

Measuring the Welfare Gains from Optimal Incentive Regulation*

Jose Miguel Abito[†]

July 22, 2019

Abstract

I empirically measure the welfare gains from optimal incentive regulation in the context of electric utilities facing both emissions and rate of return regulation (RORR). I provide evidence that RORR induces lower fuel efficiency, leading to greater coal consumption and higher emissions abatement costs. Replacing RORR with the optimal mechanism of Laffont and Tirole (1986) yields annual welfare gains of \$686 million or a 11% reduction in electricity prices. I construct a much simpler two-contract menu that can achieve more than 65% of these welfare gains.

1 Introduction

The theory of regulation has made important contributions to understanding the nature of regulatory problems and the design of optimal mechanisms (Laffont and Tirole, 1993; Laffont, 1994a; Armstrong and Sappington, 2007; Nobel Media AB, 2014). Central to the theory are informational constraints that limit the regulator’s ability to steer the firm towards the former’s objective (Laffont and Tirole, 1993). The optimal mechanism explicitly takes into account the regulated firm’s informational advantage and anticipates how it can strategically exploit it (Laffont, 1994a). Although much progress has been made in the characterization of the optimal mechanism, few empirical studies measure their actual benefits. First, since optimal mechanisms can be quite complex and require a high level of sophistication, they are rarely observed in practice hence making program

*This paper is based on the first chapter of my PhD Dissertation at Northwestern University and was previously circulated under the title “Welfare Gains from Optimal Pollution Regulation.” I would like to thank the editor, Aureo de Paula, and three anonymous referees for their time and valuable feedback. I would also like to thank my advisors Aviv Nevo, David Besanko and Robert Porter, and numerous colleagues who have helped me throughout this project. Data acquired from SNL Financial was funded by the TGS Graduate Research Grant and by the Center for the Study of Industrial Organization (CSIO) at Northwestern University, and the Dean’s Research Fund at Wharton.

[†]Department of Business Economics and Public Policy, Wharton School, University of Pennsylvania. Email: abito@wharton.upenn.edu.

evaluation infeasible. Second, a structural approach that evaluates the optimal mechanism as a counterfactual requires estimation of elements unobserved by the regulator and the econometrician. Identification in structural regulatory papers often *assumes* that the optimal mechanism is already implemented in the data, which then precludes any normative analysis.

I empirically measure the welfare gains from optimal incentive regulation using data from U.S. electric utilities that were subject to both price regulation known as *rate of return regulation* (henceforth, RORR), and environmental regulation in the form of the Acid Rain Program—a market-based sulfur dioxide (SO₂) emissions program that targeted the dirtiest coal plants. I provide evidence that RORR induced lower fuel efficiency, leading to greater coal consumption and higher SO₂ emissions abatement costs. I estimate a structural model of electricity production and emissions abatement, and simulate a counterfactual where RORR is replaced by Laffont and Tirole’s (1986) optimal mechanism.

The first contribution of the paper is to empirically quantify substantial welfare gains from the optimal mechanism. The second contribution is to show that a large fraction of these welfare gains can be achieved through a well-designed contract that is simple enough to implement.¹

The electricity industry is one of the key industries with a long history of economic regulation. In the U.S., electricity was mostly provided by vertically-integrated electric utilities that own and operate the generation, transmission and distribution of electricity.² These utilities were legal monopolists protected from essentially any threat of entry. In return for the right to operate as a monopolist, electricity prices were regulated by the state public utility commission.

Recent studies of the U.S. electricity industry provide empirical evidence of inefficiencies arising from RORR.³ To explain these inefficiencies, the “new economics of regulation” (Laffont, 1994a) emphasizes the role of asymmetric information in the regulator-firm interaction. In RORR, the

¹Rogerson (2003) computes a numerical example to show that a simple menu composed of a fixed price contract and full cost-reimbursement captures over 75% of welfare under the more complex optimal mechanism. Chu and Sappington (2007), and Bose, Pal, and Sappington (2011) investigate the performance of simple linear contracts using numerical examples as well. To the best of my knowledge, my paper is the first to empirically show the performance of simpler contracts. There is an active theory literature that investigates the efficacy and robustness of simpler contracts (e.g. Chassang, 2013; Garrett, 2014; Carroll, 2015; Dai and Toikka, 2017).

²My sample covers a period before the wave of electricity restructuring. For restructured states, utilities were required to divest the the generation side of their business with the rest of the supply chain still owned by the utility and regulated under RORR. There are currently 32 states that did not restructure. For these states, the generation side remains under RORR.

³Recent examples that compute the “average treatment effect” of RORR include Fabrizio et al. (2007) on non-fuel input costs, Davis and Wolfram (2012) on nuclear plant operations and maintenance, Cicala (2015) and Han et al. (2018) on coal procurement, and H. S. Chan et al. (2017) on fuel efficiency, procurement costs, and capacity factors of coal-fired plants. These papers essentially compare plants under RORR with similar plants in restructured states that compete in a wholesale market.

regulator needs to set regulated prices to cover the utility’s cost. However, cost is directly affected by the utility’s actions which are not perfectly monitored by the regulator. The inability to perfectly monitor the utility combined with the requirement that the utility has to be reimbursed for its cost creates an *agency problem*. The utility is no longer the residual claimant of cost-reducing effort and will not have the incentive to operate efficiently. Thus, asymmetric information and the resulting agency problem serve as important constraints that shape regulatory instruments and limit the ability of these instruments to achieve their intended outcomes.

Because of the electricity industry’s reliance on coal and natural gas, inefficient electricity generation not only creates pecuniary negative effects brought about by higher electricity prices, but also exacerbates the negative externality problem associated with burning fossil fuels (Cicala, 2015; H. S. Chan et al., 2017; Han et al., 2018). The externality problem cannot be overstated since the electricity industry is arguably the single largest source of harmful emissions. Therefore, it is necessary to understand the interaction between economic and emissions regulation in order to make sound policies.

Figure 1 illustrates the interaction between economic and emissions regulation. This example shows how the negative incentives from RORR lead to inefficiencies even with an otherwise efficient market-based mechanism (Buchanan, 1969). Suppose there is a single electric utility producing a fixed quantity of electricity q . Producing electricity generates emissions; assume that in producing q the firm chooses to abate a tons of SO₂ emissions. Under RORR, the total cost of producing the pair (q, a) is $\hat{C}(q, a)$ with marginal abatement cost denoted by $\hat{C}_a(q, a)$. $\hat{C}(q, a)$ maps (q, a) to some dollar amount and need not reflect cost minimization. If the firm faces a competitive electricity price, it will have an incentive to minimize cost and so the total cost of (q, a) will be given by the optimal $C^*(q, a)$. On the other hand, if RORR does not provide incentives to minimize cost, then it would be more expensive to produce (q, a) in total and possibly at the margin, i.e. $C_a^*(q, a) \leq \hat{C}_a(q, a)$. Consider a Pigouvian tax set equal to a constant marginal externality of emissions (or marginal benefit of abatement), τ . Facing the Pigouvian tax, the firm abates \hat{a} under RORR, which is lower than the efficient level of abatement a^* . Thus, distortions due to RORR lead to inefficient outcomes despite implementing an otherwise efficient market-based environmental mechanism.

Incentive regulation theory gives us guidance on how to think about regulatory problems like in the previous example, and how to design optimal regulatory mechanisms to handle these problems (Laffont and Tirole, 1993; Laffont 1994b). The key insight is that when provision of incentives is costly due to asymmetric information, it may be optimal to distort allocative efficiency to decrease information rents. The optimal form of regulation may involve abatement levels that do not equate marginal damages from emissions with marginal abatement costs, up to the point that

the first order gain from reducing information rents just balances the second order cost of inefficiency. Despite the simplicity of this insight, policies inspired by incentive regulation have rarely been implemented. The design and implementation of such mechanisms require significant effort and sophistication in terms of information gathering, rigorous auditing and analyses (Kahn, 1988; Joskow and Schmalensee, 1986; Joskow, 2008). Moreover, uncertainty over the actual benefits and costs of mechanisms, and the subsequent negative political and economic consequences in cases where such attempts are unsuccessful, make it difficult to convince policy-makers to adopt untested mechanisms. Therefore the goal of the paper is to quantify the theoretical gains from the optimal incentive mechanism, and see whether we can design more practical yet effective alternatives.

To quantify the gains from the optimal mechanism, I adopt Laffont and Tirole's (1986) model of monopoly regulation under asymmetric information. Cost is a function of a firm's *intrinsic* type and endogenous effort. Although a higher type means higher cost, firms can exert effort to reduce cost. Exerting effort is costly, and this cost is captured by the disutility of effort. In the model, the regulator observes cost but cannot separately identify type and effort: a firm with high cost may either be intrinsically inefficient or did not exert sufficient effort.

I use a novel identification strategy that exploits the variation in cost induced by the timing of RORR to separately identify type and effort. RORR is characterized by alternating periods when prices are set by the regulator (i.e. *rate case*), and when prices are essentially fixed until the next rate case (i.e. *regulatory lag*). Using panel data on electric utility operations, I find that, all else equal, firms' fuel efficiency drops by more than 4% during rate case years compared to non-rate case years. Lower fuel efficiency leads to higher marginal abatement cost since, when switching to lower sulfur coal, less efficient plants suffer a greater loss in electricity production even for a small reduction in SO₂ emission rates.

The rate case regulatory institution and the observed pattern in fuel efficiency and cost are consistent with the following story that I exploit for identification. During rate cases, costs are fully reimbursed and so the firm has an incentive to increase cost as much as possible. The firm optimally sets a *minimum* level of effort and this allows identification of the type distribution from observed costs during rate case years.⁴ Next, the firm becomes the residual claimant to cost-reducing effort during regulatory lag years and therefore has incentives to exert an *optimal* level of effort characterized by the equality of the firm's marginal benefit of exerting effort with its marginal disutility. Thus, I can use the implied level of effort from taking the difference between cost during the rate case and during the regulatory lag to invert the optimality condition and identify the effort

⁴The minimum level of effort can vary across firms, e.g. how strict the regulator is in auditing the firm. Since I normalize the minimum level of effort to zero for all firms, the estimated type absorbs the variation in minimum effort.

disutility function.

The shape of the type distribution is an important determinant of the complexity of the optimal mechanism and the size of the welfare gains. For example, if the type distribution is degenerate at a single type, then the optimal mechanism is just a fixed price contract equal to this type's cost under optimal (first best) effort. The type distribution I estimate exhibits significant heterogeneity and a particularly long tail. This observation alludes to potentially large welfare gains from allowing firms flexibility in terms of optimally choosing emissions and effort. At the same time, we might want to heavily distort effort downwards for extremely inefficient types to reduce information rents.

Using the estimated type distribution and effort disutility function, I simulate emissions under counterfactual regulatory contracts. To isolate the impact of RORR on an otherwise efficient market-based mechanism, I vary the form of economic regulation they face given a counterfactual Pigouvian tax on SO₂ emissions.⁵

I have the following results. First, assuming a constant marginal benefit from abatement equal to \$100 per ton of SO₂, I find that observed emissions are 32% higher compared to the counterfactual first best (i.e. competitive pricing and Pigouvian tax) level of emissions. Second, replacing RORR with the Laffont and Tirole (1986) optimal mechanism achieves welfare gains of \$686 million, which in magnitude is equivalent to a 11% reduction in electricity prices. Compared to the first best, emissions under the optimal mechanism are 5% higher. The optimal mechanism introduces some distortion in efficiency (higher emissions) in exchange for lower information rents. Although emissions are distorted upwards under the optimal mechanism, this level is far below the observed upward distortion in emissions from RORR. Third, since the optimal mechanism is likely infeasible to implement in practice,⁶ I construct a simple menu composed of two choices: a fixed price contract or full cost-reimbursement.

I find that the optimal simple menu sets the fixed price equal to the first best cost of the 25th percentile of the intrinsic cost distribution. Firms that have efficiency above the 25th percentile opt for the fixed price contract and exert the first best level of effort. On the other hand, firms that have

⁵The Acid Rain Program is implemented via a cap-and-trade instead of a Pigouvian tax. Although theoretically the equilibrium SO₂ emissions price under cap-and-trade should be equal to the Pigouvian tax, in practice these diverge. The cap set for the Acid Rain Program involved several concessions to achieve political consensus (Joskow and Schmalensee, 1998), hence the resulting equilibrium price is likely to be lower than the marginal externality. Another reason is that prices in the Acid Rain Program were volatile due to regulatory uncertainty and speculation about the future of the program (see for example the *Wall Street Journal* article "Cap and No More Trade", July 12, 2010). Finally, there are some concerns about the efficiency of the emissions permit market, although empirical evidence is mixed (e.g. Bailey (1998); Joskow, et al. (1998); Schmalensee et al. (1988); Ellerman et al. (2000); Ellerman and Montero (2001)).

⁶Laffont (1994b) shows that the optimal mechanism can be implemented via differentiated emissions taxes and transfers, the complexity of which depends on the degree of heterogeneity across regulated firms.

efficiency below the 25th percentile opt for full cost-reimbursement and exert zero effort. Similar to the optimal mechanism, incentives for low efficiency types are distorted to reduce information rents paid to types that are receiving the fixed price. This simple menu can achieve at least 65% of the welfare gains under the optimal mechanism. The ability of the simple menu to achieve a sizable fraction of the welfare gains is tied to the shape of the estimated type distribution. Simply “dropping” the long tail consisting of the most inefficient firms significantly reduces information rents and brings us closer to the optimal mechanism.

