Strategic Uncertainty, Competition and Asymmetric Pass-Through: Evidence from Electricity Markets

Harim Kim*

August 11, 2026

Abstract

Asymmetry is a pervasive feature of price transmission and economic fluctuations, yet identifying its underlying source remains challenging. I study strategic sources of asymmetric price adjustment to cost changes, arising from asymmetries in firms' strategies and beliefs, in the context of electricity markets where alternative sources of asymmetry are ruled out. Using a research design suited to test asymmetry, I find that prices respond incompletely to cost decreases than to increases, with firms' bidding strategies exhibiting corresponding asymmetry. Structural estimation shows that firms' beliefs about competitors' strategies exhibit asymmetry, with greater dispersion following cost decreases, contributing to incomplete downward price adjustment.

1 Introduction

Asymmetric price adjustment to cost changes, a specific form of downward price rigidity, is pervasive across various markets and industries, making it an important feature of price dynamics.[1] Given that many aggregate economic variables also fluctuate asymmetrically, asymmetry in micro-level pricing patterns may play a key role in shaping the transmission of shocks across the economy.[2]

Despite being widely documented, this asymmetry is puzzling because standard price theory treats cost increases and decreases symmetrically, implying equal price responses in both directions. As a result, research has focused on identifying conditions under which such asymmetry naturally arises. Prior research has explored a variety of potential explanations, including consumer-side factors – such as consumer inattention or search costs – as well as factors beyond direct production cost changes, such as inventory management or price adjustment costs.[3]

Interestingly, this paper documents asymmetric price adjustment to cost changes even in a setting where many of the previously identified sources of asymmetry are absent. This raises the question:

*Department of Economics, University of Connecticut, Email: harim.kim@uconn.edu. I would like to thank Alex Imas, Joel Flynn, Steve Cicala, Steven Puller, Gordon Leslie, Karthik Sastry, Guoxuan Ma, Eric Lewis, Nicolas Schutz, Ryan Kellogg, Ying Fan, Larry Samuelson and participants at IIOC 2025, RCEA 2026 for helpful comments.

[1]The asymmetry refers to the phenomenon where prices respond quickly and often completely to cost increases but respond slowly and incompletely to cost decreases. Peltzman (2000) provides empirical evidence of the presence of asymmetric price responses across a wide range of product markets. Beyond product markets, asymmetric pass-through has also been documented in value-added tax transmission (Benzarti et al., 2020), bank deposit rates (Neumark and Sharpe, 1992), and municipal bonds (Green et al., 2010).

[2]For a detailed discussion of asymmetric fluctuations in aggregate economic variables, see McKay and Reis (2008).

[3]For discussions of consumer-side explanations, see Yang and Ye (2008), Tappata (2009), and Lewis (2011). For other explanations involving firm-side frictions, see Borenstein et al. (1997) on inventory effects and Ball and Mankiw (1994) on menu costs.

what fundamentally differs between upward and downward cost movements that could generate such asymmetry? To address this question, I explore a new supply-side mechanism – rooted in inherent asymmetries in firms' strategies and beliefs depending on the direction of cost change – through which asymmetric price adjustment can naturally arise in a strategic setting.

In doing so, I exploit a unique empirical setting in the wholesale electricity market that enables testing for asymmetry with minimal assumptions and structural constraints. The market I study has a simple structure in which previously proposed sources of asymmetry – especially those related to consumer behavior – cannot explain the observed asymmetric price adjustment patterns. This finding suggests that asymmetry can emerge from strategic interactions among suppliers, independent of consumer-side factors.

To uncover the strategic basis of this asymmetry, I draw on rich firm-level bidding data that reveal the strategic decision-making of electricity-generating firms. I find that underlying bidding strategies differ systematically between samples with similar fundamentals (e.g., cost and demand conditions) but opposite cost-change directions – a pattern that may stem from asymmetries in firms' beliefs about their competitors' strategies. To investigate this further, I structurally estimate firms' subjective beliefs – modeled as ex ante expected bids – that rationalize their observed best-response bids. While estimating beliefs from observed decisions is often challenging, I develop a structural model that infers the beliefs revealed in firms' bidding decisions. I find that following a cost decrease, these beliefs become more dispersed, covering a wider range of strategies, including bid adjustments that do not fully reflect the cost reduction. This pattern is consistent with the underlying insight that when costs fall, a wider set of strategies remains rationalizable than when costs increase. In the downward case, firms' pricing decisions are driven mainly by strategic considerations, whereas in the upward case, they need to raise prices at least partially to avoid losses, independent of competitors' strategies. Consequently, greater strategic uncertainty associated with downward cost movements widens the range of strategies consistent with profit-maximizing behavior, allowing incomplete price adjustments to emerge as a market outcome.

Strategic sources of asymmetry are often difficult to isolate from consumer-side explanations, but the empirical setting in this study – where previously proposed sources can be excluded and symmetric cost changes generate samples suitable for testing asymmetry – allows for the identification of a supply-side source of asymmetry. The use of high-frequency, firm-level microdata further enables me to investigate how uncertainty over strategies embedded in firms' decision-making relates to asymmetric strategic behavior and asymmetric price responses.

Throughout the paper, I leverage a unique cost variation in the electricity market driven by natural gas prices, which provides a proper empirical setting to test for asymmetry. Natural gas prices exhibit upside volatility, creating an environment in which the marginal generation costs of gas-fired power plants rise sharply and then return to pre-increase levels within a short period. This variation creates an ideal setting for examining asymmetric pass-through, as it ensures symmetry in the cost range over which pass-through rates are averaged, with only the direction of cost change differing (Benzarti et al., 2020). Instead of estimating pass-through using the cost variation, I take a slightly different approach by comparing prices between samples before and after a cost increase event – referred to as the pre-

increase and post-increase samples – to test whether prices return to their pre-increase levels once marginal costs decrease back to their original level. Since these two samples are not fundamentally different, with market conditions such as marginal costs and demand being similar, prices should return to their initial levels if adjustment to cost change is symmetric. However, I find that prices in the post-increase sample are about 7% higher than in the pre-increase sample, suggesting that output prices do not fully return to their initial levels immediately, indicating asymmetric pass-through.

In addition to providing a clean empirical setting, the wholesale electricity market offers the major advantage of not being subject to the commonly cited sources of asymmetric pass-through. First, because there are no consumers in the conventional sense in this market, consumer-side explanations for asymmetric price responses can be disregarded. Electricity demand in the wholesale market comes from electric utilities and is perfectly inelastic to prices. Thus, there are no end consumers who would learn about cost changes on the supply side and adjust their demand accordingly, ruling out consumer search-cost and inattention channels. Second, the market clears every day and hour, requiring supply-side firms to compete and make swift decisions in response to exogenous market variables such as fuel prices and demand, which vary stochastically over time. This feature ensures that firms' strategic decision-making remains effectively static, leaving them with little incentive to delay adjustments of prices or production. Moreover, because prices are adjusted every hour and all firms are required to participate in these markets, explanations based on differences in price adjustment cost (i.e., menu costs) or adjustment frequency (Gopinath and Itskhoki, 2010) can be ruled out. Lastly, supply-side firms procure natural gas based on daily needs, since it cannot generally be stored, and the spot natural gas prices they observe convey fairly complete information about competitors' costs. This eliminates information asymmetry between firms, particularly regarding costs, as well as inventory effects that could delay price responses (Borenstein et al., 1997).

Therefore, the asymmetric price responses documented in this paper cannot be explained by widely known sources of asymmetry. To investigate the underlying mechanisms, I leverage high-frequency bidding data at the firm and generating-unit levels, which are available because the wholesale electricity market is cleared via a multi-unit uniform auction. Comparing the pre- and post-increase samples, I find that price bids are systematically higher in the post-increase sample and become more dispersed. Notably, the stylized *peak-hour bidding* pattern – where strategic firms typically bid higher in peak hours than in off-peak hours, a pattern generally observed in electricity markets and in the pre-increase sample – largely disappears in the post-increase sample. Because peak hours are typically where firms' bidding decisions are most sensitive to strategic considerations, this breakdown in the hourly bidding pattern is difficult to reconcile with a change in market fundamentals alone. Moreover, these differences in bidding patterns across the pre- and post-increase samples vary across firms with different strategic positions, potentially reflecting the role of firms' strategic ability in shaping this difference.

Given that observable market fundamentals are similar across samples except for the direction of cost change, the adoption of different bidding strategies in the post-increase sample suggests a change in firms' beliefs about competitors' strategies. To explore this further, I structurally estimate each firm's beliefs – modeled as ex-ante expected bid distributions of competitors – that rationalize the observed bidding behavior. In doing so, I only impose the minimal rationality condition that observed

bids are best responses to estimated beliefs, rather than assuming a particular equilibrium concept. This is feasible due to not having to estimate structural primitives such as marginal costs jointly with beliefs, as the marginal cost of electricity generation can be accurately measured from data. This particularly attractive feature of electricity markets enables the estimation of beliefs without committing to a specific equilibrium concept, and is compatible with both equilibrium and non-equilibrium interpretations of firm behavior (Hortaçsu et al., 2019). Therefore, the beliefs estimated here are not necessarily restricted to equilibrium beliefs but can be interpreted more broadly as subjective beliefs defined over a set of rationalizable strategies of competitors and are consistent with any equilibrium or non-equilibrium model that could have generated the observed market outcomes. In other words, if firms' behavior indeed reflects a particular equilibrium or behavioral framework, the estimated parameters would capture beliefs consistent with its underlying mechanisms.

The estimates allow me to assess whether there are systematic differences in beliefs – captured by the mean and variance of ex-ante expected bid distributions – between the pre- and post-increase samples. While the mean of belief distributions is slightly higher in the post-increase sample, the more notable change is an increase in dispersion: the standard deviations are higher than in the pre-increase sample for the majority of firms. This change is more pronounced during peak hours, though a considerable change also occurs in off-peak hours, consistent with the more dispersed bidding patterns observed in the bid data.

The key insight behind this finding is that the set of rationalizable strategies differs depending on the direction of cost change. For example, if firms do not increase prices by the full amount of marginal cost increase, they incur per-unit profit loss regardless of competitors' actions; thus, price responses smaller than the cost increase can be rationally excluded from the strategy set. On the other hand, when costs decrease, a firm's decision to decrease prices is driven solely by strategic considerations – only when it believes competitors to lower their prices – because the per-unit profit loss depends on competitors' strategies. Consequently, strategies in which prices are not fully adjusted downward by the size of the cost decrease are not excluded from a firm's rationalizable strategy set and can be a best response to some beliefs, resulting in a much wider strategy set under downward cost movements.[4] In the empirical analysis, I find that when firms face competitors whose strategies they become uncertain about, they tend to raise their bids on average, although the exact pattern and direction of bid adjustment vary across both firm and competitor types. These findings are consistent with strategic uncertainty and differences in firms' strategic capabilities shaping the direction and magnitude of their strategic price adjustments.

The findings of this paper highlight the strategic sources of asymmetric price responses by demonstrating – under minimal assumptions about equilibrium behavior – that beliefs and rationalizable strategy sets differ with the direction of cost change. Empirically detecting and estimating this asymmetry in beliefs is informative, as it demonstrates that distinct outcomes can emerge across samples under plausible rational behavior, a step toward understanding the microeconomic foundations of asymmetric pass-through.

[4]While this is a simple illustrative example assuming zero initial markups prior to cost adjustment, the asymmetry in rationalizable strategy sets also holds when initial markups are positive. I discuss this in detail in Section 5.

Literature Review This paper contributes to several strands of literature. First, it identfies an additional source of asymmetric price responses beyond those considered in the existing literature. A large body of work has documented asymmetric pass-through and investigated its underlying causes, both theoretically and empirically. The most widely discussed explanations focus on consumer-side mechanisms – such as inattention or higher search costs following a negative cost shocks – which suppliers can exploit to delay price reductions (e.g., Yang and Ye, 2008; Tappata, 2009; Lewis, 2011; Wu et al., 2024). On the production-side, researchers have investigated inventory effects (Borenstein et al., 1997) and market power (e.g., Neumark and Sharpe, 1992; Borenstein and Shepard, 2002) as potential sources of asymmetry. Although not explicitly examining the underlying source, Benzarti et al. (2020) study asymmetric pass-through of symmetric VAT increases and decreases using an empirical design similar to that of my paper and provide evidence of persistent asymmetry. Despite the large body of literature, relatively few papers have explored the role of strategic interaction in explaining asymmetric responses, a gap to which my paper contributes by examining strategies and beliefs using detailed micro-level analysis.

While I attribute the observed asymmetry to differences in strategies and beliefs, the analysis does not attempt to identify which specific equilibrium model best characterizes the firms' behavior underlying the observed price responses to cost decreases. Nonetheless, the finding that rationalizable strategy sets and the supporting beliefs differ across samples remains informative, as it implies that distinct outcomes can emerge even without specifying a particular equilibrium. Moreover, because the analysis does not impose belief consistency, the findings are also compatible with non-equilibrium frameworks. In this sense, the paper broadly relates to the literature on firm behavior that departs from standard equilibrium assumptions. For instance, Hortaçsu et al. (2019) introduce bounded rationality through a cognitive hierarchy framework to explain anomalous bidding behavior in wholesale electricity markets, while Doraszelski et al. (2018) adopt a learning model to explain unusual bidding behavior in the early stages of market formation. Although my paper does not estimate a specific non-equilibrium model, it shares the motivation of rationalizing observed behavior that departs from Nash predictions and extends these insights to explain asymmetric pass-through more generally.

More broadly, asymmetric price responses are not confined to a specific market or industry but are closely related to the well-known phenomenon of downward price rigidity – an important feature of price dynamics relevant across many fields of economics. Although not explicitly focusing on asymmetry, a large literature has examined incomplete pass-through in relation to price rigidity. Examples include studies on price adjustment frequency (e.g., Gopinath and Itskhoki, 2010; Bils and Klenow, 2004), price adjustment costs (e.g., Ball and Mankiw, 1994), and markup adjustments (e.g., Goldberg and Hellerstein, 2013; De Loecker et al., 2016; Amiti et al., 2019). If these mechanisms operate differently depending on the direction of cost changes, they may generate asymmetric pass-through. My empirical setting, however, is largely free from these sources and, at the same time, provides a more direct test of asymmetry.

Lastly the paper also contributes to the literature on wholesale electricity market competition and strategic interactions within empirical auction frameworks. Prior research has examined strategic bidding behavior in this market (e.g., Wolfram, 1998; Hortaçsu and Puller, 2008) and used empirical

auction models to study various aspects of competition (e.g., Wolak, 2007; Reguant, 2014). More closely related to this paper are studies that combine auction frameworks with analyses of pass-through in electricity markets (e.g., Fabra and Reguant, 2014; Kim, 2022).

2 Empirical Design and Setting

2.1 Challenges in testing the asymmetry in price responses

Price adjustment to cost change is typically measured with the pass-through rate. The cost pass-through rate ρ, defined as the response of prices to a *marginal* change in cost, depends on the state at which the marginal change is evaluated. This implies that ρ is heterogeneous and is not generally a constant parameter. For instance, the marginal pass-through rate evaluated at cost level c_0 may differ from that evaluated at c_1 (where $c_0 \neq c_1$). In other words, $\rho(c_0) \neq \rho(c_1)$.

Since the marginal pass-through rate is a point on the $\rho(\cdot)$ function, the rate itself does not vary with the direction of the cost change from the point evaluated, meaning that the marginal rate ρ cannot be asymmetric in principle. Therefore, as demonstrated in Benzarti et al. (2020), testing for price response asymmetry involves using the average pass-through, the average of ρ over a range of cost values. In my context, for a cost change from c_0 to c_1, the average pass-through is represented as:

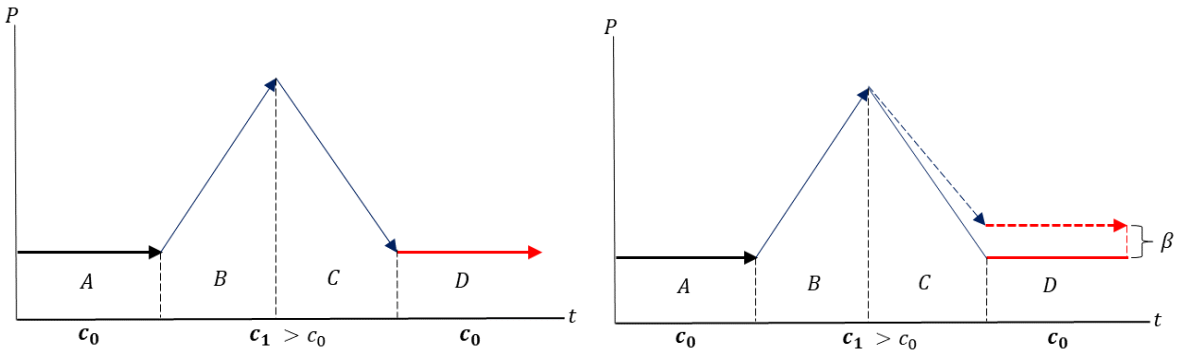
$$\bar{\rho} = \frac{1}{c_1 - c_0} \int_{c_0}^{c_1} \rho(c)\, dc$$

This expression clearly shows that the average pass-through rate varies with the range of costs as well as the start and end points over which the non-constant marginal pass-through ρ is averaged. Different sizes of $\Delta c\,(= c_1 - c_0)$ result in different levels of $\bar{\rho}$, and even with the same Δc, if the range includes different values of c_t, the corresponding ρ values will differ.

Therefore, to properly test for asymmetry, the range of cost changes in both directions should be identical, and the ranges themselves should be symmetric – with cost-increase and cost-decrease ranges given by $[c_0, c_1]$ and $[c_1, c_0]$, respectively. A simple comparison of two average pass-through rates estimated from subsamples that differ in the direction of cost change, but have different composition of cost values or lengths, may not capture the true asymmetry, as these average rates would differ by construction.[5]

To address this problem, one needs an ideal setting in which a cost change in one direction is followed by an equivalent change in cost in the opposite direction. Figure 1 describes such cost change events ideal for testing asymmetry. The figure shows the changes in the output prices when the cost, initially at c_0, increases to c_1 (block B) and then decreases back to c_0 (block C). Since both the range and the end points of the cost changes are identical in blocks B and C, asymmetry can be tested by comparing the average pass-through rates from each block, i.e., testing whether $\bar{\rho}_B = \bar{\rho}_C$. If the average pass-through rates are the same, as shown in Panel (a), the price response is symmetric. Panel (b) shows the case where the average rates differ, with the price change in block C (dashed line) smaller than in block B ($\bar{\rho}_B > \bar{\rho}_C$), indicating asymmetric pass-through. A similar empirical design was used

[5]Despite this limitation and related concerns, I estimated average pass-through rates for samples differing in the direction of cost changes, as shown in Table A1 in the Appendix.



(a) Symmetric Case (b) Asymmetric Case

Figure 1: Price and Cost Relationship: Symmetric vs. Asymmetric Price Adjustment Cases

in Benzarti et al. (2020), which studies asymmetric tax incidence by estimating average pass-through rates for two events in which the tax rate decreased and then returned to its previous level.

However, finding such an empirical setting is generally challenging. Even when one exists, another difficulty is ensuring that other factors affecting equilibrium prices remain constant between these cost-change events (blocks B and C), so that the only difference is the direction of the cost change. For example, if the events are years apart or occur in different markets or regions, there is no assurance that economic conditions are the same between these events, and thus one cannot conclude that differences in the average pass-through rates between blocks B and C provide evidence of asymmetric price response. Moreover, observing multiple such pair of events with equivalent cost ranges in real data is rare, often leaving the analysis with insufficient variation for statistical inference.[6]

2.2 Empirical Approach of this Paper

To overcome the challenges described in Section 2.1, I take a different approach to test asymmetry in price response, leveraging the unique empirical setting and data. Instead of using the cost-change events in blocks B and C of Figure 1, I focus on blocks A and D, denoted as the *Pre-increase* and *Post-increase* samples. Prices in block A are those under a cost of c_0, before the cost increase, while block D shows prices after the cost decreases back to c_0. Although the costs in both blocks are the same, the two samples differ in the direction of the preceding cost movements.

I test whether prices return to the same level as before the cost increase once the cost decreases back to c_0. That is, if other conditions are the same between blocks A and D, symmetric price response to a cost change implies that prices should be identical in both blocks, as in Panel (a) of Figure 1, because the costs are the same. If the price response to a cost change is asymmetric, prices in block D will be higher than that in block A, indicating that prices do not fully return to their previous level despite the cost returning to c_0.

The extent of asymmetry is captured by the difference in price levels between the pre-increase and post-increase samples (blocks A and D), represented by the parameter β in Panel (b). This is

[6]Both long lags between events and limited observations posed challenges in Benzarti et al. (2020), where the decrease and subsequent increase in tax rates were five years apart and only one such ideal event was observed.

analogous to testing whether the average pass-through rate in block B (price change when cost rises from c_0 to c_1) equals that in block C (price change when cost decreases from c_1 to c_0). Specifically, a positive β corresponds to downward price rigidity or the rockets-and-feathers effect.

This empirical approach, which compares outcomes across two distinct samples, can be extended beyond prices to variables such as markups or firm-level bids, enabling further analysis of underlying mechanisms. Moreover, focusing on the pre- and post-increase samples has the advantage of yielding more homogeneous data for the test, since both samples share the similar cost level and market fundamentals, thereby increasing observations and statistical power. By contrast, focusing on cost-change events (blocks B and C) would complicate the analysis, since the magnitude of cost change (i.e., Δc) may vary across events and the cost-change period itself introduces additional market changes that would need to be controlled for. One caveat is that this approach places greater weight on downward cost movements to test asymmetry, as it does not directly estimate the price response to an upward cost movement (block B).[7] This does not weaken the test, however, as testing asymmetry using blocks B and C is empirically equivalent, from a design perspective, to testing using blocks A and D.

2.3 Empirical Setting: Natural Gas Price Variation and Electricity Prices

Wholesale Electricity Market The wholesale electricity market is organized as a multi-unit uniform auction, cleared every hour of the day. Generators and demand-side companies submit supply and demand bids, respectively, which the market operator aggregates into supply and demand curves. The market is cleared by finding the price at the intersection of these curves – the price bid of the marginal unit sets the market clearing price. Each generator's price bid reflects its marginal cost of electricity generation and a possible strategic markup. Because fuel costs are the main component of marginal cost, changes in fuel prices are directly incorporated into marginal costs, making this market a useful setting for studying cost pass-through from fuel price fluctuations.[8]

Natural Gas Prices and Cost Changes For the analysis, I exploit variation in natural gas prices and study its impact on wholesale electricity prices. Since natural gas is an essential input in electricity generation, a change in gas prices implies a cost change, which is ultimately reflected in output electricity prices. The variation in gas prices is mainly driven by weather-related gas pipeline congestion due to increased residential heating demand, making it exogenous to electricity market conditions.[9]

The variation in natural gas prices creates a cost pattern similar to the ideal setting illustrated in

[7]Because the upward cost movement part is not the focus of the baseline specification, the test does not speak directly to whether pass-through of a cost *increase* is complete or incomplete. However, an estimate provided in Table A1 of the Appendix shows that pass-through is close to one in the cost-increasing case.

