Colorado School of Mines
Earth • Energy • Environment

Division of Economics and Business
Working Paper Series

Bayesian Learning and Regulatory Deterrence: Evidence from Oil and Gas Production

Peter Maniloff

Working Paper 2016-04
http://econbus.mines.edu/working-papers/wp201604.pdf

Colorado School of Mines
Division of Economics and Business
1500 Illinois Street
Golden, CO 80401

April 2016

Colorado School of Mines
Division of Economics and Business
Working Paper No. **2016-04**
April 2016

Title:
Bayesian Learning and Regulatory Deterrence:
Evidence from Oil and Gas Production*

Author(s):
Peter Maniloff
Division of Economics and Business
Colorado School of Mines
Golden, CO 80401-1887
maniloff@mines.edu

ABSTRACT

This paper proposes a Bayesian learning model of regulatory enforcement. Firms exert compliance effort based on their belief about a regulator's effort level. Firms use regulatory actions to learn about the regulator and update their own compliance efforts accordingly. This theoretical model suggests that deterrence will be most effective when regulators have discretion or when firms are inexperienced. Econometric analysis of inspections of Pennsylvania oil and gas wells supports these hypothesis. This work provides a causal mechanism for the commonly observed phenomenon of general deterrence in which regulatory actions towards one firm lead other fims to increase their own compliance.

JEL **classifications: D22, K32, L51, L71, Q58**
Keywords: enforcement, deterrence, reputation oil and gas, hydraulic fracturing

*I thank Dietrich Earnhart, Jonathan Eyer, Harrison Fell, Ian Lange, as well as seminar participants at Colorado State University and the Colorado Environmental and Resource Economics Workshop for their comments. I also thank Francisco Sotomayor and Sean Ericson for research assistance, Environmental Defense Fund and the Colorado Oil and Gas Conservation Commission for sharing contextual expertise, and a private foundation for funding this work. All errors are of course my own.

1 Introduction

One of the core questions of applied microeconomics is how to design efficient regulations. However, regulations do not merely specify acceptable and unacceptable behavior. They also typically provide for enforcement when the rules are broken. Consider traffic safety. Absent any rules, individuals with a high taste for speed or individual value of time would drive quite rapidly, potentially imposing a safety externality on others. A state may pass speed limits in response. But a speeder who has never been ticketed, and never even heard of an acquaintance being ticketed, may well continue to speed. If the speeder is ticketed, she is likely to slow down in the future. She is even likely to slow down if she hears of a friend being ticketed for speeding. An efficient enforcement policy balances the safety benefits of increasing the number of policemen working on traffic enforcement with the cost of paying those policemen (or pulling them from other duties).

Alternatively, consider a firm which is subject to environmental regulations. The firm can choose whether or not to comply as an intentional decision weighing benefits of violation (such as reduced compliance costs) and costs of violation (such as fines). The firm may be deterred from violating if it is subject to regulatory action (specific deterrence) or if it observes others being subjected to regulatory actions (general deterrence) (Polinsky and Shavell 2000; Shimshack 2014). Understanding the magnitudes and determinants of deterrence is a key aspect of designing economically efficient policy. General deterence has also been described as representing a regulator's reputation (Shimshack and Ward (2005)). This language implies a causal mechanism that depends on firms learning about the regulator. However, learning about the regulator has not been substantially observed in the literature.

The key difficulty is observing variation in firm knowledge about the regulator which can be disentangled from potentially unobserved characteristics of both the firm and the regulator. The key difficulty is identification - the econometrician would need to observe the same firms, interacting with the same regulator, with variation in the firms' level of information about the regulator. The econometrician would also need to observe the same firms repeatedly to identify an unobserved propensity to violate (to control for regulatory targeting of known-to-regulators bad actors). Ideally the data set would also observe firms across jurisdictions and time, and have variation in firm information across all of these axes.

The U.S. onshore oil and gas drilling industry provides a novel opportunity to address these identification concerns. Firms operate a number of wells across different enforcement decisions, across time, and enter at different times. Thus in any given time and regulatory jurisdiction, firms have potentially distinct information sets about the regulator. This provides variation in compliance outcomes across information sets, firms, time, and regulators.

Prior literature has found evidence of deterrence in a variety of settings, including violent crime (Levitt 1997), parking violations (Fisman and Miguel 2007), as well as industrial settings including paper mills (Magat and Viscusi 1990; Gray and Shadbegian 2005), procurement (DiTella and Schargrodsky 2003), manufacturing (Hanna and Oliva 2010), and electricity generation (Keohane, Mansur, and Voynov 2009). Studies typically find that firms facing emissions regulations respond to actual or potential inspections and penalties by reducing emissions or violation rates (Earnhart 2004). There is evidence that firms respond to general deterrence (Shimshack and Ward 2005) and that this response is limited to enforcement actions conducted within the same jurisdiction (Gray and Shadbegian 2007).

The extent to which individuals follow the law can be quite heterogenous, even under a shared regulator. This can be due to variation in cultural norms or motivation to comply, cost of compliance, financial ability to comply, or beliefs about the regulator. Fisman and Miguel (2007) find that violations of parking laws are highly correlated with national corruption scores when enforcement is lax, but that compliance is broad when enforcement is stringent. DiTella and Schargrodsky (2003) find support for the Becker-Stiglitz model of corruption in which agents who earn efficiency wages are less likely to violate out of fear of losing their jobs. In addition to heterogeneity in the cost of penalty, heterogeneous costs of compliance or ability to finance compliance activities can induce variation in compliance (Earnhart 2004; Earnhart and Segerson 2012). Perhaps the most closely related work examines variation in regulated entities' knowledge about the regulator (Earnhart and Friesen 2013). This work uses a hybrid stated-preference approach and find that firms do increase compliance after penalties, despite the absence of any change by the regulator. They term this *experiential learning*.

This paper extends the prior literature by considering heterogeneous compliance due to imperfect knowledge about the regulator. I first discuss a theoretical model of firm behavior in which firms have imperfect knowledge about the regulator and use observable regulatory actions to update their Bayesian beliefs about the regulator. Firms choose a level of compliance based on their belief, or alternatively change their compliance level based on the change in their belief. Heterogeneous firm decisions arise from variation in firm beliefs about the regulator. I then test the model on a novel data set and support for the Bayesian updating hypothesis of belief formation.

This paper's primary analysis explores how firms respond to enforcement actions *directed towards other firms*, often termed "general deterrence" (Polinsky and Shavell 2000;

Shimshack 2014). The econometric analysis considers a panel setting in which firms operate many different facilities in a capital-intensive industry: oil & gas production in Pennsylvania. The main results provide evidence that firms form their beliefs about regulator behavior in a Bayesian updating process in which inspections act as an informative signal about regulator behavior. Inexperienced firms respond to a one standard deviation increase in nearby inspections by reducing their propensity to violate by up to forty percent. For experienced firms, this general deterrence effect is near zero.

A core finding of the deterrence literature is the importance of firm's imperfect information about the regulator (Polinsky and Shavell 2000). If the regulator behaves strategically or changes behavior over time, that could lead to even less accurate firm information. These findings jointly motivate a need to understand how firms learn about regulator behavior.

