Can Learning Explain Deterrence? Evidence from Oil & Gas Production

Peter Maniloff*

October 16, 2018

Abstract

This paper tests a learning model of regulatory deterrence. Firms exert compliance effort based on their belief about a regulator's effort level at detecting violations. Firms use regulatory actions to learn about the regulator and update their own compliance efforts accordingly. This theoretical model suggests that deterrence will decrease with experience. Econometric analysis of inspections of Pennsylvania oil and gas wells supports the model. Econometric results show that inexperienced firms are substantially more deterred than experienced firms. These results are robust to regulatory targeting in inspections and different measures of experience and deterrence.

Keywords: Enforcement, Deterrence, Reputation, Oil and Gas, Hydraulic Fracturing
JEL codes: D22, K32, L51, L71, Q58

*Division of Economics and Business, Colorado School of Mines, 1500 Illinois St, Golden, CO, 80401. phone: 303-273-3481 fax: 303-273-3416 email: maniloff@mines.edu. I thank Will Duey, Dietrich Earnhart, Jonathan Eyer, Harrison Fell, Dan Kaffine, Ian Lange, Jay Shogren, Marty Smith, as well as seminar participants at Colorado State University, the Colorado Environmental and Resource Economics Workshop, Environmental Defense Fund, Iowa State University's workshop on the Economics of Water and Energy, and North Carolina State University 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

A core question of applied microeconomics is how to design efficient regulations. However, rational firms may violate regulations if the marginal benefits of violation exceed the marginal costs. This suggests that designing enforcement mechanisms also holds a key role in policy.

Many papers have found evidence that firms 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). However, the causal mechanisms of general deterrence are poorly understood. Shimshack and Ward (2005) describe general deterrence as representing a regulator's reputation, which implies a causal mechanism that depends on firms learning about the regulator. However, firm learning about the regulator has not been observed in the literature.

This paper asks whether firm learning can explain general deterrence. A theoretical model in the spirit of Sah (1991) shows that firms can use inspections of other firms to learn about a regulator's priorities, resulting in general deterrence. This model has several implications, including that the magnitude of deterrence decreases with experience. The paper then tests this implication on a data set of inspections of oil and gas wells in Pennsylvania. The core econometric finding is that the magnitude of deterrence decreases with experience, providing empirical support for the hypothesis that learning can explain deterrence amongst firms.

Prior literature has found evidence of deterrence for both individuals and firms. Studies of individuals have shown deterrence of violent crime (Levitt 1997), property crime (Levitt 1998), parking violations (Fisman and Miguel 2007), fee shirking (Rincke and Traxler 2011), tax avoidance (Andreoni, Erard, and Feinstein 1998), and other types of lawbreaking. Sah (1991) proposes a theoretical model of deterrence in which individuals learn about their local criminological environment, and shows that learning can lead to deterrence. Lochner (2007) provides econometric support for the importance of individual beliefs about the regulator in the context of crime.

Studies of firms and organizations have also found evidence of deterrence in industries as varied as 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 and firms follow the law can be quite heterogeneous, even under a shared regulator. Prior work has shown the importance of heterogeneity in cultural norms (Fisman and Miguel 2007), violation penalties (DiTella and Schargrodsky 2003), compliance costs (Earnhart 2004), and financial status (Earnhart and Segerson 2012).

This paper makes two distinct contributions. First, it provides econometric evidence that supports applying the learning model of Sah (1991) to firms. The main results provide evidence that firms form their beliefs about regulator behavior in a learning 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 be caught violating by half. For very experienced firms, this general deterrence effect is nearly halved. Second, this learning model provides an additional explanation for heterogeneity in deterrence. If firms have varying initial beliefs or experience levels, then they will revise their beliefs differently and thus exhibit different levels of deterrence.

An important consequence of the theoretical model is that it implies that in settings with sufficiently experienced firms or sufficiently precise beliefs about the regulator, econometricians will not be able to reject a null hypothesis of zero deterrence. This will hold even if firms are exerting substantial care to avoid being detected violating.

The oil and gas industry provides an excellent opportunity to study deterrence. Firms operate a number of wells across different enforcement jurisdictions, across extended time periods, and enter at different times. Their production technologies are similar, and production largely depends on fixed geological characteristics. Outputs are largely homogeneous across firms. Thus in any given time and regulatory jurisdiction, firms face largely homogeneous production functions and markets but have varying information about the regulator.

The paper continues as follows. Section 2 briefly outlines the industrial context, including some of the major regulations that firms face and how they can comply with these regulations. Section 3 develops a theoretical model of Bayesian learning and general deterrence and derives testable predictions. Section 4 discusses the identification strategy and data. Section 5 discusses an econometric model to test the theoretical model and presents estimation results. Section 6 explores whether firm learning is socially beneficial. Sections 7 provides evidence that the core results are robust to regulatory targeting or sample selection. Section 8 concludes.

2 Background Information

2.1 Oil and Gas Regulation and Compliance

Pennsylvania oil and gas producers face a variety of environmental regulations under both oil and gas laws and broader environmental laws.[1] This section provides a brief overview of some of the major environmental regulations and on steps operators can take to comply.

Chapter 78 of Pennsylvania's state code provides detailed specifications on the construction and operation of a well. This chapter was issued under Pennsylvania's Oil and Gas Act, Oil and Gas Conservation Law, and other laws. One set of requirements specifies technical requirements for well construction, for example specifying permeability and strength of the cement used in well construction. The code also provides regulations on testing the physical integrity of the well during operation, and on appropriate plugging of wells when a firm has completed production and is ready to permanently close the well (State of Pennsylvania a).

A variety of environmental regulations also apply. The Pennylvania Clean Streams Law prohibits discharge of industrial wastes "into any waters of the Commonwealth" (State of Pennsylvania b). Specifications for adequate containment of waste fluids include rules such as a requirement that storage pits have walls at least two feet higher than the maximum fluid level, along with specifications such as the thickness of materials used to line pits and

[1] They also face federal laws including occupational safety laws, minimum wages, tax obligations, et cetera.

of any subbase below the liner. Operators are required to mitigate sediment runoff during construction phases (for example, leveling the ground) via best management practices as described in the Oil And Gas Operators Manual.

Several regulations are designed to minimize harm in case of an incident at a well. If a well has associated large storage tanks, operators are required to build adequate containment dikes around the tanks. Additionally, several requirements are meant to aid first responders to any incident - these include posting adequate signage and regularly mowing any high grass.

Finally, there are a variety of reporting and notification requirements. Operators are required to submit reports on oil and gas production, to promptly notify authorities of spills, and to provide the state with plans to control erosion.

For requirements on the construction of the well itself, firms can effectively only comply at the time of drilling. Altering well characteristics like the type of cement used in the walls of the well after completing the well would be technically and financially challenging. Firms can take relatively direct steps to comply with most other rules at any stage of the production lifecycle, and can do so with moderate cost and time requirements. Compliance with these surface regulations typically requires construction work (for example to improve a containment berm or plug an abandoned well), but can be as minor as posting signage, mowing, or filing paperwork.

These regulations are primarily enforced by the Pennsylvania Department of Environmental Protection (PA DEP). The PA DEP conducts inspections of well sites and can issue fines in response to any violations it finds. There are three regional offices which manage inspections (Eyer 2018).

2.2 The Pennsylvania Oil and Gas Industry

There is a long history of oil and gas production in Pennsylvania. Since the Titusville oil rush of 1859, oil and gas production has continued at relatively modest levels to the current day. By the 1990's, production had gradually declined to about 5,000 barrels of oil and 400 mcf of natural gas per day.

In the 2000's, the fracking boom came to Pennsylvania.[2]. This led to a dramatic increase in the pace of drilling, and by the late 2000's and 2010's led to a boom in production. New well drilling during this period was primarily of unconventional (fracking) wells. Unconventional wells faced new technical and compliance problems deep underground, but shared many surface and operational compliance issues with conventional wells.

3 Theoretical Model

In the classic model of enforcement, firms choose a "level of care" to satisfy an objective function. A typical objective function would be to minimize the expected sum of cost of care and cost of violation penalties. Care, which describes a firm's effort to avoid violating rules or to avoid being caught violating rules, is costly.[3] Regulators choose their effort towards detecting violations and the stringency of penalties for detected violations in order to maximize a social welfare function such as minimizing the sum of expected damages from violations, regulatory agency costs, and firm compliance costs. A variety of models explore different aspects of enforcement (Becker 1968; Polinsky and Shavell 2000; Langpap 2007; Huang and Rios 2016). This section develops a model of one particular aspect of the firm problem - how firms understand the regulator's degree of effort at detecting violations.

The model focuses on violation detection and assumes penalties are fixed. Penalties conditional on a violation are largely fixed by statute in my study sample.[4] This means that there is little ambiguity or scope for learning about penalties. Penalties are also relatively rare and typically small in this setting. Eckert (2004) shows that inspections can be an effective deterrent even in the face of low financial penalties. Blundell, Gowrisankaran, and Langer (2018) suggests that this may be because firms want to avoid repeat-violator status, which would lead to higher costs or penalties.

[2]Technical advances in both horizontal drilling and hydraulic fracturing of underground rocks enabled profitable production from shale formations. These shale formations were previously known, but production had been technically infeasble. This combination of techniques is commonly referred to as "fracking" or "fracing". For a detailed discussion of the fracking boom, see Gold (2014) or Raimi (2017).

[3]Here I use "care" in the sense of Segerson and Tietenberg (1992)'s "level of care": the actions that the firm takes to avoid being found in violation. Being found to be in violation does not imply necessarily carelessness in the plain language sense.

[4]The formulae for calculating penalties are described in State of Pennsylvania (2002).

I model the firm's learning process as passive Bayesian updating - meaning that the firm does not actively seek new information, and that the firm integrates new information via a Bayesian updating framework. Bayesian learning is a common theoretical framework for modeling firm learning (Nagypál 2007; Pastor and Veronesi 2009). There is econometric evidence that firms do behave as Bayesian learners in settings including retail (Pakes and Ericson 1998), finance (Pástor and Veronesi 2006), oil and gas production (Levitt 2009), and beliefs about monetary policy (Bianchi 2012). Moreover, a larger literature has provided evidence of firm learning, albeit often without explicitly modeling the learning process (Wright 1936; Argote and Epple 1990; Kellogg 2011; Levitt, List, and Syverson 2013).

