

Dynamic Responses to Carbon Pricing in the Electricity Sector

Paige Weber

Preliminary draft: latest version available at <https://paige-weber.weebly.com>

August 15, 2019

Abstract

This paper investigates the impact of a regulation that prices carbon on an industry with dynamic production decisions. While carbon pricing is designed to reduce greenhouse gas (GHG) emissions, it also impacts the emissions of local air pollutants. I demonstrate that with dynamic production decisions it is theoretically possible for a carbon price in the electricity sector to increase pollution in some areas. However, the outcome depends on the empirical cost structure of the regulated firms. I develop a dynamic model of electricity production and generating unit efficiency investment decisions to analyze how firm decisions in California's electricity sector are affected by the state's carbon price. I use the model to simulate market outcomes across policy scenarios with and without carbon prices, as well as a location-specific Pigouvian tax on air pollutants. I show that under current carbon prices, the regulation leads to minimal production reallocation and small changes in the spatial distribution of local air pollutants compared to a no carbon price scenario. Higher carbon prices lead to production reallocation and co-benefits from changes in the spatial distribution of air pollutants; though, the co-benefits occur predominantly outside of state's heavily polluted regions. On the other hand, the location-specific Pigouvian tax on local pollution concentrates these benefits in heavily polluted regions. In addition, I demonstrate that under plausible conditions, private and social returns are largest when efficiency investments occur in the cleanest, most frequently utilized units.

Acknowledgements: This project is supported by Alfred P. Sloan Foundation Pre-doctoral Fellowship on Energy Economics, awarded through the National Bureau of Economic Research (NBER).

1 Introduction

Firm responses to regulation depend on both their own costs, and in many cases, the cost structure of the regulated industry. In some settings, the impact of a regulation on market outcomes depends on whether firm decisions are static or dynamic. Understanding a regulation’s impact on market outcomes and the mechanisms driving those impacts is critical to evaluating the efficiency, and equity, of a given policy.

In this paper, I investigate how a regulation impacts firm decision making in a dynamic competitive equilibrium. I study this in the context of a carbon price in California’s wholesale electricity sector, and I develop a dynamic model of electricity generating unit (unit) production and efficiency investment decisions in a market regulated with a carbon price. A carbon price could change the spatial distribution of local air pollution to the extent that it re-allocates production across units. I use the model to estimate how a carbon price impacts production decisions, and thus the local air pollutants emitted by firms across space. The effect of a carbon price on the spatial distribution of air pollutants has important distributional consequences, as low income and minority communities in many cases are more likely to reside in regions that bear disproportionate pollution burdens. I use this model to simulate and compare the market and distributional consequences of alternative policy scenarios, including combining the carbon price with a location-specific Pigouvian tax on local air pollution.

Greenhouse gas (GHG) regulations aim to alleviate a global burden—future damages from climate change—yet they affect the emissions of co-pollutants that have local impacts as well. Climate change policy is often motivated by both the benefits from mitigating global warming as well as the benefits human health from the reduction of local air pollutants. While pricing carbon is often the preferred approach to addressing global climate change, its impact on the distribution of co-pollutants that create local air pollution is not well understood and depends on the cost structure of the regulated firms and industry. Critics of carbon pricing contend that the policy’s local impacts lead to worse air quality

outcomes for regions that are already heavily polluted. Yet, understanding to what extent this occurs requires understanding how pricing carbon interacts with firm decision making in the regulated industry.

The regulation impacts firm decision making through two key mechanisms. First, the carbon price could lead to a change in aggregate supply curves in hourly markets for wholesale electricity. Thus, the share of total electricity demanded that is provided by each unit (unit market share) could change. I demonstrate that conditional on firm efficiency, production reallocation will only occur with dynamic production decisions, which occur in this setting from the presence of fixed “start-up costs” to turn units on. Second, the carbon price could lead to changes in unit efficiencies through investment decisions. The carbon price enters analogously to an increase in the cost of inputs, and thus, investments in efficiency should weakly increase following the introduction of the carbon price.

This reallocation of production from either of the mechanisms above creates scope for changes in the spatial distribution of local air pollution. While the total quantity of local air pollutants emitted by the electricity sector may be lower, some firms may increase production following the carbon price, leading local air pollutants to increase in some regions compared to a no policy scenario. The estimates of location-specific marginal damages from local air pollutants by Muller and Mendelsohn (2007) demonstrate how these damages vary across space, and the welfare effects of changes in air pollution depend on the spatial distribution of these changes.

This paper contributes to the existing literature in several ways. In terms of methodology, I develop a model of dynamic competition to characterize near- and medium-term responses to regulation with endogenous firm efficiencies and hourly production decisions. The model builds on other dynamic competitive equilibrium frameworks such as Jovanovic (1982) and Hopenhayn (1992) and is most similar to Cullen and Reynolds (2017). The model here differs from Cullen and Reynolds (2017) in that I allow firms to vary in efficiency, and this characteristic is endogenous. In addition, because the goal of the model is to capture near-

and medium- term responses to the policy, I do not consider firm entry or exit as they do.¹ Consistent with their approach, I leverage the correspondence between the competitive equilibrium and a cost minimization problem to facilitate the empirical analysis.

Overall, there is little work connecting climate change policy to distributional impacts in local air quality, despite an on going policy debate in California that centers around this question. One policy report (Cushing et al. 2016) finds descriptive evidence that suggests that the state’s carbon policy is leading to worse air quality outcomes for heavily polluted, low-income and minority communities. Meng (2017) weighs into this debate with a reduced-form econometric approach and finds that emissions do not increase in disadvantaged communities as a result of the policy. However, the paper does not address the dynamic aspects of production decision making, which create the potential for redistribution of local air pollution as a result of the carbon price. The analysis in this paper contributes to this debate, illustrating theoretically how a carbon price could lead to poorer air quality outcomes in some regions, as well as using an empirical analysis to estimate these outcomes in a dynamic setting.

In terms of policy implications, this work contributes to several strands of literature regarding the distributional consequences of energy, electricity, and climate change policy. With respect to electricity policy, Borenstein (2012) reviews the distributional impacts of electricity pricing programs on household bills, and Borenstein (2013) provides a policy framework for how to equitably implement dynamic electricity pricing. Borenstein (2017) examines the variation in incentives to adopt residential solar energy across income groups. Fowlie et al. (2012) study the distributional impacts of an emissions trading program for NO_x in Southern California, finding that the program’s benefits are not significantly different across neighborhood demographics. On GHG regulations and equity, Borenstein and Davis (2016) find that clean energy tax credits are likely to be more regressive than market-based GHG regulations, and Fullerton et al. (2012) and recent work by Goulder et al. (2018)

¹That said, as will be shown in the modeling section, the repeated production decisions in this model can be likened to a repeated entry and exit model.

review the relative regressivity of a carbon tax across alternative policy designs. Finally, perhaps most relatedly, Knittel and Sandler (2018) discuss the efficiency and equity consequences of uniform policy instruments when the level of externality generated varies across agents. They show that a uniform gasoline tax is a second-best policy when the externalities from gasoline consumption are heterogeneous across drivers. This paper considers a setting where the local air pollution externalities are heterogeneous across electricity generating units and demonstrates how a uniform carbon price fails to fully internalize these air quality externalities.

I use the model to simulate market outcomes across policy scenarios without a carbon policy, with a carbon policy at observed and higher carbon prices, and with a carbon policy at observed prices combined with a location-specific Pigouvian tax on local air pollution. The results of the model indicate that the cost structure of California's fossil-fuel electricity portfolio leads to minimal production reallocation at current carbon prices, and thus, the policy has a limited effect on the spatial distribution of air pollutants emitted by the electricity sector. While this may quell concerns about the negative unintended consequences of the carbon price's impact on the distribution of local air pollutants, the analysis also reveals that the carbon policy does little to improve air quality in regions of the state that have been identified as disproportionately polluted. Counterfactual simulations of local air quality outcomes under higher carbon prices increase production reallocation across units compared to a no carbon policy scenario, leading to aggregate co-benefits of the carbon policy from its impact on NO_x emissions compared to no carbon policy, with 30 percent of these benefits occurring in disproportionately polluted regions. I compare this to a scenario with a carbon policy at current prices combined with a location-specific Pigouvian tax on local air pollution, which doubles the benefits from avoided NO_x damages compared to a carbon policy alone and concentrates these benefits in the relatively more polluted regions.

The rest of the paper is organized as follows. Section 2 describes the model of production and efficiency investment decisions, and reviews the theoretical predictions of the model

regarding the carbon price’s impact on production and investment. Section 3 describes the empirical strategy. Section 4 reviews data sources and the empirical setting. Section 5 presents the results, and Section 6 concludes.

2 Dynamic Model of Investment and Production

In this model decisions are made at the electricity generating unit level, where a unit consists of a heat engine that converts fuel to energy. Units are assumed to act competitively in this model and make decisions as single agents. Accordingly, I treat each unit as acting as an individual firm and use the terms ‘unit’ and ‘firm’ interchangeably.² Firms make two decisions. First, the firm makes a one-time decision of whether or not to invest to improve its efficiency. A firm’s efficiency is measured by its heat rate, which is the amount of fuel required to be combusted to produce one unit of electricity. Consequently, lower heat rates translate to improved efficiency and lower input costs. This investment decision then establishes the firm’s efficiency for subsequent repeated hourly operation decisions. I model the investment decision as myopic to market outcomes beyond three years of hourly market outcomes. In practice, this three-year time period could be modified: I select three years as I am interested in the near- and medium- term impact of the carbon price on production and efficiency investments in the first three years of the cap-and-trade program, 2013 - 2015. In the production stage, firms make repeated binary operation decisions, deciding whether or not to produce electricity in each hourly market. Operation is more costly if the firm was not operating in the prior period due to the presence of start-up costs, which are incurred from the additional costs of fuel, auxiliary power, water, additives, chemicals, and unit wear and tear required to bring a unit online (Kumar et al. 1992). Firms make operation decisions based on their expectations of future demand. Hourly demand is assumed to be inelastic to wholesale electricity prices, and demand shocks are modeled as an AR(1) Markov process,

²I discuss this assumption in section 4.1.

conditional on hour of the day.³

2.1 Investment Decision

For the investment decision, firms decide how much to improve their efficiency by reducing their current heat rate ω' , measured in Btu per KWh. Firms choose j from a discrete set of activities in set \mathcal{J} of size J , which includes $j = 0$ corresponding to no investment. This modeling choice flexibly captures an array of investment activities, from capital investments such as upgrading turbine technology, to human capital such as hiring additional staff to improve efficiency through changes to operations and maintenance, without specifying the activity itself. Further, modeling investment choices as discrete is supported by the data, which exhibits lumpy investment behavior.⁴ Following the investment decision, firm i 's heat rate ω_i is:

$$\omega_i = \omega'_i(1 + \tilde{\delta}) - j_i, \tag{1}$$

where $\tilde{\delta} \in (0, 1)$ is an exogenous depreciation rate that decreases unit efficiency, i.e., increases the unit's heat rate, and corresponds to the time in between the investment decision and the production decision. Investments to improve efficiency have costs that increase in the size of the efficiency improvement and are denoted Γ :

³Residential customers pay retail rates, which allows me to treat hourly residential demand as inelastic to the price firms are paid, the wholesale price for electricity. There are some caveats with this approach. One, industrial customers may pay rates closer to wholesale prices. Two, some retail customers may have opted into demand-side management programs that provide incentives to change demand in response to wholesale prices. There is a separate question about changes in demand over time, that is, the question of whether we should expect shifts in the demand curve over time. We could imagine that a growing economy leads to an increase in demand over time; alternatively, we could foresee that an increase in renewable energy reduces the residual quantity demanded by this fossil portfolio. In this paper I do not take a stand on future demand, and consider the response of the fossil portfolio for a given level of demand.

⁴See section A.5 for more information about empirically observed investment.

$$\Gamma(j_i, v_i, \gamma) = \gamma j_i + v_i, \quad (2)$$

where v_i is a stochastic shock to investment costs, drawn from an extreme value type 1 distribution.

2.2 Operating and Production Decision

Firms also make binary operating decisions $a_{it} \in [0, 1]$, conditional on ω_i chosen in the investment decision. The quantity firm i produces in each hour q_{it} follows from its operating decision, based on the following decision rule:

$$q_{it} = \begin{cases} q_{i,max} & \text{if } P_t \geq mc_i \text{ and } a_{it} = 1 \\ q_{i,min} & \text{if } P_t < mc_i \text{ and } a_{it} = 1 \\ 0 & \text{if } a_{it} = 0. \end{cases} \quad (3)$$

P_t denotes the wholesale equilibrium prices in hour t ; mc_i denotes the marginal costs of production, defined below; and $q_{i,min}$ and $q_{i,max}$ denote the firm's minimum and maximum operation levels.⁵ Modeling production as a discrete choice is supported by the data, which shows that units generally operate at one of a discrete set of production quantities, as shown in Figure 18 in Appendix 8.8.⁶ Then, the per period profits for the firm are defined by the products below:

⁵The included simulations set $q_{min} = 0.75 \cdot q_{max}$, where q_{max} is the reported operation capacity of the unit in MW. I compare this approach to estimating $q_{i,min}$ and $q_{i,max}$ from the data using finite mixture models, and find similar estimates when averaging over the unit type categorizations that are used in the empirical analysis.

⁶The number of discrete operating levels varies across the units; limiting the number of operating levels to two as is done here is a simplifying assumption.

$$\pi_{it} = \begin{cases} q_{it}(P_t - mc_i) & \text{if } a_{i,t-1} = 1 \\ q_{it}(P_t - mc_i) - \kappa_i & \text{if } a_{i,t-1} = 0, \end{cases} \quad (4)$$

where κ_i denotes start-up costs. Marginal costs are a function of unit efficiency (heat rate), ω_i ; costs of fuel, c^f ; emissions intensity of the fuel, e^f ; and price of GHG permits, τ :⁷

$$mc_i = \omega_i(c^f + e^f\tau). \quad (5)$$

The above formulation shows that when c^f and e^f are positive and τ is non-negative, mc_i is increasing in heat rate, $\frac{\partial mc_i}{\partial \omega_i} > 0$. Further, the formulation demonstrates that investments to reduce heat rate reduce marginal costs by decreasing both fuel costs, $\omega_i \cdot c^f$, and compliance costs, $\omega_i \cdot e^f\tau$. In this sense, the GHG program can be seen as an increase in the cost of inputs to production, analogous to an increase in natural gas prices, as discussed in Mansur (2008).⁸ In this model, both fuel costs and carbon prices are known and exogenous, though future research could explore allowing these two input costs to differ in terms of how firms make expectations over future natural gas versus carbon prices. For the purposes of defining the state variables for the model, define marginal input costs, ic , as the sum of marginal fuel and GHG prices, $c^f + e^f\tau$, where the second term is first converted to \$ per Btu.

2.3 Firm's Dynamic Programming Problem

I now specify the state variables and their transitions and formulate the firm's dynamic programming problem. In the production decision, the firm observes state variables, \mathbf{s} , for

⁷The unit formulation of (5) is $\frac{\$}{KWh} = \frac{Btu}{KWH} * \frac{\$}{Btu} + \frac{Btu}{KWH} * \frac{emissions}{Btu} * \frac{\$}{emissions}$. For units also regulated by the Acid Rain Program and/or the NOx Budget Program, the marginal cost calculation includes compliance costs for these programs. These compliance cost are calculated as the permit costs in the respective program times NO_x and SO_2 emissions measured by the unit's continuous emissions monitoring system.

⁸Mansur (2008) leverages this similarity between incentives from natural gas costs and GHG prices to evaluate carbon abatement costs in electricity markets using changes in natural gas prices.

the current period demand shock, hour, its lagged operating state, its heat rate, and the input costs, $\mathbf{s} = \{\eta_t, h_t, a_{t-1}, \omega, ic\}$. A firm makes expectations about future period demand shocks, and chooses a_t to maximize the sum of future discounted profits. Firms are assumed to have rational beliefs about future demand, and demand is modeled as an AR(1) process conditional on hour of the day, h_t . Hourly demand is assumed to be inelastic to wholesale prices, and characterizing firms' expectations of demand shocks is sufficient to characterize their beliefs about wholesale prices, P_t . The lagged operating state, a_{t-1} , equals 1 when the unit was on in the last period and 0 otherwise. Input costs, ic , are exogenous, time invariant, and known to the firms, and the hour of day evolves as $h_{t+1} = h_t + 1 - \mathbb{1}(h_t = 24) \cdot 24$. The value for a firm in the production stage can be formulated as a choice-specific value function indexed by the investment decision, j :

$$V^{2j}(\eta_t, h_t, a_{t-1}, \omega^j, ic) = \max_{a_t \in \{0,1\}} \left\{ \sum_{t=0}^{\infty} \delta^t [q_t(P_t - mc(\omega^j)) - \mathbb{1}(a_{t-1} = 0, a_{it} = 1) * \kappa] \right\}, \quad (6)$$

where the second term on the right-hand side reflects start-up costs incurred for every period t that the firm operates when it was not operating in $t - 1$.