This paper makes three distinct contributions to the literature. First, I develop a simple analytic model to explain regulatory deterrence as a Bayesian updating process and use the model to generate testable predictions. Second, I test these predictions on a novel data set of oil and gas well sites in Pennsylvania. I find that inspections are effective for general deterrence, but fines are not, which is consistent with one of my testable predictions. I additionally find that general deterrence decreases with firm experience, which is also consistent with a testable prediction from my model of learning. Finally, I show that firms primarily respond via capital decisions at the time of well construction and not via changes in operational practices. This finding is grounded in the particular institutional details of oil and natural gas production.

The paper continues as follows. Section 2 develops a theoretical model of Bayesian learning and general deterrence and derive testable predictions. Section 3 discusses the identification strategy and data. Section 4 discusses an econometric model to test the theoretical model and presents estimation results. Section 5 concludes.

2 Theoretical Model

Consider an oil and gas production firm which can engage in costly effort to reduce its probability of being fined or suffering other regulatory enforcement actions. In a rational expectations framework, it will do so as long as the marginal cost is lower than the expected marginal cost of an enforcement action. In this framework, the firm can exert effort to reduce its violation probability. The regulator exerts enforcement effort, which is unknown to the firm.[1] The firm can learn about regulator effort by observing regulatory action, and chooses its own effort level accordingly. The model implies that firm will change their effort based on observations of the regulator's behavior. If the firm increases its believed level of regulator effort, it will increase its safety effort and the actual probability of violation will decrease.

This process can be broken down into a series of discrete steps as below

1. First, the regulator takes an action. This might be inspecting a facility, inspecting many facilities, fining violators - or doing nothing.

2. Second, the firm observes the regulator's action.

3. Third, the firm updates its belief about the regulator's effort level.

[1] The literature on regulatory enforcement typically distinguishes between the probability of being penalized and the severity of a penalty (Polinsky and Shavell 2000). I abstract from this to a single scalar measure of the regulator's effort or stringency. Extending this model to consider multiple dimensions of regulatory behavior (such as probability of catching violators and severity of penalty) is straightforward but adds little intuition to the model. Moreover, in this setting sanctions for substantial violations are generally fixed by statute.

4. Fourth, the firm updates its own effort level towards safety in light of its new belief.

5. Fifth, the change in the level of safety effort changes the probability of a violation.

The remainder of the section models this sequence more formally. Section 2.1 models the firm choice in step four. Section 2.2 models steps two and three. Section 2.3 combines steps two through four to show step five. The regulatory decision process is taken as exogenous to the model.[2]

2.1 The firm's choice of safety effort

Consider a firm which can choose an effort level $\rho \in \left[\underline{\rho}, \bar{\rho}\right]$ of promoting safety. A higher level of safety reduces the probability of a violation occurring, but safety effort is costly, so in the absence of regulatory enforcement the firm would choose $\rho = \underline{\rho}$. The firm has beliefs $\hat{\theta}_t$ about the regulator's degree of enforcement effort θ. The probability of a violation is a function of $\frac{\theta}{\rho}$ and a random error ε such that the probability of a penalty is

$$p(penalty) = \left\{ \begin{matrix} 1 & & \frac{\theta}{\rho} + \varepsilon \geq 0 \\ & if & \\ 0 & & \frac{\theta}{\rho} + \varepsilon < 0 \end{matrix} \right\} \tag{1}$$

[2] A related literature explores how regulators make decisions - how they decide what sites to inspect, what penalties to levy, and the stringency of regulations (Gray and Shadbegian 2004; Barrage 2015). These questions are of substantial interest, but beyond the scope of this study. I do plan to explore regulator behavior in future work. In some states, decisions on what oil & gas sites to inspect are largely based on priority ranking formulae which are a matter of public record. These rankings typically specify that wells which are the subject of public complaints are inspected immediately, suggesting an instrument which is plausibly exogenous to regulator knowledge. This literature finds that regulatory effort (typically measured by inspection probability, inspection count, or penalty size) increases with the potential harms from a violation, with the income of the nearby residents, and with the sector's profitability (Gray and Shadbegian 2004; Barrage 2015; Konisky 2009).

If the regulator is exerting much more (less) effort than the firm, a penalty is (un)likely.This implies that

$$p(penalty) = F\left(\frac{\theta}{\rho}\right) \tag{2}$$

where $F()$ is the cumulative distribution function of ε. Following the standard crime literature and normalizing such that a unit of firm effort has a cost of one, a risk-neutral firm chooses an optimal effort level ρ^* to minimize the sum of safety effort expenses and expected fines. For the risk-neutral which seeks tp minimize the sum of compliance costs and expected penalties (i.e., firm problem $\min_{\rho}(\rho + E[fine])$), the first order condition is

$$\frac{\partial F(\hat{\theta}_t/\rho^*)}{\partial \rho^*} = -\frac{\partial E[penalty]}{\partial \rho^*} \tag{3}$$

describes the first order condition. Equation 3 implies that $\partial\rho^*/\partial\hat{\theta}_t > 0$. Intuitively, if a firm learns that a regulator is more aggressive than previously believed, the firm will exert more effort to avoid penalty. Because the probability of penalty depends on the true value θ and not the firm's belief $\hat{\theta}_t$, if the firm updates its belief in $\hat{\theta}_t$ positively, the actual probability of a violation will decrease. That is $\frac{\partial p(penalty)}{\partial \hat{\theta}_t} < 0$.

2.2 The firm's learning problem

The firm has an initial noisy belief about θ, and the firm updates this belief as it observes regulatory actions such as inspections or fines.[3] The firm's initial belief $\hat{\theta}_t$ about θ is a random variable distributed normally with mean $\hat{\theta}_0$ and variance σ_0^2. The firm interprets

[3]The Bayesian learning model is directly based on the Pastor-Veroni Bayesian updating model of learning in financial markets. I substantially follow the notation and exposition of Pastor and Veronesi (2009).

regulator actions such as inspections or fines as noisy informative signals about θ. Each signal s_t is distributed normally with mean θ and variance $\sigma^2 > 0$. If a firm has observed T signals, then the Bayesian updated prior is

$$\hat{\theta}_T = \hat{\theta}_0 \frac{\frac{1}{\sigma_0^2}}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}} + \bar{s} \frac{\frac{T}{\sigma^2}}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}} \tag{4}$$

where $\bar{s}$ is the average value of the signals. The firm's belief θ_T is the average of the initial belief and average signal, weighted in proportion to the number of signals and (un)certainty in both initial belief and signals. We see immediately that as T becomes large, θ_T approaches $\bar{s}$. The variance in the updated prior is

$$\hat{\sigma}_T^2 = \frac{1}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}} \tag{5}$$

We can also describe the change in a firm's belief as it observes as single signal. If the firm has already observed T signals, then the change from one more $\Delta\hat{\theta}_{T+1} \equiv \theta_{T+1} - \theta_T$ is

$$\Delta\hat{\theta}_{T+1} = \frac{s_{T+1} - \theta_T}{1 + \sigma^2/\hat{\sigma}_T^2} \tag{6}$$

The numerator provides us the intuitive result that the firm will adjust its prior upwards if $s_{T+1} > \theta_t$ and downwards if $s_{T+1} < \theta_t$. More interestingly, the denominator tells us that the magnitude of the adjustment depends on the ratio of the noise (variance) in the signal to the uncertainty (variance) in the firm's prior. While the adjustment will always be partial (because the denominator is greater than one), the adjustment will be smaller the noisier the signal.