[8]Cost pass-through in the wholesale electricity markets has been studied extensively, including Fabra and Reguant (2014) on emissions costs and Kim (2022) on fuel price pass-through.

[9]The congestion is primarily due to higher residential heating demand during cold weather, rather than increased demand from the power generation sector. While cold weather may raise electricity use among residential customers, the increase is not large enough to cause congestion. For example, in Figure B6 in Online Appendix, where the electricity demand (load) and electricity prices are plotted together, demand levels for the winters 2012-2014 are not significantly higher than in other winters despite the electricity prices being substantially higher during these winters due to natural gas price increases. More discussion about the exogeneity of the natural gas price variation can be found in Kim (2022).

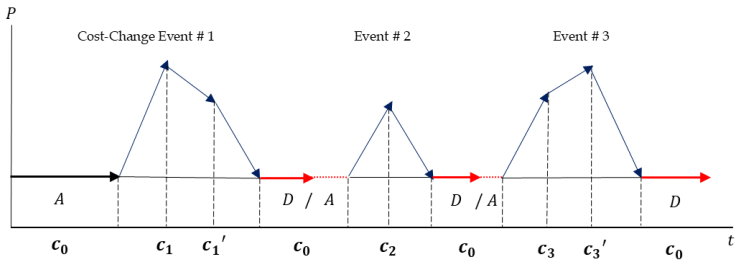


Figure 2: Example of Cost-Change Events in the Sample: Symmetric Case

Figure 1. A distinctive feature of this variation is its upward volatility: gas prices rise above the base level and then return to it. Marginal costs for gas-fired electricity generators follow the same pattern, creating a unique event in which a cost increase is followed by a decrease within a relatively short time frame.

Specifically, gas prices – and thus marginal costs for gas-fired generators – remain close to the base level with minimal change across days and hours, similar to block A in Figure 1, where marginal costs are at c_0. As gas prices rise above the base level, marginal costs increase to c_1. After this cost-increasing event, gas prices return to the base level and marginal costs return to c_0, as depicted in block D in Figure 1. Thus, marginal costs are at the same level (or in a similar range) in blocks A and D. I provide a more detailed discussion and visualization of gas prices in Section 3.2.1.

While this pattern (shown in Figure 1) is generally preserved in the data, a more realistic pattern is shown in Figure 2, where multiple cost-change events (e.g., about 20) occur over time. Observing many such cost-change events increases the size of the sample (blocks A and D), thereby enabling statistical inference on the asymmetry parameter β, estimated as the difference in prices between blocks A and D.[10] Within each cost-change event, multiple higher-cost days (with $c > c_0$) may occur before costs return to c_0, as illustrated in events 1 and 3 in Figure 2. A relatively short duration of these cost-change events (which lasts several days) further ensures that no fundamental difference in market conditions occurs between blocks A and D. Details on how samples A and D are defined are provided in Section 3.2.1.

2.4 Why Other Sources of Asymmetry Do Not Apply in This Market

The wholesale electricity market provides a particularly clean setting in which commonly discussed sources of asymmetry can be ruled out. A large literature explores the causes of asymmetric pass-through, with the most common explanation being consumer search costs on the demand side (e.g., Yang and Ye, 2008; Tappata, 2009; Lewis, 2011). The wholesale electricity market, however, does not have an active demand side that responds to prices – the demand side consists of electric utilities and retail electricity providers that must acquire electricity regardless of price, making demand almost perfectly inelastic.[11] As there are no end consumers in this market whose consumption decisions

[10] This contrasts with Benzarti et al. (2020), where only a single pair of such events (a one-time tax decrease and increase) is observed. For this reason, statistical inference was not possible in Benzarti et al. (2020).

[11] The electricity sector is divided into wholesale and retail markets, which are largely disconnected. Retail companies operate in both, but because they are obligated to meet residential electricity demand, their goal in this wholesale market

depend on information about supply-side cost changes, demand-side explanations for asymmetric pass-through can be excluded.

Differences in price adjustment frequency (Gopinath and Itskhoki, 2010) or adjustment costs (i.e., menu costs) also cannot explain the documented asymmetry in this market. In this market, price adjustment occurs every hour of the day, with all firms facing the same adjustment frequency and effectively no adjustment costs, as firms are required by market design to compete and receive profits based on cleared prices and quantities in each hourly market.[12]

Explanations involving intertemporal dependencies also do not apply in this market. Firms must compete and earn profits in daily markets, making strategic interaction effectively static – firms maximize same-day profits and have no incentive to take actions today to influence competitors' future strategies.[13]

Moreover, key market variables relevant to firms' pricing and decision-making vary stochastically and are exogenously determined. Natural gas prices, which form the basis of marginal costs, are set in a separate competitive gas market that operates independently of the electricity market. These prices fluctuate daily and are not perfectly predictable – while they may appear similar across days, this reflects typical competitive price levels rather than prices being fixed over time. Similarly, wholesale electricity demand is exogenously driven by retail consumption and varies stochastically over time. Retail consumers are unaware of wholesale market conditions, with consumption driven mainly by household activities, weather, and retail electricity prices disconnected from wholesale prices. Demand is therefore not perfectly predictable, even if intra-day patterns exhibit similar profiles.

Given the high-frequency nature of market clearing and the exogenous variation in key market variables, the strategic environment can be characterized as a repeated static game with a randomly evolving state. In each period, firms play a game given the current state, which evolves stochastically and independently of firms' actions. Firms base their decisions on current-period payoffs rather than expected discounted payoffs over time.[14] Thus, dynamic pricing incentives are largely absent in this market, which, if present, could slow price adjustment to cost changes (MacKay and Remer, 2022).

Because the strategic environment features stochastic state transitions rather than stable repeated interactions, explanations involving collusion may be ruled out. Collusion relies on intertemporal incentives, such as punishment strategies linking current actions to future outcomes, which are undermined by the stochastic nature of the state. Furthermore, firms cannot observe competitors' bids even after market clearing, as detailed bid-level information is released only after a substantial delay of four months – preventing timely detection of and response to deviations.[15] Market pricing and bidding are

is to secure supply regardless of price. For this reason, they do not specify a price when submitting demand bids.

[12]This follows from the market's high-frequency, auction-based design, where market clearing and pricing occur hourly and firms must submit bids and make swift production decisions based on current market conditions.

[13]There is a possibility of dynamic production decisions across hours due to fixed costs, but this generally does not create dynamic dependencies across days. A discussion of dynamic plant operation and its impact on the results is provided in Section 3.3.4.

[14]An exception applies to firms operating assets with dynamic production, such as those with large fixed costs. However, even for these firms, dynamic production occurs across hours rather than across days.

[15]To prevent market manipulation and tacit collusion, FERC Order 719-A requires grid operators to delay the publication of individual supply and demand bids for several months.

also regularly monitored by multiple agencies, further reducing the feasibility of sustained collusion.[16]

Finally, the uniqueness of natural gas as an input allows me to exclude several additional channels of asymmetry. In settings where inputs can be stored, inventory management creates a lag between input price changes and cost responses: an input price change today affects costs only after existing inventories are depleted. Such dynamics, combined with asymmetric inventory adjustment costs, have been identified as one explanation for asymmetric pass-through in gasoline prices (Borenstein et al., 1997). In contrast, natural gas cannot be viably stored, meaning procurement occurs on a just-in-time basis – this eliminates inventory effects and ensures that the spot gas price on the day of production is immediately reflected in generation costs of gas-fired plants. This also rules out cost-related information asymmetry: the spot gas price observed by firms conveys nearly complete information about competitors' costs, since it is common knowledge that dispersion in individual transaction-level spot gas prices is minimal in a liquid gas market (see Kim (2022) for details).[17]

3 Empirical Analysis

3.1 Data

The main dataset used in the analysis is from the New England wholesale electricity market (ISO-NE), providing high-frequency data on bids submitted by market participants, as well as corresponding price and production data. The dataset covers the years 2012 to 2014 and includes supply-side energy offer bids, demand bids, and system clearing prices (i.e., the energy component) from the day-ahead wholesale electricity market.[18] The analysis focuses on days from October to April, when natural gas price variation is most pronounced.

I obtain data on power plant ownership, fuel types, and capacities from the ISO-NE Seasonal Claimed Capability data and the EIA-860 form. Although power-generating firms operate plants using various fuels, natural gas was the dominant energy source during the study period.[19]

The natural gas price data used in the analysis are spot price indices (from *Natural Gas Intelligence*), a commonly used daily measure of natural gas spot prices.[20] I measure marginal costs using generator heat rates and daily fuel (natural gas) prices. I obtain generating unit-level heat rates (physical efficiencies) from a separate estimation described in Appendix B.1.3, and combine them with daily gas price indices to construct generator-level marginal costs.[21]

[16] A more detailed discussion is provided in Appendix B.4.2.

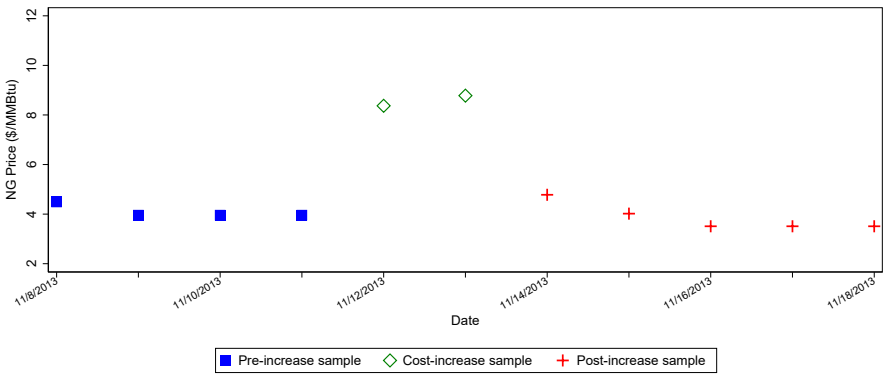
[17] Firms could purchase gas through long-term contracts, but this is not an attractive procurement channel for power generators due to inflexibility and the need to pre-commit to fixed amounts – as shown in Kim (2022), only about 20% of gas generators use long-term contracts. Moreover, long-term prices tend to be similar to spot price levels when the spot gas market is liquid (Levine et al., 2014). More importantly, firms cannot easily change their procurement channel within the short event window, so procurement methods are the same across pre- and post-increase samples, and endogenous changes in procurement do not threaten identification.

[18] As demand bids are price-insensitive, I use only aggregate demand – measured as the sum of the quantity components of demand bids. I subtract the hourly net interchange from aggregate demand to adjust for imports and exports. Transmission conditions are not incorporated into the analysis.

[19] The generation share of natural gas exceeded 50 % during the time period. Note that the generation share of renewables for this time period was very small, less than 5%.

[20] I used Algonquin citygate gas price index data. I also use index data for other fossil fuels, such as petroleum products and coal, although these are not the primary focus of the analysis.

[21] Because heat rate represents physical efficiency, it remains constant throughout the sample period. A single heat



Notes: The figure presents a subset of the dataset, selected to visualize how days are classified into each sample group. Daily natural gas prices are shown together with sample classifications: squares: pre-increase sample; crosses: post-increase sample; diamonds: cost-increase sample.

Figure 3: Visualization of Sample from Data

3.2 Samples

3.2.1 Pre-Increase and Post-Increase Samples

I divide the sample into three parts – Pre-Increase, Cost-Increase (i.e., cost-change events), and Post-Increase – based on both the *direction* of gas price changes and the gas price *level* relative to the baseline. The empirical analysis uses only the Pre-increase and Post-increase samples, with cost-increase events used only for dividing these samples. The sample division is exogenous, since natural gas prices vary for reasons independent of electricity market conditions.[22]

As shown in Figure A1, natural gas prices typically stay around a base level of \$4.5/MMBtu before rising to a level distinct from this floor and eventually returning to it. I classify a day as part of a cost-increase event if gas prices exceed the baseline threshold κ of \$5.5/MMBtu and are moving in a consistent direction – either rising or falling.[23] A more detailed discussion is presented in Appendix B.1.1.

Figure 3 shows a simplified visualization of the sample division using daily natural gas spot prices. The *Pre-Increase* sample (block A of Figure 1) includes consecutive days with base-level gas prices before any significant cost increase. The *Cost-Increase* sample includes days when gas prices rise above the base level.[24] The *Post-Increase* sample (block D of Figure 1) consists of days when gas prices return to the base level following a cost increase event. Because the data include multiple cost-increase events (20 in total), there are several post-increase blocks, which are pooled into the post-increase sample. I also consider an alternative categorization that reclassifies some post-increase

rate parameter is estimated for each gas-fired generating unit. The sample used for heat rate estimation does not overlap with the samples used in this paper. More details can be found in Appendix B.1.3.

[22] See Section 2.3 for details.

[23] I set κ at \$5.5/MMBtu because all observed values in the baseline period remain below this level, fluctuating within a narrow range around the baseline average. Formally, any day t satisfying $\kappa < P_{ng}(t-1) < P_{ng}(t)$ (e.g., the increasing segment) or $P_{ng}(t-1) > P_{ng}(t) > \kappa$ (e.g., the decreasing segment) is classified as part of the cost-increase event sample. This sample division approach is commonly applied in the asymmetric pass-through literature (e.g., Borenstein et al., 1997; Peltzman, 2000), where authors construct an indicator variable based on the direction of the cost change.

[24] The duration of these events varies – from a single day to several consecutive days – as do price levels within them, depending on the severity of pipeline congestion.

Sample	Stats	NG price ($/MMBtu)	MC ($/MWh)	Demand (MW)
Pre-Increase	Mean	4.30	37.5	13,230
	Med	4.27	37.7	13,317
	Min	3.90	24.1	6,197
	Max	5.18	70.0	23,049
	SD	0.27	6.9	3,430
	N	1,340	1,324	1,340
Post-Increase	Mean	4.55	39.3	13,356
	Med	4.64	40.1	13,175
	Min	3.51	22.1	6,092
	Max	5.16	73.8	23,081
	SD	0.37	7.2	3,181
	N	1,517	1,509	1,517

Notes: NG price is the daily spot gas price index, MC is the marginal cost of a non-dual marginal units, and demand is the hourly aggregate demand. The sample includes day-hour pairs where non-dual natural gas generating units set the market price.

Table 1: Summary Statistics: Pre-Increase and Post-Increase Samples

days immediately preceding a subsequent cost-increase event into the pre-increase sample, producing qualitatively similar and in some cases stronger results – see Section 3.3.3 and Appendix B.1.2 for details. Figure A2 in the Appendix shows visualizations of both the main and alternative sample divisions for the extended sample.

Having defined the samples, I verify that the pre- and post-increase samples are balanced in terms of market variables. Table 1 presents summary statistics for gas prices, marginal costs of gas-fired units, and demand separately for each sample. The two samples are well balanced, exhibiting similar distributions across key variables. Beyond summary statistics, the institutional features of wholesale electricity markets and the short duration of the event window ensure that characteristics relevant to firms' strategic decision-making remain unchanged between samples. For example, the number of market participants is fixed throughout the sample due to electricity market rules, and firms' fuel procurement channels cannot be changed within such a short time period.[25] As a result, any cross-sectional differences in firm characteristics are preserved between the pre- and post-increase samples and do not threaten identification. I also confirm this at the generator level by examining generator availability and offered quantity for all generators, finding no significant difference between the pre- and post-increase samples (see Table A2). Section 3.3.4 provides further discussion and robustness checks.

3.2.2 Peak and Off-Peak Hours

In electricity market, competition and pricing strategies vary across hours. During off-peak hours, low demand and sufficient idle capacity intensify competition, leaving firms with limited scope for strategic behavior. During peak hours, high demand and scarce capacity allow firms with greater strategic ability – whether from scale or a larger share of strategically usable capacity – to exercise

[25]While firms may procure natural gas through different channels – such as spot markets or long-term contracts – they cannot readily adjust procurement arrangements over short periods, as long-term contracts are set over longer time horizons. Moreover, spot prices and long-term contract prices tend to be similar because contract prices are often linked to gas price index levels, although discrepancies increase with volatility (Levine et al., 2014).

market power by submitting higher bid markups. I exploit these differences in strategic behavior between peak and off-peak hours throughout the analysis. In doing so, I index each market by (t, h) and categorize hours from 12 AM to 7 AM as off-peak and 8 AM to 11 PM as peak.[26]

3.3 Main Analysis

3.3.1 Price difference between pre- and post-increase samples

The first set of analyses examines whether there is a statistically significant difference in market prices of electricity between the *pre-increase* and *post-increase* samples. Because this market clears through a uniform-price auction, all firms receive the same market-clearing price – set by the bid of the *marginal* unit – in contrast to settings with firm-level pricing; Section 4 conducts a firm-level analysis of individual bids directly.[27] The main specification is shown below in equation (1):

$$P_{th} \;=\; \alpha \;+\; \beta\,\text{Post-Increase}_t \;+\; \gamma X_{th} \;+\; \eta + \;\varepsilon_{th} \tag{1}$$

P_{th} is the hourly wholesale electricity price that clears the entire system (i.e., the energy component price). The main variable of interest is the indicator Post-Increase_t, which distinguishes the pre- and post-increase samples. The control variables X_{th} include market fundamentals that affect prices, such as hourly aggregate demand (Demand_{th}) and the marginal cost of the marginal unit clearing the market (mc^*_{th}).[28] I exclude markets in which dual-fuel natural gas generators are the marginal unit, to avoid contamination from potential fuel-switching between gas and oil by these units.[29] Year and hour fixed effects are specified in η. Standard errors are computed using the Newey-West estimator to account for potential autocorrelation and heteroskedasticity in hourly market prices.

This empirical design, which aligns with the framework in Figure 1, is best understood as a test of the structural restriction implied by the standard price-setting model, rather than a conventional treatment-effect estimate: conditional on market fundamentals such as marginal cost and demand, a statistically significant and sizable β indicates a systematic price difference attributable to the direction of the cost change – the feature that distinguishes the samples by construction. That said, for this test to be valid, assignment to the post- versus pre-increase sample must be unrelated to other determinants of price – a requirement supported by the fact that the underlying gas-price variation is driven by factors outside the electricity market. I further examine whether other factors, such as generator availability and operational conditions, could differentially affect the two samples, and find these to be balanced across samples, with details provided in Section 3.3.4.

Table 2, column (1), presents the full-sample estimation result: prices in the post-increase sample are, on average, approximately \$3.17/MWh higher than in the pre-increase sample. Column (2),

[26] As a robustness check, I also use an alternative categorization that splits peak hours into 8 AM - 3 PM and 4 PM - 11 PM, treating the latter as more peak. The results are qualitatively similar.

[27] See Appendix B.2.4 for further discussion of how the uniform-price mechanism relates to firm-level pricing frameworks.

[28] Because the market price is set by the price bid of the last generating unit at the intersection of the system's supply and demand curves (i.e., marginal unit), I include only the marginal cost of this unit in the regression, as it is relevant for the market-clearing price.

[29] I find qualitatively similar results when including dual-fueled units, as presented in Table 4 of Section 3.3.4.

	(1)	(2)	(3)	(4)	(5)
		Full Sample		Off-peak	Peak
	P_{th}	$\log(P_{th})$	P_{th}	P_{th}	P_{th}
Post-Increase$_t$	3.170***	0.072***	2.599***	2.739***	3.697***
	(0.480)	(0.011)	(0.452)	(0.373)	(0.542)
$Peak$			-0.918		
			(0.674)		
Post-Increase$_t \times Peak$			0.940**		
			(0.395)		
Hour FE	✓	✓	x	✓	✓
Year FE	✓	✓	✓	✓	✓
N	2,833	2,833	2,833	1,114	1,719
aR^2	0.841	0.834	0.838	0.754	0.795

Notes: Standard errors in parentheses. * $p \leq 0.10$, ** $p \leq 0.05$, ***$p \leq 0.01$. Newey-west standard errors are computed. The sample consists of markets where non-dual natural gas-fueled generating units are marginal. Marginal cost and demand variables are included in all specifications.

Table 2: Main Price Analysis Results

which uses logged prices, indicates this corresponds to a 7.2% difference. Columns (3)-(5) report the estimate separately by hour type: column (3) includes an interaction with a peak-hour indicator, while columns (4) and (5) present subsample regressions run separately for off-peak and peak hours, respectively. The price difference is evident in both off-peak and peak hours, but larger during peak hours, suggesting that the factor driving asymmetric price responses is present across both hour types but more pronounced during peak hours, when strategic decision-making plays a greater role in price determination.

As an additional check, Appendix Table A1 reports estimates from a conventional pass-through regression using the cost-change sample, with indicators for the direction of cost changes. This regression also shows evidence of asymmetric pass-through: the pass-through rate is 11.6 percentage points lower for cost decreases than for cost increases. However, these estimates should be interpreted with caution, as they depend on the composition of observations along the increasing and decreasing cost paths, and therefore should not be viewed as the primary test of asymmetric pass-through in this paper.[30]

3.3.2 Markup difference

Since the equilibrium price is composed of cost and markup, the fact that prices do not return to pre-increase levels implies that markups are similarly not restored. While the main analysis indirectly captures the markup component by examining price differences conditional on marginal costs, I explicitly test how the markups of price-setting marginal units differ across samples using a variant of equation (1) with markup, $P_{th} - mc^*_{th}$, as the dependent variable. The markup here is the *market-level*

[30] As discussed in Section 2.1, average pass-through depends on the cost interval over which the marginal pass-through function is averaged. The regression comparing positive and negative cost changes uses all observations along the upward and downward paths, so the estimated pass-through rates depend on the composition and weighting of intermediate cost realizations, whose distribution generally differs between the increasing and decreasing samples. As a result, these estimates are not directly comparable to those obtained from the main empirical design.

	(1)	(2)	(3)	(4)
		Full sample	Off-Peak	Peak
			Markup_{th}	
Post-Increase_t	1.374***	0.871***	0.835**	1.833***
	(0.359)	(0.327)	(0.330)	(0.492)
$Peak$		-0.116		
		(0.357)		
$\text{Post-Increase}_t \times Peak$		0.843**		
		(0.370)		
N	2,833	2,833	1,114	1,719
aR^2	0.856	0.855	0.894	0.838

Notes: Standard errors in parentheses ($*p < 0.10, **p < 0.05, ***p < 0.01$). Standard errors are computed using newey-west method. All specifications include hour, year, and marginal-unit fixed effects. The sample consists of markets where non-dual natural gas-fueled generating units are marginal.

Table 3: Marginal Markup Analysis Results

markup, defined as the bid markup ($b_t^* - mc_{th}^*$) of the marginal firm-unit that sets the market price in day t, hour h.

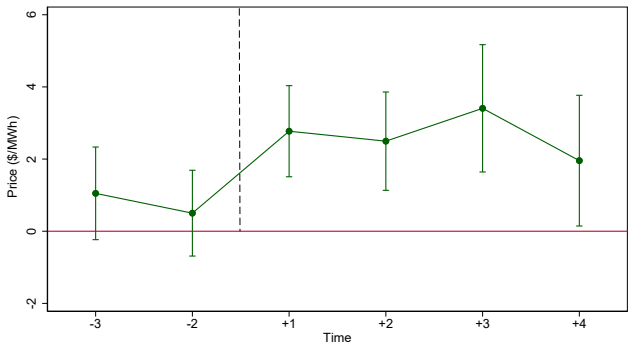
Table 3, column (1), shows that markups at the margin are on average higher in the post-increase sample than in the pre-increase sample. When further dividing the sample into off-peak and peak hours, as shown in columns (2)-(4), markups increased more in peak hours than in off-peak hours, though the increase in off-peak hours is also noticeable.

