Evaluating the Role of Information Disclosure on Bidding Behavior in Wholesale Electricity Markets

by

David P. Brown[†], Daniel O. Cajueiro[‡], Andrew Eckert[*], and Douglas Silveira[**]

Abstract

Real-time information has the potential to improve market outcomes in wholesale electricity markets. However, transparency can also facilitate coordination between firms, raising questions over the appropriate extent of information disclosure. Despite this ongoing debate, there is a lack of understanding of the information employed by firms when bidding in wholesale electricity markets. We use data from Alberta's wholesale market and leverage machine learning techniques to evaluate the real-time information firms use when forming their bidding decisions. We find that aggregate market-level variables emerge as important predictors, while detailed firm-specific information does not lead to a material improvement in predicting firms' bidding decisions. These results suggest that firm-specific information, which has raised concerns because of its potential use in facilitating coordinated behavior, may not be required to promote efficient market outcomes.

Keywords: Machine Learning, Electricity, Price Forecasting, Competition Policy
JEL Codes: D43, L13, L50, L94, Q40

† Department of Economics, University of Alberta, Edmonton, Alberta, Canada (dpbrown@ualberta.ca).

‡ University of Brasilia, Department of Economics, Brasilia, Federal District, Brazil (danielcajueiro@unb.br).

∗ Department of Economics, University of Alberta, Edmonton, Alberta, Canada (aeckert@ualberta.ca).

∗∗ Department of Economics, University of Alberta, Edmonton, Alberta, Canada
Federal University of Juiz de Fora, Department of Economics, Juiz de Fora, Minas Gerais, Brazil (dsilveir@ualberta.ca, douglas.silveira@ufjf.br).

Acknowledgements. This research project received support from the Government of Canada's *Canada First Research Excellence Fund* under the Future Energy Systems Research Initiative (Brown, Eckert) and the Social Sciences and Humanities Research Council's Canada Research Chair program (Brown). Daniel O. Cajueiro thanks CNPQ (Grant 302629/2019-0) and FAPDF (Grant 00193.00001796/2022-85) for financial support.

1 Introduction

Supply and demand must be balanced at every moment to ensure a reliable and stable supply of electricity. Historically, this was achieved by calling upon fossil-fueled power plants. Increasing concerns over the environmental implications of greenhouse gas emissions has motivated the enactment of policies to decarbonize the electricity sector. This has led to the rapid deployment of renewable resources (IEA 2022). However, this has also heightened concerns over market uncertainty due to the unique characteristics of renewable resources.[1]

In a setting with uncertain market conditions, the release of near real-time information may allow firms to adjust their behavior to the changing market environment (Kuhn and Vives 1994). Consequently, information disclosure can have an important impact on both competition and market efficiency in the electricity sector. On one hand, the existing literature has highlighted that increased transparency and public information can increase competition under certain conditions (Holmberg and Wolak 2018). In contrast, concerns have been raised that the availability of a high degree of granular information released in real-time can facilitate coordinated behavior (von der Fehr 2013; Brown and Eckert 2022). With these countervailing impacts of information disclosure, it is critical to balance information revelation to promote market efficiency while avoiding disclosing too much information that could facilitate coordination.

A challenge in evaluating arguments to increase information or limit its availability to facilitate collusion is a lack of a clear understanding of the information employed by firms. In this paper, we use data from Alberta's wholesale electricity market to fill this gap. In this market, firms submit price-quantity bids in an hourly uniform-price multi-unit auction. Generators can adjust their offers up to two hours before the hour in which the market clears. Firms have access to a wide array of information, including the real-time market price, current generation of every individual asset, import/export intertie availability, province-level demand, and forecasted pool prices, among numerous other factors.

In the presence of a vast quantity of information available to firms when they make their decisions, an important question is: what information do they use? For example, do firms rely primarily on real-time information about aggregate market conditions such as realized and forecasted demand and prices, or do they leverage detailed firm-level information such as the availability of their rivals' specific assets? If firms rely primarily on market-level information,

1. For a detailed analysis of renewable resource intermittency and uncertainty, see Weber and Woerman (2022).

then this raises questions as to whether firms need disaggregated real-time information of specific rivals which could potentially be used to facilitate coordinated outcomes.

We leverage machine learning techniques to parse out the observable pieces of information that best predict how firms make their bidding decisions. Machine learning is ideally suited for this setting, where many factors could be used in complex and varying ways to inform firms' bidding decisions over time. We use detailed data from Alberta's wholesale electricity market for the year 2021 to predict the pricing decisions of large strategic firms and evaluate the most useful factors in this prediction. We focus on the bids of coal and natural gas generators that are most likely to be used to bid strategically.[2] We break our analysis into two specifications. First, we only include aggregate market-level variables observable in near real-time, such as the realized and forecasted market demand and price, available capacity by fuel type, and a firm's private information such as its lagged bids. Second, we disaggregate a number of factors into firm-specific variables such as firm-level capacity availability and observed output, unit-level outages, and firm-specific ancillary service (AS) output.

We are interested in addressing the following questions. Do we observe a sizable increase in the prediction of firms' pricing decisions when we allow machine learning algorithms to consider rival firm-specific information? If the price prediction quality increases, this suggests that firms actively use this disaggregated real-time information to make their bidding decisions. If not, this suggests that the market-level information is primarily used to facilitate firms' responses to a changing market environment.

We use both individual and grouped Permutation Importance (PI) techniques to better understand the key factors used in forming a firm's pricing decisions (Altmann et al. 2010; Plagwitz et al. 2022). When forecasts are made based on several input variables of different sources, the PI quantifies the variables – individually or grouped – that best contribute to the machine-learning algorithm's ability to form accurate forecasts by randomly shuffling the variable's values and observing the resulting impact on the quality of predictions. The use of grouped permutation importance in prediction policy problems, especially in time-series forecasting with a large set of explanatory variables, is not fully explored. The standard approach involves adding individual variable importance scores within the same group (Medeiros et al. 2021; Albuquerque, Cajueiro, and Rossi 2022; Masini, Medeiros, and Mendes 2023).[3] Although this procedure provides insights into the explanatory power of individual

2. The remaining units in the market largely consist of wind and solar units whose supply is exogenous and gas-fired cogeneration facilities associated with large industrial firms where electricity is a by-product of a behind-the-meter industrial process. These units systematically bid at a price of $0/MWh.

3. For linear models, these studies utilize the sum of the absolute values of the predictors' coefficients to

variables, it may result in misleading interpretations when these variables are highly correlated. Grouped permutation importance addresses this issue by simultaneously randomizing all variables within the same group, providing valuable insights and interpretation (Molnar et al. 2020; Au et al. 2022).

In our analysis, no single algorithm emerges as superior at predicting price changes, with both random forest and lasso yielding the best performance, depending on firm, specification, and performance measure. Regardless of the precise machine learning algorithm used, we find no appreciable increase in the performance of the price prediction when we include rival firm-specific information. Instead, it tends to introduce additional noise and variance into the predictions. This suggests that firms rely primarily on market-level data to determine their pricing dynamics rather than leveraging firm-specific information.

Our grouped permutation importance analysis with the random forest algorithm indicates the most important market-level predictors of price changes are wholesale market demand and price variables, supply and generation unit availability by fuel type, and lagged changes in the firm's own price variable. The first two capture core metrics that determine the level and nature of available supply and demand, while the latter indicates persistence in firms' own pricing decisions. The individual permutation importance analysis reinforces these findings and demonstrates that factors such as forecasted and observed market demand are among the most valuable individual variables in forming price predictions. Hour-of-day also plays a role in both the individual and grouped permutation importance analyses. This is consistent with the observation that firms systematically adjust their bids at certain times of the day (e.g., in peak and off-peak hours). Finally, for both permutation importance approaches, the inclusion of firm-specific rival variables has no material impact on the key variables used in the price predictions by the algorithms.[4]

Our analysis contributes to ongoing policy debates over information disclosure. Our results suggest placing a lower priority on the release of real-time firm-specific information (such as available unit-specific capacity or generation), as such dis-aggregated information does not serve as an important determinant of firms' bidding decisions. Further, this type of information may be used to facilitate coordinated behavior leading to elevated market prices. Our analysis indicates that key market-level factors related to forecasted and real-

measure their relative importance.

4. The grouped and individual permutation importance analysis for the lasso algorithm place more importance on a firm's own lagged pricing variables and hourly indicates, and less on market-level generation availability by fuel type and market demand variables. Despite these differences, we continue to find that adding rival firm-specific information does not result in a material improvement in the price prediction.

ized demand, as well as market-level generation availability by fuel type, are key pieces of information used by the firms in our sample. With the goal of taking a balanced approach to information and data transparency in an uncertain environment, our results suggest that regulators should focus on providing real-time market-level information to participants.

The paper proceeds as follows. Section 2 provides a summary of our contributions to the literature. We describe the key features of Alberta's wholesale market in Section 3. Section 4 summarizes the data and variables used in our analysis. The empirical methods are presented in Section 5. Section 6 presents the results of our empirical analysis. Conclusions are discussed in Section 7.

2 Literature Review

Our paper contributes to several distinct strands of the literature. There exists a literature on the potential benefits of increased information transparency in electricity markets. Holmberg and Wolak (2018) develop a multi-unit auction model in which firms have asymmetric information about costs and argue that increased transparency of cost data, financial market outcomes, and historical bid and production data is likely to increase the competitiveness of wholesale electricity markets.

Empirically, Bergheimer, Cantillon, and Reguant (2023) consider the context of the New Zealand wholesale electricity market, where each half-hourly spot market is preceded by a 36-hour pre-dispatch market period, during which firms can update their bids. Market-clearing prices and quantities are established and reported every two hours for the upcoming half-hour. The authors demonstrate that pre-dispatch prices and quantities are informative regarding ultimate outcomes. They argue that pre-dispatch outcomes provide valuable information to participants, allowing generators to schedule their assets more effectively, leading to efficiency gains. In another recent example, Bunn and Kermer (2021) develop an analytical model of an electricity balancing market and calibrate the model to the Austrian context to examine arbitrage behavior and the role of time delays and information transparency. The authors conclude that the more timely release of information on system imbalance forecasts and wind and solar generation can reduce system costs.

In contrast, it has been argued in von der Fehr (2013) that increased information transparency can facilitate coordination among firms. The high-frequency release of firm-level information (such as bids), for example, could allow firms to quickly detect deviations of rivals from cooperative behavior. Such granular firm-specific high-frequency information could allow firms to communicate through patterns or signals in their bids or other choice variables.

The latter concern was the focus of policy attention in Alberta. In August 2013, the Alberta Market Surveillance Administrator (MSA) released a report alleging that certain firms were using patterns in their bids into the wholesale market, which operates as an hourly uniform price auction (MSA 2013). At that time, the complete set of bids by suppliers for the hour was made public at the end of the hour, but with the identity of the associated firm and generating asset removed, in a publication called the Historical Trading Report (HTR). The MSA argued that by 'tagging' their bids, firms could reveal themselves to rivals as the source of certain bids, and indicate their intent to maintain high prices. These allegations led to a hearing of the Alberta Utilities Commission, which ordered an end to the public release of the HTR (AUC 2017). Patterns in bids in the Alberta market over the time period of the allegations are documented in Brown, Eckert, and Lin (2018). Brown et al. (2023) employ machine learning methods to evaluate whether firms could identify their rivals through de-identified bids. The authors find that before the release of the MSA's August 2013 report, the firm associated with a bid could be predicted with a high degree of accuracy; after the report, accuracy decreased considerably. Brown and Eckert (2022) show that firms employing patterns could have increased profits through unilateral deviations and that for one firm, the potential for profitable deviations was greater on days when tagging was observed.

The use of patterns in prices or other public information to facilitate possible coordination has been documented in other industries. In the case of airlines, see Borenstein (1998) regarding airline fare codes; Aryal, Ciliberto, and Leyden (2022) consider whether airline managers employed keywords in their quarterly earnings calls to signal to rivals a willingness to reduce the number of seats being offered. In the case of retail gasoline, Lewis (2015) finds in U.S. markets that price endings of 5 and 9 are associated with higher and more stable prices, and suggests that such price endings may be used to establish focal prices. Price leadership and signaling through station-level gasoline prices released publicly through Western Australia's Fuelwatch price transparency program is discussed in Byrne and Roos (2019). In other settings, see Christie and Schultz (1999) for the case of odd-eight quotes in Nasdaq stocks, and Abrantes-Metz, Villas-Boas, and Judge (2011) for a discussion of the distribution of the second digit of the Libor rate.

The current analysis contributes to both strands of the literature that emphasize the countervailing effects of information disclosure (i.e., efficiency versus coordination). Focusing on the Alberta electricity market over the year 2021, we provide empirical evidence on aggregate market-level information that assists firms in making their bidding decisions while evaluating the extent to which firms use rival firm-level disaggregated information that may facilitate coordination. Despite the end of the HTR, bidders continue to have access to large amounts of high-frequency data, including real-time generation and available capacity for individual generation units; hence, although the potential for signaling has been reduced, the use of public information to facilitate coordination, for example by allowing the quick detection of deviations, remains.[5] To the best of our knowledge, we provide the first empirical evaluation of the types of information firms use when making their bidding decisions in the electricity sector. This work has the potential to help inform regulations on how to strike the right balance on information disclosure.

Finally, our paper contributes to the large literature on forecasting electricity prices; see Murthy et al. (2014) and Sapio (2021) for reviews of time series and econometric approaches, Jedrzejewski et al. (2022) for an overview of machine learning approaches, and Sai et al. (2023) for a recent example applying machine learning techniques.[6] An important distinction between the existing literature and the current paper is that a key challenge faced by the existing literature is forecasting prices in a setting where stable market-level electricity prices are interrupted by large price spikes. These challenges are less relevant in our analysis, which focuses on the prices at which firms offer specific assets into the market, rather than forecasting the equilibrium market price.