Next, I formulate the value function for a firm in the investment decision. The firm makes investment decisions by comparing the costs of investment to the payoffs in production from the reduction in fuel and compliance costs as a result of improved efficiency, less the costs of investment:

$$V^1(\mathbf{s}) = \max_{j \in \mathcal{J}} \{ \tilde{\delta} E[V^{2j}(\mathbf{s})] - \Gamma(j, v, \gamma) \}, \quad (7)$$

where the initial demand shock in the stage is same as the initial demand shock in the production stage. The optimal policy for production and investment is characterized by the value functions in 6 and 7 above.

2.4 Cost Minimization Problem

To evaluate the solution the firm’s two-stage dynamic programming problem, I leverage the correspondence between the profit-maximizing choices in a competitive equilibrium and the solution to a production cost minimization problem (henceforth, cost minimization problem). This correspondence is demonstrated to hold in this setting by Cullen and Reynolds (2017). The correspondence follows intuition in earlier work in dynamic competitive equilibriums (Robert E. Lucas and Prescott 1971; Jovanovic 1982; Hopenhayn 1992), and Cullen and Reynolds (2017) extend these results to a setting with repeated entry and exit into hourly electricity markets and non-convexities in the aggregate production technology. Accordingly, establishing that the correspondence holds in this setting requires the additional assumption that firms are small relative to the total size of the market. With this additional small firms assumption, the model is operationalized by finding the solution to a cost minimization problem.

The cost minimization problem is to find the allocations of production and efficiency investment across all units in the electricity portfolio that minimize the costs of meeting hourly electricity demand. Initially, costs will include carbon prices and exclude local air quality damages as in the empirical setting; subsequently, a counterfactual simulation will modify costs to include local air pollution damages as well. To review this problem, I introduce additional notation to collect production quantities and investments across all units into vectors. Define \mathbf{j} as a vector of investment choices for firms $i \in \{1, \dots, U\}$, where $j_i \in \mathcal{J} \forall i$, and denote the set of feasible \mathbf{j} as \mathcal{J} . Similarly, let \mathbf{q} be a U -sized vector of production quantities corresponding to each firm, where $q_i \in \{q_{i,min}, q_{i,max}\} \forall i$, and let \mathcal{Q} denote the set of feasible \mathbf{q} . Then, an allocation $\{\mathbf{j}, \mathbf{q}\}$ is feasible if $\mathbf{j} \in \mathcal{J}$ and $\mathbf{q} \in \mathcal{Q}$. Let \mathbf{a} be a U -sized vector with binary elements $a_i \in \{0, 1\}$ indicating whether firm i was on in the last period. The costs, G , associated with electricity demanded for a given demand state, η_t , are defined:

$$G(\eta_t, h_t, \mathbf{a}_{t-1}, \boldsymbol{\omega}, \mathbf{q}) = \sum_{i=1}^U [mc_i q_i - \mathbb{1}(a_{i,t-1} = 0, a_{it} = 1) * \kappa_i]. \quad (8)$$

As with the exposition of the firm problem, I formulate the investment choice-specific value functions for the production decision for a given \mathbf{j} :

$$W^{2j}(\eta_t, h_t, \mathbf{a}_{t-1}, \boldsymbol{\omega}^j, ic) = \max_{\mathbf{q} \in \mathcal{Q}} \{-G(\eta_t, h_t, \mathbf{a}_{t-1}, \boldsymbol{\omega}, \mathbf{q}) + \delta E[W^{2j}(\eta_t, h_t, \mathbf{a}_{t-1}, \boldsymbol{\omega}^j, \mathbf{q})]\}. \quad (9)$$

The stage one problem is to choose the optimal investment vector \mathbf{j} where Γ now corresponds to the sum of investment costs associated with investment vector \mathbf{j} and where \mathbf{v} denotes the U -sized vector of stochastic shocks to investment costs:

$$W^1(\mathbf{s}) = \max_{\mathbf{j} \in \mathcal{J}} \{\tilde{\delta} E[W^{2j}(\mathbf{s})] - \Gamma(\mathbf{j}, \mathbf{v}, \gamma)\}. \quad (10)$$

2.5 Theoretical Analysis

Impact of carbon price on firm market share with static efficiencies

In this section I review the impact of the carbon price on firm market share, $\zeta_i = q_i/q^d$. I review this holding efficiencies fixed, i.e. without considering investment, so as to isolate the impact of the carbon price all else equal. Understanding whether, and to what extent, market shares change as a result of the regulation is critical to this paper's analysis as changes in the spatial distribution of local air pollutants are driven by production reallocation across firms. Without production reallocation, the relative distribution of pollutants is unchanged even if there are aggregate local air quality co-benefits from reducing GHGs.⁹ I demonstrate

⁹This is driven by the feature of this empirical setting whereby reducing GHGs requires reducing fuel burned, and therefore the regulation changes co-pollutants as well as GHGs. On the other hand, adoption of carbon capture and storage technologies would change this analysis, since this would allow firms to reduce GHGs without impacting other co-pollutants.

that with static decision making the carbon price does not lead to reallocation of production across units, and thus does not provide for a change in the distribution of local air pollutants. Then, I show that in a dynamic setting production reallocation can occur. I illustrate that in a dynamic setting, market shares are weakly increasing among the more efficient units (i.e. decreasing in ω), $\frac{\partial \zeta_i}{\partial \omega_i} \leq 0$.

To start, let us review the carbon price’s impact on marginal costs. Equation 5 shows that with $\tau > 0$, $\frac{\partial^2 mc_i}{\partial \omega_i \partial \tau} = e^f$. Since $e^f > 0$, we have that the increase in marginal costs following the introduction of a non-zero τ is increasing in the firm’s heat rate, ω_i . That is, the increase in marginal costs is larger for the relatively less efficient firms. Now, let us review the impact across static and dynamic settings. In the static case, we can formulate the supply curve for each hourly electricity market by ranking the units by marginal costs.¹⁰ The result above demonstrates that the carbon price increases the slope and intercept of the supply curve, but does not change the ranking of the units by marginal costs.¹¹ Figure 1 shows how an illustrative supply curve changes following the introduction of $\tau > 0$. Note, as in the model above, firms supply discrete quantities in this market, so that the empirical supply curve is a step function comprised of discrete quantities of generation from each unit. As earlier, hourly electricity demand is assumed to be inelastic to wholesale electricity prices, so the quantity demanded following the carbon price in the near-term is the same. Then, at a given level of demand, q_d , the share of demand provided by each firm, ζ_i , is the same following the imposition of the carbon price, while the wholesale electricity price is higher. Thus, a static model of electricity production predicts no production reallocation across units as a result of the carbon price, and resultantly, no redistribution of local air pollution.

Next, I examine this with dynamic production decisions, where a firm’s participation in hourly electricity markets is a function of both marginal costs and start-up costs. As shown in Section 2.3, we can formulate the firm’s operation decision as a dynamic

¹⁰The ranking of units by costs is called the “merit order” in electricity markets.

¹¹This result is driven by the cost structure of this industry, where the portfolio of units with a GHG compliance obligation are natural gas, for which the only way to reduce GHG compliance costs is to burn less fuel (in a world where carbon capture and storage is not cost-competitive).

programming problem. Cullen (2015) demonstrates that we can solve the firm’s dynamic production problem by evaluating choice-specific value functions, where firms operate when the choice-specific value of operating for a given state, s , is greater than not operating. Denote the choice-specific value functions of operating and not operating conditional on state s as V_i^1 and V_i^0 respectively, and define the probability of operating $\phi_i = \Pr[V_i^1 > V_i^0]$. Consider a set of convex indifference curves that map marginal costs on the x-axis and start-up costs on the y-axis to equal values of V^1 to the firm, where curves closer to the origin correspond to higher values. Suppose that V_1^1 and V_2^1 lie on the same indifference curve, with $mc_1 < mc_2$, $\kappa_1 > \kappa_2$, and $\frac{\phi_1}{\phi_2} = \alpha$. Following the introduction of the carbon price, the firms with higher marginal costs experience larger increases in costs, and thus larger losses in value, compared to firms with lower marginal costs. Then, assuming that the change in the value of not operating, V_i^0 , is weakly less than the change in V_i^1 , the ratio of operating probabilities α increases following the introduction of $\tau > 0$. I demonstrate that this assumption holds in the Appendix 10. An increase in α for a given level of demand implies that the market share for firm one is weakly increasing following the introduction of the carbon price. The intuition for this result is that the carbon price increases the weight of marginal costs relative to start-up costs, where the change is larger for units that are less efficient (higher marginal costs). Extending this result to the portfolio of units provides the theoretical prediction that market shares are decreasing in heat rate (increasing in efficiency) $\frac{\partial \zeta_i}{\partial \omega_i} \leq 0$. The amount that market shares change, however, will depend on the empirical distribution of marginal costs and start-up costs across firms in a given electricity portfolio.

Impact on incentives to invest

Here I review two theoretical predictions relating to the impact of the carbon price on investments in efficiency. The first is that investments to improve efficiency (reduce ω) are increasing in τ . This is straightforward from the formulation of the marginal cost function.

Since $\frac{\partial mc_i}{\partial \omega_i} = c^f + e^f \tau$, $\tau > 0$ increases the returns to reducing ω , so that investments to improve efficiency will be weakly increasing in τ .

Next, I demonstrate that under certain conditions, private and social returns to investment are increasing in ζ_i . Let the private net returns to investment j be the production cost savings from efficiency improvements, and let the social net returns be avoided damages from GHG emissions. As shown above in the model, investment j corresponds to new efficiency $\omega_i = \omega'_i(1+\delta) - j_i$. For a given P and q^d , private net returns in a given period are $(\omega_i - \omega'_i)c^f q_i$ and social net returns are $(\omega_i - \omega'_i)\tau e^f q_i$. It is clear then, that the discounted sum of both private and social net returns over future periods is increasing in q_i and therefore ζ_i .

Now, compare the payoffs of investment across firms. If investment costs are not a function of current heat rate, as formulated in this model, then it is clear from the above that higher payoffs are achieved from improving the efficiency of the units with higher market shares. Next, assume that it is more costly to improve the relatively more efficient units, so that costs are now a function of pre-investment efficiency with $\frac{\partial \Gamma_i}{\partial \omega_i} < 0$. Suppose that firms one and two have $\omega_1 < \omega_2$; investment decisions j'_1 and j''_2 cost the same, i.e., $\Gamma(\omega_1, j_1) = \Gamma(\omega_2, j_2)$; and $\omega_1 - j'_1 < \omega_2 - j''_2$. In this setting it is still optimal to invest in firm one when:

$$\begin{aligned} (\omega_1 - \omega'_1)(c^f q_1 + \tau e^f) &> (\omega_2 - \omega'_2)(c^f q_2 + \tau e^f) \\ q^1 &> \frac{1}{c^f} \left(\left[\frac{\omega_2 - \omega'_2}{\omega_1 - \omega'_1} \right] (c^f q_2 + \tau e^f) - \tau e^f \right). \end{aligned} \tag{11}$$

This result illustrates that if the market share of a given firm is sufficiently large such that the above holds for the observed distribution of pre-investment efficiencies, other firm market shares, and cost parameters, a policy subsidizing efficiency improvements yields the highest private and social returns when targeting higher market share firms.

3 Empirical Strategy

The model above is used to characterize the firm’s optimal investment decision and recover its optimal policy function for production (dispatch policy function). First, I group units into N representative firm types. Then, I use the optimal investment decision and dispatch policy functions to simulate hourly market outcomes of production for each unit type under alternative policy scenarios. To do so, I solve the cost minimization problem in two stages. First, I find the cost minimizing solution to meeting hourly demand across different investment portfolios \mathbf{j} . That is, I recover investment choice-specific policy functions for production for J^N investment portfolios.¹²

Next, I use each of these policy functions to simulate different sequences of hourly market outcomes over three years, and I sum the discounted production costs associated with each outcome. Then, I use the estimates of investment costs (discussed subsequently), to find the optimal investment decision \mathbf{j}^* given the estimated production costs across each scenario. Then, I use simulations from the dispatch policy function associated with \mathbf{j}^* to compare market outcomes across alternative input cost states. I discuss this process in further detail below.

3.1 Dispatch Policy Function

The first step is to recover investment choice-specific policy functions for dispatch. To do so, I first group units into representative unit types based on efficiency and size, as these characteristics describe the heterogeneity across units that enters the cost minimization problem. I use k-means and scree plot analysis to establish unit type groups, and details on this process are described in Appendix 8.4. Then, let i now refer to unit type rather than individual units, $i \in \{1, \dots, N\}$, with $N = 10$. Let \mathbf{a} now be a vector with elements a_i

¹²To reduce the computational burden, I evaluate $N = 5$, that is, I consider five unit investment type groups. Further, I set $J = 2$ so that each unit investment type has the option to invest or not. When a given unit investment type invests, all firms of that type invest and reduce their heat rate by 1.5 percent, which is the average heat rate improvement observed in the data among investing units.

indicating the number of units of type i that were operating in the last period. The optimal dispatch policy function for each investment choice \mathbf{j} is denoted $\sigma^{\mathbf{j}}(\eta, h, \mathbf{a}, \omega^{\mathbf{j}}, ic)$. The policy function maps state variables to \mathbf{q} , which has N rows and two columns, where the entry in row i column one corresponds to the number of units of type i operating at their minimum generation level, and the entry in row i column two corresponds to the number of units of type i that are operating at their maximum generation level. The total number of units on at either operating level is constrained by the total number of units of that type, and the optimal dispatch policy for each state will satisfy condition (7) without violating any constraints.

I use policy function iteration to find the optimal policy function for each $\mathbf{j} \in \mathcal{J}$, using an initial estimate of unit type start-up costs (further discussed subsequently) and an exogenous discount rate.¹³

3.2 Optimal Investment Choice

After the optimal dispatch policy $\sigma^{*\mathbf{j}}(\mathbf{s})$ is recovered for each \mathbf{j} investment scenario, the optimal investment choice \mathbf{j}^* is made by solving (7), using the sum of discounted production costs over three years as a measure of the value of each investment scenario.

3.3 Structural Cost Parameters

The structural cost parameters in this model that are not completely observed by the econometrician are the firm's start-up costs and the cost of investment activities.

For start-up costs, I use the calibrated estimate used by Cullen and Reynolds (2017) of \$80 per MW. In addition, I estimate start-up costs with the following identification argument and estimation approach. In the context of the firm problem, start-up costs can be identified by the difference in a firm's willingness to operate across two states that differ only in

¹³To set a discount rate, I use a one year interest rate of 4.1 percent, implying an hourly discount rate of 0.99954130.

the lagged operating state. In the cost minimization problem, under the assumption that the empirically observed dispatch is cost minimizing, start-up costs can be identified by comparing the difference in dispatch implied by solving the cost minimization problem in a given state with an initial guess of start-up costs, κ_i^0 , and the empirically observed dispatch for the same state. Accordingly, I estimate start-up costs, $\hat{\kappa}_i$, by evaluating the difference between cost minimizing dispatch implied by the dispatch policy function recovered for some κ_i^0 and the empirically observed dispatch. Specifically, I use a generalized method of moments approach (GMM) to find the $\hat{\kappa}_i$ that minimizes deviations between the simulated dispatch and empirical counterparts across like states. Details about this procedure are provided in Appendix 8.3.

To estimate investment costs, I first use SNL Financial (SNL) data on capital expenditures to develop an estimate of the cost per percent heat rate improvement for units that invest. SNL data is discussed further subsequently, and details on how I identify evidence of persistent heat rate improvements that indicate investment are provided in Appendix 8.5. Next, I use a simulated method of moments (SMM) approach with the following identification argument. In the cost minimization problem, under the assumption that empirically observed investments are cost minimizing, investment costs can be identified by the difference in investment behavior observed empirically, and that implied by finding the investment decision that minimizes production costs given an initial guess of investment costs, γ^0 .