Several takeaways emerge from the markup analysis. In a static, competitive environment, there is no clear reason why markups are systematically higher in the post-increase sample and do not return to previous levels, conditional on all other market variables influencing firms' markup being the same. The literature typically attributes such patterns to consumer inattention and weak demand response – however, as demonstrated in Section 2.4, these frictions are absent in our setting, suggesting another factor is at play.

The finding that the average markup difference between the pre- and post-increase samples is smaller than the price difference in Table 2 suggests a non-uniform shift in the firm-level bid markup distribution. If all previously marginal firms had uniformly raised their markups in the post-increase sample, the overall bid distribution and composition of marginal units would have remained unchanged, and the markup difference would have matched the price difference. The smaller markup difference, therefore, reflects both a non-uniform shift in firm-level bid markups and changes in the identity and composition of firms near the margin.[31] This compositional shift itself is a meaningful finding that points to differences in firms' strategic responses across samples, especially since other confounding factors, such as market conditions or operational availability, are balanced across the two samples (see Section 3.3.4).

Market-level markups capture only the ex-post marginal firm's markup, which depends on both exogenous market variables and the bid distribution of all potentially marginal firms – and therefore do not reflect the overall picture of competition. To obtain a complete picture, Section 4 conducts a micro-level analysis of firm-specific bidding strategies.

[31] Figure B7 in the Appendix shows the change in the composition of marginal units between the pre- and post-increase samples.



Notes: The graph shows the estimated price difference ($\hat{\beta}_k$) for each period (*Time*) $k = \{-3, -2, +1, +2, +3, +4\}$, relative to the baseline comparison period $k = -1$. Time $k < 0$ is part of the pre-increase sample, and $k > 0$ is part of the post-increase sample. Confidence intervals are constructed using Newey-West standard errors (at the 5% significance level).

Figure 4: Price Dynamics: Price change $\pm k$ days since the cost-increase event

3.3.3 Price dynamics within the post-increase sample

Asymmetry in the speed of price adjustment or in the size difference? The estimate β captures the *average* price difference between two samples that differ in the direction of cost changes. A natural question is whether prices in the post-increase sample converge toward pre-increase levels, which relates to the literature on the *speed* of convergence. For instance, some studies define asymmetry as prices declining more slowly than they rise (adjustment speed)(Borenstein et al., 1997; Lewis, 2011), while other studies quantify asymmetry in terms of level differences in prices or pass-through rates (Peltzman, 2000; Benzarti et al., 2020). It is therefore important to distinguish whether the asymmetry I estimate reflects adjustment speed or the magnitude of price differences.

I argue that it reflects both. The observed price difference between the pre- and post-increase samples reflects differences in adjustment speed because market clearing occurs hourly, while the gas prices that determine marginal costs update daily: if prices after a cost decrease do not return to pre-increase levels despite firms optimizing on that same day, this indicates a slower adjustment. This differs from industries in which firms choose *when* to adjust prices, where the speed is measured by the time between adjustments. In my setting, because frequent adjustments are required by market design, the speed depends on how closely pricing tracks the underlying cost change once firms have the opportunity to adjust. Thus, adjustment speed depends not only on physical time but also on how pricing behavior aligns with market dynamics. Nonetheless, I present additional analyses of price adjustments over time in the following subsection.

Do prices in the post-increase sample decrease over time? I employ an event-study analysis to examine the price dynamics in the post-increase sample, treating each cost-change period as an event. I estimate the following specification:

$$P_{th} = \Big(\sum_{k \in \{-3,\ldots,+4\}} \beta_k D_{t-k} \Big) + \gamma X_{th} + \eta + \varepsilon_{th} \tag{2}$$

The coefficients β_k capture price patterns relative to the baseline $k = -1$ (the day before the event, part of the pre-increase sample). Coefficients for $k < 0$ reflect pre-increase prices, while those for $k > 0$ reflect post-increase prices.[32]

Figure 4 plots the estimated β_k, showing that post-increase prices are higher than the baseline for several days ($\beta_k > 0$ for $k > 0$), declining slightly over time though not strictly monotonically, while pre-increase prices are statistically indistinguishable from the baseline ($\beta_k \approx 0$ for $k < -1$). Overall, these results indicate that a systematic price difference between the pre- and post-increase samples persists for several days, with evidence of a gradual and non-monotonic decline. Although not the primary focus, this event study also helps identify the duration of post-increase pricing patterns, which informs the alternative post-increase sample categorization described in Section 3.2.1.

Incorporating cost-increase period information The cost-increase sample serves only as an exogenous event for dividing the data. Although competition in this market is static and market variables evolve stochastically – making past cost conditions not directly relevant to current pricing decisions – the characteristics of the preceding cost-increase episode may contribute to incomplete price adjustment in the post-increase sample through its effect on firms' strategic decision-making. To assess this possibility, I incorporate the magnitude and timing of cost declines into the price analysis using only the post-increase sample, finding that post-increase prices are higher when costs decline from a higher level and when the return to the base level is immediate. Details are provided in Appendix B.2.5.

3.3.4 Robustness Checks

The crucial condition underlying the analysis is that the pre- and post-increase samples share similar market environments, aside from the direction of the cost change. I conduct several robustness checks to verify this, with additional checks available in Appendix B.2.

The first check re-estimates the price regression excluding the first day of the post-increase period, addressing the possibility that bids and prices may not fully reflect the gas price decrease immediately. This concern may arise especially given that the analysis relies on data from day-ahead market – since the day-ahead market requires bids to be submitted roughly a day before actual generation. While this concern is mitigated by gas trading also occurring on a day-ahead schedule, columns (1) and (2) of Table 4 confirm that the estimates remain consistent with the main results, indicating the findings are not driven by a lagged generator response.

The second check addresses small differences in gas prices between samples. Although the main regression controls for marginal cost, I further restrict both samples to days and hours that closely match on gas price.[33] Columns (3) and (4) of Table 4 show that estimates are statistically significant

[32] This event study does not include a control group, unlike a conventional event study designed to estimate treatment effects. Instead, it characterizes the dynamics of prices around cost-change events. The cost-change period ($k = 0$) is excluded from the analysis.

[33] Specifically, for each post-increase day, I identify the most similar pre-increase day in terms of gas price, demand, daytime, and other factors, and repeat the process symmetrically for each pre-increase day.

	a. First-day Dropped		b. Sample Matched		c. Alternative cat		d. Unavailables		e. Duals included	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	P_{th}	$\log(P_{th})$	P_{th}	$\log(P_{th})$	P_{th}	$\log(P_{th})$	P_{th}	$\log(P_{th})$	P_{th}	$\log(P_{th})$
Post-Increase$_t$	3.081***	0.069***	3.627***	0.086***	3.222***	0.075***	3.368***	0.079***	2.991***	0.068***
	(0.491)	(0.012)	(0.680)	(0.017)	(0.436)	(0.011)	(0.477)	(0.012)	(0.447)	(0.011)
Hour FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
N	2,543	2,543	481	481	2,106	2,106	2,816	2,816	3,594	3,594

Notes: Standard errors in parentheses. * $p \leq 0.10$, ** $p \leq 0.05$, ***$p \leq 0.01$. Standard errors computed using Newey-west method. Price regressions in columns (1)-(8) use market outcomes when the non-dual natural gas units are marginal.

Table 4: Price Analysis: Robustness Checks

and larger in magnitude despite the restricted samples having virtually identical average gas prices, as shown in Table B7 in the Appendix.[34]

The next robustness check repeats the main analysis using an alternative post-increase sample categorization described in Section 3.2.1, where the first n days of each post-increase block are retained and the remaining days reassigned to the pre-increase sample of the subsequent event. The choice of $n = 4$ is guided by the price dynamics analysis in Section 3.3.3, with an example shown in Figure A2. Columns (5) and (6) of Table 4 show statistically significant results with larger magnitudes than the main regression. Further details are provided in Appendix B.2.

I also examine whether operational differences between samples can confound the documented price differences. The first concern is related to differences in generator availability between samples, which may arise if generators heavily used during the preceding high-gas-price period enter maintenance in the post-increase sample. I address this by checking the number and capacity of unavailable units, identified from bidding data, and find that they are not substantially different across samples – if anything, slightly larger pre-increase – as summarized in Appendix Table A2. Moreover, controlling for unavailable capacity in the regression yields estimates consistent with the main result, as shown in Table 4, Columns (7)-(8). A second concern applies to dual-fuel units that switched to oil during the high-gas price period and may continue using oil into the post-increase sample.[35] In this case, their marginal cost may not be accurately captured by the gas-based measure. To address this possibility, the main price analysis excludes observations where dual-fuel units are the marginal (price-setting) unit. However, even when including observations where dual-fuel units set prices, the result is not substantially different from the main estimates, as shown in Table 4, Columns (9) and (10) – suggesting that dual-fuel units' operational issues are not a significant concern. As a final check related to generator operational concerns, I specify marginal-generator fixed effects in the regression to absorb generator-specific cost or availability heterogeneity, and find a smaller but still statistically significant estimate, the magnitude of which is consistent with the markup difference between the samples analyzed in Section 3.3.2.[36] See Appendix Table A3 for details.

[34]Average gas prices are \$4.45/MMBtu for the pre-increase and \$4.43/MMBtu for the post-increase sample. Combining both checks – using the matched sample and dropping the first post-increase day – yields even larger estimates: β =3.35 for the price regression and 0.079 for the logged price regression.

[35]See Kim (2022) for details on fuel-switching behavior of dual-fuel units during the cost-increase period.

[36]This attenuation is expected, as this specification absorbs compositional shifts in which unit becomes marginal

4 Investigating the difference in the bidding strategy

In this market, firms influence the market price only indirectly, by strategically adjusting their bids ex ante through a multi-unit uniform price auction.[37] Because the market-clearing price is an outcome of firms' underlying bidding strategies, I examine whether firms' bidding behavior differs between the two samples. The dataset provides hourly observations of all firms' bidding decisions, allowing me to investigate changes in strategies at both the firm and generating-unit levels.

Description of Bidding In the market indexed by (t, h), firm i submits a bid schedule, $\langle b_{ijth}, q_{ijth} \rangle$, comprising price bid b_{ijth} and quantity bid q_{ijth} for each unit j. The bid reflects willingness to supply q_{ijth} at a price greater than or equal to b_{ijth}. Quantity bids are typically fixed, so firms primarily adjust price bids to respond strategically, maximizing expected profits while accounting for competitors' bids. A formal representation of optimal bidding is detailed in Section 6.

Although I observe all bids submitted by participating firms, only bids with a positive probability of setting the market price (i.e., ex-ante marginal bids) are strategically relevant in a multi-unit uniform price auction. Therefore, I examine bids of potentially marginal units, defined as those with price bids within a close range around the ex-post equilibrium price (e.g., ± \$5/MWh), most of which are natural gas-fired generators.[38]

Grouping firms by size and generation mix To capture heterogeneity in bidding patterns, I group firms by size and generation mix. I categorize firms into four size groups based on total generation capacity. Firms that operate a single small-scale unit below 100 MW are categorized as *Fringe* firms, known for nonstrategic bidding.[39] *Small*, *Medium*, and *Large* firms have total capacities of 100 – 500 MW, 500 – 1000 MW, and above 1,000 MW (up to 4,000 MW), respectively. All non-fringe firms operate multiple generating units and submit multi-step bids.

While firm size is the most frequently used indicator of firm strategy and market power, I further categorize firms by generation mix, focusing on the share of strategic gas capacity.[40] *Heavy Base* firms operate large-scale baseload units (e.g., nuclear) alongside some gas-fired marginal generators, though their strategic capacity share is relatively small. *Low Base* firms combine base generation (e.g., coal) with marginal generation (e.g., gas or oil), holding less base capacity and more gas generation than Heavy Base firms. *Gas Large* firms operate only gas-fired units, which are marginal in nearly every hour. *Gas/Oil Mix* firms run both gas- and oil-fired generators, relying mainly on gas units and using

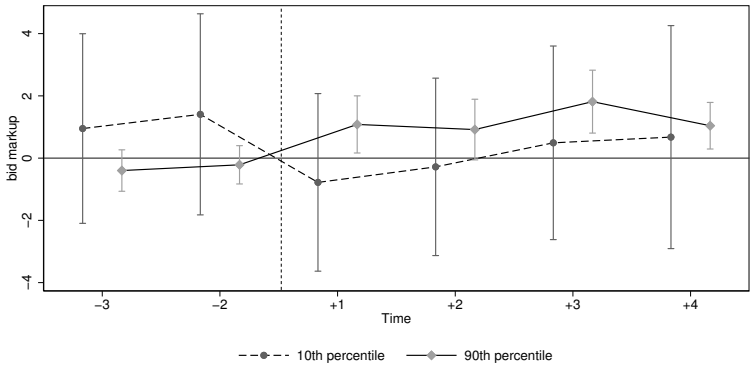
across samples. Since unit-level availability is balanced across samples, as shown above, this compositional shift is not attributable to differential availability, rather consistent with it reflecting differences in strategic bidding responses.

[37] This differs from markets in which each firm sets and receives a price for its own product, as explained in Appendix B.2.4. Nevertheless, competition in this setting is similar to price competition, as firms choose price bids that maximize their profits from an ex ante perspective, anticipating that they will receive that price ex post.

[38] I also computed the probability of being marginal for each firm-generator pair, $\partial P_{th}/\partial b_{ijth}$, by constructing residual demand curves from ex-post observed bids. Since both approaches yield similar sets of marginal units, I use the simpler method in the main analysis.

[39] Fringe suppliers typically submit single-step constant bids invariant across days and hours. I use bids b_{ijth} instead of bid markups for these firms due to small sample size and lack up bid markup variation, which makes estimating heat rates (marginal costs) difficult.

[40] The importance of generation mix for understanding strategic interactions in wholesale markets has been emphasized in Kim (2025) and Fabra and Llobet (2025).



Notes: The graph shows the estimated bid markup difference ($\hat{\beta}_{1,k}$) for 10th and 90th percentiles for each period (*Time*) $k = \{-3, -2, +1, +2, +3, +4\}$, relative to the baseline comparison period $k = -1$. Time $k < 0$ is part of the pre-increase sample, and $k > 0$ is part of the post-increase sample. Confidence intervals are computed using standard errors clustered at generating unit level (5% significance level).

Figure 5: Bid Markup Dynamics: ± k days since the cost-increase event (10^{th} and 90^{th} percentiles)

oil units only in peak hours. Finally, *Peakers* are those bidding high enough to be dispatched only during peak hours, earning scarcity rents in the hours when they are pivotal.[41]

4.1 Bidding differences between pre- and post-increase samples

I first estimate a regression using the bid markup ($b_{ijth} - mc_{ijt}$) of potentially marginal units as the dependent variable, shown in equation (3):[42]

$$b_{ijth} - mc_{ijt} \;=\; \alpha \;+\; \beta\,\text{Post-increase}_t \;+\; \gamma\,\text{Demand}_{th} \;+\; \eta \;+\; \varepsilon_{ijth} \tag{3}$$

The analysis includes bid markups of gas-fired generators, whose marginal costs are directly affected by gas price changes. To ensure consistency with the price regression, I use only markets (t, h) in which a (non-dual) gas-fired unit sets the price.

Table 5 shows the results of separate regressions of equation (3) for the full-hour sample (row A), off-peak hours (row B), and peak hours (row C). Column (1), row (A), shows a positive and significant coefficient for Post-Increase$_t$, indicating systematically higher bid markups relative to the pre-increase sample, with the increase more pronounced in peak hours (row C) than in off-peak hours (row B).[43] To examine how this increase is distributed across firms and over time, I also estimate a quantile regression of bid markups on event-time dummies, similar to Equation (2). Figure 5 reveals that the increase is not uniform: high-quantile markups are significantly higher than the pre-increase baseline, while low-quantile markups remain unchanged, indicating greater dispersion in bidding strategies across firms in the post-increase sample.[44]

[41] This pattern is consistent with strategic positioning near the margin, though it could also potentially reflect high start-up or ramping costs. Since these firms and units are identified from bidding patterns rather than verified engineering or cost data, I cannot distinguish between these explanations.

[42] A formal expression of bid markup is presented in Equation (7) in Section 6.

[43] Pooled regression results including a peak/off-peak indicator and the interaction term are available in Appendix Table A4.

[44] Appendix Figure B5 shows that this dispersion is more pronounced during peak hours.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Full	Firm Size				Firm Generation Mix				
		Fringe	Small	Med	Large	Hvy Base	Low Base	Gas Large	Gas/Oil	Peaker
Hours	$b_{ijth} - mc_{ijth}$	b_{ijth}	$b_{ijth} - mc_{ijth}$			$b_{ijth} - mc_{ijth}$				
(A) Full										
Post-Increase$_t$	1.024***	0.240	1.080***	0.917**	1.034***	0.416	1.008**	0.952***	1.314***	2.914**
	(0.350)	(0.741)	(0.378)	(0.399)	(0.333)	(0.470)	(0.429)	(0.315)	(0.429)	(1.076)
(B) Off Peak										
Post-Increase$_t$	0.399	-0.106	0.554	0.172	0.488*	-0.079	-0.374	0.601**	0.775*	–
	(0.265)	(0.595)	(0.331)	(0.315)	(0.261)	(0.354)	(0.625)	(0.231)	(0.374)	–
(C) Peak										
Post-Increase$_t$	1.510***	0.999	1.254**	1.736***	1.484***	1.950*	1.341**	1.611***	1.478**	3.521**
	(0.425)	(1.436)	(0.459)	(0.573)	(0.452)	(0.920)	(0.491)	(0.508)	(0.515)	(1.266)
N(full)	55,214	333	10,632	14,132	29,468	4,665	7,671	19,210	10,381	499
N(offpeak)	21,594	200	2,549	6,288	12,542	2,721	1,650	10,614	2,414	–
N(peak)	33,620	133	8,083	7,844	16,926	1,944	6,021	8,596	7,967	486

Notes: Standard errors in parentheses. $*p < 0.10, **p < 0.05, ***p < 0.01$. Standard errors are clustered at day (t) and firm (i) level. Day-hour (th) samples only when the non-dual natural gas generators set the price (i.e. marginal) are used for estimation. Demand variable (demeaned within each firm group, unit MWh) is also specified but not reported in this table. Hour, year, firm fixed effects are specified in all regressions.

Table 5: Bidding Differences: Pre- vs. Post-Increase

To further investigate heterogeneity across firms, I run separate regressions of equation (3) for each firm group. Table 5, columns (2)-(5), report results by firm size. Overall bid markups are higher in the post-increase sample for all firm groups except fringe firms (row A), driven mainly by peak hours (row C). Large firms also show higher post-increase bid markups in off-peak hours (row B). Notably, the pattern of higher bid markups is not fully explained by firm scale alone, motivating the generation-mix-based groupings that follow.

Columns (6)-(10) of Table 5 report results by generation mix, revealing greater heterogeneity across groups than the size-based grouping above. Gas Large and Gas/Oil Mix firms apply higher bid markups in both off-peak and peak hours, whereas Heavy Base and Low Base firms show increases mainly in peak hours. Peaker firms, which operate primarily during peak hours, also show higher post-increase bid markups. Overall, these firm-level results show that the increase in bid markups is widespread across firm types, though the timing and magnitude of the increase vary systematically with firms' generation mix and strategic capacity.

4.2 Bidding differences across hours within sample: peak-hour bidding pattern

Another way to investigate differences in bidding strategies is to examine whether the stylized *peak-hour bidding* pattern – where bid markups tend to be higher in peak than off-peak hours – is preserved in the post-increase sample. As discussed in Section 3.2.2, strategic considerations and market power are more pronounced in peak hours, leading firms to submit higher markups than in off-peak hours.

I run equation (4) separately for pre- and post-increase samples and by firm group, using bid markups of potentially marginal gas-fired units.

$$b_{ijth} - mc_{ijt} \;=\; \alpha \;+\; \beta\,\text{Peak}_h \;+\; \gamma\,\text{Demand}_{th} \;+\; \eta + \varepsilon_{ijth} \tag{4}$$

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Full		Firm Size			Firm Generation Mix			
		Fringe	Small	Med	Large	Hvy Base	Low Base	Gas Large	Gas/Oil
	$b_{ijth} - mc_{ijt}$	b_{ijth}	$b_{ijth} - mc_{ijt}$			$b_{ijth} - mc_{ijt}$			
(A) Pre-Increase									
Peak_h	0.786*	0.743	0.689*	0.599	0.898**	-0.205	0.727*	1.049***	0.732*
	(0.389)	(0.905)	(0.382)	(0.486)	(0.352)	(0.760)	(0.379)	(0.396)	(0.401)
(B) Post-Increase									
Peak_h	-0.103	-2.158	0.637	-0.014	-0.467	0.457	0.040	-0.335	0.204
	(0.364)	(1.579)	(0.434)	(0.385)	(0.314)	(0.518)	(0.454)	(0.310)	(0.393)
N(Pre)	27,585	203	5,249	7,082	14,749	2,207	4,206	8,972	5,510
N(Post)	27,629	130	5,383	7,050	14,719	2,458	3,465	10,238	4,871

Notes: Standard errors in parentheses. $*p < 0.10, **p < 0.05, ***p < 0.01$. Column (2) regression (fringe firms) uses the price bids (b_{ijth}) as a dependent variable due to the difficulty of estimating marginal cost (heat rates) for these units. Standard errors for regression (1) are clustered at the day (t) and firm (i) level, and standards errors for the subsample regressions (2)-(9) are clustered at day (t) level. Day-hour (th) samples only when the non-dual natural gas generators set the price (i.e. marginal) are used for estimation. Demand variable (demeaned within each firm group, unit MWh) is also specified in the regression but not reported. Hour, year, firm fixed effects are included.

Table 6: Bidding Differences: Peak-hour bidding pattern

Table 6 reports the results, with rows (A) and (B) corresponding to the pre- and post-increase samples, respectively. A positive and statistically significant estimate for Peak_h indicates higher bid markups during peak than off-peak hours, consistent with the typical peak-hour bidding pattern in this market.

The central finding is that peak-hour bidding is present in the pre-increase sample but absent in the post-increase sample, indicating a substantial change in firms' bidding strategies. As shown in Column (1) of Table 6, which reports average effects in the full sample, the estimate is positive and significant in the pre-increase sample (row A) but insignificant in the post-increase sample (row B).

Notably, the same result holds when the sample is broken down by firm group. Columns (2)-(5) of Table 6 show the same breakdown holds at the firm-size level – large and small firms both exhibit peak-hour bidding in the pre-increase sample, but not in the post-increase sample, including for the largest firms, which have the greatest strategic capability.[45] The difference between samples is even more pronounced when firms are grouped by generation mix, as shown in Columns (6)-(9) of Table 6. In the pre-increase sample, Low Base, Gas Large, and Gas/Oil Mix firms exhibit peak-hour bidding, most pronounced among Gas Large firms with the largest share of strategically usable gas-fired units, but this pattern disappears in the post-increase sample.