2.3 Theoretical Implications

Equation 6 yields three conclusions:

Lemma 1. *The marginal effect of a signal on a firm's change in probability of violation will be smaller the more signals a firm has observed.*

Lemma 1 arises because that the magnitude of observed deterrence depends on a firm's experience level. An experienced firm will revise its prior by less than an inexperienced firm which observes the same signal. That is, $\Delta\hat{\theta}_{T+1}$ is smaller in absolute value for the experienced firm, if both firms have the same $\hat{\theta}_T$. This smaller adjustment in belief leads to a smaller adjustment in safety effort.

Lemma 2. *If a firm has exact knowledge of regulatory effort ($\hat{\sigma}_T^2 = 0$) and exerts more than the minimal possible effort, then regulatory effort may be an effective deterrent even if the firm is never observed to change its behavior when it observes regulatory action.*

Imagine a firm with exact knowledge of regulatory effort ($\hat{\sigma}_T^2 = 0$) which exerts safety effort $\rho^* > \underline{\rho}$. If the regulator conducts enforcement actions, the firm will not update its belief $\hat{\theta}$ because it already knows θ. This implies that the firm will also not change ρ and there will be no observable change in the probability of violation. Nonetheless, the firm is exerting more safety effort than it would absent regulatory enforcement effort. In essence, this is a distinction between the *marginal* deterrence effect of enforcement actions and the *total* deterrence effect of enforcement.

Corollary 1. *The error in a firm's belief about the regulator ($\hat{\theta}_T - \theta$) will be smaller the less uncertainty about regulator effort.*

Corollary 1 follows a similar logic to Lemma 2. If there is less initial uncertainty about regulatory effort, then more weight is assigned to the initial belief and subsequent adjustments are smaller.

Sections 3 and 4 test the model econometrically. For derivations of each Lemma and Corollary, see Appendix A.

3 Identification Strategy and Data

3.1 Identification Strategy

This paper's identification strategy is based on using intra-region, intra-firm, time-varying variation in intensity of regulatory effort θ and in firm experience level to identify changes in the probability of violations. As violation probability reflect changing firm effort and fixed regulatory effort, changes in violation probability reflect changes in firm effort level.

I construct measures of θ for every individual well and time period in Pennsylvania. I also observe all well site inspections by the Pennsylvania Department of Environmental Protection, which serve as my unit of observation. For each inspection, I regress whether or not a violation was observed is regressed on a measure for θ and the interaction of θ and a measure of the number of signals the firm has observed T. Based on my theoretical model, I expect θ to have a negative effect on likelihood of violation, but that this effect will be smaller in magnitude for more experienced firms.

For every inspection, I construct a measure of previous inspections of other firms' wells in the same jurisdiction. This measure acts an informative signal to the firm about regulatory effort θ in that region . The key identification assumption for this measure is that regulatory propensity to inspect other firms wells is not correlated with a firm's propensity to violate, after inclusion of control variables. This seems likely because while regulators seem to target both high violation probability firms and areas with high damages from violations, regulators have detailed information about who controls a given well and thus informational spillovers seem unlikely.

I construct several measures of a firm's experience or number of signals observed, including the time that the firm has been operating in Pennsylvania, the number of wells the firm has, and the number of times the firm has been inspected. The length of time that a firm has been operating in a given jurisdiction has the closest conceptual linkage to the signal count T of the theoretical model. Conceptually, each day (or time period) a firm can observe how intensively the regulator is monitoring in the area around each of its' wells. It can then update its belief $\hat{\theta}$ accordingly. However, it may be that firms learn more about the regulator during direct inspections, perhaps because they then have the opportunity to speak directly with inspections staff or because direct inspections are more salient (Earnhart and Friesen 2013). Therefore I also construct an alternative measure of experience based on the number of inspections a firm has experienced. I will construct each of these measures at a variety of jurisdictional and temporal scales.

By focusing on general deterrence monitoring, or the the inspections of other firms, I use a regulatory enforcement action which is observable, variable, subject to considerable discretion by the regulator, and plausibly exogenous. The regulator has considerable latitude over which wells to inspect, which implies that there is uncertainty about the regulator's true priority (the converse would be the case in which regulatory decisions were prescribed by law, in which case the firm would have correct knowledge of θ and thus would not update its beliefs).

Firms operate across jurisdictions and time, with multiple firms per jurisdiction. This allows me to flexibly control for unobserved firm-level propensity to violate by including firm-level fixed effects. There is also evidence that regulators target enforcement based on spatial factors such as presence of cities. I can control for these factors with jurisdictional or sub-jurisdictional fixed effects (For example, including a county-level fixed effect would effectively also control for regional enforcement office, which cover multiple counties and follow county lines.).

3.2 Data

My primary data set consists of inspection records for all inspections of natural gas well sites conducted by the Pennsylvania Department of Environmental Protection (DEP) from 2000-2014. The DEP enforces environmental rules for gas production. They employ approximately 76 who inspectors[4] who conduct field inspections of wells. They inspect wells for a variety of reasons - fifty-five percent are random routine inspections, while most of the rest are conducted at the time of drilling, due to complaints, or at the time of well closure and site restoration. If violations are found, the agency can work with the firm to develop a plan for the firm to return to compliance, levy fines, or both.

[4]http://projects.propublica.org/gas-drilling-regulatory-staffing/states/PA.html, retrieved 02/19/2016

Figure 1: Inspections and Violations By Year

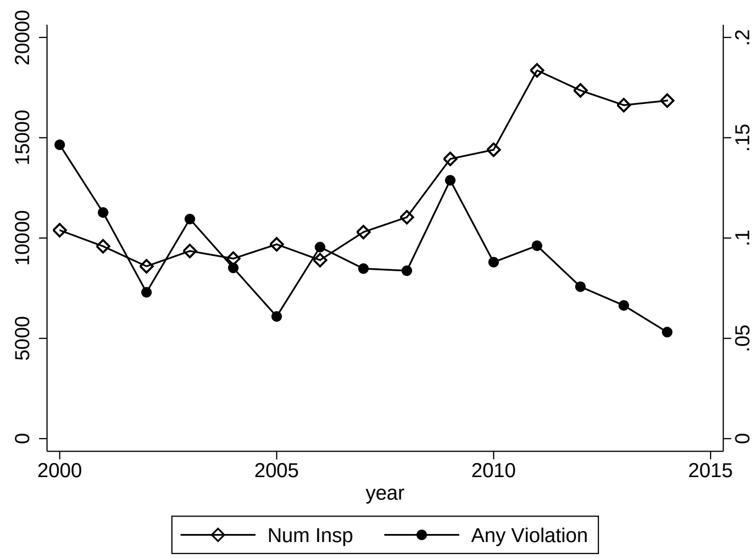

Inspectors conducted 184,367 inspections from 2000-2014, of which 8.9% found violations. Figure 1 shows the number of inspections per year as well as the fraction which find a violation or a major violation. The number of inspections has increased over time (as both the number of wells and inspectors have increased), while the violation rate has decreased over time.

An inspector might be on site during the drilling or fracing processes, in which case she would be able to observe downhole construction techniques such as what cement is used to seal the well. Otherwise, inspections are largely confined to above-ground observation. Inspections include above-ground well equipment, spill control facilities such as earthen berms, as well as more minor attributes such as whether grass in the area around the well was mowed and whether signs to the well were clear.

The unit of observation is an individual inspection of a single oil & gas location. Each observation includes an identifier for the inspection, the well inspected, the well's operator or controlling firm, the location of the well, inspection date, the type of inspection (routine, complaint, etc), and information about the outcome of the inspection. The data set additionally includes fines levied in response to violations, if any.[5] Approximately 8.9 percent of inspections find violations. Of violations, 8.0 percent result in fines. Conditional on being nonzero, the median fine is \$6000 while the mean is \$42676. Fine are right skewed, dominated by a small number of very large penalties.