To implement the SMM approach, I first estimate investment conditional choice probabilities (ICCPs) from the observed investment decisions, by unit investment type. Then, I simulate investment moments using these ICCPs. I compare these simulated investment decisions to investment decisions from solving equation 10 with an initial guess of investment costs, γ^0 . The value function in equation 10 is approximated with the three-year production costs associated with the investment scenario, plus the costs of investment. The production cost estimates are calculated from the forward simulation performed with the recovered investment choice-specific dispatch policy functions, $\sigma^{*2j}(\cdot)$. As discussed earlier, I assume

the investment cost shock comes from an extreme value type 1 distribution, and I estimate the location and scale distribution parameters for the shock from the SNL data. Then, I estimate $\hat{\gamma}$ by finding the parameter that minimizes the distance between the simulated moments using the ICCPs from the data and the investment decisions implied by the model. The advantage of using the SMM approach here is that it provides additional investment moments to match, which is useful in this setting as otherwise there are very few moments.¹⁴ Additional detail about this procedure is provided in Appendix 8.2.

3.4 Estimating the Distribution of Demand Shocks

To estimate the model, I must first estimate distributions of hourly demand shocks, which are used to construct transition matrices for each hour. Demand shocks are modeled as a function of hour of the day fixed effects and last period’s demand shock. This modeling approach is not only parsimonious, but it explains a significant amount of hour-to-hour variation in demand. For illustration, the OLS estimation of $\eta_t = \alpha + \beta\eta_{t-1} + \rho h_t + \epsilon$ has an R^2 of 0.95 (see Appendix 8.7). Accordingly, I use the data to construct distribution scale parameters for the demand shock, conditional on hour of the day and last period shock, and I use these conditional distributions to construct transition matrices. Details of this procedure are available in Appendix 8.7.¹⁵

Of note, the presence of must-take resources in the California creates a setting where the natural gas units studied here supply electricity up to the quantity of “residual” electricity demanded. Residual demand in this setting equals the total domestic demand less energy supplied by zero marginal cost resources, such as solar and wind energy, and lower cost resources such as from hydro-electric resources, nuclear energy, and electricity imports. In this way, a positive shock to renewable energy supplied in any hour, for example, would impact dispatch the same way as a negative shock to demand. Future work could explore

¹⁴On the other hand, GMM is sufficient to recover start-up costs, as in that setting there are many dispatch moments available across different states.

¹⁵The simulations currently used collapse the 12 demand shocks to two, for computational parsimony.

the impact of increasing renewables on market outcomes, as well as the impact of the carbon price on the quantity of electricity imported.

4 Empirical Setting

4.1 Data Sources

Estimating this model requires observing size, efficiency, and location, as well as hourly production quantities, fuel inputs, and emissions for all units regulated by the cap-and-trade program in California.¹⁶ Unit-characteristics and hourly production and emissions quantities are obtained from the subscription-based data provider SNL, which compiles data collected through various federal reporting requirements, including continuous emissions monitoring systems (CEMS) installed at the units. In addition, reported measures of capital expenditures from SNL are used to estimate investment costs, further described in Appendix 8.2. To connect changes in unit emissions of local air pollutants to damages from human health impacts, I use marginal damage per ton estimates from the Air Pollution Emission Experiments and Policy Analysis Model (APEEP) developed by Muller and Mendelsohn (2007).¹⁷

The fossil-fuel portfolio operating in this period in CA consists of around 200 producing units across steam, gas, and combined cycle combustion turbines, with summary statistics available in Table 1.¹⁸ To reduce the computational burden, these units are grouped into 10 unit types as discussed previously, based on unit heat rates and sizes, with summary statistics provided in Table 2. The process to establish unit type groups that explain the most variation in units along these two dimensions is explained in Appendix 8.4.

The model in this paper assumes that the market is competitive, where the actions of the electricity firms do not impact prices. While the well-known Borenstein, Bushnell,

¹⁶Hourly prices are also obtained from publicly available data maintained by the California Independent System Operator (CAISO) to estimate inferred profits in Figure 2.

¹⁷The current results use the APEEP estimates for damages from pollutants emitted in 2011.

¹⁸The coal in 2013 are from out of state units.

Wolak (BBW) paper finds evidence of market power in the California electricity market, the market has undergone significant changes since the study period evaluated in that paper. In particular, in early 2009 the CAISO implemented a host of market reforms intended to improve price transparency and prevent market manipulation as part of its Market Redesign and Technology Upgrade (MRTU). MRTU improved the grid operator’s ability to manage real-time congestion with day-ahead generation and transmission schedules; increasingly local marginal pricing by moving from three to 3,000 price nodes; and created an integrated forward market for electricity, transmission capacity, and reserves (CAISO 2009). Through MRTU and other reforms the grid operator in California has evolved significantly in its ability to detect and prevent market manipulation since BBW’s analysis. For example, CAISO currently maintains a process to submit local market power mitigation bids in locations and times that may present opportunities for local market power, such as in heavily congested times and regions. Further, the Herfindahl-Hirschman Index (HHI) for the natural gas units is 0.06, and the majority of units provide less than two percent of total fossil-fuel generation. While these facts do not entirely rule out the presence of non-competitive behavior, the modeling approach in this paper can be understood as predictive of market outcomes when units behave competitively, which, as the market monitoring reports from California’s grid operator suggest, has been the case for the large majority of hours studied in this paper.¹⁹

4.2 Descriptive Assessment

Understanding the impacts of the GHG policy on local air quality outcomes requires characterizing changes in market shares and efficiencies as a result of the carbon price, as well as the spatial distribution of these changes. Regarding changes in market shares, Section 2.5 provides the theoretical prediction that market shares would be weakly increasing among more efficient units following the introduction of the cap-and-trade program. For a

¹⁹CAISO’s quarterly market monitoring reports for 2014 and 2015 find the overall combined wholesale cost of energy was around (including slightly below, close to equal, and slightly above) their simulated competitive baseline prices under competitive conditions, with some price spikes in a small set of intervals in the second half of 2015 and in Q4 2015 (CAISO 2014, 2015).

descriptive assessment of what has happened in the program, I estimate changes in unit market share by comparing the average share of total hourly demand provided by each unit in 2015 to that provided before the cap-and-trade program started in January 2013. I plot these market share changes by the unit's average heat rate over 2012 - 2015. The graph in Figure 3 shows a negative relationship between market share changes and heat rate, supporting the theoretical prediction, though with a small magnitude.

To connect changes in market share to the spatial distribution of damages from local air pollutants, I plot market share changes by the APEEP estimate of county-specific NO_x damages where the unit is located. The graph in Figure 4 exhibits no clear pattern in the empirical distribution of market share changes and local damages from NO_x . The implication of these two descriptive results is that the program leads to some market share changes in the direction expected, but the lack of obvious trend between market share changes and local damages would provide for minimal co-costs and co-benefits as a result of the program's impact on the distribution of local air quality. Next, I plot the changes in unit efficiencies by the air pollution damage estimate where the unit resides in Figure 5. As with the plot of market share changes and damages, the figure shows no obvious trend with respect to heat rate improvements and the local marginal damages from air pollution.

The next section presents the results from my model which isolates the impact of carbon prices on market outcomes of production and efficiency changes holding other market features fixed and incorporating dynamics in decision making.

5 Results

5.1 Model Fit

The recovered policy functions for dispatch and the estimates of start-up and investment costs are used to simulate market outcomes in counterfactual policy scenarios, as well as to compare the predictions of the model to empirically observed dispatch. I review this later

simulation now. To recall, in the production decision, unit type efficiencies enter exogenously. As such, for the model fit simulation I use the dispatch policy function recovered when the exogenous heat rates correspond to the heat rates observed by unit type in 2015. Then, I simulate hourly market outcomes of generation by unit type for one quarter. Details on the simulation procedure are provided in Appendix 8.1. The results from the simulation are shown in Figure 19 in Appendix 9, which compares the simulated dispatch by unit type to that observed empirically in Q1 2015. The figure shows that the model and simulation procedure perform fairly well in predicting generation output by unit type, and t-tests for market shares by unit type show that the differences in dispatch between the simulated and empirically observed dispatch are not statistically different from zero.

The largest differences between the simulated and empirical market shares are for two unit types that are relatively more expensive, with both higher marginal costs and start-up costs. These units are dispatched empirically but not dispatched in the simulated market. This result could stem from a couple of sources. One, the differences could be driven by inaccuracy in the calibrated start-up costs used for these unit types.²⁰ Alternatively, if the start-up costs are correct, these differences could reflect that the observed dispatch is not strictly cost-minimizing and/or that there are other constraints that the California electric grid operator faces in scheduling units for dispatch that the model does not capture. For example, local transmission constraints could lead more costly units to be scheduled to meet local demand in congested regions and hours.

For an additional review of how the model is performing, I simulate market outcomes without start-up costs, and the results are shown in Figure 21 in Appendix 9. Confirming intuition, in comparison to the simulation with start-up costs, the simulation without start-up costs moves production away from high start-up cost units with relatively lower marginal costs. We also see a reduction in generation among a unit type with lower start-up costs, but with units that have much smaller per unit capacity. The intuition around this result is

²⁰Recall, the calibrated start-up cost is in terms of \$ per MW; multiplying by average unit type MW sizes provides type-specific start-up costs.

that meeting demand with these units requires turning on more individual units, and hence incurring more total start-up costs, while the per unit start-up costs are lower.

5.2 Stringent Carbon Policy Scenario

A climate change policy with a tighter GHG emissions cap requires higher levels of emissions reductions and leads to higher carbon prices.²¹ Accordingly, to evaluate the market outcomes across policy stringencies, I compare production and emissions outcomes across alternative carbon prices including a permit price equal to zero (henceforth, no carbon policy scenario), and a policy that leads to permit prices of \$42 per ton of CO_2e , corresponding to the Environmental Protection Agency’s (EPA) central estimate of the social cost of carbon (SCC) in 2020 with a 3 percent discount rate (henceforth, the SCC carbon policy scenario). The average fuel cost, c^f , over 2012 - 2015 was \$3.6 per MMBtu, and the average permit price over this period translates \$0.70 per MMBtu (\$13 per ton CO_2e). The SCC carbon policy scenario translates to an additional \$2.2 per MMBtu of input cost. As shown in equation 5, in the model the carbon price enters the firm’s profit maximization problem (equivalently, the cost minimization problem) by increasing the cost of inputs. Accordingly, I evaluate market outcomes across three inputs cost states $ic = c^f + \tau$, $\tau = \{0, 0.70, 2.2\}$.

In simulating outcomes across policy stringencies, I leverage the dual decision framework of the model to evaluate the impact of carbon prices together with endogenous efficiency investments that respond to these prices. To do so, at the beginning of the cap-and-trade program I find the cost minimizing investment decision for all unit type efficiencies. Equivalently in the profit maximization framework, units decide whether to invest to improve their efficiency or not, assuming that all other unit types make optimal investment choices. Investment decisions are made by finding the investment portfolio that is cost minimizing with respect to a market outcomes over a time period $T =$ three years. If the California market is stationary, then some T provides a reasonable comparison of the value functions across

²¹Recall that in California the cap-and-trade program is broader than the electricity industry, so that equilibrium permit prices are determined by abatement costs across all regulated sectors.

investment portfolios, and this decision will satisfy equation 10, corresponding to the decision to maximize value over an infinite time horizon. However, I consider that the California market will evolve over longer time horizons, for example, from increased renewable energy penetration, and I view this modeling choice as consistent with the near- and medium- term objectives of this paper’s analysis. Further, I view this choice as reasonable in the context of firm decision making in this setting, noting that this assumes that firms make efficiency investment decisions myopic to outcomes after T .

To simulate endogenous investment at alternative carbon prices I perform repeated simulations of investment costs, $\gamma j_i + v_i$, over a grid of γ developed based on the central estimates of the SNL data and the SMM approach.²² I average investment costs for each evaluated investment portfolio over the simulations and then choose the investment portfolio in each carbon policy that minimizes costs. Figure 6 shows how the quantity of generation provided by each unit type in the current and SCC carbon policy scenario compares to a no carbon policy scenario including endogenous investment. The current carbon prices lead to minimal production reallocation across unit types. This result corresponds to minimal reallocation of local air pollutants, and minimal co-benefits (and co-costs), as a result of the policy. The higher carbon price in the SCC carbon policy scenario leads to more noted production reallocation. In particular, this scenario leads to an increase in market share among lower marginal and higher start-up cost units and a decrease in higher marginal lower start-up cost units by around 1.5 percent each. This reallocation confirms our theoretical predication, as it corresponds to an increase among units with relatively higher start-up costs and lower marginal costs.

Figure 7 shows how the changes in production allocation in the SCC carbon policy scenario compare to a no carbon policy scenario map to changes in NO_x damages. This calculation is done by connecting changes in generation by unit type to the location and local marginal NO_x damages for units in that type. The SCC carbon policy reduces NO_x

²²In particular, I simulate investment costs starting with 2 million and increasing by steps in 0.025 million intervals, and average investment costs over 500 simulations.

damages by around \$1 million over one quarter. The reductions are driven by the market shares changes shown in Figure 6.

To understand which communities are receiving the benefits in NO_x damages avoided, I overlay changes in unit NO_x emissions on the CA Air Resources Board (ARB) map of disadvantaged communities (DAC) in California. Of the over 8,700 Census tracts in California, the ARB identifies 2,005 tracts as disadvantaged if they are in the top 25 percent of CalEnviroScreen 3.0 scores, as well as census tracts in the highest 5 percent of CalEnviroScreen's Pollution Burden without an overall CalEnviroScreen score. The CalEnviroScreen Model uses indicators of pollution burdens and population characteristics to develop a numeric rating for each census tract reflecting the relative pollution burdens and vulnerabilities of the tract compared to the rest of California (OEHHA and CalEPA 2017). I find that the majority of the co-benefits from avoided NO_x damages occur outside of the state's heavily polluted regions, with only 30 percent of the co-benefits occurring in the DACs. Similarly, an OLS regression of change in emissions of NO_x on the local damages from NO_x , denoted ϱ , fails to reject the hypothesis that $\frac{\partial}{\partial \varrho} \frac{\partial NO_x}{\partial \tau} = 0$. That is, the benefits are not occurring in regions with systematically higher or lower marginal damages from air pollution. The results of this regression across different carbon price scenarios are provided in Table 7.

5.3 Command-and-Control Efficiency Scenario

Minimum efficiency standards and other command-and-control type of policy mechanisms are frequently encountered in the energy sector as alternatives to market-based GHG regulation. For example, the Corporate Average Fuel Economy (CAFE) standards for vehicles and the U.S. National Ambient Air Quality Standards (NAAQS) for stationary sources establish maximum thresholds of pollution intensities. In this section I consider the impact of this type of policy in this setting by evaluating the private and social returns from alternative portfolios of efficiency improvements. The portfolios could stem from a policy that, for example, established minimum carbon intensity standards or subsidized efficiency investments

for certain types of units.

To compare alternative investment scenarios, I evaluate the set of heat rate vectors that would result from different combinations of unit type investments. I group the units into five types based on heat rates and the previous unit type grouping, and I will refer to these five types as investment unit types.²³ I allow each investment unit type to choose to invest or not, where investment results in a 1.5 percent reduction in heat rate, which is the average heat rate reduction observed in the set of units identified to invest over the 2013 - 2015 time frame.²⁴ The result is 2^5 alternative investment decision scenarios j , mapping to $J = 32$ different 10 by 1 vectors of unit type heat rates where $j = 0$ corresponds to no investment and $j = 31$ corresponds to all units investing.²⁵ The policy iteration that was used in the alternative carbon policy counterfactuals provides the dispatch policies for each of these investment scenarios, and here I evaluate market outcomes simulated over three years for each $j \in J$ investment scenarios. I sum the discounted costs incurred across each j investment scenario, averaged over simulations S , $\frac{1}{S} \sum_{t=1}^T \beta^{t-1} G(\mathbf{q}_t^{\mathbf{j}^*} | \omega^{\mathbf{j}})$, where $\mathbf{q}^{\mathbf{j}^*}$ is determined by the recovered policy function for investment scenario \mathbf{j} .