4.3 Summary and discussion

These findings have several takeaways. First, firms on average add higher bid markups in the post-increase sample, with the increase concentrated mostly in peak hours, where firms' bidding decisions are more influenced by strategic considerations. Second, the distinct peak-hour bidding pattern observed in the pre-increase sample disappears in the post-increase sample. Taken together, overall bid markup

[45]This result is not perfectly correlated with firm size, as medium-sized firms do not display a distinct bidding pattern.

levels are higher post-increase even as the distinctive peak-hour pattern breaks down, suggesting that bidding strategies differ from pre-increase strategies in a way that reflects greater dispersion across hours.

The results also reveal differences in bidding strategies across firm types. Although firm size is often associated with strategic bidding, firms' generation profiles – which shape their strategic positions – also appear to be strongly associated with higher bid markups in the post-increase sample. Specifically, Gas Large and Gas/Oil Mix firms – which hold a larger share of strategically usable generation – show a larger bid markup difference between the pre- and post-increase samples than Heavy Base and Low Base firms, with this difference observed across a broader range of hours.

Since market fundamentals remain similar across samples, the observed changes in bidding behavior likely reflect changes in firms' beliefs about competitors' strategies. In particular, greater bid dispersion and higher overall bids are consistent with firms anticipating a wider range of competitor strategies, potentially reflecting increased uncertainty about how competitors will bid. Section 5 presents a simple conceptual framework for why firms' beliefs and strategy sets may differ depending on the direction of the cost change, providing a basis for this interpretation. However, observed bids reveal beliefs only indirectly, through the strategic choices they generate, limiting what can be concluded from bidding patterns alone. Section 6 therefore develops a structural approach to estimate these beliefs directly and examine whether they reflect increased uncertainty about competitors' strategies.

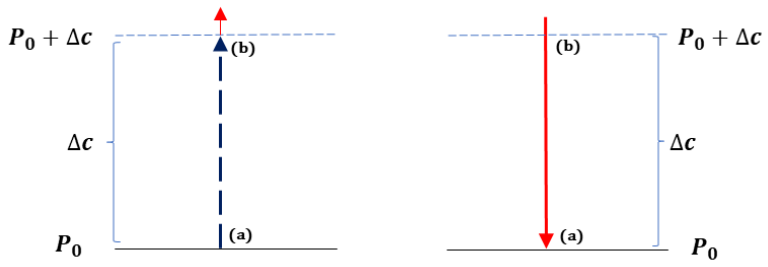
Discussion: Other possible causes Before turning to the conceptual framework and belief estimation, I consider whether alternative explanations for these patterns can be ruled out.

In relation to the pass-through literature, where incomplete pass-through is often attributed to markup adjustment, asymmetric price responses could also be seen as another realization of markup adjustment. Yet, if two periods face identical marginal costs and market conditions, firms should apply the same markup in both. Explaining the observed price asymmetry through markup adjustment alone thus requires an additional rationale for why markups differ between the pre- and post-increase samples – an asymmetry in incentives which is otherwise difficult to justify.

Another possible explanation is tacit collusion, noted as a hypothesis in Borenstein et al. (1997), where prices are sticky downward because past prices serve as focal points for coordination. Several features of my setting make tacit collusion unlikely. First, post-increase periods last less than a week, too short for collusion to develop. Moreover, firms cannot observe competitors' bids even after the market clears, since this information is disclosed only after four months.[46] Price leadership is also unlikely to emerge, since firms submit bids simultaneously rather than sequentially, and any attempt by one firm to lead would be unobservable to rivals in real time. Finally, the static nature of strategic interaction and stochastic variation in market conditions limit the formation of focal points or sustained collusion.

The empirical results also contradict the tacit-collusion explanation: coordination should be most detectable in peak hours, where detection risk is lower and market power is greater, implying a larger

[46]This market design follows FERC regulation, though the actual disclosure requirement varies across jurisdictions. While daily market-clearing prices and quantities are observable, the uniform-price mechanism means these data do not reveal any individual firm's bid.



(a) Cost Increase by Δc (b) Cost Decrease by Δc

Figure 6: Rationalizable Strategies Under Cost-Increase and Cost-Decrease Cases

peak/off-peak gap under collusion. Instead, this gap weakens in the post-increase sample, and bids become more dispersed across firms and hours – the opposite of convergence on a focal point. The overall price difference is also too small to plausibly reflect collusion.[47]

5 Belief and Strategy Differences as Sources of Asymmetric Price Responses

5.1 Asymmetry in the strategy set by the direction of the cost change

As discussed in Section 2.4, the empirical setting allows me to exclude many potential factors that could cause an asymmetric price response. The documented price differences can therefore be attributed to supply-side behavior in a non-cooperative strategic environment. Supporting this, the bidding analysis in Section 4 shows that firms adopt different bidding strategies in the post-increase sample, with bids becoming more dispersed and, on average, higher, departing from the stylized peak-hour pattern documented in the pre-increase sample (see Section 4.3 for a summary).

To understand this, I focus on differences in strategy sets between the cost-increase and cost-decrease cases, which are associated with firms' beliefs. Firms' beliefs about competitors' strategies are shaped by the competitive environment, implying that strategic uncertainty can vary across situations.[48] A change in input costs represents a shift in the competitive environment that can prompt firms to revise their beliefs about competitors' strategies. Although competition in this market is largely static, there is an intertemporal component to belief formation: by observing current and past gas prices, firms can infer the direction of cost changes – whether costs are rising or falling – and revise their beliefs about the strategic environment accordingly. Thus, differences in beliefs imply differences in the rationalizable strategy sets supported by those beliefs.

I propose a simple conceptual framework arguing that these rationalizable strategy sets inherently differ by the direction of the cost change: under a cost decrease, the set of rationalizable strategies is wider than under a cost increase. This asymmetry in strategy sets implies that downward cost movements are associated with greater strategic uncertainty than upward movements.

[47] Appendix B.3.2 provides further discussion of other possible causes (e.g., forward contracts and fixed costs).

[48] Strategic uncertainty can affect market outcomes, particularly in cases where multiple equilibria exist. For further discussion, see Morris and Shin (2002), Van Huyck et al. (1991), among others. Empirical studies have examined strategic uncertainty in the contexts of new market entry (Doraszelski et al., 2018) and uncertainty related to distant competitors' entry or exit (Aguirregabiria and Magesan, 2020).

To fix ideas, suppose firms initially charge a price P_0 when marginal cost is c_0, with an initial markup of zero. When cost increases by Δc, from c_0 to $c_1(= c_0 + \Delta c)$, it is common understanding that every firm must raise its price by at least Δc to maintain per-unit profitability. Raising the price by less than Δc results in a profit loss per unit sold, regardless of competitors' pricing decisions, and competitors would rationally expect the firm to increase its price by at least that amount. Therefore, any price below $P_0 + \Delta c$, i.e., $P_1 \in [P_0, P_0 + \Delta c)$, can be rationally excluded (strictly dominated) from the firm's strategy set, as such prices do not cover the increased marginal cost. Prices above this threshold, i.e., $P_1 \in [P_0 + \Delta c, \infty)$, reflect purely strategic considerations and belong to the firm's rationalizable strategy set. Given this set, firms are likely to pass through the cost increase fully or to a large extent. Figure 6 Panel (a) depicts this upward cost change situation.

On the other hand, when costs decrease, a firm's profit increases if everything else remains constant. Therefore, whether firms lower their prices is entirely a strategic decision, not driven by the need to avoid losses from cost changes. For example, suppose a firm's current price is $P_1 = P_0 + \Delta c$ and cost decreases by Δc, from c_1 to $c_0(= c_1 - \Delta c)$. The marginal cost reduction immediately increases the firm's per-unit margin, unless competitors also lower their prices. The firm will reduce its price only if it believes competitors to do so, in which case the profit loss arises from reduced residual demand. Thus, the expected profits associated with prices (strategies) in the range $P \in [P_0, P_0 + \Delta c]$ depend critically on competitors' strategies. Depending on firm's beliefs, setting a price in this range is not rationally excluded from the firm's strategy set, meaning that prices yielding incomplete pass-through is part of the rationalizable strategy set. Figure 6 Panel (b) depicts this downward cost change situation.

This simple example demonstrates that the set of rationalizable strategies differs depending on whether the firm's cost increases or decreases, highlighting the asymmetry in strategy sets.[49] When cost rises from c_0 to c_1, any price adjustment smaller than Δc is strictly dominated, leaving $P \in \left[P_0 + \Delta c, \bar{P}\right]$ as the rationalizable strategies. When cost decreases from c_1 to c_0, only prices below P_0 are strictly dominated. Any price in the range $P \in [P_0, P_0 + \Delta c]$ could be a best response to some belief about the competitor's strategy – for example, to a belief that the competitor may not reduce its price fully. Since rationalizability does not require beliefs be correct, these strategies are not strictly dominated. Thus, more strategies remain rationalizable compared to the cost-increase case.

While this example assumes an initial markup of zero (denoted as $\tilde{m} = 0$), the asymmetry persists even with a positive initial markup ($\tilde{m} > 0$). Unlike in the zero-markup case, a positive initial markup allows firms to sustain a cost increase without raising prices. The maximum per-unit profit loss a firm can absorb without altering the price is bounded by its initial markup, $\tilde{m}$ (or $\Delta\pi \leq \tilde{m}\tilde{q}$). Conversely, when costs decrease, the per-unit markup immediately rises by Δc without any price adjustment, making the firm's markup $\tilde{m} + |\Delta c|$. Firms will lower prices only if they believe competitors will do so, but even then, they can continue to avoid lowering the price longer than in the cost-increase case,

[49]The set of strategies that can be supported as best responses to some belief about competitor's strategy is termed the rationalizable strategies (Bernheim 1984; Pearce 1984). A firm may assign a probability distribution over its competitor's strategies, and any strategy to which it assigns positive probability is one it believes could possibly be played by the competitor, as the strategy has not been eliminated as strictly dominated. The firm's own rationalizable strategies, then, consist of best responses to such beliefs. Importantly, these beliefs do not have to be correct in the sense of matching the actual strategy of the competitor (i.e., consistency); rationalizability only requires that each strategy be optimal under *some belief*.

as they can sustain a per-unit profit loss up to $\tilde{m} + |\Delta c|$ (or $\Delta\pi \leq (\tilde{m} + |\Delta c|)\tilde{q}$).

5.2 Equilibrium Concepts

This reasoning in Section 5.1 does not depend on any particular equilibrium concept. Regardless of the solution concept employed, the direction of cost change – whether an increase or a decrease – can result in different rationalizable strategy sets, thereby shifting the domain from which any equilibrium strategy must be drawn. The specific equilibrium that emerges depends on the configuration of the game – that is, the underlying behavioral assumptions and structural features. Nevertheless, not all equilibrium concepts are consistent with the empirical findings – certain ones can be ruled out. For example, the simplest symmetric Bertrand model, which admits a unique Nash Equilibrium, predicts symmetric pricing responses regardless of the direction of cost change and therefore cannot rationalize empirical observations in which different pricing outcomes emerge in the cost-decreasing phase.

The most relevant model for this electricity market is a multi-unit uniform-price auction, which is known to admit multiple Bayesian Nash Equilibria. In this setting, the documented asymmetry in price responses across cost-increasing and cost-decreasing samples could reflect a shift from one equilibrium to another. If this is the case, analyzing the rationalizable strategy sets clarifies how the broader range of strategies in the cost-decreasing phase gives rise to equilibrium beliefs and outcomes that differ from those in the cost-increasing phase.

More importantly, the comparison between cost-increasing and cost-decreasing situations need not be limited to equilibrium beliefs, because rationalizability permits any subjective belief that makes the observed choice a best response, which need not coincide with equilibrium beliefs. Therefore, throughout the analysis, rather than committing to a particular equilibrium concept – which would impose an additional selection mechanism on top of the underlying strategy space – I focus directly on strategies and beliefs, which are common primitives across many equilibrium frameworks. To do so, I estimate a parametric distribution of beliefs over competitors' strategies by leveraging the condition that observed choices are best responses – an empirical approach that does not rely on specifying a particular equilibrium concept and is compatible with a broad class of equilibrium and non-equilibrium models. This approach is related to estimation methods that rely on best-response conditions rather than equilibrium assumptions (e.g., Hailey and Tamer 2003). In the electricity market, this approach benefits from the fact that marginal costs are relatively transparent and can be accurately measured from the data, which further facilitates the estimation of structural primitives such as beliefs without imposing equilibrium restrictions (see Section 6).

Because estimation based on the best-response condition does not require an equilibrium assumption, it also accommodates non-equilibrium models. For instance, non-equilibrium frameworks such as the level-k model or fictitious play have been studied in the context of electricity markets (Hortaçsu et al., 2019; Doraszelski et al., 2018), suggesting that such frameworks are empirically relevant in this setting. However, distinguishing between equilibrium and non-equilibrium explanations is not the objective of this paper, as the estimated beliefs are informative about the strategic environment without commitment to a specific behavioral model. In particular, these beliefs can be consistent with different forms of strategic reasoning – whether arising from full equilibrium behavior, limited

strategic sophistication, or other behavioral processes.[50]

To summarize, the goal of this paper is to show that the strategy sets rationalizing observed choices differ by cost direction – a difference that reflects different strategic environments and constitutes the strategic source of the asymmetric price response documented in this paper. Rather than identifying a specific equilibrium model that generates this asymmetry, this paper establishes the strategic foundations that any such model would need to accommodate.

6 Belief Estimation

In this section, I estimate firms' beliefs about competitors' strategies – modeled as ex ante bid distributions – and analyze how beliefs and strategy sets differ between the pre- and post-increase samples. Specifically, I recover beliefs that rationalize observed bids as best responses, without committing to any particular equilibrium concept (as discussed in Section 5.2). The only maintained assumption is that firms choose bids that maximize expected profit given their beliefs, which allows me to leverage the necessary condition of profit maximization to identify beliefs from observed bidding behavior.

Relaxing the equilibrium assumption is feasible in this setting because I do not need to jointly estimate other structural primitives, such as marginal costs, alongside beliefs. In the nonparametric estimation of auction models (e.g., Guerre et al. 2000; Hortaçsu and McAdams 2010), researchers impose a Bayesian Nash Equilibrium assumption to infer ex-ante beliefs (or the distribution of competitor valuations) from observed bids. This assumption is necessary to account for the uncertainty firms face about competitors' bids while simultaneously identifying other primitives of the model. However, in the electricity market, the marginal costs of electricity generators can be treated as known parameters, since they can be accurately measured from the data. This feature allows me to estimate the ex ante subjective distribution of beliefs over competitors' bids without relying on equilibrium conditions to map observed bids to beliefs.

Optimal Bidding Model I present a model characterizing a firm's bidding decisions in a multi-unit uniform-price auction, adapted to the wholesale electricity market context.[51] Suppose there are $i = \{1, \ldots, N\}$ firms, each operating J_i units, indexed by $j = \{1, \ldots, J_i\}$, that can generate electricity using multiple energy sources. In the day-hour auction, a firm submits price bids (b) and quantity bids (q) for each of its generating units. Given the market-clearing price P_{th}, the (ex-post) profit function of firm i in auction (t, h) is:

$$\pi_{ith}(\mathrm{b}_{ith}, \mathbf{b}_{-ith}) = P_{th}(\mathrm{b}_{ith}, \mathbf{b}_{-ith}) \left(Q_{ith}(P_{th}(\mathrm{b}_{ith}, \mathbf{b}_{-ith}))\right) - \sum_{j=1}^{J_i} C_{ijt}(\, q_{ijth}(P_{th}(\mathrm{b}_{ith}, \mathbf{b}_{-ith}))\,) \tag{5}$$

where b_{ith} denotes firm i's bid vector across its units and $\mathbf{b}_{-ith}$ collects bid vectors of all other firms. Q_{ith} is the total quantity generated by firm i's dispatched units, and q_{ijth} is the output of unit j, both depending on the market price $P_{th}(\mathrm{b}_{ith}, \mathbf{b}_{-ith})$.[52] Following the literature, the generation cost of unit

[50] See Appendix B.4.1 for further discussion.

[51] The model is closest to those in Kim (2022), Reguant (2014), and Wolak (2007).

[52] Forward-contracted quantity is omitted from the generation quantity for simplicity, though it is included in the estimation.

j is assumed linear in quantity, $C_{ijt}(q_{ijth}) = mc_{ijt}q_{ijth}$, implying a constant marginal cost mc_{ijt}.

Since firms submit bids simultaneously, they face uncertainty over competitors' bids ($\mathbf{b}_{-ith}$) and forms *beliefs* about their distribution ($\widetilde{\mathbf{b}}_{-ith}$). Firm i chooses the optimal bid b_{ith} that maximizes expected profit:

$$\max_{\mathrm{b}_{ith}} \ \mathbb{E}_{-ith}[\ \pi_{ith}(\mathrm{b}_{ith}, \widetilde{\mathbf{b}}_{-ith})\ ]$$

where the expectation is taken over the given belief, $\widetilde{\mathbf{b}}_{-ith}$.

Because firms lack incentives to bid strategically when they cannot influence the price, the optimality condition applies only to their ex-ante marginal units with a positive probability of setting the market price.[53] Although only one unit sets the price ex-post, multiple ex-ante marginal units can exist in a single auction due to uncertainty over competitors' bids and the market-clearing price.

Given the bid of firm i's unit j, denoted $\mathrm{b}_{ijth} = \langle b_{ijth}, q_{ijth} \rangle$, is ex-ante marginal in auction (t, h), the necessary first-order condition of profit maximization is shown in equation (6).

$$\mathbb{E}_{-ith}\left[\frac{\partial P_{th}}{\partial b_{ijth}}\left[\ (Q_{ith}(P_{th}))\ +\ \left(b_{ijth} -\ mc_{ijt}\right)\frac{\partial RD_{ith}}{\partial P_{th}}\ \right]\right] = 0 \tag{6}$$

Rearranging (6) yields (7), which expresses a firm's price bid as marginal cost plus a strategic markup.

$$\text{bid markup}_{ijth} = b_{ijth} - mc_{ijt}\ = -\ \frac{\mathbb{E}_{-ith}[\frac{\partial P_{th}}{\partial b_{ijth}}\ Q_{ith}\ ]}{\mathbb{E}_{-ith}[\frac{\partial P_{th}}{\partial b_{ijth}}\ \partial RD_{ith}/\partial P_{th}\ ]} \tag{7}$$

Since marginal cost mc_{ijt} can be measured from available data – a feature widely used in the literature (e.g., Wolfram, 1998; Hortaçsu and Puller, 2008; Hortaçsu et al., 2019), the bid markup in (7) is observable and becomes a function of firms' beliefs $\widetilde{\mathbf{b}}_{-ith}$. Specifically, the markup depends on the expected slope of the firm's residual demand and expected quantity – both determined by firm i's beliefs, $\widetilde{\mathbf{b}}_{-ith}$, through the expectation terms $\mathbb{E}_{-ith}(\cdot)$.

Estimation of beliefs about the competitors' bids I implement a simulated method of moments (SMM) estimation to estimate firms' beliefs, using hour-day bids submitted by firms and their generating units. The high-frequency bidding data provide rich variation in bid markups both within and across firms. I construct a moment condition such that the expected difference between observed and simulated bid markups is zero. Because the expectation over $\widetilde{\mathbf{b}}_{-ith}$ must be approximated via simulation, I employ SMM estimation (Pakes and Pollard, 1989).

The objective is to estimate parameters of firm i's belief distribution for each competing unit, such that the observed bids of firm i's ex-ante marginal units are rationalized as optimal. Firm i's belief distribution for each competing unit is parameterized as a log-normal distribution, consistent with existing simulation-based empirical auction studies (e.g., Laffont et al., 1995; Cantillon and Pesendorfer, 2006) and supported by the data.[54]

[53] An analytical expression for the probability of being an ex-ante marginal unit, $\frac{\partial P_{th}}{\partial b_{ijkth}}$, is provided in Appendix B.5.4.

[54] Figure B10 in the Appendix shows that log-transformed bid prices of a generating unit are approximately normally distributed.

Estimation method In each hour-day market (t,h), firm i faces a set of competitors indexed by $m=1,\ldots,M$. The price bids of competitors of firm i is $\mathbf{b}_{-ith} = \{b_1, b_2, \ldots, b_M\}$ with corresponding quantity bids $\{q_1, q_2, \ldots, q_M\}$. Throughout the analysis, quantity bids are fixed at the observed levels in market (t,h), as participants tend to adjust mainly the price bids when bidding.

The uncertainty about competitors' strategies implies that firm i forms a distribution over $\widetilde{\mathbf{b}}_{-ith}$, with draws from this distribution indexed by $s=1,\ldots,S$:

$$\widetilde{\mathbf{B}}_{-ith} = \begin{bmatrix} \widetilde{\mathbf{b}}^1_{-ith} & \widetilde{\mathbf{b}}^2_{-ith} & \ldots & \widetilde{\mathbf{b}}^S_{-ith} \end{bmatrix}^\top, \quad \widetilde{\mathbf{b}}^s_{-ith} = (b^s_{i,1}, \ldots, b^s_{i,M})$$

I assume that firm i's belief about the distribution from which competitor m's bid is drawn follows a log-normal distribution, *i.e.*, $\phi_{i,m} : \log(b_{i,m}) \sim N(\mu_{i,m}, \sigma^2_{i,m})$.[55] Then, the bid vector of competitors generated from the s draw is given by:

$$\widetilde{\mathbf{b}}^s_{-ith} = [\exp(\mu_{i,1} + \sigma_{i,1} Z^s_1), \ldots, \exp(\mu_{i,m} + \sigma_{i,m} Z^s_m), \ldots, \exp(\mu_{i,M} + \sigma_{i,M} Z^s_M)]$$

with each standard normal draw of $\mathbf{Z}^s = [Z^s_1, \ldots, Z^s_M]$ producing one realization of $\{b^s_{i,m}\}_{m=1}^M$.

For each simulated bid vector $\widetilde{\mathbf{b}}^s_{-ith}$, I compute the corresponding bid markup using equation (7), conditioning on firm i believing that the price bid of its marginal unit sets the ex-post market price (i.e., $P_{th} = b_{ith}$). Specifically, the simulated bid pairs $\langle b^s_{i,m}, q_m \rangle$ are ordered by increasing values of $b^s_{i,m}$ to form firm i's residual demand. I then compute the slope of the residual demand, $RD'(\widetilde{\mathbf{b}}^s_{-ith}; P_{th} = b_{ith})$, and total output, $Q_{ith}(\widetilde{\mathbf{b}}^s_{-ith}; P_{th} = b_{ith})$.[56]

This procedure is repeated for each simulation draw $s=1,\ldots,S$, using $S=500$ draws, generating simulated bid vectors $\{\widetilde{\mathbf{b}}^1_{-ith}, \ldots, \widetilde{\mathbf{b}}^S_{-ith}\}$. The simulated bid markup of firm i in market (t,h) is then computed as follows:

$$\text{bid markup}^{sim}_{ith} = \frac{\mathbb{E}_{-it}[Q_{ith}(\widetilde{\mathbf{b}}^s_{-ith}; P_{th} = b_{ith})]}{\mathbb{E}_{-it}[RD'_{ith}(\widetilde{\mathbf{b}}^s_{-ith}; P_{th} = b_{ith})]} \tag{8}$$

where details on how the expectation terms $\mathbb{E}_{-it}(\cdot)$ are approximated provided in Appendix B.5.1.

