

Dynamic Response to Environmental Regulation in the Electricity Industry

Joseph A. Cullen *

August 18, 2014

Abstract

In the effort to curb carbon dioxide emissions, the electricity sector will be subject to substantial environmental regulation. The response of electricity generators to regulation occurs in markets with volatile electricity prices where generators face substantial costs of starting up and shutting down generators. These dynamic considerations may affect the responsiveness of generators to regulations that price carbon. This paper recovers the cost parameters of the industry and solves for dynamic competitive equilibria under different environmental policies. The results show that for modest carbon prices, short-run reductions in emissions are negligible. However, the impact on firm profits is substantial and will shape generator investment in the long-run.

Climate change is becoming an important political and economic issue. Governments seeking to curb carbon dioxide emissions are implementing environmental regulations that put a price on carbon dioxide emissions or significantly increase renewable energy production. Since electricity generation is the largest single source of CO₂ in the economy, reducing emissions in the electricity sector is an important component of any potential policy¹.

Carbon regulations interact with electricity-generating decisions in a highly complex market. In particular, generators incur significant costs of starting up and shutting down generators in the face of volatile prices within and across days. Generators essentially face a repeated, high-frequency, entry/exit decision with entry costs. Firms must weigh the costs associated with starting up (entry) or shutting down (exit) generators against the expected future profits from their decisions.

Unlike most entry/exit models, firms in this market are persistent and cannot be considered identical. Firms exit the wholesale electricity market with the

*I wish to thank Gautam Gowrisankaran, Ariel Pakes, Allan Collard-Wexler, Alex Shcherbakov, Daniel Akerburg, Gregory Crawford, Keisuke Hirano, Mo Xiao, Stan Reynolds, Price Fishback, Jonah Gelbach, Glenn McDonald, and Bart Hamilton, as well as the participants in many seminars, for their helpful comments on drafts of this paper.

¹The contribution from sectors in the US: electricity (40%) transportation (33%), direct industrial emissions (17%), direct commercial emissions (4%), direct residential emissions (6%)(EIA 2008)

intention of participating again at some point in the future when the conditions are favorable. Firms are also known to have very different production costs and entry costs. Since there are a large number of generators, I develop a competitive entry/exit model that allows for the estimation of firm-specific heterogeneity and also accommodates firm persistence.

Modeling the dynamics induced by start-up costs is centrally important when considering potential environmental regulation. For example, environmental policies which encourage the development of wind power will increase the difference between peak and off-peak residual demand for generation which, in turn, increases the need for generators to start and stop². Alternatively, policies which explicitly price carbon will increasingly make high start-up cost generators the marginal producers.

Figure 1 illustrates the how dynamics may affect the operation of a generator as its marginal costs increase due to a price on carbon. In the top row of diagrams, the generator's marginal costs are low relative to price. It operates in every period regardless of whether or not it faces startup costs. In every period it receives positive operating profits as shown by the shaded area. However, as a price on carbon increases its marginal costs, it will operate differently in a static setting where it does not face startup costs than in a dynamic setting with startup costs. Without startup costs, the generator will shut down when the market price dips below its marginal cost. However, in the dynamic setting the generator continues to operate in every period even though it incurs operating losses during some time periods. The firm is willing to incur losses to avoid paying the startup cost. Finally, when marginal costs are sufficiently high relative to prices, the generator will choose not to operate at all even though the static model would predict that there are profits to be made in short, peak price periods.

This intuition is born out when simulating output with a model. Figure 2 shows the predicted emissions from static and dynamic models for a coal generator facing significant start-up costs³. A static model without start-up costs would predict a gradual decline in output as the price of carbon increases. This reflects the gradual increase in the number of periods in which prices are below marginal cost where the firm chooses not to operate. A dynamic model, on the other hand, shows that the firm is much less responsive to carbon prices when prices are low. The firm finds it optimal to continue operating as usual, despite incurring losses in some periods, in order to avoid the costs of starting

²Wind farms, which are on shore, have the highest output during times of off-peak demand and have little output during high demand periods. Wind power thus reshapes the residual demand curve by increasing the difference in demand between on and off peak periods. Wind farms which are built offshore will have the opposite effect of residual demand since the wind usually blows off shore during peak demand periods when energy is most needed.

³The graph shows the aggregate expected output over a three month period of a Texas, coal-fired power plant at different levels of carbon prices with a \$150,000 start-up cost. For this illustration, the plant faces the baseline wholesale electricity price at each carbon price rather than the equilibrium price path that would exist if all firms had to pay for carbon emissions. Thus the output reduction in response to carbon prices in the diagram is much more dramatic than the output reduction that would exist in equilibrium.

up and shutting down. However, at some critical point, the firm switches to rarely operating. Thus, depending on the price of carbon emissions, a simple static model may over- or under estimate the impact of carbon regulation on output and profitability. However, it is important to note that this simulation holds fixed electricity prices. Equilibrium prices under a carbon price may either mitigate or exacerbate the behavior shown here.

The model builds on John Rust (1987) to estimate generator-specific distributions of start-up costs using a detailed dataset on generator output and energy prices. These start-up costs are then used to simulate outcomes under alternative policy environments. A central contribution of the paper is showing that the optimal policies of dynamically-optimizing single agents in the model can be aggregated together to solve for new price equilibria. In each counterfactual equilibrium, the market clears in each period and firms' expectations for price are consistent with the distribution of prices in the new equilibrium.

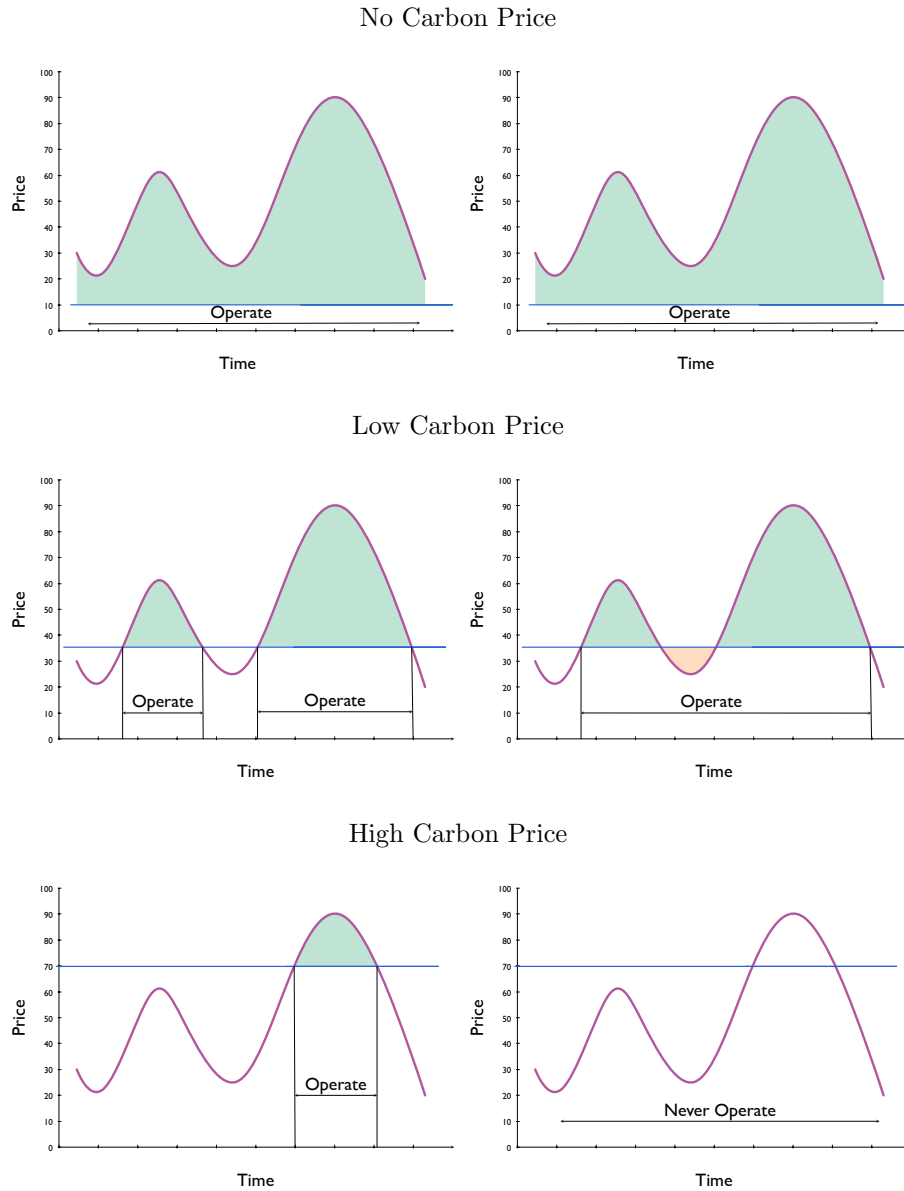
Using this new framework, the model is solved under two potential policy environments: 1) putting an explicit price on carbon and 2) increasing the number of wind farms. The results show that short-run emissions reductions from pricing carbon are negligible for modest carbon dioxide prices of \$20/ton. Introducing more wind power into the grid has a much more immediate effect on emissions reductions; boosting the share of wind power from 2% to 10% results in a 6% reduction in carbon dioxide emissions.

I also find that counterfactual emission reductions are not dramatically different when comparing the results from the dynamic model with a static formulation. However, a static model substantially underestimates price volatility and overestimates the profitability of renewables. Higher price volatility in the dynamic model means lower prices for wind farms when they are most productive.

The results reflect the short-run response of the electricity industry to environmental policies. They do not include any emissions reduction from investments in new generating technologies that may occur in the long run. However, properly characterizing the short-run equilibrium is critical for understanding how regulation will shape investment in the future. Firms' profits are determined by the day-to-day operation of generators and their competitive interaction. Modeling short-run changes to profitability due to environmental regulation is an important first step to understanding the long-run investment response. Also, even though investment in new facilities is outside of the scope of the model, "exit" by firms can occur when, given the new equilibrium prices, firms choose not to operate indefinitely.


Identifying the immediate impact of environmental policies is also valuable in and of itself. For example, understanding the emissions response in the near future sets appropriate expectations for regulators and policy makers. Even though economists extol the virtues of market-based instruments, regulators are often hesitant to implement regulations with uncertain emission benefits. Properly characterizing the short-run response to market-based policies can reduce regulator uncertainty. Also, the lag time for investment in this industry is significant. Efficient combined cycle power plants take three or more years

Figure 1: STATIC AND DYNAMIC RESPONSE TO CARBON PRICING



to build, not including the time spent in the permitting process. New nuclear facilities construction times are six to eight years. With the lengthy permitting process, time to build can extend well past a decade (Kehlhofer et al. 2009) (AEP 2014) (NEA 2014). On the other hand, existing infrastructure, such as

Figure 2: Emission Response to Carbon Pricing



coal plants with large sunk costs, but relatively low costs going forward, may continue operating for some time. Thus, understanding equilibrium outcomes given existing generator infrastructure will be important for years to come.

Relatively little attention has been paid to dynamic considerations in electricity markets with most work adopting static formulations for production and profits. There are several notable exceptions. First, work by Mansur (2008) highlights the importance of dynamics when gauging the competitiveness of electricity markets. Using a reduced form model, Mansur finds that ignoring dynamics substantially overestimates welfare losses associated with deregulation of the market. Second, a forthcoming paper by Reguant (2012) evaluates the welfare effects of allowing firms to use complex bids to better reflect their start-up costs. She estimates start-up costs using detailed generator-level bid information in a finite-horizon dynamic model. She finds that more complicated bidding structures do not increase the use of market power, but do lead to increased productive efficiency. In this paper, I extend the literature by developing an infinite-horizon dynamic framework that can both estimate adjustment costs and solve for counterfactual equilibria using readily available data. I apply the model to a regional electricity grid to simulate dynamic industry response to environmental regulation.

This model has several advantages over reduced form or engineering approaches to analyzing counterfactual outcomes in the electricity industry. Since the model explicitly solves each generator's dynamic problem, it is possible to simulate equilibrium outcomes that are very different from observed equilibrium outcomes. In contrast, reduced form approaches are not able to effectively deal with counterfactual equilibria which are too far out of sample. Since we do not observe carbon prices or high levels of wind power, we need a structural model to simulate these counterfactual scenarios⁴. Second, the structural approach is more appropriate for simulating situations with increasingly volatile equilibrium prices. The reduced form approach cannot handle such situations since the firms' reactions are known only for the observed level of volatility in the market. As prices become increasingly volatile, startup costs will become increasingly important for understanding firm behaviour. Highly volatile prices favor flexible, low-startup-cost generators which can exploit price spikes while avoiding low price periods. Inflexible generators will have to decide whether to ride out the price swings, or stay out of the market. If uncertainty increases in step with volatility, then the option value created by uncertainty will also affect the behaviour of generators differentially depending on their startup costs. The structural model explicitly incorporates uncertainty and option values into the firm's decision that are typically not incorporated into engineering planning models. Of course, a structural approach is not without its downside. Explicit assumptions must be made regarding firm profits, beliefs, and the nature of competition. In addition, there is significant computational overhead associated with estimating and solving a dynamic, structural model.

⁴In the absence of carbon prices, recent work by Cullen & Mansur (2014) has used the unprecedented drop in natural gas prices in a reduced form model as a proxy for pricing carbon.

The rest of the paper proceeds as follows. In sections 2 and 3, I describe the necessary details of the market and the data used for estimation. Then in section 4 the model is presented. Sections 5 and 6 detail the estimation of start-up costs. Finally, in section 7, the estimated parameters are used to solve for counterfactual equilibrium under different carbon prices and levels of wind penetration.

1 Electricity Market

This paper uses the Texas grid, which is managed by the Electricity Reliability Council of Texas (ERCOT), as a laboratory for the analysis. Before presenting the model, it is useful to understand the basic characteristics of power systems and the institutional details of ERCOT that motivate the modeling approach.

1.1 Power System Basics

Electricity is an unusual commodity in several ways. First, demand for electricity is almost perfectly inelastic in the short-run; very few consumers of electricity are willing or able to adjust consumption in response to changing market conditions. Second, the quantity of electricity demanded at a given price varies cyclically over the course of a day and throughout the year. Peak demand can be twice that of off-peak periods within the same day. Finally, electricity is unusual because it cannot be stored in meaningful quantities⁵. Electricity production and consumption on a grid must be balanced on a second-by-second basis. If more power is being consumed than is being produced, then the reliability of the grid is threatened. Sufficient imbalances result in brownouts (dropping electrical frequency) or blackouts (complete loss of electrical service). Given that short-run demand is inelastic and highly variable, the lack of energy storage puts high demands on generators to preserve the reliability of the grid by adjusting output to follow changing demand.

As generators follow demand, they face several output constraints. First, generators are capacity constrained. The maximum output capability of a generator is determined at the time of its construction and generally remains fixed over its life. Generators also face minimum output constraints. The minimum output constraint is the lowest level of sustained output the firm can generate without shutting down. Operating below the minimum output level results in large inefficiencies and can damage generating equipment.

Generators also face costly adjustments to output. Foremost among these are start-up costs. Start-up costs are incurred when bringing a generator online after a period of zero production. Bringing the generator online requires fuel

⁵Chemical storage of electricity in conventional batteries is too costly to be used to store any meaningful amount of electricity in a system. Technologies do exist to turn electrical energy into potential mechanical energy which is storable such as compressed air or pumped hydro electrical storage. These technologies do make minor contributions on some grids, but such technologies have not been implemented on the electrical grid in this study.

to heat up equipment and bring the turbine up to speed as well as additional labor to supervise the process. In addition, start-ups are hard on equipment leading to increased maintenance costs in the long run. In fact, engineering studies estimate that wear and tear on generating equipment may account for the majority of start-up costs (Chow et al. 2002).

Start-up costs vary widely by generation technology and the age of the unit. For example, gas combustion turbines have lower start-up costs than coal fired steam plants. Likewise, large plants will have higher start-up costs than smaller generators even though they may use the same technology. Also, as generators age, they degrade in efficiency which will increase start-up costs.

The costs associated with starting up generators are significant. Engineering estimates of start-up costs range from hundreds of dollars to hundreds of thousands of dollars per start depending on the size and technology of the generator. Consequently, a generator with high start-up costs may continue to run during low price periods to avoid start-up costs. Likewise, a generator may not start up even though prices exceed its marginal cost of production if it believes that the profits will not be sufficient to cover its start-up costs.

Concrete information on start-up costs is generally unavailable to researchers and policy makers. The information is considered proprietary and thus is not made publicly available. In addition, there may be substantial non-engineering costs that would be left out of an engineering cost model. This paper provides a method to estimate start-up costs given publicly available information.

1.2 ERCOT

ERCOT operates as a deregulated electricity market which serves most of the state of Texas. It operates almost independently of other power grids, with very few connections to outside markets. Electricity generation and retailing are deregulated, while the transmission and distribution of energy remains regulated to ensure that competitors in the generation and retailing markets have open access to buy and sell power. Unlike many regulated and even restructured markets, companies in this market are vertically separated. There are no vertically integrated firms that control generating, transmitting, and retailing resources.