For each inspection, I construct measures of specific and general deterrence for that well at that time, due to both prior inspections and fines. Thus each inspection has four treatments. For specific inspection deterrence, I divide the number of inspections of the firm's wells over the firm's total number of wells. This yields the rate of inspections during previous months. For specific fine deterrence, I divide the total fines to a given firm to date by the firm's total number of fines for an average fine conditional on a fine occurring. By omitting zero-fine inspections from the fine average, I can distinguish between the likelihood of enforcement action and its severity. General deterrence measures are analogous, except that they are the rate of inspection and mean fine for other firms' wells in the same county as the inspected well. I convert these measures to a z-score with mean zero and standard deviation one by subtracting the mean and dividing by the standard deviation. This allows some degree of comparability across treatments - regression coefficients will describe the impact of a one standard deviation change in treatment.

Table 1 provides summary statistics. Approximately 8.89 percent of observations find

[5]I further observe fines paid, which is typically a substantially smaller number. I leave exploring the implications of this finding for future research.

violations, and approximately 0.7 percent incur fines (approximately 8 percent of violations). Conditional on a fine occuring, the average is $42,676. Table 1 also describes the average number of inspections, wells, and years of experience that a firm has at varying jurisdictional scales at the time of each inspection. The finest jurisdiction scope I use is the county. Pennsylvania has 67 counties, which have an average area of 687 square miles. If enforcement is based on local socioeconomic characteristics, regulators may target within counties (or firms may believe that they will). Treatments are also aggregated to the jurisdictional level as deterrence spillovers may only apply within the boundaries of an enforcement jurisdiction (Gray and Shadbegian 2007). The Pennsylvania Department of Environmental Protection has three oil and gas enforcement districts, which strictly follow county lines (Pennsylvania Department of Environmental Protection 2013). Finally, treatments are aggregated to the state level.

For some observations, no fines have been issued in the region by the time of the observation. This means that the average fine is undefined. For my core analyses, this is a

Table 1: Summary Statistics

VARIABLES	(1) mean	(2) sd
Found Violation	0.0889	0.285
Was Fined	0.00753	0.0865
Fine	42,676	97,898
Firm inspections in county (000's)	0.318	0.471
Firm inspections in district (000's)	0.683	0.917
Firm inspections (000's)	1.263	1.831
Firm experience in county (years)	5.151	4.040
Firm experience in district (years)	6.125	4.177
Firm experience (years)	6.517	4.203
Firm well count in county (000's)	0.106	0.139
Firm well count in district (000's)	0.220	0.255
Firm well count (000's)	0.405	0.498
Inspections of others in county	4,951	4,617
Number of observations	184367	
Number of firms	727	

relatively small number of observations, so I drop these observations. In subsequent analyses I look at smaller time periods, and the number of observations without defined fine averages is larger. I therefore drop fine averages as a regressor in those regressions. In each case, results are robust to including fine average z-scores (dropping observations without defined averages) or omitting fine measures (and including all observations).

4 Econometric Analysis

I begin by confirming that firms do seem to be deterred by regulatory action. The first analysis estimates whether or not an inspection noted a violation on measures of regulatory effort, stringency of penalty, and a rich set of control variables and fixed effects. This analysis finds a negative, statistically significant, and economically substantial impact of monitoring effort on violation probability.

This analysis estimates equation 7. The dependent variable y_i is a boolean indicator for whether or not a violation was noted at inspection i of a well site operated by firm o conducted in year t. Deterrence is measured by an array of treatments Z_{ijk}, discussed in detail in section 3.2. Indicator variables for type of inspection (routine, complaint, etc), whether or not the well is a tracking well, and whether or not it is the first inspection for a site comprise X_i. This specification allows for unobservables at the firm, year, and inspection level (ν_o, μ_t, and ε_i respectively).

This analysis estimates equation 7 as a linear probability model because many firms have no violations.[6] In a logistic regression model with fixed effects, these firm are dropped and thus cannot be used to identify county and time effects. A panel probit model assumes that the firm-level effects ν_o are uncorrelated with other independent variables, an assumption which is not satisfied in linear probability estimates.

$$y_i = \sum_{j \in [spec,gen]} \sum_{k \in [insp,fine]} \alpha_{jk} Z_{ijk} + X_i \beta + \nu_o + \mu_t + \varepsilon_i \qquad (7)$$

Results for estimating equation 7 are presented in Table 2.[7] We see across specifications that a one standard deviation in inspection treatment reduces the likelihood of a violation approximately 3.3 percentage points. Recalling that the baseline probability of finding a violation is 8.9 percent percent, a one standard deviation in inspection decreases the probability of violations by approximately 37 percent.[8]

Column 3 of Table 2 reports results for specific deterrence ($j = spec$). The point estimate is small in magnitude and statistically insignificant. This coefficient may be biased upwards if regulators preferentially target sites where they believe violations are likely.[9] This is a well known problem in the literature. However, this endogeneity does not seem likely to materially bias the parameters of interest for two reasons. First, the correlations between

[6]I omit consideration of imperfect detection. If regulators do not observe observe all violations which are committed (due to incomplete monitoring, for example), then estimates of regulatory deterrence can be biased (Feinstein 1990; Helland 1998). By assuming that the probability of detection of a violation is unity, the actual probability of violation can be underestimated. In this case, estimates of $\alpha_j k$ can be biased. However, the key parameter of interest in this analysis is not $\alpha_j k$ itself but how $\alpha_j k$ varies with firm experience. If the probability of detecting a violation conditional on one occuring is correlated with experience, after controlling for firm, time, and jurisdictional characteristics, then estimates of the relationship between $\alpha_j k$ and experience may be biased. This point follows from Theorem 2.2 of Feinstein (1990).

[7]Estimates for control variables are included in Appendix B.

[8]3.3/8.9

[9]Conversations with state regulators confirm that they consider some firms to be less careful or more likely to violate environmental regulations than others, and use this information in deciding what facilities to inspect.

Table 2: Average General Deterrence Effects

VARIABLES	(1) Any	(2) Any	(3) Any
General Inspection	-3.612***	-3.305***	-3.243***
	(0.202)	(1.002)	(0.996)
Own Inspection			0.104
			(0.069)
Observations	164,544	164,544	164,544
R-squared	0.111		
Operator FE		Y	Y
R^2 of FEs		0.472	0.472
R^2 of X's		0.0762	0.0792

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include yearly fixed effects as well as additional control variables. "Operator FE" describes firm-level fixed effects. All specifications use heteroskedasticity-consistent standard errors.*

specific and general deterrence measures are low (as reported in Table 3). Second, results are robust to including or excluding specific enforcement measures (as reported in Columns 2 and 3 of Table 2).