This paper's model of Bayesian learning is based on the results of Pakes and Ericson (1998) and Covert (2015). Pakes and Ericson (1998) develop models of both passive Bayesian learning and active experimentation and show that each model seems to apply to different industries. Covert (2015) shows that oil producers in North Dakota primarily learn passively instead of actively experimenting. While oil producers in North Dakota and oil and gas producers in Pennsylvania are relatively similar industrial settings, I would not take the strong stance that Covert (2015) and Pakes and Ericson (1998) clearly show that passive Baysian learning is the sole correct model. Instead, these papers suggest that passive Bayesian learning is a plausible model for oil and gas producers, and the following theoretical model explores further implications of passive Bayesian learning in a compliance setting.

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 nonstrategic regulator chooses a level of effort towards detecting violations which is unknown to the firm.[5] The firm learns about regulator stringency by observing regulatory activities or signals, then the firm chooses its own compliance effort level accordingly. The model implies that firms will change their effort based on observations

[5]The literature on regulatory enforcement typically distinguishes between the probability of being penalized and the severity of a penalty (Polinsky and Shavell 2000). I focus on the probability of being penalized due to characteristics of my data set. Extending this model to consider multiple dimensions of regulatory behavior (such as jointly modeling probability of catching violators and severity of penalty) is straightforward but adds little intuition to the model.

of the regulator's behavior. If the firm increases its believed level of regulator effort, it will increase its own 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 sends a signal about its effort level. Operationally, this might be inspecting a facility, inspecting many facilities, or doing nothing.

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

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

4. Fourth, the firm updates its own effort level towards avoiding being detected violating in light of its new belief.

5. Fifth, the firm's change in effort changes the probability of a violation.

The remainder of the section models this sequence more formally. Section 3.1 models the firm choice in step four. Section 3.2 models steps two and three. Section 3.3 combines steps two through four to show step five. The regulatory decision process is taken as exogenous to the model.[6]

3.1 The Firm's Choice of Effort

Consider firm i with experience n which can choose an effort level $\rho_{i,n} \in \left[\underline{\rho}, \bar{\rho}\right]$ towards avoiding penalty. This effort can entail reductions in the probability of violating rules (adopting safe but costly practices) or exerting effort to avoid being caught breaking rules.[7] In either

[6] In the context of this study, total regulatory effort level is largely determined by agency budgets and staffing, which are effectively set on an annual level by legislators. There may be discretion at the agency level about what areas or types of wells to target. A related literature explores how regulators make decisions - how they decide what sites to inspect, what penalties to levy, and the stringency of regulations. These questions are of substantial interest, but beyond the scope of this study. I do plan to explore regulator behavior in future work. Prior 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; Konisky 2009; Barrage 2018).

[7] Firm effort to avoid being caught is typically referred to as "avoidance" activity. This model will not distinguish between avoidance activity and actual changes in practices; they will all be considered choices a firm can make to reduce its probability of penalty. Section 6.2 provides suggestive evidence distinguishing between avoidance and real compliance effort, but I leave direct tests for further study.

case, a higher level of effort reduces the probability of a violation occurring. Effort is costly, so in the absence of regulatory enforcement the firm would choose $\rho_{i,n} = \underline{\rho}$. The firm has beliefs $\hat{\theta}_{i,n}$ about the regulator's degree of enforcement effort θ. The probability of a violation being detected is a function of θ and $\rho_{i,n}$ such that the probability of a detected violation is

$$P\left(\theta, \rho_{i,n}\right) \tag{1}$$

where $P_\rho\left(\theta, \rho_{i,n}\right) < 0$, $P_\theta\left(\theta, \rho_{i,n}\right) > 0$, $P_{\rho\rho}\left(\theta, \rho_{i,n}\right) > 0$, and $P_{\rho\theta}\left(\theta, \rho_{i,n}\right) < 0$. That is, increasing firm care reduces the probability of penalty, increasing regulatory effort increases the probability of penalty, and increasing regulatory effort reduces the effectiveness of firm care. For the risk-neutral firm which seeks to minimize the sum of compliance costs and expected penalties, the firm solves the minimization problem

$$\min_{\rho_{i,n}} \left(C(\rho_{i,n}) + P(\hat{\theta}_{i,n}, \rho_{i,n})R \right) \tag{2}$$

$$s.t. \quad \rho_{i,n} \in \left[\underline{\rho}, \bar{\rho}\right] \tag{3}$$

where $C(\rho_{i,n})$ is the cost of care and R is the fixed positive penalty. Note that the firm is choosing $\rho_{i,n}$ based on its belief $\hat{\theta}_{it}$ and not on the true value of θ. $C_\rho(\rho_{i,n})$ and $C_{\rho\rho}(\rho_{i,n})$ are assumed to be positive, continuous, and bounded. Differentiating with respect to $\rho_{i,n}$ and noting that Equation 4 defines the optimal choice $\rho^*_{i,n}$, we find the first order condition

$$C_\rho(\rho^*_{i,n}) + P_\rho(\hat{\theta}_{i,n}, \rho^*_{i,n})R = 0 \tag{4}$$

Applying the implicit function theorem to Equation 4 yields

$$\frac{\partial \rho^*_{i,n}}{\partial \hat{\theta}_{i,n}} = \frac{-P_{\rho\theta}\left(\hat{\theta}_{i,n}, \rho^*_{i,n}\right) R}{C_{\rho\rho}(\rho^*_{i,n}) + P_{\rho\rho}(\rho^*_{i,n})R} \tag{5}$$

which is positive.

Equation 5 states that if a firm learns that a regulator exerts more (less) effort towards detecting violations than previously believed, the firm will exert more (less) effort to avoid penalty. Because the probability of penalty depends on the unchanged true value θ and not the firm's belief $\hat{\theta}_{i,n}$, if the firm updates its belief in $\hat{\theta}_{i,n}$ positively, the actual probability of a violation will decrease. Put formally, $\frac{\partial P\left(\theta,\rho^*_{i,n}\right)}{\partial\hat{\theta}_n} < 0$.

Equation 2 assumes that the firm choice of $\rho_{i,n}$ is separable from other firm choices such as production level. While this is a special case in general, it is a reasonable simplifying assumption for oil and gas producers. In this industrial setting, production levels and input decisions are largely determined by geologic endowments. Implications of this assumption are discussed in Section 3.3.

3.2 The Firm's Learning Problem

The firm has an initial noisy belief about θ, and the firm updates this belief as it observes informative signals such as inspections or fines.[8] The firm's initial belief $\hat{\theta}_{i,n}$ about θ is a random variable distributed normally with mean $\hat{\theta}_0$ and variance σ_0^2. The firm interprets regulator actions such as inspections or fines as noisy informative signals about θ. Each signal $s_{i,n}$ is distributed normally with mean θ and variance $\sigma^2 > 0$. If a firm has observed n signals, then the Bayesian updated prior is

$$\hat{\theta}_{i,n} = \hat{\theta}_0 \frac{\sigma^2}{\sigma^2 + n\sigma_0^2} + \bar{s}_{i,n} \frac{n\sigma_0^2}{\sigma^2 + n\sigma_0^2} \tag{6}$$

where $\bar{s}_{i,n}$ is the average value of the signals. The firm's belief $\hat{\theta}_{i,n}$ 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. If the variance in the initial belief σ_0^2 is large, then the average signal $\bar{s}_{i,n}$ is accorded relatively more weight, whereas if the variance in the signals σ^2 is large, the initial belief $\hat{\theta}_0$ is given more weight. We see immediately that as n becomes large, $\hat{\theta}_{i,n}$ approaches $\bar{s}_{i,n}$.

The variance in the updated prior is[9]

$$\hat{\sigma}^2_{i,n} = \frac{1}{\frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}} \tag{7}$$

[8]Equations 6-8 are based on standard models of Bayesian learning, described in detail in textbooks including Lynch (2007). This model has been adapted in other contexts, such as Pastor and Veronesi (2009).

[9]This is equivalent to $\frac{\sigma^2\sigma_0^2}{\sigma^2+n\sigma_0^2}$, in notation more like that of equation 6.

We can also describe the change in a firm's belief as it observes a signal. If the firm has already observed n signals, then the change from one more $\Delta\hat{\theta}_{n+1} \equiv \hat{\theta}_{n+1} - \hat{\theta}_n$ is

$$\Delta\hat{\theta}_{i,n+1} = \frac{s_{i,n+1} - \hat{\theta}_{i,n}}{1 + \sigma^2/\hat{\sigma}^2_{i,n}} \tag{8}$$

The numerator provides us the intuitive result that the firm will adjust its prior upwards if $s_{i,n+1} > \hat{\theta}_{i,n}$ and downwards if $s_{i,n+1} < \hat{\theta}_{i,n}$. 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. As the firm's uncertainty $\hat{\sigma}^2_{i,n}$ approaches zero, the magnitude of the adjustment will shrink.

3.3 Implications of the Theoretical Model

This model has several implications:[10] First, the marginal effect of a signal on a firm's change in probability of being caught violating will approach zero as the number of signals a firm has observed approaches infinity.

Lemma 1. *$lim_{n\to\infty}\Delta P \to 0$ where $\Delta P \equiv P(\hat{\theta}_{i,n+1}, \rho^*_{i,n+1}) - P(\hat{\theta}_{i,n}, \rho^*_{i,n})$,*

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, $\Delta\hat{\theta}_{i,n+1}$ is smaller in absolute value for the experienced firm, if both firms have the same $\hat{\theta}$. (That is, if $\hat{\theta}_{i,n} = \hat{\theta}_{j,m}$ for firms i and j where $m \neq n$.) This smaller adjustment in belief leads to a smaller adjustment in firm effort.

As a corollary, an external event which reduces uncertainty about the regulator will also reduce the marginal effect of a signal on a firm's change in probability of violation.

Corollary 1. $\frac{\partial \Delta P}{\partial \sigma_0^2} > 0$

Intuitively, if the initial belief about the regulator has low variance (σ_0^2 is small), then variance in belief $\hat{\sigma}^2_{i,n}$ will be small and therefore the adjustment $\Delta\hat{\theta}_{i,n+1}$ will be small (holding all else equal).

[10] For derivations of each Lemma and Corollary, see Appendix A.

Lemma 2 describes a limitation of reduced form econometric approaches to estimating deterrence. This Lemma tells us that if a firm has exact knowledge of regulatory effort ($\hat{\sigma}^2_{i,n} = 0$) 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.

Lemma 2. *Under certain conditions, $\Delta P = 0$ and $\rho^*_{i,n} > \underline{\rho}$.*

Imagine a firm with exact knowledge of regulatory effort ($\hat{\sigma}^2_{i,n} = 0$) which exerts effort $\rho^*_{i,n} > \underline{\rho}$. If the regulator conducts enforcement actions, the firm will not update its belief $\hat{\theta}_{i,n}$ because it already knows θ. This implies that the firm will also not change $\rho_{i,n}$ and there will be no observable change in the probability of violation. Nonetheless, the firm is exerting more 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.