1.2.1 Generators

There are approximately 500 generators owned by 80 firms which supply electricity in ERCOT. The ownership of facilities is fairly diffuse, as shown in table 1. The two largest firms account for 24% and 16% of total capacity, respectively, with the residual dispersed among many smaller firms. Major generation technologies include coal, nuclear, natural gas, and wind power, with very little hydropower.

Each generator sells its energy to buyers either through bilateral contracts or through ERCOT's real-time spot market called the Balancing Market. Approximately 95% of energy produced is sold through bilateral contracts. The

Table 1: Ownership Shares

Owner	Capacity	Share
TXU Generation LLP	17720	24%
Texas Genco II	12301	16%
City Public Service San Antonio	4301	6%
Austin Energy	3523	5%
ExTex LaPorte	2404	3%
Lower Colorado River Authority	1911	3%
Other firms < 2% share	31748	43%

remaining 5% is allocated through the Balancing Market. In general, firms are not required to meet their bilateral energy contracts using their own generators. Instead they can buy power in the Balancing Market. Likewise, any available excess capacity is easily bid into the Balancing Market. Due to this operational flexibility, the Balancing Market price represents a generator's opportunity cost of production. Further details on the market mechanisms in ERCOT can be found in Appendix C.

Balancing Energy prices can be quite volatile as shown in figure 3. The three lines show representative hourly price paths with daily variation in the 25th, 50th, 75th percentiles of price variance for 2006. All three days exhibit higher prices during peak demand periods; the highest variance price path shown has peak prices that are twenty times that of off-peak periods.

1.2.2 Transmission Congestion


Most of the time ERCOT operates as a single market with a single spot price for wholesale electricity⁶. In these uncongested periods, ERCOT does not differentiate between remote generators, which make extensive use of transmission lines, and those which are located in close proximity to demand. During periods when transmission congestion does arise, the grid is split geographically into zone specific markets each with their own spot prices. For example, zones that are net importers of electricity would see higher prices in order to spur local production while exporting zones would see lower prices in order to relieve demands placed on inter-zonal transmission lines. In this way, interzonal congestion is relieved through pricing mechanisms in the Balancing Market.

1.2.3 Demand

As in most electricity markets, demand in ERCOT does not respond directly to wholesale price signals. Residential and commercial users purchase electricity at fixed prices which are constant for period of time ranging from one month to several years. As such, they have no incentive to reduce consumption during

⁶A single price prevailed in ERCOT in 98.5% of the hours in 2006.

Figure 3: Representative Daily Price Variation by Percentile



high price periods in the wholesale markets⁷. Over a longer period of time, if average prices in the wholesale markets rise, this information will eventually be passed along to consumers in the form of higher retail rates. However, in the short run demand for electricity is inelastic.

1.2.4 Comparison with other markets

Electricity markets are complex and highly idiosyncratic. Although each electricity grid has its own unique mix of generators and market rules, it is useful to place Texas within the larger context of US and European electricity markets. Table 2 compares share of demand by consumer type while tables 3 and 5 compare electricity generating capacity and electricity production in Texas with the US as a whole, Europe and the world.

First on the demand side, the share of demand by residential, commercial and industrial customers in Texas is similar to the US and Europe, but slightly more industrial focused. With 68% of demand coming from industrial customers, Texas is more similar to world wide demand shares than to the US or Europe. In commercial demand, Texas demand share lies in between that of US and Europe, but has a residential market share is that below both.

Table 2: Share of Demand by Customer Type

Fuel Types	Texas	US	Europe	World
Residential	17	24	27	19
Commercial	15	19	13	11
Industrial	68	57	60	70
Total	100	100	100	100

On the supply side, Texas represents just a little over 10% of all generating capacity in the US. However, nearly all of its capacity, 91%, uses fossil fuel as an energy source. It has limited amounts of nuclear and renewable capacity. Texas has negligible hydroelectric capacity; nearly all of its renewable capacity comes from wind turbines. The US as a whole relies less on fossil fuel generators (77%) due to higher shares of nuclear and hydro capacity. Europe is even less fossil fuel dependent (53%) with even higher shares of renewables and nuclear.

When looking at fuel types within fossil fuel capacity, Europe and the US have roughly similar mixes, as shown in table 4. Texas has almost no liquid

⁷Some large industrial consumers do curtail electricity use when reserve capacity becomes short, but they do not directly respond to fluctuations in the price of electricity in the wholesale market. These large industrial users negotiate lower energy prices by agreeing to have their supply of electricity temporarily interrupted in emergency situations when generating reserves on the grid reach critical levels. Industrial users with interruptible loads are called Loads Acting As Resources (LaaRs). In the event of an unexpected change in load, electricity delivery to the LaaR will be interrupted to maintain the frequency on the grid. Approximately half of responsive reserve services are supplied by LaaRs (MF7). Again, it is important to note that LaaRs respond to events that threaten the reliability of the grid, not to price changes in the wholesale market.

petroleum generators, but has a larger share of gas capacity and a smaller share of coal capacity. Texas' large gas capacity implies that most, if not all, generation could be provided by less polluting gas fired power plants. This will be important to keep in mind when viewing the counterfactual results.

Table 3: Percent Capacity by Fuel

Fuel Types	Texas	US	Europe	World
Fossil	91	77	53	67
Nuclear	5	10	16	9
Hydro	1	8	17	19
Renewable	3	5	14	6
Total	100	100	100	100

Table 4: Percent Fossil Capacity by Fuel

Fuel Types	Texas	US	Europe	World
Coal	22	41	47	48
Gas	78	43	41	38
Liquids	<1	16	12	14
Total	100	100	100	100

When looking at electricity production rather than capacity, we see that Texas production largely comes from fossil fuel facilities. Nuclear and renewables provide 13% of production with residual 87% coming from fossil fuels. This is higher than the US fossil share of 71% and much higher than Europe's 53% production share. Again the difference comes primarily from the lack of hydro power and the limited nuclear facilities in Texas.

More importantly, Texas has a cleaner energy mix within fossil fuels than either the US or Europe, as shown in table 6. Only 43% of its fossil power comes from coal, compared with 54% in Europe and 70% in the US. The remaining 57% of fossil fuel generated electricity in Texas comes from natural gas. This compares with 39% in Europe and 28% in the US as a whole.

Table 5: Percent Production by Fuel

Fuel Types	Texas	US	Europe	World
Fossil	87	71	53	66
Hydro	<1	7	14	17
Renewable	3	3	5	2
Nuclear	10	20	27	15
Total	100	100	100	100

Table 6: Percent Fossil Production by Fuel

Fuel Types	Texas	US	Europe	World
Coal	43	70	54	61
Gas	57	28	39	31
Liquids	<1	2	7	8
Total	100	100	100	100

These comparisons have implications for how the Texas response to environmental regulation may differ from that of other grids. First, since Texas already produces a larger share of its electricity from cleaner gas fired facilities, the maximum possible percentage change in emissions may be lower in Texas than in other electricity generating areas where the fossil generating mix is more coal based. Second, Texas may be more likely to achieve its maximum possible reduction, or may achieve it sooner, due to the availability of gas fired generating capacity. Despite the differences between Texas and the rest of the world, the dynamic incentives and constraints that generators face when operating remain largely the same. Generators in Texas and elsewhere will face increasing costs and price volatility as the result of environmental regulation.

2 Data

The data used for this analysis span the period May 1, 2006 to August 31, 2006. Over this time frame, the output of each generator is observed every hour. The zone-specific market clearing price for Balancing energy is also observed hourly⁸. Other generator-level characteristics include the maximum and minimum output capability for each generator, the age of the generator, outage status, fuel type and its location.

These data are supplemented with information from the Environmental Protection Agency (EPA) and the Energy Information Administration (EIA) on generator heat rates and emission rates. A generator's heat rate measures its productive efficiency in terms of the heat input from fuel necessary to produce 1 MWh of electricity. Emission rates measure the average quantity of the SO₂, NO_x, and CO₂ emitted per MWh of output. Additional information on the cost of fuels and pollution permits are collected from the EIA, EPA and ICE⁹.

⁸Output and price data are available at fifteen-minute intervals, but data at the hourly level are used in this study. This approach is taken for two reasons. First, very few generators have the technical ability to turn on or off within a fifteen-minute period. Thus, although one may observe a high price this period, a generator may not be able to respond to that price. Looking at the hourly prices averages out some of the noise introduced by temporary price spikes. Second, averaging over an hourly period more closely matches the scheduling decisions of firms, which are submitted at the hourly level.

⁹For fuel costs for coal plants, monthly data is collected from EIA form 423 which gives the delivered quantity and cost of fuel for coal in Texas. The quantity-weighted average coal price

The data are used to construct the marginal cost of electricity production, which is the marginal cost of fuel plus the marginal costs for emissions for each generator.

$$MC_i = FuelCost * HeatRate_i + SO_2Cost * SO_2Rate_i + NO_xCost * NO_xRate_i$$

The data do have some limitations. First, the electricity prices used in the model are not necessarily the prices the firm received for its output since most energy in this market is sold via bilateral contracts with unobserved prices. However, as previously discussed, spot prices do represent the opportunity cost of production for the firm. A firm can always shut down production and fulfill its contract by buying power in the balancing market. Market analysis by ERCOT also suggests that forward contract prices for energy follow balancing price quite closely (Potomac 2007).

Second, some generators are paid to provide ancillary services for the market such as regulation, capacity reserve, or out-of-merit-order energy. These generators respond to price signals that are not observed. For example, particular generators may be called on to produce electricity in order to alleviate local congestion, even though the prevailing price for electricity would not merit production. Although data is not available on these deployments, conversations with operators indicate that such deployments are unpredictable and do not account for a large share of generation.

3 Model

The model builds off of the work of Rust (1987) in single agent dynamics to estimate start-up costs. Once startup costs are estimated, the model is used to solve for counterfactual competitive equilibria. In the model, the decision maker is a firm deciding when to operate or shutdown its generator.

In this competitive model, firms are assumed to be price takers. That is, they take prices in the market as exogenously given; they do not think strategically about manipulating the price for electricity with their operation decisions.

Price taking is not an innocuous assumption, especially considering the active literature on the exercise of market power in electricity markets (Severin Borenstein, James B. Bushnell & Frank A. Wolak 2002), (Mansur 2008), (Ali Hortacsu & Steven L. Puller 2008). There are several conditions specific to ERCOT that make this assumption more plausible, though not unassailable. First, ownership rules limit a firm's ownership of generation facilities to 20% of the total generation capacity in any zone¹⁰. Second, most of the energy is sold

as the price for coal for all generators in the market for that month is used as the monthly coal price. For the cost of fuel for gas powered plants, the average spot prices for natural gas transactions on the Intercontinental Exchange (ICE) are used to create a monthly price index. For pollution permits, prices are collected from EPA auctions for both SO₂ and NO_x permits in 2006. Carbon dioxide is currently unregulated so there is no cost associated with CO₂ emissions.

¹⁰This rule does seem to be violated for one incumbent utility, TXU, which has several

via bilateral contracts. Since most of the energy is not sold at the spot price, this reduces the incentives for a firm to withhold production to increase the energy price in the spot market if it were to exercise market power (Wolak 2000), (Bushnell, Mansur & Saravia 2008). That said, price taking is an important and possibly restrictive simplifying assumption of the model.

Price taking allows the firm's decision problem to be modeled as a single-agent dynamic problem, rather than as a more complex, dynamic game. Solving for counterfactual equilibria, in a dynamic game with many strategic players would be computationally infeasible. Price taking also renders the ownership of power plants irrelevant¹¹. Thus in a competitive market, each generator on the grid can be thought as a separate firm maximizing its own profit, regardless of ownership status. Accordingly, the terms "generator" and "firm" will be used interchangeably throughout the paper.

As is standard in the literature on electricity markets, marginal costs are assumed to be constant and known. The heat rate of generators, though not entirely constant, is relatively flat within the operating range of the generator. Accurate data on heat rates, fuel costs, and emissions prices allow the marginal cost to be calculated rather than estimated. There are components of the marginal cost that are left out of the standard calculation. These include transmission costs or other variable input costs such as water for steam plants. However, these deviations from the standard assumption are likely to be of second order importance.

I also assume that firms are not constrained by local transmission bottlenecks when optimizing with respect to price. This still allows for the primary paths of congestion, namely congestion between zones, to be represented by the model since this type of congestion is alleviated in ERCOT via price mechanisms. However, this assumption does rule out congestion within a zone. Due to data limitations, the model cannot account for local congestion directly.

Although the model incorporates costly output adjustment through startup cost, it abstracts away from adjustment costs if a generator is already operating. That is, it is assumed that a generator can costlessly adjust output within its operating range. In practice, all generators in the model have the technical capability of moving between their minimum and maximum operating levels in 5-10 minutes with a few needing as long as 20 minutes to make the move. With this assumption, the generator's decision collapses from a continuous choice of output level to a discrete choice of whether to operate or not.

Given these assumptions, each generator is modeled as a single firm with the following dynamic problem. In each period, the firm observes the price in

smaller companies that appear to be owned by a single holding company. It is not clear to what degree the companies are considered separate entities. The total capacity share for the holding company is 24%, though each smaller company has capacity shares much smaller than that.

¹¹Consider a generator with a set of cost characteristics. Given beliefs over the future path of prices, there is only one optimal way to operate the facility. Any rational firm in the market would operate the generator in exactly the same profit-maximizing way. This is not the case when firms can exercise market power. Each firm would use the characteristics of their entire generator portfolio to optimally exercise market power.

the market and the hour of the day. The firm can take one of two actions which are notated as:

$$a_{it} = \begin{cases} 1 & \text{if operate in } t \\ 0 & \text{if not operate in } t \end{cases} \quad (1)$$

where i indexes the generator
 t indexes each hourly period

If the firm decides to operate, the firm's output will be decided by the market price. If the price in the market is greater than the firm's marginal cost, then it will produce at maximum capacity. If the price is below marginal cost, then the firm will produce at its minimum possible level. If the price is equal to the firm's marginal cost, the firm may operate at any point along its operating range.

$$\begin{aligned} q_{it} &= \max \text{ if } P_t \geq c_i \text{ and } a_{it} = 1 \\ q_{it} &= \min \text{ if } P_t < c_i \text{ and } a_{it} = 1 \end{aligned} \quad (2)$$

where c_i = constant marginal cost of generator i
 P_t = price for electricity in the generator's zone

Each period when the firm is operating, its profits are simply the price-cost differential earned on every unit produced minus any start-up costs associated with operating. The per period profit function for the generator is then:

$$\Pi(P_t, q_{it}, a_{it}) = \begin{cases} (P_t - c_i)q_{it} & \text{if } a_{it} = 1 \text{ and } L_{it} = 1 \\ (P_t - c_i)q_{it} - START_i & \text{if } a_{it} = 1 \text{ and } L_{it} = 0 \\ 0 & \text{if } a_{it} = 0 \end{cases} \quad (3)$$

where $START_i$ = cost of starting up generator i
 $L_{it} = a_{it-1}$ = the lagged operating state

No profit is earned when the generator is not operating. A start-up cost is incurred only if the firm decides to operate, but was not operating last period ($L_{it} = 0$). The structural parameters of the model are c_i and $START_i$. The constant marginal cost of production, c_i , is calculated from the generator heat rate, fuel prices, and emission costs as previously discussed. The structural parameter $START_i$ is not known and will be the object of the estimation procedure. For notation simplicity, the i subscript will be dropped for the remainder of the paper since each generator is always modeled separately as a single agent.

In the dynamic model, the firm's expectations over future prices must be explicitly modeled. Prices are assumed to follow a conditional AR(1) Markov process described by the distribution $F(P_t|P_{t-1}, H_{t-1})$ where H_t is an indicator for each hour of the day. Note that because of the price taking assumption, the evolution of price does not depend on the action of the generator.

One might argue that a simple Markov process is not sufficiently rich to accurately model the expectations of the firm. Indeed, firms have more information than simply the lagged price and time of day with which to form expectations for price in the next period. For example, firms may have expectations over future temperatures, load levels, and congestion. In addition they may use a

long price history when predicting future prices. The extent to which the proposed model for the evolution of price is adequate depends on the degree to which lagged price summarizes all of the other components of the expectations of price. I investigate the price process in detail after the empirical specification for price expectations is introduced in section 4.

Given the specification of the transition and the profit function, the state space for the dynamic problem will be the price, the hour of day, and the lagged operating state, (P_t, H_t, L_t) . The Bellman equation representing the dynamic problem can be written as:

$$V(P_t, H_t, L_t) = \max_{a_t \in \{0,1\}} \{ \Pi(P_t, L_t, a_t) + \beta E[V(P_{t+1}, H_{t+1}, L_{t+1} | P_t, H_t, L_t)] \} \quad (4)$$

$$\text{where } \begin{aligned} H_{t+1} &= H_t + 1 - 1(H_t = 24) * 24 \\ L_{t+1} &= a_t \end{aligned}$$

where the expectation is taken with respect to P_{t+1} according to the distribution $F(P_{t+1} | P_t, H_t)$. The parameter β is a fixed discount factor.

The optimal policy for this dynamic problem is a cutoff rule in P_t for every pair of (H_t, L_t) . This implies that the firm should take same action whenever it encounters the same state (P_t, H_t, L_t) . This creates a problem for estimating structural parameters from the data as the firm will invariably deviate from what appears to be the optimal policy. For example, on one day we may observe an idle generator starting up at a price of \$50/MWH and 8:00AM. However, on another day at the same price and time of day we might find that the generator does not start up. As is, the model cannot rationalize this behavior.