The paper contributes to three main literatures. First, I contribute to the structural empirical regulation literature pioneered by Wolak (1994). Wolak (1994) and Brocas et al. (2006) use the normative models of Baron and Myerson (1982) and Besanko (1985) to provide a link between observed behavior and the firm’s private information. This approach assumes that the actual regulatory institution can be modeled “as if” the optimal form of regulation was being implemented.⁷ In contrast to the normative approach, I exploit the actual regulatory institution in my identification strategy, similar to Gagnepain and Ivaldi (2002), and more recently, Lim and Yurukoglu (2018).

Second, I contribute to the growing empirical literature on environmental regulation. My work is most related to recent papers studying the interaction between market-based mechanisms, and existing market structure and regulatory institutions.⁸ Closest to the regulatory institutions in my paper is Fowlie’s (2010) analysis of the interaction between RORR and the NO_x budget program, which is a cap-and-trade program to regulate NO_x emissions. Comparing restructured states to those still under RORR, she finds that firms facing RORR tend to utilize compliance methods that are more capital-intensive. Unlike Fowlie (2010), I do not focus on the choice of abatement method since firms mainly switch to lower sulfur coal during my sample period. In fact, the issue of information asymmetry is unlikely to arise with respect to the choice of abatement method since the regulator observes coal prices and the capital cost of end-of-pipe equipment to remove SO₂ from the emissions stream; thus my focus is on agency problems arising from the difficulty of monitoring firms’ effort in managing fuel efficiency.

Third, I contribute to the empirical literature that examines the effects of economic regulation

⁷The optimal mechanism characterizes a mapping between the firm’s private information and observed regulatory variables (e.g. price and rate of return) which can then be inverted to identify and estimate the firm’s primitives (Perrigne and Vuong, 2011). Although Perrigne and Vuong (2011) allow observed regulatory variables to deviate from the ones specified by the optimal mechanism, this deviation has to be unsystematic, i.e. unrelated to the firm’s primitives.

⁸For example, looking at the U.S. cement industry, Fowlie, Reguant and Ryan (2016) find empirical support for Buchanan’s (1969) argument that only under conditions of perfect competition would a Pigouvian tax be “unambiguously hailed as welfare-improving.” In fact, their estimated welfare losses under a Pigouvian tax are substantial in this highly concentrated market: \$18 billion with a carbon tax of \$30 per ton.

on the electricity industry. Recent literature focuses on agency problems and attempts to estimate the effect on technical efficiency by comparing plants located in restructured states with plants in states that remained under RORR (e.g. Fabrizio et al., 2007; Davis and Wolfram, 2012; Cicala, 2015; H. S. Chan et al., 2017; Han et al., 2018). In contrast, I provide firm-level evidence of the impact of agency problems on fuel efficiency by looking at the variation in incentives induced by the timing of rate cases.

The paper is organized as follows. The next section provides a brief overview of SO₂ emissions regulation and RORR. Section 3 examines RORR both theoretically and empirically. I construct a three-period model of RORR with Laffont and Tirole's (1986) information structure to study how the timing of rate cases affect the utility's incentives. I then move to the data and show that variation in cost and fuel efficiency induced by the timing of rate cases is consistent with the theory. The section ends with a discussion of the assumptions needed to exploit this variation for identification of the firm's type and disutility of effort. Section 4 is devoted to identification and estimation while Section 5 uses the estimated parameters to conduct counterfactual welfare analysis. Finally, Section 6 concludes. For the benefit of the reader, Figure 2 shows the key steps involved in the estimation and counterfactual analysis.

2 Background

2.1 Pollution regulation

The electricity industry is arguably the single largest source of harmful emissions. In the U.S. for example, the industry accounts for 38% of CO₂ emissions (Environmental Protection Agency, 2014) and 65% of SO₂ emissions (Environmental Protection Agency, 2001). Because of its significant contribution to pollution, the industry has constantly been the prime target of environment regulations. I focus on the time period 1988-1999 which covers the design and first Phase (1995-1999) of the Acid Rain Program (ARP).

The ARP is the George H. W. Bush administration's answer to the demands for serious SO₂ emissions regulation of electric utilities, and is enforced by the Environmental Protection Agency's (EPA) under the Clean Air Act. ARP is a federal cap-and-trade program that establishes a market for SO₂ emission permits. While generally lauded as a success (G. Chan et al., 2012), the legislative history of ARP illustrates that implementation of the program largely hinged on the ability to provide concessions to states and their local economies (Joskow and Schmalensee, 1998; Ellerman et al, 2000 Ch 3; G. Chan et al., 2012; Schmalensee and Stavins, 2012). Concessions were in the form of initial permits directly given to affected electric utilities for free, with the hope of eventually passing the benefit of these free permits to consumers via lower electricity prices. Joskow and

Schmalensee (1998) estimate the value of these permits to be about \$600 million to \$1.8 billion.

Phase I of the ARP regulates the top 300 dirtiest coal-fired plants in the U.S. in terms of SO₂ emissions as designated in Table A of the statute. Although the electric utilities that own these Table A plants burn other types of fuel (or generate electricity from non-fossil fuel sources), these firms primarily rely on coal to produce electricity: the average ratio of coal consumption to total fuel burned is about 92%.

Coal contains sulfur and SO₂ is released to the atmosphere as a by-product of the fuel-burning process. Sulfur content ranges from about 0.2 pounds per heat input (lbs/MMBtu) to about 7 lbs/MMBtu (Perry et al., 1997). There is a tradeoff between heat and sulfur content: bituminous coal tends to have a higher heat content—hence a higher ability to generate electricity for the same mass of coal—but also high sulfur content compared to sub-bituminous coal. Hence, all else equal, plants tend to burn coal with higher sulfur content absent pollution regulation.

It is important to point out that distance of the plant from coal mines is another factor that determines coal choice since transportation costs are a significant component of delivered prices. The dirtiest plants in terms of SO₂ are those that are located far from sources of lower sulfur coal. Rail deregulation and falling delivered prices of sub-bituminous coal from the Powder River Basin (PRB) made this type of coal more competitive. However Ellerman et al. (2000, p. 89) note that although the competitiveness of PRB coal led to an overall decrease in contracted prices of coal, long-term contracts continued delivering high sulfur coal to Table A plants.

Two primary forms of reducing SO₂ emissions are fuel-switching and installation of a flue-gas desulfurization (FGD) unit, also known as a scrubber. Fuel-switching involves using coal with lower sulfur content or blending different types of coal with varying sulfur contents. In contrast, a plant can install an FGD which is an end-pipe control technology installed near the plant's emission stacks. The plant can still burn high sulfur coal, and the FGD will “scrub” SO₂ from the emissions stream. Compared to installing FGDs, fuel-switching was the more popular abatement method during Phase I. Although FGDs can remove almost 100% of SO₂ emissions, fuel-switching still accounted for 54% to 60% emissions reductions (Ellerman and Montero, 2007, Table 5). Plants with FGDs represent only 20% of all the plants during my sample period. Moreover, out of the 150 plants with scrubbers, only 15 plants installed scrubbers specifically in response to ARP. The rest of the plants installed scrubbers to satisfy SO₂ regulations that were in place well before the ARP.

For fuel-switching, the cost of reducing emissions is a function of the relative prices of coal with different sulfur contents, and also of fuel efficiency. If the coal-burning process is more efficient, then less coal is needed to produce the same amount of electricity. Consequently, the cost of changing to lower sulfur coal will be less with greater fuel efficiency, the reason being that lower sulfur coal contains less heat. In my sample, shifting to coal that complies with the ARP's implicit emissions

standard of 2.5 lbs per MMBtu leads to a “heat penalty” of 2%. That is, to produce the same amount of electricity, coal consumption has to increase by 2%. In this case, abatement cost is proportional to the amount of coal originally needed, which is higher the less efficient the firm is.

2.2 Rate of return regulation

Against the backdrop of SO₂ emissions regulation, electric utilities were also facing state-level price regulation in the form of *rate of return regulation* (RORR). During the sample period, the electric utilities were operating as vertically-integrated monopolists responsible for the generation, transmission and distribution of electricity for a service area in a given state.

In exchange for the right to be a monopolist, electricity prices are set by a state-level regulator known as the Public Utility Commission (PUC) based on information about cost and operations that the firm provides. Legally, the PUC is tasked to provide prices that reflect a “fair” rate of return on an electric utility’s invested capital and that would adequately recover the utility’s operating cost. As to what constitutes a fair rate of return is not explicitly specified and moreover, the PUC can question the “used and usefulness” of the utility’s investments and the prudence of the firm’s operations.

Different PUC objectives sway the generosity of approved rates of return and the stringency of prudence reviews (Lim and Yurukoglu, 2018). While understanding the objectives of PUCs is important, I take these objectives as given and focus on how utilities optimally react to two procedural features of RORR. The first is the actual price-setting process, which is conducted via a quasi-judicial proceeding called the *rate case*. The second feature is the *regulatory lag* that occurs in between two consecutive rate cases, where prices remain fixed.

The primary goal in the rate case is to set the revenue requirement, which is the total amount to be collected from consumers to compensate the firm for providing electricity. The revenue requirement is the sum of operating expenses and a return on the assets of the firm. The return is the product of the *rate base*, which is the value of the firm’s investment, and an allowed rate of return.

Since the regulator needs information to determine what revenue requirement to authorize, the rate case serves as a platform for the firm to provide information about its operating cost and environment. A hearing takes place where the firm and concerned parties (e.g. consumer interest groups) participate and provide testimony on the rationale of the proposed changes and the potential impacts these may have on consumer welfare. The firm, consumer groups, and the commission staff testify to support their position and to refute opposing arguments. A discovery phase also occurs where data and evidence are presented. If a settlement between concerned parties is not reached, the PUC commissioners decide on the case. The decision consists of the approved

revenue requirement which often differs from the initial proposal of the firm (difference is about 4% of proposed revenue requirement, on average).

In theory, the debate and disagreement in rate cases revolve around the three components of the revenue requirement: operating expenses, the rate base, and the rate of return. In practice, major rate cases focus on the determination of the rate base and the rate of return. Reported operating expenses are often passed through as long as these abide established accounting rules (Alt, 2006). In fact in my data, an often cited reason for initiating a rate case is to adjust electric utilities' return on equity, or to include a new plant (or a previously disallowed portion of some nuclear capacity) into the rate base. Moreover 80% of the amount the PUC disapproves are rate base and return on equity-related.

Once the revenue requirement has been determined, electricity rates are then set. Rates essentially remain fixed until the next rate case, except for adjustments triggered by significant changes in fuel prices. The price of the most relevant fuel (coal) during my sample remained flat so was unlikely to trigger a significant adjustment. Thus, at least to a first order approximation, electricity prices during regulatory lags were fixed.⁹

3 Rate of return regulation: theory and evidence

I model RORR as a three-period model of regulator-firm interaction. The key aspect of RORR that the model captures is the timing of rate cases. The model shows that RORR induces the utility to (i) decrease effort and raise the regulated electricity price set during the rate case, and then (ii) increase effort to some optimal level to maximize and extract the wedge between the approved regulatory price and the firm's operating cost during the regulatory lag. I then use panel data on cost, operations and rate case variables to show that variation in cost and fuel efficiency induced by the timing of rate cases is consistent with the theory. I end the section with a discussion of the assumptions needed to exploit this variation for identification.

3.1 Incentives during rate cases and regulatory lags

The main objective of a rate case is to set the revenue requirement which I denote as RR . It is the sum of operating cost, \tilde{C} , and the return on the rate base, \bar{R} . \bar{R} is a utility's profit over and above its operating costs, and is equal to the monetary value of the firm's capital (rate base) multiplied by an allowed rate of return. The model's focus is on how \tilde{C} influences \bar{R} and I do not separately model the components of \bar{R} , i.e. the rate base and the rate of return.

⁹I show the robustness of my analysis to other institutional features of RORR including fuel adjustment clauses in the online appendix.

I adopt the information structure in Laffont and Tirole's (1986) model of monopoly regulation under asymmetric information. I assume the regulator observes operating cost \tilde{C} . Operating cost is a function of an type θ that determines the firm's intrinsic cost efficiency, and effort $e \geq 0$ which reflects activities that can improve efficiency, both of which are not separately observed by the regulator. Formally, operating cost is given by $\tilde{C}(\theta) = \exp(\theta - e)C(q, s)$ where q is the quantity of electricity produced and s is a measure of SO₂ emissions. Cost is strictly increasing in θ and strictly decreasing in e . The firm's intrinsic type θ is a draw from some distribution \mathcal{F}_θ . The firm can reduce its cost by exerting effort. However, exerting effort is costly for the firm where this cost is captured by the disutility function $\psi(e)$. I assume $\psi(\cdot)$ is strictly increasing and strictly convex.