3 Alberta's Wholesale Electricity Market

Alberta's wholesale electricity market operates as an hourly uniform-priced multi-unit auction. Firms offer up to seven price-quantity blocks for each generating asset for each hour. Offers can be adjusted up to two hours in advance of the hour. Prices are restricted to lie between \$0 and \$999.99. The Alberta Electric System Operator (AESO) calls upon (dispatches) generation units throughout the hour to meet demand, in increasing order of offer price. The price of the last block dispatched is the System Marginal Price (SMP); the time-weighted average SMP over the hour is the pool price for the hour. Generators are paid the pool price each hour for the energy they generate.

Alberta has an "energy-only" market design, meaning that generators earn revenues only

5. There is a large literature discussing the role of information in detecting deviations from collusive strategies and the implications this has on facilitating collusive outcomes (e.g. see Green and Porter (1984) and Harrington and Skrzypacz (2007)).

6. Machine learning and statistical methods are reviewed and compared in Lago et al. (2021).

from supplying electricity.[7] As a result, firms are permitted to exercise unilateral market power through economic withholding (i.e., submitting bids in excess of their short-run marginal cost).[8] Alberta's wholesale market clears with a single price; there are no nodal or zonal pricing mechanisms. Finally, there is no day-ahead market, only an hourly spot market.

During our sample period, Alberta's wholesale market was moderately concentrated; the five largest firms accounted for 70% of the generation capacity in the province as of January 2021. At the start of 2021, 32% of Alberta's generation capacity was coal-fired, 44% was natural gas fueled (including cogeneration facilities), while wind farms accounted for 11% of the provincial capacity. The remaining capacity was distributed across solar, hydro, biofuel, and storage facilities (MSA 2021).[9] Alberta has a small amount of imports and exports that arise through interties to British Columbia, Montana, and Saskatchewan.

A large amount of information about the market is available to participants in real or near-real time. Throughout the hour, the AESO posts numerous market-level factors, including the current market demand, measures of supply adequacy reflecting the amount of unused generation capacity available, imports/exports to each neighboring jurisdiction, generation and available capacity by fuel type, and the system marginal price. The AESO also posts forecasted pool prices for the next two hours and the realized pool price immediately after the hour closes. Forecasted demand for the entire day and the first six hours of the following day are posted at midnight.

The AESO also publishes a certain degree of firm-level information in real-time, including the current generation from each generating unit. Further, output providing certain AS products referred to as "contingency reserves" are reported in real-time at the unit level. In addition, generation units going on outage and returning online are published through the AESO's "event log". While firms have a sizable amount of information on the generation output from specific units, during our sample period hourly price-quantity bids of individual generating units were not released in real-time but were made available with a 60-day lag in the AESO's Merit Order Snapshot.

7. In contrast, a number of other restructured markets have capacity payment mechanisms that provide supplementary revenues based on each generation unit's capacity value (Holmberg and Tangerås 2023).

8. The proposed logic of permitting certain types of unilateral market power in the wholesale market is to facilitate long-run investment in generation capacity via fixed cost recovery. Entry is expected to discipline the market to avoid excessive rents in the long-run. For a detailed discussion of Alberta's wholesale market design, see Brown, Eckert, and Shaffer (2023).

9. Over 2021, several coal generators were converted to gas generators. By April 2022, the percentage of coal-fueled capacity had fallen to 8%, while the percentage of natural gas-based (non-co-generation) capacity had increased from 12% to 29%.

While firms are required to submit their wholesale market bids the day before market clearing, firms can adjust their price-quantity bids up to 2 hours before the market clears. This allows firms to incorporate information released by the AESO in near real-time. To illustrate the precise timing of information disclosure and how it can be incorporated into firms' bids, consider the following example of offers submitted by firms for the hour running from 8:00 - 9:00 AM (denoted as Hour Ending 9, HE9). Because bids can only be changed up to 2 hours before the hour, the latest that bids for HE9 can be changed is 6:00 AM (the start of HE7). This implies that only real-time information released before 6:00 AM (HE7) can be incorporated into bids for HE9. More generally, any information revealed on or before HE $t-3$ can be incorporated into the bids for HE t. The timing in which firms can incorporate information revealed in real-time will play an important role in our empirical methodology detailed below.

4 Data

4.1 Data Sources

Our analysis uses publicly available data for the year 2021 from the AESO. We employ data from the hourly Energy Merit Order, which indicates all price-quantity bids in each hour for each generation unit. These data are released publicly with a 60-day lag, meaning that they are not available in real-time to the firms; however, these data can be used to construct output and capacity availability measures by fuel type and the unit, which are observable to firms in real-time. We use hourly metered volumes data to estimate wind and solar output. Hourly ancillary service market data are collected to provide information on total and unit-level AS quantities. The AESO also provides data on hourly realized and forecasted market-level demand and prices, as well as export and import intertie capacity limits and supply levels. We collect temperature data from Environment and Climate Change Canada for the two major cities in the province, Edmonton and Calgary.

Table 1 provides the summary statistics for a number of key market-level outcome variables. We observe a mean pool price of \$101.95/MWh with a high standard deviation.[10] Looking at the price distribution, we can see that the pool price distribution is rightward skewed. Wind and solar output reflect a modest share of total market demand. Similarly, Alberta is a net importer, with imports exceeding exports. Import and export quantities

10. Throughout our analysis, all prices are in Canadian dollars.

are small relative to market demand, demonstrating that the majority of demand is met by within-province generation.

Table 1: Summary Statistics – Key Market-Level Variables

	Mean	Std. Dev.	25th percentile	50th percentile	75th percentile
Pool Price ($/MWh)	101.95	138.98	43.15	57.34	82.09
Market demand (MWh)	9,727.72	792.11	9,165.00	9,654.00	10,342.00
Wind Supply (MWh)	700.01	508.41	236.69	613.26	1,152.28
Solar Supply (MWh)	260.64	78.89	165.00	289.00	325.00
Import Supply (MWh)	473.88	244.37	265.00	460.00	693.00
Export Supply (MWh)	15.02	56.14	0.00	0.00	0.00

4.2 Price Variable

Our objective is to understand the factors that predict individual firms' short-term pricing decisions and evaluate whether disaggregated firm-specific factors improve these predictions beyond just using market-level variables. We focus on the four large strategic firms operating in Alberta's wholesale market: Capital Power, ENMAX, Heartland, and TransAlta, who in January 2021 accounted for 13%, 9%, 15%, and 25% of Alberta's capacity, respectively.[11] For each of these firms, their bidding decisions are driven by cost-based factors and strategic decisions regarding market power (i.e., to bid in excess of their marginal costs). The remaining firms in the market consist of over 40 smaller ("fringe") producers, who typically offer into the market at short-run marginal cost.

Within the larger firms, while some types of generating assets are sometimes offered above short-run marginal cost, others, such as renewable capacity and cogeneration, appear not to be used for strategic purposes, and have offer prices exhibiting little variation. Therefore, for each of the large strategic firms, we identify generating assets that offer bids in excess of short-run marginal cost, and that are potentially economically withheld in the exercise of market power.[12] These plants include coal generators (some of which converted to natural gas over the sample period), and natural gas plants. For each hour and large firm, we compute the quantity-weighted offer price on bids above marginal cost for these units. This approach allows us to focus on bids driven in part by strategic factors.

11. As noted in Section 3, while there are five "large" firms in Alberta's wholesale market, one (Suncor) operates only cogeneration units, and as a result can be viewed as a price-taking firm.

12. Details on marginal cost estimation, and a list of assets used to compute the price variable of each firm, can be found in Appendix C.

Table 2: Summary Statistics of Volume-Weighted Prices ($/MWh) - Strategic Assets

	Mean	**Std. Dev.**	**25th percentile**	**50th percentile**	**75th percentile**
Capital Power	396.06	168.06	291.39	388.68	504.95
ENMAX	163.30	101.47	73.79	152.91	228.77
Heartland	265.01	138.46	152.04	237.08	375.32
TransAlta	266.76	198.25	96.10	201.86	416.00

Table 2 provides summary statistics for the firm-specific quantity-weighted price variables on these strategic units that will serve as the dependent variable in our subsequent analysis. We can observe considerable dispersion in the price levels within and across firms. These prices far exceed estimates of the firm-level short-run marginal costs of these assets which systematically lie below $100/MWh.

Because our focus is on price adjustments in response to high-frequency real-time information, we are interested in predicting changes in prices. One empirical challenge that we face is that we observe relatively "sticky" pricing decisions at the firm level. That is, firms often only adjust their prices at periodic intervals throughout the day. To reduce the number of zero first differences in our price variable, we reduce the frequency of data to every three hours.[13] Specifically, we use evenly spaced data from hours ending 1, 4, 7, 10, 13, 16, 19, and 22.[14] This allows us to keep a time-series structure that captures the strategic pricing dynamics of each firm over the day. This time frame encompasses relevant changes over the on-peak and off-peak hours and provides eight evenly-spaced observations within each day.

We define our dependent (target) variable as the first difference in the logarithm of prices for every three hours. The transformation of price levels into logarithms and the computation of the first difference is a standard procedure to remove non-stationarity in time-series forecasting. See Hamilton (2020) for details.

4.3 Market and Firm-Level Predictors

A large array of market and firm-level factors are observable by firms in real-time. A primary objective of our analysis is to evaluate the extent to which firms supplement market-level variables with firm-specific variables when making their bidding decisions. To address this

13. In time-series modeling, a large number of zeros can lead to data sparsity, which often results in biased parameter estimates, decreased model sensitivity, and difficulties in adequately fitting the model, either through overfitting to zeros or underfitting impactful non-zero trends. See Syntetos and Boylan (2005), Perumean-Chaney et al. (2013), and Petropoulos et al. (2022) for detailed discussions.

14. Alberta's system operator denotes hours by the time at which the hour ended. For example, Hour Ending 1 (or HE1) refers to the hour from 12:00 - 1:00 AM.

question, we consider two specifications. In the first specification, we only include market-level variables that are observed in real-time. In a second specification, we decompose several key categories into firm-level factors to determine if our prediction of a firm's bidding decisions improves materially. To simplify the discussion, we start with the variables used in Specification I and group these variables into eight categories. Tables A.1 and A.2 in Appendix A provide a detailed list of the predictors used in specifications I and II.

First, we consider a firm's own lagged price variables. We often observe that firms adjust their bids in certain hours and hold prices at those levels for several hours. Consequently, it is important to control for lagged values of the dependent variable to capture that a firm's current pricing decisions are a function of how they bid in previous periods.

Second, a firm may adjust its bids depending on the capacity and supply available from its generation units. To control for this flexibly, we include lagged values of a firm's generation output and capacity availability in aggregate (across all units) and separated by generation technology. In addition, we include a firm's own Ancillary Services output to capture how much supply the firm has committed in other markets.

Third, we include a range of market-level generation supply and capacity availability variables. These variables capture observed output (in MWhs) and capacity availability (in MWs) measures broken down by coal, dual-fuel (gas/coal) units, combined cycle (CC) gas, simple cycle gas cogeneration, biomass/other, hydroelectric, wind, solar, and battery storage. We also include a market-level generation unit outage (in MWs) measure. These variables aim to determine if firms adjust their bids based on observed variations in market-level generation factors (e.g., large changes in natural gas CC unit availability), suggesting that only reporting real-time technology-specific measures is sufficient to inform bidding decisions.

Fourth, we include variables to capture wholesale market demand, the hourly pool price, and price variation captured by the distribution of the SMPs throughout the hour. More specifically, we include realized and forecasted market demand and pool price measures to capture backward-looking and forward-looking values. Market demand is expected to have a central impact on wholesale market outcomes, and pool prices directly capture realized and anticipated equilibrium outcomes. We compute the standard deviation and coefficient of variation of the SMPs throughout the hour. These measures allow us to capture the slope of the industry supply curve in the neighborhood of market demand. For example, if the SMP has limited variation over the hour, then we are more likely to be at a flat portion of the supply curve. This could indicate a limited scope for market power. If the SMP has considerable variation, then market-clearing is at a steep vertical section on the supply curve. Such information has the potential to impact bidding behavior.

Fifth, we include export and import quantities and intertie capacity availability measures for each neighboring jurisdiction. Sixth and seventh, we include hour dummies and temperature variables for Edmonton and Calgary, the two largest cities in the province. These variables aim to capture systematic time and temperature-based variation that could impact bidding decisions. Finally, for our eighth category, we include total AS market output to capture the quantity of reserves procured by the AESO.[15]

In Specification II, we adjust two of the eight categories of variables to include detailed firm-specific values that firms could readily observe in the information published by the AESO. For this analysis, the third category summarized above now includes rival firm-specific supply and capacity availability measures. These variables are provided separately for each other large strategic firm, and the remaining firms are aggregated as the "fringe". We include supply and available capacity measures in aggregate by firm and broken down by fuel type for each rival firm. This increases the dimensionality of information considerably but allows us to evaluate if firms are looking at the availability of capacity and supply (output) from specific technologies of certain firms when making their bidding decisions.

We also adjust the eight categories to provide information on the firm-specific AS market output. This captures the fact that firms could determine which specific rivals are supplying output in the AS market and foregoing the ability to supply this output to the wholesale market.