Figure 9 plots the savings in production costs for each of the alternative investment scenarios compared to no investment (henceforth, net investment returns) across carbon price scenarios. Changes in net investment returns are driven by two mechanisms, which I review now in the framework of the cost minimization problem. One, the investment could change a preference for one unit type over another due to the improvement in efficiency (reduction in marginal costs). Two, even if the efficiency improvement does not change preferences across units, and market shares stay the same following the efficiency improvement, private and social costs would still be reduced if some of the improved units have non-zero market share.²⁶ The figure demonstrates that, as expected, the highest net returns to investment

²³I collapse the ten unit type grouping into five unit type groups by sorting on heat rate.

²⁴Details on the procedure to identify investment are provided in Appendix 8.5

²⁵The 5-type grouping is used only for the investment decision; then, for the purposes of dispatch, the investment decision is mapped onto the 10-type grouping used earlier.

²⁶Whether or not investment in efficiency improvement changes the ranking of units in costs depends on both the quantity of efficiency improvement evaluated, as well as the distribution of efficiencies. Here all

come under high carbon price scenarios. This illustrates the theoretical point made in Section 2.5, that the carbon price increases the returns to increasing unit efficiency, since the carbon price enters the firm’s profit maximization problem (equivalently, the cost minimization problem) as an increase in the cost of inputs. The expected returns to investment can be seen as a measure of willingness to pay for efficiency improvements, and the figure shows that higher carbon prices increase firms’ willingness to pay, and thus increase the likelihood of investment.

We also see variation in net investment returns across investment scenarios within a given carbon price state. Figure 10 plots net investment returns on the y-axis for of each j investment scenario on the x-axis, differentiated by carbon price scenario. The x-axis orders investment scenarios by the number of unit types that invest, with $j = 0$ and $j = 31$ corresponding to no and all units investing, respectively. Given this ordering, this figure demonstrates that net investment returns are not strictly increasing in the number of units that invest. Rather, a key determinant of net returns is the market share of the unit type(s) investing. For example, compare the results from scenario six to seven. In scenario six, four unit types invest, and in scenario seven, those same four unit types invest, as well as two units types that are dispatched frequently pre-investment, which leads to a large increase in returns to investment.

Across the scenarios, improving the efficiency of frequently dispatched units leads to larger net investment returns compared to improving the efficiency of higher-cost, less-frequently dispatched units as these units are utilized less. Figure 11 plots net investment returns for an investment scenario on the x-axis by the market share of the unit type(s) that was simulated to invest in the respective scenario. The positive slope reflects an increase in returns when scenarios correspond to simulated investment from units with higher market share, demonstrating the second theoretical prediction made in Section 2.5. This result has

efficiency improvements improve the heat rate by 1.5 percent; whether or not such an improvement changes preferences across units for dispatch depends on how the unit that improves compares to the units close to it in terms of costs.

important implications for environmental regulation, especially in settings where minimum efficiency standards have been used to meet policy objectives. The findings here contradict the claim that policy should focus on improving the dirtiest, least efficient capital to reduce pollution. Rather, here we see larger gains from improving the lower-cost, more frequently utilized capital, which in this setting corresponds to the relatively cleaner units. Note, the results here assume the functional form of investment costs as described in Section 2.1; Section 2.5 shows the conditions on market shares and investment costs under which this result would hold elsewhere.

5.4 Local Air Pollution Tax Scenario

Next, I simulate market outcomes in a counterfactual scenario with a location-specific Pigouvian tax on local air pollution. In this scenario, units make production decisions as earlier, but with an additional component included in marginal costs corresponding to location-specific compliance costs from the local air pollution tax. The level of the tax is set to the APEEP \$ per ton damage estimate for NO_x emissions in the unit's county k , and enters marginal costs in the following way:²⁷

$$mc_{ik} = \omega_i(c^f + e^f \tau^{ghg}) + \iota_i \cdot \tau_k^n, \quad (12)$$

where ι_i is the unit's NO_x emissions intensity (ton per MWh produced), which is observed from the CEMS data, and τ_k^n is the tax on NO_x in county k , distinct from the carbon price, τ^{ghg} . To operationalize this scenario, I again leverage the correspondence between the competitive equilibrium outcomes and the solution to a cost minimization problem. In this scenario, the market outcomes are equivalent to the solution to a cost minimization problem that now includes the marginal damages from local air pollution in addition to the carbon

²⁷As the quantity of SO_2 emitted by these units is small, I evaluate a tax on NO_x only.

price.

To solve this new cost minimization problem, I first re-group units into new unit type groups. The new unit type groups are needed as the unit type needs to include all characteristics relevant for the cost minimization problem, which now includes the local marginal air pollution damages where the unit is located. The procedure to establish this new unit type grouping is described in detail in Appendix 8.6. Then, I use the process described earlier to recover policy functions for dispatch across J scenarios of investment with the new unit type groups. I use the previously estimated investment costs to make an optimal investment choice at the start of a three-year dispatch period and then simulate market outcomes of production and emissions given this investment decision.

Figure 12 illustrates how this change in the functional form of marginal costs from the introduction of the tax impacts the ranking of the unit types by marginal costs, and resultant, market shares by unit type. The graph on the left of Figure 12 ranks the unit types by marginal costs before and after the tax. The change in marginal cost ranking drives the changes in market shares shown on the panel on the right. As before, the market share changes drive changes in the distribution of NO_x and thus NO_x damages. I sum NO_x damages in this scenario and compare them to the case where units only pay a carbon price, where the carbon price state corresponds to the current average observed price.

The simulated market outcomes in this counterfactual show a reduction in NO_x damages by \$8.3 million over one calendar quarter of hourly market outcomes, compared to a carbon policy only scenario at currently observed carbon prices.²⁸ To see which communities receive these benefits, I again overlay changes in unit NO_x emissions on the CA Air Resources Board (ARB) map of disadvantaged communities in California. I find that the local air pollution tax concentrates benefits in the disadvantaged communities, with 89 percent of the \$8.3 million saved occurring from NO_x emissions reductions in units located in disadvantaged

²⁸ NO_x damages in the scenario with observed carbon prices and no local air pollution tax are \$34.3 million over one calendar quarter, compared to damages with a carbon price and a local air pollution tax of \$26.0 million.

communities.

6 Conclusion

The results from the simulations present several important findings. Principally, I find that the structure of the fossil-fuel electric portfolio in terms of marginal costs, start-up costs, and the location of the units, does not provide for a systematic redistribution of local air pollutants under a carbon pricing policy at current prices. Higher prices provide some redistribution of air pollution and some co-benefits from local air quality improvements, but the benefits occur predominantly outside of heavily polluted regions. On one hand, the results indicate that siloing climate change policy from local air pollution regulation does not lead to significant co-costs from the redistribution of local air pollutants. On the other hand, they show that the GHG regulation provides minimal improvements in local air pollution in the disproportionately polluted regions of California. To improve the air quality of these heavily polluted regions, other regulatory instruments will need to be leveraged. The counterfactual simulation demonstrates one such instrument, a location-specific tax on local air pollution. I find that combining the carbon policy with the local air pollutant tax provides double the benefits from avoided NO_x damages compared a carbon policy alone. Further, these savings occur predominantly from emissions reductions in areas that bear relatively higher pollution burdens in the state.

These findings also provide insights for other industries with dynamic production decisions. First, they suggest that industries where fixed costs are small compared to marginal costs are less exposed to market share changes following an increase in input costs, a result which has important implications for analyzing production and emissions leakage potential from environmental regulation in other industries. Second, they are instructive in evaluating alternative investments in efficiency. The imposition of minimum efficiency standards is a common tool in environmental and air quality regulation, yet in this setting, such a standard

would lead to far fewer savings in production costs and NO_x damages avoided as compared to a regulation that improved the efficiency among units with lower pollution intensities. This finding is driven by a key characteristic of the firms studied here — for the portfolio of natural gas units, pollution externalities are decreasing in production efficiency, and this finding may be generalizable to other industries with this feature.

References

- Severin Borenstein. The redistributive impact on nonlinear electricity pricing. *American Economic Journal: Economic Policy*, 4(3):56–90, 2012.
- Severin Borenstein. Effective and equitable adoption of opt-in residential dynamic electricity pricing. *Review of Industrial Organization*, 42(2):127–160, 2013.
- Severin Borenstein. Private net benefits of residential solar PV: The role of electricity tariffs, tax incentives, and rebates. *Journal of the Association of Environmental and Resource Economists*, 4(S1):S85–S122, 2017.
- Severin Borenstein and Lucas Davis. The distributional effects of U.S. clean energy tax credits. *NBER Tax Policy and the Economy*, 30(1):191–234, 2016.
- CAISO. Quarterly report on market issues and performance. *Department of Market Monitoring*, 2014, 2015.
- Joseph Cullen. Dynamic response to environmental regulation in the electricity industry. *Working Paper*, 2015.
- Joseph Cullen and Stanley Reynolds. Market dynamics and investment in the electricity sector. *Working paper*, 2017.
- Lara J. Cushing, Madeline Wander, Rachel Morello-Frosch, Manuel Pastor, Allen Zhu, and James Sadd. A preliminary environmental equity assessment of california’s cap-and-trade

- program. *Program for Environmental and Regional Equity, University of Southern California Dornsife*, 2016.
- Meredith Fowlie, Stephen P Holland, and Erin T Mansur. What do emissions market deliver and to whom? Evidence from Southern California NOx trading program. *American Economic Review*, 102(2):965–993, 2012.
- Don Fullerton, Garth Heutel, and Gilbert E. Metcalf. Does the indexing of government transfers make carbon pricing progressive? *American Journal of Agricultural Economics*, 94(2), 2012.
- Lawrence H. Goulder, Marc A. C. Hafstead, Gyurim Kim, and Xianling Long. The incidence of carbon taxes in u.s. manufacturing: Lessons from energy cost pass-through. *Working Paper*, 2018.
- Hugo A. Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150, 1992.
- Boyan Jovanovic. Selection and the evolution of industry. *Econometrica*, 50(3):649–670, 1982.
- Christopher Knittel and Ryan Sandler. The welfare impact of second-best uniform-pigouvian taxation: Evidence from transportation. *American Economic Journal*, 10(4):211–242, 2018.
- N. Kumar, P. Besuner, S. Lefton, D. Agan, and D. Hilleman. Power plant cycling costs. *National Renewable Energy Laboratory*, 1992.
- Kyle Meng. Is cap-and-trade causing more greenhouse gas emissions in disadvantaged communities? *Working paper*, 2017.
- Nicholas Z. Muller and Robert Mendelsohn. Measuring the damages of air pollution in the

united states. *Journal of Environmental Economics and Management*, 54(1):1–14, 2007.
ISSN 00950696.

OEHHA and CalEPA. CalEnviroScreen 3.0 Factsheet. *CA.gov*, 2017.

Jr. Robert E. Lucas and Edward C. Prescott. Investment under uncertainty. *Econometrica*, 39(5):659–681, 1971.

7 Tables and Figures

Table 1: Fossil Unit Summary Statistics, CA 2012-2015

	2012	2013	2014	2015
Units producing	221	197	207	201
Steam Turbine	50	41	39	37
Gas Turbine	90	85	87	87
Combined Cycle	81	71	81	77
Natural Gas	221	193	207	201
Coal	0	4	0	0
Retired	2	0	0	1
Put in Service	11	26	1	0
Mean Capacity MW	139	160	134	136
Total Capacity GW	30.6	31.5	27.8	27.2
Num. Units with Capacity Change Up		5	11	7
Mean MW Capacity Up		4	7	7
Num. Units with Capacity Change Down		5	6	9
Mean MW Capacity Down		10	2	4
Mean Heat Rate (Btu per KWh)	14318	12797	14046	12244
Prct of Hours Operating	.35	.31	.35	.35

Table 2: Unit Characteristics by Type

Type Num.	Num. Units	Size MW	2012 HR	MC Rank	Start-up Cost*	Start-up Cost Rank
1	7	121	7308	1	9680	8
2	9	145	7565	3	11600	9
3	7	94	12783	8	7520	4
4	13	95	13567	10	7600	5
5	31	170	7362	2	13600	10
6	22	74	10535	5	5920	1
7	10	76	9911	4	6080	2
8	23	107	12823	9	8560	7
9	31	90	10543	6	7200	3
10	30	105	11889	7	8400	6

(*) Using calibrated estimate of \$80 per MW

Figure 1: This figure shows the impact of the carbon price on an illustrative supply curve for a given hour t . Following the carbon price, the portion of the supply curve comprised of natural gas units shifts up and increases in slope from mc_t to mc_t' .

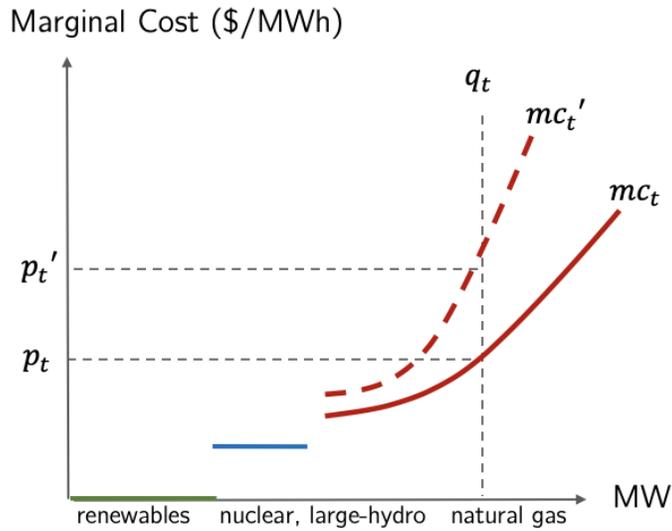


Figure 2: This figure plots average profits per hour without start-up costs, i.e. $q(P - mc)$, for natural gas units on the y-axis, over hour of the day on the x-axis. It illustrates that in the early hours of the morning some units operate with $P < mc$, which allows them to avoid shutting down and incurring additional start-up costs later.

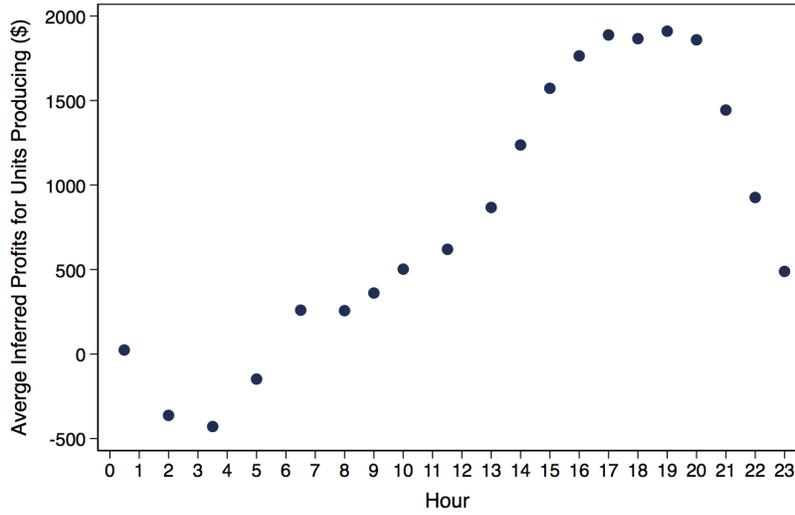


Figure 3: This figure compares market shares, ζ_i , pre- and post- the carbon price (2012 to 2015) on the y-axis, over the average unit heat rates, ω_i , over the same period on the x-axis. It shows some evidence that the pattern of market share changes support the theoretical prediction, $\frac{\partial \zeta_i}{\partial \omega_i} < 0$.