The goal is to examine how RORR affects a utility's incentive to reduce its cost. To do this, I consider a three-period model. In period $t = 1$, a rate case is held where \tilde{C} is observed and the regulator determines \bar{R} *potentially* using cost information. Once the rate case concludes, we enter period $t = 2$, which is the regulatory lag. Consistent with the idea of a regulatory lag, I assume that the regulator commits to setting the revenue requirement in period 2 to be equal to the revenue requirement from the preceding rate case (i.e. period 1): $RR_2 = RR_1 = \tilde{C}_1 + \bar{R}_1(\tilde{C}_1)$. Finally, period $t = 3$ is a new rate case and the model ends once this case concludes. I allow the return in period 3 to depend on current and past cost, i.e. $\bar{R}_3(\tilde{C}_1, \tilde{C}_2, \tilde{C}_3)$. Thus in period 3, $RR_3 = \tilde{C}_3 + \bar{R}_3(\tilde{C}_1, \tilde{C}_2, \tilde{C}_3)$.

The relationship between \bar{R}_t and $\tilde{C}_{t'}$ determines a utility's incentive to exert effort at each time t, t' . Suppose that the regulator completely ignores \tilde{C}_1 in determining both current and future \bar{R}_t , say because the regulator only uses information directly pertaining to the rate base or return on equity. Then exerting effort to lower \tilde{C}_1 will only decrease RR_t without any additional benefit to the utility. Therefore it will be optimal for the utility to exert zero effort in this case.

The following proposition summarizes conditions such that we would expect the utility to exert (i) zero effort during rate cases, and (ii) optimal "first best" effort during regulatory lags.¹⁰

Proposition 1 *If $\frac{\partial \bar{R}_1(\tilde{C}_1)}{\partial \tilde{C}_1}$, $\frac{\partial \bar{R}_3(\tilde{C}_1, \tilde{C}_2, \tilde{C}_3)}{\partial \tilde{C}_1}$, $\frac{\partial \bar{R}_3(\tilde{C}_1, \tilde{C}_2, \tilde{C}_3)}{\partial \tilde{C}_3} \geq 0$, and $\frac{\partial \bar{R}_3(\tilde{C}_1, \tilde{C}_2, \tilde{C}_3)}{\partial \tilde{C}_2} = 0$, then (i) $e_1^* = 0$, (ii) $e_2^* > 0$ and solves $\exp(\theta - e_2^*)C(q, s) = \psi'(e_2^*)$, and (iii) $e_3^* = 0$.*

The conditions in Proposition 1 reflect the non-accountability of costs in periods 1 and 3, consistent with the idea that costs are fully passed through to electricity prices during rate cases (i.e. the notion of cost-of-service regulation). This means that at the margin, the benefit of exerting effort to reduce costs during periods 1 and 3 is zero. Since exerting positive effort is always costly at the margin, it will be optimal for the firm to exert zero effort in those periods. However, incentives are different in period 2. Since now cost in period 2 affects period 2 profits because the utility is

¹⁰The online appendix contains the proof.

the residual claimant of cost-reducing effort by virtue of \bar{R}_2 being fixed, it is worthwhile for the firm to exert positive effort at the margin even if it incurs some disutility in doing so. Specifically, the utility will set effort such that the marginal benefit of exerting effort is equal to its marginal disutility.

3.2 Evidence on rate case timing and fuel efficiency

I now turn to the data to see how rate case timing is related to cost and fuel efficiency. My sample consists of electric utilities that have coal plants covered by Phase I of the EPA's Acid Rain Program. I use operations and fuel data from the Energy Information Administration's (EIA) Form 767 and the Federal Energy Regulatory Commissions' (FERC) Form 423. I combine these with financial and rate case data from SNL Financial and the Regulatory Research Associates (RRA). Details on data construction are in the online appendix.

Table 1 contains operations and cost statistics for my sample. The number of firm-year observations is 351. The operating and maintenance (O&M) variable cost measure is the sum of fuel and non-fuel O&M expense related to electricity generation. Fuel expense accounts for about 75% of O&M expense, on average. Moreover, on average, coal accounts for about 93% of total fuel consumption while about 5% and 3% are oil and natural gas respectively. Average O&M cost per kilowatt-hour of net generation is about 1.8 cents. Fuel efficiency is inversely measured by heat rate, which is the amount of fuel burned per unit of output. Firms that have *higher* heat rates are *less* fuel efficient. The mean heat rate is 10.67. In terms of efficiency, this heat rate corresponds to about $3.412/10.67 = 32\%$, which is typical for coal plants during this time period.

Table 2 contains rate case statistics for the utilities. On average, a rate case lasts just over a year but can extend for 3 years. The average number of years from the time a rate case is authorized to when a new rate case is proposed (i.e. the regulatory lag) is 2 but can be as long as 6 years. Majority of utilities in my sample experienced no more than 3 rate cases during 1988-1999.

Recall that the return on the rate base (RRB) is the return that the utility gets from its investment, net of operating cost. It is the product of the monetary value of its assets (rate base) and a rate of return (ROR). The firm proposes the rate base and ROR at the beginning of the case, and the regulator authorizes a rate base and ROR at the end. Although in theory, one can go over the individual rate case reports and pinpoint which are the specific expense categories that were not allowed by the regulator, I construct a summary measure instead. Specifically, I define a *disallowance* as the difference between what the firm proposes and what is eventually authorized by the regulator. A disallowance arises either because a specific expense was disallowed, or there was a disagreement on the monetary value of a given expense. On average, about 4% of the proposed revenue requirement is disallowed. This disallowance ranges from no disallowance to 12%. The

Table 1: Summary statistics of operations and costs data (electric utility-level)

Variable		Mean	Std dev	Min	Max
O&M var cost	\$M	328	253	23	1198
Net generation	MWh	2.4×10^7	2.2×10^7	558739	9.9×10^7
Ave O&M cost / net gen	cents / kwh	1.8	0.9	0.5	5.4
Heat rate	MMBtu/MWh	10.67	1.30	6.60	19.63
SO ₂ Emission rate	lbs/MMBtu	1.82	1.04	0.38	7.22
Capacity	MW	5010	4897	232	23227
FGD dummy	FGD dummy	0.33	0.47	0	1
Salary	\$000/emp/mo	15.8	8.2	.6	52.3
Price coal	\$/ton	33.11	9.93	12.48	53.80
Price oil	\$/barrel	22.53	4.94	10.06	37.38
Price gas	\$/MMBtu	2.96	0.97	1.38	15.48

Notes: Capacity is nameplate capacity which is the amount of electricity (in MW) the plant produces at 100% load.

average RRB disallowance, i.e. the difference between proposed and authorized RRB, as measured as a percentage of proposed RRB is 7%, and ranges from no disallowance to as high as 31%. The average RRB disallowance as expressed as a percentage of the total disallowance in the revenue requirement is 80%, so most disallowances are due to disagreements on the rate base rather than on operating costs.

A necessary implication of Proposition 1 is that, all else equal, we will observe higher cost during rate cases compared to during the regulatory lag. To check whether this holds in the data, I regress the log of O&M variable costs on the log of output (electricity and emissions), input prices (labor, coal, oil and gas) and capital (nameplate and an indicator for FGD), together with indicator variables for whether the observation comes from years when the rate case is ongoing. I construct three indicator variables. The first dummy is equal to one if the observation occurs during the rate case, i.e. from proposed to authorized year, inclusive. The second dummy is equal to one if the observation occurs in the year immediately after the authorization year. Finally the third dummy is equal to one if neither of the two dummies are one. In the regression, the omitted category is the second one so coefficients measure the % difference relative to the year after the rate case concludes. Table 3 contains the regression results.¹¹

¹¹I include specifications where I use state-level electricity demand as an instrument for electricity output, and regional prices for low and high sulfur coal as instruments for emission rates. Low sulfur coal is defined as coal with sulfur content below 1.2 lbs/MMBtu, while high sulfur coal is defined as coal with sulfur content above 3 lbs/MMBtu. First stage F-statistics are 164 and 27 for electricity output and emission rates respectively. The first stage regression

Table 2: Summary statistics of rate case data

Variable		Mean	Std dev	Min	Max
Rate case duration	Years	1.2	0.7	0	3
Regulatory lag	Years	2.3	1.9	0	6
Percent disallow Rev Req	% of Prop Rev Req	4	3	0	12
Proposed RRB	\$M	320	406	7	1868
Authorized RRB	\$M	296	374	7	1617
Percent disallow RRB	% of Prop RRB	7	4	0	31
Proposed rate base	\$M	2580	3272	73	15963
Authorized rate base	\$M	2453	3090	66	14485
Proposed ROR	%	10.2	0.9	7.9	12.2
Authorized ROR	%	9.8	1.0	7.4	11.8

Focusing on the estimates for the rate case dummy, we see that costs are 4% to 6% higher during a rate case compared to the year after, when the regulatory lag starts. Moving to the “neither” dummy coefficient estimate, I find no statistically significant difference in cost among non-rate case years, i.e. regulatory lag years. These results hold even when looking at within firm and within firm-rate case variation in cost. Including firm-rate case fixed effects in the regression allows us to measure the effect of rate case timing on cost for a given firm and a particular *rate case event*.¹²

The observed pattern in cost can be explained by a firm strategically initiating a rate case and locking-in high prices when it knows that costs will be (temporarily) high. Thus, this pattern can be consistent with a story about asymmetric information on *exogenous* fluctuations on cost, and not necessarily an unobserved *endogenous* “effort” story.¹³ More importantly, I need to find a link

coefficients are as expected. Specifically, I find that state demand and electricity output are positively related. In terms of coal prices, I find that while emissions rate is both positively related with the prices of both low and high sulfur coal, the coefficient on low sulfur coal is significantly larger than the coefficient on high sulfur coal.

¹²Firm-rate case fixed effects refer to firm-rate case events. Consider the following example. Suppose we have annual data on PECO’s cost and operations from 1990 until 1999, and that PECO had a rate case in 1991 and then in 1996. I tag PECO’s observations from 1991 to 1995 as “PECO 1991 rate case” (i.e. 5 years worth of observations), and observations from 1996 to 1999 as “PECO 1996 rate case.” For “PECO 1991 rate case,” rate case cost is cost in 1991 while non-rate case cost are costs from 1992, 1993, 1994 and 1995.

¹³A related alternative story to strategic initiation of rate cases is a story about cost-padding. In this case, a firm can exert effort to artificially increase its observed cost during the rate case. Laffont and Tirole (1992) extends Laffont and Tirole (1986) to allow for cost-padding and explore issues such as auditing and collusion. In their extension, *net effort* is equal to cost-reducing effort less cost-padding effort. Ultimately what is important in my empirical exercise is that economic regulation (i.e. the way the regulator compensates the regulated utility) affects fuel efficiency and operating cost endogenously. Thus in the cost-padding case, what is important is net effort and not its individual

Table 3: Regression results: O&M variable cost and rate case dummies.

log O&M var cost	(1)	(2)	(3)	(4)	(5)
Rate case	0.037 (0.030)	0.054** (0.024)	0.056*** (0.019)	0.052** (0.020)	0.052** (0.020)
Neither Rate case nor Year after	-0.001 (0.067)	-0.006 (0.018)	-0.024 (0.025)	-0.023 (0.027)	-0.023 (0.027)
Year FE	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	No	No
Firm-Rate Case FE	No	No	Yes	Yes	Yes
IV for electricity	No	No	No	Yes	Yes
IV for emission rate	No	No	No	No	Yes
R ²	0.89	0.98	0.99	0.99	0.99
Num. Obs.	351	351	351	293	293

Notes: Standard errors are either clustered at the firm level. Regression via OLS except when indicated. Additional regressors are a dummy for FGD; the logs of electricity output, emission rate, input prices (labor, coal, oil and gas), and nameplate capacity. I use log of state electricity demand as an IV for electricity output and regional prices for low (< 1.2 lbs/MMBtu) and high (> 3 lbs/MMBtu) sulfur coal for emission rates. Significance level: * 10%, ** 5%, *** 1%.

between RORR and SO₂ emissions. SO₂ is a byproduct of burning coal to produce electricity. Conditional on coal quality (heat, ash and sulfur content), a more efficient fuel-burning process leads to lower SO₂ emissions. Thus to examine the link between RORR and SO₂ emissions, I check whether the same pattern in costs we found earlier arises for heat rates.

Higher heat rates mean less efficient production since the firm burns more coal to produce the same amount of electricity. Short-run variations in heat rates are more likely due to temporary changes in how a firm operates its plants, or how the load is distributed across its plants.¹⁴ I regress the log of heat rate on the log of electricity generated, log of capital, an indicator for FGD, and the rate case dummies. Table 4 contains results of this regression.

Heat rates are about 4% to 6% higher during rate cases relative to the year after. Hence, during

components.

¹⁴Although it is not important whether the mechanism is reduction of efficiency within plants or re-allocation of output across plants for my empirical exercise, I provide some suggestive evidence for the latter in the online appendix.