An important implication of Lemma 2 is that an econometrician who studies the effect of regulatory effort on deterrence may find a null result even in the presence of significant actual deterrence if firms are experienced. This could lead to underestimates of the real effectiveness of enforcement effort.

These theoretical predictions do have several important limitations. First, the prediction that a firm would not update its belief at all is quite extreme. It relies on the assumption that a firm's experience has gone to infinity while the regulator has not changed its effort level at all. In practice, regulations, agency budgets, and enforcement priorities can all change over time. This suggests that in practice we might see that the magnitude of firm adjustments decrease during a given legal-budgetary-priority regime, but that this might change when regimes change.

Another important limitation is that this model assumes that the level of care ρ^* is separable the firm's output decision. While output is largely determined by geologic characteristics (Economides, Hill, Ehlig-Economides, and Zhu 2013), it is possible that an increased level of care would be associated with reduced output. In some cases, we can straightforwardly consider the reduced output to be simply the cost of care - for example, if compliance does not affect output quantity but would defer output by some number of days, then the present value of the project would be reduced by discounting. This lost present value would be part of the compliance cost. However, care could also reduce firm output quantity. Nonseperability of care and output has several implications. First, it generally reduces the efficient level of environmental policy stringency because of the policy's effect on surplus in the output market. More subtly, assuming that output and care are seperable restricts the set of compliance options available to firms.[11] Models of compliance may thus over-estimate compliance costs, and rules based on separability may result in higher costs than more flexible rules (Schwabe 1999).

In this context, non-separability could imply that a firm which wishes to increase compliance would change its production mix. This could then lead to a selection problem. While it is primarily motivated by concern about targeting, Section 7 explores the impact of potential sample selection and finds results consistent with the paper's core results.

3.4 Testing the Theoretical Model

Let us conduct a thought experiment to consider how this model might be observed in data. Consider firm i with experience n which observes the regulator conducting a large number k of nearby inspections, which suggests that the regulator has a higher effort level than the firm initially believed. The firm would increase its belief in the regulator's effort from $\hat{\theta}_{i,n}$ to $\hat{\theta}_{i,n+k}$. It will then increase its level of care to $\rho^*_{i,n+k}$. This will lead to a reduction in the actual probability of being caught violating. So in the data, we would expect to see a reduction in the probability of being caught violating after an increase in nearby inspections. This is learning, manifested as general deterrence.

Now consider firm j with experience m ($m > n$) which operates an identical well to the first firm and has the same belief of regulatory effort ($\hat{\theta}_{j,m} = \hat{\theta}_{i,n}$). If the second firm observes the same k nearby inspections as the first, it will also increase its belief in regulatory effort and thus increase its level of care, decreasing its true probability of being caught violating.

[11] Rubin and Kling (1993) provide an example of non-seperable care in which automakers respond to an environmental constraint by changing the mixture of models they sell, not just by improving the environmental performance of each model.

However, due to its experience, it will adjust its belief by a smaller amount than the first firm, resulting in a smaller reduction in the probability of being caught violating.

This stylized example of two firms suggests estimating the probability of being caught violating as a function of (a) how many nearby inspections the firms have recently seen and (b) the firms' experience levels. This model suggests that violation probability will be decreasing in the number of recently observed inspections, but that the magnitude of this effect will decrease with experience.

Sections 4-7 develop and implement econometric tests for the theoretical model. Testing the model requires measures of the probability of being found violating $P(.)$, number of signals n, and the average of observed signals $\bar{s}_{i,n}$. The probability of a detected violation will be measured with an indicator variable for whether an inspection detected a violation. The number of signals, or experience, will be measured in two different ways. The first is the length of time a firm has been operating in the area. In this conception, a firm acquires a signal each day based on the how many similar wells the regulator inspects that day. In a distinct measure of experience, I use the number of previous inspections the firm has had. If a firm is able to gather information from interactions with regulatory staff during an inspection, then the number of inspections may be a more appropriate measure of experience. Finally, the signal average $\bar{s}_{i,n}$ is measured by the fraction of similar wells which have been inspected recently, where similarity is defined as being nearby (in the same county). Each of these measures is discussed in more detail in Section 4.

4 Identification Strategy and Data

4.1 Identification Strategy

This paper's identification strategy follows the thought experiment of Section 3.4. Ideally, I would observe the probability of being caught violating at wells controlled by firms with the same belief, but different levels of experience. I could then compare the two firms' reaction to new information. While I cannot directly observe either the true probability of being caught violating or firms' beliefs, I can observe whether or not a firm was caught violating at periodic inspections as well as proxy variables for the set of signals a firm has received. I will thus regress indicator variables for whether a violation was observed on proxy measures for firms' information sets.

I construct daily measures of $\bar{s}_{i,n}$ for every oil and gas well in Pennsylvania. I also observe all well inspections by the Pennsylvania Department of Environmental Protection. For each inspection, I regress whether or not a violation was observed on a measure for $\bar{s}_{i,n}$ and the interaction of $\bar{s}_{i,n}$ and a measure of the number of signals the firm has observed n. Based on my theoretical model, I expect $\bar{s}_{i,n}$ to have a negative effect on likelihood of violation, but that this effect will be smaller in magnitude for more experienced firms.

Identification relies on variation in treatment as well as fixed effects at the firm, county, and year level. By observing firms across jurisdictions, I can observe the same firms with different experience levels and information sets, allowing me to control for unobserved firm heterogeneity. I also observe multiple firms entering each jurisdiction at different times, which provides variation in both experience and observed signals. Fixed effects for jurisdictions, firm, well, and time control for unobserved factors. The effect of experience is thus identified based on variation in experience levels between firm-dates, controlling for firm, time, and jurisdiction effects.

The first identifying assumption is that a firm's decision to drill a well in a county is random, conditional on observables and fixed effects. This is a strong assumption, particularly given that papers on the impacts of oil and gas production typically instrument for the location of oil and gas wells (Weber 2012; Feyrer, Mansur, and Sacerdote 2017; Maniloff and Mastromonaco 2017). I thus relax this assumption by allowing for well-level fixed effects in some specifications. The identifying assumption with well-level fixed effects thus becomes that time-variant well unobservable violation probability is uncorrelated with observables such as experience and deterrence measures. The second assumption is that inspections are conditionally random - in particular, that they are uncorrelated with unobserved propensity to violate.

Section 7 relaxes both assumptions and tests whether inspections are non-randomly targeting likely violators or whether time-varying characteristics are otherwise endogenous, and finds little evidence that inspections and outcomes are correlated. This finding is somewhat striking, given literature such as Harrington (1988) which shows that it can be optimal for regulators for inspect violators more frequently than nonviolators. However, this null result is consistent with Eyer (2018)'s finding that firms' past violation history is not associated with increased inspections by the Pennsylvania DEP.

4.2 Data

My primary data set consists of inspection records for all inspections of oil and natural gas wells conducted by the PA DEP from 2000-2014. The unit of observation is an individual inspection of a single oil or gas well. 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.

PA DEP enforces environmental rules for oil and gas production. They employ approximately 76 inspectors who conduct field inspections of wells.[12] 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, at the time of well closure and site restoration, or for similar reasons. 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.

Inspectors conducted 184,367 inspections from 2000-2014, of which 9% found violations. Figure 1 shows the annual probability of inspection (number of inspections divided by number of wells), as well as the fraction of inspections which result in a violation. The inspection rate initially decreased as more wells were drilled, and then increased as the state was able to do more inspections. By contrast, the violation rate has generally decreased over time.

An inspection record also includes fines levied in response to violations, if any.[13] Of violations, 9 percent result in fines (which is approximately 0.8% of total inspections). Conditional on being nonzero, the median fine is $6000 while the mean is $44,384. Fines are right skewed, dominated by a small number of very large penalties.

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

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

Figure 1: Per Well Inspections and Violations

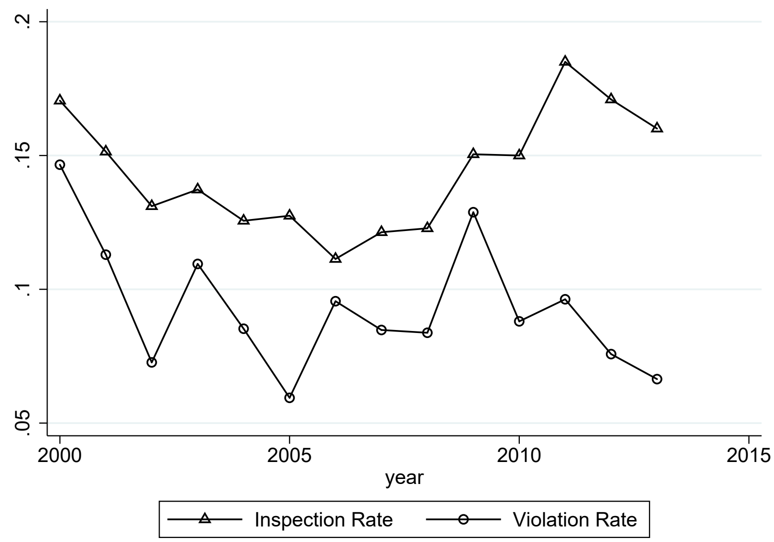

Firms are able to learn about inspections of other firms' wells. Inspection records are released publicly on DEP's website. While updating the database may take days to weeks, firms can observe what wells have been inspected, and how recently, by downloading this data from DEP's website.

4.3 Description of Constructed Variables

The econometric analysis relies on constructing measures of both firm experience and of the signals firms have observed. Table 1 lists the four variables used to measure signals $\bar{s}_{i,n}$ as well as the two measures of experience n. Each is discussed below.

Table 1: Constructed Variables

Signals	
Specific Deterrence	Predicted Own Inspection Rate
Specific Deterrence	Own Fines
General Deterrence	Others' Inspection Rate
General Deterrence	Others' Fines
Experience	Number of Inspections
Experience	Years Since First inspection

First let us discuss the construction of signals. There are two measures of specific deterrence - a measure based on past inspections, and a measure based on past fines. For specific inspection deterrence, I predict the probability of inspection for every well-date. I then calculate the running sum of predicted inspections, and divide the number of predicted inspections of the firm's wells over the firm's total number of wells. This yields a predicted rate of inspections during previous months. I use the predicted rate of inspections instead of the actual rate to avoid the well known problem of bias which arises in estimates of specific deterrence if regulators preferentially inspect likely violators (Helland 1998; Stafford 2002). I calculate predicted inspections based on the predicted values of the model in Column (1) of Table 5. 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. Construction of general deterrence signals is analagous, except that they are the actual rate of inspection and mean fine for other firms' wells in the same county as the inspected well.[14]

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 deterrence measures - regression coefficients will describe the impact of a one standard deviation change in deterrence.