As described in Section 3, we need to account for the fact that firms can only incorporate information revealed on or before hour ending $t-3$ due to the rules on bid adjustments running up to market-clearing. In addition, as noted in Section 4.2, we have eight evenly-spaced observations within each day, split into 3-hour intervals. Consequently, we include data on the predictor variables summarized above from hours $t-3$, $t-6$, and $t-9$ in our analysis to consider how information released in near real-time impacts firms' strategic bidding decisions in hour t. Two exceptions apply. We include additional lags on the firm's own lagged price variable to provide a more flexible relationship between a firm's prior pricing decisions and current bidding behavior, with lags at 6, 9, 12, 15, 18, 21, and 24 hours. In addition, for the forecasted market demand variable, we use demand forecasts for HE t, $t+3$, and $t+6$ to account for the fact that firms have this forward-looking information available

15. Supplying output to certain ancillary services products limits generators' abilities to provide this output in the wholesale market. Brown, Eckert, and Silveira (2023) develop a model to demonstrate the strategic linkages between the provision of output in AS and wholesale markets.

when they make their bidding decisions for HE t.[16]

We apply a data standardization procedure in the set of predictors. Generally, machine learning algorithms benefit from this approach as it can improve their predictive power.[17] To reach that aim, we first apply a logarithm transformation to all the numerical variables in our set of predictors and take the first difference to achieve a stationary distribution.[18,19] We follow Gelman and Pardoe (2007) and standardize all the numerical variables by dividing them by two times their standard deviation. In data settings with both numerical and dummy variables, this procedure allows all the variables to have a similar scale and magnitude and, therefore, has the potential to improve the performance of our machine-learning models.

5 Empirical Methods

In this section, we present the models that will be used to predict the quantity-weighted bids for the four strategic firms in Alberta. We generate the results of our forecasting exercises using a rolling window that includes six months of training data with eight evenly-spaced observations within each day, split into 3-hour intervals. In other words, we begin by training the model using data from January to June 2021 in order to make predictions for t+3. Illustratively, with the information available in HE22 of the last day of June, we first make predictions for HE1 of the first day of July, then advance three hours, retrain the model, and forecast for HE4. We continue this step-by-step process, and as a result, our findings are based on predictions made from July to December 2021.

Section 5.1 begins with a simple random walk model to serve as a benchmark for our more sophisticated machine learning models. Section 5.2 presents the ridge and lasso regularized linear models, which are techniques used to help control the complexity of the forecasting model and enhance its predictive performance, especially in situations where there are potentially many predictor variables with varying levels of importance. Section 5.3 introduces two non-linear regression-tree-based machine learning methods – random forest and XGBoost.

16. While the AESO provides a day-ahead pool price forecast, this information is only released for hour $t+1$ and $t+2$. Firms are unable to build in this information into their bids for hour $t+3$. This differs from the demand forecasts which are posted publicly for the entire day. As a result, unlike the demand forecast, we do not include forward-looking measures on the day-ahead pool price forecast.

17. In particular, regularization machine learning algorithms, such as lasso and ridge, require data standardization procedures to ensure their correct functioning. See Hastie et al. (2009) for details.

18. The exception is for our hourly dummy variables, which are included as indicator variables.

19. It is worth mentioning that the majority of the variables are non-stationary. Overall, the Augmented Dickey-Fuller unit root test indicates that 81% (75%) of the variables are non-stationary at a significance level of 0.01 (0.05).

These methods are known for their ability to provide a mix of accuracy, flexibility, interpretability, and robustness. Section 5.5 presents the Permutation Importance (PI) technique we use for assessing the most informative explanatory variables. Section 5.4 introduces the evaluation metrics used to assess the predictive power of our machine learning models.

5.1 Random Walk

The random walk model forecasts all horizons as the last value observed for the variable as in $\hat{y}_{t+h} = y_t$, where h is the forecast horizon. This model is commonly used as a benchmark in time series forecasting analyses due to its simplicity and intuitive nature. In our approach, this naive model provides a baseline for comparison against the machine learning models we apply, allowing us to evaluate the effectiveness of the algorithms in capturing the nuanced patterns and dynamics present in the data. By contrasting the predictions of our machine learning models with those of the random walk, we can assess the extent to which machine learning methods enhance forecasting accuracy and offer valuable insights beyond the simplistic (and naive) random walk approach.

5.2 Regularized Linear Models

Time-series forecasting of the target variable y_{t+h} with linear models have the following functional form:

$$y_{t+h} = \beta_0 + \mathbf{X}_t\hat{\beta} + \varepsilon_{t+h}, \tag{1}$$

where β_0 is the intercept and $\hat{\beta}$ is the vector containing the estimates of the model coefficients. The matrix $\mathbf{X}_t$ represents the set of predictors, and its dimension is given by $T \times n$, where T is the number of observations and n is the number of predictors, respectively. The error term is given by ε_{t+h}.

Typically, linear regression models depend on OLS estimates of $\hat{\beta}$ to minimize the sum of squared residuals:

$$\hat{\beta} = \arg\min_{\beta} \left[\sum_{t=1}^{T} (y_{t+h} - \beta_0 - \mathbf{X}_t\beta)^2 \right]. \tag{2}$$

When the explanatory variables are highly correlated, the least-squares estimate may become sensitive to random errors, producing a large variance and affecting the accuracy and interpretation of the coefficient estimates.

In addition, another source of poor prediction accuracy is its low bias combined with high variance. Balancing bias and variance is key to improving the overall prediction accuracy. In many cases, it may be necessary to tolerate bias to lower variance and increase the predictive performance (Hastie et al. 2009).

Intending to overcome these drawbacks, regularized linear models add penalty terms to eq.(2) as follows:

$$\hat{\beta} = \arg\min_{\beta} \left[\sum_{t=1}^{T} (y_{t+h} - \beta_0 - \mathbf{X}_t \beta)^2 + p(\beta, \lambda) \right]. \tag{3}$$

We can describe the regularization term p as a function of the coefficients β, and a tuning parameter $\lambda \geq 0$ to penalize overfitting. As β_0 depends on the starting value of the target variable, it is not considered in the regularization function. In the following subsections, we introduce the regularization functions and the main features of the ridge and lasso regression methods – two algorithms widely used in prediction policy problems (Medeiros et al. 2021; Silveira et al. 2022).

5.2.1 Ridge

Ridge Regression is a method that aims at reducing the mean error by increasing bias and reducing the variance of OLS estimators to deal with heavily correlated regressors and increasing the reliability of the estimates. The ridge regression method penalizes the OLS estimators by introducing the squared ℓ_2 norm of the coefficients in eq.(3):

$$p(\beta, \lambda) = \lambda ||\beta||_2^2 = \lambda \sum_{i=1}^{n} \beta_i^2. \tag{4}$$

The ℓ_2 norm penalizes complexity, and the coefficients are rarely set to zero – i.e., the solution to the minimization problem is not sparse (Hoerl and Kennard 1970). Typically, ridge regression fits best in circumstances where the least squares estimates have high variance.

5.2.2 Lasso

The Least Absolute Shrinkage and Selection Operator (lasso) regularizes the coefficients in eq.(3) by introducing the ℓ_1 norm:

$$p(\beta, \lambda) = \lambda ||\beta||_1 = \lambda \sum_{i=1}^{n} |\beta_i|. \tag{5}$$

The regularization term in eq.(5) provides a sparse solution. In other words, the lasso algorithm concomitantly performs shrinkage and variable selection (Tibshirani 1996).[20]

5.3 Nonlinear Models

We introduce in this section the random forest and Gradient Tree Boosting algorithms. These are well-known regression tree-based (nonlinear) methods applied in machine learning time-series forecasting.

Regression trees are a powerful tool for approximating unknown nonlinear functions using localized predictions and recursive partitioning of the covariate space. Izenman (2008) explains that the tree-building process begins with the root node and proceeds as follows. Each node within the tree corresponds to a specific region in the covariate space. This region represents a subset of the entire set of predictor variables. The core principle is that each node aims to approximate the relationship between explanatory variables (the predictors) and the variable we want to predict (the target variable). We can achieve this approximation by calculating a single value, which represents the average of the target variable for all observations included within that particular node. The algorithm iteratively seeks the best way to split non-terminal nodes into two child nodes based on a chosen criterion, often using a boolean condition on predictor values as a threshold. As the tree grows, nodes divide into child nodes until predefined stopping criteria are met. When a node no longer splits, it is referred to as a "terminal node" or "leaf." The output value for any data point falling into a terminal node is the average of the target variable values for all the observations within that specific leaf node.

The are many ways to control the size of the tree, but the most common is the maximal depth of the tree. We can find the correct choice of this parameter using a cross-validation approach.[21] While deep trees can overfit the data, shallow trees may not be enough to recognize patterns in the data structure.

5.3.1 Random forest (RF)

Random forests for regression are specialized ensemble models (combination of models) that average multiple individual regression tree models (Breiman 2001). To enhance its predictive power it is crucial to ensure that the constituent trees exhibit a low correlation with each

20. In the context of the Lasso, shrinkage refers to the process of reducing the magnitudes of the coefficients towards zero by adding a penalty term to the regression objective function.

21. Cross-validation refers to a technique used to assess how well a machine learning model generalizes its predictions to unseen data, i.e., observations not used during the training phase.

other. Achieving this requires the implementation of two distinct strategies: (1) to bootstrap the sample, which involves creating different trees by training them on multiple subsets of the training data generated through random sampling with replacement. (2) to consider a random subset of variables rather than considering all available variables in the construction of a specific tree. The growing process of each tree typically involves recursive binary splitting of the data based on a selected feature and threshold that minimizes a splitting criterion such as the mean squared error. The final prediction of the random forest is achieved by taking the average of all the trees.

Random forests can perform notably well in representing and recognizing complex relations in the data set. Besides, it requires simple fine-tuning adjustments, contrary to methods like deep neural networks (Athey and Imbens 2019).

5.3.2 Gradient Tree Boosting

Gradient-boosting tree models (Friedman 2001) are another category of specialized ensemble models for regression. While they share the ensemble concept with random forests, gradient boosting takes a different approach. Instead of constructing independent trees, it forms a sequence of trees in which each new tree seeks to correct the errors made by the previous ones. Gradient-boosting tree models iteratively optimize the underlying tree structures. The goal is to find the best approximating parameters for each tree, including the parameters associated with splitting regions and the combination of each tree in the sequence of models. This optimization is achieved by minimizing a chosen loss function, often the square loss function, which quantifies the difference between the predicted and the actual target values. With each iteration, a new tree is built to focus on the residual errors of the ensemble up to that point. It effectively learns to improve upon the mistakes of the preceding trees. The process continues until a predefined number of iterations is reached or until the loss function reaches a satisfactory minimum. In our paper, we use a specific implementation known as the XGBoost algorithm (Chen and Guestrin 2016) that is a scalable efficient gradient-boosting solution.

5.4 Evaluation Metrics

When evaluating the overall performance and accuracy of machine-learning regression models on time-series data, it is important to use metrics that account for the temporal nature of the data. In this section, we introduce some common evaluation metrics for time-series forecasting.

Mean absolute error (MAE)

MAE is scale-independent and provides a straightforward measure of the average absolute difference between predicted and actual values and is less sensitive to extreme outliers compared to squared error metrics. Also, MAE treats all deviations equally, regardless of the magnitude of the error.

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n} |y_i - \hat{y}_i|$$

Mean absolute scaled error (MASE)

MASE is a scale-independent metric. In our setting, MASE calculates the mean absolute error of each ML algorithm (MAE_{ML}) divided by the mean absolute error of the Random Walk (naive) forecast (MAE_{RW}). MASE < 1 indicate that the evaluated machine learning algorithms perform better than the Random Walk forecast.

$$\text{MASE} = \frac{\text{MAE}_{ML}}{\text{MAE}_{RW}}$$

Mean squared error (MSE)

MSE measures the average of squared errors between predicted and actual values

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n} (y_i - \hat{y}_i)^2.$$

As the MSE formula reveals, the squared errors can be heavily influenced by extreme values, making the mean squared error less robust in the presence of outliers.

5.5 Variable Importance Measures

In prediction policy problems, it is vital to pinpoint the specific predictors that play a key role in enhancing the effectiveness of machine learning algorithms in detecting patterns within a dataset. This not only enhances interpretability but also provides valuable insights for the use of data-driven approaches to guide public policy decisions (Kleinberg et al. 2015;

Athey and Imbens 2019; Doumpos et al. 2023; Wallimann, Imhof, and Huber 2023; Silveira et al. 2023; Brown et al. 2023).

Within our research, we employ Permutation Importance (PI) to highlight the pivotal variables that influence firms' offer prices. Using PI allows for a fair comparison between different models, whether linear or nonlinear. Essentially, PI is a technique that quantifies a variable's influence on a model's predictive capacity by randomly shuffling the variable's values and observing the resulting impact on the quality of predictions (Altmann et al. 2010).

The process of PI involves the following steps: first, training the model with all input variables and evaluating its performance on test data; then, shuffling the values of a single input variable in the test data while keeping others constant and reassessing the model's performance on the shuffled data; next, comparing the original and shuffled performance metrics to determine the impact, i.e., the increase in prediction error; and finally, ranking the variables based on this impact, where larger increases in prediction error signify greater importance. For example, a variable with a PI of 0.1 denotes a 0.1 increase in the model's prediction error when its values are shuffled. In other words, this metric can be used to compare variables within the same model, as it allows both variable selection and the identification of the most crucial variables for precise predictions.

The individual permutation importance technique has been used as a way to increase our ability to understand and explain the results generated by machine learning models (Genuer, Poggi, and Tuleau-Malot 2010; Aras and Lisboa 2022; Ji et al. 2022; Silveira et al. 2022; Silveira et al. 2023; Brown et al. 2023). However, this approach has its limitations. In a database with a large number of predictors, some of them may be highly correlated, causing collinearity issues that can compromise the accurate measurement of each predictor's individual contribution to the model's performance. As a result, individual permutation importance may not fully reflect the combined impact of correlated predictors on the model (Gregorutti, Michel, and Saint-Pierre 2017; Lundberg et al. 2020; Mi et al. 2021; Molnar et al. 2023).