Table 3: Correlation between General and Specific Deterrence Measures

Variables	Own Inspection	Own Fine	General Inspection	General Fine
Own Inspection	1.000			
Own Fine	-0.005	1.000		
General Inspection	0.002	0.033	1.000	
General Fine	0.012	0.017	0.148	1.000

4.1 Deterrence decreases with experience

This section tests Lemma 1 by exploring the hypothesis that firm experience with the regulator reduces the magnitude of the deterrence effect. Here the magnitude of deterrence to varies with a proxy variable N_{ot} for the number of signals that firm o has undergone at time t. I estimate the linear probability model of Equation 8. All other variables are the same as in Equation 7. Some specifications allow this interaction to include polynomial effects to allow for nonlinear effects as implied by the theoretical model.

$$y_i = \sum_{j \in [spec, gen]} \sum_{k \in [insp, fine]} \sum_{m=0}^{p} \alpha_{jkp} Z_{ijk} N_{ot}^{m} + X_i \beta + \nu_o + \mu_t + \varepsilon_i \tag{8}$$

Results of a linear interaction ($p = 1$) are shown in Table 4 for three different measures of firm experience. Columns 1-3 measure experience as the number of inspections that a firm has received in a county, PA DEP district, or statewide. Columns 4-6 measure experience as the time that a firm has been operating in a county, PA DEP district, or statewide. In all cases, we see that a completely inexperienced firm has a deterrence effect of 4-6 percentage points. In all specifications, a one standard deviation increase in experience is associated with a 0.7-1 percentage point decrease in the magnitude of the deterrence effect.

My preferred specification in Column 2 measures N_{ot} as the number of times the regulator has previously inspected a firm's facilities (divided by 1000), by each district office of the PA DEP. This is preferred because there is evidence that firms react more strongly to firms at their own facilities than inspections of others (Earnhart and Friesen 2013). We see that a completely inexperienced firm has a deterrence effect of -4.427 percent - i.e., a one standard deviation increase in nearby monitoring effort is associated with a 4.427 percentage point reduction in the probability of a violation. A firm which has experienced the mean value of 683 inspections would have a deterrence of $-4.427 + 1.098 * 0.683$ or -3.7 percentage points, which is comparable to the average estimate of -3.3 percentage points from Columns

2 and 3 of Table 2. Put somewhat different, we see that the deterrence effect decreases by approximately 25% per thousand inspections (1.098/4.427).

An alternative measure of experience is presented in Columns 4-6. Here experience N_{ot} is measured by the number of years that a firm has been operating in a county, district, or statewide. This measurement of experience may be more comparable to the measure T from the theoretical model in which a firm observes a signal about regulatory effort in each time step than the number of inspections. Here we see a qualitatively similar response. In Column 5, a firm with the average value of 6.1 years of experience in a district has a deterrence effect of approximately -4.2 percentage points. This is equivalent to the deterrence effect decreasing by approximately 3% per year.

Table 4 also includes the effect of own enforcement actions (starting with the first row, "Own Inspection"). We see that these are generally small in magnitude, statistically insignificant, and have an unexpected sign suggesting that enforcement action increases violation probability. This may reflect targeting by the regulator - if the regulator targets likely violators for inspection, then the coefficient on own inspections will be biased upwards. Thus we focus on the coefficents on inspections of others and note again that there is little correlation between the two as shown in Table 3.

The next specifications allow for nonlinear interactions. For ease of interpretation, instead of presenting a full table of polynomial interaction terms, Figure 2 shows the marginal effect of deterrence $\partial y_i / \partial Z_{ijk}$ with respect to a firm's experience (or number of signals T). For each graph, the outcome variable y_i is the change in probability of violation after observing a one standard deviation increase in nearby monitoring. We see that for each measure of firm learning, deterrence is more effective for inexperienced firms than for more experienced

Table 4: The Effect of Firm Experience on General Deterrence

	Inspections			Time		
	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	County	District	All	County	District	All
Own Inspection	0.052	0.074	0.039	0.079	0.077	0.076
	(0.071)	(0.068)	(0.075)	(0.073)	(0.073)	(0.073)
General Inspection	-4.606***	-4.427***	-4.046***	-5.065***	-5.228***	-5.707**
	(0.955)	(0.953)	(0.975)	(1.443)	(1.680)	(2.215)
Own Insp * Experience	3.235*					
	(1.784)					
Gen Insp * Experience	1.993***					
	(0.451)					
Own Insp * Experience		1.391				
		(1.174)				
Gen Insp * Experience		1.098***				
		(0.222)				
Own Insp * Experience			2.547			
			(1.922)			
Gen Insp * Experience			0.365***			
			(0.117)			
Own Insp * Experience				0.038		
				(0.026)		
Gen Insp * Experience				0.173**		
				(0.080)		
Own Insp * Experience					0.040	
					(0.026)	
Gen Insp * Experience					0.165*	
					(0.098)	
Own Insp * Experience						0.038
						(0.026)
Gen Insp * Experience						0.193
						(0.135)
Observations	164,544	164,544	164,544	164,544	164,544	164,544
Own Enforcement	Y	Y			Y	Y
R-2 of FEs	0.474	0.475	0.475	0.477	0.476	0.489
R-2 of X's	0.0820	0.0822	0.0823	0.0816	0.0816	0.0815

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include operator fixed effects, yearly fixed effects, and additional control variables. All specifications use heteroskedasticity-consistent standard errors.*

firms. In all specifications, a completely inexperienced firm reduces its violation probability by 4-5 percentage points after a one unit deterrence. In the top row, experience is measured as the number of times the firm has been inspected (at the county, district, or state level in columns 1,2, and 3). By this measure, the most experienced firms cannot be said to act at all - their deterrence magnitude is statistically indistinguishable from zero. In the bottom row, experience is measured as the number of years a firm has been operating (again at a county, district, or state level). Again we see that more experienced firms are less deterred by monitoring signals, but this learning seems to be weaker in magnitude than learning which has occurred via direct inspection.

Figure 2: Deterrence Effect versus Number of Signals Received

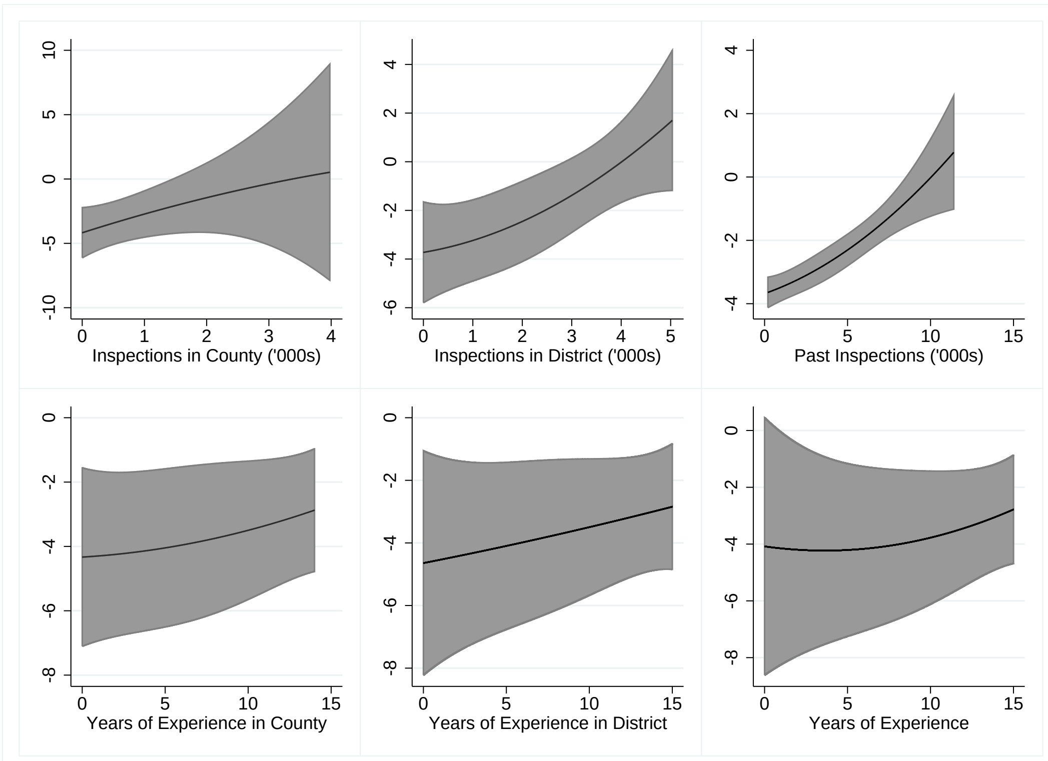

Each graph measures the effect of a one standard deviation increase in deterrence on probability of violation versus the number of signals a firm has recieved. In the top row, the firm's signal stock is measured by the number of times it has been inspected. In the bottom row, the signal stock is measured as the length of time the firm has been operating. Both measures are calculated at the level of county (first column), regulatory district (second column), and state (third column).