Moving to experience, I construct two measures of a firm's experience at the time of an inspection. The first is the number of years since the firm's first inspection in the jurisdiction. As wells are generally inspected at the time of drilling, time since first inspection is a close proxy for actual time operating. The second experience measure is the count of inspections that a firm has experienced at the time of inspection, again at the jurisdictional level. This may provide a more precise measure of experience if inspections provide the opportunity to speak directly with inspections staff or because direct inspections are more salient than observation (Earnhart and Friesen 2013).

Experience variables are constructed at the firm level. This means that if a firm drills its

[14]For observations with no fines to date, the average fine level is undefined. I drop this small number of observations in my primary analyses. Appendix B includes no fines as zeros in the average calculation. Results are consistent with the core results.

first well in a district on Jan 1, 2006 and drills its second well on Jan 1, 2007, and the second well is inspected for the first time on Feb 1, 2007, then the firm's level of experience for that inspection (as measured in time) is 1 1/12 years. Similarly, if the first well had been inspected twice by Feb 1, 2007, then the firm's experience measure for the Feb 1, 2007 inspection (as measured by inspections) would be two. This assumes that knowledge resides at the firm level, so knowledge gained at one well can be applied at subsequent wells.

4.4 Summary Statistics

Table 2 provides summary statistics. Approximately 9 percent of observations find violations, and approximately 0.8 percent incur fines (approximately 9 percent of violations). Conditional on a fine occurring, the average is $44,384. Table 2 also describes the average number of inspections, wells, and years of experience that a firm has at the time of each inspection. Experience and wellcount are measured at the jurisdictional level, defining jurisdictions as PA DEP oil and gas districts.[15] Due to the long history of oil & gas production in Pennsylvania, only about 15.9% of inspections are of unconventional (fracking) wells. The remainder are of conventional wells.

Table 2: Summary Statistics

VARIABLES	(1) Mean	(2) Std Dev
Found Violation	0.088	0.283
Levied Fine	0.008	0.088
Fine	$44,384	$100,785
Firm inspections (000's)	0.714	0.943
Firm experience (years)	6.50	4.10
Firm well count (000's)	0.230	0.261
First inspection of well	0.338	0.473
Well inspection number	4.39	6.38
Unconventional Well	0.159	.366
Number of observations	164109	
Number of firms	1790	

This sector is not highly concentrated - the largest 10 firms are primarily regional operators such as Range Resources, National Fuel, Dominion, and Chesapeake, which between

[15] Section F provides summary measures at the county and state level.

them control only 27% of the wells. The Herfindahl index for firms (based on the number of wells they are observed operating in Pennsylvania) is 109, far below threshold of 1500 used to define concentrated markets (US DOJ 2010, p.19).

5 Results

5.1 Firms are Deterred by Inspections

This analysis estimates equation 9. The dependent variable y_{wort} is 1 if a violation was noted at an inspection of a well w operated by firm o in jurisdiction r at time t and 0 otherwise. Deterrence enters as a vector of measures Z_{rjkt}, discussed in detail in section 4.3. Indicator variables for type of inspection (routine, complaint, etc), whether or not the well is a fracking well, and whether or not it is the first inspection for the well comprise X_{wt}. This specification controls for unobservables at the firm, year, and month of year levels (ν_o, μ_y, and γ_m respectively). Additional specifications will replace ν_o with a well effect η_w.

This analysis estimates equation 9 as a linear probability model. Limited dependent variable models require assumptions about the distribution of the error term, whereas OLS is consistent without such an assumption, and in practice typically yields similar average treatment effects (Angrist and Pischke 2008, Sec 3.4.2). Appendix C provides evidence that the linear probability model is a reasonable fit to the data.

In Equation 9, the dependent variable y_{wort} describes the outcome of an inspection of well w of operator o in jurisdiction r at time t, and j and k index regulatory actions towards the operator and other firms, and inspections and fines respectively.

$$y_{wort} = \sum_{j \in [spec, gen]} \sum_{k \in [insp, fine]} \alpha_{jk} Z_{rjkt} + X_{wt}\beta + \nu_w + \mu_y + \gamma_m + \varepsilon_{wort} \tag{9}$$

Results for estimating equation 9 are presented in Table 3. The dependent variable has been transformed so that the coefficients describe the change the probability of detecting a violation from a 1-unit change in the independent variable. In Column (1), a one unit change in inspection frequency is associated with a 3.393 percentage point decrease in the probability of detecting a violation. We see across specifications that a one standard deviation increase in inspection treatment reduces the likelihood of a violation approximately 3 percentage points, on average. Recalling that the baseline probability of finding a violation is about 9 percent, a one standard deviation increase in inspection treatment decreases the probability of violations by approximately one third. This holds across specifications which control for specific deterrence, which include firm fixed effects, and which include well fixed effects.

We also see that a one standard deviation increase in specific deterrence is associated with about a 0.7-1 percentage point decrease in likelihood of a violation.

Table 3: Average Deterrence Effects

VARIABLES	(1)	(2)	(3)	(4)	(5)
General Inspection	-3.393***	-3.334***	-3.241***	-3.066***	-2.289***
	(0.981)	(0.988)	(0.989)	(0.987)	(0.338)
Own Inspection		-0.698**		-0.708**	-1.032***
		(0.343)		(0.310)	(0.184)
Observations	164109	164109	164109	164109	164109
R^2	0.109	0.109			
Within R^2			0.072	0.076	0.137
Intraclass Correlation			0.476	0.477	0.459
FE			Firm	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include additional control variables. Columns 1-4 cluster standard errors at the firm level. Column 5 clusters standard errors at the well level. "FE" refers to "Fixed Effect". LHS transformed to take on values of 0 or 100.*

5.2 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 varies with a proxy variable N_{ort} for the number of signals that firm o has received in jurisdiction r at time t. If a firm has wells in multiple jurisdictions, the N_{ort} describes the firm's experience in the jurisdiction of well w. In my core results, jurisdictions are measured at the PA DEP Oil & Gas District level. I estimate the linear probability model of Equation 10, which adds the interaction term and experience to Equation 9. All other variables are the same as in Equation 9.

$$y_{wort} = \sum_{j \in [spec,gen]} \sum_{k \in [insp,fine]} (\alpha_{jk} Z_{rjkt} + \psi_{jk} Z_{rjkt} N_{ort}) \\ + \rho N_{ort} + X_{wt}\beta + \nu_o + \mu_y + \gamma_m + \varepsilon_{wort} \tag{10}$$

Results are shown in Table 4.[16] Columns 1 and 2 measure N_{ort} as the number of inspections that a firm has received in a PA DEP district.[17] Columns 3 and 4 measure experience as the time that a firm has been operating in a PA DEP district. In all cases, we see that a completely inexperienced firm has a substantial and significant deterrence effect. In all specifications, an increase in experience is associated with a decrease in the magnitude of the deterrence effect. This effect is statistically significant. Using a one-tailed t-test, the effect of experience on deterrence is significant at the 1% level in specifications 1,2, and 4, with a p-value of 0.07 in specification 3.[18] It is also economically substantial - the effect of general deterrence for a very experienced firm is approximately half that of an inexperienced firm.

Using the results of Column 1, we see that a completely inexperienced firm has a deterrence effect of -4.247 percent - i.e., a one standard deviation increase in nearby monitoring effort is associated with a 4.247 percentage point reduction in the probability of a violation. A firm which has experienced the mean value of 714 inspections would have a deterrence effect of $-4.247 + 0.714 * 1.067$ or -3.5 percentage points, which is a 18 percent reduction relative to a new entrant and comparable to the average deterrence estimate of about 3 percentage points from Table 3. A firm at the 90th percentile of experience, which has encountered 1,873 inspections in the county, has a deterrence effect of -2.2 percentage points, a reduction of nearly half relative to a new entrant. Put somewhat differently, we see that the deterrence effect decreases by approximately 25% per thousand inspections (1.067/4.247).

[16]Coefficient estimates for other variables are shown in Appendix D. Appendix E presents a variety of alternative specifications, including the time scale of deterrence, time lags, and a logit model.

[17]Appendix F repeats the same analysis, but measuring experience at the county and state levels instead of the district level. Results are similar.

[18]A one-tailed test is appropriate for the effect of experience because Lemma 1 predicts that the sign will be positive. Therefore the appropriate null hypothesis is that the coefficient is less than or equal to 0. Under a more conservative two-tailed test, the specification in Column 3 is significant at the 10% level and all others are significant at the 1% level. All remaining specifications base p-values on one-tailed tests for the coefficient on the interaction of inspections and experience.

Table 4: Learning and General Deterrence

	Inspections		Time	
	(1)	(2)	(3)	(4)
General Inspection	-4.247***	-4.989***	-5.005***	-8.420***
	(0.930)	(0.516)	(1.631)	(0.873)
Gen Insp * Experience	1.067***	1.432***	0.151*	0.454***
	(0.232)	(0.129)	(0.098)	(0.059)
Observations	164109	164109	164109	164109
Within R^2	0.079	0.141	0.078	0.140
Intraclass correlation	0.479	0.471	0.479	0.473
FEs	Firm	Well	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. P-values for the interaction term are based on a one-tailed test. All specifications include additional control variables. All specifications cluster standard errors at the individual level. LHS transformed to take on values of 0 or 100.*

An alternative measure of experience is presented in Columns 3 and 4. Here experience N_{ort} is measured by the number of years that a firm has been operating. This measurement of experience may be more comparable to the measure n from the theoretical model in which a firm observes a signal about regulatory effort in each time step. Here we see a similar response, although the effect of experience is smaller than in Columns 1 and 2. In Column 3, an inspection of a firm with the average value of 6.49 years of experience has a deterrence effect of approximately -4.0 percentage points, a reduction of 20%. This is equivalent to the deterrence effect decreasing by approximately 3% per year. A firm at the 90th percentile of experience, which has been operating for 12 years, would have a deterrence effect of approximately 3.2 percentage points, more than a third lower than that of a new entrant.

To compare the magnitudes of the effects of learning from inspections to learning over time, we can multiply the learning coefficients by the standard deviation of the experience measures. With standard deviations of .943 (thousand) inspections and 4.10 years, we see that a one standard deviation increase in inspection-based experience is associated with a reduction in general deterrence of 1.01 (1.35), whereas a one-standard deviation increase in years is associated with a 0.619 (1.86) reduction for operator and well fixed effect models, respectively.[19] There is some spread, particularly between the different fixed effect models, but a one standard deviation increase in both measures of experience have broadly similar

[19]Calculations are $.943 * 1.067$, $0.943 * 1.432$, $4.10 * 0.151$, and $4.10 * .454$, respectively

effects.

