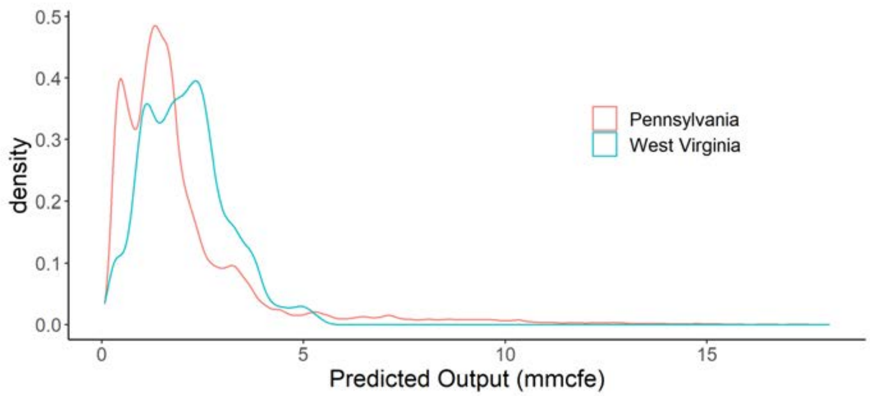

(b) Grids leased by 2005

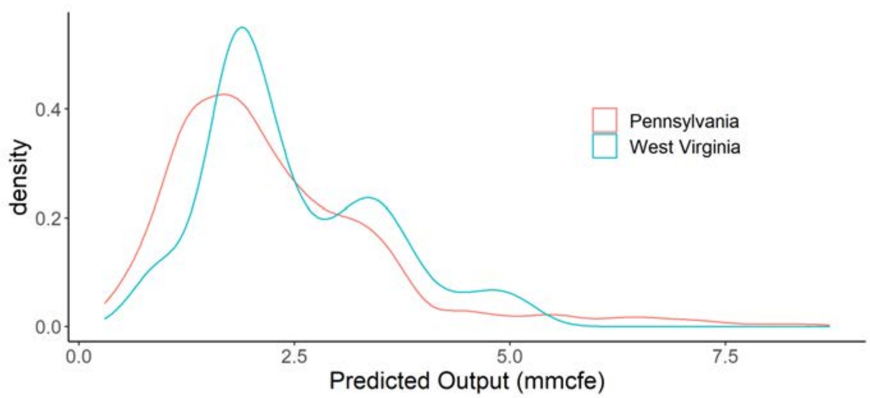

Output predictions generated using the krigging procedure outlined in section 5.

Table 1: 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

Appendix G Additional Tables and Figures

Figure 3: Border Grids, Pre-2010 Leasing and Drilling

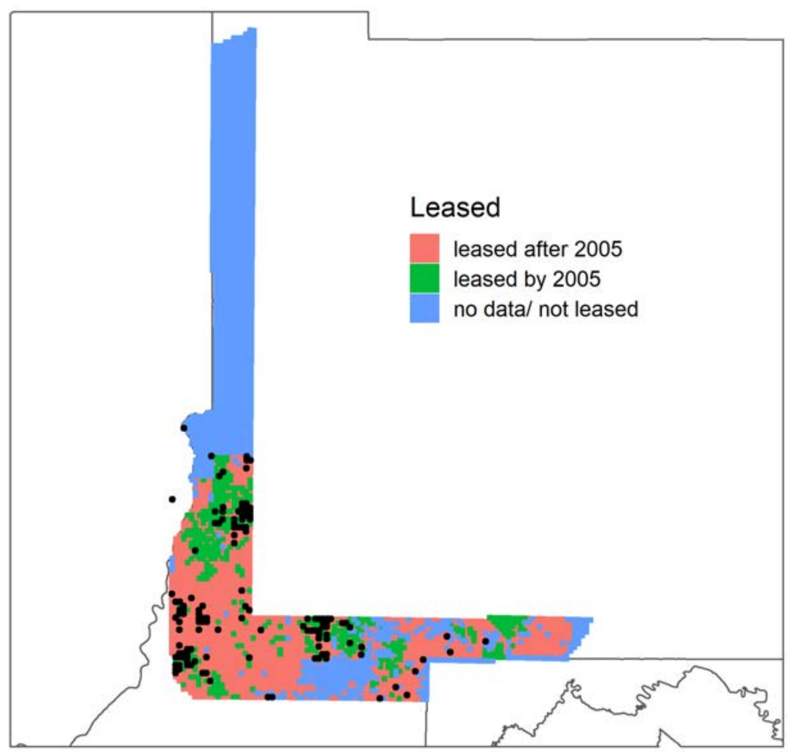

Black circles reflect square-mile grids which were drilled before the end of 2010.