4.2 Robustness Checks

This section describes a series of robustness checks, which all support the primary analysis. The first set of robustness checks varies the time horizon of deterrence signals. It is not clear *a priori* whether firms discard old signals, and if so at what rate. In these analyses, the measures of deterrence treatment are calculated over one month, one year, and two years. Results are consistent with previous results, although somewhat less precise. The next check considers the possibility that firm size, instead of firm experience, decreases deterrence. This hypothesis is unsupported. The final check considers a 2012 change in some oil and gas regulations by estimating the core models separately before and after the rule change.

4.2.1 Core Results are Robust to The Time Horizon of Deterrence

The core results that monitoring decreases violations and that this effect decreases with learning are robust to different time horizons. This section repeats the above analyses, but instead measures of deterrence treatment as occurring over a previous limited time period. That is, the treatment Z_{ijk} describes the number of inspections over the previous 1, 12, or 24 months, again scaled such that the coefficient describes the impact of a one standard deviation change.[10]

The results are presented in Tables 5-7. Column 1 of each Table repeats the first analysis confriming that monitoring does deter violations (as in Column 3 of Table 2). We see that point estimates are substantially smaller, approximately 1.2% instead of 3.3% found before. This is consistent with the operational structure of the industry. In oil and gas production, many operational decisions are made at the time the well is drilled and are prohibitively expensive to alter later (because altering them entails underground equipment). A firm's ability to respond to new beliefs about the regulator is limited in the short-run because well

[10] Previous analyses also included the average fine (conditional on fines occuring) as a measure of the cost of sanction. They are omitted from this analysis because there are larger a number of county-months (or years) in which no fines occurred. In these observations, the average fine is undefined. This analysis omit fine average instead of dropping observations from the regression. This is why the number of observations is larger than in Table 4. Estimates which include the variable and drop observations with no fines produce analogous results and are available upon request.

decisions were already made at the time of drilling. We should therefore expect short-run responses to be smaller than long-run responses.

While firm's ability to respond in the short run is limited, we still see evidence of learning. Columns 2 and 3 reestimate the models in which firms's knowledge is captured by their number of inspections (column 2) or time operating in a county (column 3). In each case, we see that inexperienced firms are somewhat more deterred than average, with the magnitude of deterrence decreasing with experience. Again the results for learning via direct experience with the regulator are strongly statistically significant. Point estimates for learning over time have the expected sign and again suggest that the magnitude of deterrence decreases by several percent per year, but are not statistically significant. In each case, the relative magnitude of learning is comparable to the primary estimates of section 4.1, albeit larger in most specifications. The magnitude of deterrence decreases by 3-5% per year, or by approximately 15% per thousand inspections.

Table 5: Deterrence from Monitoring Within One Month

VARIABLES	(1) No Learning	(2) Experience	(3) Time
Gen Insp	-1.144***	-1.294***	-1.384***
	(0.299)	(0.343)	(0.512)
Gen Insp * Experience		0.195***	0.038
		(0.068)	(0.043)
Observations	184,367	184,367	184,367
R^2 of FEs	0.468	0.469	0.469
R^2 of X's	0.0115	0.0129	0.0124

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include operator fixed effects, yearly fixed effects, and additional control variables. All specifications use heteroskedasticity-consistent standard errors.*

Table 6: Deterrence from Monitoring Within One Year

VARIABLES	(1) No Learning	(2) Experience	(3) Time
Gen Insp	-1.218***	-1.362***	-1.562***
	(0.301)	(0.350)	(0.521)
Gen Insp * Experience		0.201***	0.051
		(0.072)	(0.043)
Observations	184,367	184,367	184,367
R^2 of FEs	0.468	0.468	0.469
R^2 of X's	0.0116	0.0126	0.0125

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include operator fixed effects, yearly fixed effects, and additional control variables. All specifications use heteroskedasticity-consistent standard errors.*

Table 7: Deterrence from Monitoring Within Two Years

VARIABLES	(1) No Learning	(2) Experience	(3) Time
Gen Insp	-1.201***	-1.345***	-1.535***
	(0.306)	(0.354)	(0.527)
Gen Insp * Experience		0.198***	0.050
		(0.072)	(0.044)
Observations	184,367	184,367	184,367
R^2 of FEs	0.468	0.468	0.469
R^2 of X's	0.0116	0.0126	0.0125

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include operator fixed effects, yearly fixed effects, and additional control variables. All specifications use heteroskedasticity-consistent standard errors.*

4.2.2 The effect of firm size

An alternative hypothesis could be that larger firms systematically have more information or more precise information than smaller firms. If the results of Section 4.1 were caused by variation in firm knowledge stock due to resource levels (for example, because large firms could hire more legal staff), then the magnitude of deterrence should be smaller for large firms. This hypothesis predicts that observed deterrence would decrease (in magnitude) with firm size. This section tests this hypothesis. Columns 1-3 of Table 8 interact the general deterrence measure with a measure of firm size. Firm size is measured as the number of wells a firm has at the time of the inspection in the county, DEP district, or state.

Here firm size is measured as the firm's number of wells (in thousands). A variety of measures of firm size exist, such as capitalization, number of employees, or revenues. In the oil and gas industry, measures of firm value are highly correlated with the value of in situ oil and gas reserves. Well counts are both closely related to reserves and available for all firms in my sample.

Table 8: Several Robustness Checks

VARIABLES	(1) County	(2) District	(3) State	(4) Pre Act 13	(5) Post Act 13
General Inspection	-3.298***	-3.694***	-3.608***	-5.868***	-2.634
	(1.135)	(1.051)	(1.023)	(1.310)	(3.435)
Gen Insp * County Wellcount	1.305				
	(1.647)				
Gen Insp * District Wellcount		1.820			
		(1.324)			
Gen Insp * State Wellcount			0.823		
			(0.621)		
Gen Insp * District Experience				0.874**	1.057
				(0.423)	(0.975)
Observations	164,544	164,544	164,544	123,137	41,331
Own Enforcement	Y	Y	Y	Y	Y
R^2 of FEs	0.473	0.474	0.474	0.464	0.554
R^2 of X's	0.0813	0.0813	0.0812	0.103	0.113

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include operator fixed effects, yearly fixed effects, and additional control variables. All specifications use heteroskedasticity-consistent standard errors. Wellcounts are measured in thousands.*

Columns 1,2, and 3 of Table 8 all have a baseline deterrence effect of 3.3-3.7 percentage points, comparable to the average effects of Table 2. The interaction terms with size are all statistically insignificant. A firm with an average number of wells in a district at the time of inspection (220) has a deterrence effect of -3.3 percent, 10 percent lower in magnitude than the smallest firm which has a single well Analogous calculations based on the specifications of Columns 1 and 3 yield similar results. Across specifications, the impact of firm size is stastically indistinguishable from zero and small in magnitude.