To address this issue, I allow for a shock each period which affects the decision of the firm to operate. If the firm is not operating, this shock, plus the start-up cost, can be interpreted as a draw from a distribution of start-up costs with mean $START_i$ and variance σ . If the generator is operating, the shock can be interpreted as a draw from the distribution of exit or shutdown costs. Since the model will not be able to differentiate between start-up and shutdown costs, the mean of the exit cost distribution is normalized to 0, but with the same variance as the start-up cost distribution¹².

The shock introduces an additional state variable into the dynamic problem which is observed to the firm, but unobserved to the econometrician. Technically, this is implemented as a choice specific shock to profit each period and is notated as $\epsilon_t(a_t) \in \{\epsilon_t(0), \epsilon_t(1)\}$, as in Rust (1987). Like in Rust (1987), the

¹²A firm that is considering starting up today, knows that it may shutdown sometime in the future. Thus it will take into account any shut down cost when it makes its startup decision. From the firm perspective, the main difference between the two costs is that one is incurred now and one is incurred later. In principle this implies that the discount factor could be used to separate the two costs. For example, if the firm discounted the future heavily enough or if the shutdown decision is expected to be very far in the future, then the impact of the shutdown cost on the startup decision would be negligible. On the other hand, if the future is not discounted at all then the firm would realize the full shutdown cost when it made its startup decision. Since the discount factor in this setting is close to one, I do not try to separately identify startup and shutdown costs.

shock is assumed to be an independent and identically distributed stochastic process that introduces noise into the underlying decision process. For computational simplicity, the distribution for $\epsilon_t(0)$ and $\epsilon_t(1)$ is assumed to be extreme value type I. This allows for analytical integration over the unobserved shocks. Because the level of profit can be calculated for each state, the variance of the error process can be estimated, unlike in most discrete choice models. Accordingly, the choice-specific shock becomes $\sigma\epsilon_t(a_t)$. Let the vector of unknown structural parameters be $\theta = (START, \sigma)$. With the unobserved state variable, the Bellman equation now becomes:

$$V_\theta(P_t, H_t, L_t, \epsilon_t(a_t)) = \max_{a_t} \{ \Pi(P_t, L_t, a_t) + \sigma\epsilon_t(a_t) + \beta EV(P_t, H_t, L_t, a_t) \} \quad (5)$$

where the function

$$EV_\theta(P_t, H_t, L_t, a_t) \equiv \int \int \int V(P_{t+1}, H_{t+1}, L_{t+1}, \epsilon_{t+1}(a_{t+1}) | P_t, H_t, L_t) dG(\epsilon_{t+1}(0))G(\epsilon_{t+1}(1))F(P_{t+1} | P_t, H_t) \quad (6)$$

Note that because of the independence assumption, the current draw of the error does not directly affect the future payoffs; it only affects the future through its impact on the current choice. Since the value function does not have an analytic solution, it is solved for a discrete set of values in the state space¹³. Also since the econometrician does not observe $\epsilon_t(a_t)$, it is convenient to represent the solution to the value function in terms of choice specific value functions. Let the function V_θ^1 represent the value for the firm of operating while V_θ^0 represents the value for the firm of not operating, both net of the structural errors. These can be viewed as the value a firm sees from operating and not operating before realizing the draws of $\epsilon_t(a_t)$.

$$\begin{aligned} V_\theta^0(P_t, H_t, L_t) &= \Pi(P_t, L_t, a_t = 0) + \beta EV(P_t, H_t, L_t, a_t = 0) \\ V_\theta^1(P_t, H_t, L_t) &= \Pi(P_t, L_t, a_t = 1) + \beta EV(P_t, H_t, L_t, a_t = 1) \end{aligned} \quad (7)$$

The optimal policy of the firm will be to choose the option that gives the greater value after accounting for the realized draw of the error.

$$a_t^*(P_t, H_t, L_t, \epsilon_t(0), \epsilon_t(1)) = \operatorname{argmax}_{a_t} \left\{ \begin{array}{l} V_\theta^0(P_t, H_t, L_t) + \sigma\epsilon_t(0), \\ V_\theta^1(P_t, H_t, L_t) + \sigma\epsilon_t(1) \end{array} \right\} \quad (8)$$

The value function then is the value created by making the optimal choice. This is identical to equation 5, only written in terms of choice specific value functions.

$$V_\theta(P_t, H_t, L_t, \epsilon_t(a_t)) = \max_{a_t} \left\{ \begin{array}{l} V_\theta^0(P_t, H_t, L_t) + \sigma\epsilon_t(0), \\ V_\theta^1(P_t, H_t, L_t) + \sigma\epsilon_t(1) \end{array} \right\} \quad (9)$$

¹³The states H_t and L_t are already discrete, but P_t must be discretized or at least evaluated at a discrete set of points. The resulting state space could be quite large depending on how finely P_t is discretized. The dimension of H_t is 24 since there are 24 operating hours in each day. The operating state last period, L_t , is a binary outcome. The size of the state space is then $DP * 24 * 2$ where DP is the number of discrete prices used. For one hundred discrete prices, the total size of the state space would be 4,800 which is large, but computationally feasible.

Since the econometrician does not observe the error, the optimal policy can be viewed as a probability that the firm will operate in a given state of the world. As the errors are distributed as extreme value type I, the operating probability for a generator has the following analytic form.

$$p(a_t = 1|P_t, H_t, L_t; \theta) = \frac{e^{\frac{V_\theta^1(P_t, H_t, L_t)}{\sigma}}}{e^{\frac{V_\theta^0(P_t, H_t, L_t)}{\sigma}} + e^{\frac{V_\theta^1(P_t, H_t, L_t)}{\sigma}}}$$

Conditional on a set of parameters, the probability of operation can then be used to construct a likelihood function for a generator.

$$L(\theta) = \prod_{t=1}^{t=T} p(a_t|P_t, H_t, L_t; \theta)p(P_t|P_{t-1}, H_{t-1}) \quad (10)$$

Here $p(P_t|P_{t-1}, H_{t-1})$ is derived from the conditional distribution $F(P_{t+1}|P_t, H_t)$, which is the probability of transitioning from one discrete price to another given the interval of the day. It should be noted that $p(a_t|P_t, H_t, L_t; \theta)$ implicitly depends on the transition probability matrix given by $p(P_t|P_{t-1}, H_{t-1})$ through the solution to the value function.

Since the transition probabilities do not depend on the vector of unknown parameters θ , they can be flexibly pre-estimated outside of the likelihood function. The simplified likelihood function can then be written as simply a function of the operating probability in each period.

$$L(\theta) = \prod_{t=1}^{t=T} p(a_t|P_t, H_t, s_t; \theta) \quad (11)$$

4 Estimation

The vector of unknown structural cost parameters $\theta = (START, \sigma)$ for each generator is estimated via maximum likelihood. While conceptually straightforward, solving for the structural parameters which maximize the likelihood function can be quite computationally intensive.

4.1 Price Transition

A necessary input for the maximization of the likelihood is a set of price transition matrices which capture the firm's expectations about future prices at any state. Since the price transitions do not depend on the action of the firm in a price taking model, the transition matrix can be estimated outside of the likelihood function. Given that the conditional transition probabilities, $p(P_t|P_{t-1}, H_{t-1})$, depend on both the last period's price and the hour of the day, the size of each time-specific matrix depends entirely on how finely price is discretized. For example, if price were discretized into 100 bins, then each transition matrix has 10,000 elements. With 24 intervals in each day, this means that 240,000 conditional probabilities would need to be estimated. The large number of conditional probabilities renders nonparametric estimation of the transition matrices infeasible even for very modest levels of price discretization.

Consequently, a semi-parametric method is used to estimate the conditional probabilities. I regress next period's price on a cubic polynomial expansion of the current price with dummies for each hour of the day.

$$P_{t+1} = \beta_0 + \beta_1 P_t + \beta_2 P_t^2 + \beta_3 P_t^3 + \mathbf{D}\alpha_0 + P_t \mathbf{D}\alpha_1 + P_t^2 \mathbf{D}\alpha_2 + P_t^3 \mathbf{D}\alpha_3 + \epsilon_t. \quad (12)$$

where \mathbf{D} is a vector of 24 hour of day dummies. The interactions between the cubic polynomial in prices and the hour of day dummies allow the coefficients to be completely flexible across each hour of the day. This is important because the current price for electricity may have different implications for future prices depending on the time of day. For example, observing \$50/MWH at 5:00pm may mean that prices will be even higher next hour, while the same price late in the evening may signal that prices can be expected to decrease the following hour. Since price expectations are a critical component of the firm's dynamic programming problem, it is important that the model of beliefs appropriately captures the dynamic nature of prices. The performance of this specification as a model of beliefs, as well as robustness of the results to alternative specifications, is explored in detail in appendix B.

The parameter estimates from the above equation yield $E[P_t|P_{t-1}, H_{t-1}]$ ¹⁴. To create a conditional distribution around $E[P_t]$, I use the empirical distribution of the residuals from the estimation procedure. The probability matrix $p(P_t|P_{t-1}, H_{t-1})$ is then created by integrating over the errors for each discrete price.

4.2 Structural Parameter Estimation

Once the transition matrix is defined, the structural parameters can be estimated using the dynamic model. Models of this type are traditionally estimated using a nested fixed point method. Typically the dynamic programming problem is solved using value function iteration while searching over the parameter space. Alternatively, two-step dynamic estimators can be used with much lower computational overhead. (V.J. Hotz & R.A. Miller 1993, V. Aguirregabiria & P. Mira 2002, P Bajari, L Benkard & J Levin 2007). Two-step methods leverage the equilibrium observed in the data to estimate the structural parameters. While this approach would be perfectly appropriate estimating the model, these estimators do not allow for the computation of counterfactuals. Since counterfactual computation is a key part of the research question, two-step methods are not employed.

Applying the nested fixed point method using value function iteration is not computationally feasible for this model. The solution time for value function iteration, which uses the contraction mapping property of the Bellman equation, depends on the discount factor β . As β nears one, the time to convergence

¹⁴The parameters are estimated using the observed continuous prices, and then the conditional probabilities are calculated given the number of discrete prices. Alternatively, the price space could be discretized before estimating the parameters with some loss of precision. As the number of discrete prices increase, the results would converge to the continuous price estimates.

increases exponentially. Since a generator’s operating choice and price are observed every hour, the discount factor in the model is very close to one rendering value function iteration impractical¹⁵.

Instead, policy function iteration is applied inside the nested fixed point algorithm to solve the dynamic problem and estimate the parameters. Importantly, the solution time for solving the value function by policy function iteration does not depend on the discount factor. However this method does require inverting a potentially large probability matrix, which in some cases may be computationally infeasible. However, this particular application is able to take advantage of the fact that the transition matrix is quite sparse, reducing the computation time significantly. Importantly, I can use the same solution method to solve for counterfactual equilibria later.

Additional data, outside of the four-month time window, from May 2006 through August 2006, is available for estimation. This subset is used for several reasons. First, within this period fuel costs are relatively constant. With fixed fuel costs, they do not enter the state space of the dynamic problem. If a longer time period were used with changing fuel costs, the model would need to explicitly incorporate each firm’s expectations for future fuel costs by introducing it as a new dimension into the state space. Second, by focusing on a compact time period, I am able to use a relatively simple, uniform price process over the estimation sample. Extending the dataset would necessitate expanding the state space to account for seasonality, new generator capacity, and demand growth as these would change the price transition probabilities. The increased size of the state space would make computation infeasible. Third, the variability in demand, and thus prices, during this time period is sufficient to induce the generators to start up and shut down. Prices need to be high enough in some periods such that high cost generators startup and also need to be low enough in other time periods such that lower cost generators shut down. There is not sufficient price variability in other time periods to estimate startup costs. Finally, by using these months of data, I avoid maintenance periods for generators. When generators have scheduled maintenance, their hourly operating decisions are not motivated by price signals. Interpreting maintenance periods as a reaction to current prices could lead to overestimating start-up costs. Using a four month time period prevents the model from becoming overly complicated and allows for relatively clean and simple estimation of the structural startup cost parameters.

4.3 Identification

The arguments for the identification of the structural parameters are fairly straightforward. First, the generator’s start up cost is identified by the difference in the willingness to operate between two states with the same price and interval, but with a differing operating state last period. Consider the

¹⁵If the annual discount rate is 0.95, this translates into a discount factor of approximately 0.9999945 every hour.

price/interval combination ($P_t = 50, H_t = 20$). The start-up cost is identified by the difference in the firm's behavior at ($P_t = 50, H_t = 20, L_t = 1$) versus ($P_t = 50, H_t = 20, L_t = 0$). Start-up costs imply that the probability of operation will be higher in the first case. In a world with no start-up costs, the behavior of the firm would be identical when faced with either of those states. The scale of the variance, σ , of the cost shock is identified by the willingness to operate in states outside of the cutoff rule implied by the deterministic model. More extreme or frequent deviations from the cutoff rule imply a higher σ and thus a wider distribution of start-up and shutdown costs.

The parameters for certain generators will not be identified. In order for start-up costs to be identified, a generator needs to turn on/off voluntarily in response to price signals. Some baseload generators, such as nuclear plants or some coal generators, may only shutdown for scheduled maintenance or an equipment breakdown. For such generators, start-up costs cannot be point identified, although a lower bound on start-up costs might be obtained. A lower bound would be identified by the lowest levels of observed prices under which the generator continues to operate. This paper does not attempt to bound start-up costs on baseload generators, but rather uses calibrated parameters for these generators.

5 Results

The estimation and counterfactual efforts are focused on a representative subset of the more than three hundred generators in ERCOT. Although the estimation and counterfactual methods in this paper are feasibly applied to the entire grid, the primary purpose of this paper is to explore the role of dynamics over a wide range of possible counterfactual scenarios. Focusing on a representative subset of generators enables a much wider set of counterfactuals to be explored. Specifically, parameters are estimated for all generators in ERCOT's West zone.

The West zone encompasses 22 fossil fuel generators representing all major technologies and is the smallest of ERCOT's four congestion zones¹⁶. As is the case for ERCOT as a whole, most generating capacity in this zone is gas fired and includes both combined cycle and simple cycle gas generators. Table 7 compares the composition of fossil fuel generators in the West zone to the ERCOT grid as a whole. Overall, the West zone has less coal capacity but more gas capacity, particularly gas turbine capacity, than ERCOT as a whole. The West zone looks roughly similar to the rest of the grid in terms of production, though the West zone still utilizes gas plants more than ERCOT as a whole.

¹⁶Wind generators are not included in the model since they lack the capability to increase production in response to price variations. They also do not reduce output during low price periods since their marginal cost of production is near zero. I also exclude one small hydroelectric plant for the analysis because it cannot increase aggregate production. Nuclear power plants are not considered in this paper. With high start-up costs, low marginal costs, and zero carbon emissions, their operation is unlikely to be changed in counterfactual policies as they are already operating at maximum capacity.

Table 8 shows detailed generator-level characteristics of generators in the West zone.

Table 7: Grid Generator Characteristics

Technology	Capacity		Generation	
	West	All Texas	West	All Texas
Gas GT	22%	9%	2%	3%
Gas ST	30%	28%	11%	8%
Gas CC	34%	40%	54%	44%
Coal	14%	23%	33%	45%

The parameters for each generator in the zone are estimated independently¹⁷. Table 9 shows the estimates of start-up costs for all but two generators in the West zone¹⁸. The first two columns of the table show the estimated start-up costs and scale of the operating cost shock for each generator. Standard errors are shown in parenthesis below the estimates. The third column indicates what type of technology is used at the plant.

The estimates of start-up costs are higher than typically considered. The start-up costs estimated start at \$26,000 for gas turbines and are as high as \$460,000 for very large combined cycle plants¹⁹. Traditionally, engineering parameters for start-up costs have been roughly an order of magnitude smaller than these estimates. However, recent work in engineering has more rigorously measured increased maintenance, decreased plant life spans, and opportunity

¹⁷There is one exception to this rule. In order to model each generator as a single agent, each needs to be able to react independently to prices. Combined cycle gas generators violate this rule since they run multiple turbines in sequence. In a combined cycle plant, a simple combustion turbine is first used to burn the natural gas. The exhaust of this turbine is used to heat water which powers a secondary steam turbine. Thus, the operation of the steam turbine is closely linked to the operation of the combustion turbine. Some plants may have two or three combustion turbines which all feed a single steam turbine. Such plants can run in multiple configurations, such as with one or two combustion generators feeding the steam turbine rather than all three. Since the cost of starting up the steam turbine may be high, a plant may operate one gas turbine at its minimum capacity to avoid the start-up costs associated with restarting the steam turbine. If the gas turbine were modeled as a single agent, this would overstate the start-up cost of this generator. To alleviate this problem, I aggregate the output of all generators which are part of a combined cycle plant. In doing this I assume that the economically important start-up costs are incurred when the entire plant starts production and abstract away for ramping costs associated with the output capacity of the plant.

¹⁸Start-up costs could not be estimated for two of the generators. One coal plant never shut down during the sample; one gas plant never operated. Since I do not observe start-up or shutdown decisions for these plants, I cannot obtain a point estimate for their start-up costs.