Table 4: Regression results: Heat rates and rate case dummies

log heat rate	(1)	(2)	(3)	(4)
Rate case	0.065*** (0.024)	0.051** (0.020)	0.046* (0.026)	0.042** (0.020)
Neither Rate case nor Year after	0.003 (0.020)	-0.014 (0.035)	-0.016 (0.041)	-0.012 (0.033)
Year FE	Yes	Yes	Yes	Yes
Firm FE	Yes	No	Yes	No
Firm-Rate Case FE	No	Yes	No	Yes
IV for electricity	No	No	Yes	Yes
R ²	0.92	0.96	0.89	0.95
Num. Obs.	351	351	293	293

Notes: Standard errors are clustered either at firm level. Regression via OLS except when indicated. Additional regressors include a dummy for FGD and the logs of electricity output and nameplate capacity. I use log of state electricity demand as an IV for electricity output. Significance level: * 10%, ** 5%, *** 1%.

rate cases, a utility burns as much as 6% more coal to produce the same quantity of electricity. Since the quality of coal delivered does not vary with the timing of the rate case, higher heat rates directly imply more SO₂ emissions during rate cases. Finally, I do not find statistically significant differences in heat rates when I compare the year after the rate case and the succeeding non-rate case years. Thus the decrease in fuel efficiency during the rate case is only temporary and disappears immediately during the regulatory lag.

3.3 Interpreting low and high incentives: Assumptions on effort

The goal of the empirical exercise is to measure the welfare gains from the optimal mechanism. In order to measure these gains, I need to predict how firms will behave when facing a counterfactual regulatory regime. Firm behavior is driven by type θ and the disutility of exerting effort $\psi(\cdot)$, which I estimate from the data.

The previous subsection shows that during rate cases, utilities tend to systematically burn more coal per unit of electricity generated, and hence operating cost during rate cases are higher relative to cost during the regulatory lag. This pattern in cost induced by RORR is a useful source of

variation that provides snapshots of how a utility behaves when it faces low and high incentives. However, we need to be able to understand how “low” or “high” incentives map to our model to back out θ and $\psi(\cdot)$.

Let ω_{RC} be a given utility’s cost during a rate case after conditioning out observables such as electricity generated, emissions and fuel prices. Similarly let ω_{RL} be the cost during the regulatory lag. Following Laffont and Tirole (1986), we have $\omega_{RC} = \theta_0 - e_{RC}$ and $\omega_{RL} = \theta_0 - e_{RL}$ where θ_0 is the given firm’s true intrinsic cost type, and e_{RC} and e_{RL} are the effort levels exerted during the rate case and regulatory lag respectively. Since $\omega_{RC} > \omega_{RL}$, we have $e_{RC} < e_{RL}$.

If the conditions of Proposition 1 where in fact true in the data, then we have $\omega_{RC} = \theta_0$ and e_{RL} is equal to type θ_0 ’s first best level of effort $e^{FB}(\theta_0)$, where $e^{FB}(\theta_0)$ solves $MB(e^{FB}(\theta_0)) = \psi'(e^{FB}(\theta_0))$ and $MB(\cdot)$ denotes the marginal benefit from effort. In this case, “low” incentives basically map to zero effort while “high” incentives map to first best effort. We can then use these relationships for identification: (i) the distribution of costs during rate cases (after conditioning on observables) identifies the distribution of θ_0 ’s, and (ii) one can derive the implied first best level of effort by looking at the difference in a utility’s cost between rate cases and regulatory lags, and use this implied level of effort to identify the marginal disutility function.

The online appendix provide suggestive evidence consistent with the conditions of Proposition 1. Data limitations preclude more convincing evidence. Moreover, even if we can find strong empirical evidence, these conditions likely only hold *approximately* in practice, and it is useful to assess how results will change if such strong conditions do not obtain. Therefore in what follows, I explicitly *assume* Proposition 1 to aid identification, but show that these assumptions are, in a sense, innocuous in that the welfare gains from the optimal incentive mechanism that I compute is actually a *lower bound* of the true welfare gains.

I make the following two assumptions:

Assumption 1 *Effort during the rate case is zero.*

Assumption 2 *Effort during the regulatory lag equates the marginal benefit of effort with its marginal disutility.*

Assumption 1 is a normalization since type θ is only identified up to location and scale (Perrigne and Vuong, 2011). The question is how this normalization affects the interpretation of my results. Suppose $e_{RC} > 0$. For example, it is possible that the regulator can monitor fuel efficiency during rate cases, albeit imperfectly, which then means the utility exerts some minimum positive level of effort. Denoting my type estimate as $\hat{\theta}$, I will then have $\hat{\theta} \equiv \theta_0 - e_{RC} < \theta_0$. By assuming effort is zero during rate cases, the estimate of θ_0 will be biased downwards, i.e. the distribution of

utilities will be more efficient than the true distribution. In using the distribution of $\hat{\theta}$ to compute welfare gains, I essentially treat factors that induce a positive level of effort during the rate case (e.g. imperfect monitoring of fuel efficiency) as “free” ways for the regulator to incentivize the utility. As long as these “free incentives” exist and do not differ in their effect under RORR and the optimal mechanism, then the assumption is just pure normalization.

I now turn to Assumption 2. Since the utility is still technically under economic regulation during the regulatory lag, it may not fully receive the benefits of cost reductions from fuel efficiency. Moreover, there may also be some risk that observed decreases in cost during regulatory lags will be passed on to consumers through lower electricity prices. These factors dampen the utility’s incentive to exert effort during the regulatory lag.

The marginal benefit from effort is just the cost reduction from a small change in effort, i.e. $\left| \partial \tilde{C} / \partial e \right| = \exp(\theta - e) C(q, s)$. Note that this marginal benefit is strictly decreasing in e . Next, let $\psi_0(\cdot)$ be the true effort disutility function with marginal disutility given by $\psi'_0(\cdot)$. If $e^{FB}(\theta)$ is type θ ’s first best level of effort, then the true marginal disutility function is identified through the equation $\exp(\theta - e^{FB}(\theta)) C(q, s) = \psi'_0(e^{FB}(\theta))$.

Suppose for each type θ , the estimated effort is below the first best, i.e. $e(\theta) < e^{FB}(\theta)$. The level of effort is not of interest *per se* but is just used to identify the disutility function. Specifically, from the first order condition of effort, a lower estimated effort implies a higher marginal disutility, holding the marginal benefit curve fixed. Therefore, the marginal disutility function I will estimate lies *above* (or to the left of) the true marginal disutility function as in Figure 3.

Effort *under the optimal mechanism* equates the marginal disutility with a *distorted* marginal benefit of effort, the idea being that effort is distorted below first best levels to reduce information rents. The online appendix shows that the distortion is an increasing function of the difference in marginal disutilities between two adjacent types, i.e. $\psi'(e) - \psi'(e - \Delta\theta)$. Let $\Psi(\cdot)$ be the distortion in marginal benefit that I estimate and let $\Psi_0(\cdot)$ be the true distortion. If $\Psi(e) \geq \Psi_0(e)$ for all e , then, as Figure 3 illustrates, the estimate for effort under the optimal mechanism e^* is a lower bound for the true level of effort e_0^* .

To find conditions leading to $\Psi(e) \geq \Psi_0(e)$, I assume that the marginal disutility function is increasing in a parameter γ . Denote the true value of γ as γ_0 and the estimated one as $\hat{\gamma}$. Since the marginal disutility function that I estimate lies above the true one, then $\hat{\gamma} > \gamma_0$. Therefore, one condition that guarantees $\Psi(e) \geq \Psi_0(e)$ is for $\psi'(e; \gamma)$ to exhibit increasing differences:

$$\psi'(e; \hat{\gamma}) - \psi'(e - \Delta\theta; \hat{\gamma}) > \psi'(e; \gamma_0) - \psi'(e - \Delta\theta; \gamma_0).$$

The online appendix shows that this is indeed the case under the assumed functional form for $\psi(e; \gamma)$.

Finally, as long as $\psi(e; \gamma) - \psi(e - \Delta\theta; \gamma)$ is increasing in γ —which is the case with the assumed functional form—, the social planner’s value function when facing asymmetric information (Laffont, 1994b) will be decreasing in γ . This implies that the welfare gains from the optimal mechanism that I estimate is a lower bound of the true welfare gains.

4 Structural model of abatement costs

In the previous section, I have provided evidence of how RORR affects strategic behavior of regulated electric utilities. In particular, electric utilities strategically have lower fuel efficiency during rate cases to increase operating costs and raise regulated electricity prices. During regulatory lags, electric utilities’ fuel efficiency significantly improves which lowers operating cost and allows the utility to extract rents from the previously set regulated price.

I use these observations together with Assumptions 1 and 2 to identify and estimate a structural model of SO₂ abatement costs. For fuel-switching, the cost of reducing emissions is a function of fuel efficiency and the relative prices of coal with different sulfur contents. If the coal burning process is more fuel efficient, then less coal is needed to produce the same amount of electricity. Hence, the cost of switching to (or blending with) lower sulfur coal—which contains less heat to produce electricity and is relatively more expensive for the coal plants in my sample¹⁵—will be lower for more fuel efficient firms.

Similar to Carlson et al. (2000), I estimate a multiproduct cost function from which I derive a measure of marginal abatement costs for electric utilities. Note that the cost function I estimate is a behavioral cost function in that it does not necessarily represent cost minimization. I perform the analysis at the utility-level since RORR involves the firm as a whole. Ellerman et al. (2000, p. 301) remark that compliance decisions are often made at the utility-level even if pollution regulation *per se* is at the plant-level. Finally, I restrict attention to costs, output, emissions and input choices related to coal, oil and gas plants that the firm owns.

I assume the following stochastic specification for realized operations and maintenance (O&M) variable costs of producing electricity and emissions. For firm i at year t , realized O&M variable cost is given by

$$\tilde{C}_{it} = \exp(\omega_{it})C(q_{it}, s_{it}, p_{lit}, p_{fit}, N_{it}, d_{FGDit}, d_{PBRit}, d_{Lit}, d_{Mit}, d_{95it}, t; \beta) \exp(\varepsilon_{it}) \quad (1)$$

¹⁵Table A coal plants had existing long-term coal contracts that continued to deliver high sulfur coal despite the economy-wide decrease in low sulfur coal from the Powder River Basin (Ellerman et al., 2000). Such long-term contracts distort the relative prices of high and low sulfur coal in a way that makes the former more attractive for Table A plants, at the margin.

where

$$\begin{aligned}
\omega_{it} &= \theta_{it} - e_{it} \\
p_{fit} &= (p_{cit}, p_{oit}, p_{git}) \\
\beta &= (\beta_0, \beta_N, \beta_{FGD}, \beta_{YEAR}, \beta_q, \beta_s, \beta_{sd}, \beta_l, \beta_c, \beta_o, \beta_g) \\
C(q, s, p_l, p_f, N, d_{FGD}, d_{PBR}, d_L, d_M, d_{95}, t; \beta) &= \exp(\beta_0) N^{\beta_N} q^{\beta_q} s^{\beta_s} p_l^{\beta_l} p_c^{\beta_c} p_o^{\beta_o} p_g^{\beta_g} \\
\beta_0 &= \tilde{\beta}_0 + \beta_{YEAR}t + \beta_{FGD}d_{FGD} + \beta_{PBR}d_{PBR} \\
\beta_s &= \tilde{\beta}_s + \beta_{95}d_{95} + \beta_{sd}d_{FGD} + \beta_Ld_L + \beta_Md_M.
\end{aligned}$$

The term $\exp(\omega) = \exp(\theta - e)$ captures cost efficiency where θ is the firm's intrinsic type and e is unobserved effort. The utility knows θ and chooses e . The function

$$C_{it}(\beta) \equiv C(q, s, p_l, p_f, N, d_{FGD}, d_{PBR}, d_L, d_M, d_{95}, t; \beta)$$

is the *baseline cost function* of the utility where q is net generated electricity, s is the SO₂ emission rate, p_l is the average salary for full-time employees related to electricity generation, p_f is a vector composed of fuel prices for coal, oil and gas averaged across the utility's plants, N is the sum of nameplate capacities of the utility's plants, d_{FGD} is a dummy equal to one if the utility has at least one plant with a flue-gas desulfurization (scrubber) unit installed, d_{PBR} is a dummy equal to one if the utility is subject to Performance-Based Regulation (PBR), d_L is a dummy equal to one if the emission rate of the utility is below 1.2 lbs/MMBtu (i.e. low sulfur or compliance coal), d_M is a dummy equal to one if the emission rate is greater than 1.2 lbs/MMBtu but less than 2.5 lbs/MMBtu (the latter being the target emission rate in Phase I and I refer to this as mid sulfur coal), and d_{95} is a dummy equal to one if $t \geq 1995$, i.e. the year the Acid Rain Program (ARP) was implemented. Finally, although there is no explicit SO₂ price before 1995, utilities were still subject to SO₂ regulation under the previous Clean Air Act Amendment.

The baseline cost component captures differences in O&M costs that can be explained by differences in input prices, outputs and capital. The vector β contains the parameters of the baseline cost function that need to be estimated. I allow firms in states with PBR to have a different average baseline cost. Moreover, I let the coefficient on the emission rate β_s to depend on the emission rate s (whether the firm's emission rate is low, mid or high), to whether a firm has a scrubber, and whether observations occur pre and post ARP implementation. Finally ε is a mean zero stochastic error term that summarizes other factors that affect realized costs. I assume ε is unanticipated by the firm when making its input choices and uncorrelated with the regressors.