4.2.3 A change in regulatory status

Pennsylvania's Act 13 of 2012 substantially changed state laws for oil and gas production. These changes included increasing well bonding requirements, increased minimum well distances from sensitive sites, and increasing possible penalties for violations. Act 13 also restricted local jurisdictions abilities to restrict drilling, providing clarity to operators. The reader may be concerned that Act 13 changed the process of deterrence in Pennsylvania. To test for this, the preferred deterrence specification of Column 2 in Table 4 is reestimated on the subsamples of the data set before and after the passage of Act 13. Results are presented in Columns 4 and 5 of Table 8. In each case, the hypothesis that the coefficient on the interaction between experience and inspections is equal to the original value from Table 4. The uninteracted effect of general inspections is larger before the passage of Act 13 than after, but these estimates are not statistically different from each other or the original estimate. The effect of general deterrence might be smaller after the passage of Act 13 to the extent that Act 13 provided regulatory clarity or to the extent that firms have learned over time, but the primary results of Table 4 do not seem to be due to Act 13.

4.3 Mechanisms of firm responses

This final analysis explores what actions firms take to prevent violations. In particular, whether firms alter fixed capital deployment decisions or operational practices when changing their safety effort level. The oil & gas industry is particularly fruitful for this research question because capital decisions are primarily made at the time the well is drilled. Altering capital after this point is rare and expensive because it is primarily located far beneath the surface of the earth, down a narrow cement and steel pipe. Installed equipment is typically not changed over the course of a well's lifetime and thus can be characterized with a well-level fixed effect.

Equation 9 is similar to previous models with the addition of a well-level fixed effect ω_w:

$$y_i = \sum_{j \in [spec,gen]} \sum_{k \in [insp,fine]} \delta_{jk} Z_{ijk} + X_i \beta + \omega_w + \mu_t + \varepsilon_i \tag{9}$$

If a firm acts on deterrence by constructing more safe sites, that decision will be captured the well level fixed effect ω_w and the resulting coefficients δ_{jk} will be zero. On the other hand, if the firm responds by altering time-varying operational decisions and not altering infrastructure decisions, the δ_{jk}'s will equal the α_{jk} coefficients of equation 2. If a firm responds with some combination of infrastructure and operational adjustments, then the treatment effects will be between zero and α_{jk}'s.

Results are presented in Table 9.[11] The first two columns present familiar specifications, with a well-level fixed effect in place of the prior firm-level fixed effects. The third column includes a fixed effect for for each well-firm pair *wo*. Some wells do change hands, perhaps

[11] Estimates for control variables are included in Appendix B.

due to firm acquisitions, raising the possibility of knowledge sharing and motivating this specification.

The key finding is that the magnitude of deterrence is somewhat smaller than previous

Table 9: Deterrence Effects, Controlling for Capital

VARIABLES	(1) Any	(2) Any	(3) Any
General Inspection	-2.663***	-2.653***	-1.668***
	(0.339)	(0.340)	(0.344)
Own Inspection		0.172	-0.013
		(0.136)	(0.138)
Observations	164,544	164,544	164,544
Well FE	Y	Y	
R-2 of FEs	0.452	0.454	0.482
R-2 of X's	0.137	0.138	0.125
Own Enforcement		Y	Y
Operator x Well FE			Y

*Notes: *,**,*** indicate statistical significance at the 5%, 1 %, and 0.1% level, respectively. All regressions also include year, county, and type of inspection fixed effects. The fracking indicator is dropped from these regressions because it is perfectly correlated with the well fixed effects.*

estimates, but still economically substantial. After controlling for fixed characteristics, the average deterrence effect is 1.7-2.7%, substantially smaller than the 3.3% of Table 2. This suggests that firms do have some operational margins on which they can increase compliance effort, but that fixed capital decisions are a major avenue for firm effort at avoiding regulatory penalty.

5 Conclusions and Directions of Future Research

In summary, this paper contributes to the literature finding that regulatory enforcement actions deter both enforced-upon and other firms. I propose a model of firm behavior in which firms have uncertain beliefs about regulators and use observations about regulator behavior to learn about the true state of the regulator. This model generates the broadly observed pattern of enforcement spillovers (general deterrence) in which regulatory action towards one firm also induces other firms to increase compliance effort. I then test this model on a novel data set which allows me to identify intra-firm, intra-juristiction, and intra-time variation in firms' levels of knowledge about the regulator. My econometric analysis supports the hypothesis that firms learn about the regulator and that learning is a mechanism of observed general deterrence.

This suggests that regulators could recognize firms' learning and strategically design enforcement programs to provide more informative signals to firms, particularly if firms underestimate regulatory concern about particular violation types. An additional insight from this model is that it emphasizes that econometricians only observe *marginal deterrence* - that is, firms' changes in violation probability in response to enforcement actions. If firms were initially perfectly aware of the likelihood of stringent fines and invested in compliance to avoid being fined, that would constitute deterrence in the conventional sense of the word but would present a null result in an econometric framework. This suggests that econometric estimates of deterrence effects may represent a lower bound on the actual effectiveness of regulatory enforcement at deterring violations.

References

Barrage, L. (2015). Is regulatory rulemaking state-contingent? an empirical analysis. AERE Summer Conference.

DiTella, R. and E. Schargrodsky (2003). The role of wages and auditing during a crackdown on corruption in the city of buenos aires. *The Journal of Law and Economics 46*(1), 269–292.

Earnhart, D. (2004). Panel data analysis of regulatory factors shaping environmental performance. *Review of Economics and Statistics 86*(1), 391–401.

Earnhart, D. and L. Friesen (2013). Can punishment generate specific deterrence without updating? analysis of a stated choice scenario. *Environmental and Resource Economics 56*(3), 379–397.

Earnhart, D. and K. Segerson (2012). The influence of financial status on the effectiveness of environmental enforcement. *Journal of Public Economics 96*(9), 670–684.

Feinstein, J. S. (1990). Detection controlled estimation. *The Journal of Law & Economics 33*(1), 233–276.

Fisman, R. and E. Miguel (2007). Corruption, norms, and legal enforcement: Evidence from diplomatic parking tickets. *Journal of Political economy 115*(6), 1020–1048.

Gray, W. B. and R. J. Shadbegian (2004). Optimal pollution abatement: whose benefits matter, and how much? *Journal of Environmental Economics and Management 47*(3), 510–534.

Gray, W. B. and R. J. Shadbegian (2005). When and why do plants comply? paper mills in the 1980s. *Law & Policy 27*(2), 238–261.

Gray, W. B. and R. J. Shadbegian (2007). The environmental performance of polluting plants: A spatial analysis. *Journal of Regional Science 47*(1), 63–84.

Hanna, R. N. and P. Oliva (2010). The impact of inspections on plant-level air emissions. *The BE Journal of Economic Analysis & Policy 10*(1).

Helland, E. (1998). The enforcement of pollution control laws: Inspections, violations, and self-reporting. *Review of Economics and Statistics 80*(1), 141–153.

Keohane, N. O., E. T. Mansur, and A. Voynov (2009). Averting regulatory enforcement: Evidence from new source review. *Journal of Economics & Management Strategy 18*(1), 75–104.