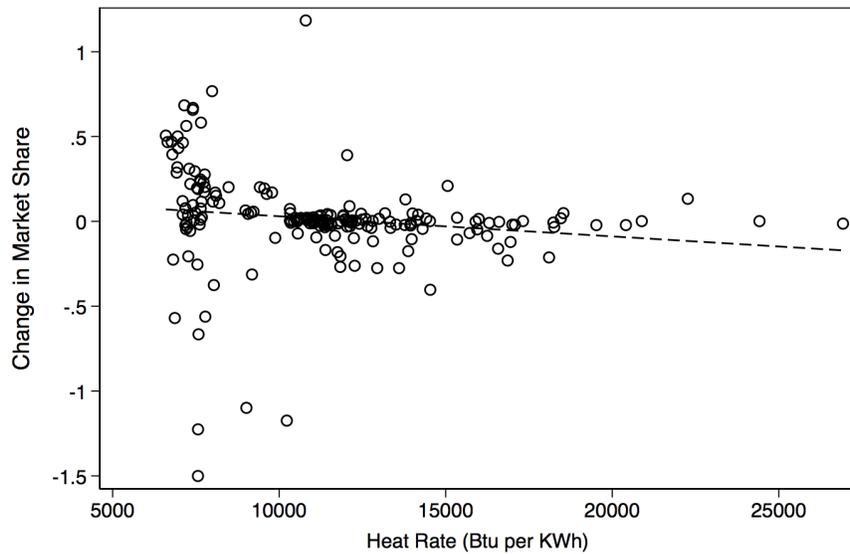


Figure 4: This figure plots each unit's change in market share pre- and post- the carbon price (2012 to 2015) on the y-axis, by the APEEP estimate of marginal damages from NO_x on the x-axis. The figure shows no clear pattern in the empirical distribution of unit market share changes and local marginal damages from air pollution.

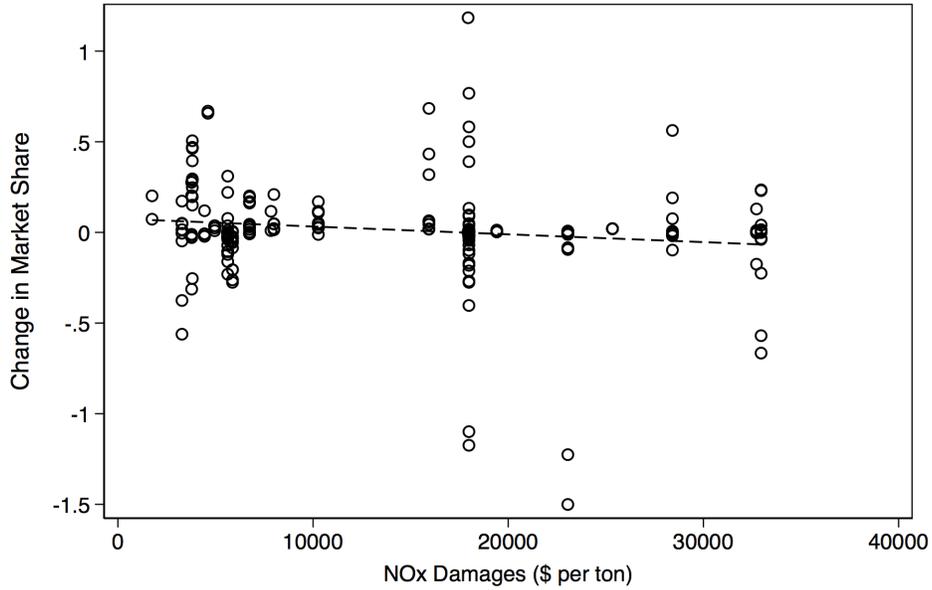


Figure 5: This figure plots each unit's change in efficiency pre- and post- the carbon price (2012 to 2015) on the y-axis, by the APEEP estimate of marginal damages from NO_x on the x-axis. The figure shows no clear pattern in the empirical distribution of unit efficiency changes and local marginal damages from air pollution.

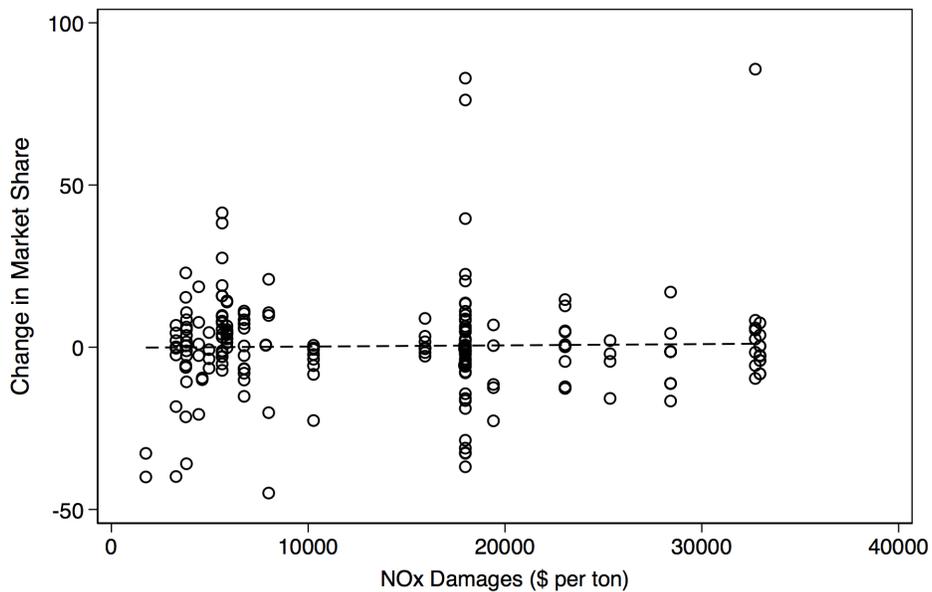


Figure 6: This graph shows the change in generation provided by unit type at current and high carbon price scenarios, compared to generation provided in a no carbon price scenario.

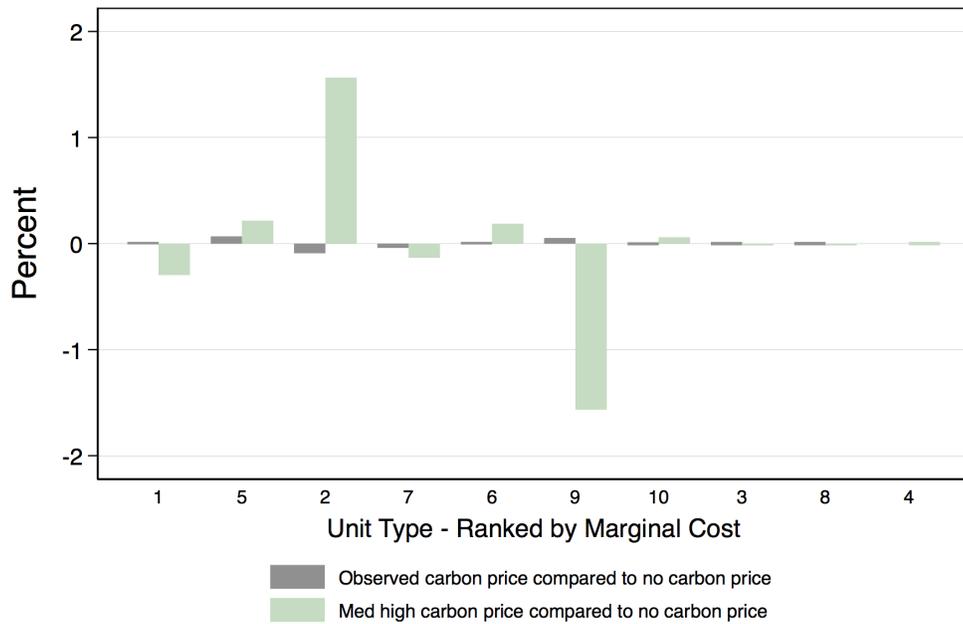


Figure 7: This graph shows the change in NO_x damages by unit type at current and high carbon price scenarios compared to a no carbon price scenario. Damages are summed over one calendar quarter of market outcomes.

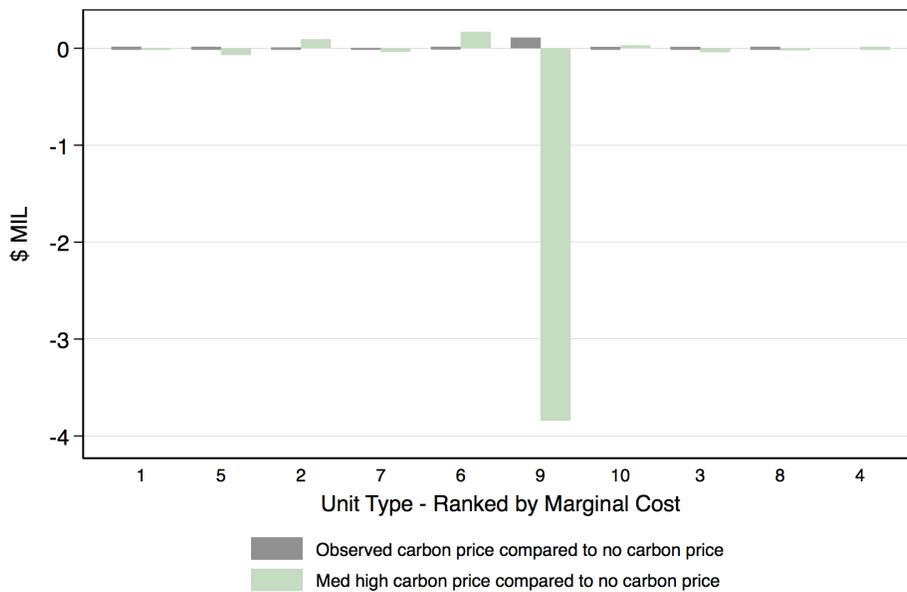


Figure 8: This figure shows changes in market share under the high carbon price scenario compared to no carbon price scenario, plotted by the cost characteristics of the unit type. The size of the circles represent the relative size of market share change in a unit type, with teal reflecting an increase and red a decrease. The circles are plotted according to the unit type's marginal costs on the y-axis and fixed-costs (start-up costs) on the x-axis.

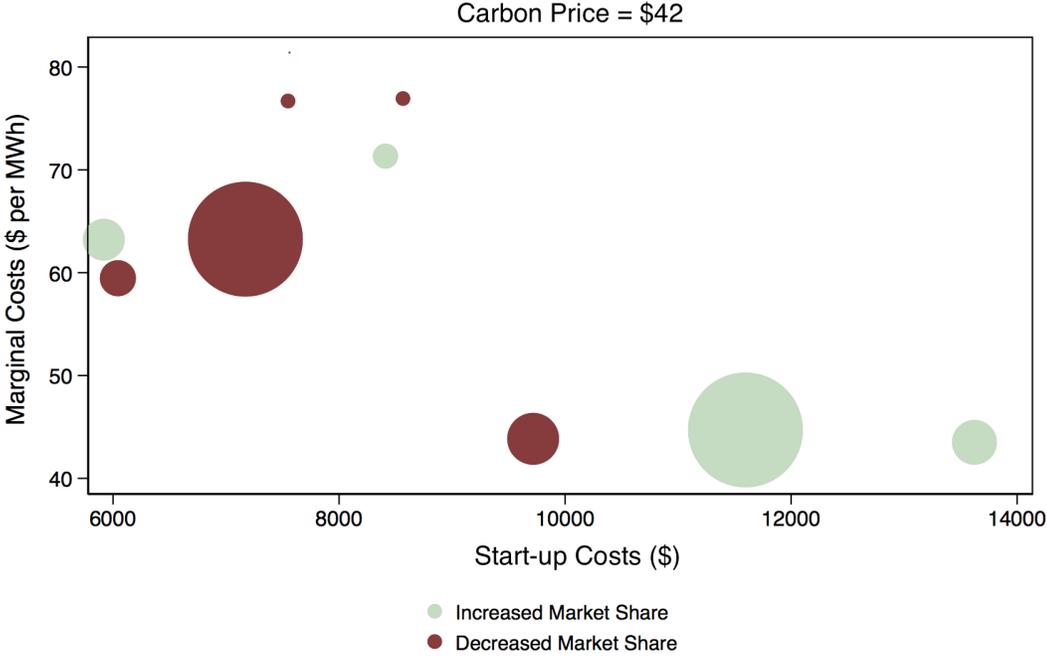


Figure 9: This figure plots the net returns from investment compared to no investment across four carbon price scenarios, where each symbol represents one investment scenario.

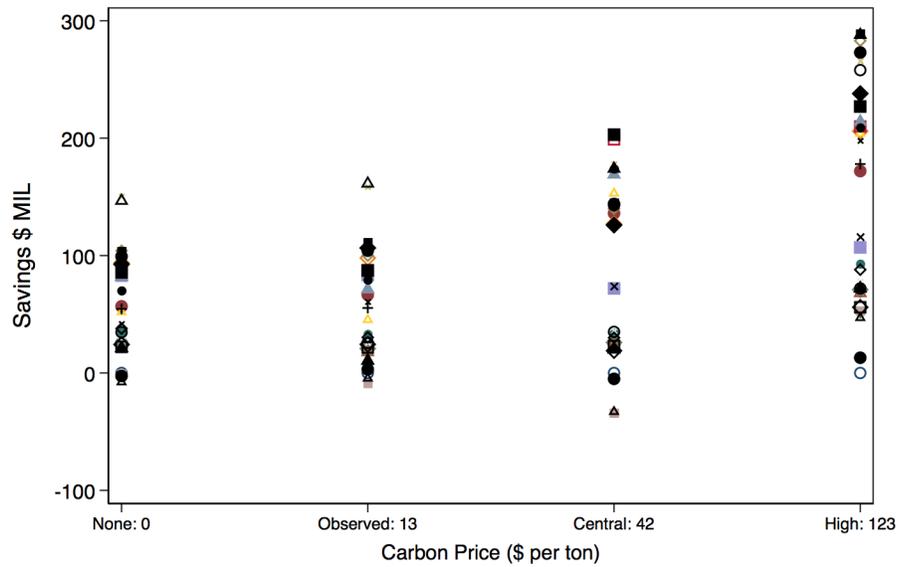


Figure 10: This graph plots the savings in production costs with optimal dispatch across 31 investment scenarios compared to no investment. That is, 3-yr production costs in investment scenario j are compared to investment scenario 0, which sets heat rates to average observed heat rates in 2012. Savings are shown across 4 different scenarios of carbon pricing, corresponding to scenarios with no carbon price, carbon price equal to average price observed 2012 - 2015, and carbon price equivalent to a \$42 and \$123 social costs of carbon.

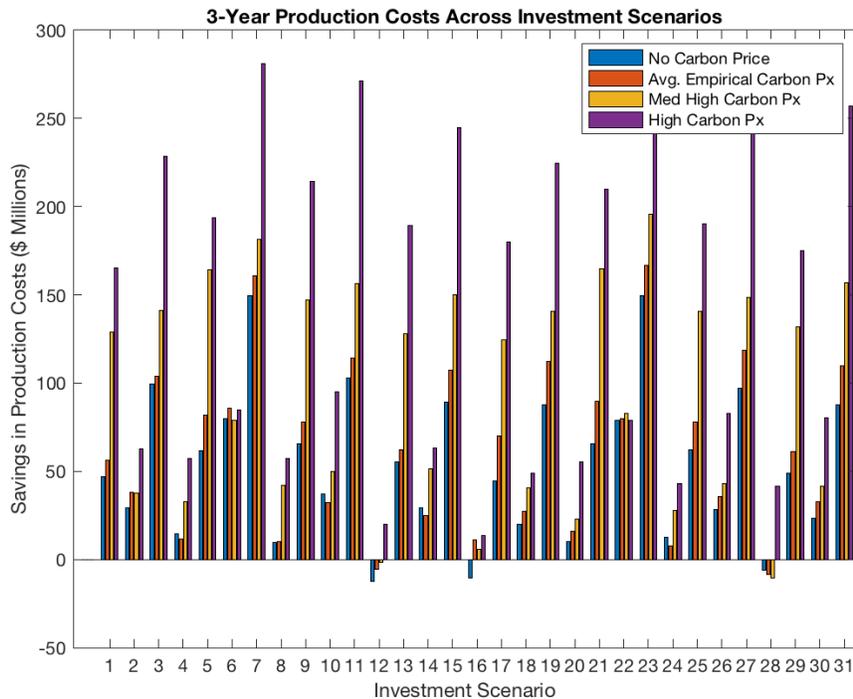


Figure 11: This figure plots the net returns in a given investment scenario on the x-axis by the market shares of the unit types simulated to invest on the y-axis. The positive slope illustrates higher returns from investment scenarios in which unit types with higher market shares invest.