Group-based permutation importance involves assessing the importance of groups of variables rather than individual variables, providing a more comprehensive view of the importance of predictor groups, especially in situations where predictors exhibit strong correlations. The process begins by defining groups of related variables based on domain knowledge or patterns observed in the data. Then, similar to individual variable permutation importance, the values of the variables within a group are shuffled simultaneously, and the model's performance is evaluated based on these shuffled values. By comparing the model's performance with the original unshuffled data, we can determine the impact of shuffling variables within a group on the quality of predictions (Au et al. 2022; Plagwitz et al. 2022). This approach allows us to understand the combined influence of related variables within a group. Additionally, it provide a more detailed view of how the categories of predictors – and their underlying specifications we presented in Section 4.3 – may improve the prediction of firms' bidding decisions. In other words, by applying the grouped permutation importance, we gain deeper insights into how predictors influence machine-learning models and uncover potential synergistic effects among variables within a group, giving a more extensive understanding of the collective impact of the categories of variables of specifications 1 and 2 on predictive performance.

We define eight different categories in our group-based permutation importance approach. These eight categories were determined by grouping variables based on that capture similar market dynamics. Many of these variables are moderately to highly correlated. The categories include the following. "Own Firm Lagged Price Variables" controls for lagged values of the dependent variable to capture the influence of previous bidding on current pricing decisions. "Other Own Firm Lagged Variables" includes lagged values of a firm's outputs and capacity availability, separated by generation technology. "Wholesale Market Demand and Price Variables" captures measures that potentially impact bidding behavior. "Supply and Availability by Fuel Type" comprises market-level generation supply and capacity availability variables to determine if firms adjust their bids based on observed variations. "Imports and Exports" consists of import and export quantities and intertie capacity availability measures for neighboring jurisdictions. "AS Market" includes the total Ancillary Services market output to capture the quantity of reserves procured by the AESO. "Temperature" and "Hour" comprise temperature variables for Edmonton and Calgary and hour indicators, respectively, aiming to capture time and temperature-based variations impacting bidding decisions. By applying grouped permutation importance one group at a time, we can discern which sets of variables contribute most significantly to our understanding of bidding behaviors and pricing decisions in the market.

6 Results

Table 3 presents the values of the evaluation metrics for each model. Note, first, that all the evaluated ML algorithms perform better than the random walk, indicating that the data provided to the algorithms are useful in predicting price changes. There are limited differences in the performance across algorithms for all four firms; while the non-linear tree-based algorithms (random forest and XGBoost) yield lower MAE and MASE scores than Lasso and Ridge for all four firms and both specifications, the MSE suggests that Lasso exhibits stronger performance for three of the firms. Throughout our results, we will focus on the results from the random forest algorithm and briefly summarize results from the Lasso algorithm. Additional results for the other algorithms will be presented in the Appendix to demonstrate that our key conclusions are robust.

Table 3: MAE, MASE and MSE

	Capital Power			ENMAX			Heartland			TransAlta		
	MAE	MASE	MSE	MAE	MASE	MSE	MAE	MASE	MSE	MAE	MASE	MSE
Random Walk	0.494	-	0.587	0.633	-	0.843	0.259	-	0.177	0.463	-	0.537
Specification I												
Lasso	0.289	0.585	**0.195**	0.295	0.466	**0.166**	0.167	0.645	0.076	0.313	0.677	**0.211**
Ridge	0.304	0.615	**0.195**	0.301	0.476	0.167	0.188	0.723	0.080	0.337	0.729	0.231
RF	**0.282**	**0.571**	0.211	**0.279**	**0.441**	0.179	**0.153**	**0.591**	**0.075**	**0.283**	**0.612**	0.217
XGBoost	0.286	0.579	0.204	0.285	0.450	0.178	0.170	0.656	0.078	0.296	0.639	0.216
Specification II												
Lasso	0.298	0.603	**0.193**	0.301	0.476	**0.169**	0.164	0.633	0.077	0.323	0.698	0.218
Ridge	0.315	0.637	0.200	0.314	0.496	0.175	0.197	0.761	0.084	0.348	0.752	0.240
RF	0.283	0.573	0.211	**0.284**	**0.449**	0.181	**0.152**	**0.587**	**0.076**	**0.282**	**0.610**	0.217
XGBoost	**0.277**	**0.561**	0.200	0.286	0.452	0.173	0.172	0.664	0.080	0.294	0.635	**0.215**

Notes. MAE, MASE, and MSE are calculated based on the first difference in the logarithm of prices over three hours. The best fit for each metric among the four models is indicated in bold.

Table 3 does not indicate a large performance improvement when disaggregated firm-level supply and available capacity variables are used. While Specification II does result in slightly lower errors for some firms and performance measures, these differences are small and not consistent across algorithms or performance measures. Our analysis suggests that the disaggregated information does not provide a considerable increase in information to predict firms' pricing decisions. Instead, it tends to introduce more noise and greater variance into the predictions. In forecasting models, variance is a source of error. Therefore, our results show that the additional variance resulting from using firm-level and more disaggregated data usually leads to greater predictive error. In addition, this suggests that firms rely primarily on market-level data rather than firm-specific information on rivals' generation units when setting prices.

In addition to providing common evaluation metrics for time-series forecasting, we also examine how the analysis performs in predicting each firm's pricing decisions by looking at actual and forecasted price levels. Figure 1 provides a time-series plot of the actual and predicted quantity-weighted offer price (i.e., our main target variable) using the random forest algorithm and Specification I.[22] While there is considerable variation in the firms' pricing behavior both over the entire sample and within a given month, we can see that our analysis can track the observed pricing patterns reasonably well.[23]

Figure 1: Price Level Forecast - Random Forest

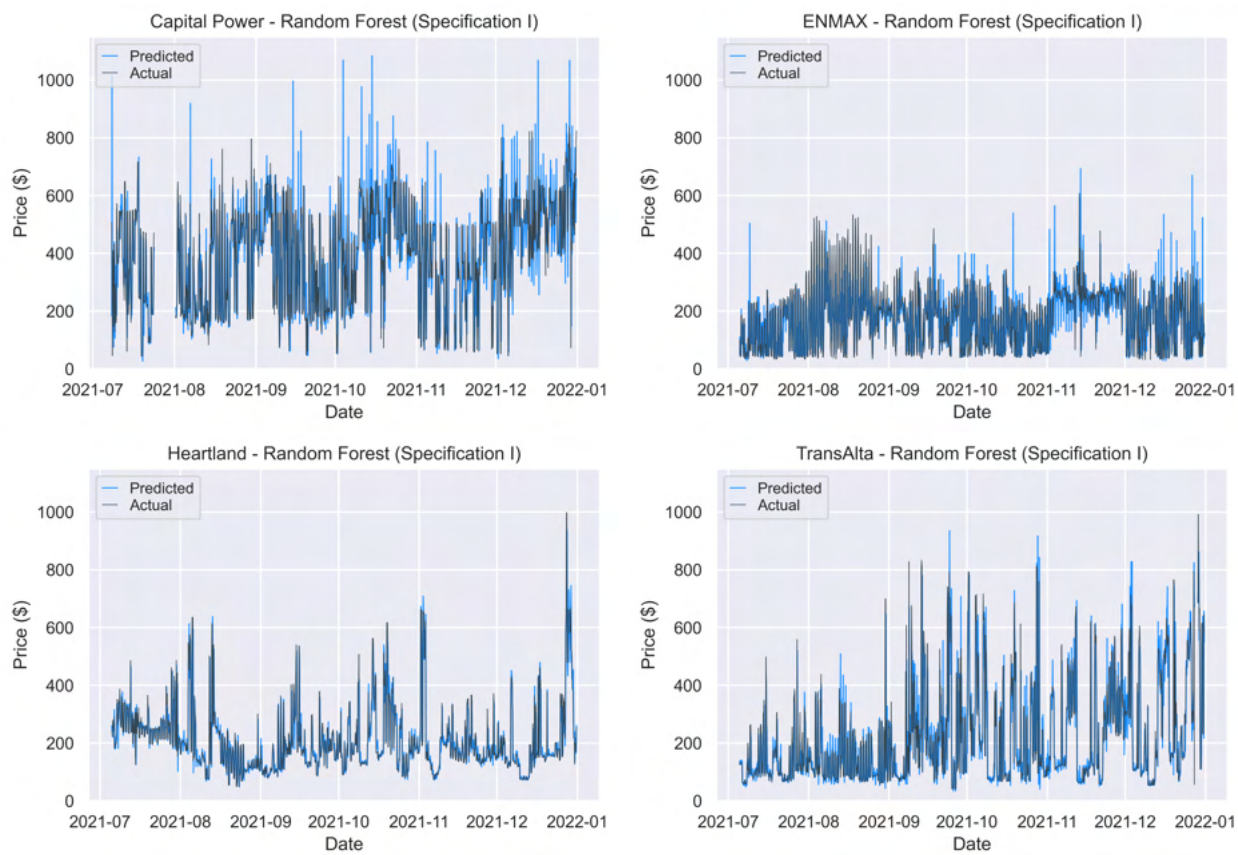

Figure 2 presents the distribution of the differences between the actual and predicted price variable for each firm using the random forest algorithm and Specification I.[24] While the range is broad, the vast majority of observations lie closely centered around 0. This continues to demonstrate that our analysis captures firms' pricing dynamics on their strategic offers (i.e., offers in excess of marginal cost).

22. There is a gap in Figure 1 for Capital Power and smaller gap for TransAlta. This is due to the fact that our analysis focuses on the bids of generation units that exceeded short-run marginal cost (recall Section 4.2). Capital Power and TransAlta did not have bids above their marginal cost during these periods.
23. See Figure B.1 in Appendix B for the price level forecasts for all the four firms using Lasso.
24. See Figure B.2 in Appendix B for the distribution of the differences using Lasso.

Figure 2: Histogram - Random Forest

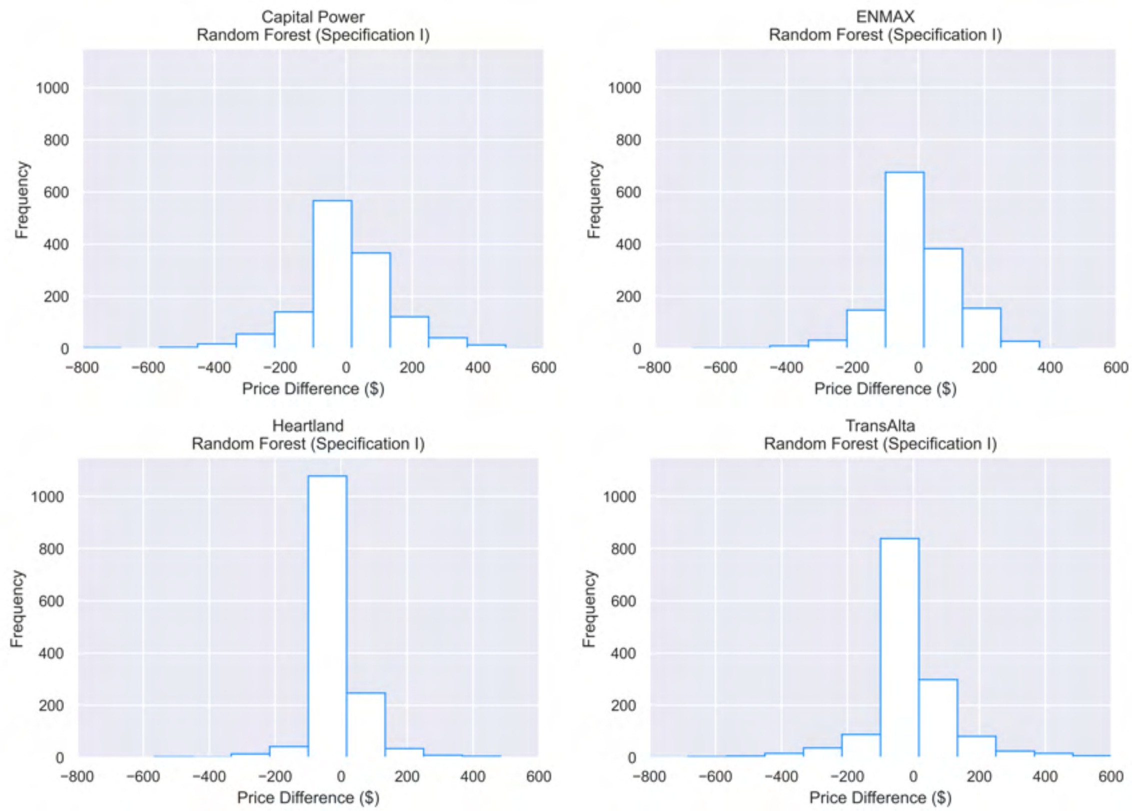

To gain insight into which variables are more important for prediction, we now consider the grouped permutation importance measures. We use box plots to illustrate the distribution of an increase in prediction error when the original value of individual and grouped variables is replaced by random counterparts. The box represents the interquartile range, i.e., it describes how spread out the variable importance is, and the line within the box indicates the median. Outliers can highlight variables with exceptional individual or grouped importance over the sample period. Thus, the box plots allow us to visualize variations in importance levels and compare between different variables – either individually or grouped.

Figure 3 presents the grouped permutation importance based on the predictions made by the random forest algorithm. The Figure presents results for both Specifications I and II to evaluate the extent to which the inclusion of firm-specific variables impacts the analysis. Random forest primarily uses a combination of information derived from the own-firm-lagged variables representing a firm's previous bidding behavior, wholesale market demand and price variables, and supply and availability variables by fuel type. In addition, hourly indicators have an important impact on the predictions for three of the four firms. This is consistent with daily patterns in prices, in which bids are raised for peak hours and lowered in off-peak hours. In contrast, variables related to imports and exports, ancillary services, or temperature appear to play only a minor role in prediction.