¹⁹The fuel and emission segments of these start-up costs can be separated from other costs using EPA's Continuous Emissions Monitoring System (CEMS). The EPA tracks heat input and emissions output for generators on a continuous basis. Thus, it is possible to calculate average fuel usage and emissions releases over the period when a generator is starting up. The data reveal that the cost of fuel and emissions alone range from \$500 for small combustion gas turbines to \$55,000 for large gas steam turbines. The residual part of start-up costs must be attributed to long run maintenance costs or other costs associated with changing output.

Table 8: Generator Characteristics: West Zone

Name	Fuel	Type	Max (MW)	Min (MW)	Capacity Share	Generation Share
Calenergy	Gas	CC	212	84	4.4%	5.3%
Graham 1	Gas	ST	229	46	4.8%	1.4%
Graham 2	Gas	ST	377	26	7.8%	3.2%
Morgan Creek 5	Gas	ST	127	15	2.6%	0.1%
Morgan Creek 6	Gas	ST	450	90	9.4%	0.0%
Morgan Creek A	Gas	GT	83	30	1.7%	0.2%
Morgan Creek B	Gas	GT	85	30	1.8%	0.1%
Morgan Creek C	Gas	GT	83	30	1.7%	0.1%
Morgan Creek D	Gas	GT	85	30	1.8%	0.2%
Morgan Creek E	Gas	GT	83	30	1.7%	0.1%
Morgan Creek F	Gas	GT	84	30	1.7%	0.1%
Odessa-Ector	Gas	CC	960	145	20.5%	44.2%
Oklaunion 1	Coal	ST	630	312	13.1%	34.3%
Permian Basin 5	Gas	ST	116	7	2.4%	0.5%
Permian Basin 6	Gas	ST	492	45	10.2%	6.1%
Permian Basin A	Gas	GT	65	40	1.4%	0.2%
Permian Basin B	Gas	GT	65	40	1.4%	0.3%
Permian Basin C	Gas	GT	65	40	1.4%	0.2%
Permian Basin D	Gas	GT	65	40	1.4%	0.2%
Permian Basin E	Gas	GT	65	40	1.4%	0.1%
Sweetwater	Gas	CC	266	173	4.8%	2.6%
Wichita Falls	Gas	CC	75	22	1.6%	0.6%

costs of forced outages due to start-up and shutdown. These studies indicate that start-up costs can range from \$300-\$80,000 for gas turbines and \$15,000-\$500,000 for steam gas and coal plants(Phil Besuner & Steven Lefton 2006). The wide range of possible startup costs underscores the significant heterogeneity that exists across generators. Though high, the estimated startup costs are not outside of the range of reason. However, they still are higher than estimates of startup costs found in other papers²⁰.

The performance of the model is examined against the data by comparing the operating state implied by the model with that observed in the data for each generator and time period. Since the model gives only a probability of operating, a generator is considered "operating" for this exercise if the probability of operating is greater than 50%. Table 10 gives the proportion of periods where the operating predictions of the model and the observed behavior in the data are aligned. In only 8% of the periods do the predictions of the dynamic model differ from the observed behavior. For comparison we can contrast this with predictions from a static model, which has the same marginal costs for generators, but no startup cost. Without startup costs, the model predicts that a generator will operate if and only if the price of electricity is greater than its marginal cost. Applying the static model to the data shows that the predictions disagree with the data in 22% of the generator periods. Despite the fact that the only parameters estimated are start-up costs and the variance of the error, the dynamic model is able to fit the data reasonably well. I also examine the out of sample fit of the model, by using the model to predict outcomes in the month before and the month after the estimating sample²¹. The results, shown in table 11, indicate that the predictive power of the model holds even outside of the sample. Predictions of the dynamic model differ from the observed behavior only 7% of the time in out of sample periods.

²⁰In particular, Reguant (2012) finds average startup costs in the Spanish electricity market to be around \$30,000 USD for both coal plants and combined cycle gas plants. Her modeling approach is quite different as it uses bidding data combined with a 5-day finite horizon dynamic model to estimate both marginal costs and startup costs in a strategic setting. The standard errors on the estimates are relatively large, especially for the gas plants, due to their dependence on first stage estimates which have wide standard errors. The author urges caution in interpreting the point estimates for gas plants in particular. However even looking at the 95th percentile of the confidence interval of Reguant estimates would only imply a startup cost of a little more than \$100,000 for a combined cycle gas plant. This is still much lower than the estimates obtained in this paper.

²¹The same optimal policy function is used to make model predictions in sample and out of sample. The optimal policy is the solution to the dynamic programming problem at the estimated parameters with beliefs estimated from the estimating sample.

Table 9: West Zone Results

Unit	$START_i$	σ_i	Capacity in MW	Type
Calenergy	\$246,210 (313,630)	\$24,965 (32,232)	212	CC
Graham 1	\$59,727 (3,871)	\$10,419 (686)	229	ST
Graham 2	\$33,297 (2,337)	\$6,802 (429)	377	ST
Morgan Creek A	\$36,490 (4,275)	\$6,092 (670)	83	GT
Morgan Creek B	\$34,525 (3,802)	\$5,733 (598)	85	GT
Morgan Creek C	\$34,325 (3875)	\$5,599 (594)	83	GT
Morgan Creek D	\$33,460 (3,709)	\$5,632 (589)	85	GT
Morgan Creek E	\$35,056 (3,890)	\$5,782 (607)	83	GT
Morgan Creek F	\$36,767 (4,335)	\$6,064 (670)	85	GT
Morgan Creek 5	\$61,463 (5,860)	\$8,585 (873)	127	ST
Morgan Creek 6	— —	— —	512	ST
Odessa	\$461,330 (211,420)	\$67,006 (28,957)	960	CC
Oklaunion	— —	— —	630	COAL
Permian Basin A	\$26,456 (2,221)	\$4,533 (389)	65	GT
Permian Basin B	\$28,601 (2,425)	\$4,925 (425)	65	GT
Permian Basin C	\$30,637 (2,840)	\$5,146 (491)	65	GT
Permian Basin D	\$35,802 (3,570)	\$5,709 (589)	65	GT
Permian Basin E	\$43,755 (4,764)	\$6,409 (716)	65	GT
Permian Basin 5	\$183,210 (45,756)	\$31,178 (8,031)	116	ST
Permian Basin 6	\$119,520 (16,079)	\$21,847 (3,057)	492	ST
Sweetwater	\$56,857 (5,153)	\$10,996 (873)	226	CC
Witchita	\$58,006 (8,118)	\$10,042 (1,422)	75	CC

Table 10: Operating State Fit

Model		Data	
		Off	Operate
Dynamic	Off	0.66	0.02
	Operate	0.06	0.26
Static	Off	0.63	0.13
	Operate	0.09	0.15

Table 11: Out of Sample Fit

Time period		Data	
		Off	Operate
Month Prior	Off	0.78	0.01
	Operate	0.06	0.14
Month After	Off	0.76	0.01
	Operate	0.06	0.17

6 Counterfactual

Given estimated parameters, the structural model can be used to simulate operating decisions for any counterfactual path of prices. In particular, we want to simulate market behavior with equilibrium market prices that would emerge under potential environmental policies such as carbon taxes or renewable mandates. To do so, we need first to find the path of prices that constitutes an equilibrium in the electricity market. This section defines the equilibrium conditions and proposes a method for solving for such counterfactual price equilibria.

6.1 Counterfactual Supply

In this section, I describe how the model is used to construct the supply curve for any counterfactual price path for a given generator. First, we must create the price transition matrix characterizing firm's expectations for future prices at each possible state. Given the counterfactual price path \mathbf{P} , we can construct the beliefs, $f(P_t|P_{t-1}, H_{t-1})$, which are consistent with \mathbf{P} . Second, we solve the dynamic programming problem conditional on beliefs. The resulting optimal policy function dictates the optimal operating decision for any state which could be encountered. This optimal policy is denoted as $a_{it}(P_t, H_t, a_{it-1}, \epsilon_{it})$. The quantity of electricity supplied at a given state is simply the operating decision multiplied by the quantity produced. Recall, that the quantity produced depends on the price of electricity and the generator's marginal cost. The supply of the generator for any state is then:

$$s_{it}(P_t, H_t, a_{it-1}, \epsilon_{it}) = a_{it}(P_t, H_t, a_{it-1}, \epsilon_{it})q_{it}(P_t) \quad \forall t \in \{1, 2, \dots, T\} \quad (13)$$

While P_t and H_t are fixed for a given counterfactual price path, ϵ_{it} and a_{it-1} are ex-ante unknown. To determine the operating decision of the generator in each period, we need to know both of these values. Since the error is never observed, we need to draw from the distribution of the error for the generator in each time period. Finally, once the error is known, we can calculate optimal decision in the first period (a_{i1}) using an initial condition (a_{i0}). The optimal decision in subsequent time periods can then be determined by using the lagged operating decision in the previous time period. By solving forward in this manner we can determine the supply of the generator in each time period for the counterfactual price path \mathbf{P} . However, the decision of the generator will be specific to the particular sequence of draws of the errors used. It is relatively simple though to repeat the process using a new sequence of errors. The dynamic programming problem does not need to be resolved since the distribution of the errors and of the counterfactual price path have remained unchanged. Rather, the same optimal policy can be reseeded with new errors and solved forward in the same manner as before. In this way it is possible to simulate many potential decision sequences for the same generator without incurring a large computational penalty. In the counterfactuals presented later, I simulate one hundred

markets, each with unique, generator-specific draws for the errors and report average outcomes over markets.

The industry supply in each time period is then simply the sum of all the generator specific supply functions.

$$S_t(P_t, H_t, \mathbf{a}_{t-1}, \boldsymbol{\epsilon}_t) = \sum_{i=1}^N s_{it}(P_t, H_t, a_{it-1}, \epsilon_{it}) \quad (14)$$

It is important to note that the generator supply function, and thus the industry supply function, is not smooth or continuous in P_t . It is not smooth since it is a step function with steps at the firm's startup point and marginal cost. The generator will produce nothing until it starts up, at which point it will produce at least at its minimum level. Likewise, it will move from producing at its minimum level to its maximum level when price exceeds marginal cost. Only when price equals marginal cost, will the generator produce in between its maximum and minimum output level. The supply functions are not continuous since the firm cannot produce below its minimum output level. This leaves a gap in the supply function between the zero and minimum output level. Both of these characteristics will need to be taken into consideration when solving for a market equilibrium.

6.2 Counterfactual Demand

For the counterfactual simulations, a model of demand for electricity is needed. Rather than estimate a demand side model, a simple demand function is calibrated with parameters taken from the literature. Counterfactual price equilibria will be calculated using both an inelastic demand function and a calibrated demand curve which reflects the long-run response of consumers to electricity prices.

In the very short run, the demand for electricity is almost perfectly inelastic. This is because consumers of electricity generally face constant prices for electricity over some time horizon, ranging from one month to several years, which are invariant to changes in wholesale prices of electricity. Thus, consumers generally have no incentive, or even available information, to change consumption as wholesale prices change. This demand function highlights the ability of the supply side to reduce emissions in response to environmental regulation in the counterfactual simulations. It is also the demand side model that, from a conceptual view, is most consistent with the short run supply side model which holds generating capital fixed.

Although inelastic demand is a realistic assumption for the very short run, it will not fully capture the new market equilibrium which will determine the profitability for different technologies going forward. Even though consumers do not respond immediately to wholesale price changes, changes in the average wholesale price for electricity will eventually filter down to the prices consumers face. The long-run response of consumers to average electricity prices is characterized by a constant elasticity demand function with a calibrated elasticity

parameter. In particular, it is assumed that demand in each period is characterized by $D_t = K_t * p_c^\alpha$ where D_t is the observed hourly demand for electricity in time period t , p_c is the average price consumers face for electricity, α is the demand elasticity parameter, and K_t is a positive constant. It is assumed that consumers face the average wholesale price for electricity over the simulation period. The long literature on electricity demand reports a wide range of demand elasticities depending on the time horizon²². However, many studies identify the long-run elasticity for electricity demand to be somewhere around $\alpha = -0.7$ (Bohi 1981)(Espey & Espey 2004)(EIA 2008)

Using a longer-run demand elasticity is somewhat inconsistent with the supply model since the model assumes that the supply side is not able to adjust its capital; this implies that consumers can change capital much more quickly than electricity generators. However, just as inelastic demand gives a lower bound on short-run emissions reductions, long-run demand provides an upper bound on the emissions reductions that could be achieved by environmental policies holding electricity generating capital constant.

6.3 Counterfactual Equilibria

The equilibrium in the model is defined by a price vector \mathbf{P} that equates aggregate supply from the model with demand in each period. In addition, the firms' expectations for prices at each state are required to be consistent with distribution of equilibrium prices. Appendix D shows that there exists a unique market clearing price vector conditional on beliefs and discusses the existence of the equilibrium. Unlike general equilibrium price taking models, such as Hugo Hopenhayn & Richard Rogerson (1993), where the stationary equilibrium price is the same in each time period, here the equilibrium prices remain uncertain and fluctuate each period due to aggregate demand shocks. However, the equilibrium is stationary in the sense that for a given state, firms expect the same distribution for future prices regardless of how they arrived at that state.

In short, a candidate price path will be an equilibrium price path if:

1. Each firm is acting optimally with respect to its price expectations.
2. The equilibrium prices clear the market in each period.
3. Firms' expectations for prices are consistent with equilibrium price path.

²²Dynamics exist on the demand side which limit consumers' response to price changes in the medium run versus the long run. Just as owners of SUVs are temporarily "locked in" to a higher gas usage even as the prices of gasoline rise, consumers of electricity must make costly adjustments to capital in order to fully optimize with respect to prices. Purchasing more efficient appliances, upgrading heating/cooling systems, or insulating a home will allow consumers to reduce consumption more in the long run than in the short run given higher electricity prices.

6.4 Algorithm

In this section, I delineate the algorithm used to solve for the equilibrium in each counterfactual. Since realizations of the errors are necessary to calculate generator outcomes, I solve for equilibrium outcomes in one hundred simulated markets simultaneously. In each simulated market, I draw a sequence of errors for each generator. Although, the equilibrium in each market is solved separately, the beliefs are estimated using the distribution of prices in all simulated markets. This allows the beliefs to represent the complete distribution of equilibrium prices including the variation due to the draws of the errors. Estimating common beliefs for generators in all simulated markets also allows me to solve the dynamic programming problem just once for each iteration of the algorithm rather than solving the different dynamic programming problem for each simulation.

To solve for the equilibrium for a given counterfactual, I use the following algorithm. Let P^0 be a candidate equilibrium price vector for each simulated market.

1. Change structural parameters or demand function as determined by the policy change²³.
2. Draw a sequence of errors for each generator for each simulation.
3. Estimate the price transition matrix from the distribution of prices across all simulations, $p(P_t^0|P_{t-1}^0, H_{t-1})$.
4. Given the price transitions, solve the dynamic programming problem for each generator.
5. For each simulated market, choose a new path of prices, P' , starting with the initial conditions a_0 , such that supply equals demand in each period, $S_t(P_t|H_t, \mathbf{a}_{t-1}, \boldsymbol{\epsilon}_t) = D_t$.
6. Re-calculate D_t given the new average price, $E[P']$. The average price is taken across all simulated markets.
7. Re-estimate the price transitions $p(P'_t|P'_{t-1}, H_{t-1})$.
8. Return to 4 and iterate.

The algorithm reaches an equilibrium for a counterfactual when the price vector does not change between iterations. At this point, the price vector clears the market in each period and the firms are optimizing with expectations consistent with the price vector. It should be noted that convergence with this algorithm is not guaranteed, though in practice it seems to work well²⁴.

²³In this paper, marginal costs are the structural parameters that are changed when a carbon tax is implemented. For the wind power counterfactuals, the residual demand is reshaped by the amount of wind power increase.

²⁴The iterative price search is continued until the average difference in the price vectors between iterations is less than \$0.01.

Table 12: Equilibrium Price Dynamics Comparison

	Actual	Dynamic	Static
Mean	\$63.03	\$65.13	\$62.96
Total Variance	1154	1014	207
Within Day	928	934	196
Between Day	225	79	10

6.5 Counterfactual Results

The counterfactuals are simulated using the estimated parameters from the previous section²⁵. Given the parameters, the model can be solved for equilibrium market clearing price for any potential sequence of demand realizations provided that demand does not exceed the total capacity of the generators in the market.


To demonstrate the methodology, I first solve for counterfactual prices for the time period used for estimation. This helps to evaluate how well the estimated parameters together with the counterfactual solution method can predict the observed data. To do this, I take the aggregate production of all generators each hour as the realized "demand" and solve for the prices to meet this path of demand realizations. The results are shown in table 12. The first column shows the statistics for the prices observed in the data over the sample period. The second column describes the equilibrium prices solved by the counterfactual algorithm using the estimated structural parameters. The final column shows the equilibrium prices predicted by a static model where generators are deployed in merit order based on marginal cost. In the static model, generators have the same marginal costs as in the dynamic case, but no start-up costs. Without start-up costs, there are no dynamic implications for firms' decisions in the current period. Since their behavior is determined only by the current price, they will operate if and only if the price is greater than their marginal costs of production in that period. Although simple, this serves as a useful benchmark against which to illustrate the role of dynamics in the market. Additional details on the static counterfactual can be found in Appendix A.