Figure 2: Gas Region Map

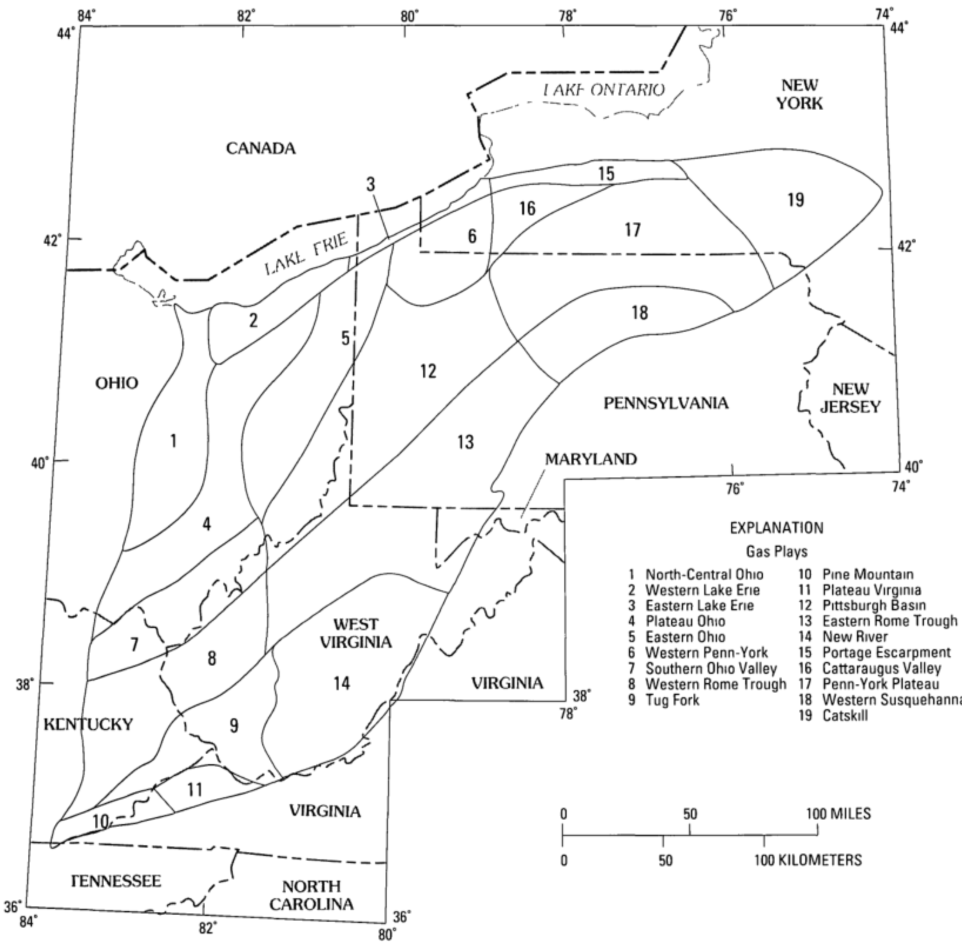

Figure 2. Devonian shale gas plays in the Appalachian basin

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 4: Predicted Output (mmcfe)

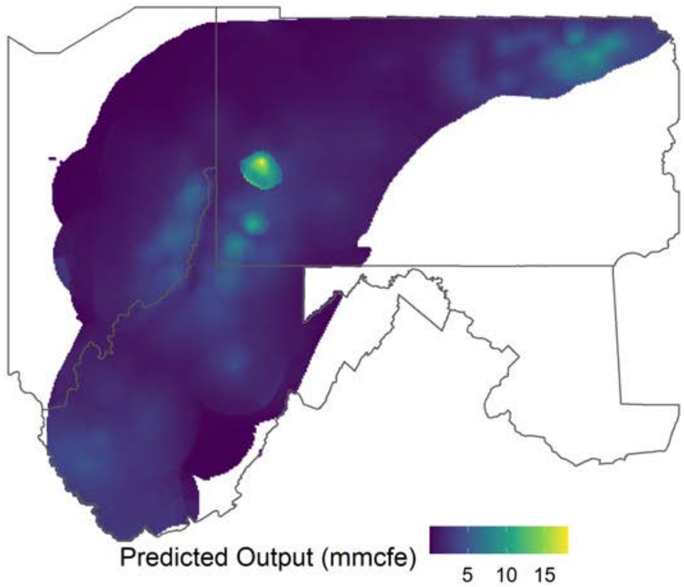

Output predictions generated using the kriging procedure outlined in section 5.

Appendix H Additional theoretical results and proofs

H.1 Optimal single agent payoffs are less valuable than CS equilibrium payoffs

Let s_0 be the cutoff signal at which a single agent would choose to drill, so that $0 = \mathbb{E}\left[\pi(X) \mid s_0\right]$. We know $u_2 < s_0$ because a firm with signal u_2 can only profitably drill under the news that its rival's signal was above u_1, and we know $s_0 < u_1$ because a firm with the cutoff signal must get strictly positive profits to drilling. Firms with signals at or above u_1 will drill early in the CS game and will drill in a single agent problem. Firms with signals between s_0 and u_1 will drill in a single agent problem and will choose to wait in a CS game, so those firms must be better off in the CS game by revealed preference. Firms with signals between u_2 and s_0 will drill after waiting with positive probability in the CS game, earning positive profits in expectation, while they won't drill at all in the single agent problem. Thus for all signals, the payoffs to playing a CS game are equal or higher than the payoffs in a single agent problem.

Appendix I Additional estimation details

I.1 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 $P = \$3.10/mcf$.

I.2 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 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 t_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. After substituting in our functional form assumptions for the conditional distribution of signals, and for the profit function, we obtain:

$$0 = \frac{1}{N}\sum_i ((P(1-\text{royalties}-\text{taxes})-\text{o\&m})X_i-K)\left(1-\delta\left(1-\Phi\left(\frac{u_1-\log X_i}{\nu}\right)\right)\right)\phi\left(\frac{u_1-\log X_i}{\nu}\right)$$

We can construct similar expressions for u_2, and for the cutoffs in the full disclosure game.