Figure 3: Box Plot of Grouped Permutation Importance - Random Forest

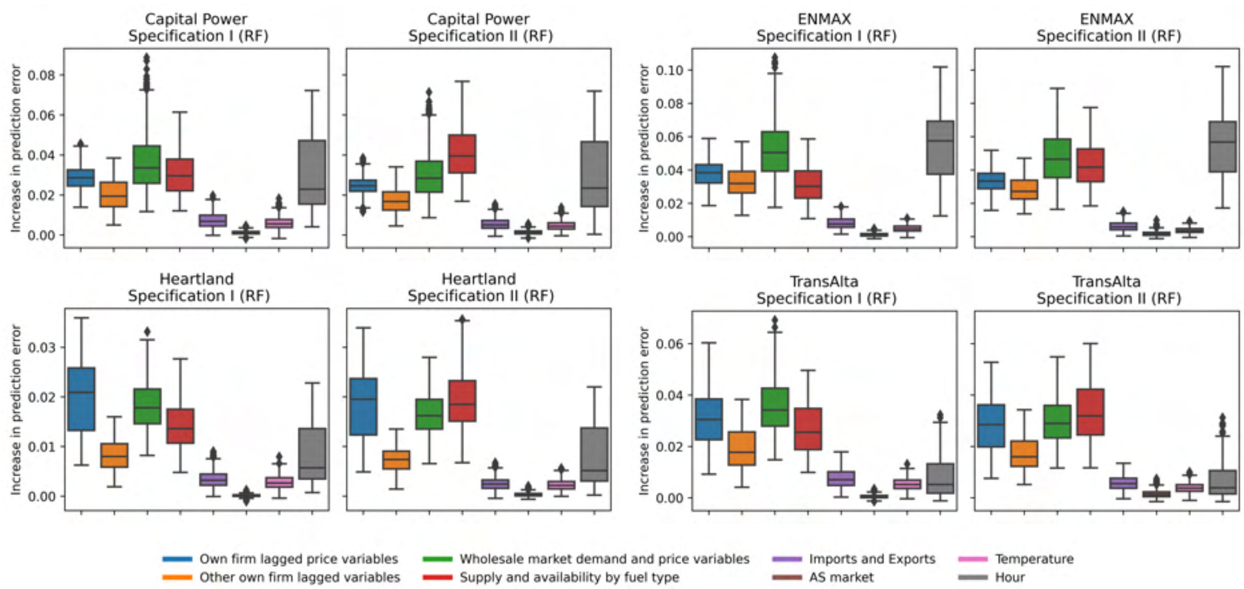

Comparing the results for Specifications I and II, we observe a modest increase in the importance of the supply and available capacity variables. However, recall from Table 3 that this does not translate into a substantial improvement in predictive performance. This continues to suggest that firms rely primarily on aggregate market-level variables to respond to supply and demand dynamics when making their bidding decisions.

Figure 4 presents the grouped permutation importance analysis using the Lasso algorithm. Different from the random forest algorithm, the results using Lasso indicate that this algorithm places considerably more importance on lagged own-price variables and hourly indicators; the importance of market price and demand and supply and availability variables is reduced compared to Figure 3. This is possibly driven by the linear nature of the Lasso algorithm, which may be less able to exploit these variables than one that permits their effect to be non-linear.[25] Despite these differences, we continue to observe only modest differences between Specifications I and II when disaggregated firm-level variables are included. This continues to support the conclusion that firms rely primarily on aggregate market-level

25. See Figures B.3 and B.4 in Appendix B for the results obtained with both the XGBoost and ridge models, respectively.

variables when forming their pricing decisions.

Figure 4: Box Plot of Grouped Permutation Importance - Lasso

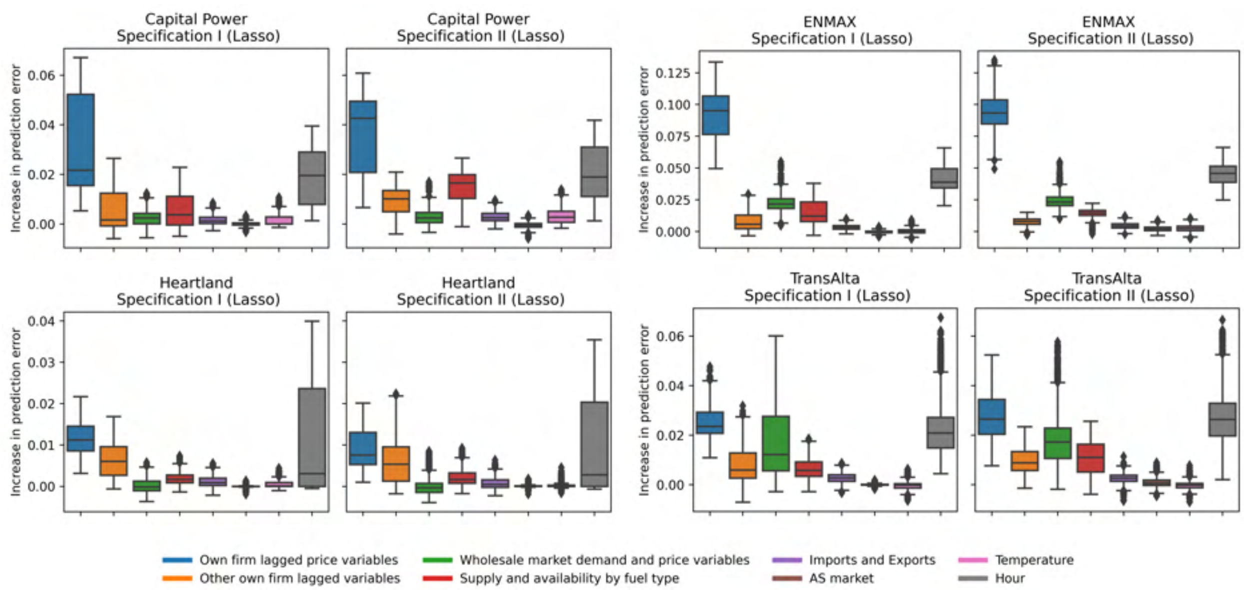

We now carry out our individual permutation importance analysis to explore the contribution of each individual variable in the price prediction algorithm. While the individual permutation importance approach faces challenges when variables are highly correlated (recall the discussion in Section 5.5), it can provide useful insights to the extent that it can identify the importance of individual variables that were previously nested within groups in the grouped permutation importance approach.

Figure 5 shows the ten variables that most increased the random forest's prediction error after having their values replaced by randomly shuffled counterparts for both Specifications I and II. Focusing on Specification I, for two of the four firms, the single most important variable for prediction is the forecast of the Alberta internal load (demand) for the current hour (ForecastAIL F0); for Capital Power and Heartland, this variable has the second highest permutation importance score. Overall, we notice that Alberta's internal load variables represent two to three of the top ten variables by permutation importance for our four firms. Additionally, lagged target variables (i.e., lagged price variables) or hourly dummy variables account for six of the top ten most important variables for all four large firms.

The combination of lagged prices, hourly indicators, and market demand forecast variables suggests that daily pricing patterns and changes in observed or expected load are the most important variables for predicting offer price adjustments. Other individual variables in the top ten for at least one of the firms include the actual pool price, SMPs coefficient of variation, and dual fuel supply. Observed price levels and the price variance measure may be used to provide information about the shape of the market supply curve, supplementing the important market demand variables. The importance of dual fuel supply could be capturing the fact that these are recent coal-to-gas converted units that are often setting the market-clearing price (i.e., are on the margin) in peak hours during our sample period.

Figure 5: Box Plot of Top 10 Explanatory Variables - Random Forest

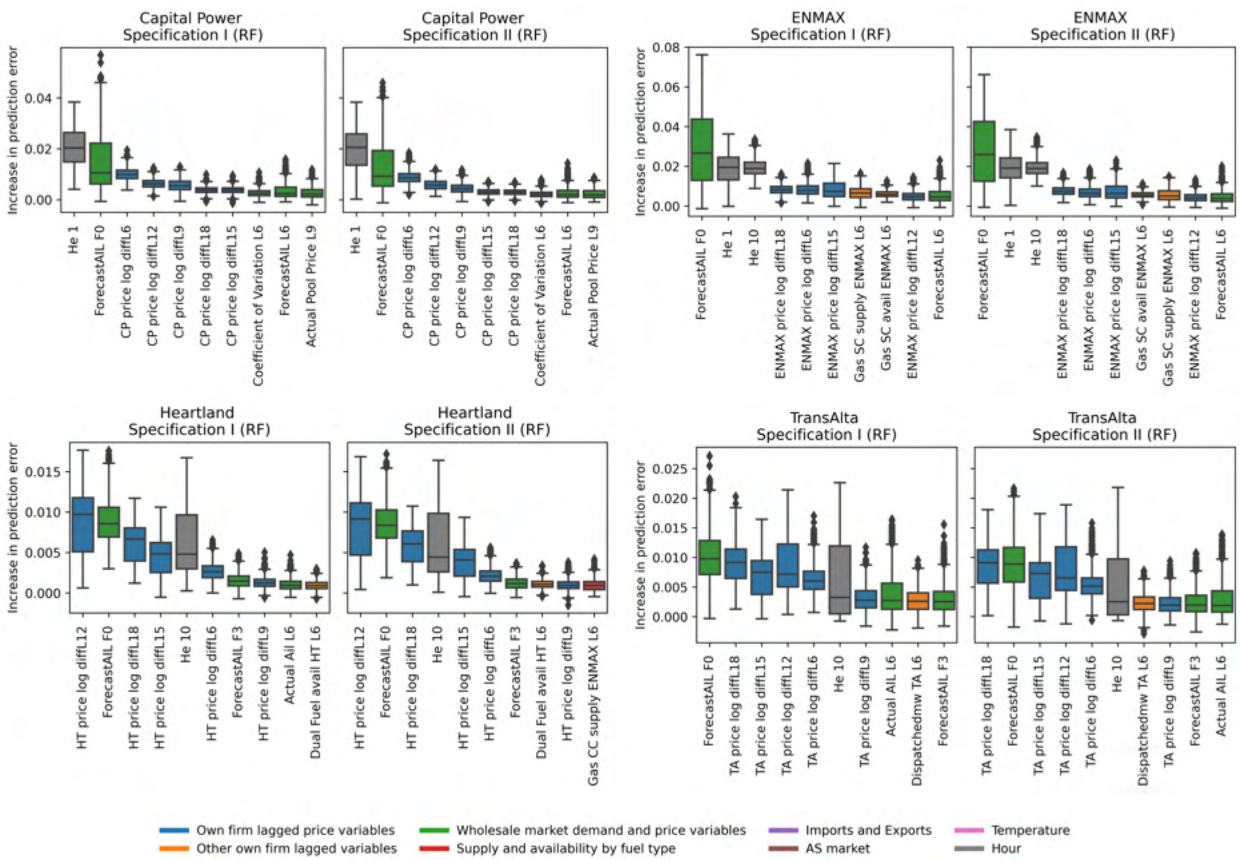

Figure 5 demonstrates that the conclusion regarding the relative importance of daily patterns, observed demand, and demand forecasts carries over to Specification II, where lagged own-price variables, hour dummies, and market demand variables account for seven up to nine of the top ten variables, depending on the firm.[26] In fact, only a single rival firm-specific variable shows up in the top 10 looking across all four firms; ENMAX's supply

26. From Figure B.5 in the Appendix B, we can see that the predictions made by XGBoost are based on patterns similar to those used by random forest.

from combined cycle (CC) gas units is the 10th most important single variable for Heartland. This continues to support the fact that the consideration of firm-specific variables does not have a substantive impact on the price predictions or conclusions drawn from the analysis.