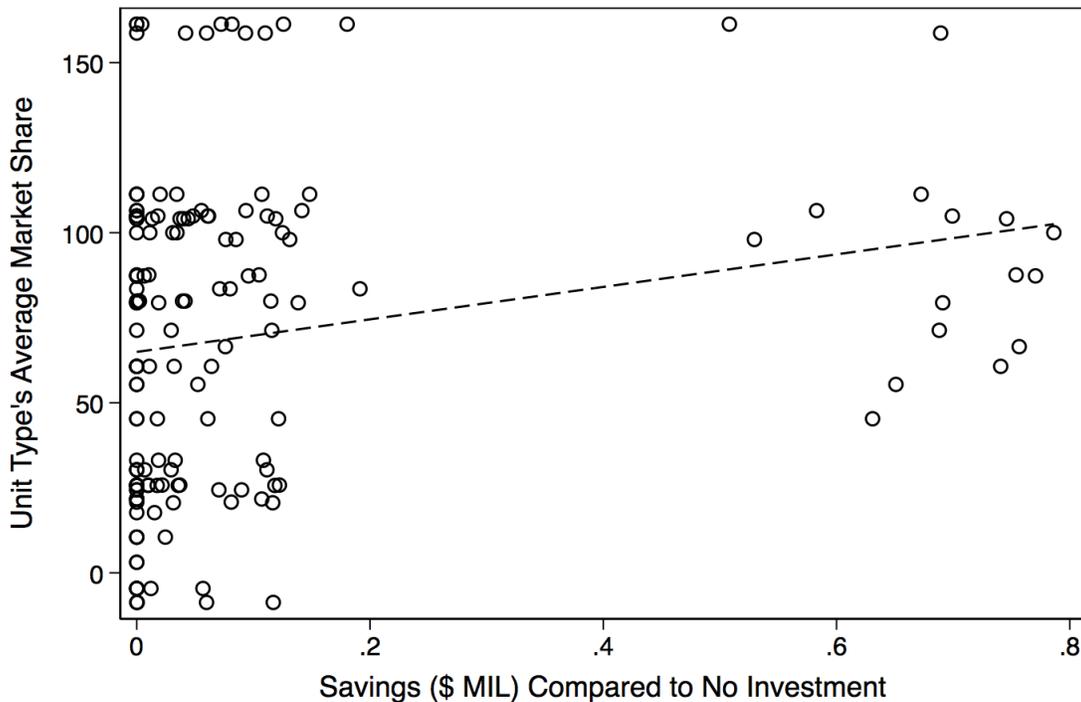


Table 3: Distribution of Benefits by Census Tract Type

	Num. Tracts	Pct. with Benefits Obs. Carbon Price, Tax	Pct. with Benefits SCC Carbon Price, No Tax
Disadvantaged Tract	2,005	.89	.33
Non-Disadvantaged Tract	6,777	.11	.67
Total Tracts	8782	1.0	1.0

Benefits are compared to scenario with no carbon policy and no local air pollution tax.

Figure 12: This figure on the left shows how the scenario with a local air pollutant tax in addition to the carbon policy changes marginal costs across unit types compared to a carbon policy alone. The figure on the right shows the changes in market share in the carbon price plus tax scenario compared to a carbon policy alone.

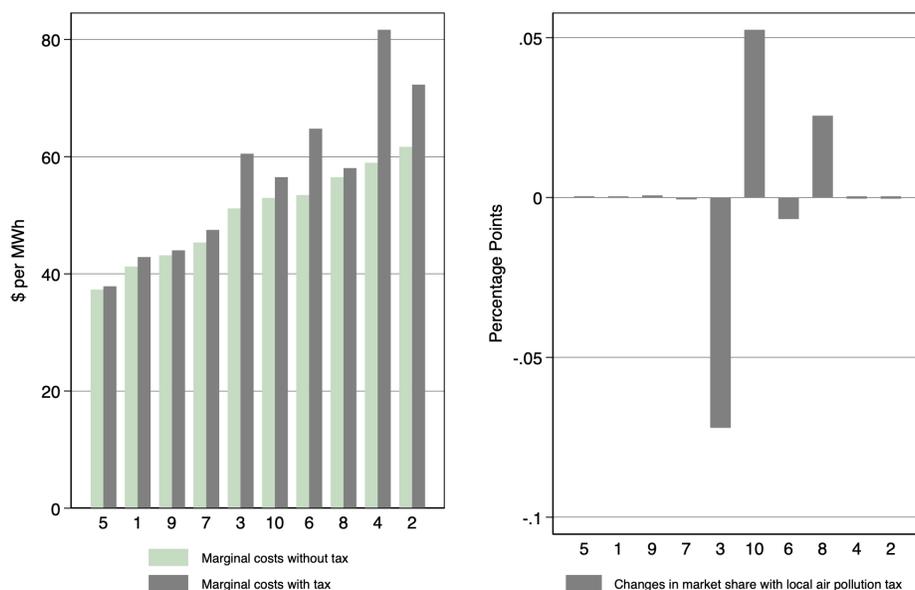


Table 4: Change in Quantity of Pollutants Emitted on Pollutant's Local Marginal Damages

	Observed Carbon Px	Med. High Carbon Px	High Carbon Px	Observed Carbon Px	Med. High Carbon Px	High Carbon Px
NOx \$ per ton	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)			
SO2 \$ per ton				0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
R-squared	0.007	0.003	0.001	0.016	0.016	0.020
N	182	182	182	182	182	182

Standard errors shown in parenthesis. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

8 Appendix

8.1 Simulation Procedure

This section outlines the simulation procedure used to generate market outcomes. The first simulation presented in the paper is done for model fit. This simulation uses the recovered policy functions $\sigma^*(\cdot)$ to forward simulate market outcomes over one calendar quarter. For this simulation, the initial heat rate vector used corresponds to average heat rates by unit types in 2015, ω^{2015} , and the results from the simulation are compared to average hourly market outcomes empirically observed for Q1 2015, shown in Appendix 9.

Simulation for Model Fit

1. For a given initial state, set $t=0$, and draw demand shock η_t .
2. Let $\omega = \omega^{2015}$ and let ic = empirically observed ic in 2015 Q1.
3. Use $\sigma^*(\eta_t, \mathbf{a}_0, ic, h_t | \omega)$ to find optimal dispatch for this state \mathbf{q}_t^* .
4. Simulate η_{t+1} as the next step of a Markov chain with transition matrix $F\eta$.²⁹
5. Update the elements $a_{t+1} \in \mathbf{a}_{t+1} = 1$ when elements $q_t \in \mathbf{q}_t > 0$ and 0 otherwise.
6. Update $h_{t+1} = h_t + 1 - \mathbb{1}(h_t = 24) * 24$.
7. Update $\mathbf{q}_{t+1} = \sigma(\eta_{t+1}, \mathbf{a}_{t+1}, ic, h_{t+1})$.
8. Update $t = t + 1$ and repeat until $t = 7160$.

Simulation of Production Costs for three-year Market Outcomes

For the simulations that calculate production costs over three years, the procedure above is

²⁹ $F\eta$ is estimated from the empirical distribution of demand shocks. The Markov chain simulation proceeds by drawing a random variable η with the distribution according to the row in $F\eta$ corresponding to η_t . Then, the discrete inverse transform method is used to simulate η_{t+1} , which involves drawing a random variable u from a uniform distribution and comparing to the probability entries in row η_t of $F\eta$.

used to generate quarterly outcomes. For each outcome, total production costs are calculated as the sum of marginal costs and start-up costs based on the units that are dispatched in each hourly period. The production costs are then stacked over three years, using a quarterly discount rate of 0.0034, which corresponds to an annual interest rate of 1.34 percent. The policy functions are recovered for four different input cost scenarios corresponding to average fuel costs plus no carbon price, a carbon price equal to the average price observed 2012 - 2015, and a carbon price equivalent to a \$42 and \$123 social cost of carbon; in other words setting the permit price $\tau = \{0, 13, 42, 123\}$ in dollar per ton CO_2e . The policy scenario-specific estimates of production costs use the optimal dispatch decisions indicated by the recovered dispatch policy function $\sigma^*(\cdot|ic)$ for the respective ic scenario.

Simulation of NO_x Damages

To connect market outcomes to damages to human health, the NO_x and SO_2 damages associated with the simulated market outcomes are estimated.³⁰ In the simulation above, the number and type of units that are dispatched in each hour are recorded, y_t^{sim} . Each of the 10 unit types represent a set of units for which I observe each of their locations in longitude and latitude, the NO_x and SO_2 emissions intensities from the unit's continuous emissions monitors (CEMS), and the local marginal damages of air pollutants based on APEEP estimates for the unit's county. I allocate the total production assigned to the unit type in the simulation across the units of that unit type equally. That is, if there are seven units in type one, and type one is dispatched to produce 700 MWh of electricity in a given hour, each unit is allocated 100 MWh of production in that hour. Then, I multiply each unit's allocated production by its emissions intensity for NO_x and SO_2 respectively. Finally, I estimate damages to human health by multiplying unit emissions by county- and pollutant-specific APEEP damage estimates.

³⁰As these units emit small amounts of SO_2 , the SO_2 results are not reviewed.

8.2 Estimation of Investment Costs

The SNL Energy Platform has data on gross capital and fixed production costs for a subset of power plants. This subset of power plants corresponds to 14 (28 percent) of the units flagged as investing based on the criteria outlined in Appendix 8.5. I use these data together with the heat rate improvements observed for these units over 2012 to 2015 to construct an estimate of mean investment costs per percent heat rate reduction of 2.3 million, with a standard deviation of 3.1 million.

I also estimate investment costs using a simulated method of moments approach (SMM) as discussed in the text. I recover the dispatch policies across J investment scenarios as described earlier. Then, I use the procedure described in A.1 to forward simulate market outcomes, and I calculate the total production costs associated with three years of hourly market outcomes for each scenario, V^j . Next, I draw an initial investment cost γ^0 , and I select the optimal investment policy based on the simulated production costs, V^j , and the investment costs associated with each scenario $\Gamma(\mathbf{j}, \mathbf{v}, \gamma)$:

$$\mathbf{j}^*(\gamma^0) = \arg \max_{j \in J} (V^j + \Gamma(\mathbf{j}, \mathbf{v}, \gamma^0)). \quad (13)$$

As in the cost minimization problem described in the text, \mathbf{j} and \mathbf{v} refer the vector of investment decisions and cost shocks for each unit type, respectively. The investment costs $\Gamma(\mathbf{j}, \mathbf{v}, \gamma)$ are calculated $\sum_{i=1}^N j_i \gamma + v_i$, where $j_i \in \{0, 1\}$; γ denotes costs of investing to improve heat rate by 1.5 percent; and as in the text, v_i is an investment cost shock drawn from an extreme value distribution type 1 with location and scale parameters estimated from the SNL data. The investment cost shocks are type-specific, drawn for each input cost state, and known to the firms before making investment decisions.

Next, I use the data to estimate the probability of investment across c different unit investment types.³¹ I use this probability to simulate S investment decisions, corresponding

³¹The use of subscript c here instead of n , which in the text corresponds to unit types, is to distinguish between unit investment type groups and unit type groups. The estimation approach collapses the 10 unit

to c -length vectors capturing investment decisions for each unit type. \mathbf{j}_{sim} denotes the matrix with c rows and S columns of simulated moments. Let $g(\cdot, \gamma^0) = (\mathbf{j}_{\text{sim}} - \mathbf{j}^*(\gamma^0))^2$, a matrix with entries corresponding to the squared deviations from the simulated investment moments and the investment choice made by solving equation 13 given investment costs γ^0 . Denote $M = c * S$, which corresponds to the total number of moments. I reshape $g(\cdot, \gamma^0)$ into a M -sized vector of moments, and estimate $\hat{\gamma}$:

$$Q(\gamma) = g(\cdot, \gamma)' \hat{W} g(\cdot, \gamma) \tag{14}$$

$$\hat{\gamma} = \arg \min_{\gamma \in \Theta} Q(\gamma),$$

where Θ is the set of positive real numbers, and \hat{W} is estimated as $(g(\hat{\gamma})g(\hat{\gamma})')^{-1}$. Using this approach, I find that estimates of $\hat{\gamma}$ are sensitive to the initial condition, but in a range comparable to the SNL data, $\hat{\gamma} \in \{2, 9\}$ in \$ million. Accordingly, when using these estimates to simulate endogenous investment, I create a grid of $\hat{\gamma}$ based on the estimates from the simulated method of moments approach and the SNL data. Then, I simulate investment costs at each grid point by taking the sum of the value at that grid point and a draw of \mathbf{v} from an extreme value type 1 distribution with scale parameters estimated from the SNL data. I repeat this for 500 simulations for each grid point. Then, the investment decision evaluates equation 13 using an estimate of $\hat{\gamma}$ equal to the mean of the simulated investment costs.

8.3 Estimation of Start-up Costs

The estimation of start-up costs compares production (dispatch) decisions implied by the model for a given start-up cost κ^0 to empirically observed dispatch. The parameter estimate, $\hat{\kappa}$, is a scalar value of \$ per MW size of the unit, which maps to type-specific start-up costs

type groups to five unit investment type groups for the purposes of investment decisions for computational speed.

through the variation in unit size. The simulation procedure in section 8.1 uses the recovered policy function for dispatch $\sigma^*(\cdot)$ to simulate market outcomes over the Q1 2015 period using 2015 average heat rates. Let the N -length vector \mathbf{q}^* capture the outcomes for all unit types where q_n^* is the dispatch outcome for a single unit type n .

The empirical counterparts are assembled by categorizing each period in the data by the discretized state variables used in the model which include demand shock, lagged operating state, input cost, and hour: $s = \{\eta, \mathbf{a}, ic, h\}$. In constructing the set of moments, note that dispatch outcomes are not observed empirically for all states in the model. Further, to simplify the linking of lagged states in the model to their empirical counterparts, the moments used are for states where all units were either on or off in the last period. Denote S as the number of states used for moments. I assemble N -length vectors corresponding to empirically observed dispatch by unit type in each state, $\mathbf{q}^e(s)$. I construct a S -length vector of moments $g(s, \kappa^0) = \sum_{i=1}^N (\mathbf{q}^*(s, \kappa^0) - \mathbf{q}^e(s))^2$, that is, the sum of deviations for all unit types in a given state s . I estimate $\hat{\kappa}$ as:

$$\begin{aligned} Z(\kappa) &= g(s, \kappa)' \hat{W} g(s, \kappa) \\ \hat{\kappa} &= \arg \min_{\kappa \in \varkappa} Z(\kappa), \end{aligned} \tag{15}$$

where \varkappa is the set of positive real numbers, and \hat{W} is estimated as $(g(s, \hat{\kappa})g(s, \hat{\kappa}))^{-1}$.

8.4 Unit Type Groups

This section explains the process to group units into unit types. The goal is to group the units into a computationally tractable number of unit types, where groups are determined by unit characteristics that are relevant to the cost minimization problem. The cost minimization problem is solved using a vector of unit sizes and efficiencies as inputs. Accordingly, I use k-means to cluster the units that provide non-zero generation to California in 2015 into

groups according to unit size and efficiency (heat rate). I perform clustering with alternative numbers of groups, from 0 to 20, and then examine the tradeoff between the number of groups and the amount of variation explained by the groups. Intuitively, explained variation should be weakly increasing in the number of groups.