I assume the firm's type θ_{it} is a draw from the distribution \mathcal{F}_θ . Ideally \mathcal{F}_θ would be conditioned on variables such as firm's capacity or portfolio of plants, but to make estimation more tractable

later, I assume \mathcal{F}_θ is only a function of the rate case year. Next, I allow firms' type to vary across different rate cases but assume θ remains constant between consecutive ones. Formally, let t_τ be the time index (year) for a specific firm's rate case τ . For example, if firm i has three rate cases during the sample period, then $\tau \in \{1, 2, 3\}$ which occurs on years t_1 , t_2 and t_3 respectively. Thus for each i , t and τ ,

$$\theta_{it} = \begin{cases} \theta_{it_\tau} & \text{if } t \in [t_\tau, t_{\tau+1}) \\ \theta_{it_{\tau+1}} & \text{if } t = t_{\tau+1}. \end{cases}$$

I assume for each firm i , θ_{it_1} is a draw from \mathcal{F}_θ , θ_{it_2} is a draw from $\mathcal{F}_{\theta|\theta_{it_1}}$, θ_{it_3} is a draw from $\mathcal{F}_{\theta|\theta_{it_2}}$, etc. The distribution of types I identify will thus reflect the distribution of costs *during the τ -th rate case* of firms that appear in my data, and not the distribution of costs for *all* rate cases of these firms.

In the next subsection, I discuss identification of the baseline cost function parameters β , the disutility function $\psi(\cdot)$, and the distribution of types \mathcal{F}_θ .

4.1 Identification

There are two interrelated challenges for identification. The first challenge is an endogeneity problem in identifying the cost parameters β . The second challenge involves extracting the distribution of the unobserved type θ from the variation in realized costs that is unobserved by the econometrician, i.e. $\exp(\theta_{it} - e_{it}) \exp(\varepsilon_{it})$. The first challenge arises primarily because effort e_{it} is an endogenous variable chosen by the firm. Variables that enter the firm's baseline cost affect the choice of effort since the baseline cost captures the marginal benefit from effort, i.e. cost reduction from a small increase in effort. The firm's cost efficiency ($\omega = \theta - e$) in turn affects costs, electricity output (since regulated electricity prices are based on reported expenses), and potentially, input prices (Cicala, 2015). If ω_{it} were observed by the econometrician, then we can directly control for it, and identify the vector β . However ω_{it} is not observed.

Once we have identified the vector of baseline cost parameters β , we then need to identify the distribution of θ . This distribution is a critical input in the counterfactual analysis since it determines the degree of heterogeneity in marginal abatement costs, the complexity of the optimal mechanism, and the size of the welfare gains. The second challenge then, is to extract the distribution of θ from the unobserved variation $\exp(\theta_{it} - e_{it}) \exp(\varepsilon_{it})$.

My identification strategy involves two parts. First, to identify the parameters of the empirical model, I use similar techniques from the production function literature to extract ω_{it} out of the estimating equations and identify the cost parameters β . In this step, the variation in incentives and cost induced by the RORR is essential in pinning down ω_{it} for different time periods which allows me to use the techniques from the literature. Second, to identify the distribution of types θ ,

I recast the problem in the framework of measurement error with repeated measurements (e.g. Li and Vuong (1988)) and use the deconvolution result of Kotlarski (1967).

4.1.1 Parameters

The potential endogeneity problem in estimating the parameters β is due to simultaneity and selection. Simultaneity may arise if regressors in Equation 1 are correlated with $w_{it} = \exp(\theta_{it} - e_{it}) \exp(\varepsilon_{it})$, which is unobserved by the econometrician. The theoretical model of the paper *assumes* there is correlation since effort e_{it} is chosen by the firm, and this choice is driven by cost-related information. Finally, endogeneity due to selection may also arise since majority of rate cases are initiated by utilities. All else equal, firms may strategically initiate a rate case when operating costs are anticipated to be high.

Problems of simultaneity and selection arise in the estimation of production functions in that unobserved productivity affects input choices—i.e. what is known as “transmission bias” (Griliches and Mairesse, 1998)—, and that firms that we observe are the ones who do not exit (e.g. Olley and Pakes (1996)). To address transmission bias, I follow the dynamic panel approach (Arellano and Bond, 1991; Blundell and Bond, 1998, 2000) by adopting functional form assumptions that exploit the panel structure of the data to difference out unobserved productivity (type, in my case). To handle selection bias, I combine the dynamic panel approach with a simple Heckman (1979) correction.¹⁶ As Kyriadizou (2001) first points out, time-differencing not only eliminates unobserved individual effects but *may* also eliminate the effect of sample selection.

Under Assumptions 1 and 2, high costs during rate cases reflect zero effort and thus $\omega_{it_\tau} = \theta_{it_\tau}$. During the regulatory lag (time $t_\tau + 1$), the firm is the residual claimant to cost-reducing effort and hence exerts the first best level: $\exp(\omega_{it_\tau+1}) C_{it_\tau+1}(\beta) = \psi'(e_{it_\tau+1})$. To determine what ω_{it} is during the regulatory lag, I impose the following functional form for $\psi(\cdot)$ similar to Gagnepain and Ivaldi (2002):

Assumption 3 *The disutility of effort is given by*

$$\psi(e_{it}, v_{it}) = \frac{1}{\gamma} \exp(\gamma e_{it} + v_{it}) - \frac{1}{\gamma}$$

where γ is a parameter and v_{it} 's are mean zero shocks that are uncorrelated with

$$z_{it} = (q_{it}, s_{it}, pl_{it}, p_{fit}, N_{it}, d_{FGDit}, d_{PBRit}, d_{Lit}, d_{Mit}, d_{95it}, t)$$

and *iid* across i and t .

¹⁶Related literature on the control function approach to estimating production functions include Olley and Pakes (1996); Levinsohn and Petrin (2003); Akerberg, Caves and Frazer (2015); Gandhi, Navarro and Rivers (2013); Doraszelski and Jaumandreu (2013).

Finally, to characterize how ω_{it_τ} evolves from one rate case to another, I assume that θ_{it} follows an AR(1) process across two consecutive rate cases similar to Blundell and Bond (1998):

Assumption 4 For each i and τ , a firm's types across two rate cases τ and $\tau-1$ evolve according to

$$\theta_{it_\tau} = \rho\theta_{it_{\tau-1}} + \xi_{it_\tau}$$

where ρ is a parameter and ξ_{it_τ} is iid across i and t_τ .

These assumptions allow me to write the log of realized cost (i) during a rate case ($t = t_\tau$), (ii) during the corresponding regulatory lag ($t = t_\tau+1$), and (iii) during the next rate case ($t = t_{\tau+1}$), all as functions of type θ_{it_τ} :

$$\ln \tilde{C}_{it_\tau} = \theta_{it_\tau} + \ln C_{it_\tau}(\beta) + \varepsilon_{it_\tau} \quad (2)$$

$$\ln \tilde{C}_{it_{\tau+1}} = \frac{\gamma}{1+\gamma}(\theta_{it_\tau} + \ln C_{it_{\tau+1}}(\beta)) + \frac{1}{1+\gamma}v_{it_{\tau+1}} + \varepsilon_{it_{\tau+1}} \quad (3)$$

$$\ln \tilde{C}_{it_{\tau+1}} = \rho\theta_{it_\tau} + \xi_{it_{\tau+1}} + \ln C_{it_{\tau+1}}(\beta) + \varepsilon_{it_{\tau+1}}. \quad (4)$$

By taking (quasi) differences of Equations 2, 3 and 4, I construct the following moment conditions to identify the parameters (β, γ, ρ) :¹⁷

$$E[\eta_{1it_\tau}] = 0 \quad (5)$$

$$E\left[\eta_{2it_\tau} \cdot \begin{pmatrix} z_{it_{\tau-1}} \\ \ln \tilde{C}_{it_{\tau-1}} \end{pmatrix}\right] = 0 \quad (6)$$

$$E\left[\eta_{3it_\tau} \cdot \begin{pmatrix} 1 \\ \ln \tilde{C}_{it_{\tau+1}} \end{pmatrix}\right] = 0. \quad (7)$$

where

$$\begin{aligned} \eta_{1it_\tau} &= \xi_{it_\tau} + \varepsilon_{it_\tau} - \rho\varepsilon_{it_{\tau-1}} \\ \eta_{2it_\tau} &= \frac{1}{1+\gamma}v_{it_{\tau+1}} + \varepsilon_{it_{\tau+1}} - \frac{\gamma}{1+\gamma}\varepsilon_{it_\tau} \\ \eta_{3it_\tau} &= \frac{\gamma}{1+\gamma}(\xi_{it_\tau} + \varepsilon_{it_\tau}) - \rho\left(\frac{1}{1+\gamma}v_{it_{\tau-1}+1} + \varepsilon_{it_{\tau-1}+1}\right). \end{aligned}$$

For the moment conditions to be valid, it suffices to (i) have the shocks η , ε and ξ to be iid across i and t , and unanticipated by the firm, i.e. independent of z_{it-j} for $j > 1$, and (ii) that rate case initiation is not related to anticipated increases in operating cost. Reading through detailed

¹⁷The subsequent discussion shows that the order condition for identification is satisfied. However I have not formally established the rank condition for identification. This requires proving that the system of equations defined by Equations 5, 6 and 7 has a unique solution.

summaries of about 50 randomly picked rate cases in my sample, I find that rate cases are often formally initiated for reasons related to the rate base rather than on anticipated operating cost increases. This observation is consistent with the fact that 80% of rate case disallowances are non operating cost-related. Examples of reasons for initiating rate cases include adjustments of the allowed return on equity, disagreement about how the rate base will be adjusted after a nuclear plant is decommissioned, and inclusion of a previously disallowed plant from the rate base. In these situations, the timing of the rate case is likely to be orthogonal to anticipated increases in operating cost.¹⁸ Nevertheless, I allow for selection in what follows.

I assume that the firm initiates a rate case at t if

$$\theta_{it} - \theta_{it-1} > -\mathbf{Z}'_{it}\boldsymbol{\alpha}, \quad (8)$$

where \mathbf{Z} is a vector of variables related to the firm's perception of how likely it would succeed in a future rate case and $\boldsymbol{\alpha}$ is a vector of unknown coefficients. In other words, initiation of rate cases is driven by a firm anticipating a significant increase in its cost type θ relative to its previous draw, beyond some threshold that depends on rate case variables. Using Assumption 4, we can rewrite the difference in subsequent types as

$$\theta_{it} - \theta_{it-1} = \xi_{it} - (1 - \rho) \sum_{j=1}^t \rho^{j-1} \xi_{it-j}.$$

Letting $u_{it} = -\xi_{it} + (1 - \rho) \sum_{j=1}^t \rho^{j-1} \xi_{it-j}$, we have that the firm initiates a rate case at t if $\mathbf{Z}'_{it}\boldsymbol{\alpha} > u_{it}$. This form of selection invalidates the above moment conditions through the correlation between u_{it} and ξ_{it} . Specifically, moment conditions 5 and 7 are no longer valid. To parameterize the selection equation, I assume the following:

Assumption 5 *Assume that ξ 's are drawn from a normal distribution and let*

$$\mathbf{Z}'_{it}\boldsymbol{\alpha} = \alpha_0 + \alpha_D \text{DISALLOW}_{it} + \alpha_{L1} \text{LAST}_{it} + \alpha_{L2} \text{LAST}_{it}^2$$

where *DISALLOW* is the fraction of the proposed revenue requirement that was disallowed in the previous rate case, and *LAST* is the number of years since the last case.

Given Assumption 5, the selection error u_{it} is also normally distributed. Following Heckman's (1979) two-step approach, I first estimate $\boldsymbol{\alpha}$ using a Probit regression of a dummy variable indicating

¹⁸Suppose the correlation between the timing of rate cases and increase in cost was driven by exogenous improvements in technology and observable decreases in overall cost over time. A well-informed regulator that is aware of these trends *would not* allow future prices to be locked-in based on high costs reported during the rate case. The fact that the data is inconsistent with the behavior of a well-informed regulator goes to show its informational disadvantage and the resulting inefficiency with RORR.

rate case initiation on the vector \mathbf{Z} , and then include the inverse Mills ratio evaluated at $\mathbf{Z}'\boldsymbol{\alpha}$ as an additional regressor in moment conditions 5 and 7.

4.1.2 Type distribution

Given the parameters and using Assumption 4, I can rewrite realized cost during two consecutive rate cases as¹⁹

$$\begin{aligned}\frac{\ln \tilde{C}_{it_\tau} - \ln C_{it_\tau}(\beta)}{\rho} &= \theta_{it_{\tau-1}} + \frac{\xi_{it_\tau} + \varepsilon_{it_\tau}}{\rho} \\ \ln \tilde{C}_{it_{\tau-1}} - \ln C_{it_{\tau-1}}(\beta) &= \theta_{it_{\tau-1}} + \varepsilon_{it_{\tau-1}}.\end{aligned}$$

The problem of finding the distribution of θ can be recast in the framework of measurement error with repeated measurements. Let $(\xi_{it_\tau} + \varepsilon_{it_\tau})/\rho$ and $\varepsilon_{it_{\tau-1}}$ be the “measurement errors” while $\theta_{it_{\tau-1}}$ is the latent variable. The two measurement errors and the latent variable are all mutually independent and this follows from the assumptions on ξ_{it_τ} and the unanticipated cost shocks. The key intuition for identification is that the distribution of the latent variable can be extracted from the joint distribution of the two observed measurements. For example, the covariance between the two measurements reveal the variance of the latent variable under the independence assumptions. Higher moments can be generated from functions of the two measurements, which then reveal higher moments of the latent variable. The online appendix provides more discussion on identification.