Figure 6: Box Plot of Top 10 Explanatory Variables - Lasso

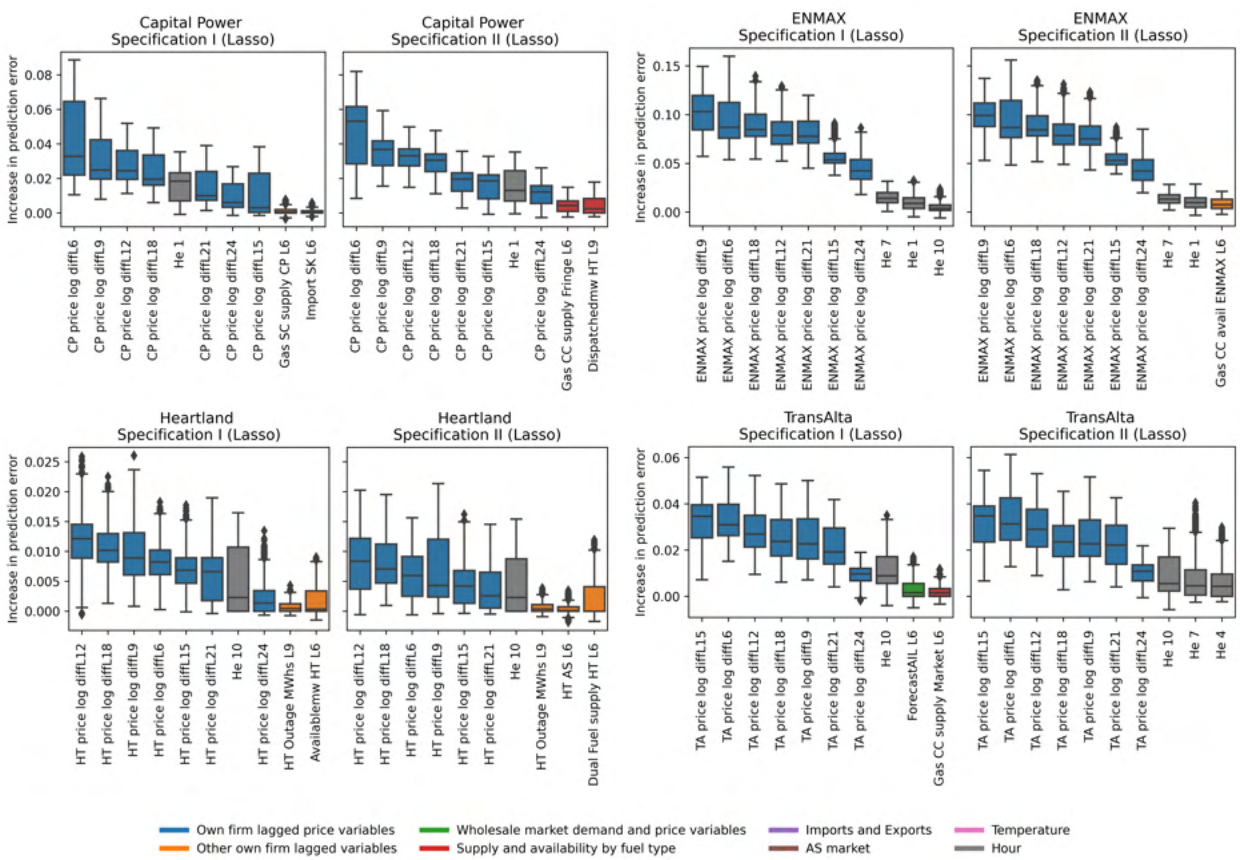

Comparatively, we see through Figure 6 that for both specifications I and II, lasso extensively uses lagged price predictors combined with hourly indicators and other own-firm-specific lagged variables to predict price changes.[27] This parallels the findings in Figure 4 when using the grouped permutation importance approach. As suggested above, this could be driven by the linear nature of the lasso algorithm. Overall, analyzing the 10 individual variables with the highest permutation importance, we can see that the algorithms are capable of assimilating, albeit in subtly different ways, the daily patterns of wholesale prices. Further, the inclusion of disaggregated firm-specific variables in Specification II does not have a considerable impact on the key variables with the highest permutation importance.

27. Figure B.6 in Appendix B displays the box plots of the top ten variables for the ridge model.

Figure 7: Box Plot of Finer Group Analysis - Random Forest

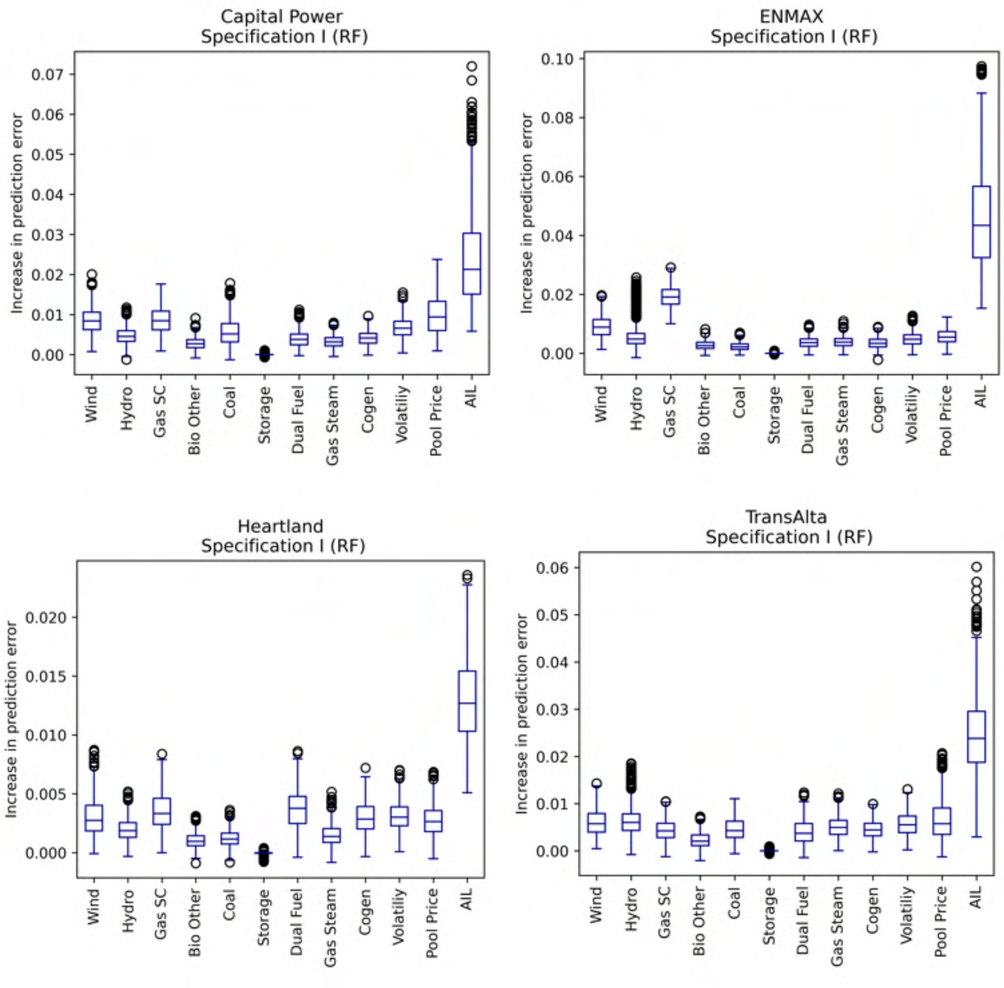

The analysis above indicates that while the "Wholesale market demand and price" and "Supply and availability by fuel type" groups are important for prediction in the random forest algorithm, the only variables in these categories to show up consistently in the top ten as shown in Figure 5 are measures of forecasted and observed AIL. However, other variables in these categories may not have high individual permutation importance scores if they are highly correlated with other variables in the same group. To explore this further, focusing on the random forest algorithm for Specification I, we divide the supply/availability variables into groups based on technology, with each group containing the lags of the available and supply variables for that specific technology. The market demand and price variables are divided into three groups; "Volatility" includes lags of both SMP standard deviation and coefficient of variation, "Pool Price" consists of lags of day-ahead and actual pool price variables, and "AIL" is made up of lags and leads of both forecast and actual AIL variables. This approach represents an intermediate strategy where we retain the grouped structure of variables but with a more granular subdivision. This enables us to harmonize the broad overview offered by the grouped permutation importance method with the need for correctly addressing multicollinearity issues inherent in the individual permutation importance method.

The results of the grouped permutation importance for these finer groups are shown in Figure 7. For all four firms, the AIL variables are the most important, while pool price levels and volatility measures play a minor role in prediction. Of the supply and availability categories, ENMAX specifically highlights simple cycle gas (Gas SC) as significant among other technologies, while no single technology stands out for the other firms.[28] Hence, this analysis suggests that the previous finding, that AIL variables are the primary driver of prices, is robust to collinearity concerns, while no single technology seems to be important for forecasting.

7 Conclusion

An important consideration in the design of restructured wholesale electricity markets is the degree of transparency in the market and the amount of information available to market participants in real or near-real time. While on the one hand, a high degree of information availability can promote efficient market outcomes by allowing firms to respond optimally to supply and demand circumstances, there is also a potential that large amounts of firm-level information could facilitate coordination among firms.

In this paper, we use data from Alberta's wholesale electricity market for 2021 to consider which information best predicts firm-level offer prices and the reliance of firms on market-level vs disaggregated firm-level data. Specifically, we use machine learning algorithms to predict changes in the bid prices of the four largest firms, over a three-hour period, focusing on those assets that are used for economic withholding. We ask whether these algorithms are more successful in predicting price adjustments when they have access to data on the generation and available capacity of specific rival firms, compared to when only market-level generation and capacity data are included. Furthermore, we explain and interpret the

28. Figure B.7 in Appendix B shows box plots of the detailed group analysis for the other evaluated algorithms. Similar to the random forest algorithm, the AIL variables are important predictors for the forecasts made with all the algorithms, especially for XGBoost. Furthermore, technologies such as Coal, Gas SC, and Dual Fuel may also impact the predictions made with both lasso and ridge regression.

main factors that drive firms' bidding behavior using Permutation Importance techniques on individual and grouped data. This technique is proving to be a valuable supplementary tool for addressing prediction policy problems and is playing a significant role in dispelling the notion that machine learning models are "black boxes."

In general, we find that firm-specific information does not improve the prediction performance of our algorithms, suggesting that firms primarily make use of market-level information when formulating price changes. We find that hour indicators and lagged own-price changes are important predictors, suggesting strong daily patterns in price adjustments. Group-level permutation importance scores applied to the results for the non-linear tree-based algorithms indicate that variables in the market demand and price category and the supply and available capacity category are important predictors. Dividing these latter categories into smaller groups suggests that key predictors of firm-level price changes are measures of Alberta Internal Load, the primary measure of electricity demand in Alberta. We find limited indication that the supply or availability of specific technologies are important for prediction. Finally, our results suggest that certain categories of variables (ancillary services, imports and exports, own non-price lagged variables) are not important for predicting changes to firm-level prices.

In terms of policy implications, our results suggest that a focus on providing firms with more disaggregated firm-level real-time information may be misguided. Our results indicate that firms rely primarily on market-level information, particularly on market demand, when choosing bids; as a result, the potential gains from more disaggregated data may be low relative to the increased risk of coordination. Rather, it may be more beneficial to provide more accurate information on market-level indicators such as observed and forecasted load and market prices. This may involve improving forecast methodologies for demand and the pool price. While our results do not indicate that wind and solar supply/availability measures are of particular importance for predicting price changes, one might expect this to change as these technologies represent a higher proportion of total generation capacity. In this event, improved wind and solar generation forecasting methods may be of value.

References

Abrantes-Metz, Rosa M, Sofia B Villas-Boas, and George Judge. 2011. "Tracking the Libor Rate." *Applied Economics Letters* 18 (10): 893–899.

Albuquerque, Pedro C, Daniel O Cajueiro, and Marina DC Rossi. 2022. "Machine Learning Models for Forecasting Power Electricity Consumption Using a High Dimensional Dataset." *Expert Systems with Applications* 187:115917.

Altmann, André, Laura Toloşi, Oliver Sander, and Thomas Lengauer. 2010. "Permutation Importance: A Corrected Feature Importance Measure." *Bioinformatics* 26 (10): 1340–1347.

Aras, Serkan, and Paulo JG Lisboa. 2022. "Explainable Inflation Forecasts by Machine Learning Models." *Expert systems with applications* 207:117982.

Aryal, Gaurab, Federico Ciliberto, and Benjamin Leyden. 2022. "Coordinated Capacity Reductions and Public Communication in the Airline Industry." *Review of Economic Studies* 89:3055–3084.

Athey, Susan, and Guido W Imbens. 2019. "Machine Learning Methods that Economists Should Know About." *Annual Review of Economics* 11:685–725.

Au, Quay, Julia Herbinger, Clemens Stachl, Bernd Bischl, and Giuseppe Casalicchio. 2022. "Grouped Feature Importance and Combined Features Effect Plot." *Data Mining and Knowledge Discovery* 36 (4): 1401–1450.

AUC. 2017. "Application by the Market Surveillance Administrator Regarding the Publication of the Historical Trading Report. Decision Proceeding: 21115–D01-2017." Alberta Utilities Commission.

Bergheimer, Stefan, Estelle Cantillon, and Mar Reguant. 2023. "Price and Quantity Discovery without Commitment." *International Journal of Industrial Organization* 90:102987.

Borenstein, Severin. 1998. "Rapid Communication and Price Fixing: The Airline Tariff Publishing Company Case." In *The Antitrust Revolution: The Role of Economics,* edited by John Kowka and Lawrence White. Oxford University Press.

Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45 (1): 5–32.

Brown, David, Daniel Cajueiro, Andrew Eckert, and Douglas Silveira. 2023. "Information and Transparency: Using Machine Learning to Detect Communication Between Firms." *Stanford Computational Antitrust* 3:198–231.

Brown, David P, and Andrew Eckert. 2022. "Pricing Patterns in Wholesale Electricity Markets: Unilateral Market Power or Coordinated Behavior?" *The Journal of Industrial Economics* 70 (1): 168–216.

Brown, David P, Andrew Eckert, and James Lin. 2018. "Information and Transparency in Wholesale Electricity Markets: Evidence from Alberta." *Journal of Regulatory Economics* 54 (3): 292–330.

Brown, David P, Andrew Eckert, and Blake Shaffer. 2023. "Evaluating the Impact of Divestitures on Competition: Evidence from Alberta's Wholesale Electricity Market." *International Journal of Industrial Organization* 89:102953.

Brown, David P, Andrew Eckert, and Douglas Silveira. 2023. "Strategic Interaction between Wholesale and Ancillary Service Markets." *Competition and Regulation in Network Industries* 24 (4): 174–198.

Bunn, Derek W., and Stefan O.E. Kermer. 2021. "Statistical Arbitrage and Information Flow in an Electricity Balancing Market." *The Energy Journal* 42 (5): 19–40.

Byrne, David P., and Nicolas de Roos. 2019. "Learning to Coordinate: A Study in Retail Gasoline." *American Economic Review* 109 (2): 591–619.

CASA. 2004. *A Study on the Efficiency of Alberta's Electrical Supply System.* Project # CASA-EEEC-02-04. Clean Air Strategic Alliance, Prepared by JEM Energy.

Chen, Tianqi, and Carlos Guestrin. 2016. "Xgboost: A Scalable Tree Boosting System." In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,* 785–794.