Selection of a Competitor Set Although strategic uncertainty includes uncertainty about competitors, I treat this information as fairly predictable and therefore do not endogenize the set of potential competitors in the estimation. Instead, I use the competitor set from the day-hour observation most closely resembling the candidate market (t,h), implicitly assuming that firms form beliefs based on past days with similar market conditions. A similar day-hour $(\bar{t},\bar{h})$ is selected from the pre-increase sample based on the closeness of daily gas prices, daily demand, hour, real-time demand, and daytime.[57] Given the selected similar day-hour $(\bar{t},\bar{h})$, I identify close competitors for each candidate

[55] To alleviate computational burden, I assume uncorrelated bid belief distributions across competitors – an independence restriction commonly imposed in private-value auction settings. Also, I assume no correlation of beliefs among multiple units owned by the same firm as only one generating unit of a firm appears as a competitor in each market (t,h).

[56] The slope RD' is obtained by fitting a linear specification to the logged price bids ($\log(b_{i,m})$) and corresponding quantity bids (q_m). This linear approximation is appropriate because only a small number of close competitors (five units on average) are used, making the slope a local estimate. To ensure that the marginal bid of firm i is close to the ex-post equilibrium price, I restrict (t,h) to markets in which firm i's marginal bid lies within ± \$5/MWh of the ex-post equilibrium price.

[57] For $(t,h) \in$ Post Increase, the similar day-hour is selected from pre-increase sample days. For $(t,h) \in$ Pre-Increase, the similar day-hour is chosen from pre-increase days prior to the candidate (t,h) (i.e., $\bar{t} \leq t$).

market (t, h) as follows.[58] First, I construct firm i's residual demand curve, $RD_{i\bar{t}\bar{h}}$, using the ex-post bids observed in market $(\bar{t}, \bar{h})$. I then select $M = 5$ close competing units $u^i_{m,th}$ for $m = 1, \ldots, M$ based on proximity to the ex-post equilibrium point of $RD_{i\bar{t}\bar{h}}$. Each competing unit corresponds to a firm-generator-step combination.[59]

Estimation The moment condition is the difference between the simulated bid markup in Equation (8) and the observed bid markup, denoted bid markup$^{obs}_{ith}$, in market (t, h). The sample mean of this difference averaged across days $(t = 1, \ldots, T)$ for a given hour h is:

$$\bar{g}^h_{i,T}(\theta_i) = \frac{1}{T} \sum_{t=1}^{T} \left[\text{bid markup}^{sim}_{ith}(\theta_i) - \text{ bid markup}^{obs}_{ith}\right]$$

I estimate belief distribution parameters, θ_i, only for a subset of competitors u^i_l $(l = 1, \ldots, L)$, selected based on their frequency in the estimation sample. The parameter set for firm i is $\theta_i = \{\mu_{i,1}, \ldots, \mu_{i,L}, \sigma_{i,1}, \ldots, \sigma_{i,L}\}$, which collects the distribution parameters describing firm i's beliefs about its competitors, where $\log(b_{i,l}) \sim N(\mu_{i,l}, \sigma^2_{i,l})$.

Although firm i faces M competitors in each market (t, h), denoted $U^i_{th} = \{u^i_{1,th}, \ldots, u^i_{M,th}\}$ for $t = 1, \ldots, T$ and $h = h_1, \ldots, h_H$, only a fixed set of competitors $U^i = \{u^i_1, \ldots, u^i_l, \ldots, u^i_L\}$ is parameterized and estimated for each firm i. For unparameterized competitors, I use bid data from the similar market $(\bar{t}, \bar{h})$, implicitly assuming firm i's beliefs about these competitors are degenerate at their ex-post bids.

The moment conditions and objective function are:

$$\bar{g}_{i,T}(\theta_i) = \left[\bar{g}^{h_1}_{i,T}(\theta_i) \quad \bar{g}^{h_2}_{i,T}(\theta_i) \quad \cdots \quad \bar{g}^{h_H}_{i,T}(\theta_i)\right]^\top$$

and,

$$\hat{\theta}_i = \underset{\theta_i}{\text{argmin}} \;\; \bar{g}_{i,T}(\theta_i)' \, W_T \, \bar{g}_{i,T}(\theta_i)$$

where W_T is the optimal GMM weighting matrix.

I estimate θ_i separately for off-peak and peak hours, and for the pre- and post-increase samples, with $L_{\text{off-peak}} = 4$ and $L_{\text{peak}} = 5$ competitors parameterized, respectively.[60] This yields four parameter sets: $\{\theta_{i,pre}, \theta_{i,post}\}_{h \in \text{off-peak}}$ and $\{\theta_{i,pre}, \theta_{i,post}\}_{h \in \text{peak}}$. Standard errors are computed using the GMM variance estimator adjusted for simulation error.[61]

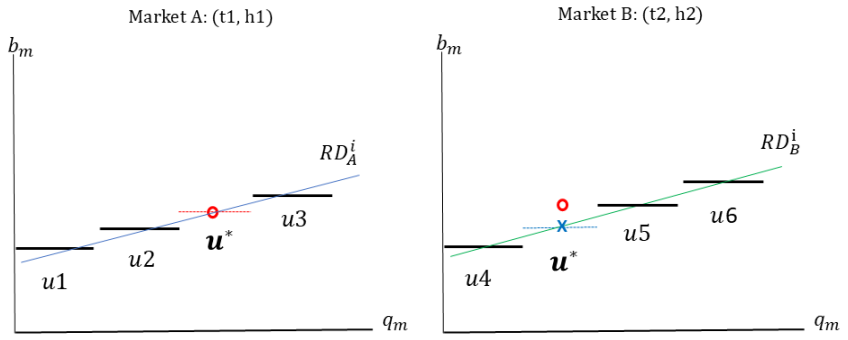
Identification Identification relies on variation in firm i's bid markups, competitor set U^i_{th}, and competitors' bids $b_{m,th}$ $(m \neq l)$ across day (t) and hour (h) markets where parameterized competitor u^i_l appears as a close competitor. Variation in firm i's bid markups reflects variation in the expected residual demand slope, which is the average outcome of the belief distribution. Identifying the mean $(\mu_{i,l})$ and dispersion $(\sigma_{i,l})$ of the belief distribution for u^i_l requires observing variation in firm i's bid markups across markets where the same competitor u^i_l appears as part of U^i_{th}, but alongside different

[58] See Appendix B.5.2 for a summary of competitors.

[59] Although I omit the step subscript k throughout the model for brevity, the smallest bidding unit in this auction is a step, and each firm-generator-step combination is treated as an independent unit.

[60] I restrict the number of parameterized competitors to satisfy the order condition – the number of parameters must not exceed the total number of moment conditions, which equals the total number of hours in the sample.

[61] The asymptotic GMM variance estimator is multiplied by $(1 + \frac{1}{S})$, where S is the number of simulation draws.



(a) Market A: $(t1, h1)$ observation (b) Market B: $(t2, h2)$ observation

Figure 7: Identification: Variation in Residual Demand and Competitors

sets of other competitors. Variation in the set of competitors other than the parameterized unit – namely, $u^i_{m,th}$ $(m \neq l)$ when $u^i_l \in U^i_{th}$ – is particularly important for identification.

Figure 7 illustrates the identification strategy with a simple example. Panels (a) and (b) show firm i's competitor set U^i_{th}, where u^* is the only parameterized competitor, and the corresponding (reversed) residual demand (RD^i), in markets A and B, respectively. In market A, firm i's competitor set is $U^i_A = \{u_1, u_2, u_3, u^*\}$. Unit u^*'s average price bid, denoted as $\bar{b}^*_A$ (red circle in panel (a)), is identified as the bid value b^* that minimizes the difference between the bid markup implied by RD^i_A and the observed bid markup.

If only the mean (μ) of u^*'s parameterized bid distribution is of interest, it can be identified from this single market observation. However, identifying both μ and σ from a single observation is not feasible, since different (μ, σ) can generate the same expected average bid consistent with RD^i_A and firm i's markup. For example, the identified average bid could be rationalized by either a degenerate distribution with $(\mu, 0)$ or a distribution with parameters $(\tilde{\mu}, \tilde{\sigma})$.

This problem is resolved by observing market B, where firm i faces a different competitor set $U^i_B = \{u_4, u_5, u_6, u^*\}$, also including u^*. A different set of competitors, together with firm i's bid markup, identifies the average price bid $\bar{b}^*_B$ (blue cross in panel (b)) as the bid value that minimizes the difference between the markup implied by RD^i_B and the observed markup. The intuition behind the identification is as follows. If the distribution with $(\mu, 0)$ rationalizing market A implies an average bid that does not coincide with $\bar{b}^*_B$ – shown as a red circle in panel (b) – this observation identifies the existence of σ and determines its magnitude. Observing multiple markets in which u^* appears as a competitor of firm i, alongside varying sets of other competitors, allows unique identification of both μ and σ.

Identification of (μ, σ) thus relies on observing multiple markets in which u^* competes with varying sets of competitors, generating variations in firm i's residual demand and optimal bid markups.[62] The

[62]Since the estimation pools observations across days, one may ask whether the small variation in gas prices – and thus marginal costs – across days is reflected in the estimated belief dispersion. This concern is unlikely to be valid, for two reasons. First, the GMM objective matches bid markups, net of the firm's own marginal cost. Second, identification relies on the residual demand slope constructed from competitors' bids: a small shift in the average level of marginal cost across days shifts competitors' bids roughly in parallel without necessarily altering the slope, which depends on how bids vary across competitors and quantities rather than on the common cost level they are centered around. The modest difference in daily gas-price variability between the pre- and post-increase samples is therefore unlikely to mechanically

Variable	Mean	Min	Median	Max	Std. Dev.
$\hat{\mu_{i,l}}$	3.5186	2.7068	3.5215	4.1276	0.1603
$\hat{\sigma_{i,l}}$	0.0147	0.0001	0.0032	0.1874	0.0302
Price bid Mean	3.5538	3.2933	3.5616	4.0444	0.1735
Price bid S.D.	0.1033	0.0607	0.0919	0.1816	0.0388

Notes: Summary statistics of the distribution parameter estimates (θ_i) at the competing unit level (u_l^i) are presented in the first two rows of the table. In the last two rows, I compute the mean and standard deviation of the competing unit's price bids during the sample period and report the summary statistics.

Table 7: Model Fit: Parameter Estimates and Bid Data Comparison

electricity market provides such variation naturally, as exogenously determined hourly demand levels generate changes in competitor sets across markets.

6.1 Results

I estimate a separate set of parameters $\theta_i = \{\mu_{i,1}, \ldots, \mu_{i,L}; \sigma_{i,1}, \ldots, \sigma_{i,L}\}$ characterizing firm i's beliefs about each of its L competitors, separately for the pre- and post-increase samples and for off-peak and peak hours. Although it is difficult to evaluate model fit given that the estimates represent subjective belief distributions, I compare the estimated $\mu_{i,l}$ and $\sigma_{i,l}$ to the mean and standard deviation of actual price bids of the corresponding competing unit. The estimated beliefs fall within a reasonable range of values: summary statistics of these estimates align closely with the empirical bid statistics, as shown in Table 7.[63]

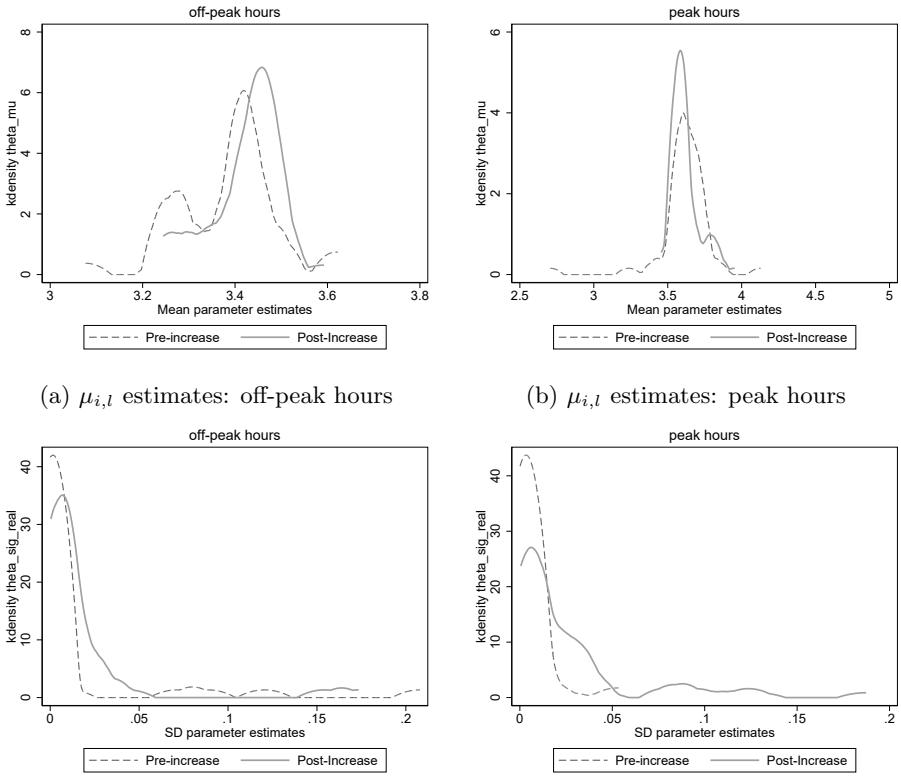
Figure 8 presents the distribution of estimates $\hat{\theta}_{i,\text{Pre}}$ and $\hat{\theta}_{i,\text{Post}}$, plotted separately for off-peak and peak hours. Panels (a) and (b) show estimates of the mean parameter ($\mu_{i,l}$) of firms' belief distributions, which are slightly larger in the post-increase sample than in the pre-increase sample, indicated by a rightward shift in the distribution. A more distinct difference appears in the standard deviation parameter ($\sigma_{i,l}$) of firms' belief distributions, shown in panels (c) and (d). In the pre-increase sample, estimated $\sigma_{i,l}$ values are concentrated around lower values in both off-peak and peak hours, indicating little dispersion in firms' belief distributions. In the post-increase sample, however, the estimates spread toward higher values, particularly during peak hours, indicating that beliefs become more dispersed following the cost decrease.

To show how much of these parameter differences translate into bid range differences, Figure 9 plots the range of price bids of a firm's competitor l that showed the largest increase in bid dispersion – that is, the competitor with the highest increase in the σ parameter among firm i's competitors.[64] I plot the 95% range of the selected competitor's logged price bids ($\log(b_{i,l})$) for both pre- and post-increase samples, separately for off-peak and peak hours, shown in panel (a) and (b), respectively. The bid ranges of these selected competitors become wider in the post-increase sample than in the pre-increase sample, extending beyond the range implied by the pre-increase belief estimates – though the extent

generate the estimated difference in belief dispersion.

[63]Because the estimated belief parameters are recovered from firm i's own bidding behavior – not from competitor l's bids directly – a close correspondence between the estimated parameters and competitor l's realized bid distribution is not mechanical; rather, it suggests that the model and estimation procedure yield reasonable belief estimates.

[64]Figure B4 in the Appendix shows the same graph, overlaid with the competitors' actual ex-post price bids (data), averaged within the sample.



(a) $\mu_{i,l}$ estimates: off-peak hours

(b) $\mu_{i,l}$ estimates: peak hours

(c) $\sigma_{i,l}$ estimates: off-peak hours

(d) $\sigma_{i,l}$ estimates: peak hours

Figure 8: Mean and Standard Deviation Estimates: Pre- and Post-Increase Samples

of the increase in dispersion varies across firms and hours.

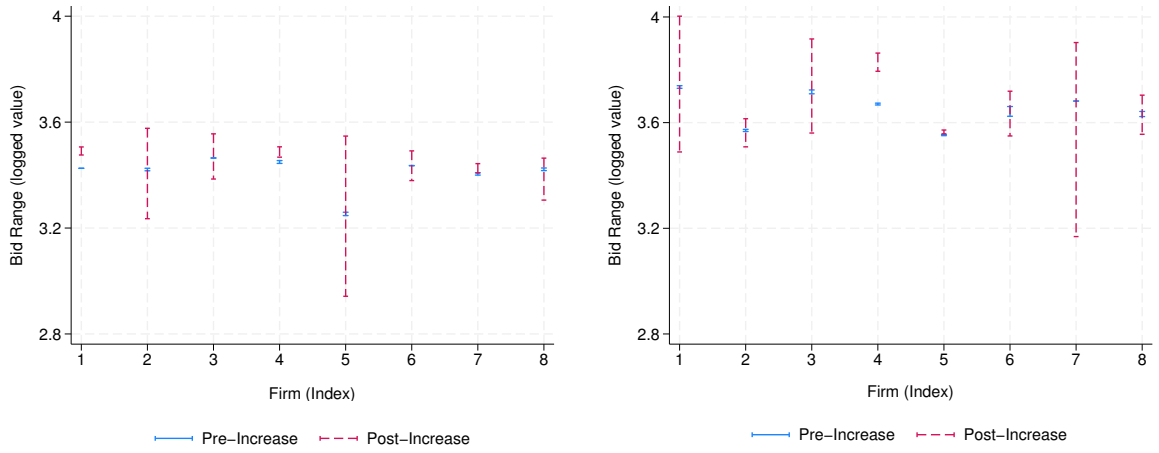
Lastly, I examine how firms adjust their bids toward competitors whose strategies they have become more uncertain about. These competitors are identified as those for whom firm i's belief dispersion increased the most between the pre- and post-increase samples (i.e., the largest $\Delta\hat{\sigma}$). I define $\text{Competitor}_{i,th}$ at the market (t,h) level, assigning a value of 1 if the competitor with the largest increase in $\hat{\sigma}_{i,l}$ appears as a close competitor of firm i.[65]

I estimate the following specification,

$$\begin{aligned} b_{ijth} - mc_{ijt} = \alpha + \beta\ \text{Post-increase}_t + \rho\ \text{Competitor}_{i,th} + \delta\ \text{Post-increase}_t \times \text{Competitor}_{i,th} \\ + \gamma\ \text{Demand}_{th} + \eta + \varepsilon_{ijth} \end{aligned} \tag{9}$$

using bid markup values from the data and the same control variables as in specification (3), with the addition of the $\text{Competitor}_{i,th}$ variable and its interaction with the post-increase sample indicator. The coefficient of interest is δ, on the interaction term, which captures whether the change in bid markups between the pre- and post-increase samples is different when firm i faces this high-uncertainty competitor than when it does not. Table 8 shows the results, where I find $\hat{\delta}$ is positive and statistically

[65]Because beliefs are formed based on ex-ante similar markets $(\bar{t},\bar{h})$ throughout the belief estimation, I assign a value of 1 to market (t,h) if the most uncertain competitor of firm i appears as a close competitor in market $(\bar{t},\bar{h})$. I do not include the estimated parameters in the regression directly – they are used only to identify the most uncertain competitor from firm i's perspective.



(a) Off Peak (b) Peak

Notes: The graph shows the 95% distribution range of the selected competitor's logged price bids, assumed to follow a normal distribution ($\mu \pm 1.96\sigma$). Firm Index 1-3: Gas Large firms, 4-5: Gas/Oil Mix firms, 6: Low Base, and 7-8: Heavy Base firms. Firms 1,2,3 and 6 are Large firms and firms 4,5,7, and 8 are medium firms.

Figure 9: Bid Range of the Competitor with the Largest Change in Belief Dispersion: by Firm

significant: firms, on average, bid higher when facing a competitor whose strategies they have become more uncertain about, relative to when they do not. Notably, bids in the pre-increase sample – when facing the same competitor – do not show a significant change, as reflected in $\hat{\rho}$, indicating that this differential response emerges only once the competitor's belief dispersion has increased, not as a pre-existing pattern.

When examining equation (9) at the firm level, reported in Table A5 in the Appendix, I find bid response patterns consistent with the full-sample analysis: firms adjust bids upward more when facing uncertain competitors. However, these patterns are more heterogeneous across firms, with the hours and magnitude of bid responses varying depending on the level of uncertainty they face. For example, as shown in Figure 9, Firm 1 faces greater uncertainty against its selected competitor during peak hours than during off-peak hours, whereas Firm 2 faces greater uncertainty in the off-peak sample (measured by the length of the bid range). This difference in uncertainty across hours is reflected in the bid responses shown in Table A5: $\hat{\delta}$ is positive and significant for Firm 1 in the peak sample, while Firm 2's estimate is positive and significant mainly in the off-peak sample. Moreover, the magnitude and significance of bid responses depend on firms' strategic ability: firms with limited strategic capacity (in terms of both size and idle gas-fired capacity) show weaker and, in some cases, statistically insignificant responses, even as the competitor's belief dispersion increases.

7 Conclusion

Asymmetric price responses to cost changes are a widespread phenomenon documented across many industries and product markets. More broadly, such asymmetry is related to downward price rigidity, an important feature of price dynamics that shapes the aggregate behavior of economic variables. Despite its prevalence and importance, understanding the sources of such asymmetry has

	$b_{ijth} - mc_{ijt}$		
	Full Hours	Off-peak Hours	Peak Hours
Post-increase$_t$ (β)	0.746**	0.305	1.456***
	(0.354)	(0.258)	(0.450)
Competitor$_{i,th}$ (ρ)	0.032	−0.046	0.285
	(0.277)	(0.235)	(0.704)
Post-increase$_t$ × Competitor$_{i,th}$ (δ)	1.424***	0.298	1.811*
	(0.431)	(0.385)	(0.938)
N	32,656	15,526	17,130
R^2	0.584	0.456	0.611

Notes: Standard errors in parentheses. $*p < 0.10, **p < 0.05, ***p < 0.01$. Standard errors for regression are clustered at the day (t) level. Day-hour (th) samples only when the non-dual natural gas generators set the price (i.e. marginal) are used for estimation. Demand variable (demeaned) is also specified in the regression but not reported. Hour, year, firm fixed effects are specified.

Table 8: Bid markup when facing a high $\Delta\sigma$ competitor

been challenging.

In this paper, I document asymmetric price responses to cost changes in a unique empirical setting where all previously studied sources of asymmetry are absent and show that this asymmetry can arise from firms adopting different strategies. I attribute this finding to inherent asymmetries in beliefs and strategies that depend on the direction of cost change. Specifically, a cost decrease leaves firms with a broader set of rationalizable strategies than a cost increase – a claim supported by the analysis of firms' bidding strategies and the structural estimation of underlying beliefs that rationalize ex-post bids. This, in turn, implies that incomplete responses to cost decreases can emerge solely from strategic considerations when uncertainty over competitors' strategies increases. Although this finding relies on specific context and setting, the strategic basis of asymmetric pass-through examined in this paper has broader implications, as strategic interaction among firms is a fundamental force in economics.

References

Aguirregabiria, V. and A. Magesan (2020). Identification and estimation of dynamic games when players' beliefs are not in equilibrium. *Review of Economic Studies 87*.