4.2 Estimation and results

To estimate the parameters,²⁰ I use the sample analog of the moment conditions given by equations (5), (6) and (7). Ideally I would use the same set of firms to construct the three moment conditions. However these moment conditions taken together require each firm in the sample to have at least two rate cases that are initiated *and* completed in the period 1988-1998. This leaves me with just 26 firms. The vector β contains 11 elements and therefore I need to estimate 13 parameters in total (i.e. β , γ and ρ).

Rate cases that were not completed by 1998 have missing regulatory lag observations. Although these rate cases cannot be used for moment conditions (6) and (7), they can still be used for (5) since it only depends on rate case observations. I define a dummy variable m_i equal to 1 if regulatory lag data is missing and 0 otherwise, and then weight moment conditions (6) and (7) by $(1 - m_i)$. This approach is similar in spirit to the approach proposed by Abrevaya and Donald (2017) to handle

¹⁹The error term ξ appears in the first equation hence is subject to the same selection bias discussed earlier. I correct this equation using Assumption 5.

²⁰I discuss estimation of the type distribution in the online appendix.

missing data using GMM, where (5) serves a similar purpose to the assumed linear projection of regressors with missing observation onto regressors with complete data.

Since types are firm-rate case specific, I treat the same firm in two different rate cases as if they were different firms in order to increase the number of firms in my sample. For example, if PECO had a rate case in 1991 and then followed by another rate case in 1996, I consider $\theta_{it_\tau} =$ “PECO 1991 rate case” and $\theta_{it_{\tau+1}} =$ “PECO 1996 rate case” as two separate “firms.” Although there is dependence across these two firms, this dependence is fully captured by ω_{it} which we differentiate out.

Table 5 presents the parameter estimates.²¹ The coefficient on the emission rate imply that during ARP implementation (i.e. from 1995-1999) for a 10% decrease in emission rates, O&M variable cost increases by about 5.5% if emission rates are low (less than 1.2 lbs/MMBtu). Interestingly, the coefficient on emission rate goes down in magnitude after the implementation of ARP. This is consistent with improved techniques in fuel-switching practices (Ellerman, et al., 2000). The increase in O&M variable cost decreases to 3.5% with a 10% decrease in emission rates within the mid range of 1.2 lbs/MMBtu to 2.5 lbs/MMBtu, although not statistically significant.

To assess model fit, I compare the distribution of the (residualized) log cost in the data with the distribution generated by the model. The model predicts a cost distribution based on the estimated distribution of θ , the effort disutility function parameter γ , and the baseline cost function parameters β . Specifically, in panel (a) of Figure 4, the distribution of costs predicted by the model is the distribution of *non rate case* costs that takes observations *during the rate case* as given, applies the model (optimal effort), and then predicts the distribution of costs if firms exerted first best effort instead. The predicted distribution in panel (b) instead takes *non-rate case* observations as given, predicts what costs would look like if firms exerted zero effort instead of the first best, and compares it with the observed distribution of costs *during the rate case*. For non-rate case cost (panel (a)), the Kolmogorov-Smirnov statistic is 0.1290 with p -value of 0.130, and therefore we cannot reject that the two distributions are the same. A difference in means test has a t -statistic

²¹The coefficient estimates of the first-step Probit regression are: $\alpha_0 = 1.11$ (SE = 0.05), $\alpha_D = -4.1$ (SE = 1.78), $\alpha_{L1} = -0.91$ (SE = 0.10) and $\alpha_{L2} = 0.08$ (SE = 0.02). The estimates suggest some interesting results. First, the likelihood of initiating a rate case declines as the firm experiences a larger disallowance in the past. Second, for every observation in the data where the next rate case is at least one year apart, I compute a positive second derivative of the probability of initiation with respect to time since rate case, suggesting a U-shaped relationship between the likelihood of rate case initiation and time since the last case. New rate cases that occur almost immediately after the previous case often reflect appeals which is thus consistent with the likelihood of a new case having an initial negative relationship with the time since the last rate case. However, as time passes by, an appeal is less likely to occur and the time since the last rate case is now positively correlated with the likelihood of a new case. Finally, the coefficient on the inverse Mills ratio is -1.09×10^{-7} with SE equal to 0.002.

Table 5: Parameter estimates

log O&M variable cost	Model	
	Est	SE
log emission rate	-0.479**	0.142
log emission rate*FGD	0.298**	0.095
log emission rate*LOW	-0.546*	0.196
log emission rate*MID	-0.346	0.262
log emission rate*95	0.612**	0.177
log Electricity output	0.675***	0.140
log Price of coal	0.647***	0.110
log Price of oil	0.177***	0.052
log Price of gas	0.112***	0.041
log Price of labor	0.063***	0.046
log Nameplate	-0.213**	0.121
FGD	0.768*	0.561
PBR	-0.019**	0.287
Year	-0.003***	0.003
Disutility (γ)	19.236**	3.15
Type evolution (ρ_1)	0.658***	0.121
Inverse Mills	-1.09×10^{-7} **	0.002

Notes: Standard errors for these are computed using bootstrap, where sampling is over firm-rate case. Significance levels are determined using the bootstrap confidence intervals. Significance level: * 10%, ** 5%, *** 1%.

of 0.23 and p -value of 0.819. For rate-case cost (panel (b)), the Kolmogorov-Smirnov statistic is 0.2261 with p -value of 0.001 while a difference in means test (data versus model) has a t -statistic of 2.08 and p -value of 0.038. Although the model slightly under-predicts rate case costs for the lower part of the cost distribution, the fit improves after.

The distribution of θ implies a distribution of marginal abatement costs (MACs). I compute what MACs the model predicts under the parameters of Phase I of the Acid Rain Program. Figure 5 plots the distribution of marginal abatement costs (MACs) assuming (i) all firms have emission rate of 2.5 lbs per MMBtu, (ii) observable variables (electricity output, input prices and fuel burned) are at their mean values, and (iii) firms do not have FGDs installed (i.e. $d_{FGD} = 0$). The emission rate of 2.5 is the implicit emission standard under Phase I of the Acid Rain Program. Thus this distribution is the model's prediction of the distribution of MACs if the SO₂ regulation were implemented through a uniform emission standard. Additionally, I compute the distribution of MACs assuming zero effort (panel (a)) and assuming optimal first best effort (panel (b)).

The estimated MAC distributions exhibit significant heterogeneity and a long tail despite controlling for observable variables. Butraw (1999) presents a range of average MAC estimates from various studies in the literature (both engineering-based and econometrically estimated). These MACs range from \$291 to \$760. Note that these estimates from the literature are affected by variation in output and input prices, while the distribution of MACs I estimate are fully generated by the estimated distribution of θ .

5 Counterfactual welfare analysis

I use the estimated structural parameters to simulate SO₂ emissions under counterfactual regulatory regimes. I take a random sample of cost efficiency types from the estimated distribution of θ . Denote this sample as Θ . Next, I rewrite the baseline cost function as a function of emission rate s , i.e. $C(s) = \Xi s^{\beta_s}$, where $\beta_s < 0$ and $\Xi = \exp(\beta_0) N^{\beta_N} q^{\beta_q} p_l^{\beta_l} p_c^{\beta_c} p_o^{\beta_o} p_g^{\beta_g}$. Since Ψ is a function of electricity output and input prices, I take the mean of Ξ in the data and use this in the simulations. Thus all the variation in the simulations is due to variation in θ . Finally, to get aggregate numbers, I calculate the implied number of firms such that aggregate emissions in a regulatory regime with a uniform emission standard is equal to the number of freely distributed emission permits in Phase I of the ARP. The basic rule for freely allocated permits multiplies the emission rate of 2.5 lbs/MMBtu with historical fuel consumption, and this amounts to about 5.456 million tons of SO₂ (Joskow and Schmalensee, 1998), which is, by construction, the level of observed SO₂ emissions.

Throughout the simulations, I assume firms face a Pigouvian tax on SO₂ emissions. Denote the Pigouvian tax as p . Under competitive pricing, the firm exerts optimal first best effort $e^{FB}(\theta) > 0$

and chooses emissions $s^{FB}(\theta)$ such that the efficient first best marginal abatement cost is equal to the Pigouvian tax. Assuming $p = 100$,²² emissions under the first best scenario is 4.211 million tons, which is 30% lower than observed emissions. This is consistent with first best marginal abatement cost being below the marginal abatement cost under RORR, as in Figure 1. Finally, emissions under a scenario assuming zero effort (i.e. as in during rate cases) is equal to 5.883 million tons. Thus, effort under RORR is much closer to zero effort than first best effort.

I now consider scenarios where electricity prices remain regulated. Since there is full cost reimbursement during rate cases with RORR, the firm will exert zero effort. On the other extreme, if the economic regulator sets an electricity price that is fixed regardless of the firm's reported cost, i.e. a fixed price contract, then the firm has an incentive to exert first best effort, as if it were facing a competitive electricity price. Since firms should receive at least zero profits as required by economic regulation, the fixed price should be set high enough to accommodate even the least efficient firm. This participation constraint leads to strictly positive economic profits (information rents) for all but the least efficient firm. In the context of economic regulation, strictly positive economic profits are costly due to the social cost of public funds since these profits could have been used to replace funds raised through distortionary taxes (Laffont and Tirole, 1993; Goulder et al., 1997). Thus, while the fixed price contract maximizes efficiency, it ignores these information rents which can be substantial to the point of eliminating the efficiency benefits of high powered schemes such as fixed price contracts. In fact, when these information rents are large, it may not be desirable or even feasible to implement the first best allocations (Spulber, 1998). Laffont and Tirole (1986) characterize the optimal mechanism that balances the tradeoff between efficiency distortion and information rent extraction.

I focus on emission rates as the regulatory variable, taking the quantity of electricity, capital and input prices as exogenously given. To make analysis easier, I characterize a regulatory regime as a direct revelation mechanism that specifies a bundle (s, \tilde{C}, t) for each type θ . The bundle consists of an emission rate s , a realized and observable cost \tilde{C} , and a lump-sum transfer t (= revenue requirement RR). Finally, the mechanism can be implemented as follows: the firm reports its emissions rate and realized cost to the regulator, and then the regulator provides a transfer given this report.

Following Laffont (1994b), the planner maximizes social welfare given by

$$W = \int \{V(q(\theta)) - D(s(\theta)) - (1 + \lambda)t(\theta) + \Pi(\theta)\} d\mathcal{F}(\theta), \quad (9)$$

where $V(q)$ is the surplus from consuming electricity, $D(s)$ is the damage from emissions, λ is the

²²This value for the Pigouvian tax is consistent with a marginal damage under a planned 2.5 lbs/MMBtu standard, at least as revealed *ex post* by the SO₂ permit market with prices settling around this value.

cost of public funds and $\Pi(\theta)$ is firm's profit given by

$$\Pi(\theta) = t(\theta) - [\tilde{C}(\theta) + \psi(e(\theta))]. \quad (10)$$

In choosing the optimal mechanism, the planner faces a couple of constraints. First, the planner needs to satisfy participation constraints which require leaving firms with nonnegative economic profits: $\Pi(\theta) \geq 0$ for all θ . As in Laffont (1994b) and Laffont and Tirole (1986, 1993), I assume the social planner observes realized cost but not the firm's type and effort. This informational constraint leads to incentive compatibility constraints

$$\Pi(\theta) \geq t(\theta') - \left\{ \tilde{C}(\theta') + \psi[e(\theta', \theta)] \right\}$$

for all θ and $\theta' \neq \theta$ where $e(\theta', \theta) = \theta - \ln \frac{\tilde{C}(\theta')}{\tilde{C}(s(\theta'))}$. These constraints ensure that a particular type θ does not have an incentive to pick some other type's bundle.

Given a regulatory regime, I compute welfare as

$$\widetilde{W}(p, \lambda) = \frac{1}{N} \sum_{\theta \in \Theta} \{-p \cdot \mathcal{S}(s(\theta)) - (1 + \lambda)t(\theta) + \Pi(\theta)\}.$$

This welfare metric \widetilde{W} does not include the surplus from electricity consumption and thus I focus on welfare differences across different regulatory regimes. The linear function $\mathcal{S}(s(\theta))$ converts an emission rate $s(\theta)$ to tons of SO₂ emissions using the mean amount of fuel burned. Note that I impose a linear pollution damage function so, with some abuse of notation, $p > 0$ represents the constant marginal damage from a ton of pollution (as well as the earlier Pigouvian tax). The variable $\lambda > 0$ is the social cost of public funds. I treat (p, λ) as simulation parameters and I compute \widetilde{W} for different combinations of (p, λ) and different regulatory regimes. Throughout, dollar amounts are in annual 1995 USD.

5.1 Optimal mechanism

The optimal mechanism maximizes \widetilde{W} subject to the participation and incentive compatibility constraints. I provide more detail and closed-form solutions in the online appendix. For reference, I compare the optimal mechanism with a fixed price contract. The fixed price contract is of interest since it implements the first best emissions and effort levels, which minimizes the sum of damages and abatement costs. However, because firms have private information about their intrinsic cost types and effort, the planner has to pay sufficiently high profits to induce firms to reveal their information, i.e. information rents.