Christie, William G, and Paul H Schultz. 1999. "The Initiation and Withdrawal of Odd-Eighth Quotes Among Nasdaq Stocks: An Empirical Analysis." *Journal of Financial Economics* 52 (3): 409–442.

Doumpos, Michalis, Constantin Zopounidis, Dimitrios Gounopoulos, Emmanouil Platanakis, and Wenke Zhang. 2023. "Operational Research and Artificial Intelligence Methods in Banking." *European Journal of Operational Research* 306 (1): 1–16.

Friedman, Jerome H. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." *Annals of Statistics,* 1189–1232.

Gelman, Andrew, and Iain Pardoe. 2007. "Average Predictive Comparisons for Models with Nonlinearity, Interactions, and Variance Components." *Sociological Methodology* 37 (1): 23–51.

Genuer, Robin, Jean-Michel Poggi, and Christine Tuleau-Malot. 2010. "Variable Selection Using Random Forests." *Pattern Recognition Letters* 31 (14): 2225–2236.

Green, Edward J., and Robert H. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica* 52 (1): 87–100.

Gregorutti, Baptiste, Bertrand Michel, and Philippe Saint-Pierre. 2017. "Correlation and Variable Importance in Random Forests." *Statistics and Computing* 27:659–678.

Hamilton, James D. 2020. *Time Series Analysis.* Princeton University Press.

Harrington, Joseph, and Andrzej Skrzypacz. 2007. "Collusion under Monitoring of Sales." *RAND Journal of Economics* 38 (2): 314–331.

Hastie, Trevor, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Vol. 2. Springer.

Hoerl, Arthur E, and Robert W Kennard. 1970. "Ridge Regression: Biased Estimation for Nonorthogonal Problems." *Technometrics* 12 (1): 55–67.

Holmberg, Pär, and Thomas Tangerås. 2023. "A Survey of Capacity Mechanisms: Lessons for the Swedish Electricity Market." *The Energy Journal* 44 (6): 275–304.

Holmberg, Pär, and Frank A Wolak. 2018. "Comparing Auction Designs Where Suppliers Have Uncertain Costs and Uncertain Pivotal status." *The RAND Journal of Economics* 49 (4): 995–1027.

IEA. 2022. *Renewables 2022 Analysis and Forecast to 2027.* International Energy Agency.

Izenman, Alan Julian. 2008. "Modern Multivariate Statistical Techniques." *Regression, Classification and Manifold Learning* 10:978–.

Jedrzejewski, Arkadiusz, Jesus Lago, Grzegorz Marcjasz, and Rafal Weron. 2022. "Electricity Price Forecasting: The Dawn of Machine Learning." *IEEE Power & Energy Magazine* 20 (3): 24–31.

Ji, Gang, Jingmin Yu, Kai Hu, Jie Xie, and Xunsheng Ji. 2022. "An Adaptive Feature Selection Schema Using Improved Technical Indicators for Predicting Stock Price Movements." *Expert Systems with Applications* 200:116941.

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. 2015. "Prediction Policy Problems." *American Economic Review* 105 (5): 491–495.

Kuhn, K.-U., and X. Vives. 1994. *Information Exchanges Among Firms and their Impact on Competition.* European Commission Working Paper.

Lago, Jesus, Grzegorz Marcjasz, Bart De Schutter, and Rafal Weron. 2021. "Forecasting Day-ahead Electricity Prices: A Review of State-of-the-art Algorithms, Best Practices and an Open-access Benchmark." *Applied Energy* 293:116983.

Lewis, Matthew S. 2015. "Odd Prices at Retail Gasoline Stations: Focal Point Pricing and Tacit Collusion." *Journal of Economics & Management Strategy* 24 (3): 664–685.

Lundberg, Scott M, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. "From Local Explanations to Global Understanding with Explainable AI for Trees." *Nature machine intelligence* 2 (1): 56–67.

Masini, Ricardo P, Marcelo C Medeiros, and Eduardo F Mendes. 2023. "Machine Learning Advances for Time Series Forecasting." *Journal of economic surveys* 37 (1): 76–111.

Medeiros, Marcelo C, Gabriel FR Vasconcelos, Álvaro Veiga, and Eduardo Zilberman. 2021. "Forecasting Inflation in a Data-rich Environment: The Benefits of Machine Learning Methods." *Journal of Business & Economic Statistics* 39 (1): 98–119.

Mi, Xinlei, Baiming Zou, Fei Zou, and Jianhua Hu. 2021. "Permutation-based Identification of Important Biomarkers for Complex Diseases via Machine Learning Models." *Nature Communications* 12 (1): 3008.

Molnar, Christoph, Gunnar König, Bernd Bischl, and Giuseppe Casalicchio. 2023. "Model-agnostic Feature Importance and Effects with Dependent Features: A Conditional Subgroup Approach." *Data Mining and Knowledge Discovery,* 1–39.

Molnar, Christoph, Gunnar König, Julia Herbinger, Timo Freiesleben, Susanne Dandl, Christian A Scholbeck, Giuseppe Casalicchio, Moritz Grosse-Wentrup, and Bernd Bischl. 2020. "General Pitfalls of Model-agnostic Interpretation Methods for Machine Learning Models." In *International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers,* 39–68. Springer.

MSA. 2013. "Coordinated Effects and the Historical Trading Report: Decision and Recommendation." *Market Surveillance Administrator.*

———. 2021. "Market Share Offer Control Report Data." *Market Surveillance Administrator.*

Murthy, Girish, Vijayalakshmi Sedidi, Ajaya Kumar Panda, and Badri Narayan Rath. 2014. "Forecasting Electricity Prices in Deregulated Wholesale Spot Electricty Market: A Review." *International Journal of Energy Economics and Policy* 4 (1): 33–42.

Perumean-Chaney, Suzanne E, Charity Morgan, David McDowall, and Inmaculada Aban. 2013. "Zero-inflated and Overdispersed: What's One to Do?" *Journal of Statistical Computation and Simulation* 83 (9): 1671–1683.

Petropoulos, Fotios, Daniele Apiletti, Vassilios Assimakopoulos, Mohamed Zied Babai, Devon K Barrow, Souhaib Ben Taieb, Christoph Bergmeir, Ricardo J Bessa, Jakub Bijak, John E Boylan, et al. 2022. "Forecasting: Theory and Practice." *International Journal of Forecasting* 38 (3): 705–871.

Plagwitz, Lucas, Alexander Brenner, Michael Fujarski, and Julian Varghese. 2022. "Supporting AI-Explainability by Analyzing Feature Subsets in a Machine Learning Model." In *MIE,* 109–113.

Sai, Wei, Zehua Pan, Siyu Liu, Zhenjun Jiao, Zheng Zhong, Bin Miao, and Siew Hwa Chan. 2023. "Event-driven Forecasting of Wholesale Electricity Price and Frequency Regulation Price Using Machine Learning Algorithms." *Applied Energy* 352:121989.

Sapio, Alessandro. 2021. "Econometric Modelling and Forecasting of Wholesale Electricity Prices: Fundamentals and Applications for Engineers and Energy Planners." In *Handbook of Energy Economics and Policy,* edited by Alessandro Rubino, Alessandro Sapio, and Massimo La Scala, 595–640. Elsevier.

Silveira, Douglas, Lucas B de Moraes, Eduardo PS Fiuza, and Daniel O Cajueiro. 2023. "Who Are You? Cartel Detection Using Unlabeled Data." *International Journal of Industrial Organization* 88:102931.

Silveira, Douglas, Silvinha Vasconcelos, Marcelo Resende, and Daniel O Cajueiro. 2022. "Won't Get Fooled Again: A Supervised Machine Learning Approach for Screening Gasoline Cartels." *Energy Economics* 105:105711.

Syntetos, Aris A, and John E Boylan. 2005. "The Accuracy of Intermittent Demand Estimates." *International Journal of forecasting* 21 (2): 303–314.

Tibshirani, Robert. 1996. "Regression Shrinkage and Selection Via the Lasso." *Journal of the Royal Statistical Society: Series B (Methodological)* 58 (1): 267–288.

von der Fehr, Nils-Henrik M. 2013. "Transparency in Electricity Markets." *Economics of Energy & Environmental Policy* 2 (2): 87–106.

Wallimann, Hannes, David Imhof, and Martin Huber. 2023. "A Machine Learning Approach for Flagging Incomplete Bid-rigging Cartels." *Computational Economics* 62 (4): 1669–1720.

Weber, Paige, and Matt Woerman. 2022. *Intermittency or Uncertainty? Impacts of Renewable Energy in Electricity Markets.* CESifo Working Paper No. 9902.

Appendix A Additional Tables

Table A.1: Set of Explanatory Variables used in Specification I

	Capital Power (CP)	ENMAX (EN)	Heartland (HT)	TransAlta (TA)
Own firm lagged price variables	CP_price_log_diffL6, CP_price_log_diffL9, CP_price_log_diffL12, CP_price_log_diffL15, CP_price_log_diffL18, CP_price_log_diffL21, CP_price_log_diffL24	EN_price_log_diffL6, EN_price_log_diffL9, EN_price_log_diffL12, EN_price_log_diffL15, EN_price_log_diffL18, EN_price_log_diffL21, EN_price_log_diffL24	HT_price_log_diffL6, HT_price_log_diffL9, HT_price_log_diffL12, HT_price_log_diffL15, HT_price_log_diffL18, HT_price_log diffL21, HT_price_log_diffL24	TA_price_log_diffL6, TA_price_log_diffL9 TA_price_log_diffL12, TA_price_log_diffL15, TA_price_log_diffL18, TA_price_log_diffL21, TA_price_log_diffL24
Other own firm lagged variables	AS_Output_CP, Avaliablemw_CP, Coal_avail_CP, Coal_supply_CP, Dispatchedmw_CP, Gas_CC_avail_CP, Gas_CC_supply_CP, Gas_SC_avail_CP, Gas_SC_supply_CP, Wind_avail_CP, Wind_supply_CP	AS_Output_EN, Avaliablemw_EN, Dispatchedmw_EN, Gas_CC_avail_EN, Gas_CC_supply_EN, Gas_SC_avail_EN, Gas_SC_supply_EN, Wind_avail_EN, Wind_supply_EN	AS_Output_HT, Avaliablemw_HT, Cogen_Total_avail_HT, Cogen_Total_supply_HT, Dispatchedmw_HT, Dual_Fuel_avail_HT, Dual_Fuel_supply_HT, Gas_SC_avail_HT, Gas_SC_supply_HT, HT_Outage	AS_Output_TA, Avaliablemw_TA, Coal_avail_TA, Coal_supply_TA, Dispatchedmw_TA, Dual_Fuel_avail_TA, Dual_Fuel_supply_TA, Gas_Steam_avail_TA, Gas_Steam_supply_TA, Hydro_avail_TA, Hydro_supply_TA,
(Wholesale Market) Supply and Availability by fuel type	Bio_Other_avail_Mkt, Bio_Other_supply_Mkt, Coal_avail_Mkt , Coal_supply_Mkt Cogen_Total_avail_Mkt, Cogen_Total_supply_Mkt, Dual_Fuel_avail_Mkt, Dual_Fuel_supply_Mkt, Gas_CC_avail_Mkt, Gas_CC_supply_Mkt, Gas_SC_avail_Mkt, Gas_SC_supply_Mkt, Gas_Steam_avail_Mkt, Gas_Steam_supply_Mkt, Hydro_avail_Mkt, Hydro_supply_Mkt, Mkt_Outage, Solar_avail_Mkt, Solar_supply_Mkt, Storage_avail_Mkt, Wind_avail_Mkt, Wind_supply_Mkt	Bio_Other_avail_Mkt, Bio_Other_supply_Mkt, Coal_avail_Mkt , Coal_supply_Mkt Cogen_Total_avail_Mkt, Cogen_Total_supply_Mkt, Dual_Fuel_avail_Mkt, Dual_Fuel_supply_Mkt, Gas_CC_avail_Mkt, Gas_CC_supply_Mkt, Gas_SC_avail_Mkt, Gas_SC_supply_Mkt, Gas_Steam_avail_Mkt, Gas_Steam_supply_Mkt, Hydro_avail_Mkt, Hydro_supply_Mkt, Mkt_Outage, Solar_avail_Mkt, Solar_supply_Mkt, Storage_avail_Mkt, Wind_avail_Mkt, Wind_supply_Mkt	Bio_Other_avail_Mkt, Bio_Other_supply_Mkt, Coal_avail_Mkt , Coal_supply_Mkt Cogen_Total_avail_Mkt, Cogen_Total_supply_Mkt, Dual_Fuel_avail_Mkt, Dual_Fuel_supply_Mkt, Gas_CC_avail_Mkt, Gas_CC_supply_Mkt, Gas_SC_avail_Mkt, Gas_SC_supply_Mkt, Gas_Steam_avail_Mkt, Gas_Steam_supply_Mkt, Hydro_avail_Mkt, Hydro_supply_Mkt, Mkt_Outage, Solar_avail_Mkt, Solar_supply_Mkt, Storage_avail_Mkt, Wind_avail_Mkt, Wind_supply_Mkt	Bio_Other_avail_Mkt, Bio_Other_supply_Mkt, Coal_avail_Mkt , Coal_supply_Mkt Cogen_Total_avail_Mkt, Cogen_Total_supply_Mkt, Dual_Fuel_avail_Mkt, Dual_Fuel_supply_Mkt, Gas_CC_avail_Mkt, Gas_CC_supply_Mkt, Gas_SC_avail_Mkt, Gas_SC_supply_Mkt, Gas_Steam_avail_Mkt, Gas_Steam_supply_Mkt, Hydro_avail_Mkt, Hydro_supply_Mkt, Mkt_Outage, Solar_avail_Mkt, Solar_supply_Mkt, Storage_avail_Mkt, Wind_avail_Mkt, Wind_supply_Mkt
Wholesale Market Demand and Price Variables	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP
Imports and Exports	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import
Hour	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22
Temperature	Temp_Calgary, Temp_Edmonton	Temp_Calgary, Temp_Edmonton	Temp_Calgary, Temp_Edmonton	Temp_Calgary, Temp_Edmonton
AS Market	AS_Market_Output	AS_Market_Output	AS_Market_Output	AS_Market_Output