Appendix F shows that results are similar if learning is measured either at the county level (to control for spatial variation) or at the state level (if regulatory stringency is similar across DEP jurisdictions.) Models which include both operator and well fixed effects (based on Correia (2016)) yield results nearly identical to those which include only well fixed effects, which is unsurprising as only 15% of wells are observed to change operators.

In sum, the effect of learning is economically substantial. For a new entrant with no experience, a one standard deviation increase in nearby inspections is associated with a reduction of violation likelihood of nearly one quarter[20] while a firm at the 90th percentile of experience has a general deterrence effect that is approximately half that of a new entrant.

6 Is Learning Socially Beneficial?

Firm learning may be beneficial if the regulator's stringency is efficient and the firm is adjusting effort by engaging in better environmental practices. Alternatively, learning may be harmful if the regulator's stringency is inefficient or if the firm is adjusting effort by engaging in regulatory avoidance. A complete answer is beyond the scope of this paper and is a promising topic for future research. However, I will try to provide some suggestive evidence. I do this in two ways: First I consider whether regulators preferentially inspect high-marginal-damage wells. Second I try to distinguish between operational improvements and strategic compliance.

The first analysis tells us that inspectors do generally seem to inspect more in areas with higher damages. This is consistent with efficiency, but not dispositive of it.

The second analysis finds that more experience is associated with a reduction in spills but has no effect on "paperwork" violations. This is consistent with at least some real improvement in environmental performance.

[20] $.943 * 1.067/4.247 \approx 0.25$

6.1 Where are regulators inspecting?

Testing whether or not regulators are inspecting efficiently would require knowledge of the marginal damages of violations and marginal costs of inspections, each of which are unobserved. Instead I follow the suggestive approach of Gray and Shadbegian (2004), who show that pulp and paper mills emissions decrease in settings where we might expect particularly high marginal damages of emissions: in areas with dense populations or where a larger fraction of the population is very young or very old. In this case, I will test whether inspection frequency increases with population, metro area status, surface water prevalence, or economic reliance on recreation. Higher populations and metro areas might be associated with larger marginal damages from environmental incidents simply because more people are exposed, while water and recreational activities might be associated with increased damages due to fears of water contamination. I also include measures of whether inspectors inspect other wells in the same county on the same day day (the number of inspections and an indicator for whether there are any) as measures of costs - if inspectors are nearby, then the travel cost of inspection is lower.

Table 5 presents results. We see that most of these factors are associated with a higher inspection frequency. Keeping in mind that the unconditional probability of inspection on a given single day is about 0.4%, these effects are substantial. Being in a metro area or adjacent county is associated with an increase in inspection probability of about 25-50%, and being in a recreation intensive non-metro county is associated with an increase in inspection probability of about 50-75%. Inspectors inspecting another well in the county is associated with an increase in the inspection probability, while wells in more densely populated areas are also more likely to be inspected. The density of surface water in a county does not seem to affect inspection likelihood, perhaps suggesting that regulators do not view surface water as particularly vulnerable (or that this is too imprecise a measure of water risk). Taken together, these results suggest that regulators are inspecting frequently in counties where environmental incidents would have higher marginal damages, a necessary but not sufficient condition for an efficient inspection regime.

These point estimates are generally qualitatively consistent across probit, logit, and linear probability model specifications. The major exception is that relative to the logit and probit, the LPM puts more emphasis on the number of (other) inspections in a day in a county and less on whether there are any other inspections.

Table 5: Regulator's choice of inspections

	(1) LPM	(2) Logit	(3) Probit
Num Insp In County	0.0016***	0.0002***	0.0003***
	(0.0002)	(0.0000)	(0.0000)
Any Insp In County	0.0029***	0.0104***	0.0085***
	(0.0011)	(0.0006)	(0.0006)
Pop Density (1000 People / Acre)	1.9269***	1.4154***	1.3978***
	(0.4934)	(0.3328)	(0.3373)
Metro Area	0.0024***	0.0012**	0.0013**
	(0.0006)	(0.0006)	(0.0007)
Metro Adjacent	0.0023***	0.0011*	0.0012**
	(0.0007)	(0.0006)	(0.0006)
Percent Water	0.0000	0.0000	0.0000
	(0.0000)	(0.0000)	(0.0000)
ERS Recreation County	0.0031***	0.0023***	0.0020**
	(0.0007)	(0.0008)	(0.0008)
Observations	43735178	2186759	2186759

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications cluster standard errors at the firm level. Logit and probit specifications report average marginal effects.*

6.2 Avoidance or Improvement?

Thus far, I have not addressed whether firms who increase their effort level are making real improvements in complying with regulations, or whether they are engaging in regulatory avoidance practices which reduce the probability of penalty without reducing the probability of incidents. There is evidence of avoidance in other settings, such as firms choosing a size just below a threshold which would incur greater regulatory stringency (Sneeringer and Key 2011; Almunia and Lopez-Rodriguez 2018).

I do not observe either compliance or avoidance effort directly. However, I can perform some suggestive tests. For some violations, I can identify whether they amounted to paperwork, or whether they constituted direct environmental hazards. I identify three different types of violations: whether the violation was for failing to post sufficient signage at the well site, whether the violation was for an insufficient erosion control plan, or whether the violation was for a spill. These types of violations sare all highly observable. Signage and

Table 6: The Impact of Experience on Violation Type

	(1) Spill	(2) Plan	(3) Signage
Experienced Firm	0.586**	0.810	1.136
	(0.143)	(0.175)	(0.151)
Probability of Violation	0.002	0.002	0.010
N	158004	162727	147477

Reported values are odds ratios. Standard errors in parentheses. Reported values are odds ratios. T-tests are based on a null hypothesis that the odds ratio is one.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

plans can be directly observed, and spills typically leave observable marks on the ground.

In this section, I test whether experienced firms are more or less likely to incur these potential violation types. Table 6 reports the reports of logit regressions for whether an inspection yields a signage violation, a sediment plan violation, or a spill violation on an indicator for whether an operator is in the top quarter of experience at the time of the inspection. Experience is measured as time operating in the district. Odds ratios are reported because the treatment variable is an indicator variable, and significance tests are based on the null hypothesis that the odds ratio equals one.

We see that experienced firms are 0.586 times as likely as inexperienced firms to have a spill, and we can reject the null hypothesis that this ratio is one ($p < 0.05$). This suggests that there may be real improvements in environmental performance with experience. In contrast, we cannot reject the null hypothesis that experienced firms and inexperienced firms are equally likely to incur plan or signage violations. Anecdotally, plan and signage violations are easy to comply with. A lack of compliance typically indicates pressure on corporate managers to meet internal deadlines related to commencing production even if plans or signage is not yet ready. In conjunction, these results suggest that at least some reductions in violations are due to real improvements, and that firms are considering both penalties and business needs in their effort decisions.

7 Sample Selection

One concern with a data set based on inspections is sample selection. If regulators preferentially inspect firms that they think have low ability to comply with regulations, ignoring the selected nature of the sample will yield positively biased estimates (Helland 1998; Stafford 2002). Intuitively, we can frame this as a selection problem by conceiving of the relevant population as all possible well-days. If regulators target inspections and I only observe well-days with inspections, then the sample of observed well-days will be selected based on an unobservable-to-the-econometrician likelihood of violation. The analyses of Section 5 address this by in essence instrumenting for specific deterrence. This section introduces an additional test of selection based on the Altonji, Elder, and Taber (2005)-Oster (2017) approach of modelling selection bias as an omitted variable problem. In this framework, we can think of each well-day's baseline level of ability and willingness to comply with regulations as an omitted variable. If the sampling process (regulators choosing which sites to visit on a given day) is based on the omitted variable, that would imply sample selection bias. This test finds that general deterrence results are much less vulnerable to omitted variable bias than specific deterrence.

Consider the selection problem as being an omitted variable problem. In particular, assume that the model of Equation 10 omits variables W.[21] These omitted variables might include the firm's ability to comply with regulation. If $Corr(W, Z) \neq 0$, then the coefficients in Section 5 will be biased. Altonji, Elder, and Taber (2005) provides a method of analysing the effect of omitted variable bias on coefficient estimates based on the effect of the omitted variables. This can be used either to correct for point estimates based on *a priori* knowledge of $Corr(W, Z)$, or can be used as a sensitivity analysis to calculate what $Corr(W, Z)$ would have to be to for the coefficient of interest to take on particular values (like zero).

The Altonji, Elder, and Taber (2005) model, extended by Oster (2017), decomposes the error term into omitted control variables W and pure stochasticity ε. It is assumed that observable variables provide some information about the unobserved W - in particular, it is assumed that the correlation between a single treatment variable of interest and the

[21] In this section I omit subscripts for simplicity. I also assume that omitted variables are normalized to have coefficients of one without loss of generality.

Table 7: Previous R^2 Estimates

Paper	Context	R^2
Earnhart (2004)	municipal wastewater treatment plants	0.166-0.5166
Gray and Shadbegian (2007)	manufacturing emissions	0.13-0.218
Muehlenbachs, Cohen, and Gerarden (2013)	offshore oil rig environment & safety	0.28-0.33$^+$
Muehlenbachs, Staubli, and Cohen (2016)	offshore oil rig safety	0.086-0.322

Notes: $^+$ *: Pseudo-*r^2 *reported*

unobserved W is proportional to the correlation between the treatment variable and other observed variables. This proportion is denoted δ.

In this section, I calculate the value of δ which would be required to for the true values of the coefficients estimated in Section 5 to be zero. Put differently, this value of δ would yield the estimated coefficients, given that the true value was zero.

The Oster (2017) model leads to somewhat more empirically based modeling than the Altonji, Elder, and Taber (2005) model, although the researcher does need to choose a plausible R^2_{max} to represent the R^2 of a hypothetical regression which included both observed and unobserved control variables. In this case, we would want to know all factors that predict firm performance which are within firm control, while omitting variables such as weather shocks which might stochastically alter compliance. I survey several papers which use firm environmental performance as a regressand and report R^2s of their regressions in Table 7. The Muehlenbachs, Cohen, and Gerarden (2013) pseudo-R^2 is specifically for non-weather-related incidents - that is, by construction those regressions omit a major source of stochasticity. As such I consider that a plausible upper bound, and then round up to let $R^2_{max} = 0.35$. This is larger than the values from studies in related contexts, although smaller than that in more highly controlled settings of municipal water treatment plants. One important limitation is that current econometric methods are based on estimating a single treatment effect, therefore I estimate δ separately for specific and general deterrence.

Table 8 presents estimates of δ. We see that if the true value of general deterrence were zero, selection on unobservables would have to be between 0.56 and 1.62 times as important as selection on observables. For specific deterrence, these ratios are 0.33 and 0.54. These results suggest that selection on unobservables is less important for general deterrence than for specific deterrence.