Amiti, M., O. Itskhoki, and J. Konings (2019). International shocks, variable markups and domestic prices. *Review of Economic Studies 86*(6).

Ball, L. and N. G. Mankiw (1994). Asymmetric price adjustment and economics fluctuations. *The Economic Journal 104*(March), 247–261.

Benzarti, Y., D. Carloni, J. Harju, and T. Kosonen (2020). What goes up may not come down: Asymmetric incidence of value-added taxes. *Journal of Political Economy 128*(12).

Bernheim, D. (1984). Rationalizable strategic behavior. *Econometrica 52*(4).

Bils, M. and P. J. Klenow (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy 112*(5).

Borenstein, S., A. Cameron, and R. Gilbert (1997). Do gasoline prices respond asymmetrically to crude oil price changes? *Quarterly Journal of Economics*.

Borenstein, S. and A. Shepard (2002). Sticky prices, inventories, and market power in wholesale gasoline markets. *RAND Journal of Economics 33*(1), 116–139.

Cantillon, E. and M. Pesendorfer (2006). Combination bidding in multi-unit auctions. *working paper*.

De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica 84*(2), 445–510.

Doraszelski, U., G. Lewis, and A. Pakes (2018). Just starting out: Learning and equilibrium in a new market. *American Economic Review 108*(3), 565–615.

Fabra, N. and G. Llobet (2025). Fossil fuels and renewable energy: Mix or match? *RAND Journal of Economics forthcoming*.

Fabra, N. and M. Reguant (2014). Pass-through of emissions costs in electricity markets. *American Economic Review 104*(9), 2872–99.

Goldberg, P. and R. Hellerstein (2013). A structural approach to identifying the sources of local currency price stability. *Review of Economic Studies 80*(1).

Gopinath, G. and O. Itskhoki (2010). Frequency of price adjustment and pass-through. *Quarterly Journal of Economics 125*(2), 675–727.

Green, R., D. Li, and N. Schurhoff (2010). Price discovery in illiquid markets: Do financial asset prices rise faster than they fall. *Journal of Finance 65*(5).

Guerre, E., I. Perrigne, and Q. Vuong (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica 68*(3), 525–574.

Hailey, P. and E. Tamer (2003). Inference with an incomplete model of english auctions. *Journal of Political Economy 111*(1), 1–51.

Hortaçsu, A., F. Luco, S. Puller, and D. Zhu (2019). Does strategic ability affect efficiency? evidence from electricity markets. *American Economic Review 109*(12), 4302–4342.

Hortaçsu, A. and D. McAdams (2010). Mechanism choice and strategic bidding in divisible good auctions: An empirical analysis of the turkish treasury auction market. *Journal of Political Economy 118*(5).

Hortaçsu, A. and S. 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.

Kim, H. (2022). Heterogeneous impacts of cost shocks, strategic bidding, and pass-through: Evidence from new england electricity market. *American Economic Journal: Microeconomics 14*(2), 370–407.

Kim, H. (2025). Cleaner but volatile energy? the effect of coal plant retirement on market competition in the wholesale electricity market. *Working paper*.

Laffont, J.-J., H. Ossard, and Q. Vuong (1995). Econometrics of first-price auctions. *Econometrica 63*(4), 953–980.

Levine, S., P. Carpenter, and A. Thapa (2014). Understanding natural gas markets. *American Petroleum Institute*.

Lewis, M. (2011). Asymmetric price adjustment and consumer search: An examination of the retail gasoline market. *Journal of Economics and Management Strategy*.

MacKay, A. and M. Remer (2022). Consumer inertia and market power. *working paper*.

McKay, A. and R. Reis (2008). The brevity and violence of contractions and expansions. *Journal of Monetary Economics 55*(4).

Morris, S. and H. S. Shin (2002). Measuring strategic uncertainty. *working paper*.

Neumark, D. and S. A. Sharpe (1992). Market structure and the nature of price rigidity: Evidence from the market for consumer deposits. *Quarterly Journal of Economics 107*(2), 657–680.

Pakes, A. and D. Pollard (1989). Simulation and the asymptotics of optimization estimators. *Econometrica 57*(5), 1027–1057.

Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica 52*(4).

Peltzman, S. (2000). Prices rise faster than they fall. *Journal of Political Economy 108*(3).

Reguant, M. (2014). Complementary bidding mechanisms and startup costs in electricity markets. *Review of Economic Studies 81*(4), 1708–1742.

Tappata, M. (2009). Rockets and feathers: Understanding asymmetric pricing. *RAND Journal of Economics 40*(4), 673–687.

Van Huyck, J. B., R. C. Battalio, and R. O. Beil (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics 106*(3), 885–910.

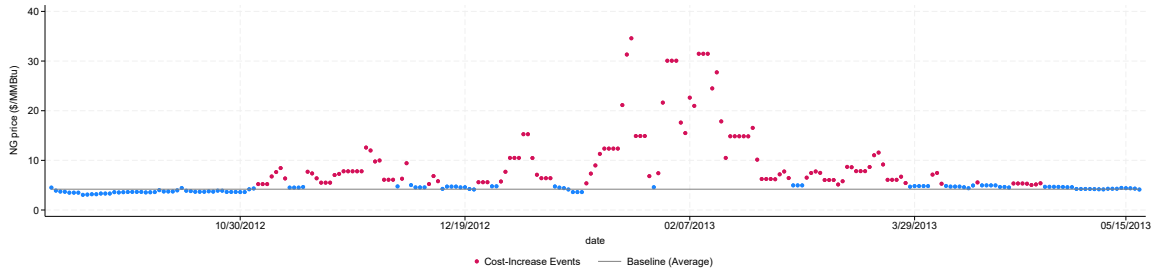
Wolak, F. (2007). Quantifying the supply-side benefits from forward contracting in wholesale electricity markets. *Journal of Applied Econometrics 22*(7), 1179–1209.

Wolfram, C. D. (1998). Strategic bidding in a multiunit auction: An empirical analysis of bids to supply electricity in england and wales. *RAND Journal of Economics 29*(4), 703–725.

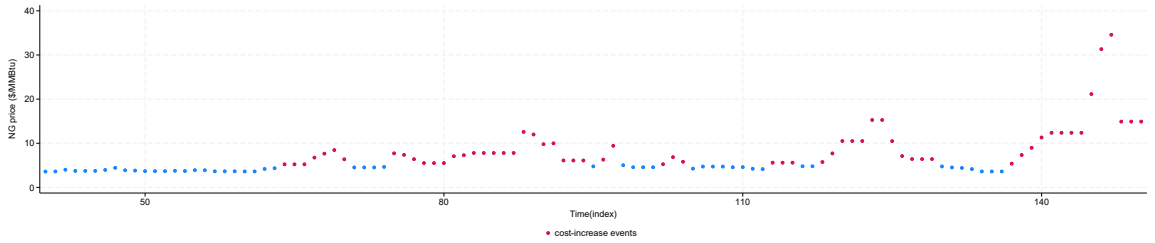
Wu, X., M. S. Lewis, and F. A. Wolak (2024). Search with learning in the retail gasoline market. *RAND Journal of Economics 55*(2).

Yang, H. and L. Ye (2008). Search with learning: understanding asymmetric price adjustments. *RAND Journal of Economics 39*(2), 547–564.

Appendix



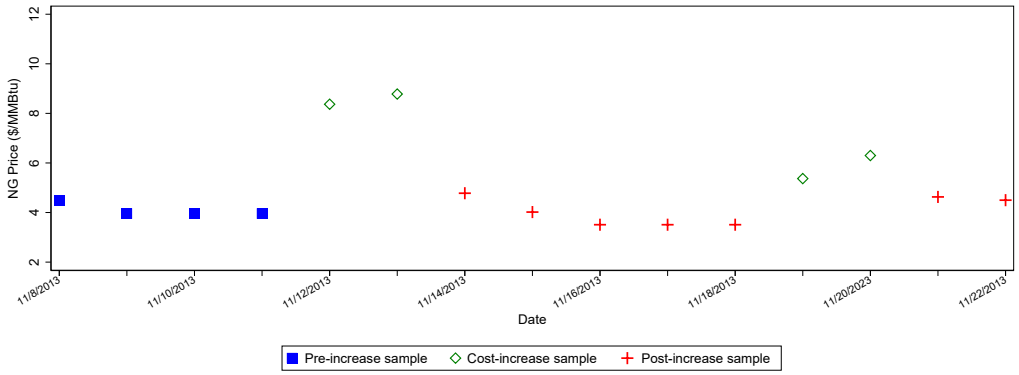
(a) Natural Gas Prices: 2012-2013 Sample



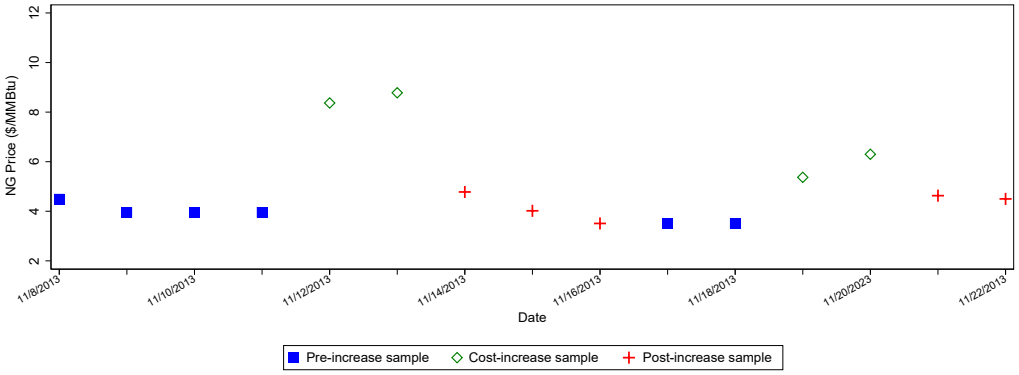
(b) Sub Sample (expanded)

Notes:The figure presents a time-series plot of daily natural gas spot prices, where cost-increase events are highlighted in red and the remaining sample – which is further divided into pre- and post-increase periods – is shown in blue. Panel (b) is a zoomed-in version of the subsample shown in Panel (a).

Figure A1: Cost-Increase Events Identified from the Data



(a) Main Sample Division



(b) Alternative Sample Division

Notes: The figure presents a subset of the dataset, selected to visualize how days are classified into each sample group. Daily natural gas prices are shown together with sample classifications: squares: pre-increase sample; crosses: post-increase sample; diamonds: cost-increase sample. Panel (a) shows the main sample division, and panel (b) shows the alternative sample division.

Figure A2: Visualization of Sample from Data: Main vs. Alternative Categorization

	(1)
	P_{th}
mc_{ijth}	0.996^{***}
	(0.0257)
$D_+ \times mc_{ijth}$	0.116^{**}
	(0.0479)
N	2,128
adR^2	0.924

Notes: Standard errors are reported in parentheses. Pass-through is estimated using only the *cost-change* sample, in which natural gas prices differs from the base level. $D_+ = 1$ indicates that the natural gas price is on an increasing path ($P_{ng}(t) > P_{ng}(t-1)$), and $D_+ = 0$ indicates a decreasing path ($P_{ng}(t) < P_{ng}(t-1)$). Variable mc_{ijth} denotes the implied marginal cost of the marginal unit (ij) in market (t, h), constructed following Kim (2022). The interaction term $D_+ \times mc_{ijth}$ captures the difference in average pass-through between increasing and decreasing cost paths. Both mc_{ijth} and its interaction with D_+ are instrumented using the natural gas spot price. All specifications control for demand and include year and hour fixed effects. Standard errors are clustered at the marginal-firm level.

Table A1: Pass-through using cost-change sample: Increasing vs decreasing, IV regression

Sample	Variable	Mean	Med (p50)	p25	p75	SD
Pre	Number of Units	18.93	19.0	16.0	22.0	3.18
	Total Size (MW)	725.71	463.2	338.4	1009.0	497.34
Post	Number of Units	16.74	16.0	14.0	19.0	3.84
	Total Size (MW)	636.34	355.4	329.0	1009.0	456.04

Notes: Generators undergoing maintenance (or with status marked as *unavailable*) are identified from the bidding data, since generators must indicate unavailability when bidding in this market, by adjusting the econMax parameter or price/quantity bids. *Number of Units* is the count of unavailable generating units, and *Total size (MW)* is the offered quantity that could have been generated by these units, approximating their effective unavailable capacity. Unavailability is assessed at the daily level.

Table A2: Unavailable Generating Units: Number and Size by Pre- and Post- samples

	(1)	(2)	(3)	(4)
	Unit FE		Firm FE	
	P_{th}			
Post-Increase$_t$	2.014^{***}	2.000^{***}	2.878^{***}	2.863^{***}
	(0.420)	(0.428)	(0.428)	(0.431)
Marginal Unit FE	✓	✓	x	x
Marginal Firm FE	x	x	✓	✓
Hour FE	x	✓	x	✓
Year FE	✓	✓	✓	✓
N	2831	2831	2833	2833
adj. R^2	0.891	0.893	0.862	0.864

Notes: Standard errors in parentheses. $*p < 0.10, * * p < 0.05, * * *p < 0.01$. Standard errors for regression are clustered at the day (t) and hour (h) level. Day-hour (th) samples only when the non-dual natural gas generators set the price (i.e. marginal) are used for estimation.

Table A3: Price Regression with Marginal Unit and Marginal Firm Fixed Effects

	(1)	(2)
	$b_{ijth} - mc_{ijt}$	
Post-Increase$_t$	1.024**	0.552*
	(0.350)	(0.321)
Peak$_h$		−0.110
		(0.372)
Post-Increase$_t$ × Peak$_h$		0.784**
		(0.340)
Hour FE	✓	x
Year FE	✓	✓
Firm FE	✓	✓
N	55,214	55,214
R^2	0.656	0.652

Notes: Standard errors in parentheses. $*p < 0.10, **p < 0.05, ***p < 0.01$. Standard errors clustered at the day (t) and firm (i) level. Day-hour (th) samples only when the non-dual natural gas generators set the price (*i.e.*, marginal) are used for estimation. Demand variable is demeaned (unit: MWh).

Table A4: Bid Markup Regression: Full Sample

				$b_{ijth} - mc_{ijt}$				
Firm Type:		Gas Large		Gas/Oil Mix		Low Base	Heavy Base	
Firm Size:		Large		Medium		Large	Medium	
Firm (Index)	Firm 1	Firm 2	Firm 3	Firm 4	Firm 5	Firm 6	Firm 7	Firm 8
Sample:	$\hat{\delta}$: Post-Increase$_t$ × Competitor$_{i,th}$							
Full	2.272***	1.178**	0.373	−1.952	−1.977	1.251	1.787**	4.151***
	(0.576)	(0.501)	(0.630)	(1.412)	(2.436)	(1.081)	(0.899)	(0.985)
Off peak	0.544	0.912**	−0.496	1.374	−1.098	0.299	−1.297*	2.151
	(0.576)	(0.449)	(0.784)	(1.225)	(2.544)	(0.922)	(0.766)	(2.188)
Peak	3.684***	2.457	1.725**	−6.677***	−2.530	2.719*	5.523***	2.930*
	(1.376)	(2.935)	(0.772)	(1.587)	(3.247)	(1.522)	(2.060)	(1.717)
Sample:	$\hat{\rho}$: Competitor$_{i,th}$							
Full	−1.041**	−0.063	0.475	1.793*	−1.052	0.356	−1.337**	1.297**
	(0.423)	(0.365)	(0.526)	(1.039)	(1.915)	(0.543)	(0.649)	(0.524)
Off peak	-0.281	−0.076	0.732	−1.063	−1.031	0.718	−0.043	0.266
	(0.326)	(0.317)	(0.551)	(1.054)	(1.939)	(0.561)	(0.538)	(0.619)
Peak	−0.784	−0.436	−1.242**	3.191**	−1.137	−0.506	−3.081**	2.619**
	(1.120)	(2.733)	(0.540)	(1.538)	(1.834)	(0.633)	(1.258)	(1.304)

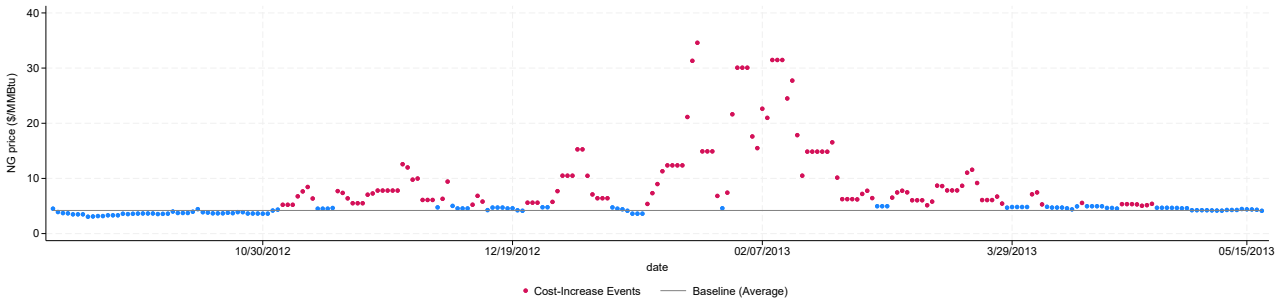
Notes: Standard errors in parentheses. $*p < 0.10, **p < 0.05, ***p < 0.01$. Standard errors for regression are clustered at the day (t) level. Hour and year fixed effects are included. Day-hour (th) samples only when the non-dual natural gas generators set the price (i.e. marginal) are used for estimation. For Firm 2, the most peak hours (16-22h) among the original peak-hour category were used. Unit fixed effect is specified for Firm 6 due to its large number of generating units in the sample.

Table A5: Firm-level Regression: Bid Markup with a high $\Delta\sigma$ competitor present

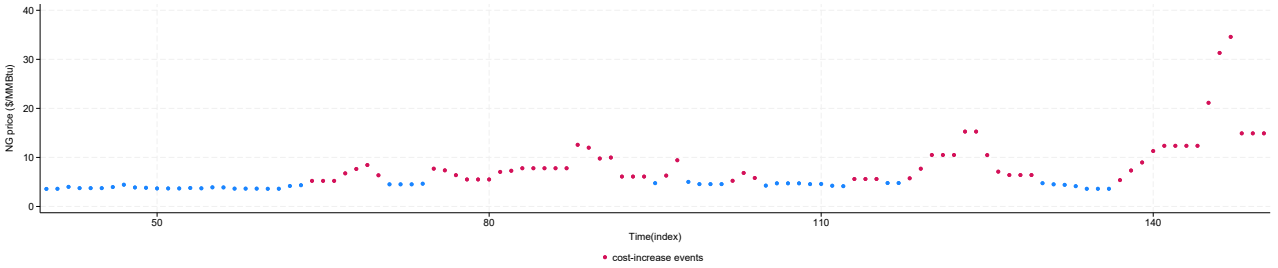
B Online Appendix

B.1 Data and Sample

B.1.1 Cost-Increase Event Categorization Used for Sample Division



(a) Natural Gas Prices: 2012-2013 Sample



(b) Sub Sample (expanded)

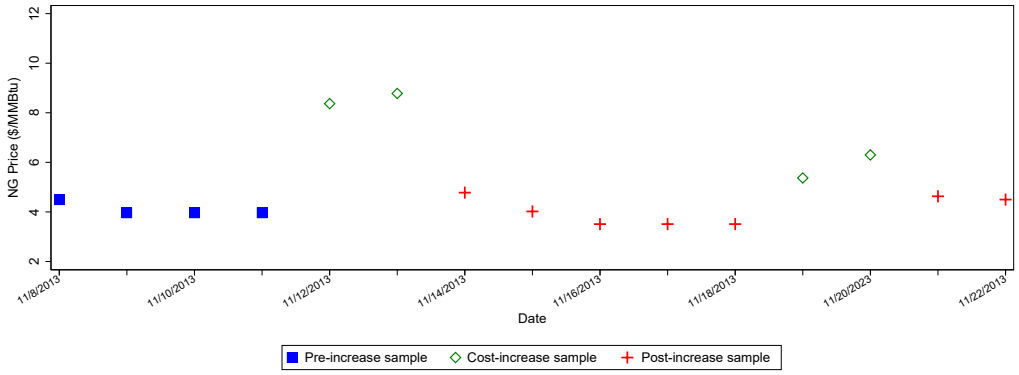
Notes:The figure presents a time-series plot of daily natural gas spot prices, where cost-increase events are highlighted in red and the remaining sample – which is further divided into pre- and post-increase periods – is shown in blue. Panel (b) is a zoomed-in version of the subsample shown in Panel (a).

Figure B1: Cost-Increase Events Identified from the Data

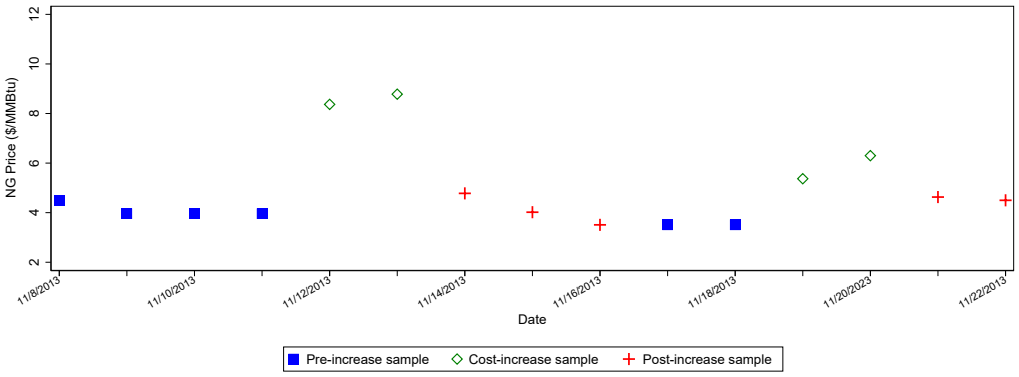
The benefit of the natural gas variation I exploit is that the upside volatility feature is distinguishable even from visual inspection of the time series of natural gas prices, which eases identifying the cost-increase events: there is a normal range at the lowest base level of gas prices, and any departure from this level can be categorized as a cost-increase (or cost-change) event. Figure B1 Panel (a) shows the natural gas price index data for 2012-2013, with a zoomed-in subsample presented in Panel (b). From both panels, a baseline level of approximately \$4.5/MMBtu can be identified as the prevailing floor of gas prices – roughly corresponding to the average value marked by the horizontal line in the figure. Although gas prices fluctuate slightly around this baseline level (between \$3.9 – 5.2/MMBtu) they cluster tightly around this baseline. Since all observed values remain below the threshold of \$5.5/MMBtu, I set the maximum threshold κ at \$5.5/MMBtu.

When defining cost-change events (also denoted as cost-increase events in the text), I consider the direction of gas price changes as the primary criterion, alongside the level of gas prices relative to the baseline. Specifically, a given day is classified as part of a cost-change event only if the gas price on that day exceeds the baseline threshold κ and the gas price is moving in a consistent direction – either rising or falling – relative to the previous day. I categorize day t that satisfies the condition of $\kappa < P_{ng}(t-1) < P_{ng}(t)$ (i.e., the increasing part of the cost-change event) or $P_{ng}(t-1) > P_{ng}(t) > \kappa$ (i.e., the decreasing part of the cost-change event) into the cost-change event sample. This method of dividing sample observations into cost increases and decreases is commonly used in the asymmetric pass-through literature (e.g., Borenstein et al., 1997; Peltzman, 2000), where authors typically assign

dummies by the direction of cost change. Using this approach, I identified a total of 19 cost-increase events.