We see that the dynamic model matches average equilibrium prices fairly well, but does predict that prices will be slightly higher (\$65) than what we observe in the data (\$63)²⁶. The variance of prices predicted by the dynamic model also matches observed prices quite well. The static model however vastly underestimates the price variance. Decomposing the variance into within day and across day variance shows a similar picture. The dynamic model does

²⁵For the two generators which did not operate or did not shut down during sample period, it was not possible to estimate parameters. These generators include one large coal generator which operated continuously throughout the estimation period and one older gas-steam plant which never operated. As a lower bound on the start-up costs of the coal plant, the estimated start-up costs of the largest gas plant were used. The older gas plant which never operated was not included in the simulation

²⁶The average price is the demand weighted average across all time periods. For the dynamic model, the demand weight average price is also averaged across market simulations.

Figure 4: Estimating Sample Equilibrium Price Histogram



particularly well in matching the within day fluctuations in prices, though it underestimates the across day variance. Prices from the static model have very low variance both within and across day when compared with the observed prices. Comparing the distribution of prices in figure 4 reveals that prices from the dynamic model, though not as diffuse as observed prices, exhibit similar tails in the high and low price regions. The model predicts prices well above \$200 and below \$0.

We now move from simulating the observed estimating sample to simulating counterfactual outcomes. The startup costs that are estimated using the four-month sample can be applied to other time periods if the start-up costs do not change over time. Thus, rather than focusing the counterfactual exercises on a four month time period, the entire year of 2006 is used as the simulation period. This allows the results to take into account seasonal variability in demand. For example, lower demand periods in the fall and spring may be more responsive to pricing carbon since there is unused capacity that can be reallocated. Likewise, wind farms are most productive in lower demand periods which will influence their impact on emissions and profitability. Simulating outcomes over an entire year provides a more complete picture of market equilibrium under counterfactual policies.

In order to avoid problems with an increasing state, the model is solved separately for each quarter of the year. In each quarter, firms are assumed to face a stationary distribution of prices, but the price distribution can vary flexibly across quarters reflecting changes in the level and distribution of demand across seasons. Also, fuel prices in each quarter are set to the average yearly fuel prices. This negates the need for firms to predict fuel prices²⁷. Relative fuel prices of gas and coal are important parameters in determining the impact of carbon pricing on outcomes. In 2006, the average price of gas was \$6.40/MMBTU while the average price for coal was \$1.49/MMBTU²⁸. Neither of these prices is in the extreme of its historical price distributions, and they are both within the range of long term forecasts for gas and coal prices(EIA 2012). Finally, the same set of generators is carried through the entire year, with no entry or exit.

For the results that follow, all outcomes are compared to a baseline specification of the model, reflecting a no-regulation scenario. In addition to the details above, it is necessary to choose a baseline sequence of demand realizations for each hour of each day (i.e. a sequence of 8750 demand realizations) that generators will face. For representativeness, observed system-wide ERCOT demand for each hour in 2006 is used as the foundation for the baseline, rather than demand only in the West Zone. Since total system-wide demand would outstrip the productive capacity of West zone generators, it is scaled by the

²⁷In 2006 there is actually relatively little price variation across quarters. The average price for delivered gas in quarters one through four is: \$7.23, \$6.10, \$6.14, \$6.12. The average price for delivered coal in quarters one through four is: \$1.48, \$1.49, \$1.43, \$1.55.

²⁸Recent developments in natural gas production in the US have caused gas prices in the US to fall dramatically. Whether this is a temporary or a permanent change remains to be seen. However, low gas prices mean that a carbon tax may have an impact at a lower price than under historical fuel prices.

ratio of total West zone generation to total system generation in 2006. The result is a representative distribution of demand realizations that can be met by the generators in the sample²⁹.

Wind power production in the baseline follows a similar strategy. Baseline wind power potential is calculated using observed system-wide wind power production with the same scaling factor as for demand. Since the technological characteristics of wind turbines allow them to costlessly change their utilization of available wind, they are not subject to the same start/stop dynamics as conventional generators. Consequently, wind farms will curtail production when the price drops below their effective marginal costs. The operational costs of wind turbines are close to zero, but with the addition of federal and state production subsidies, marginal costs are effectively negative. For the counterfactuals, a marginal cost of -\$30/MWh is used for wind farms to reflect the approximately \$20/MWh federal production tax credit and \$10/MWh state renewable energy credits.

The counterfactual policy environments examined in this section include explicitly pricing carbon dioxide emissions and increasing the share of electricity supplied by wind power³⁰. Counterfactual equilibria are solved for each scenario under both inelastic and elastic demand functions. Outcomes of interest include changes in emissions, profitability, and prices relative to the base case scenario of no policy intervention. The results show the expected impact policies in the short run and how they would affect the profitability of various technologies going forward.

6.6 Carbon Price Counterfactual

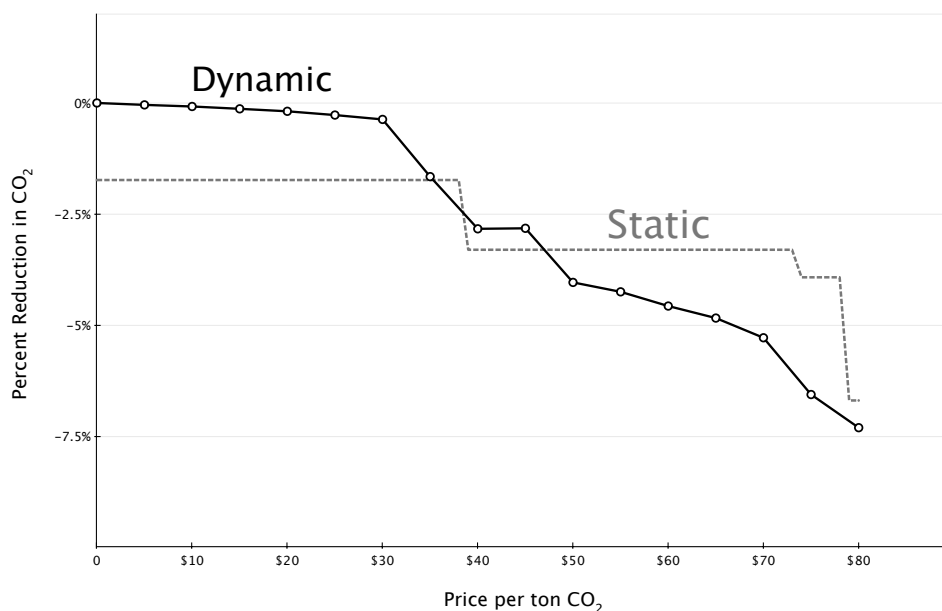
The first counterfactual examines outcomes under a price on carbon dioxide emissions ranging from \$0 to \$80 per ton of CO_2 in \$5 increments³¹. For each price point, the dynamic model was used to solve for the equilibrium price path that cleared the market in every hour over the year. Figure 5 shows the response of aggregate annual emissions in both the dynamic and static model to increasing prices for carbon using an inelastic demand curve. Note that the dynamic model has higher initial levels of CO_2 , but responds more quickly to carbon prices

²⁹Although observed demand in the West zone could be used as the baseline, it does not follow the same distribution as system-wide demand. Since the grid usually operates as a single entity, the prices that West zone generators face are in response to system-wide demand rather than localized demand in the West zone.

³⁰The model is agnostic about the source of carbon pricing which could be the result of a carbon tax or the equilibrium permit price under a cap and trade regime. Wind power development is assumed to be spurred by funding external to the electricity market such as the federal production tax credit subsidy. As such, increased renewable energy investment does not increase the cost of producing electricity. Rather, it exogenously arrives as the result of policy which incentivizes renewable energy investment. Modeling the investment process as a function of subsidization or incentives is outside the scope of this paper.

³¹Due to its computational simplicity, the static model was solved over the same range, but in \$1 price increments.

Figure 5: Carbon Counterfactual: CO_2 Response



than does the static model³². Although they follow similar trends, at some price points the reduction of carbon dioxide emissions in the dynamic model is nearly double that implied by the static model. In either case, carbon dioxide emissions are largely unphased by carbon prices at levels generally discussed in policy circles; a \$30 per ton price on carbon fails to reduce emissions by even 1%. Very little substitution occurs between high marginal cost, low emissions gas generators and low marginal cost, high polluting coal generators due to the large initial marginal cost advantage enjoyed by coal plants. A moderate carbon price still leaves coal plants as the low cost producer. However, even at \$80 per ton, aggregate carbon dioxide emissions fall by less than 7%. This can be compared to the maximum possible emissions reduction that could occur if cleaner generators were always deployed first regardless of cost. In this case, the short-run emissions reduction would be 23%. This implies that an \$80 carbon price captures one third of the total possible emissions reduction using existing generating capital.

While short run emissions show little response to increasing carbon prices, profits tell another story. Even though a \$20 carbon price has negligible impact of CO_2 emissions, it has a significant effect on the profitability of different gen-

³²Initially higher levels of CO_2 emissions in the dynamic model are driven in part by the higher underlying price volatility due to dynamics. True to observed behavior, lower efficiency simple cycle gas turbines participate in the market in the dynamic baseline simulation, but they do not in the static baseline simulation.

Table 13: Carbon Price: Profits and Prices

	Carbon Price			
	\$0	\$20	\$40	\$80
Δ Profit:				
Coal	–	-27%	-50%	-83%
Combined Cycle Gas	–	6%	30%	88%
Steam Gas	–	-1%	2%	3%
Gas Turbine	–	-2%	-3%	-3%
Dynamic Prices:				
Average \$/MWh	\$63	\$73	\$86	\$113
Percent Change	–	15%	35%	78%
Variance	\$328	\$350	\$472	\$410
Within-Day	\$254	\$271	\$357	\$305
Across-Day	\$74	\$79	\$114	\$104
Static Prices				
Average \$/MWh	\$64	\$75	\$87	\$115
Percent Change	–	18%	36%	81%
Variance	\$115	\$158	\$197	\$112
Within-Day	\$62	\$86	\$106	\$58
Across-Day	\$52	\$72	\$91	\$53

erating technologies. At \$20/ton, the profits of coal plants drop by almost 25%, while at the same time efficient combined cycle plants experience a significant bump in profitability as shown in Table 19. This underlies the long run implications of carbon pricing; even if current production decisions remain essentially unchanged, firms may make very different future investments. Increasing the carbon prices beyond \$20 per ton quickly erodes most coal profits and is a boon to most gas generators. Startup costs paid by coal double with a \$40 carbon price and quadruple at \$80, but play a relatively small direct role in reducing profits. Increased startup costs account for roughly eight percentage points of the 83% decrease in profits for the \$80 carbon price scenario. Interestingly, the model predicts that inefficient simple cycle gas turbines experience minor drops in profitability as their marginal costs further separate from those of more efficient generators.

Equilibrium prices show corresponding increases due to pricing carbon. Average annual prices increase from \$63/MWh to \$73/MWh with a price on carbon of \$20 per ton. Since demand is inelastic in this counterfactual, the increased price of carbon is completely passed through to consumers. However, since many generators emit less than one ton of CO_2 per MWh, the average price does not increase by \$20. A very high carbon tax of \$80/MWh would increase average prices by \$50. For comparison, the equilibrium price characteristics for the static model are also shown. The average prices in static model are nearly identical to those of the dynamic model. However, the dynamic model exhibits significantly higher price variation than the static formulation. The higher variance

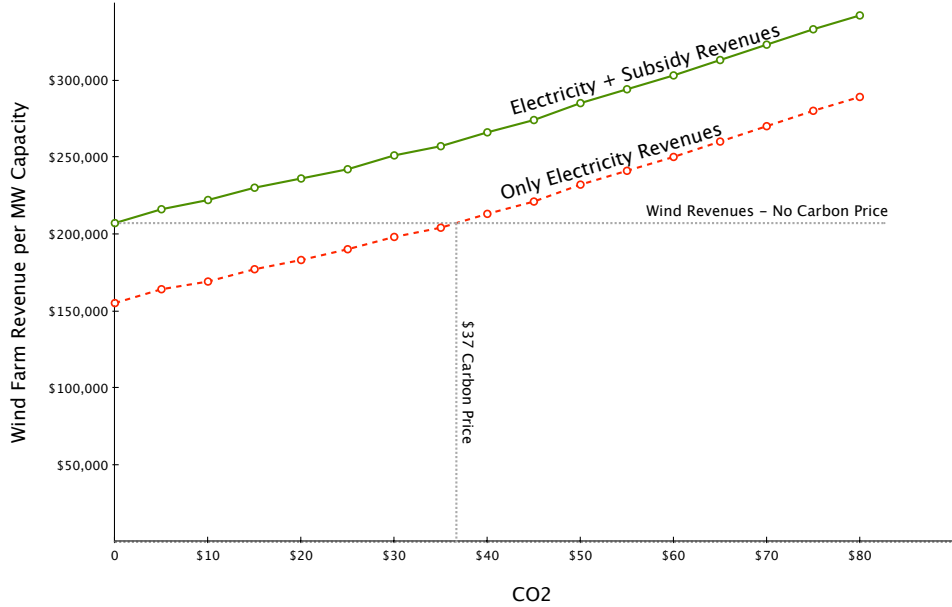
comes mostly from greater within-day price variation rather than across-day price variation. This is reflective of the need for firms to startup and shut down to follow daily load fluctuations. Higher price variation is needed to induce this behavior in the presence of startup costs. However, the higher peak prices and lower off peak prices lead to roughly similar average prices.

Wind revenues also steadily increase as a function of carbon prices. Figure 6 shows the increase in annual wind farm revenues per MW of capacity. The upper line shows revenues from both electricity sales and subsidies, while the lower line shows revenues from electricity sales only³³. The importance of subsidies for a wind generator's bottom line is immediately apparent from the graph. The results of the model identify the "break-even" carbon price for wind farms under counterfactual equilibrium prices. Without subsidies, it would take approximately a \$37/ton price on carbon to leave installed wind generators with the same revenues that they currently receive under a policy of subsidization.

Since the estimated startup costs are higher than expected, we might wonder whether the results will hold with lower start up costs. Also, even if the estimated startup costs are appropriate, firms may be able to invest in technology to lower their startup costs if it is profitable to do so. In either case, we would want to understand how the counterfactual results change with lower startup costs. To explore the sensitivity of the results to the estimated startup costs, I solved for counterfactual equilibria using much lower calibrated startup costs. The calibrated counterfactual results show a pattern of emissions reduction which is very similar to that found under the estimated parameters. The calibrated parameters and discussion of the full results can be found in appendix E.

³³Although wind farms in Texas receive approximately \$30/MWh in subsidies, the federal portion of the subsidy expires after the first 10 years of operation. Under the assumption that wind farms continue to operate after federal benefits expire and continue to receive a state subsidy, the discounted cost of the subsidy is approximately \$20/MWh over the life of the wind farm. This is the value that is used when comparing annual revenues with and without subsidies.

Figure 6: Carbon Counterfactual: Annual Wind Revenues



6.7 Wind Development Counterfactual

The model is also used to solve for equilibrium prices as the share of electricity produced by wind increases. Counterfactual equilibria are calculated as the potential share of electricity produced by wind increases from the observed level of 2% up to 30% of the total. Since wind power installations already exist in Texas, the production profile of those wind farms can be used to simulate the production of additional wind farms.

In 2006, wind farms accounted for 2% of electricity production in ERCOT. To construct counterfactual wind production, observed wind production is simply scaled up. For example, for the 10% wind counterfactual, wind production in each period over the year is scaled up until the total potential wind production could be 10% of annual electricity demand³⁴. Even if new wind farms are less productive than existing wind farms, due to being placed on less desirable properties, scaling the production patterns of existing wind farms will provide a good approximation of continued build-out as long as the diurnal and seasonal patterns of wind production are similar.

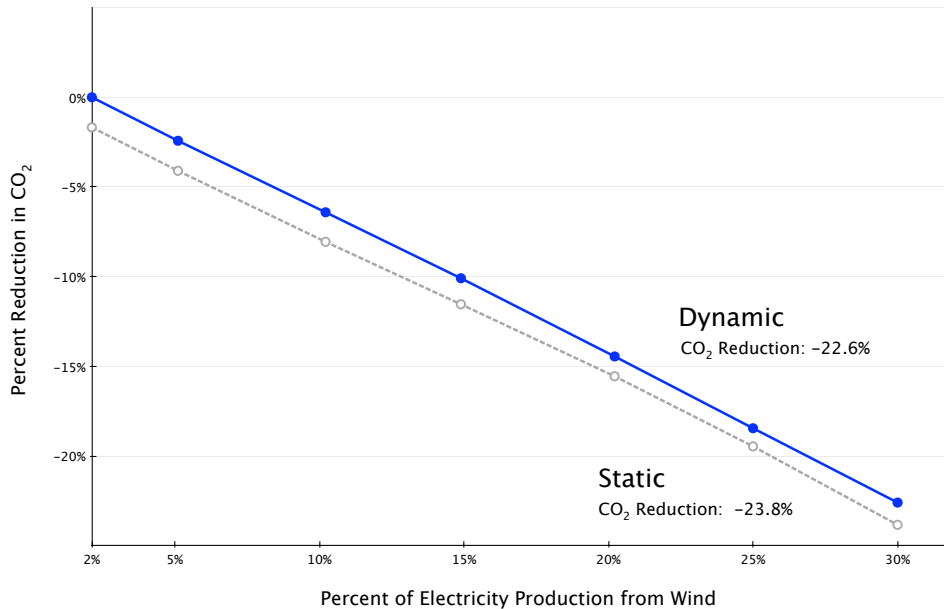
Wind farms in this region exhibit electricity production patterns that are heavily skewed toward off-peak power production as demonstrated in Cullen (2013). It is typical for on-shore wind farms to have significantly higher levels

³⁴Note that in equilibrium, wind production might not be the same as potential wind production as wind farms will curtail production if prices dip low enough.

of production at night and in the spring and fall, when demand for electricity is at its lowest levels. This pattern of production will exacerbate the volatility in the residual demand curve and resulting equilibrium prices, especially in the presence of start-up costs.