Table 8: Selection on Unobservables

	(1) Inspections	(2) Time
Specific Deterrence δ	0.54	0.33
General Deterrence δ	1.62	0.56

8 Conclusions and Directions of Future Research

In summary, this paper provides suggestive evidence for the common hypothesis that general deterrence arises from firms updating their beliefs about regulatory parameters. An implication of the firm learning model is that the magnitude of general deterrence would decline with experience. I test this hypothesis on a novel data set which allows me to identify intra-firm, intra-jurisdiction, 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.

Recent work has explored the importance of information in individuals' compliance decisions (Alm, Jackson, and McKee 2009; Rincke and Traxler 2011). This work documents the importance of information for firm behavior. This growing information literature 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. For example, transparent communication of what areas regulators prioritize would likely lead to fewer violations in those areas - but at the cost of more violations in low priority areas. However, an important limitation of this analysis is that it does not distinguish between changes in actual stringency versus changes in perceived stringency (Lochner 2007).

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' initial beliefs are strong (have low variance), then firms will be deterred from violating 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 in settings with experienced firms.

There are several limitations of this research. First, the oil and gas industry differs from other industries in important ways. It is characterized by homogeneous output, vertical contracting in the production process, and combines high fixed costs of production with low variable costs. Moreover, this study examines a single state in which regulatory activity is highly observable. Different industries or regulatory contexts might lead to substantially different learning dynamics. Second, the proposed model of learning makes strong assumptions. I assume that learning proceeds via Bayesian updating in which firms know the variance of received signals, observe regulatory signals perfectly, and in which the regulator is naive. Relaxing these assumptions may lead to a richer model of firm behavior. Moreover, the econometric tests cannot distinguish between different learning models in which firms update their beliefs in response to new information, but do so less when they are more experienced.

This paper leaves several areas for future research. First, the mechanism of firm learning is poorly understood. Levitt, List, and Syverson (2013) attempt to identify mechanisms of firm learning-by-doing in manufacturing by distinguishing between the effects of experience on workers, organizational structures, and physical capital. A similar examination of the locus of knowledge about compliance would aid understanding of firm compliance behavior. Second, determinants of the magnitudes of deterrence (and learning about regulators) are poorly understood. Finally, firms are subject to a variety of regulations by a variety of agencies - environmental rules, safety rules, financial rules, etc. If compliance effort is complementary (or substitutable) across contexts, then efficient enforcement mechanisms will depend on the actions of other agencies, and potentially be less (or more) stringent than would be predicted by ignoring contextual spillovers. Understanding the extent to which compliance spills over across contexts seems to be an important area of future research.

References

Alm, J., B. R. Jackson, and M. McKee (2009). Getting the word out: Enforcement information dissemination and compliance behavior. *Journal of Public Economics 93*(3?4), 392 – 402.

Almunia, M. and D. Lopez-Rodriguez (2018). Under the radar: The effects of monitoring firms on tax compliance. *American Economic Journal: Economic Policy 10*, 1–38.

Altonji, J. G., T. E. Elder, and C. R. Taber (2005). Selection on observed and unobserved variables: Assessing the effectiveness of catholic schools. *Journal of political economy 113*(1), 151–184.

Andreoni, J., B. Erard, and J. Feinstein (1998). Tax compliance. *Journal of economic literature 36*(2), 818–860.

Angrist, J. D. and J.-S. Pischke (2008). *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

Argote, L. and D. Epple (1990). Learning curves in manufacturing. *Science 247*(4945), 920–924.

Barrage, L. (2018). Regulatory dynamics in u.s. aviation safety: Economic determinants of faa behavior. Working Paper.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*.

Bianchi, F. (2012). Regime switches, agents' beliefs, and post-world war ii us macroeconomic dynamics. *Review of Economic Studies 80*(2), 463–490.

Blundell, W., G. Gowrisankaran, and A. Langer (2018, July). Escalation of scrutiny: The gains from dynamic enforcement of environmental regulations. Working Paper 24810, National Bureau of Economic Research.

Correia, S. (2016). Linear models with high-dimensional fixed effects: An efficient and feasible estimator. Technical report, Duke University. Working Paper.

Covert, T. R. (2015). Experiential and social learning in firms: the case of hydraulic fracturing in the bakken shale. Technical report, University of Chicago. Available at SSRN: https://ssrn.com/abstract=2481321.

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.

Eckert, H. (2004). Inspections, warnings, and compliance: the case of petroleum storage regulation. *Journal of Environmental Economics and Management 47*(2), 232–259.

Economides, M. J., A. D. Hill, C. Ehlig-Economides, and D. Zhu (2013). *Petroleum Production Systems* (2 ed.). Prentice Hall.

Eyer, J. (2018). The effect of firm size on fracking safety. *Resource and Energy Economics 53*, 101–113.

Feyrer, J., E. T. Mansur, and B. Sacerdote (2017). Geographic dispersion of economic shocks: Evidence from the fracking revolution. *The American Economic Review 107*(4), 1313–1334.

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

Gold, R. (2014). *The boom: How fracking ignited the American energy revolution and changed the world*. Simon and Schuster.

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).

Harrington, W. (1988). Enforcement leverage when penalties are restricted. *Journal of Public Economics 37*(1), 29–53.

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

Huang, J. and J. Rios (2016). Optimal tax mix with income tax non-compliance. *Journal of Public Economics 144*, 52–63.

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

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.

Langpap, C. (2007). Pollution abatement with limited enforcement power and citizen suits. *Journal of Regulatory Economics 31*(1), 57–81.

Levitt, C. J. (2009). *Learning through oil and gas exploration.* Ph. D. thesis, University of Iowa.

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

Levitt, S. D. (1998). Why do increased arrest rates appear to reduce crime: deterrence, incapacitation, or measurement error? *Economic inquiry 36*(3), 353–372.

Levitt, S. D., J. A. List, and C. Syverson (2013). Toward an understanding of learning by doing: Evidence from an automobile assembly plant. *Journal of Political Economy 121*(4), 643–681.

Lochner, L. (2007). Individual perceptions of the criminal justice system. *The American Economic Review 97*(1), 444–460.

Lynch, S. M. (2007). *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists*. Springer.

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.

Maniloff, P. and R. Mastromonaco (2017). The local employment impacts of fracking: A national study. *Resource and Energy Economics 49*, 62–85.

Muehlenbachs, L., M. A. Cohen, and T. Gerarden (2013). The impact of water depth on safety and environmental performance in offshore oil and gas production. *Energy Policy 55*, 699 – 705. Special section: Long Run Transitions to Sustainable Economic Structures in the European Union and Beyond.

Muehlenbachs, L., S. Staubli, and M. A. Cohen (2016). The impact of team inspections on enforcement and deterrence. *Journal of the Association of Environmental and Resource Economists 3*(1), 159–204.

Nagypál, É. (2007). Learning by doing vs. learning about match quality: Can we tell them apart? *The Review of Economic Studies 74*(2), 537–566.

Oster, E. (2017). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics 0*(0), 1–18.

Pakes, A. and R. Ericson (1998). Empirical implications of alternative models of firm dynamics. *Journal of Economic Theory 79*(1), 1 – 45.

Pástor, L. and P. Veronesi (2006). Was there a nasdaq bubble in the late 1990s? *Journal of Financial Economics 81*(1), 61–100.

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

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

Raimi, D. (2017). *The Fracking Debate: The Risks, Benefits, and Uncertainties of the Shale Revolution*. Columbia University Press.

Rincke, J. and C. Traxler (2011). Enforcement spillovers. *Review of Economics and Statistics 93*(4), 1224–1234.

Rubin, J. and C. Kling (1993). An emission saved is an emission earned: an empirical study of emission banking for light-duty vehicle manufacturers. *Journal of Environmental Economics and Management 25*(3), 257–274.

Sah, R. K. (1991). Social osmosis and patterns of crime. *Journal of political Economy 99*(6), 1272–1295.

Schwabe, K. A. (1999). The effects of separability on incentive-based instrument performance. *Economics Letters 63*(3), 377 – 380.

Segerson, K. and T. Tietenberg (1992). The structure of penalties in environmental enforcement: an economic analysis. *Journal of Environmental Economics and Management 23*(2), 179–200.

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

Sneeringer, S. and N. Key (2011). Effects of size-based environmental regulations: Evidence of regulatory avoidance. *American Journal of Agricultural Economics 93*(4), 1189–1211.

Stafford, S. L. (2002). The effect of punishment on firm compliance with hazardous waste regulations. *Journal of Environmental Economics and Management 44*(2), 290–308.

State of Pennsylvania. Chapter 78. oil and gas wells. `https://www.pacode.com/secure/data/025/chapter78/chap78toc.html`. [Accessed 04/12/2018].

State of Pennsylvania. The clean streams law. `http://www.dep.pa.gov/business/energy/oilandgasprograms/oilandgasmgmt/pages/laws,-regulations-and-guidelines.aspx`. [Accessed 04/12/2018].

State of Pennsylvania (2002). Civil penalty assessments in the oil and gas management program. Technical Report 550-4180-001, State of Pennsylvania.

US DOJ (2010). Horizontal merger guidelines. `https://www.justice.gov/sites/default/files/atr/legacy/2010/08/19/hmg-2010.pdf`. [Accessed 04/12/2018].

Weber, J. G. (2012). The effects of a natural gas boom on employment and income in colorado, texas, and wyoming. *Energy Economics 34*(5), 1580 – 1588.

Wright, T. P. (1936). Factors affecting the cost of airplanes. *Journal of the aeronautical sciences 3*(4), 122–128.

A Theoretical proofs

For concision, I omit firm subscripts i from the proofs below.

A.1 Proof of Lemma 1

Equation 7 (reprinted below) implies that $N \to \infty \implies \hat{\sigma}_N^2 \to 0$.

$$\hat{\sigma}_N^2 = \frac{1}{\frac{1}{\sigma_0^2} + \frac{N}{\sigma^2}}$$

Based on Equation 8 (reprinted below) $\hat{\sigma}_N^2 \to 0 \implies \Delta\hat{\theta}_{N+1} \to 0$.

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

Equation 5 implies that $\frac{\partial \rho^*}{\partial \hat{\theta}_N}$ is bounded, so, $\Delta\hat{\theta}_{N+1} \to 0 \implies \Delta\rho^* \to 0$.

Finally, $\Delta\rho^* \to 0 \implies \Delta P(\theta, \rho^*) \to 0$ because P_ρ is bounded.