(a) Main Sample Division



(b) Alternative Sample Division

Notes: The figure presents a subset of the dataset, selected to visualize how days are classified into each sample group. Daily natural gas prices are shown together with sample classifications: squares: pre-increase sample; crosses: post-increase sample; diamonds: cost-increase sample. Panel (a) shows the main sample division, and panel (b) shows the alternative sample division.

Figure B2: Visualization of Sample from Data: Main vs. Alternative Categorization

B.1.2 Alternative Pre- and Post-Increase Sample Division

As explained in the robustness check section (Section 3.3.4), I also implemented the analysis with an alternative categorization of the samples by further dividing the post-increase sample days into pre- and post-increase samples. In the main sample division, all base-level gas price days between two cost-increase events were grouped into the post-increase sample. In this new categorization, I reclassified n days before the next cost-increase event, previously part of the post-increase sample, into the pre-increase sample. For example, if a post-increase sample block consists of K days followed by another cost increase event, the main categorization assigned all K days to the post-increase sample. However, because the few days immediately preceding another cost-increase event can be viewed as part of the pre-increase sample, I instead assigned the first n days to the post-increase sample and the remaining $K-n$ days to the pre-increase sample of the subsequent cost increase event. Figure B2 illustrates this comparison between the main and alternative sample divisions using actual data. The figure displays a subsample of the dataset to clearly show how days are classified into each sample group.

This alternative categorization presents practical challenges. Determining the valid n cannot be done arbitrarily, thus the result of dynamic price pattern analysis implemented in Section 3.3.3 helps

identify up to how many days in the post-increase sample exhibit similar price patterns, thus can be grouped into post-increase sample. I find qualitatively similar results when trying $n = 2, 3, 4$, with even stronger results under this alternative categorization, which I discussed in Section 3.3.3. The results for the $n = 4$ case were presented in Columns (5) and (6) of Table 4. While the alternative categorization is feasible and yields qualitatively similar and even stronger results than under the main categorization, I opted for the main categorization due to a relatively short length of the post-increase sample block (average around five days) which implies that the underlying structure of these days within the original block may not differ significantly.

B.1.3 Marginal Cost

The empirical advantage of studying electricity markets is that marginal costs can be directly measured from data, a feature widely utilized in empirical studies. This enables direct assessment of strategic bidding behavior through markups (e.g., Wolfram 1998) and examination of deviations from equilibrium (e.g., Hortaçsu et al. 2019).

A generating unit's marginal cost comprises of its physical efficiency, i.e., heat rate, and the prices of fuel and, when applicable, emissions permits. Some studies use heat rates provided by private data sources (e.g., Hortaçsu and Puller 2008; Hortaçsu et al. 2019) or infer them from CEMS data on generation and fuel use. Structural studies instead estimate marginal costs implied by optimal bidding behavior (Wolak, 2007; Reguant, 2014; Kim, 2022).

I adopt a hybrid approach, using heat rates structurally estimated from an optimal bidding model but computing daily marginal costs by multiplying these rates by observed fuel prices. Using the first-order condition in Equation (6), I estimate heat rates via GMM with the following specification:

$$mc_{ijt} = hr_{ij}(FP_t + \tau e_{ij}) \tag{10}$$

Here, the fuel price, FP_t, and emissions cost τe_{ij} (permit price multiplied by emissions factor) are observable, so estimating marginal cost is equivalent to estimating unit-level heat rates.

For estimation, I use a separate sample with stable gas prices that does not overlap with the main analysis period, particularly the post-increase sample. This estimation sample covers October and March of 2012-2013 (days before the first cost-increase event of the winter season occurs), while the main analysis relies on November-February of years 2012-2014.

This estimation assumes Bayesian Nash Equilibria (BNE) and employs nonparameteric resampling methods that rely on consistent beliefs to approximate the expected bid distribution. Thus, it is important to use a sample without abnormal pricing or bidding behavior. Further details are provided in Kim (2022).

A single heat rate parameter is estimated for each generating unit, representing its physical efficiency, which does not generally vary across days or hours. The daily marginal cost of each unit is then computed using Equation (10) with data on daily fuel prices (Natural Gas Intelligence), emissions permit prices (Regional Greenhouse Gas Initiative), and emissions factors (EIA).

Although the belief estimation does not impose BNE, using this parameter estimated under the assumption of BNE remains consistent for the following reasons. First, the sample used for estimating the heat rate does not exhibit anomalous pricing or bidding patterns, nor any significant changes in costs, making it valid to assume equilibrium behavior. This sample also does not overlap with the pre- and post-increase samples analyzed in this study. Moreover, since the same heat rate parameter is used in both pre- and post-increase samples, any observed asymmetry is not driven by estimation error in the heat rates. As a measure of physical efficiency, the heat rate parameter is fixed in principle and is not endogenously affected by differences in strategic behavior or beliefs across samples.

B.2 Price Analysis

B.2.1 Price Regression Only With the Pre-Increase Sample

To check if the same results will be obtained when using only the days in the pre-increase sample, I implemented a price regression using hypothetical post-increase sample selected from the days in the pre-increase sample. For each day in the original post-increase sample, I selected a most similar day having the closest features regarding the following market variables; daily maximum demand, daily average demand, peak hour maximum and average demand, and gas price levels. Since gas price levels are similar across pre- and post-increase samples, the priority is given to the similarity in demand levels – the maximum and average demand of the peak hours – when selecting similar days. In this way, I have generated counterparts of post-increase sample days among the pre-increase days, and assigned a hypothetical post-increase sample indicator, Hy. Post-Increase$_t$, to this sample. I dropped the original post-increase sample in the analysis. Table B1 shows the results of this first robustness

	(1) P_{th}^{da}	(2) log (P_{th}^{da})	(3) P_{th}^{da}
Hy. Post-Increase$_t$	0.376	0.010	0.524
	(0.59)	(0.015)	(0.533)
Peak$_h$			0.824
			(0.541)
Hy. Post-Increase$_t$ × Peak			−0.189
			(0.625)
Hour FE	✓	✓	x
Year FE	✓	✓	✓
N	1,236	1,236	1,236

Notes: Standard errors in parentheses. * $p \leq 0.10$, ** $p \leq 0.05$, ***$p \leq 0.01$. Newey-west standard errors with lag $m = 17$. Using only the markets in which the non-dual natural gas units are marginal.

Table B1: Robustness Check: Using Hypothetical Post-Increase Variable

check. Despite choosing the most similar days of the original post-increase sample among the pre-increase sample, none of the estimates are significantly different from zero. This suggests that the original post-increase sample has some structurally different strategic environment that is not present in the pre-increase sample.

B.2.2 Price Regression with Subsample: Post-increase Days with Lower Gas Prices

The average gas price in the original post-increase sample is slightly higher– by about \$0.2/MMBtu. This minor discrepancy is a natural variation, as data points cannot be identical. Although it does not violate the sample division or weaken the credibility of the analysis, it is important to check whether the main results are sensitive to this difference in average gas prices. To do so, I construct a subsample of the original post-increase sample by selecting days with relatively low gas prices, such that the new post-increase subsample has a lower average gas price than the pre-increase sample. The modified post-increase sample has an average gas price of \$4.17/MMBtu, lower than the average gas price of \$4.29/MMBtu of the pre-increase sample.[66]

If the positive coefficient in the main regression were driven by the slightly higher gas prices in the post-increase sample, the estimated $\hat{\beta}$ from this subsample regression should not be positive and

[66]Because I selected days to shift the average gas price, demand became slightly more unbalanced across samples, though not substantially. Specifically, the new sample has an average demand of 12,801 MWh, lower than that of pre-increase sample's 13,255 MWh.

	(1) P_{th}^{da}	(2) log (P_{th}^{da})	(3) P_{th}^{da}
Post-Increase$_{t,h}$	1.440**	0.030*	1.309**
	(0.584)	(0.016)	(0.553)
Peak			0.313
			(0.403)
Post-Increase$_{t,h}$ × Peak			0.256
			(0.510)
Hour FE	✓	✓	x
Year FE	✓	✓	✓
N	1,774	1,774	1,774

Notes: Standard errors in parentheses. * $p \leq 0.10$, ** $p \leq 0.05$, ***$p \leq 0.01$. Newey-west standard errors with lag $m = 17$. Using only when the non-dual natural gas units are marginal.

Table B2: Robustness Check: Low gas price sub sample of the post-increase sample

significant. However, Table B2 shows that the coefficient remains positive and significant even when the post-increase sample has lower average gas prices and lower demand. This indicates that the main price regression result is not simply driven by small discrepancies in the market variables.

B.2.3 Pass-through rate estimation

I also estimated a variant of Equation 1 to capture differences in pass-through rates between the pre- and post-increase samples. The challenge with this specification is the limited variation in marginal costs within the sample, which is not ideal for estimating pass-through. Nevertheless, as a robustness check – and because the specification in Equation 11 is more common in the literature – I estimate ρ_1. Results in Table **??** show that the rate is low overall, though higher in the post-increase sample.

$$P_{th} = \beta_0 + \rho_0\, mc_{th} + \ \rho_1 mc_{th}\, D\ + \gamma X_{th} + \mu_h + \eta_t + \varepsilon_{th} \tag{11}$$

Second, I implemented a standard pass-through specification, estimating the rate of price change relative to marginal cost change separately for cost-increasing and cost-decreasing samples, using cost-increase event days instead of using pre- and post-increase samples. Table A1 shows that the estimated rate is lower for the cost-decreasing sample than for the cost-increasing sample, though the sample size is small and this estimation remains subject to the issues discussed in Section 2.1.

B.2.4 General applicability of results: comparing single market price and firm-level pricing

The pricing mechanism in this market differs from a typical setting in which firms set their own product prices – instead, the market clears through a uniform-price auction, resulting in a single market price for all firms. Two considerations justify the broader applicability of the analysis despite this difference.

First, while firms in firm-level pricing receive the prices they set, firms in a uniform price auction receive the market-clearing price, which may differ from their bids. However, the bids of interest here are those from firms expecting their bids to set the market price (i.e., ex-ante marginal bids). The market setting is therefore not fundamentally different from one where firms set prices to maximize expected profit, making the uniform-price mechanism comparable to firm-level pricing. Second, using a single market price is comparable to using representative price indices, a common approach in pricing studies. The uniform price indirectly reflects the distribution of individual firm-level prices,

much like a price index. Indeed, even in markets with firm-level pricing, most empirical studies of price movements or pass-through effects rely on aggregate or average price measures – pass-through analyses in particular frequently use sales-weighted average price indices to capture meaningful price variation. This approach is also standard in macroeconomics and trade economics research on price adjustments.[67]

B.2.5 Incorporating the cost-increase period information

Throughout the analysis, the cost-increase events served only as an exogenous event used for dividing the samples. Although firms' optimization is largely static, their beliefs and strategies adopted in the post-increase period could be influenced by the characteristics of the cost-increase period.

To assess this, I run a variant of Equation (1) to examine how the price difference across samples is correlated with the magnitude and timing of cost decrease. The first variable, gp_{t-1}, represents the gas price on the last day of each cost-increase event – before the price returned to the post-increase range at time t, which captures the magnitude of the cost decrease. However, in some cases, the gas price declined gradually rather than immediately. For example, as seen in the Event #1 of Figure 2, the price dropped over a couple of rounds before returning to the base level. I define gp_{LastInc} as the gas price on the last day before this decline began and also introduce an indicator, Im_{dec}, to capture whether the return to the base level was immediate.[68] Since these variables are undefined for days before the cost-increase event, the pre-increase sample days are assigned a value of zero for these variables. Therefore, the Post-Increase$t$ indicator is excluded and replaced with the newly generated variables.

Table B3 shows the results. As shown in column (2) regression, a higher gp_{LastInc} is associated with a statistically significant increase in the post-increase sample prices, suggesting that a cost decline from a higher level leads to a more incomplete adjustment of prices. However, $gp_{\text{t-1}}$ has no significant effect, shown in column (1), likely due to limited variation within the sample and the dissipation of effects across multiple rounds of cost declines. Column (3) shows a positive and significant coefficient of Im_{dec}, which indicates that post-increase prices are higher when the cost decrease is immediate rather than spread over several rounds.

	(1)	(2)	(3)	(4)	(5)
			P_{th}		
gp_{t-1}	0.172			-0.00711	
	(0.242)			(0.265)	
gp_{LastInc}		0.108***			0.120***
		(0.0369)			(0.0373)
Im_{dec}			1.467**	1.474***	1.592***
			(0.580)	(0.564)	(0.603)
N	2,833	2,833	2,833	2,833	2,833

Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01

Table B3: Price regression: with characteristics of the cost-increase period

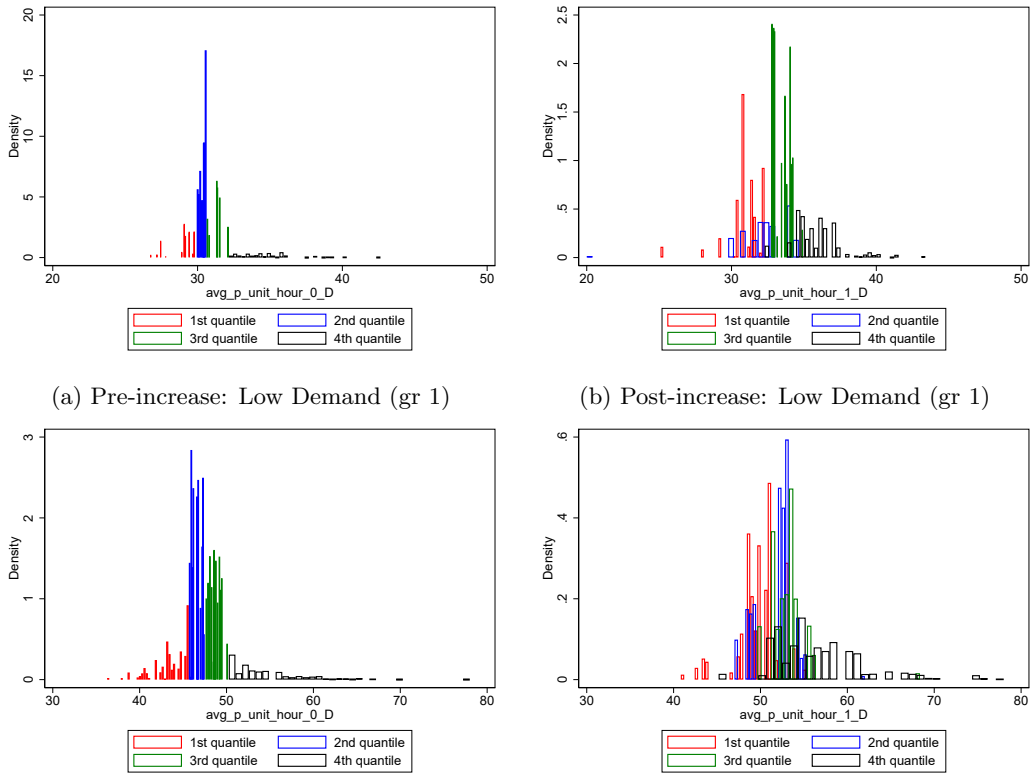
B.3 Bidding Analysis

B.3.1 Bid distribution change: quartile analysis

An initial examination of the bidding data suggests a change in bidding patterns. I compare the distribution of price bids, b_{ijkth}, between the pre- and post-increase samples, as shown in Figure B3.

[67] See Goldberg and Hellerstein (2013) for a review.

[68] Pre-increase sample days are assigned $Im_{\text{dec}} = 0$, assuming a sufficiently long gap since the last cost increase.



(a) Pre-increase: Low Demand (gr 1)

(b) Post-increase: Low Demand (gr 1)

(c) Pre-increase: High Demand (gr 4)

(d) Post-increase: High Demand (gr 4)

Figure B3: Distribution of (volume-weighted) price bids of pre- and post-increase samples: by demand

Price bids of generating units j of the pre-increase sample are ranked in increasing order, and index values from 1 to 4 are assigned to units in each quartile, separately for each demand group (e.g., low to high demand).[69] I then plot the bid distributions of these units for the post-increase sample in increasing order and examine whether their quartile indices match those of the pre-increase sample. Panels (a) and (b) show the distributions for the lowest-demand subsample (e.g., off-peak hours), while panels (c) and (d) correspond to the highest-demand subsample (e.g., peak hours).

If bidding strategies were consistent across the pre- and post-increase samples, the bids in panels (b) and (d) would be expected to align with those in panels (a) and (c). However, the distributions do not preserve their original order: units in the first and second quartiles of the pre-increase sample tend to bid higher, and bids are more dispersed across all quartiles in the post-increase sample. This suggests that the original bid rankings are not preserved – previously low-bidding units no longer bid the lowest, and vice versa. The mixing of quartile groups indicates that firms may be employing different bidding strategies in the post-increase sample.

B.3.2 Discussion: other possible causes

Forward contract Electricity-generating firms often set up forward contracts with the demand side (retail companies), which can reduce their responsiveness to wholesale market conditions and limit strategic behavior when forward obligations are binding (see Bushnell, Mansur, and Saravia, 2008). One might question whether the documented differences between samples could result from forward

[69] I also tried an alternative method of using the price bid of a single closest-to-marginal step bid of each unit j instead of the volume-weighted bids, where I find qualitatively similar graphs.

contracts. However, this can be ruled out, as forward contract terms should be the same for sample days that differ only in cost direction. Firm-level forward contract rates are typically fixed for at least a season and do not change over short periods. Since the pre- and post-increase samples in our analysis are at most a month apart, the forward contract rate can be assumed constant.[70]

Fixed costs Another factor is the fixed cost of power plant operation, such as start-up or ramping costs. Power plants incur an initial fixed cost when starting or additional costs when ramping up. These fixed costs can introduce dynamic considerations into production decisions. However, if these costs exist, they are present in both pre- and post-increase periods, not specific to either sample. Moreover, these costs typically create dynamic linkages in production across hours, but not across days. Given that the price asymmetry I observe is mainly across days in different samples, it cannot be explained by fixed costs. Finally, fixed start-up costs are typically relevant for conventional baseload generators, such as nuclear or coal, but not for gas-fired generators. Because my analysis focuses on marginal cost changes to gas-fired generators, fixed costs are not a concern and cannot account for the observed asymmetry.

B.4 Model

B.4.1 Equilibrium Concept

Determining whether the observed behaviors are better explained by an equilibrium or a non-equilibrium model would require additional testing and model comparison, similar to the approach in Hortaçsu et al. (2019). In their setting, off-equilibrium behavior by small firms could be clearly identified, making it evident that a standard Bayesian Nash Equilibrium (BNE) model could not explain all observed outcomes.[71] In contrast, such identification is much more difficult in my setting. It is not straightforward to falsify or rule out the possibility that the observed outcome reflects one of multiple equilibria. Thus, even testing whether an observation is on- or off-equilibrium is inherently challenging. As a result, model selection – determining which framework best fits the data – is complicated by the lack of clear predictions from either class of models, especially when multiple equilibria cannot be excluded. Moreover, even if one wishes to compare an equilibrium model with a non-equilibrium model, the choice of which non-equilibrium framework to use is non-trivial. For example, the level-k model used in Hortaçsu et al. (2019) requires identifying a type-0 (non-strategic) player, but in my context, it is not obvious which firm – or what behavior – should be classified as less sophisticated. Similarly, models based on fictitious play (as in Doraszelski et al. (2018)) rely on players forming beliefs based on past play, but in this setting, firms make decisions that cannot be rationalized as responses to observable past competitor behavior.

B.4.2 Why not dynamic decision-making?

The strategic interaction in this market is not dynamic in the sense that there is no dynamic feed loop between actions and state variables, though there is a dynamic element in terms of the belief formation.

[70] Using the rate instead of the actual amount of forward contracts is not an issue. Even if the rates are proportional to production, total hourly output remains nearly the same in both samples after controlling for marginal cost, demand, and other market conditions.

[71] Moreover, the analyses in Hortaçsu and Puller (2008) and Hortaçsu et al. (2019) were feasible in part because the model imposed a specific structure: uncertainty entered only through the forward contract, which induced a single equilibrium path characterized by parallel shifts in bidding strategies. The model I adopt differs from theirs in that I allow uncertainty to arise from changes in the residual demand slope.

The strategic interaction in this market is not dynamic in the conventional sense, as there is no feedback loop between firms' actions and state variables. Firms aim to maximize static payoffs (profits), which are realized each period and depend solely on contemporaneous state variables. These state variables – such as gas prices and demand – are exogenously determined, with transition probabilities given by $f(X_{t+1})$ rather than $f(X_{t+1}|Y_t, X_t)$, since both gas prices and demand are driven by weather and residential consumption behavior, factors outside the model. Consequently, X_t is stochastic, and firms have no incentive to choose actions today to influence competitors' behavior in future periods.

The game is simultaneous rather than sequential, and it is not a repeated game, as market fundamentals X_t change stochastically each period. Unlike standard dynamic games, where current actions depend on past states, actions, or observed outcomes, and players incorporate historical fundamentals into current decisions, the present setting lacks such dynamics in fundamentals. Instead, the intertemporal element arises solely through the evolution of beliefs. Specifically, firms' beliefs about competitors' strategies today, denoted $B_{i,t}$, depend on the current state variable X_t, an error term ε_t, and the direction of cost change $d_t = \text{sgn}(g_t - g_{t-1})$, where $(g_t - g_{t-1})$ is exogenously determined. In this sense, the setting retains a limited dynamic feature through belief updating, as past state variables – particularly the gas price level g_{t-1} – influence firms' expectations and perceived strategies of competitors in the current period.

B.5 Estimation

B.5.1 Estimation Method

Expectation Approximation The expectation terms that appear in Equation (8) are approximated as follows:

$$E_{-ith}\left[RD'(\widetilde{\mathbf{b}}_{-ith})\right] = \prod_{m=1}^{M}\int_0^{\infty} RD'(\widetilde{\mathbf{b}}_{-ith})\,\phi_{i,m}(b_{i,m})\,db_{i,m} \approx \frac{1}{S}\sum_{s=1}^{S} RD'(\widetilde{\mathbf{b}}^s_{-ith})$$

$$E_{-ith}\left[Q_{ith}(\widetilde{\mathbf{b}}_{-ith})\right] = \prod_{m=1}^{M}\int_0^{\infty} Q_{ith}(\widetilde{\mathbf{b}}_{-ith})\,\phi_{i,m}(b_{i,m})\,db_{i,m} \approx \frac{1}{S}\sum_{s=1}^{S} Q_{ith}(\widetilde{\mathbf{b}}^s_{-ith})$$

where $\phi_{i,m}(\cdot)$ denotes firm i's belief distribution of competitor m's bids $b_{i,m}$. I use $S = 500$ draws to construct these simulated expectation terms.