Figure 7 shows the change in CO_2 emissions as wind increases its share of production for both the static and dynamic models. Surprisingly, for the estimated start-up costs and sample of generators in the counterfactual, both models show similar reductions in CO_2 as wind production increases, despite the increased volatility in residual demand. The inclusion of dynamics does not seem to reshape the response of aggregate emissions to increased wind production. When wind produces 30% of electricity, carbon dioxide emissions are reduced by approximately 23% in either model. This is more than three times the reduction achieved by fuel switching with an \$80 price on carbon.

Figure 7: Wind Counterfactual: CO_2 Response



Profits and prices also significantly change as a result of wind production. Table 14 shows that profits for all fossil fuel generators fall across the board as a result of increased wind production. However, technologies with larger start-up costs fare far worse than more flexible technologies such as gas turbines. Prices also fall as the result of more wind power production by anywhere from 6% up to 29%. Importantly, the dynamic and static models have different price distributions. Price changes are steeper and price volatility is much higher in a dynamic model than with a static model.

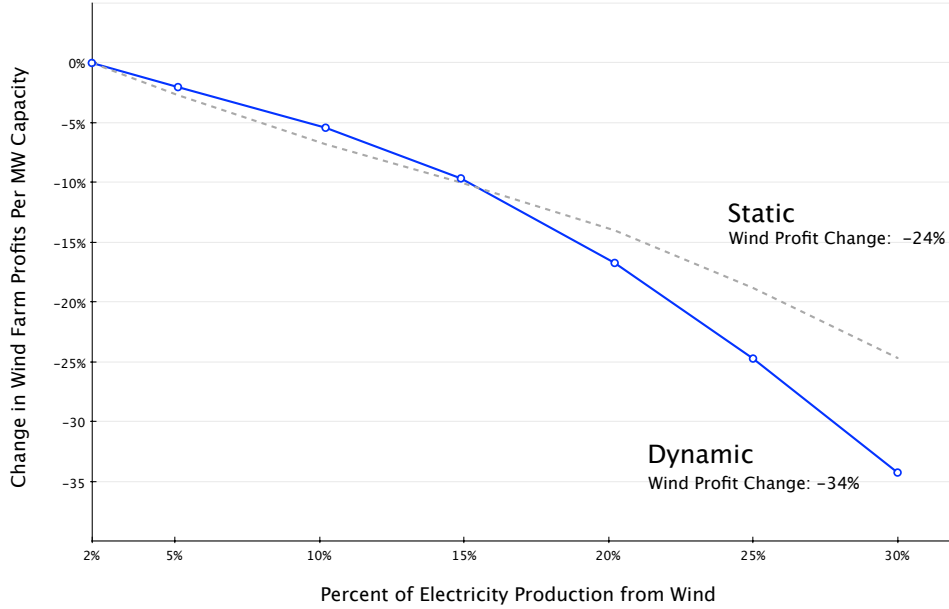
Price volatility has important implications for the profitability of wind farms,

Table 14: Wind Production: Profits and Prices

	Wind Production		
	10%	20%	30%
Δ Profit			
Coal	-9%	-22%	-37%
Combined Cycle Gas	-20%	-40%	-54%
Steam Gas	-13%	-28%	-41%
Gas Turbine	-6%	-13%	-21%
Dynamic Prices:			
Average \$/MWh	\$60	\$53	\$45
Percent Change	-6%	-16%	-29%
Variance	\$333	\$488	\$957
Static Prices:			
Average \$/MWh	\$60	\$55	\$48
Percent Change	-6%	-13%	-23%
Variance	\$144	\$221	\$561

as illustrated by figure 8. As wind accounts for larger shares of electricity production, the increased price volatility increases dramatically due to start-up costs as shown in table 14. Prices are pushed to lower levels in off-peak periods to get conventional generators to shut off production and raised in peak periods to get generators to meet increased demand. Since wind farms are most productive in off-peak periods and cannot shift their production to high price periods, their profits are much lower than would be implied by a static model. Operational dynamics reshape the potential profits from renewable technologies leading to different investment incentives.

Figure 8: Wind Counterfactual: Annual Wind Revenues

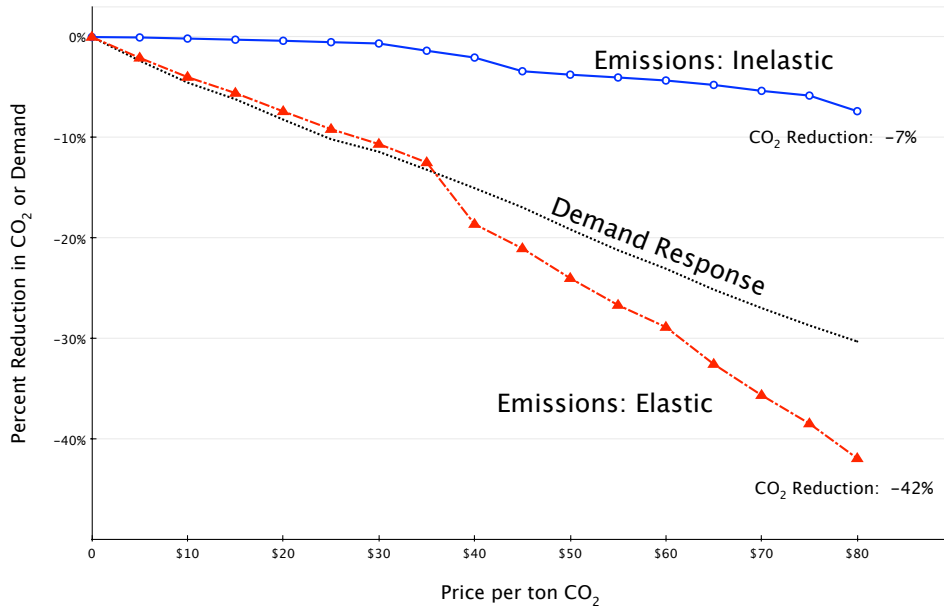


6.8 Demand Response

All the results presented so far have assumed that demand is inelastic. However, as the equilibrium price for electricity changes, the quantity demanded will change. This is likely to be an important factor in overall emissions from the sector, especially for high carbon prices. In this section, I examine the implications of the demand response on profits and emissions. As laid out in section 6.2, the demand function is characterized by consumers responding to average electricity prices with a calibrated demand elasticity of -0.7 . Figure 9 compares the former results with CO_2 reductions when firms face the elastic demand curve. The impact of CO_2 pricing is dramatically different as compared with an elastic demand curve. Emissions show an immediate response that grows quickly as the price on carbon increases. A \$20 carbon price reduces emissions by 7%, the same amount as an \$80 price with an inelastic demand curve. However, the reduction in emissions if coming exclusively from reduced demand for electricity rather than lower emission rates from generators. In fact, the average emission rate across generators for low carbon prices increases slightly since the reduced demand cuts first into cleaner gas production. Unsurprisingly, the profits of fossil fuel generators also decline more sharply with an elastic demand curve. At \$20 per ton of CO_2 , coal generators experience a 31% drop in profitability. The profits of higher cost steam and gas turbine generators also drop due to declining demand. However, efficient low emission combined cycle generators have

relatively flat or increasing profits making them more attractive investments.

Figure 9: Carbon Counterfactual: Demand Response



The impact of wind generation on emissions also changes significantly with an elastic market demand curve, as shown in figure 10. Emissions decrease by 14% rather than 22% when wind accounts for 30% of production. This is a direct result of increased demand for electricity due to lower electricity prices when wind development is funded from external sources. This demand rebound highlights the importance of incorporating the costs of renewable into electricity prices when working to achieve emissions reductions.

This also underscores the advantage carbon pricing enjoys in properly aligning incentives. Without accounting for demand adjustments, a 15% share wind power seems to be more effective at reducing emissions (-10%) than a \$20 price on carbon (<1%). However, once demand is allowed to respond to the equilibrium prices, a \$20 carbon price is more effective (-8%) than subsidizing wind power (-6%). Additionally, incentives to invest in low carbon fossil technology would only increase the advantages of pricing carbon over technology specific subsidies.

Figure 10: Wind Counterfactual: Demand Response

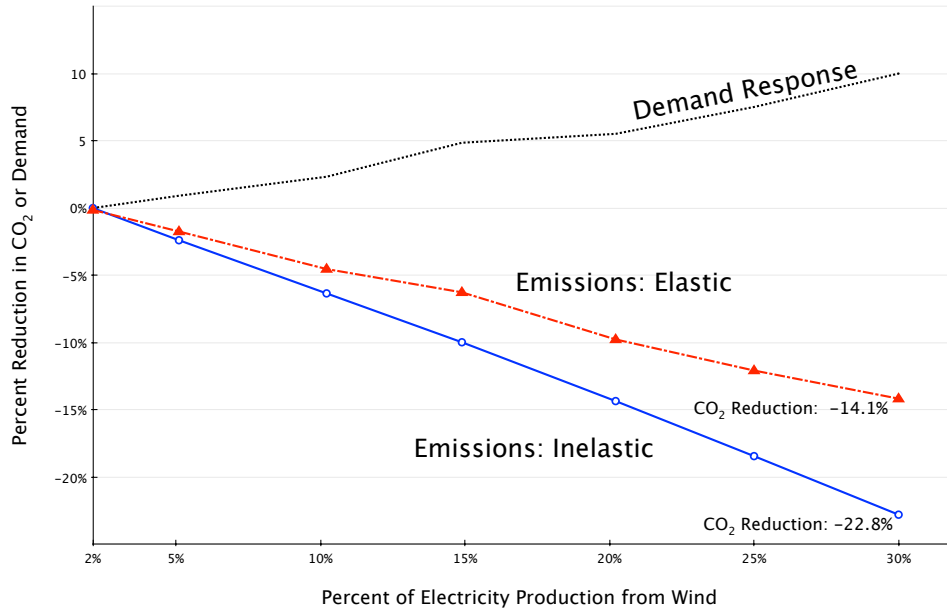


Table 15: Demand Response: Profits and Prices

	Carbon Price			Wind		
	\$20	\$40	\$80	10%	20%	30%
Δ Profit:						
Coal	-31%	-62%	-91%	-5%	-12%	-19%
Combined Cycle Gas	0%	-5%	32%	-11%	-18%	-21%
Steam Gas	-24%	-37%	-55%	-2%	-2%	9%
Gas Turbine	-9%	-14%	-21%	-1%	-2%	1%
Dynamic Prices:						
Average \$/MWh	\$71	\$79	\$103	\$60	\$57	\$53
Percent Change	8%	16%	40%	-4%	-8%	-15%

7 Conclusion

This paper builds a repeated entry/exit model in a competitive market where firms are persistent and their identities are known. The dynamic framework allows for the estimation of firm-specific distributions of entry (start-up) costs and facilitates counterfactual simulation of the market. The model aggregates the actions of a finite collection of dynamically optimizing single agents to solve for equilibrium market clearing prices such that firms' expectations for prices are consistent with the equilibrium distribution of prices.

The model is applied to the electricity industry in order to understand the short-run implications of potential environmental policies for equilibrium electricity prices, emissions, and firm profits³⁵. These changes in short-run emissions and profitability changes will shape future investment decisions.

Using a portion of the Texas grid as a test bed, start-up cost distributions are estimated for each individual electricity generator. Counterfactual simulations using the estimated parameters indicate that emissions from electricity production are largely unresponsive to pricing carbon at levels discussed in policy circles. A price of \$20/ton of CO_2 results in negligible changes in carbon emissions when demand is inelastic. However, the profitability of dirty generators is substantially reduced, and profits for clean generators increased, for even modest prices on carbon.

Incentives which spur the development of wind farms lead to more immediate reductions in emissions. However, increased wind power production lacks the same long-run incentives for technology switching that a carbon price provides. It decreases the profitability of cleaner generators more than that of dirtier ones. In addition, emissions reductions are substantially eroded away by a rebound in demand in the long run when funded by external subsidies.

Incorporating the short-run dynamic considerations facing generators yields different outcomes in a number of important aspects. First, equilibrium prices are much more volatile due to dynamics. This is particularly true as renewables gain higher market shares. Second, due to high price volatility, the profits that wind farms earn are substantially lower when dynamics are accounted for. Finally, emissions reduction in the static and dynamic frameworks do differ, though such differences were not dramatic for the estimated parameters used in the simulations.

The results of this analysis indicate that we should not expect much from pricing carbon emissions in the short run. However, even moderate carbon pricing impacts the profitability of high carbon technologies profoundly. For policy makers, this implies that meaningful reductions in carbon dioxide emission may be able to be achieved without drastic environmental policies, but a longer time horizon will be required to realize those benefits.

³⁵For any potential environmental policy, there will be short-run and long-run implications for electricity generators. In the short run, firms will re-optimize their day-to-day operating decisions given their existing generating capital. In the long-run, firms will make investment decisions based on how their profitability has changed under the new policy environment.

References

- AEP.** 2014. “Issues in Electricity: Natural Gas.” <https://www.aep.com/about/IssuesAndPositions/Generation/Technologies/docs/NaturalGasTechnology.pdf>, Accessed:2014-08-05.
- Aguirregabiria, V., and P. Mira.** 2002. “Swapping the nested fixed point algorithm: a class of estimators for discrete Markov decision models.” *Econometrica*, 1519–1543.
- Bajari, P, L Benkard, and J Levin.** 2007. ““Estimating Dynamic Models of Imperfect Competition.”” *Econometrica*, 75: 1331–1371.
- Bartlett, M.S.** 1955. *An Introduction to Stochastic Processes with Special Reference to Methods and Applications*. Cambridge, UK:Cambridge University Press.
- Besuner, Phil, and Steven Lefton.** 2006. “The Cost of Cycling Coal Fired Power Plants.” *Coal Power Magazine*, 16–20.
- Bohi, Douglas R.** 1981. *Analyzing demand behavior: a study of energy elasticities*. John Hopkins University Press for Resources for the future.
- Borenstein, Severin, James B. Bushnell, and Frank A. Wolak.** 2002. “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market.” *American Economic Review*, 92(5): 1376–1405.
- Bushnell, James B., Erin T. Mansur, and Celeste Saravia.** 2008. “Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets.” *American Economic Review*, 98(1): 237–66.
- Chow, Joe, Kevin Ho, Xiaoqiao Du, H.S. Lee, and Micheal Pearson.** 2002. “Experience from extensive two shift operation of 680MW coal/gas-fired units at castle peak power station.” *Operation Maintenance and Materials Issues*, 1.
- Cullen, Joseph.** 2013. “Measuring the Environmental Benefits of Wind-Generated Electricity.” *American Economic Journal: Economic Policy*, 5(4): 107–33.
- Cullen, Joseph, and Erin Mansur.** 2014. “Complementarity of Shale Gas and Carbon Prices.”
- Cullen, Joseph, and Stan Reynolds.** 2014. “Dynamic Competition in Wholesale Electricity Markets.”
- EIA.** 2008. “Energy Information Administration: Annual Energy Outlook 2008.”

- EIA.** 2012. “Annual Energy Outlook 2012.” U.S. Energy Information Administration.
- Espey, James A., and Molly Espey.** 2004. “Turning on the Lights: A Meta-Analysis of Residential Electricity Demand Elasticities.” *Journal of Agricultural & Applied Economics*, 36(1): 65–81.
- Hopenhayn, Hugo, and Richard Rogerson.** 1993. “Job Turnover and Policy Evaluation : A General Equilibrium Analysis.” *Journal of Political Economy*, 101(5): 915–938.
- Hortacsu, Ali, and Steven L. Puller.** 2008. “Understanding strategic bidding in multi-unit auctions: a case study of the Texas electricity spot market.” *RAND Journal of Economics*, 39(1): 86–114.
- Hotz, V.J., and R.A. Miller.** 1993. “Conditional choice probabilities and the estimation of dynamic models.” *The Review of Economic Studies*, 497–529.
- Kehlhofer, Rolf, Rolf Bachmann, Henrik Nielsen, and Judy Warner.** 2009. *Combined-cycle gas and steam turbine power plants*. . 3rd ed., PennWell.
- Kumar, Nikhil, Philip M. Besuner, Steven A. Lefton, Dwight D. Agan, and Douglas D. Hilleman.** 2012. “Power Plant Cycling Costs.” National Renewable Energy Laboratory.
- Ljung, G. M., and G. E. P. Box.** 1978. “On a Measure of a Lack of Fit in Time Series Models.” *Biometrika*, 65(2): 297303.
- Mansur, Erin T.** 2008. “Measuring Welfare in Restructured Electricity Markets.” *The Review of Economics and Statistics*, 90(2): 369–386.
- NEA.** 2014. “Issues in Electricity: Natural Gas.” <https://www.oecd-neo.org/press/press-kits/economics-FAQ.html>, Accessed:2014-08-05.
- Potomac.** 2007. “State of the Market Report for ERCOT Wholesale Electricity Markets.” Potomac Economics.
- Reguant, Mar.** 2012. “Complementary Bidding Mechanisms: An Application to Electricity Markets.”
- Rust, John.** 1987. “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” *Econometrica*, 55(5): 999–1033.
- Wolak, Frank A.** 2000. “An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market.” *International Economic Journal*, 14(2): 1–39.