A.2 Proof of Lemma 2

Assume that the firm's choice of effort has an interior solution. That is,

$$\arg\min_{\rho} \Big(C(\rho) + P(\hat{\theta}_N, \rho) R \Big) > \underline{\rho}$$

Also assume that the firm has exact knowledge of the regulator's effort level such that $\hat{\sigma}_N^2 = 0$. This can arise if either there is no initial uncertainty ($\hat{\sigma}_0^2 = 0$), perhaps because the regulator has credibly and exactly communicated its effort level. This case can also arise if $N = \infty$.

From Equation 8, $\lim_{\hat{\sigma}_N^2 \to 0} \Delta\hat{\theta}_{N+1} = 0$.

Equation 5 implies that $\frac{\partial \rho^*}{\partial \hat{\theta}_N}$ is bounded, so, $\Delta\hat{\theta}_{N+1} \to 0 \implies \Delta\rho^* \to 0$.

Finally, $\Delta\rho^* \to 0 \implies \Delta P(\theta, \rho^*) \to 0$ because P_ρ is bounded.

A.3 Proof of Corollary 1

First let us rewrite Equation 8 below.

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

Equation 7 (reprinted below) implies that $\hat{\sigma}_N^2 \rightarrow 0$ as $\sigma_0^2 \rightarrow 0$. This implies that $1 + \sigma^2/\hat{\sigma}_N^2 \rightarrow \infty$ as $\sigma_0^2 \rightarrow 0$.

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

Taken together, the above two findings imply that as $\sigma_0^2 \rightarrow 0$, $\Delta\hat{\theta}_{N+1} \rightarrow 0$.

Recalling Section 3.1, the firm responds to its change in belief $\Delta\hat{\theta}_{N+1}$ by updating its effort level. Using a differential approximation, we see that the smaller the change in $\hat{\theta}_{N+1}$, the smaller the change in ρ^*:

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

Of course, the probability of penalty is a function $P(\theta, \rho)$ of firm effort and the true value of regulator effort θ. This implies that a smaller initial uncertainty σ_0^2 will lead to a smaller change in the probability of penalty.

B Specifying the Fine Control Variable

In the core results, the average fine level is conditional on a fine occuring. The logic, following Earnhart and Segerson (2012), is that the fine measure is intended to capture the stringency conditional on a violation occuring. In this section, I instead calculate the average fine level including zero fines, on the logic that not yet observing a fine implies that fines are small or unlikely.

Results of reestimating the core model are presented in Table B.1. We see that estimates for our coefficients of interest - general deterrence and general deterrence interacted with experience are nearly equal to the core estimates in Table 4. It seems that the core results are robust to how we calculate the average fine.

We can also compare our coefficients on fines to those of our core specification, which are provided in Appendix D. These coefficients are generally small and somewhat inconsistent. One possibility is that typical fines just aren't a major deterrent. Industry representatitives tell me that they are more concerned with potential impacts to their reputation from formal Notices of (Alleged) Violations or accidents than from administrative fines. Eyer (2018) provides evidence of this, finding that operators with larger brand values are safer.

Table B.1: Learning and General Deterrence with Inclusive Fine Specification

	Inspections		Time	
	(1)	(2)	(3)	(4)
General Inspection	-4.547***	-6.327***	-5.683***	-9.695***
	(0.852)	(0.434)	(1.475)	(0.806)
Own Fine	0.657	0.983**	-0.452	0.808
	(1.189)	(0.472)	(0.927)	(0.547)
General Fine	0.504	-2.365***	0.429	-2.879***
	(0.434)	(0.308)	(0.620)	(0.407)
Gen Insp * Experience	1.098***	1.453***	0.190**	0.488***
	(0.214)	(0.112)	(0.0901)	(0.0562)
Own Fine * Experience	-3.680**	-3.393***	-0.221	-0.330***
	(1.567)	(0.600)	(0.170)	(0.0874)
Gen Fine * Experience	-0.149	0.656***	0.00108	0.177***
	(0.328)	(0.148)	(0.0862)	(0.0488)
Observations	183807	183807	183807	183807
FEs	Firm	Well	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. P-values for the interaction term are based on a one-tailed test. All specifications include additional control variables. All specifications cluster standard errors at the individual level. LHS transformed to take on values of 0 or 100.*

C Testing the Linear Probability Model

One key question in using a Linear Probability Model (LPM) for a binary outcome is how well the LPM actually predicts, or how well it fits the data. In this appendix, I provide several tests. We will see that the predicted probabilities of violations lie primarily between 0 and 1 and that the distribution of probabilities is substantially higher for inspections which actually resulted in violations.

First, I graph the density of predicted probabilities. Figure C.1 graphs the predicted probabilities of violation based on the results of Table 4. Recalling that only about 9 percent of inspections result in violations, we see several salient features. First, for inspections which did not result in a violation (the dashed line), the density of predicted values is tightly clustered near zero. Second, for inspections which did result in a violation (solid line), the probability distribution generally has higher values. Third, predicted probabilities for both distributions were primarily between zero and one.

Figure C.1: Density of Predicted Violation Probability

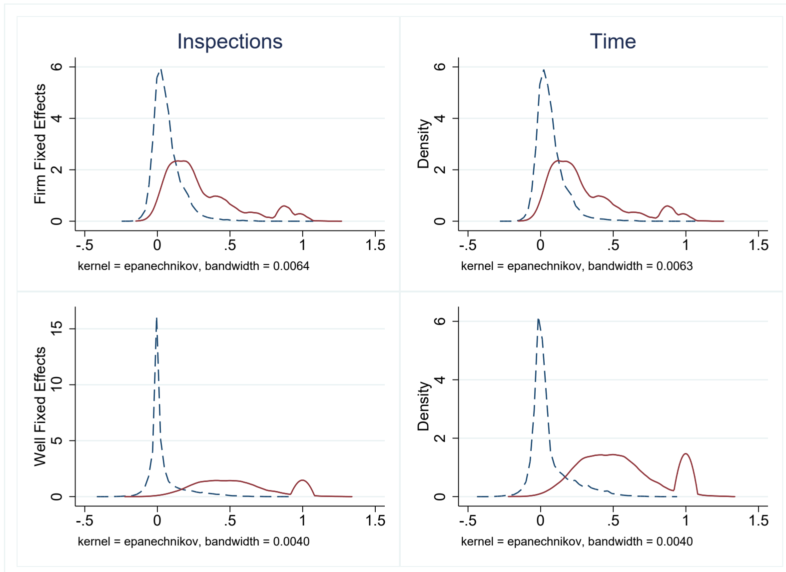

Dashed line describes the density of predicted violation probabilities for inspections which did not find violations. Solid line for inspections which did find violations. Graphs are smoothed kernel densities using rule-of-thumb bandwidths. The top row describes results with firm fixed effects, while the bottom row describes results with well fixed effects. The left column measures experience by number of inspections, while the right column measures experience by time.

Table C.1 provides several summary statistics for the distributions of predicted probabilities based on the models of Table 4. We see that the mean predicted violation probability is about 4-7% for inspections without violations and 31-55% for inspections with violations.

It bears noting that a relatively large fraction of predictions are outside of the range of 0 to 1 - a fifth to a third of predictions in most models. This is of course a weakness of linear probability models and a common motivation for discrete choice models. However, we see from inspecting these graphs that is primarily because the predicted values are centered near zero for inspections without violations. While negative probabilities are nonsensical, this does seem to be telling us that the model is accurately predicting near-zero probabilities of violation for most inspections which do not actually incur violations.

It also bears noting that there are relatively few predicted values above 0.5 - this is about

Table C.1: Summary Statistics of Predicted Violation Probabilities

		Inspections		Time	
		(1)	(2)	(3)	(4)
No Violation	Mean	0.07	0.04	0.07	0.04
	Median	0.05	0	0.05	0
	Fraction outside 0,1	0.20	0.30	0.20	0.30
	5th percentile	-0.05	-0.07	-0.05	-0.07
	95th percentile	0.24	0.32	0.25	0.32
Violation	Mean	0.31	0.55	0.31	0.55
	Median	0.23	0.52	0.23	0.52
	Fraction outside 0,1	0.03	0.02	0.03	0.02
	5th percentile	0.02	0.15	0.02	0.15
	95th percentile	0.88	1	0.88	1
FEs		Firm	Well	Firm	Well

Predicted probabilities based on the regression results of Table 4. The top panel describes the distribution of predicted probabilities for inspections with no violations. The bottom panel describes the distribution for inspections with violations.

the median for models with well fixed effects, but about the eightieth percentile for models with firm fixed effects. This suggests the model would have a relatively high false negative rate - if we used the common method of predicting a "1" outcome when the predicted probability was above 0.5, we would incorrectly predict no violation for 50-80 percent of inspections. However, given that the overall violation rate is about 9%, the mean predicted probabilities of 0.3-0.55 suggest that the model is reasonably predicting which inspections are more likely to find violations.

D Additional Regression Coefficients

Tables D.1 and D.2 list control variables for the core results of Table 4. I omit jurisdiction, firm, and inspection type fixed effects for brevity and anonymity. Table D.1 contains experience and deterrence effects, as well as a fracking indicator and well inspection count. We see some effect of fines in some specifications, although it is small and not precisely estimated across specifications. We see that fracking wells are more likely to be issued violations, and that violations are most likely at the initial inspections with violation probability decreasing with subsequent inspections. Table D.2 lists the year and month-of-year effects.