B.5.2 Competitor summary

Table B4 shows the competing units summarized by their firm type. In summarizing I grouped firm i into four types, using the same category as for the firms of competing units. Overall, the Gas Large firm types most frequently appear as competitors of most type of firms. Table B5, which is a summary by peak and off-peak hours, shows that the relative share of Gas Large type competitors is lower in the peak hours than in the off-peak hours, and increasingly more of mixed type firms (both gas/oil and base/gas) become competitors during peak hours, indicating that firms are slightly more likely to face the Gas Large firm type as a competitor in the off-peak hours.

B.5.3 Moment condition and Estimation

Since the dimension of the moment vector corresponds to the number of hours (H) in which the bids of parameterized competitors appear, the total number of estimable parameters from this moment condition does not exceed H. Note that H and the number of days (T) vary across firms and samples, as they depend on the number of markets in which the parameterized competitors are observed.

	Competing Firm Type (share)			
Firm Type	Heavy Base	Gas Large	Gas/Oil Mix	Base/Gas Mix
Heavy Base	0.08	0.49	0.17	0.11
Gas Large	0.12	0.43	0.18	0.11
Oil/Gas Mix	0.12	0.48	0.14	0.11
Base/Gas Mix	0.12	0.48	0.17	0.07
Total	0.11	0.47	0.16	0.10

Table B4: Competitor Firm Type Share: Total

	Competing Firm Type (share)			
Firm Type	Heavy Base	Gas Large	Gas/Oil Mix	Base/Gas Mix
	Off-Peak Hours			
Heavy Base	0.08	0.55	0.14	0.08
Gas Large	0.12	0.48	0.14	0.08
Oil/Gas Mix	0.11	0.54	0.11	0.08
Base/Gas Mix	0.11	0.54	0.13	0.05
Total	0.11	0.53	0.13	0.08
	Peak Hours			
Heavy Base	0.08	0.46	0.19	0.12
Gas Large	0.13	0.40	0.19	0.12
Oil/Gas Mix	0.12	0.45	0.16	0.12
Base/Gas Mix	0.12	0.45	0.19	0.09
Total	0.12	0.44	0.18	0.12

Table B5: Competitor Firm Type Share: Off-Peak and Peak Hours

One caveat of the estimation is that I pool days and hours within each sample to estimate a single set of parameters per competitor. This implies that if beliefs vary across (t, h) within the sample, the estimated parameters will reflect the average of those belief distributions. As shown in Section 3.3.3, prices in the post-increase period exhibit a slight decreasing path over time, suggesting that beliefs could be adjusted. While estimating at the daily level would be ideal, doing so would sacrifice sufficient variation in competitor sets within each firm, which is critical for identification. Therefore, the parameters estimated from the pooled data can be interpreted as the average beliefs that rationalize the observed price patterns within the post-increase sample.

B.5.4 Smoothed supply bid, residual demand, and weight

The derivatives of the supply offer curve and the residual demand curve of each firm enter the empirical analogue of the first-order condition. Since these curves are submitted as step functions, I first smooth the curves using the normal kernel smoothing approach following Wolak (2007), using a bandwidth of \$3/MWh for the Sample 0 estimations and \$6/MWh for the Sample 1 estimations. As a robustness check, I tried different bandwidths to see how sensitive the derivatives are to the bandwidth selection. Results are quite robust across bandwidths except for some days when electricity prices are extremely high.

Let firm i's unit j's step k bid to be $\mathrm{b}_{ijkht} =< b_{ijkht}, q_{ijkht} >$. Suppose the market clearing price at hour h is P_{ht}. Note that $\mathcal{K}$ and κ are the CDF and pdf of a normal distribution. Then, the smoothed supply bid curve of firm i using the bandwidth bw is represented as below:

$$\widehat{Q}_{iht}\left(P_{ht}, \mathbf{b}_{iht}\right) = \sum_{j \in J_i} \sum_{k} q_{ijkht}\, \mathcal{K}\Big(\frac{P_{ht} - b_{ijkht}}{bw}\Big)$$

The smoothed residual demand curve of firm i, using bandwidth bw is shown below:

$$\widehat{RD}_{iht}\left(P_{ht}, \mathbf{b}_{-iht}\right) = D_{ht} \; - \sum_{m \neq i} \sum_{j \in J_m} \sum_{k} q_{mjkht} \, \mathcal{K}\Big(\frac{P_{ht} - b_{mjkht}}{bw}\Big)$$

Then the derivative of the residual demand curve is:

$$\frac{\partial \widehat{RD}_{iht}}{\partial P_{ht}}\left(P_{ht}, \mathbf{b}_{-iht}\right) = -\frac{1}{bw} \sum_{m \neq i} \sum_{j \in J_m} \sum_{k} q_{mjkht} \, \kappa\Big(\frac{P_{ht} - b_{mjkht}}{bw}\Big)$$

Finally, the expression of the weight, which is the probability of bid step b_{ijkht} being the marginal unit, is shown below (Wolak, 2007):

$$\frac{\partial P_{ht}}{\partial b_{ijkht}} = \frac{\partial \widehat{Q}_{iht}(P_{ht})}{\partial b_{ijkht}} \Bigg/ \Big(\frac{\partial \widehat{RD}_{iht}(P_{ht})}{\partial P_{ht}} - \frac{\partial \widehat{Q}_{iht}(P_{ht})}{\partial P_{ht}}\Big)$$

B.6 Tables and Figures

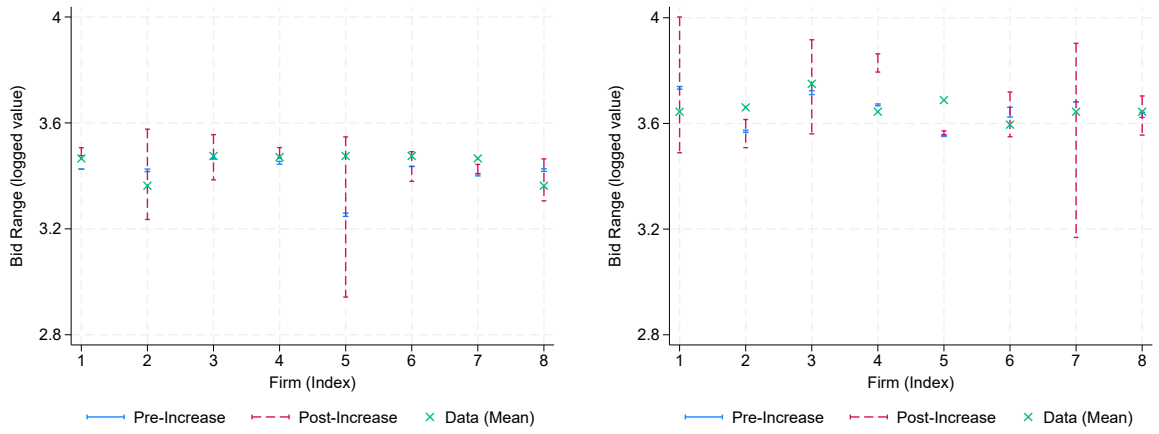
	mean	med	min	max	s.d.	N
Block's maximum k	5.05	4	1	10	2.99	19

Notes: Each block within the phase 2 have multiple days which I indexed by k. The table shows the summary of the block-level maximum k.

Table B6: Summary statistics of maximum of block-level day index, k

Adjustment	Variable	Sample	Mean	Median	S.D.	min	max	Obs
a. Sample Matched	gas price	pre-	4.44	4.49	0.26	4	5.18	173
		post-	4.43	4.50	0.29	4.02	5.16	313
c. Sample Matched + first day dropped	gas price	pre-	4.44	4.49	0.27	4	5.18	173
		post-	4.39	4.41	0.27	4.02	5.16	253
a. Sample Matched	demand	pre-	13550	14033	3446	6617	21505	173
		post-	13260	13511	2955	6529	21113	313
c. Sample Matched + first day dropped	demand	pre-	13550	14033	3446	6617	21505	173
		post	13133	13367	3012	6648	20583	253

Table B7: Summary Statistics: Robustness Checks

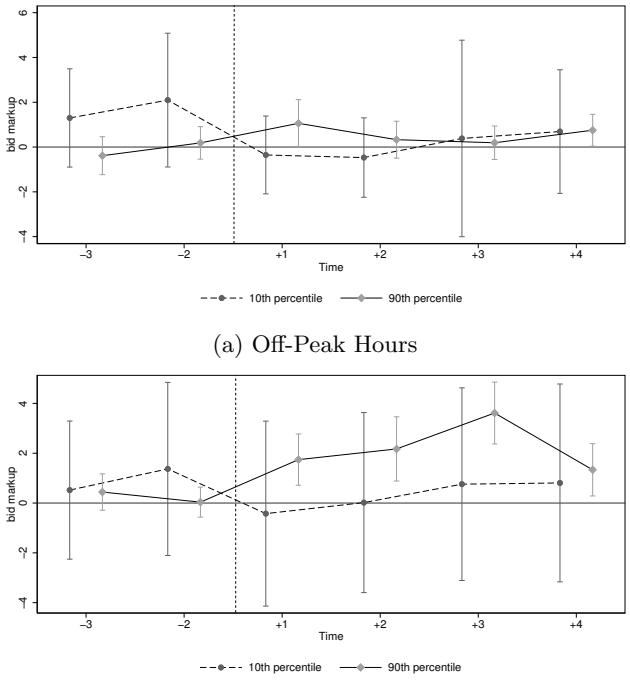


(a) Off Peak

(b) Peak

Notes: Firm Index 1-3: Gas Large firms, 4-5: Gas/Oil Mix firms, 6: Low Base and 7-8: Heavy Base firms. Firms 1,2,3 and 6 are Large firms and firms 4,5,7, and 8 are medium firms. "Data" shows the mean of the competitor's ex-post bids.

Figure B4: Bid range of a competitor with the largest change in belief dispersion: by Firm, with mean price bid



(a) Off-Peak Hours

(b) Peak Hours

Notes: The graph shows the estimated bid markup difference ($\hat{\beta}_{1,k}$) for 10th and 90th percentiles for each period (*Time*) $k = \{-3, -2, +1, +2, +3, +4\}$, relative to the baseline comparison period $k = -1$. Time $k < 0$ is part of the pre-increase sample, and $k > 0$ is part of the post-increase sample. Panel (a) shows the off-peak hours, and Panel (b) shows peak hours. Confidence intervals are computed using standard errors clustered at generating unit level (5% significance level).

Figure B5: Bid markup dynamics: Off-peak and Peak hours

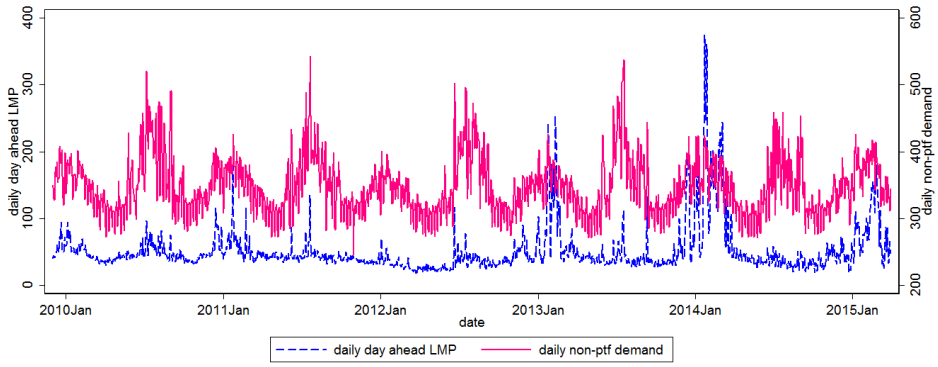
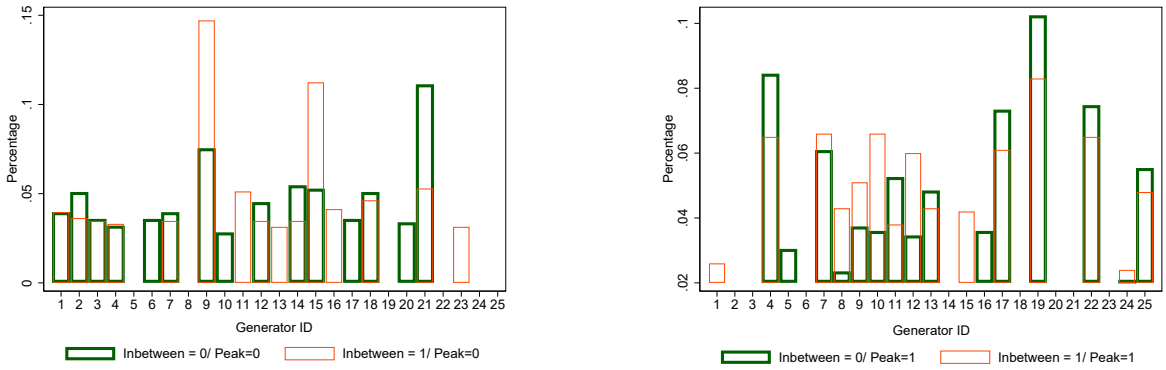


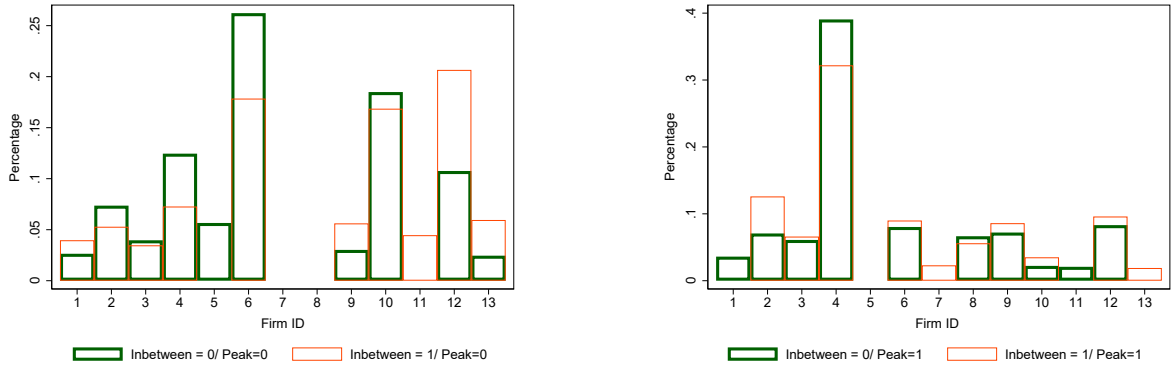
Figure B6: Day-ahead Electricity Demand and Day-ahead Electricity Prices



(a) Offpeak Hours (b) Peak Hours

Notes: The graph shows the percentage of each unit's occurrences as marginal power plant, taken separately by Pre-increase and Post-increase samples and Peak and Off-peak samples. Generation ID indicates an arbitrary index assigned to each generating unit. Figure B7 was generated by assigning unique IDs to ex-post marginal units and plotting their percentage occurrences as marginal units in the pre- and post-increase samples. These graphs were created separately of off-peak and peak hours, and in both cases, we observe shifts in the composition of marginal price setters. A similar analysis at the firm level, shown in Figure B8 in the Online Appendix, indicates composition changes, though they are less pronounced than at the unit level.

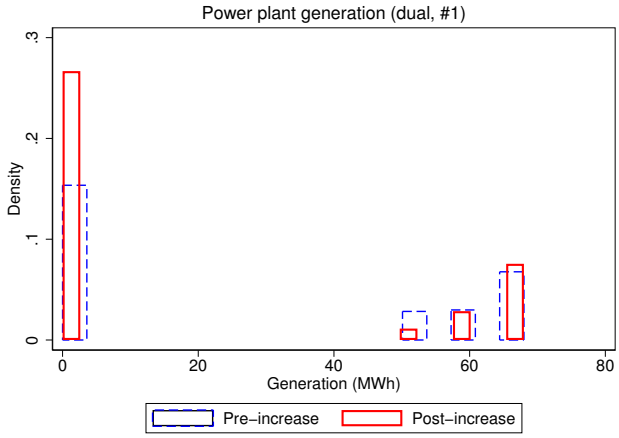
Figure B7: Marginal Units: distribution by Pre- and Post-Increase and Peak/Offpeak



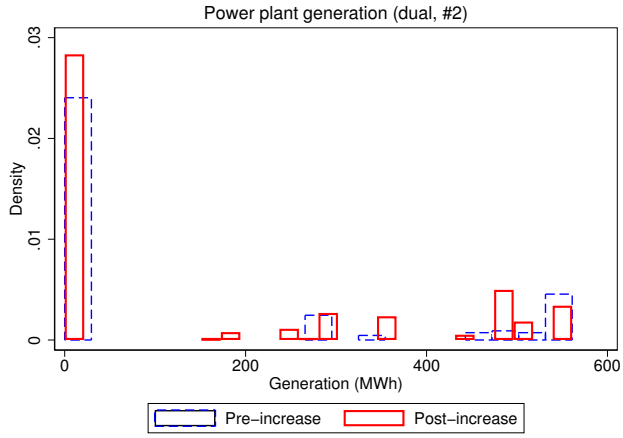
(a) Offpeak Hours (b) Peak Hours

Notes: The graph shows the percentage of each firm's occurrences as being a marginal price setter (by owning the marginal power plant), taken separately by Pre- and Post-increase samples and Peak and Off-peak samples. Firm ID indicates an arbitrary index assigned to each firm.

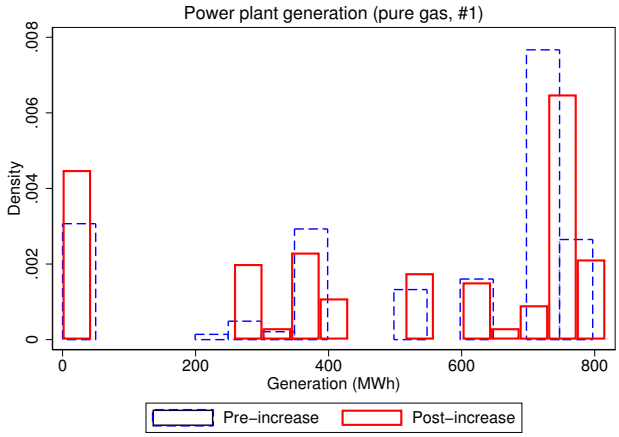
Figure B8: Marginal Firms: distribution by Pre- and Post-Increase and Peak/Off-peak



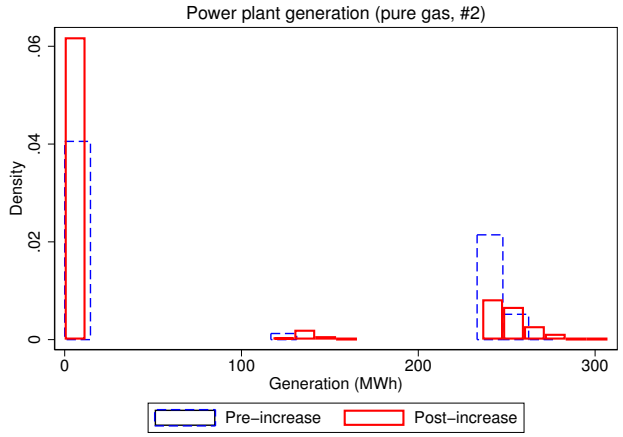
(a) Dual gas unit #1



(b) Dual gas unit #2

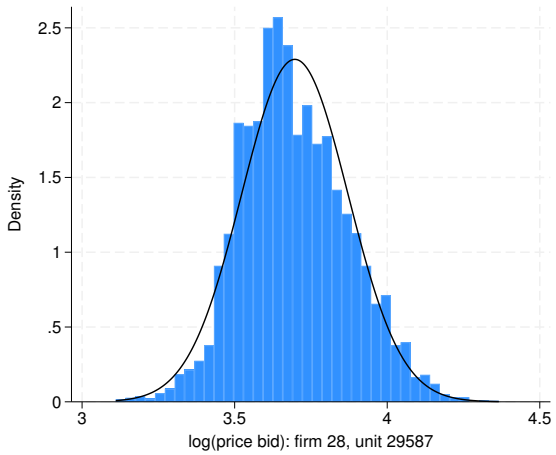


(c) Gas unit #1

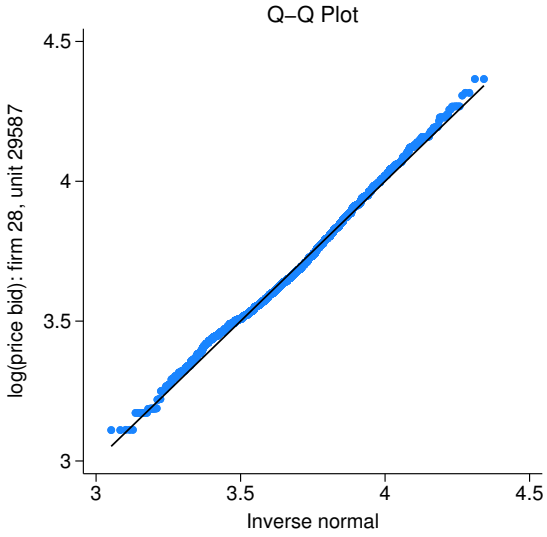


(d) Gas unit #2

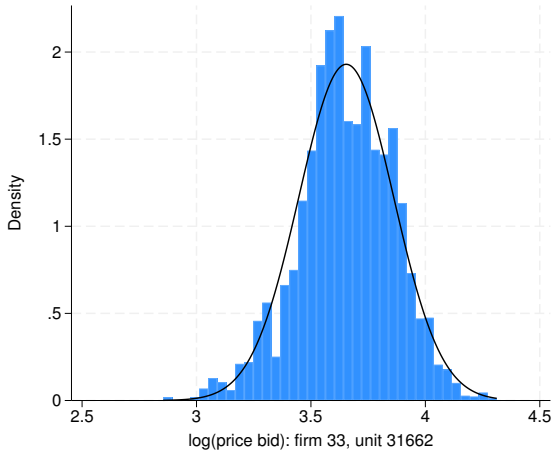
Figure B9: Histogram of generator-level quantity: pre-increase vs. post-increase



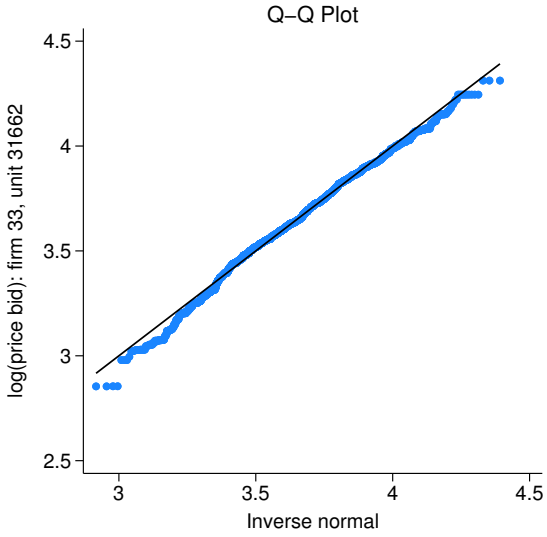
(a) Firm 28: distribution



(b) Firm 28: Q-Q plot



(c) Firm 33: distribution



(d) Firm 33: Q-Q plot

Figure B10: Log transformed Price Bids: single generating unit's of firm 28 and 33