A Static Counterfactual

The static counterfactual is carried out in the following manner. In the static setting, firms do not incur startup costs. Rather, they operate with respect to their marginal cost only. This implies that generators will operate at maximum capacity whenever price is above marginal cost, will not operate whenever the price is below marginal cost, and will operate at some point within its operating range when price equals marginal cost. Firms are still bound by their minimum and maximum output constraints.

To solve for the static counterfactual, price is set each period such that demand is just met by the total output of all the generators. The market clearing price will be the marginal cost of the marginal generator. There are no issues with multiple equilibria as long as no two firms have the exactly the same marginal cost.

There do arise cases when the minimum output constraints for the marginal generator bind. In the case, production is reduced by the next highest cost generator until there is no longer excess supply. The market clearing price is still the marginal cost of the marginal generator. However, for the purposes of calculating profits, I assume that the grid operator will have to "make whole" any inframarginal generators which are operating at less than maximum capacity. That is, the generator would receive payments for its entire productive capacity when it is asked to back down from optimal maximum production. For example, suppose an inframarginal unit must reduce production to 80% of capacity in order to clear the market. The generator would like to produce at 100% capacity since price is above its marginal cost. Therefore, revenues would be calculated as if the generator were producing at full capacity.

B Dynamic Price Process

In a structural, dynamic model, firms' expectations over prices must be explicitly modeled. Ideally, one would use each firm's actual beliefs for the evolution of prices. Absent direct information about beliefs, I estimate beliefs under a rational expectations framework. That is, I use actual equilibrium prices to estimate firm beliefs over prices.

Expectations over the future distribution of prices are modeled using a Markov AR(1) process conditional on hour of the day.

$$F(P_{t+1}|P_t, H_t)$$

Here P is the market clearing price and H is the hour of the day in time period t . This parsimonious model of beliefs assumes that this period's price and hour are sufficient for predicting the distribution of next period's price.

The AR(1) process is motivated in part, by computational constraints. The dynamic model demands that the specification be as simple as possible. Each variable that is used to predict prices adds an additional dimension to the state space. An additional dimension in the state space exponentially increases the

computational burden of solving the dynamic programming problem. It is advantageous therefore to have the simplest possible model of beliefs that reasonably captures the evolution of prices.

One might argue that this simple Markov process is not sufficiently rich to accurately model the expectations of the firm. Indeed, firms have more information than simply the lagged price and time of day with which to form expectations for price in the next period. For example, firms may have expectations over future temperatures, load levels, and congestion. In addition they may use a long price history when predicting future prices. The extent to which the proposed model is adequate depends on the degree to which the current price encapsulates other information that the firm may have. To this end, we now evaluate the model for price expectations by first examining the price process itself, and then by testing the robustness of the structural parameter estimates to changes in the model.

B.1 Price Process

The conditional markov price process is implemented in using a semi-parametric functional form. We start with a cubic polynomial expansion of the current price. Then the polynomial is interacted with dummies for each hour of the day.

$$P_{t+1} = \beta_0 + \beta_1 P_t + \beta_2 P_t^2 + \beta_3 P_t^3 + \mathbf{D}\alpha_0 + P_t \mathbf{D}\alpha_1 + P_t^2 \mathbf{D}\alpha_2 + P_t^3 \mathbf{D}\alpha_3 + \epsilon_t.$$

where \mathbf{D} is a vector of 24 hour of day dummies. The interactions between the cubic polynomial in prices and the hour of day dummies allow the coefficients to be completely flexible across the hours of the day. The specification is essentially performing a separate regression of next hour's price on current price for each hour of the day. This is important because the current price for electricity will have different implications for future prices if we are in time period of increasing demand as opposed to decreasing demand. For example, \$50/MWH at 5:00pm likely means that prices will be even higher next hour since demand is usually ramping up at that time of day. The same price late in the evening would indicate a price lower than \$50/MWH the next hour since demand is generally decreasing. By estimating the price process in cubic form separately for each hour of the day, beliefs can be modeled simply, but flexibly.

The results show that the model explains the data reasonably well. First, the model fits the data fairly well with an adjusted R^2 of 0.71. The full results from estimating the price process are shown in table 17. Second, I examine the residuals to see if a further lag of price can explain the residual variation from the model. If an AR (1) process with P_t is not a sufficient, then a additional price information, such as P_{t-1} , should be able to explain the residual variation from the estimation. To do this, I apply the same flexible model by regressing the residuals, e_{t+1} , on a cubic polynomial expansion of P_{t-1} interacted with dummies for each hour.

$$Residual_{t+1} = \beta_0 + \beta_1 P_{t-1} + \beta_2 P_{t-1}^2 + \beta_3 P_{t-1}^3 + \mathbf{D}\alpha_0 + P_{t-1} \mathbf{D}\alpha_1 + P_{t-1}^2 \mathbf{D}\alpha_2 + P_{t-1}^3 \mathbf{D}\alpha_3 + \omega_{t-1}.$$

The lagged price explains very little of the residuals with an adjusted R^2 of 0.001. The lagged price variables are not jointly significant at even the 25% level. Finally, I test the residuals for white noise using two different tests. If the residuals can be characterized as white noise, then we can rule out serial correlation. Serial correlation might indicate a missing variable that could explain persistent trends in the price process. I first apply the Bartlett test which is a periodogram-based white-noise test (M.S. Bartlett 1955). We fail to reject the null hypothesis that the residuals are white noise with a P-value of 0.43. The second test used is the Ljung-Box test (G. M. Ljung & G. E. P. Box 1978). This test requires the econometrician to specify the number of lags to be used when testing for white noise. With 10 or 20 lags, we fail to reject the null of white noise at almost any significance level. With enough lags, however, we can eventually reject the null. (At 40 lags the null can be rejected at the 5% level). Though not entirely conclusive, these results suggest that serial correlation is not a first order concern. Altogether, the price process results show strong explanatory power with residuals that show evidence of being white noise and cannot be easily explained with a additional lag in prices.

B.2 Distribution of Expected Prices

The model of the price process is used to construct beliefs over possible future prices. This is represented as a price transition matrix that is then feed into the dynamic optimization problem. In constructing beliefs, not only are expected prices important, but also the distribution of expected prices are an important part of the decision making process. High variance in the distribution of future prices can create an option value for the firm. For example, if prices could be either very high or very low next period, a firm may want to wait to startup until it observes next period's price. Consequently, I want to allow for a flexible distribution of prices as well as for state-dependent heteroskedasticity.

State dependent heteroskedasticity is introduced by allowing the variance of the distribution of prices to vary by the hour of the day. Thus high demand-high price hour of the day can have a higher variance than more predictable low demand-low price hours. This is accomplished by separately estimating the variance of prices for each hour of the day. To do this, I regress next hour's price on current price separately for each hour of the day.

$$P_{t+1} = \beta_{i0} + \beta_{i1}P_t + \beta_{i2}P_t^2 + \beta_{i3}P_t^3 + \epsilon_t.$$

if time period t is in hour i . The point estimates are identical to the specification above that includes dummies. No dummies are necessary here since the parameters are already completely flexible across hours of the day.

Rather than assuming that the errors are normally distributed, I instead use the empirical distribution of the residuals in each hour to characterize the distribution around the predicted price. This allows not only the variance to be different for each hour, but also allows the entire distribution of prices to be different in each hour. Allowing both the variance and the shape of the

distribution to vary with the hour of the day introduces significant flexibility into the model of beliefs over prices.

B.3 Robustness of Structural Parameters

Ultimately the estimated price process is used to formulate beliefs for the dynamic model. In this section, I explore how changes in the model of beliefs affect the structural parameter estimates. Given the complexity of the model, it is not a trivial task to change the model of beliefs. If more information is used to construct beliefs, such as additional lags of price, then the state space of the model is also changed. Exploring many possible specifications for beliefs is not feasible. Instead, I propose a single specification that will be in some sense, the limit of many other specifications. By more richly specifying the model of beliefs, firms will have weakly better predictions of future prices. In fact, if firms had perfect information about everything in the market, they would, in theory, be able to perfectly predict the path of prices. Thus to test the robustness of the structural estimates to better predictions over prices, I estimate the model under the assumption that firms can perfectly predict prices. This perfect foresight scenario can be thought of as the limiting case for richer and richer specifications of beliefs. Since firms know the future prices, they can perfectly account for the dynamic implications on their decisions. Table 16 shows the structural parameter estimates from the perfect foresight model in column one compared with the estimates with the preferred specification for beliefs in column two. The results show that the estimates are very similar. In fact, for most generators, estimated startup costs with perfect foresight are greater than those under preferred model for beliefs. These results illustrate that the estimates of startup costs are being driven by the specification used for beliefs.

Table 17: Price Process

	AR(1) Price Process	
P_t	.785	(.413)
P_t^2	.00225	(.00168)
P_t^3	-7.72e-06	(.0000184)
Hour 2	-14.7	(23.4)
Hour 3	-12.3	(30.5)
Hour 4	-4.62	(31.3)
Hour 5	.537	(32.6)
Hour 6	-3.89	(39.5)
Hour 7	-22.6	(24.6)
Hour 8	-3.89	(15.1)
Hour 9	-10.6	(15.1)
Hour 10	-1.85	(17.6)
Hour 11	-21.1	(20.2)
Hour 12	-11.8	(35.7)

Hour 13	-20.7	(28)
Hour 14	176	(44.8)
Hour 15	8.75	(22.9)
Hour 16	-42.7	(23.3)
Hour 17	-72.5	(21.2)
Hour 18	-15.6	(21.7)
Hour 19	-51.6	(19.4)
Hour 20	-33.7	(21.9)
Hour 21	10.3	(38.9)
Hour 22	14.6	(53.5)
Hour 23	-12.8	(54.3)
Hour 24	-71.5	(32)
P_t^* Hour 2	.462	(.888)
P_t^* Hour 3	-.0679	(1.63)
P_t^* Hour 4	-.837	(2.28)
P_t^* Hour 5	-1.18	(3.04)
P_t^* Hour 6	-.353	(4.26)
P_t^* Hour 7	.138	(2.29)
P_t^* Hour 8	-.244	(.461)
P_t^* Hour 9	.118	(.436)
P_t^* Hour 10	.52	(.576)
P_t^* Hour 11	1.16	(.678)
P_t^* Hour 12	.157	(1.49)
P_t^* Hour 13	.68	(1.05)
P_t^* Hour 14	-8.55	(1.83)
P_t^* Hour 15	-.446	(.659)
P_t^* Hour 16	1.18	(.654)
P_t^* Hour 17	2.15	(.557)
P_t^* Hour 18	.58	(.567)
P_t^* Hour 19	1.45	(.505)
P_t^* Hour 20	.841	(.598)
P_t^* Hour 21	-.91	(1.57)
P_t^* Hour 22	-1.15	(2.45)
P_t^* Hour 23	.442	(2.56)
P_t^* Hour 24	2.34	(1.1)
P_t^{2*} Hour 2	-.00875	(.0103)
P_t^{2*} Hour 3	.00521	(.0291)
P_t^{2*} Hour 4	.0275	(.0562)
P_t^{2*} Hour 5	.0423	(.0942)
P_t^{2*} Hour 6	.0114	(.148)
P_t^{2*} Hour 7	.0135	(.0755)
P_t^{2*} Hour 8	.00555	(.00283)
P_t^{2*} Hour 9	.0131	(.00308)
P_t^{2*} Hour 10	-.0102	(.00446)
P_t^{2*} Hour 11	-.0143	(.00597)

P_t^{2*} Hour 12	.00344	(.0191)
P_t^{2*} Hour 13	-.00324	(.0117)
P_t^{2*} Hour 14	.126	(.0234)
P_t^{2*} Hour 15	.0066	(.00472)
P_t^{2*} Hour 16	-.00884	(.00452)
P_t^{2*} Hour 17	-.0185	(.00311)
P_t^{2*} Hour 18	-.00794	(.00318)
P_t^{2*} Hour 19	-.0144	(.00246)
P_t^{2*} Hour 20	-.00871	(.00361)
P_t^{2*} Hour 21	.0153	(.0197)
P_t^{2*} Hour 22	.0187	(.0351)
P_t^{2*} Hour 23	.00393	(.038)
P_t^{2*} Hour 24	-.0277	(.011)
P_t^{3*} Hour 2	.0000143	(.0000429)
P_t^{3*} Hour 3	-.0000906	(.000168)
P_t^{3*} Hour 4	-.00027	(.000435)
P_t^{3*} Hour 5	-.000459	(.000912)
P_t^{3*} Hour 6	-.0000949	(.00161)
P_t^{3*} Hour 7	-.000244	(.000777)
P_t^{3*} Hour 8	.00004	(.0000559)
P_t^{3*} Hour 9	-.000126	(.0000294)
P_t^{3*} Hour 10	.0000241	(.0000214)
P_t^{3*} Hour 11	.0000302	(.0000241)
P_t^{3*} Hour 12	-.0000537	(.0000777)
P_t^{3*} Hour 13	-.0000206	(.000044)
P_t^{3*} Hour 14	-.000546	(.0000957)
P_t^{3*} Hour 15	-.0000211	(.0000215)
P_t^{3*} Hour 16	.0000189	(.0000211)
P_t^{3*} Hour 17	.000041	(.0000192)
P_t^{3*} Hour 18	.0000226	(.0000192)
P_t^{3*} Hour 19	.0000302	(.0000187)
P_t^{3*} Hour 20	.0000157	(.0000193)
P_t^{3*} Hour 21	-.0000871	(.0000785)
P_t^{3*} Hour 22	-.000104	(.000158)
P_t^{3*} Hour 23	-.0000878	(.000179)
P_t^{3*} Hour 24	.0000713	(.0000392)
Constant	12.3	(14.4)

Standard errors in parentheses

Table 16: Perfect Foresight

Unit	Perfect Foresight	AR(1)	Percent Difference
Calenergy	\$886,020 (482,990)	\$246,210 (313,630)	260%
Graham 1	\$80,361 (6,806)	\$59,727 (3,871)	35%
Graham 2	\$43,870 (4,070)	\$33,297 (2,337)	32%
Morgan Creek A	\$38,959 (4,496)	\$36,490 (4,275)	7%
Morgan Creek B	\$37,247 (4,170)	\$34,325 (3,802)	9%
Morgan Creek C	\$35,374 (3,926)	\$34,325 (3,875)	3%
Morgan Creek D	\$36,753 (4,119)	\$33,460 (3,709)	10%
Morgan Creek E	\$36,445 (3,921)	\$35,056 (3,890)	4%
Morgan Creek F	\$36,959 (4,076)	\$36,767 (4,335)	1%
Morgan Creek 5	\$98,176 (14,287)	\$61,463 (5,860)	60%
Morgan Creek 6	— —	— —	—
Odessa	\$408,060 (137,080)	\$461,330 (211,420)	-12%
Oklaunion	— —	— —	—
Permian Basin A	\$27,159 (2,554)	\$26,456 (2,221)	3%
Permian Basin B	\$30,762 (3,002)	\$28,601 (2,425)	8%
Permian Basin C	\$31,931 (3,300)	\$30,637 (2,839)	4%
Permian Basin D	\$39,462 (4,471)	\$35,802 (3,570)	10%
Permian Basin E	\$55,870 (7,752)	\$43,755 (4,764)	28%
Permian Basin 5	\$181,810 (41,021)	\$183,210 (45,756)	-1%
Permian Basin 6	\$203,970 (37,694)	\$119,520 (16,079)	71%
Sweetwater	\$55,030 (4,705)	\$56,857 (5,153)	-3%
Witchita	\$53,492 (6,754)	\$58,006 (8,118)	-8%

C ERCOT Market Mechanisms

To ensure that there is sufficient supply, ERCOT requires generators and electricity retailers to submit scheduled energy transactions a day ahead. These schedules are submitted through a Qualified Scheduling Entity (QSE) which generally submits schedules for a portfolio of generators and power purchasers. These schedules outline which generators are planning to produce power and how that power will be transmitted to end users for each hour of the day. ERCOT allows QSEs to submit day-ahead schedules which leave them in long or short positions entering the production period³⁶. QSEs are also required to submit Balancing Market bidding functions for each hour of the day. The bidding functions show the willingness of generation portfolio to deviate from its scheduled output as a function of the price in the Balancing Market. The QSE must submit its willingness to both increase and decrease the portfolio output in response to price.