Konisky, D. M. (2009). Inequities in enforcement? environmental justice and government performance. *Journal of Policy Analysis and Management 28*(1), 102–121.

Levitt, S. D. (1997). Using electoral cycles in police hiring to estimate the effect of police on crime. *The American Economic Review*, 270–290.

Magat, W. A. and W. K. Viscusi (1990). Effectiveness of the epa's regulatory enforcement: The case of industrial effluent standards. *The Journal of Law & Economics 33*(2), 331–360.

Pastor, L. and P. Veronesi (2009). Learning in financial markets. *Annu. Rev. Financ. Econ. 1*(1), 361–381.

Pennsylvania Department of Environmental Protection (2013). 2013 oil and gas annual report. Technical report.

Polinsky, A. M. and S. Shavell (2000). The economic theory of public enforcement of law. *Journal of Economic Literature 38*(1), 45–76.

Shimshack, J. P. (2014). The economics of environmental monitoring and enforcement. *Annual Review of Resource Economics 6*(1), 339–360.

Shimshack, J. P. and M. B. Ward (2005). Regulator reputation, enforcement, and environmental compliance. *Journal of Environmental Economics and Management 50*(3), 519–540.

A Theoretical proofs

A.1 Proof of Lemma 1

The key to the proof of Lemma 1 is realizing that as a firm receives more signals, the change in belief from recieving a new signal gets smaller. This follows from Equation 6, rewritten below.

$$\Delta\hat{\theta}_{T+1} = \frac{s_{T+1} - \hat{\theta}_T}{1 + \sigma^2/\hat{\sigma}_T^2}$$

I begin by noting that the definition of $\hat{\theta}_T$ in Equation 4 implies $E[\hat{\theta}_T] \to E[s_{T+1}]$ as $T \to \infty$, so $E[s_{T+1} - \theta_T] \to 0$ as $T \to \infty$.

Next I note that Equation 5 (reprinted below) implies that $\hat{\sigma}_T^2 \to 0$ as $T \to \infty$. This implies that $1 + \sigma^2/\hat{\sigma}_T^2 \to 0$ $T \to \infty$.

$$\hat{\sigma}_T^2 = \frac{1}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}}$$

Taken together, the above two findings imply that as $T \to \infty$, $\Delta\hat{\theta}_{T+1} \to \frac{0}{\infty}$, or 0.

Recalling Section 2.1, the firm responds to its change in belief $\Delta\hat{\theta}_{T+1}$ by updating its effort level. Using a differential approximation, we see that the smaller the change in $\hat{\theta}_{T+1}$, the smaller the change in ρ^*:

$$\Delta\rho^* = \frac{\partial\rho^*}{\partial\hat{\theta}_T} * \Delta\hat{\theta}_{T+1}$$

Of course, the probability of penalty is a function of firm effort and the true value of regulator effort θ.

$$p(penalty) = F(\frac{\theta}{\rho})$$

This implies that a smaller change in ρ (or ρ^*) will lead to a smaller change in the probability of penalty.

A.2 Proof of Lemma 2

Intuitively, if the signal has zero noise, then all signals are exactly θ and the firm adopts a signal value as it's belief $\hat{\theta}$. Because all signals are equal, the firm will have exact knowledge of regulatory effort and be able to set its own effort accordingly.

We know from Equation 5 (reprinted below) that as the noise in signals approaches zero ($\sigma^2 \to 0$), then the variance in the firm's belief approaches zero once the firm has observed a single signal.

$$\hat{\sigma}_T^2 = \frac{1}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}}$$

Similarly, we know from Equation 6 (reprinted below) that as the variance in the firm's belief approaches zero, its change from a new signal approaches zero. The term $\hat{\sigma}_T^2 \to 0$, which means that the denominator approaches infinity.

$$\Delta\hat{\theta}_{T+1} = \frac{s_{T+1} - \hat{\theta}_T}{1 + \sigma^2/\hat{\sigma}_T^2}$$

Therefore as signaling uncertainty goes to zero, so do changes in firm behavior.

However, the firm is choosing an optimal value ρ^* based on its knowledge of θ. As long as $\theta > 0$, $\rho^* > 0$. This implies that the firm is exerting effort at avoiding costly penalties for violation (and thus being deterred as the term is commonly understood), but that the firm will not respond to signals. As an econometrician will only observe *changes* in firm effort or violation probability in response to signals[12], the econometrician will observe no marginal effect of deterrence even while the firm is effectively being deterred.

A.3 Proof of Correlary 1

Uncertainty about regulatory behavior is measured in the variance of the signals σ^2 and in the initial beliefs by the bias $\hat{\theta}_0 - \theta$ and by the variance σ_0^2. It obtains trivially that as the bias and variance of initial beliefs decrease, $E[\hat{\theta}_0] \to \theta$. And as σ^2 decreases, $E[\bar{s}] \to \theta$.

Now recall the definition of our Bayesian belief that

$$\hat{\theta}_T = \hat{\theta}_0 \frac{\frac{1}{\sigma_0^2}}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}} + \bar{s} \frac{\frac{T}{\sigma^2}}{\frac{1}{\sigma_0^2} + \frac{T}{\sigma^2}}$$

We can rewrite this as a weighted sum

$$\hat{\theta}_T = \hat{\theta}_0 * \alpha + \bar{s} * (1 - \alpha)$$

As both $\hat{\theta}_0$ and $\bar{s}$ are closer to the true value θ when there is less uncertainty about the regular, so does $\hat{\theta}_T$.

[12] Particularly one who liberally uses fixed effects

B Additional regression results

Table 10: Control Variables for Table 2

VARIABLES	(1) Any	(2) Major	(3) Any
Fracking Well	2.834***	3.374***	3.559***
	(0.242)	(0.697)	(0.676)
Site Inspection Number	-0.138***	-0.065*	-0.065*
	(0.009)	(0.034)	(0.034)
Own Fine			-0.041***
			(0.005)
General Fine		0.712*	0.801**
		(0.366)	(0.355)
Site First Inspection	2.794***	1.052*	1.050*
	(0.169)	(0.625)	(0.616)
Observations	164,544	164,544	164,544
R-squared	0.111	0.076	0.079

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications use heteroskedasticity-consistent standard errors. All regressions also include year, county, and type of inspection fixed effects. As discussed in the main text, own inpection and fine coefficiencts may be positively biased.*

Table 11: Control Variables for Table 9

VARIABLES	(1) Base	(2) Own Deterrence	(3) Own Deter & Well-Oper FE's
Site Inspection Number	-0.163***	-0.162***	-0.210***
	(0.031)	(0.031)	(0.034)
Own Inspection		0.172	
		(0.136)	
Own Fine		-0.024***	
		(0.003)	
Site First Inspection	5.054***	5.021***	4.544***
	(0.237)	(0.237)	(0.242)
Observations	164,544	164,544	164,544
Well FE	Y	Y	
R-2 of FEs	0.452	0.454	0.482
R-2 of X's	0.137	0.138	0.125
Own Enforcement		Y	Y
Operator x Well FE			Y

*Notes: *,**,*** indicate statistical significance at the 5%, 1 %, and 0.1% level, respectively. All regressions also include year, county, and type of inspection fixed effects. The fracking indicator is dropped from these regressions because it is perfectly correlated with the well fixed effects. As discussed in the main text, own inpection and fine coefficiencts may be positively biased.*