The figure below demonstrates this tradeoff across several metrics. The graph on the top left compares within sum of squares (WSS) across different number of groups k , and the top right compares $\log(\text{WSS})$ across groups. The graph on the bottom left is the coefficient η^2 , which calculates the proportional reduction in WSS that each k provides, compared to the total sum of squares (TSS):

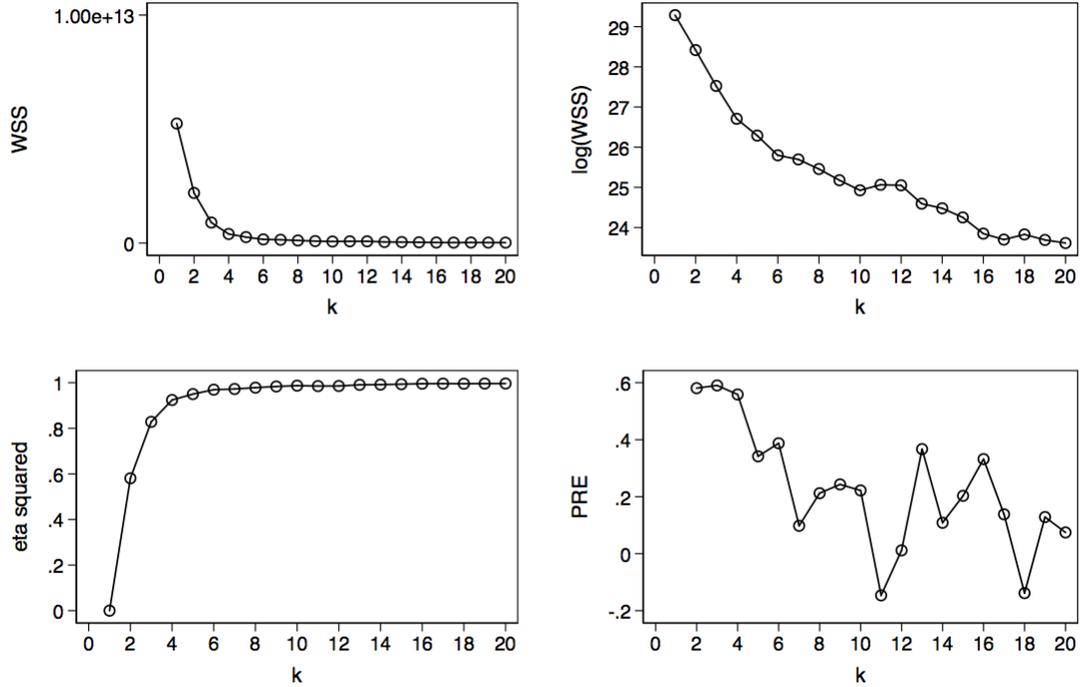
$$\eta_k^2 = 1 - \frac{WSS(k)}{WSS(1)} = \frac{WSS(k)}{TSS} \quad \forall k \in K. \quad (16)$$

The graph on the bottom right shows the proportional reduction of error (PRE) coefficient, which is defined:

$$PRE_k = \frac{WSS(k-1) - WSS(k)}{WSS(k-1)} \quad \forall k \geq 2. \quad (17)$$

The WSS plot shows a drop in WSS up to around five groups, with little additional WSS reduction thereafter. The $\log(\text{WSS})$ plot shows a steady additional reduction as k increases, with a kink at $k = 10$. The η^2 plot shows little additional proportional reduction in WSS to TSS after around $k = 6$. Finally, the PRE plot shows that moving from 10 to 11 groups does not explain additional variation, where the negative value indicates that $WSS(11) \geq WSS(10)$. Given this and the kink at $k = 10$ in the $\log(\text{WSS})$ plot, I categorize the units into 10 different unit types based on size and heat rate.

Figure 13: Performance of K-means clustering by number of groups



To operationalize the model, I calculate the mean operating capacity (MW) and heat rate for each unit type. For the dispatch model, I also need to characterize the minimum and maximum operating levels for each of the units. To identify these levels, I model unit generation as a bimodal distribution, and I use a finite mixture model to identify the two means of generation levels, which provides estimates of minimum and maximum operating levels. I also use a more straightforward approach, which assumes that minimum and maximum operating levels are equal to 0.75 and 1.0 times mean capacity (MW) of the unit type, respectively. The included simulation results use this later approach.

8.5 Identifying Investment and Investment Levels

This section reviews the process to identify investment in the data. The data provides two measures of unit heat rates, which measure efficiency. One measure is a monthly self-reported heat rate, provided by the unit pursuant to federal reporting requirements. The second

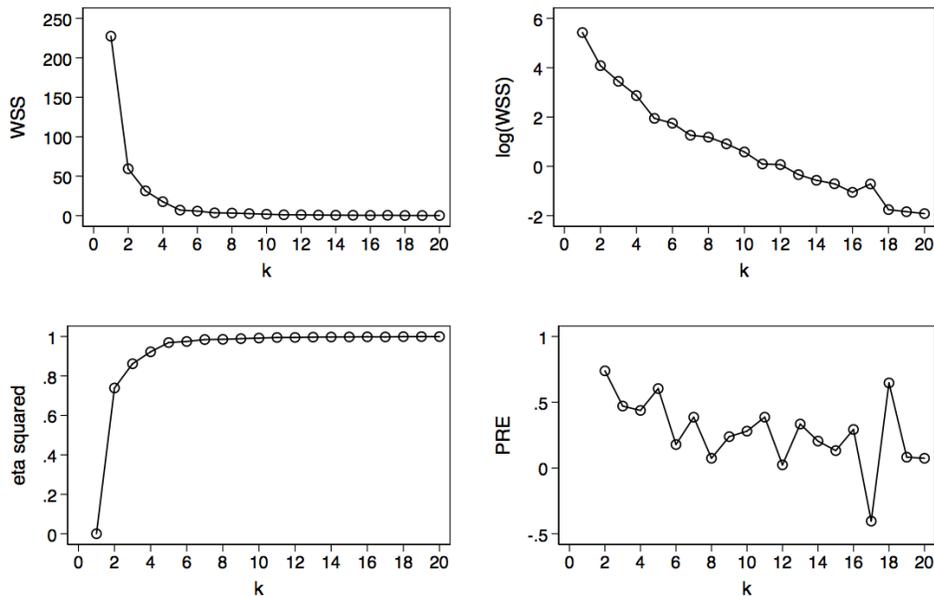
measure I call “inferred heat rate”, which is a calculation of fuel inputs in one hour, divided by electricity produced in that hour. I assume that heat rate improvements (reduction in heat rate) as indicated by both reported and inferred measures decreasing with persistence provide evidence of an efficiency investment decision. To identify cases of persistent heat rate improvements, I first calculate yearly average heat rates, separately based on reported and inferred heat rate data. Monthly average heat rates are also calculated for the inferred heat rate measure. Reported heat rates are already at the monthly level, so do not need to be averaged; missing reported monthly values are filled in based on the last month of available reported data.

Next, I calculate annual heat rate differences by subtracting average annual heat rate in year t from average annual heat rate in year $t + 1$, separately for inferred and reported measures of heat rate. Then, I identify which units decreased their heat rates from year t to $t + 1$, based on both the inferred and reported heat rate metrics, and I record the year of investment. Note, this approach of calculating heat rate improvements from annual averages identifies evidence of endogenous heat rate improvements while allowing for month to month fixed effects that may have exogenous impacts on heat rate, for example, due to seasonal changes in dispatch and weather variation. I that find 50 units observed in the 2012 - 2015 data set decreased their heat rate by both measures in this annual assessment.

Finally, I identify the percent heat rate improvement among the units flagged as investing in order to group investments into bins of heat rate improvement. For each unit that invested, I calculate the percent change in heat rate as a result of investment. I used k-means to group investment improvement levels into bins, considering the variation in improvement levels explained across different number of groups k , similarly done to establish unit type groups as discussed in A.2. To operationalize the model parsimoniously while looking at the tradeoff between increasing the number of groupss and increasing variation explained in the plots below across alternative metrics (WSS, $\log(\text{WSS})$, η^2 , PRE), I set $k = 4$ and group investment improvement levels into four bins. The mean improvement levels as a percent of

heat rate reduction for the bins from one to four are 8, 4, 0.3, and 1.5 percent respectively. The table below shows which unit types invested, and the improvement levels among those units that invested.³²

Figure 14: Performance of K-means clustering for investment by number of clusters



8.6 Unit Type Groups with Local Air Pollution Taxes

In this section I review the unit type groups used for the counterfactual scenario with the location-specific Pigouvian tax on local air pollution. To establish this unit type grouping, I follow a similar procedure as in A.4, but here the grouping variables include local marginal damage estimates for NO_x in addition to the grouping variables in the first unit type grouping, size (MW) and efficiency (heat rate). The intuition is that the grouping variables should include unit characteristics that are relevant to the cost minimization problem, which in this scenario now includes local marginal damages of air pollution. The price of SO_2 is not included, since these units emit a small amount of SO_2 .

³²The simulations included restrict investment to a binary decision for computational speed. This section of the Appendix is included to demonstrate how future versions of this model could expand the investment choice-set.

Table 5: Observed Investment Types and Levels

Type Num.	Num. Units	Avg. HR	Num. Units Invest	Invest Level (1,2,3,4)
5	31	6795	7	(2,2,2,1)
2	9	7460	1	(1,0,0,0)
1	7	8012	1	(0,0,0,1)
7	10	9627	1	(1,0,0,0)
9	31	10657	9	(0,0,6,3)
10	30	11613	13	(1,1,7,4)
6	22	12600	7	(0,1,4,2)
8	23	14403	7	(0,0,5,2)
4	13	16991	3	(0,0,3,0)
3	7	22144	1	(0,0,1,0)
Total Units			Total Invest	Total per Level
183			50	(5,4,28,13)

The graphs below illustrate the increasing variation explained by increasing the number of clusters. We see a large reduction in WSS moving to 4 cluster, and a kink in WSS and $\log(\text{WSS})$ at four clusters and again at seven and 10 clusters. These results together with the upward kink in PRE at 10 clusters lead me to choose 10 clusters for the unit type grouping for the grouping that also considers local air pollution. Note, while this is the same number of groups as the earlier unit type grouping, the composition of the groups is distinct.

Figure 15: Performance of K-means clustering by size, HR, and air pollution damages by number of clusters

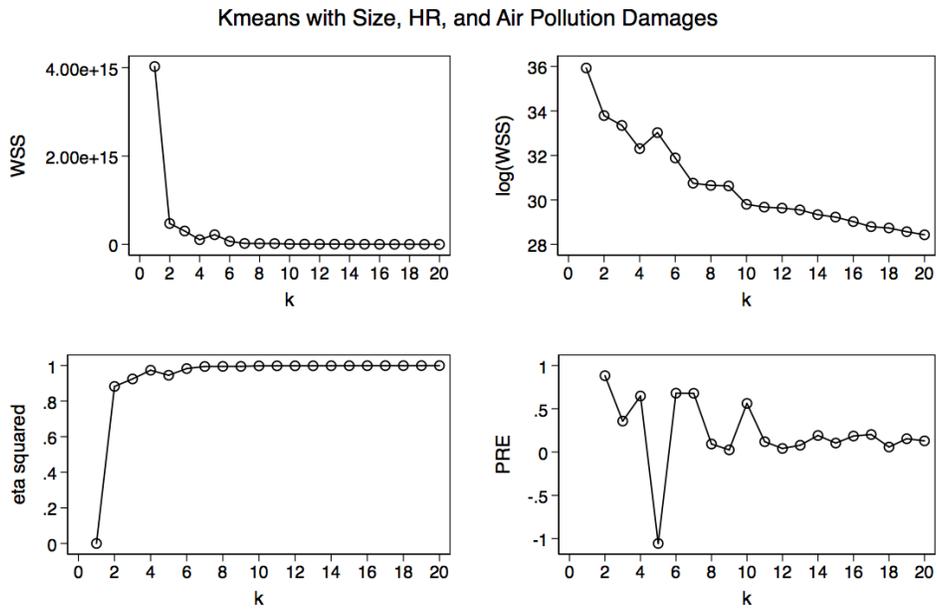


Table 6: Unit Types with Local Air Pollution

Type Num.	Num. Units	Size (MW)	Heat Rate	NOx \$ per ton
1	9	181	9,164	32,955
2	35	122	13,702	17,986
3	18	81	11,368	26,244
4	5	46	13,097	19,413
5	18	153	8,288	4,345
6	8	87	11,876	28,413
7	21	106	10,077	4,480
8	13	134	12,549	6,936
9	17	68	9,589	11,687
10	18	108	11,765	5,628

8.7 Demand Shock Process

I use the data to estimate an AR(1) demand shock process conditional on hour of the day.³³ First, I sum the total generation provided by the evaluated fossil portfolio, q^d , which is the residual demand for electricity corresponding to total domestic demand less imports and less must-take generation. Then, I calculate the average residual demand provided by the fossil portfolio \bar{q}^d for each hour of the day. Each period's demand shock is calculated $\eta = \bar{q}^d - q^d$. The figures below plot \bar{q}^d and the 25th and 75th percentiles of η by hour. Next, for each hour, I calculate the mean and standard deviation of η .

For each hour, I bin the η 's into a D -length grid of demand shocks, with the minimum and maximum shock as observed in the data and establish grid points in equal intervals. This Appendix discusses a grid with $D = 12$, though the procedure can be generalized to other D , depending on computational capacity. Since the demand shock is a state variable, each additional point in the demand shock grid increases the state space exponentially.³⁴ After the demand shocks have been discretized, I record η_{t-1} , and then calculate average demand shock, conditional on both hour of the day and the last period's demand shock. I use the assembled data to construct transition probability matrices conditional on each hour of the day, where each ij element in the matrix corresponds to $Pr[\eta_j|\eta_i]$ and where the rows sum to one.

With sufficient observations and a fine enough grid, one could simply calculate the $P_{ij}|hour = h$. With the discretized shocks, this procedure leads to sparse matrices that are indeterminate, which is un-workable as the probability matrix will need to be inverted in the policy iteration stage. Accordingly, I do the following to construct non-singular probability transition matrices: for each hour of the day and η_{t-1} , I draw a D -length vector of random variables from the conditional distribution $f(\eta|\eta_{t-1}, hour)$ using the distribution parameters calculated earlier. I sort the vector from smallest to largest and calculate the

³³The demand shocks are currently estimated using Q1 2015 data

³⁴I have experimented with different sizes of the grid: this Appendix discusses a grid with $D = 12$, and the current simulation results use a grid of 2. I will update this when grid size is finalized.

probability density function at each point, using the respective distribution's scale parameters. I rescale the vector of probabilities to sum to one, and the set each element P_{ij} as this re-scaled probability. In this way, I generate 24 non-singular probability transition matrices.

This procedure of predicting η based on hour and η_{t-1} has the benefit of being parsimonious, and it can also be shown empirically to explain much of the variation in each period's demand shock. For example, Table 8.7 below shows the results from an OLS regression of $\eta = \alpha + \eta_{t-1} + dHour_t + \epsilon$, which notably includes an R^2 of 0.95.

Table 7: Current Period Demand Shock on Lagged Shock and Hour of Day Fixed Effect

	Curent Period Demand Shock
Last Period Demand Shock	0.97*** (0.00)
Hour Fixed Effect	Yes
R-squared	0.950
N	2159

Standard errors shown in parenthesis. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Figure 16: This figure plots the hourly generation provided by the fossil portfolio in California. The gray bars represent the residual demand provided by the fossil portfolio, which is total demand less imported energy and must-take generation from renewables and other facilities with must-take regulatory status. The blue lines show values between the 25th and 75th percentiles.