Table D.1: Control Variables for Table 4 Part I

	Inspections		Time	
	(1)	(2)	(3)	(4)
Site Inspection Number	-0.092***	-0.217***	-0.079**	-0.171***
	(0.031)	(0.035)	(0.032)	(0.032)
Own Inspection	-0.606*	-0.435**	-0.905**	-1.224***
	(0.324)	(0.178)	(0.455)	(0.268)
Own Fine	-0.023	-0.013*	0.025	-0.087***
	(0.018)	(0.008)	(0.037)	(0.023)
General Fine	0.858	-2.729***	0.978	-2.853***
	(0.589)	(0.449)	(0.820)	(0.559)
Specific Inspection * Experience	-0.047	0.017	0.042	0.080*
	(0.298)	(0.133)	(0.061)	(0.041)
Specific Fines * Experience	-0.021	-0.019**	-0.010*	0.009***
	(0.025)	(0.009)	(0.006)	(0.004)
General Fines * Experience	-0.460**	0.344**	-0.056	0.065
	(0.209)	(0.154)	(0.088)	(0.056)
Experience	0.015	2.432***	-0.097	0.331***
	(0.659)	(0.276)	(0.188)	(0.096)
Is Fracking Well	3.349***	0.000	3.504***	0.000
	(0.669)	(.)	(0.706)	(.)
Is Site First Inspection	0.871	4.194***	0.908	4.529***
	(0.604)	(0.242)	(0.590)	(0.239)
Site Inspection Number	-0.092***	-0.217***	-0.079**	-0.171***
	(0.031)	(0.035)	(0.032)	(0.032)
Constant	20.329***	23.292***	18.517***	16.230***
	(4.375)	(1.264)	(4.561)	(1.393)
FEs	Firm	Well	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. P-values for the interaction term are based on a one-tailed test. All specifications include additional control variables. All specifications cluster standard errors at the individual level. LHS transformed to take on values of 0 or 100.*

Table D.2: Control Variables for Table 4 Part II

	Inspections		Time	
	(1)	(2)	(3)	(4)
February	-1.105	-0.746*	-1.116	-0.710*
	(0.844)	(0.432)	(0.848)	(0.431)
March	1.578**	0.668	1.615**	0.895**
	(0.661)	(0.456)	(0.667)	(0.455)
April	2.150*	2.190***	2.246*	2.425***
	(1.154)	(0.476)	(1.149)	(0.474)
May	1.042	-0.156	1.246	0.213
	(0.946)	(0.457)	(0.945)	(0.455)
June	1.718	1.055**	1.850	1.321***
	(1.284)	(0.469)	(1.274)	(0.466)
July	-0.561	-0.512	-0.405	-0.165
	(1.077)	(0.465)	(1.088)	(0.465)
August	0.240	-0.480	0.432	-0.095
	(1.185)	(0.481)	(1.192)	(0.482)
September	-1.342	-1.867***	-1.172	-1.530***
	(0.964)	(0.460)	(0.987)	(0.462)
October	0.368	0.051	0.506	0.348
	(1.198)	(0.415)	(1.190)	(0.420)
November	-0.992	-1.286***	-0.742	-0.929**
	(1.057)	(0.431)	(1.066)	(0.437)
December	0.596	0.669	0.783	0.989**
	(1.027)	(0.457)	(1.041)	(0.464)
2001	1.459	2.371***	2.007	4.394***
	(2.010)	(0.862)	(2.048)	(0.881)
2002	-1.122	2.714***	-0.344	5.395***
	(2.573)	(0.873)	(2.675)	(0.911)
2003	3.281	6.505***	4.236	9.778***
	(3.053)	(0.951)	(3.171)	(1.035)
2004	2.004	1.277	3.191	5.444***
	(2.217)	(1.007)	(2.527)	(1.133)
2005	-0.209	0.132	1.156	4.866***
	(2.179)	(1.030)	(2.671)	(1.214)
2006	1.550	1.901*	3.160	7.229***
	(2.352)	(1.136)	(2.944)	(1.365)
2007	2.471	3.132**	4.201	8.791***
	(2.644)	(1.226)	(3.264)	(1.483)
2008	2.557	0.978	4.437	6.952***
	(2.544)	(1.221)	(3.322)	(1.569)
2009	5.432*	5.482***	7.528**	11.569***
	(2.879)	(1.273)	(3.659)	(1.693)
20010	2.769	2.122	5.035	8.431***
	(2.652)	(1.316)	(3.754)	(1.773)
2011	3.973	2.008	6.464	8.469***
	(2.771)	(1.370)	(3.992)	(1.853)
2012	1.297	0.193	4.003	6.746***
	(2.823)	(1.453)	(4.084)	(1.951)
2013	-0.381	-1.384	2.604	5.460***
	(2.994)	(1.532)	(4.362)	(2.053)
2014	-0.023	-1.887	3.107	5.149**
	(3.109)	(1.598)	(4.524)	(2.124)
FEs	Firm	Well	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. P-values for the interaction term are based on a one-tailed test. All specifications include additional control variables. All specifications cluster standard errors at the individual level. LHS transformed to take on values of 0 or 100.*

E Econometric Specification Robustness checks

This section describes a series of robustness checks to the primary econometric specification. These are decribed below, with results presented in Table E.1. Table E.1 flips the orientation of previous tables. Here, the columns present the regression coefficients for deterrence and learning (with standard errors below them), while each row describes a different model. Unless otherwise specified, each model includes firm-level fixed effects.

The first two analyses lag deterrence variables by 60 and 365 days. This recognizes that firms may not be able to instantaneously alter their level of care. These results are labeled "60 Day Lag" and "One Year Lag". We see that the results are very similar to the core results of Table 4.

Specifications which vary the time horizon of deterrence find smaller and less precisely estimated effect sizes than the primary specification, but are consistent with the primary results. The next two models repeat the primary analyses, but instead measure deterrence treatment as occurring over a limited time period. The treatment Z_{rjkt} describes the inspections over the previous 12 or 24 months, again scaled to a z-score.[22] These are described in Table E.1 as "One Year Info" and "Two Year Info", respectively. Results are similar to the core specifications, although smaller and less precise. Additionally, lagging the "Two Year Info" by 60 days also produces similar results (denoted "Two Year Info, 60 Day Lag"; other permutations are omitted for brevity but are similar).

Specifications with shorter time horizons of deterrence do consistently yield smaller point estimates than the core results. One possible explanation is that firms are learning about

[22] Previous analyses also included the average fine (conditional on fines occurring) 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.

long-run characteristics of inspection priorities, not short-run characteristics.

The next model estimates a logit instead of a linear probability model. Table E.1 presents results as average marginal effects (labeled "Logit"). We see that results are similar to the linear probability model, although somewhat smaller.

The "No Indiv Effect" model presents results omitting firm and well fixed effects. The results are generally similar to the core results, although the time-learning coefficient is smaller than in previous models.

Finally, I present a falsification test. Instead of learning, an alternative hypothesis could be that experience serves as a proxy variable for firm lobbying effectiveness or resource levels. In that case, it seems likely that firm size would also be associated with lobbing effectiveness or resources, reducing the impact of inspections. This section reestimates Equation 10, but replaces experience with the number of wells a firm has in Pennsyvania at the time of an inspection. We see that the average level of general deterrence is similar to previous models (about 3.6), but the impact of firm size is statistically indistinguishable from zero and small in magnitude (going from the 25th to 75th percentile in firm size is associated with a decrease in deterrence of approximately fifteen percent). Following Eyer (2018), this tests omits firm level fixed effects as there is insufficient within-firm variation in size to identify the size effect in a fixed effects model.

Table E.1: Regression Robustness Checks

Model	Experience Measure	General Inspection	Gen Insp * Experience
60 Day Lag	Inspections	-3.968***	0.922***
		(0.843)	(0.209)
	Time	-4.988***	0.127*
		(1.686)	(0.080)
One Year Lag	Inspections	-3.755***	1.049***
		(0.868)	(0.226)
	Time	-4.910***	0.143**
		(1.690)	(0.085)
One Year Info	Inspections	-0.961***	0.397***
		(0.306)	(0.134)
	Time	-1.419**	0.1050**
		(0.581)	(0.053)
Two Years Info	Inspections	-0.972***	0.414***
		(0.309)	(0.131)
	Time	-1.441***	0.107**
		(0.587)	(0.053)
Two Years Info, 60 Day Lag	Inspections	-0.822***	0.308**
		(0.341)	(0.130)
	Time	-1.315**	0.104**
		(0.623)	(0.055)
Logit	Inspections	-2.644***	0.853***
		(0.234)	(0.159)
	Time	-3.274***	0.059**
		(0.372)	(0.027)
No Indiv Effect	Inspections	-4.234***	0.912***
		(1.024)	(0.226)
	Time	-3.615***	0.030
		(1.236)	(0.063)
Firm Size	Wellcount	-3.366***	1.607
		(1.017)	(1.540)

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. All specifications include operator fixed effects and additional control variables unless otherwise noted. All specifications cluster errors at the firm level. LHS transformed to take on values of 0 or 100.*

F The Effect of Proximity and Spatial Scale

This section repeats the analysis of Section 5.2, but measures experience at the county level (Table F.2) and state level (Table F.3). This provides robustness against the hypothesis that knowledge is more localized (at the county level) or that enforcement is consistent statewide and thus statewide knowledge matters. Regression results are generally similar, with one notable exception: when experience is measured by inspections (ie columns 1 and 2), the coefficient on experience * general deterrence from inspections (measuring learning) decreases with spatial scale.

This seems to reflect the variation in experience at different spatial scales. We see in Table F.1 (which shows summary statistics) that on average, at an inspection, the firm had previously faced 318 inspections in the county and 1263 in the state, with standard deviations of 484, and 1880, respectively. As an example, based on Column 1 in Tables F.2 and F.3, from moving to no experience to an average level of experience is associated with a reduction in the general deterrence effect of about 0.5-0.7, while a one standard deviation increase in experience is associated with a decrease of about 0.7-0.9.[23] There is still some range here, but the variation between spatial scales is smaller than appears from comparing coefficients.

This pattern does not show up when measuring experience with time because we do not see the same spread in time operating across spatial scales. Again from Table F.1, on average, at an inspection, a firm had been operating for 5.152 years in the county and 6.518 years in the state - a spread of only about a quarter, far less than the spread of four times in inspection experience.

[23] Example calculation: $2.079 * 0.318 \approxeq 0.7$.

Table F.1: Summary Statistics for Experience at Different Spatial Scales

VARIABLES	(1) Mean	(2) Std Dev
Firm inspections in county (000's)	0.318	0.471
Firm inspections (000's)	1.263	1.831
Firm experience in county (years)	5.152	4.041
Firm experience (years)	6.518	4.202

Table F.2: Learning and General Deterrence: County Experience Measure

	Inspections		Time	
	(1)	(2)	(3)	(4)
General Inspection	-4.480***	-4.619***	-5.137***	-7.948***
	(0.928)	(0.465)	(1.429)	(0.799)
Gen Insp * Experience	2.079***	2.196***	0.198**	0.475***
	(0.456)	(0.193)	(0.083)	(0.054)
Observations	164109	164109	164109	164109
FEs	Firm	Well	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. P-values for the interaction term are based on a one-tailed test. All specifications include additional control variables. All specifications cluster standard errors at the individual level. LHS transformed to take on values of 0 or 100.*

Table F.3: Learning and General Deterrence: Statewide Experience Measure

	Inspections		Time	
	(1)	(2)	(3)	(4)
General Inspection	-3.946***	-5.554***	-5.335**	-9.689***
	(0.957)	(0.536)	(2.134)	(0.949)
Gen Insp * Experience	0.385***	0.693***	0.161	0.513***
	(0.123)	(0.063)	(0.131)	(0.062)
Observations	164109	164109	164109	164109
FEs	Firm	Well	Firm	Well

*Notes: *,**,*** indicate statistical significance at the 10%, 5%, and 1% level, respectively. P-values for the interaction term are based on a one-tailed test. All specifications include additional control variables. All specifications cluster standard errors at the individual level. LHS transformed to take on values of 0 or 100.*