In comparing the optimal mechanism and the fixed price contract, I focus on the case with $p = 100$ and $\lambda = 0.3$ (30 cents for each dollar of public funds). The value of 0.3 for the social cost

of public funds is motivated by estimates in the public finance literature (Laffont, 2005; Ballard, Shoven and Whalley, 1985).

The optimal mechanism distorts first best emissions upwards to reduce the necessary information rents paid to firms. There are two sources of distortion. First, the optimal mechanism allows less efficient types to emit more relative to their first best levels. Second, conditional on the same level of abatement, the optimal mechanism induces less effort. These two distortions only lead to second order losses in welfare. In exchange for these distortions, first order gains are achieved by lowering information rents.

Aggregate emissions under the optimal mechanism is 0.19 million tons higher compared to the fixed price contract. Note that in both scenarios, firms face the same Pigouvian tax. Emissions are higher under the optimal mechanisms since effort levels are distorted downwards compared to the fixed price contract, leading to a higher marginal abatement cost curve under the optimal mechanism. In terms of generation and emissions abatement cost (including cost of effort), cost under the optimal mechanism is \$203 million higher compared to the fixed price contract.

Although efficiency is higher with the fixed price contract, the economic regulator needs to set the fixed price to be high enough so that even the least efficient type earns nonnegative profits. This means that types that are relatively more efficient all get strictly positive profits, which carries welfare costs as long as $\lambda > 0$. The difference in information rents between the fixed price contract and the optimal mechanism is \$540 million. The large information rents basically wipe out the efficiency advantage of the fixed price contract. At the end, welfare under the optimal mechanism is higher by \$318 million.

Now consider the other extreme with full cost reimbursement and zero effort as in during a rate case (RORR). Although information rents are zero, generation and emissions abatement cost are \$416 million higher compared to the optimal mechanism. The optimal mechanism yields annual welfare gains of \$686 million relative to full cost reimbursement.

Finally, I let λ vary and compute welfare gains from the optimal mechanism relative to RORR. Annual welfare gains range from \$578 million to \$1.06 billion. These correspond to about a 9% to 17% reduction in electricity prices assuming reduction in costs are fully passed on to consumers.

5.2 Simple menu

Laffont (1994b) shows that the optimal mechanism can be implemented using a set of emission taxes and monetary transfers, both of which are functions of reported cost. Each cost type will then choose a particular pair of emissions tax and transfer, and report the corresponding cost to the regulator. The menu is incentive-compatible by design. Such a mechanism is likely to be complicated to implement in practice since the regulator has to design and implement a full menu

of emission taxes and transfers, the complexity of which depends on the heterogeneity in types. What we often see are much simpler contracts composed of a single contract or a menu consisting of a limited number of contracts (Rogerson, 2003). Consistent with this observation, I construct a simple menu composed of two choices: a fixed price contract or full cost reimbursement.

Define the cutoff type $\hat{\theta}$. The simple menu is constructed such that types that are less efficient than $\hat{\theta}$, i.e. $\theta > \hat{\theta}$ choose full cost reimbursement, while types that are more efficient, choose the fixed price and receive the transfer \hat{T} . Types that choose full cost reimbursement exert zero effort but get zero economic profits. In contrast, types that choose the fixed price exert optimal positive effort and earn a payoff equal to the difference between \hat{T} and their true cost (i.e. includes effort cost). Using the estimated type distribution and cost of effort, I solve for the optimal fixed price \hat{T} that would maximize social welfare (subject to individual rationality and incentive compatibility constraints). It turns out that \hat{T} is equal to the true cost of the cutoff type when it exerts optimal positive effort, so the problem is equivalent to finding the optimal cutoff type.

Figure 6 plots the cutoff type and the fraction of welfare gains from the optimal mechanism captured by our simple menu as we vary the social cost of public funds. The cutoff type is decreasing with λ since larger λ means that information rents carries more welfare costs, and therefore it pays to distort effort and convince more types to choose full cost reimbursement. For $\lambda = 0.3$, the cutoff type is about 0.75. Types that have efficiency above the 25th percentile opt for the fixed price contract and exert the first best level of effort while types below the 25th percentile opt for full cost-reimbursement and exert zero effort. This observation is consistent with the substantial heterogeneity of the estimated type distribution. That is, “dropping” (inducing zero effort) the most inefficient firms substantially reduces information rents paid to everyone else who are exerting first best effort.

Finally, in terms of welfare gains captured by our simple menu, for our benchmark value of $\lambda = 0.3$, the simple menu captures 65% of the welfare gains. The value of λ where the curve hits its minimum is the λ where the optimal mechanism diverges the most from either a fixed price contract or full cost reimbursement. Despite the divergence, my results suggest that a simple two menu contract can still capture a large fraction of welfare gains from the optimal mechanism. For $\lambda \in [0, 1]$, the simple menu captures at least 65% of the welfare gains.

6 Conclusion

The paper estimates large welfare gains from a mechanism that explicitly takes into account a regulated firm’s informational advantage and the cost of incentive provision (Laffont and Tirole, 1986). Furthermore, I construct a simple mechanism that can achieve a substantial fraction of

these welfare gains.

The shape of the type distribution is an important determinant in the design of the optimal mechanism and the welfare gains it delivers. The gains from the optimal mechanism is not simply about getting *all* firms to exert more effort, but rather more about effectively mitigating the cost of the regulator's informational disadvantage. The cost of the informational disadvantage is determined by the shape of the type distribution and therefore importantly affects the measure of welfare gains. Thus, as in most studies on asymmetric information, the key challenge is to estimate this distribution.

I have exploited the timing of rate cases and how this timing induces different incentives for fuel efficiency as the main source of identification. However, the approach I adopt to specifically handle transmission and selection bias is highly parametric and crucially depends on the assumed functional forms. One such assumption is the linearity of cost with respect to the firm's type and the baseline cost function. While this assumption allows me to conveniently estimate the type distribution and parameters of the baseline function separately, it assumes away a richer model for the type distribution which can potentially explain the substantial heterogeneity in unobserved types that I estimate. For example, if the linearity assumption does not obtain, then the true type distribution can still be an explicit function of variables that both the firm and regulator observe hence reducing the latter's informational disadvantage. More generally, the assumptions I take impose a specific structure on the uncertainty that the regulator faces relative to the firm which then importantly determines the gains from mitigating this uncertainty.

The empirical exercise in this paper covers a time period (1988-1999) when almost all of the states' generating capacity were still subject to RORR. Moreover, at this time, dirty coal was still the dominant fuel source. The regulatory environment has evolved since, though a good number of states still have their generation subject to RORR, and a sizable fraction of electricity generation is still being produced using fossil fuels. Finally, the paper focuses on the efficiency of generating electricity rather than the specific abatement method. Whether the future fuel mix will be predominantly coal, natural gas, nuclear or renewables, efficiency in electricity generation will remain an important topic that regulators and policy-makers care about. Well-designed incentive mechanisms and simpler derivatives guided by theory can potentially yield large welfare gains without needing to switch fuel, install expensive equipment, or fundamentally change how we produce electricity.

References

- Akerberg, D. A., K. Caves and G. Frazer (2015), "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83:6, 2411-2451.
- Alt, L. E. (2006), *Energy Utility Rate Setting: A Practical Guide to the Retail Rate-Setting Process for Regulated Electric and Natural Gas Utilities*, Morrisville, NC: Lulu.com.
- Armstrong, M. and D. M. Sappington (2007), "Recent Developments in the Theory of Regulation," In *Handbook of Industrial Organization*, vol. III, edited by M. Armstrong and R. Porter, 1557-700. Amsterdam: Elsevier Science Publishers.
- Arellano, M. and S. Bond (1991), "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations," *Review of Economic Studies*, 58:2, 277-297.
- Arellano, M. and S. Bonhomme (2012), "Identifying Distributional Characteristics in Random Coefficients Panel Data Models," *Review of Economic Studies*, 79:3, 987-1020.
- Bailey, E. M. (1998), "Allowance Trading Activity and State Regulatory Rulings: Evidence from the U.S. Acid Rain Program," mimeo, MIT.
- Ballard, C. L., J. B. Shoven and J. Whalley (1985), "General Equilibrium Computations of the Marginal Welfare Costs of Taxes in the United States," *American Economic Review*, 75:1, 128-138.
- Baron, D. P. and R. B. Myerson (1982), "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50:4, 911-930.
- Besanko, D. (1985), "On the Use of Revenue Requirements Regulation Under Imperfect Information," in M. A. Crew (ed.), *Analyzing the Impact of Regulatory Change in Public Utilities*, Lexington, MA: Lexington Books, 39-55.
- Blundell, R. and S. Bond (1998), "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics*, 87, 115-143.
- Blundell, R. and S. Bond (2000), "GMM Estimation with persistent panel data: an application to production functions," *Econometric Reviews*, 19:3, 321-340.
- Bose, A., D. Pal, and D. E. M. Sappington (2011), "On the Performance of Linear Contracts," *Journal of Economics and Management Strategy*, 20:1, 159-193.

- Brocas, I., K. Chan and I. Perrigne (2006), "Regulation under Asymmetric Information in Water Utilities," *American Economic Review, Papers and Proceedings*, 96, 62-66.
- Buchanan, J. M. (1969), "External Diseconomies, Corrective Taxes, and Market Structure," *American Economic Review* 59:1, 174-177.
- Butraw, D. (1998), "Cost Savings, Market Performance, and Economic Benefits of the U.S. Acid Rain Program," Resources for the Future, Discussion Paper 98-28-REV.
- Butraw, D. (1999), "Cost Savings, Market Performance and Economic Benefits of the US Acid Rain Program," in S. Sorrell and J. Skea (eds.), *Pollution for Sale: Emissions Trading and Joint Implementation*, Edward Elgar Publishing, 43-62.
- Carlson, C. P., D. Butraw, M. Cropper and K. Palmer (2000), "SO₂ Control by Electric Utilities: What are the Gains from Trade," *Journal of Political Economy*, 108:6, 1292-1326.
- Carroll, G. (2015), "Robustness and Linear Contracts," *American Economic Review*, 105:2, 536-563.
- Chan, G., R. Stavins, R. Stowe and R. Sweeney (2012), "The SO₂ Allowance-trading System and the Clean Air Act Amendments of 1990: Reflections on 20 Years of Policy Innovation," *National Tax Journal*, 65:2, 419-452.
- Chan, H. S., H. Fell, I. Lange and S. Li (2013), "Efficiency and Environmental Impacts of Electricity Restructuring on Coal-fired Power Plants," *Journal of Environmental Economics and Management*, 81 (January), 1-18.
- Chassang, S. (2013), "Calibrated Incentive Contracts," *Econometrica*, 81:5, 1935-1971.
- Chu, L. Y. and D. E. M. Sappington (2007), "Simple Cost-Sharing Contracts," *American Economic Review*, 97:1, 419-428.
- Cicala, S. (2015), "When Does Regulation Distort Costs? Lessons from Fuel Procurement in U.S. Electricity Generation," *American Economic Review*. 105:1, 411-444.
- Dai, T. and J. Toikka (2017), "Robust Incentives for Teams," mimeo, MIT.
- Davis, L. W. and C. D. Wolfram (2012), "Deregulation, Consolidation, and Efficiency: Evidence from U. S. Nuclear Power," *American Economic Journal: Applied Economics*, 4:4, 194-225.
- Doraszelski, U., and J. Jaumandreu (2013), "R&D and Productivity: Estimating Endogenous Productivity," *Review of Economic Studies*, 80, 1338-1383.


- Ellerman, A. D., P. L. Joskow, R. Schmalensee, J.-P. Montero and E. Bailey (2000), *Markets for Clean Air: The U.S. Acid Rain Program*, Cambridge, UK: Cambridge University Press.
- Ellerman, A. D. and J.-P. Montero (2001), "A Puzzle in SO₂ Emissions Trading: Temporary Inefficiency or Higher Prices," mimeo.
- Ellerman, A. D. and J.-P. Montero (2007), "The efficiency and robustness of allowance banking in the US Acid Rain Program," *The Energy Journal*, 28, 47-61.
- Environmental Protection Agency (2001), *Acid Rain Program: 2001 Progress Report*, EPA-430-R-02-009.
- Environmental Protection Agency (2014, July 7) *Overview of Greenhouse Gases*, Retrived from <http://www.epa.gov/climatechange/ghgemissions/gases/co2.html>.
- Evdokimov, K. and H. White (2012), "Some Extensions of a Lemma of Kotlarski," *Econometric Theory*, 28:4, 925-932.
- Fabrizio, K., N. Rose and C. Wolfram (2007), "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency," *American Economic Review*, 97:4, 1250-1277.
- Fowlie, M. (2010), "Emissions Trading, Electricity Industry Restructuring, and Investment in Pollution Control," *American Economic Review*, 100:3, 837-869.
- Fowlie, M., M. Reguant and S. P. Ryan (2014), "Market-based emissions regulation and industry dynamics," *Journal of Political Economy*, 124:1, 249-302.
- Gagnepain, P. and M. Ivaldi (2002), "Incentive Regulatory Policies: The Case of Public Transit Systems in France," *RAND Journal of Economics*, 33:4, 605-629.
- Gagnepain, P., M. Ivaldi and D. Martimort (2013), "The Cost of Contract Renegotiation: Evidence from the Local Public Sector," *American Economic Review*, 103:6, 2352-2383.
- Gandhi, A., S. Navarro and D. Rivers (2013), "On the Identification of Production Functions: How Heterogenous is Productivity," mimeo, Univ of Wisconsin-Madison.
- Garrett, D. F. (2014), "Robustness of simple menus of contracts in cost-based procurement," *Games and Economic Behavior*, 87, 631-641.
- Goulder, L. H., I W. H. Parry and D. Butraw (1997), "Revenue-raising versus Other Approaches to Environmental Protection: The Critical Significance of Preexisting Tax Distortions," *RAND Journal of Economics*, 28:4, 708-731.