In real-time, ERCOT uses the Balancing Market to ensure adequate supply and to equate the marginal costs of production across generators. Every fifteen minutes ERCOT intersects the hourly bidding functions to arrive at a Market Clearing Price for Energy (MPCE) in each zone via a multi-unit uniform price auction³⁷. If there is no congestion between zones, then the prices are the same in each zone and the entire grid acts a single market. If congestion would occur between zones with a single MCPE, then ERCOT intersects the bidding functions separately by zone to achieve market clearing prices for each zone which do not exceed the transmission capability between zones. For example, if more power is needed in the South zone, but the transmission lines are at capacity, ERCOT will raise the prices in the South zone while lowering or keeping constant the prices in the other zones. In any case, generators respond to MCPE based on their bidding functions. The Balancing Market also helps to ensure that the lowest cost producers are generating electricity. At a low MCPE, high marginal cost firms have incentives to reduce or shut down production and satisfy their contractual obligations through energy procured from the Balancing Market. In a static, price-taking setting, the Balancing Market would ensure that only the lowest cost generators were producing energy each period. With the introduction of dynamics in the generating process, this no longer holds.

In addition to interzonal congestion discussed in the body of the paper, congestion can also arise within zones. This type of congestion cannot be resolved with market prices since there is only one price for each zone. To deal with local congestion, ERCOT deploys generators out of bid order. That is, ERCOT deploys specific generators which are not willing to increase production at current prices by offering them prices higher than the prevailing market price. The costs of deploying these resources to alleviate local congestion is covered by an output tax levied on all generators in the zone. This amounts to a uniform increase in marginal costs across all generators. Thus, transmission congestion

³⁶ERCOT also requires firms to have sufficient levels of ancillary power services

³⁷See Hortacsu & Puller (2008) for a detailed explanation of the auction process.

is either explicitly accounted for in the market price, if it occurs between zones, or it arrives as a uniform output tax on all generators in a zone.

D Equilibrium Proof

In this section, I show that an equilibrium exists for the model and that the equilibrium is unique under certain conditions. An equilibrium is defined by a finite vector of prices and demand functions such that in each period t the aggregate supply of all generators in the market meets demand. In addition, I require that firms' expectations for prices are consistent with the equilibrium price vector. Recall that the evolution of price is described by a Markov process which is conditional on the current price P and hour of day H . Let $F(P, H)$ denote a model of beliefs that describes the distribution of future prices for each relevant state (P, H) ³⁸. To show existence, I first show that there is a unique price equilibrium, P^* , for a fixed $F(P, H)$.

With fixed beliefs $F(P, H)$, there is a unique equilibrium price in a given period t , if the aggregate supply function, $S_{it}(P_t)$, is continuous and increasing in P_t and if the demand function is also continuous and strictly decreasing. Given that demand does not exceed total generation capacity, then by the intermediate value theorem there is a single intersection of supply demand. Recall that the aggregate supply function has the following form.

$$S_t(P_t, H_t, \mathbf{a}_{t-1}, \boldsymbol{\epsilon}_t) = \sum_{i=1}^N s_{it}(P_t, H_t, a_{it-1}, \epsilon_{it}) \quad (15)$$

Since the aggregate supply function is the sum of the firm-level supply functions, it is sufficient, though not necessary, to show that each firm's supply function is increasing in P_t for every state.

Referring back to the model in the paper, the generator-specific expected supply function is:

$$s_{it}(P_t, H_t, a_{it-1}, \epsilon_{it}) = a_{it}(P_t, H_t, a_{it-1}, \epsilon_{it})q_{it}(P_t) \quad \forall t \in \{1, 2, \dots, T\} \quad (16)$$

If q_{it} and a_{it} are increasing in P_t , then the firm level supply function s_{it} will also be increasing in P_t . First, q_{it} is increasing simply by definition as shown in equation 2. Second, if expectations for future prices are weakly increasing in the current price then the optimal policy function, $a_{it}(P_t, H_t, a_{it-1}, \epsilon_{it})$ will be increasing in price.

The optimal policy function, $p_i(a_t|P_t, H_t, L_t; F(\cdot))$, gives the optimal operating choice for each state and, in an abuse of notation, is written here to show its implicit dependance on beliefs $F(\cdot)$.

The optimal policy is a function of the choice specific value functions as shown below.

³⁸Note that the lagged operating state, L , does not affect the distribution of prices by assumption.

$$a_{it}(P_t, H_t, a_{it-1}, \epsilon_{it}) = (V_1(P_t, H_t, a_{it-1}) + \epsilon_{1it} \geq V_0(P_t, H_t, a_{it-1}) + \epsilon_{0it}) \quad (17)$$

where

$$\begin{aligned} V_0 &= \beta EV(P_t, H_t, a_t = 0 | F(\cdot)) \\ V_1 &= \Pi(P_t, a_{it-1}) + \beta EV(P_t, H_t, a_t = 1 | F(\cdot)) \end{aligned} \quad (18)$$

Using these equations, we can show that the optimal policy, and thus supply, is weakly increasing in the difference in the choice specific value functions. Holding fixed the expected future value $EV(\cdot)$, a_{it} is increasing in P_t since current profits are strictly increasing in P_t and are only earned if the generator is operating. However, the $EV(\cdot)$, will be increasing in P_t only if higher prices this period imply weakly higher prices next period. Thus, a sufficient condition for the optimal policy to be increasing is that $F(P', H)$ weakly first order stochastically dominates $F(P, H)$ for every $P' \geq P$. This rules out beliefs where increasing the price this period leads to a lower price next period.

I have shown that for a conditional on a_{it-1} and ϵ_{0it} and with the appropriate $F(P, H)$, s_{it} will be increasing in the current price for each generator. Thus, starting in the first period with initial conditions a_0 we can find the unique market clearing price, P_1^* , that satisfies demand.

$$D_1(P_1^*) = S_t(P_t^*) = \sum_i s_{i1}(P_1^* | H_1, a_{i0}, \epsilon_{i1}) \quad (19)$$

Using a_1 implied by P_1^* , I then solve for the next period's market clearing price, P_2^* . This is continued for all time periods resulting in a unique price equilibrium.

The result above assumes that S_t is continuous and that D_t is strictly decreasing. However, due to minimum output constraints, S_t may not be continuous. In addition, D_t may not be strictly decreasing, such as in the inelastic demand case. I will address these in turn.


First, since S_t is a non-continuous function there may be cases where the demand and supply curve do not intersect. This is illustrated in figure 11. In this case, at the "equilibrium" price there is a discontinuous jump in production due to the fact that the marginal generator cannot produce below its minimum output constraint. At this price, supply would exceed demand. To equate supply and demand, I reduce production from inframarginal generators until supply meets demand starting with the highest cost generator. In doing so, I respect the minimum output constraints of each generators. If there is still an excess supply, with all generators operating at their minimum production levels, I assume that the residual excess supply can be costlessly discarded. In this way, there is still a unique equilibrium price with production allocated such that supply and demand match exactly³⁹.

³⁹Note that firms would not be profit maximizing by producing less than their optimal levels. To maintain profit maximization of these inframarginal firms, I assume that they are made whole by the grid operator. That is, there are payments are made to the generators such

Second, in the case of an inelastic demand curve, there may be more than one equilibrium price if the quantity demand intersects a vertical portion of the step function as shown in figure 12. In this case, I use the lowest market clearing price as the equilibrium price.

After accounting for these two cases, there exists a unique equilibrium price vector, \mathbf{P}^* , which clears the market in each period conditional on the set of beliefs $F(P, H)$.

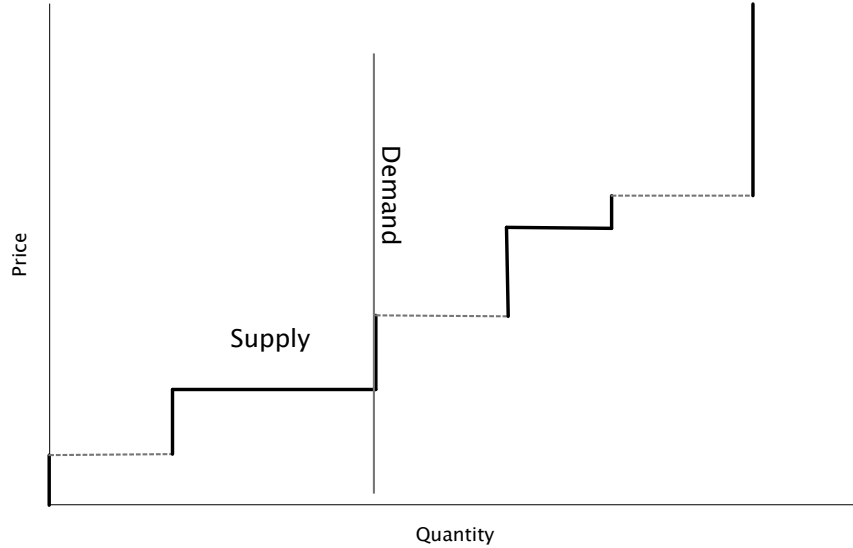
Figure 11: Non-continuous Function Case



An additional condition that I want to be satisfied is that firms' expectations over prices be consistent with the equilibrium price vector. Suppose that there is a function G which firms use to map a path of prices \mathbf{P} at hours \mathbf{H} into a model of beliefs $F(P, H)$, $G : \mathbf{P}, \mathbf{H} \rightarrow F(P, H)$. I say that an equilibrium price vector, $\mathbf{P}^*(\hat{F}(P, H))$, is consistent with beliefs $\hat{F}(P, H)$ if $G(\mathbf{P}^*, \mathbf{H}) = \hat{F}(P, H)$. That is, the beliefs that are generated by the equilibrium price vector are the same beliefs that firms were using in their optimization. In this case there is no ex-post regret on the side of the firms. Had they observed this path of prices in the past, they would not have changed their policies. This amounts to finding a fixed point between beliefs and the market clearing price vector. Due to the discontinuous and non-differentiable nature of the supply function, it is difficult to prove existence in this setting. In addition, I have not shown that such an equilibrium, if it exists, would be unique. However, subsequent work by Joseph Cullen & Stan Reynolds (2014), the authors show that when the number of firms in the market is large, such an equilibrium does exist and is in fact unique.

that their profits are the same as if they had been producing at full capacity. Alternatively, I could assume that all excess production is discarded. This would result in the same profits as above but would have needlessly higher emissions.

Figure 12: Inelastic Demand Case



E Calibrated Startup Costs

In this section, the sensitivity of the results to changes in the estimated startup costs is investigated. As noted in the paper, the estimated startup costs are much higher than startup costs that have been estimated in other work and are on the high end of the range of possible engineering costs. Fortunately, it is not necessary to estimate the startup costs in order to solve for the counterfactual equilibria using this framework. Any calibrated values for startup costs can be used in the counterfactual simulations. However, finding direct information on startup costs is challenging as it is considered propriety business information in most electricity markets.

Absent firm reported information on startup costs, I use engineering estimates of startup cost to calibrate the model. The National Renewable Energy Laboratory recently commissioned an engineering analysis of startup costs for the purpose of investigating the impact of renewables on market operations (Kumar et al. 2012). In the report, the authors estimate a lower bound on startup costs per MW of capacity for generators of various technology types. Applying these estimates to the generators in the counterfactual we are left with the calibrated startup costs shown in table 18.

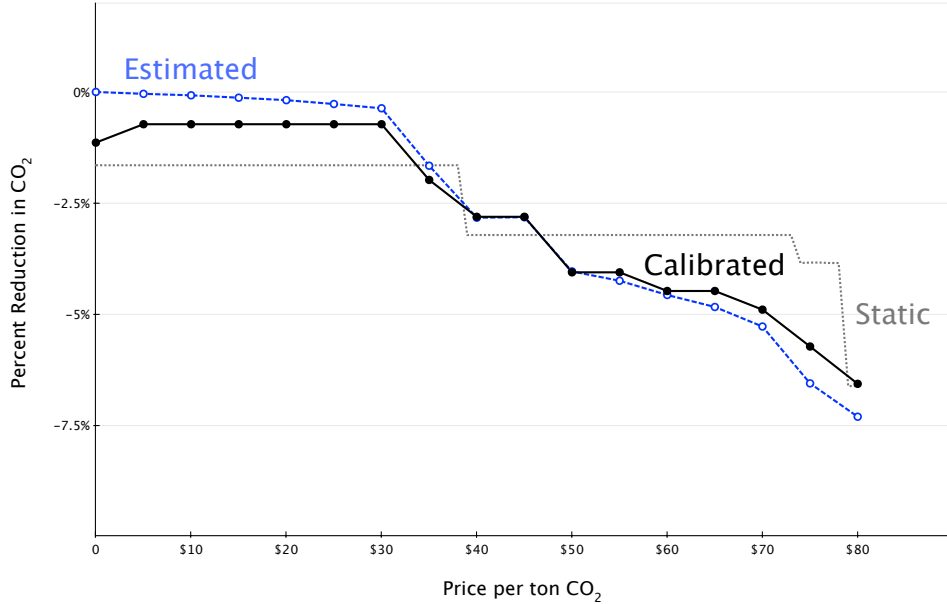
The calibrated startup costs are anywhere from two to ten times smaller than the estimated startup costs. The variance of the structural error for each generator is calibrated such that σ is 10% of the calibrated startup cost.

Despite the large differences between the calibrated and estimated startup costs, the path of counterfactual emissions reductions follow a very similar pattern. Figure 13 plots the emission response to a carbon price for the dynamic,

Table 18: Calibrated Startup Costs

Unit	Calibrated		Estimated	
	Startup	σ	Startup	σ
Calenergy	\$19,080	\$1,908	\$246,210	\$24,965
Graham 1	\$29,770	\$2,977	\$59,727	\$10,419
Graham 2	\$49,010	\$4,901	\$33,297	\$6,802
Morgan Creek A	\$3,320	\$332	\$36,490	\$6,092
Morgan Creek B	\$3,400	\$340	\$34,525	\$5,733
Morgan Creek C	\$3,320	\$332	\$34,325	\$5,599
Morgan Creek D	\$3,400	\$340	\$33,460	\$5,632
Morgan Creek E	\$3,320	\$332	\$35,056	\$5,782
Morgan Creek F	\$3,360	\$336	\$36,767	\$6,064
Morgan Creek 5	\$16,510	\$1,651	\$61,463	\$8,585
Morgan Creek 6	—	—	—	—
Odessa	\$86,40	\$8,640	\$461,330	\$67,006
Oklaunion	\$94,500	\$9,450	\$461,330	\$67,006
Permian Basin A	\$2,600	\$260	\$26,456	\$4,533
Permian Basin B	\$2,600	\$260	\$28,601	\$4,925
Permian Basin C	\$2,600	\$260	\$30,637	\$5,146
Permian Basin D	\$2,600	\$260	\$35,802	\$5,709
Permian Basin E	\$2,600	\$260	\$43,755	\$6,409
Permian Basin 5	\$15,080	\$1,508	\$183,210	\$31,178
Permian Basin 6	\$63,960	\$6,396	\$119,520	\$21,847
Sweetwater	\$20,340	\$2,034	\$56,857	\$10,996
Witchita	\$6,705	\$670	\$58,006	\$10,042

Figure 13: Calibrated Parameters: CO_2 Response



calibrated model against the estimated and static models. Emission predictions from the calibrated model lie in between the static model and the dynamic model with estimated parameters, but are more similar to the dynamic model. Baseline emissions with the calibrated parameters are about 1% lower than the model with the estimated parameters. Also, the emissions reduction at high carbon prices is not quite as large with calibrated parameters. Overall though, the calibrated counterfactual is quite similar to the estimated counterfactual. This might be a bit surprising given that the estimated startup costs are much larger than the calibrated ones. The similarity is partly due to the fact that the magnitude of the structural errors has also decreased in the calibrated model. Higher variance in the structural errors means that firms still have some probability of starting up despite higher startup costs.

Average prices with calibrated startup costs are also very similar to the counterfactuals with estimated startup costs. In each case the calibrated model predicts average prices that are about \$1 lower than before. The variance of prices, however, has decreased significantly; the variance had decreased by half in the most extreme case. However, it is still 50% higher than the variance in the static model.

Lower calibrated startup costs have a positive impact on profitability. Coal profits experience very small changes; they lose one percentage point fewer profits in each carbon price scenario. On the flip side, gas plants enjoy higher profits with lower startup costs. This is especially true for steam gas and simple cy-

cle generators who see significant increases in profits when compared with the counterfactuals with estimated parameters. Since these are the generators that do the most cycling, they are the generators that see the greatest benefit from lower startup costs.

The central theme of the results is not sensitive to calibrating the model with lower startup costs. Counterfactual emissions follow a path not unlike that was found with the estimated parameters. Also, the calibrated model still has significantly higher price variance than the static formulation. The most salient differences are found in the profitability changes of high-cost gas generators. These marginal generators had small increases or small decreases in profits with high startup costs, while under lower startup costs they benefit considerably from pricing carbon.

Table 19: Calibrated Counterfactual: Profits and Prices

	Carbon Price			
	\$0	\$20	\$40	\$80
Δ Profit:				
Coal	–	-24%	-47%	-82%
Combined Cycle Gas	–	17%	35%	99%
Steam Gas	–	23%	48%	95%
Gas Turbine	–	37%	43%	14%
Dynamic Prices:				
Average \$/MWh	\$62	\$71	\$82	\$111
Percent Change	–	15%	35%	78%
Variance	\$198	\$203	\$280	\$182
Within-Day	\$141	\$150	\$207	\$134
Across-Day	\$56	\$53	\$72	\$47