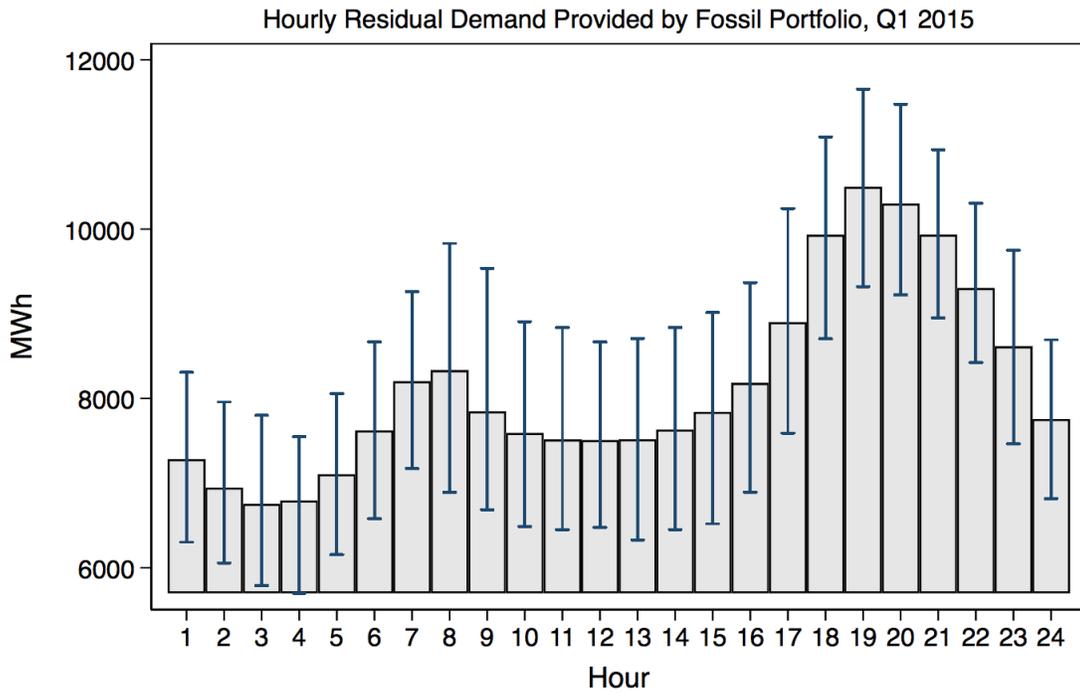
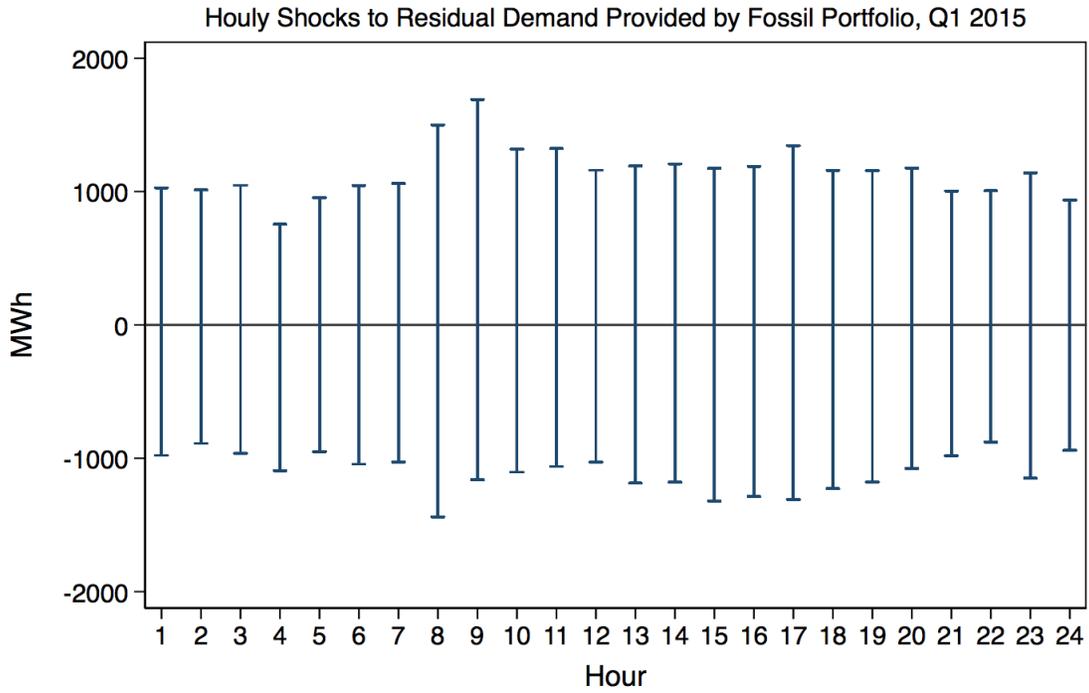
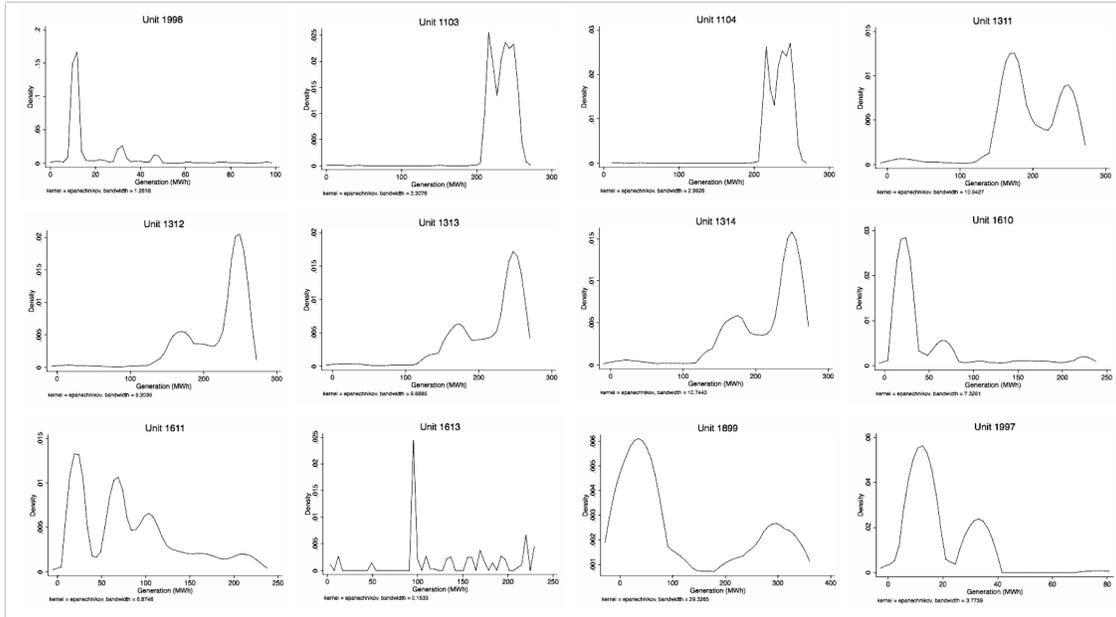


Figure 17: This figure plots the 25th to 75th percentile of hourly demand shocks. The figure shows larger shocks during the morning demand peak around hour eight, as well as in the evening peak, around hour 17.



8.8 Production Levels

Figure 18: These histograms show frequencies of production levels across a sample of units. The histograms illustrate the non-continuous nature of production in this setting, as units appear to operating most frequently across a set of discrete quantities.



9 Model fit

Figure 19: This figure compares hourly generation by unit type in simulated and empirically observed market outcomes of unit type dispatch in Q1 2015.

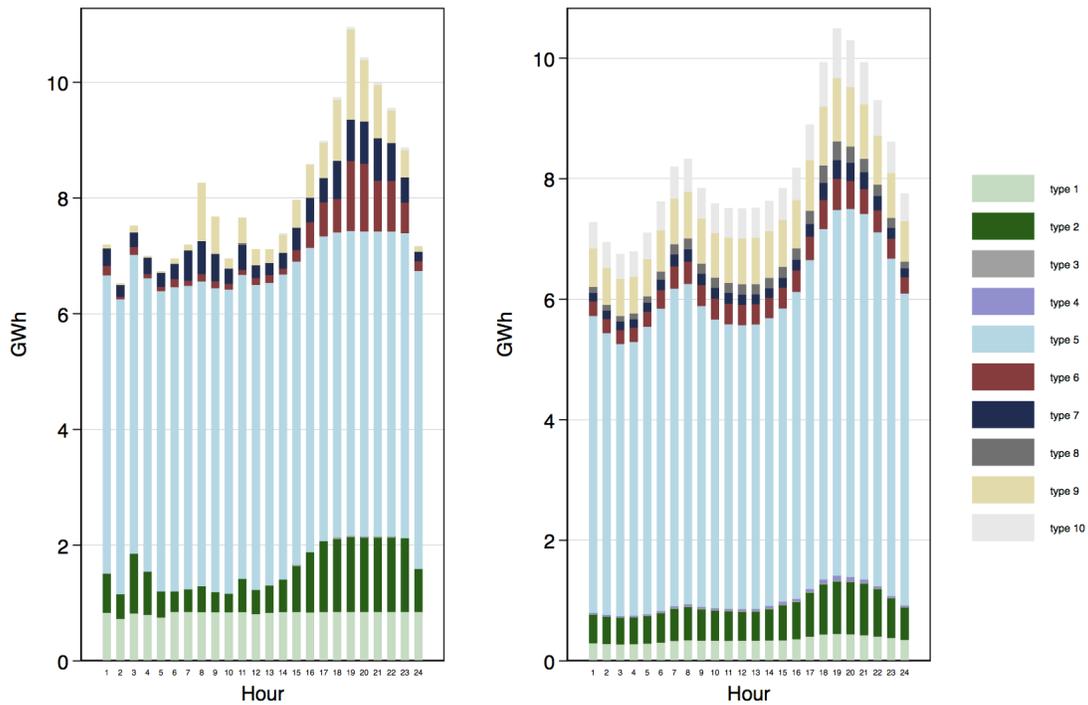


Figure 20: This table compares predicted market shares by unit types in the model to those empirically observed. The t-tests show that the average market shares by unit types are not statistically different across the model predictions and those empirically observed.

	Obs	Mean	Std. Err.	Std. Dev.	$H_o: (\text{diff} < 0)$ $\Pr(T < t)$	$H_o: (\text{diff} \neq 0)$ $\Pr(T > t)$	$H_o: (\text{diff} > 0)$ $\Pr(T > t)$
Simulated MS	21,600	.1	.0013	.1944			
Empirical MS	21,600	.1	.0012	.1825	0.5	1.0	0.5
Simulated gen.	21,600	806	10.35	1521.21			
Empirical gen.	21,600	816	10.33	1517.51	0.01	0.002	0.999

Figure 21: This figure compares the total generation provided by unit type in the simulations with no start-up costs, to those with a start-up cost of \$80 per MW. The large decrease in generation for the unit type with high start-up costs and lower marginal costs corresponds to unit type two, and confirms intuition—with start-up costs, units with lower marginal costs and higher start-up costs are used less. There is also a decrease in generation from type six units, which have lower start-up costs and marginal costs in the middle of the distribution. These units are much smaller in size, so using them to meet demand requires turning on more individual units, and hence more total start-up costs incurred. Thus, they are scheduled for dispatch less frequently in the simulation that includes start-up costs.

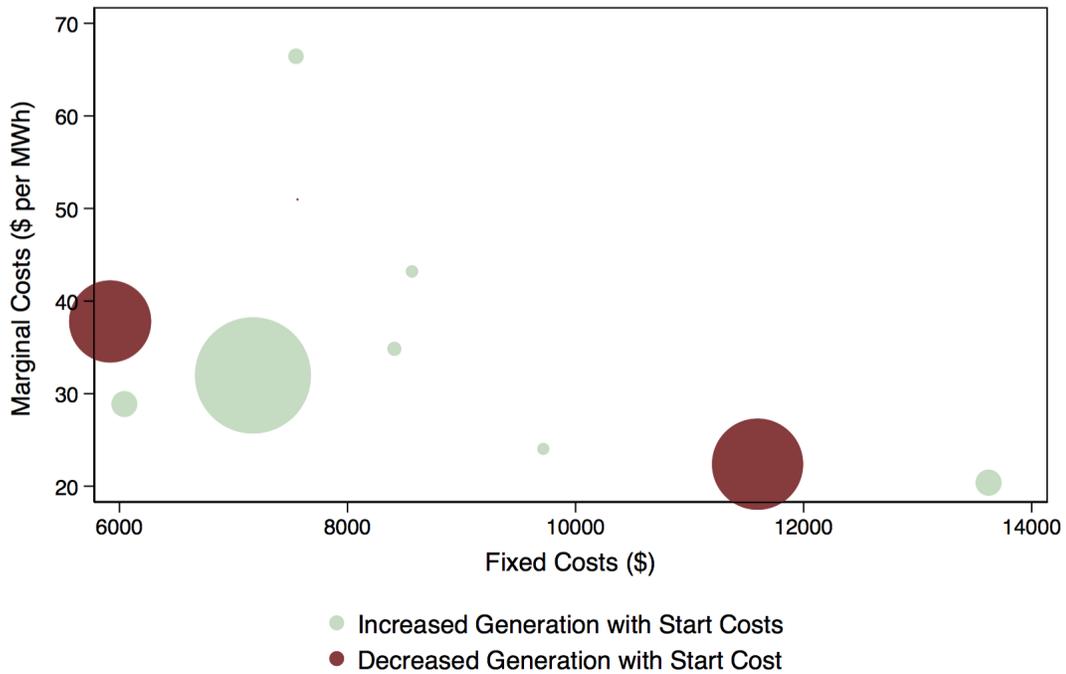
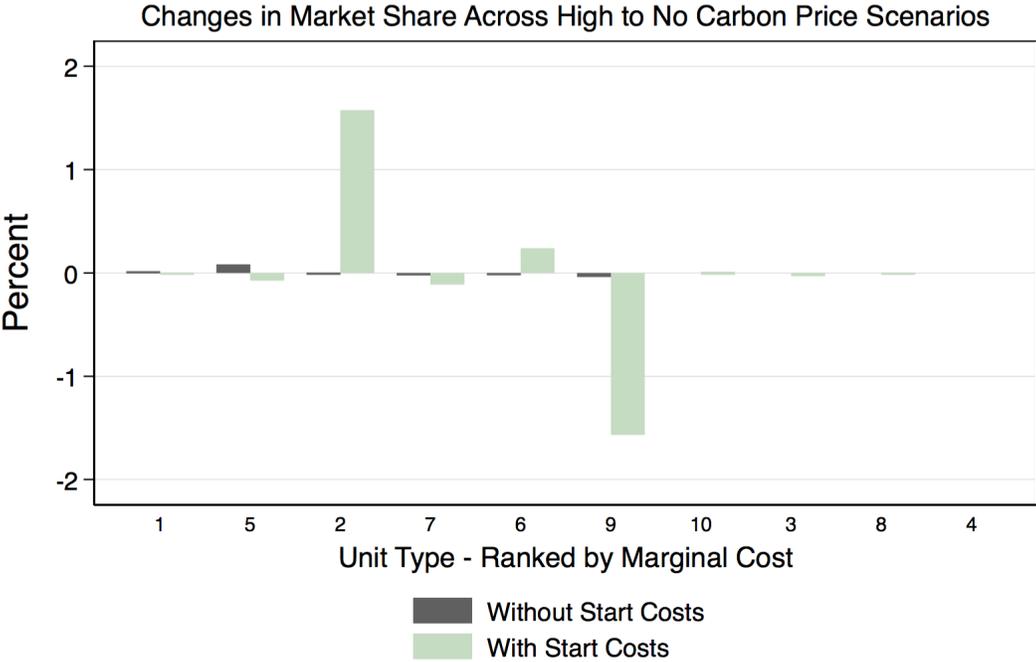


Figure 22: This figure shows the changes in market share across high carbon price scenario compared to no carbon price, for simulations where start-up costs are included and excluded. We see that excluding start-up costs leads to very minimal production reallocation: the mean (max) of the absolute value of the changes in market share for the no start-up cost scenario is 0.009 (0.083); for the scenario with start-up costs of \$ 80 per MW the mean (max) of the absolute value of market share changes is 0.11 (1.6).



10 Theoretical Analysis

In section 3.1 I demonstrate the theoretical prediction that following the introduction of a carbon price, market shares should weakly increase among the more efficient units. Demonstrating this claim requires showing $\frac{\partial V_i^1}{\partial \tau} > \frac{\partial V_i^0}{\partial \tau}$, which I show here, suppressing the i subscript for expositional simplicity.

$$\begin{aligned}
V^1 &= q(P - mc) + \beta(V^1\Pr[a = 1|a = 1] + V^0\Pr[a = 0|a = 1]) \\
\frac{\partial V^1}{\partial \tau} &= qe^f + \beta\left\{\frac{\partial V^1}{\partial \tau}\Pr[a = 1|a = 1] + V^1\frac{\partial}{\partial \tau}\Pr[a = 1|a = 1] + \right. \\
&\quad \left. \frac{\partial V^0}{\partial \tau}\Pr[a = 0|a = 0] + V^0\frac{\partial}{\partial \tau}\Pr[a = 1|a = 0]\right\}
\end{aligned} \tag{18}$$

$$\begin{aligned}
V^0 &= 0 + \beta(V^1\Pr[a = 1|a = 0] + V^0\Pr[a = 0|a = 0]) \\
\frac{\partial V^0}{\partial \tau} &= 0 + \beta\left\{\frac{\partial V^1}{\partial \tau}\Pr[a = 1|a = 0] + V^1\frac{\partial}{\partial \tau}\Pr[a = 1|a = 0] + \right. \\
&\quad \left. \frac{\partial V^0}{\partial \tau}\Pr[a = 0|a = 0] + V^0\frac{\partial}{\partial \tau}\Pr[a = 0|a = 0]\right\}
\end{aligned} \tag{19}$$

The first term in the brackets in equation 18 is bigger than the first term in brackets in 19 since the probability of operating is larger when operating in the last period. The second and fourth term in brackets in equation 18 are equal to the second and fourth terms in equation 19 since the impact of τ 's on operating probability does not vary by lagged operating state since start-up costs are not a function of τ . The third term in 18 is larger than the third term in 19 since the probability of turning off is smaller when operating in the last period.