- Han, J. S., J.-F. Houde, A. A. van Benthem, and J. M. Abito (2018), "When Does Regulation Distort Costs? New Insights from Coal Procurement," mimeo, University of Pennsylvania.
- Heckman, J. J. (1979), "Sample Selection Bias as a Specification Error," *Econometrica*, 47:1, 153-161.
- Joskow, P. L. (2008), "Incentive Regulation and its Application to Electricity Networks," *Review of Network Economics*, 7:4, 547-560.
- Joskow, P. L. and R. Schmalensee (1986), "Incentive Regulation for Electric Utilities," *Yale Journal of Regulation*, 4, 1-49.
- Joskow, P. L. and R. Schmalensee (1998), "The Political Economy of Market-Based Environmental Policy: The U.S. Acid Rain Program," *Journal of Law and Economics*, 41:1, 37-83.
- Joskow, P. L., R. Schmalensee and E. M. Bailey (1998), "The Market for Sulfur Dioxide Emissions," *American Economic Review*, 88:4, 669-685.
- Kahn, A. E. (1988), *The Economics of Regulation: Principles and Institutions*, Cambridge, MA: MIT Press.
- Kotlarski, I. I. (1967), "On Characterizing the Gamma and Normal Distribution," *Pacific Journal of Mathematics*, 20, 69-76.
- Laffont, J.-J. (1994a), "The New Economics of Regulation Ten Years After," *Econometrica*, 62:3, 507-537.
- Laffont, J.-J. (1994b), "Regulation of Pollution with Asymmetric Information," in C. Dosi and T. Tomasi (eds.), *Nonpoint Source Pollution Regulation: Issues and Analysis*, Kluwer Academic Publishers, 39-66.
- Laffont, J.-J. (2005), *Regulation and Development*, Cambridge, UK: Cambridge University Press.
- Laffont, J.-J. and J. Tirole (1986), "Using Cost Observation to Regulate Firms," *Journal of Political Economy*, 94, 614-641.
- Laffont, J.-J. and J. Tirole (1992), "Cost Padding, Auditing and Collusion," *Annales d'Économie et de Statistique*, 25/26, 205-226.
- Laffont, J.-J. and J. Tirole (1993), *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA: MIT Press.

- Levinsohn, J. and A. Petrin (2003), "Estimating Production Functions Using Inputs to Control for Unobservables," *Review of Economic Studies*, 317-342.
- Li, T. and Q. Vuong (1998), "Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators," *Journal of Multivariate Analysis*, 65, 139-165.
- Lim, C. S. H. and A. Yurukoglu (2018), "Dynamic Natural Monopoly Regulation: Time Inconsistency, Asymmetric Information, and Political Environments," *Journal of Political Economy*, 126:1, 263-312.
- Nobel Media AB (2014), "The Prize in Economic Sciences 2014 - Advanced Information," Nobel-prize.org. Web. 9 Nov 2017. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2014/advanced.html.
- Olley, S. and A. Pakes (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263-1295.
- Perrigne, I. and Q. H. Vuong (2011), "Nonparametric Identification of a Contract Model With Adverse Selection and Moral Hazard," *Econometrica*, 79:5, 1499-1539.
- Reiss, P. C. and F. A. Wolak (2007), "Structural Econometric Modeling: Rationales and Examples from Industrial Organization," in *Handbook of Econometrics*, ed. J. J. Heckman and E. E. Leamer, 4277-4415. Amsterdam and Oxford: Elsevier, North-Holland.
- Rogerson, W. (2003), "Simple Menus of Contracts in Cost-Based Procurement and Regulation," *American Economic Review*, 93:3, 919-926.
- Schmalensee, R. (1989), "Good Regulatory Regimes," *Rand Journal of Economics*, 20, 417-436.
- Schmalensee, R. and R. N. Stavins (2012), "The SO₂ Allowance Trading System: The Ironic History of a Grand Policy Experiment," Harvard Kennedy School Faculty Working Paper Series, RWP12-030.
- Schmalensee, R., P. L. Joskow, A. D. Ellerman, J. P. Montero and E. M. Bailey (1998), "An Interim Evaluation of Sulfur Dioxide Emissions Trading," *Journal of Economic Perspectives*, 12:3, 53-68.
- Spulber, D. F. (1998), "Optimal Environmental Regulation under Asymmetric Information," *Journal of Public Economics*, 35:2, 163-181.
- Wolak, F. A. (1994), "An Econometric Analysis of the Asymmetric Information, Regulator-Utility Interaction," *Annales d'Economie et de Statistiques*, 34, 13-69.

Figures

Figure 1: Distortion due to Rate of Return Regulation



Notes: $C_a^*(q, a)$ and $\hat{C}_a(q, a)$ are the marginal abatement cost under competitive pricing and RORR, respectively.

Figure 2: Empirical Analysis

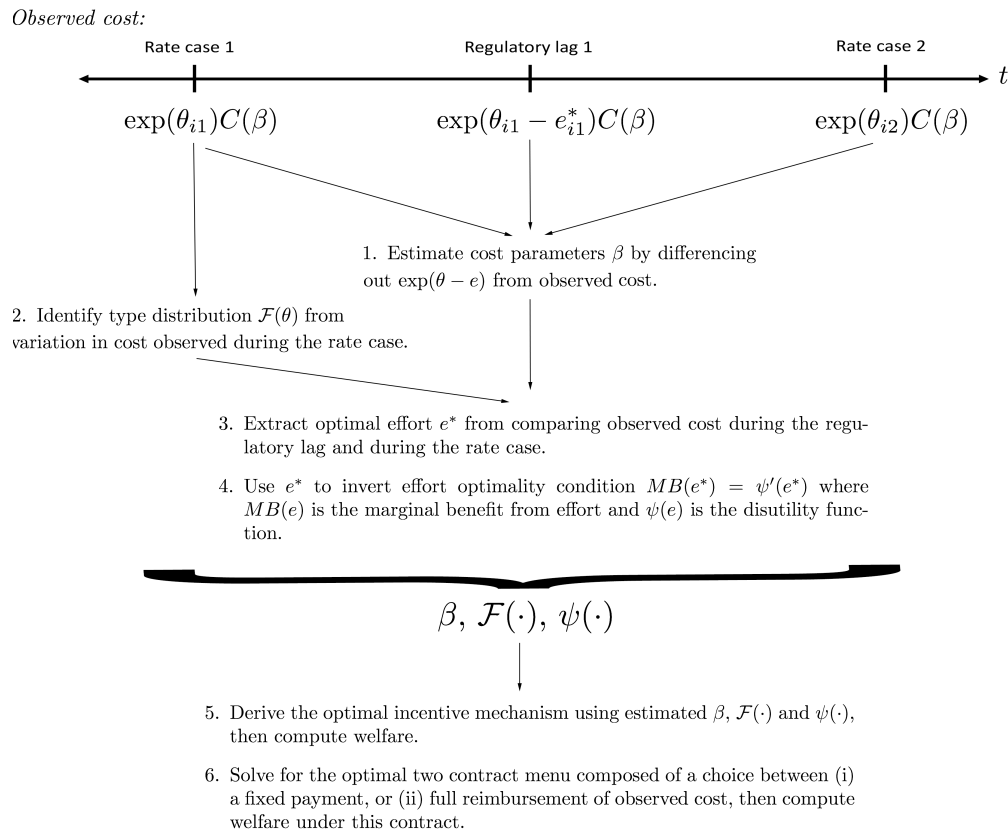


Figure 3: Bounds analysis

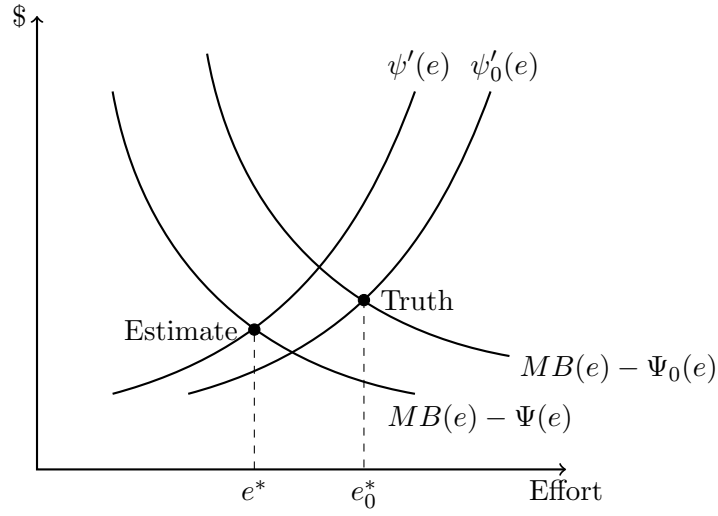


Figure 4: Model fit

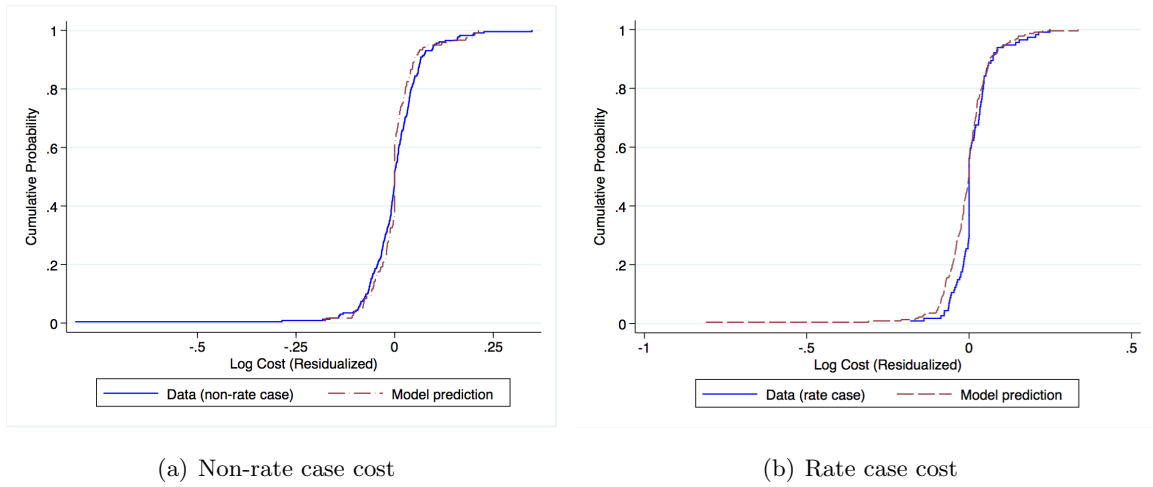
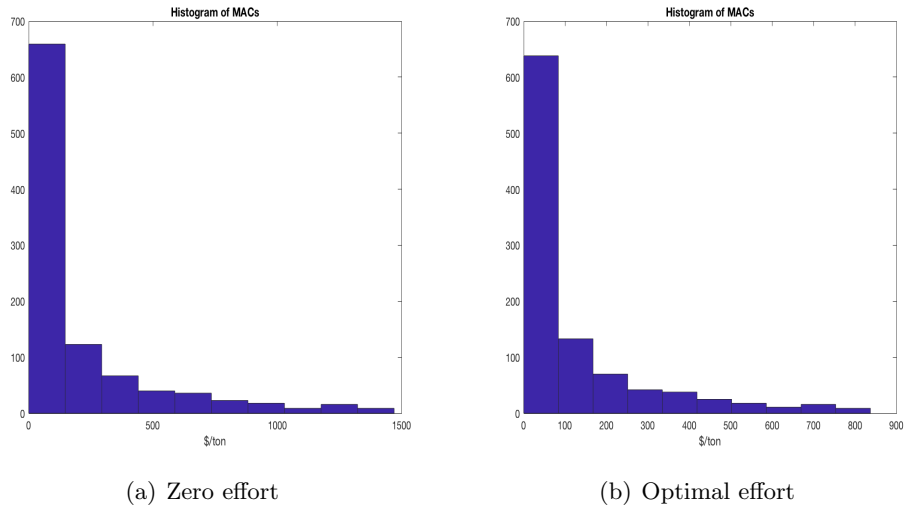


Figure 5: Histogram of marginal abatement costs in \$ per ton of SO₂ emissions under optimal and zero effort



Notes: The figure contains the histogram of marginal abatement costs (MAC) for the random sample I drew from the estimated type distribution. MACs are evaluated at an emission rate of 2.5 lbs/MMBtu and expressed in 1995\$ per ton.

Figure 6: Simple Menu