Table A.2: Set of Explanatory Variables used in Specification II

	Capital Power (CP)	ENMAX (EN)	Heartland (HT)	TransAlta (TA)
Own firm lagged price variables	CP_price_log_diffL6, CP_price_log_diffL9, CP_price_log_diffL12, CP_price_log_diffL15, CP_price_log_diffL18, CP_price_log_diffL21, CP_price_log_diffL24	EN_price_log_diffL6, EN_price_log_diffL9, EN_price_log_diffL12, EN_price_log_diffL15, EN_price_log_diffL18, EN_price_log_diffL21, EN_price_log_diffL24	HT_price_log_diffL6, HT_price_log_diffL9, HT_price_log_diffL12, HT_price_log_diffL15, HT_price_log_diffL18, HT_price_log_diffL21, HT_price_log_diffL24	TA_price_log_diffL6, TA_price_log_diffL9 TA_price_log_diffL12, TA_price_log_diffL15, TA_price_log_diffL18, TA_price_log_diffL21, TA_price_log_diffL24
Other own firm lagged variables	AS_Output_CP, Avaliablemw_CP, Coal_avail_CP, Coal_supply_CP, Dispatchedmw_CP, Gas_CC_avail_CP, Gas_CC_supply_CP, Gas_SC_avail_CP, Gas_SC_supply_CP, Wind_avail_CP, Wind_supply_CP	AS_Output_EN, Avaliablemw_EN, Dispatchedmw_EN, Gas_CC_avail_EN, Gas_CC_supply_EN, Gas_SC_avail_EN, Gas_SC_supply_EN, Wind_avail_EN, Wind_supply_EN	AS_Output_HT, Avaliablemw_HT, Cogen_Total_avail_HT, Cogen_Total_supply_HT, Dispatchedmw_HT, Dual_Fuel_avail_HT, Dual_Fuel_supply_HT, Gas_SC_avail_HT, Gas_SC_supply_HT, HT_Outage	AS_Output_TA, Avaliablemw_TA, Coal_avail_TA, Coal_supply_TA, Dispatchedmw_TA, Dual_Fuel_avail_TA, Dual_Fuel_supply_TA, Gas_Steam_avail_TA, Gas_Steam_supply_TA, Hydro_avail_TA, Hydro_supply_TA
(All firms) Supply and Availability by fuel type	Availablemw_EN, Available_HT, Available_TA, Bio_Other_avail_Fringe, Bio_Other_supply_Fringe, Coal_avail_TA, Coal_supply_TA, Cogen_NonZero_avail_Fringe, Cogen_Total_avail_HT, Cogen_Total_supply_Fringe, Cogen_Total_supply_HT, Cogen_Zero_avail_Fringe, Dispatchedmw_EN, Dispatchedmw_HT, Dispatchedmw_TA, Dual_Fuel_avail_HT, Dual_Fuel_avail_TA, Dual_Fuel_supply_HT, Dual_Fuel_supply_TA, Fringe_Outage_MWhs, Gas_CC_avail_EN, Gas_CC_avail_Fringe, Gas_CC_supply_EN, Gas_CC_supply_Fringe, Gas_SC_avail_EN, Gas_SC_avail_Fringe, Gas_SC_avail_HT, Gas_SC_supply_EN, Gas_SC_supply_Fringe, Gas_SC_supply_HT, Gas_Steam_avail_Fringe, Gas_Steam_avail_TA, Gas_Steam_supply_Fringe, Gas_Steam_supply_TA, HT_Outage_MWhs, Hydro_avail_Fringe, Hydro_avail_TA, Hydro_supply_Fringe, Hydro_supply_TA, Solar_avail_Fringe, Solar_supply_Fringe, TA_Outage_MWhs, Wind_avail_EN, Wind_avail_TA, Wind_supply_EN, Wind_supply_Fringe, Wind_supply_TA	Availablemw_CP, Available_HT, Available_TA, Bio_Other_avail_Fringe, Bio_Other_supply_Fringe, Coal_avail_CP, Coal_avail_TA, Coal_supply_CP, Coal_supply_TA, Cogen_NonZero_avail_Fringe, Cogen_Total_avail_HT, Cogen_Total_supply_Fringe, Cogen_Total_supply_HT, Cogen_Zero_avail_Fringe, Dispatchedmw_CP, Dispatchedmw_HT, Dispatchedmw_TA, Dual_Fuel_avail_HT, Dual_Fuel_avail_TA, Dual_Fuel_supply_HT, Dual_Fuel_supply_TA, Fringe_Outage_MWhs, Gas_CC_avail_CP, Gas_CC_avail_Fringe, Gas_CC_supply_CP, Gas_CC_supply_Fringe, Gas_SC_avail_CP, Gas_SC_avail_Fringe, Gas_SC_avail_HT, Gas_SC_supply_CP, Gas_SC_supply_Fringe, Gas_SC_supply_HT, Gas_Steam_avail_Fringe, Gas_Steam_avail_TA, Gas_Steam_supply_Fringe, Gas_Steam_supply_TA, Heart_Outage_MWhs, Hydro_avail_Fringe, Hydro_avail_TA, Hydro_supply_Fringe, Hydro_supply_TA, Solar_avail_Fringe, Solar_supply_Fringe, TA_Outage_MWhs, Wind_avail_CP, Wind_avail_TA, Wind_supply_CP, Wind_supply_Fringe, Wind_supply_TA	Availablemw_CP, Available_EN, Available_TA, Bio_Other_avail_Fringe, Bio_Other_supply_Fringe, Coal_avail_CP, Coal_avail_TA, Coal_supply_CP, Coal_supply_TA, Cogen_NonZero_avail_Fringe, Cogen_Total_supply_Fringe, Cogen_Zero_avail_Fringe, Dispatchedmw_CP, Dispatchedmw_EN, Dispatchedmw_TA, Dual_Fuel_avail_TA, Dual_Fuel_supply_TA, Fringe_Outage_MWhs, Gas_CC_avail_CP, Gas_CC_avail_EN, Gas_CC_avail_Fringe, Gas_CC_supply_CP, Gas_CC_supply_EN, Gas_CC_supply_Fringe, Gas_SC_avail_CP, Gas_SC_avail_EN, Gas_SC_avail_Fringe, Gas_SC_supply_CP, Gas_SC_supply_EN, Gas_SC_supply_Fringe, Gas_Steam_avail_Fringe, Gas_Steam_avail_TA, Gas_Steam_supply_Fringe, Gas_Steam_supply_TA, Hydro_avail_Fringe, Hydro_avail_TA, Hydro_supply_Fringe, Hydro_supply_TA, Solar_avail_Fringe, Solar_supply_Fringe, TA_Outage_MWhs, Wind_avail_CP, Wind_avail_EN, Wind_avail_TA, Wind_supply_CP, Wind_supply_EN, Wind_supply_Fringe, Wind_supply_TA	Availablemw_CP, Available_EN, Available_HT, Bio_Other_avail_Fringe, Bio_Other_supply_Fringe, Coal_avail_CP, Coal_avail_CP, Coal_supply_CP, Cogen_NonZero_avail_Fringe, Cogen_Total_supply_Fringe, Cogen_Zero_avail_Fringe, Dispatchedmw_CP, Dispatchedmw_EN, Dispatchedmw_HT, Dual_Fuel_avail_HT, Dual_Fuel_supply_HT, Fringe_Outage_MWhs, Gas_CC_avail_CP, Gas_CC_avail_EN, Gas_CC_avail_Fringe, Gas_CC_supply_CP, Gas_CC_supply_EN, Gas_CC_supply_Fringe, Gas_SC_avail_CP, Gas_SC_avail_EN, Gas_SC_avail_Fringe, Gas_SC_avail_HT, Gas_SC_supply_CP, Gas_SC_supply_EN, Gas_SC_supply_Fringe, Gas_SC_supply_HT, Gas_Steam_avail_Fringe, Gas_Steam_supply_Fringe, HT_Outage_MWhs, Hydro_avail_Fringe, Hydro_supply_Fringe, Solar_avail_Fringe, Solar_supply_Fringe, Wind_avail_CP, Wind_avail_EN, Wind_supply_CP, Wind_supply_EN, Wind_supply_Fringe,
Wholesale Market Demand and Price Variables	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP	Actual_AIL, Actual_Pool_Price, Coefficient_of_Variation, Day_Ahead_Pool_Price, Forecast_AIL, Std_SMP
Imports and Exports	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import	BC_Export, BC_Import, BC_ImportCapability, MT_EXPORT, MT_Import, MAT_ImportCabability, SK_EXPORT, SK_Import
Hour	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22	he_1, he_4, he_7, he_10, he_13, he_16, he_19, he_22
Temperature	Temp_Calgary, Temp_Edmonton	Temp_Calgary, Temp_Edmonton	Temp_Calgary, Temp_Edmonton	Temp_Calgary, Temp_Edmonton
AS Market	EN_AS_Output, Fringe_AS_Output HT_AS_Ouput, TA_AS_Output	CP_AS_Output, Fringe_AS_Output HT_AS_Ouput, TA_AS_Output	CP_AS_Output, EN_AS_Output Fringe_AS_Ouput, TA_AS_Output	CP_AS_Output, EN_AS_Output Fringe_AS_Ouput, HT_AS_Output

Appendix B Additional Results

Figure B.1: Price level forecast - Lasso

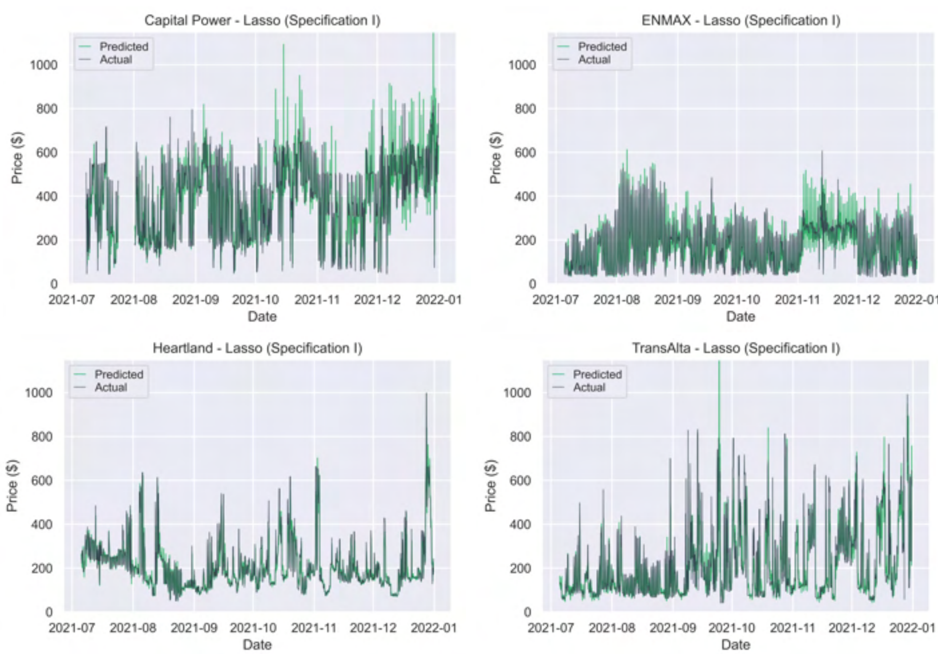

Figure B.2: Histogram - Lasso

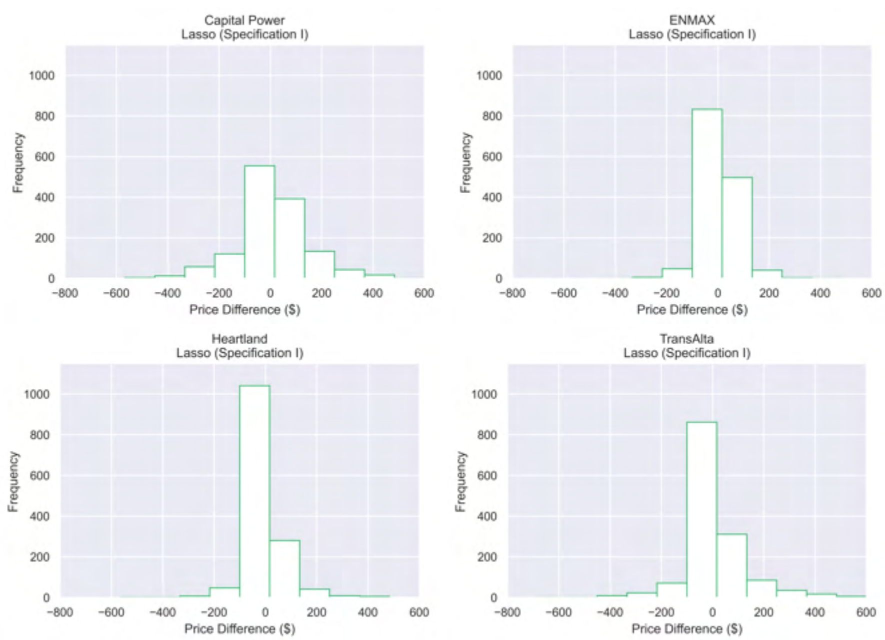

Figure B.3: Box Plot of Grouped Permutation Importance - XGBoost

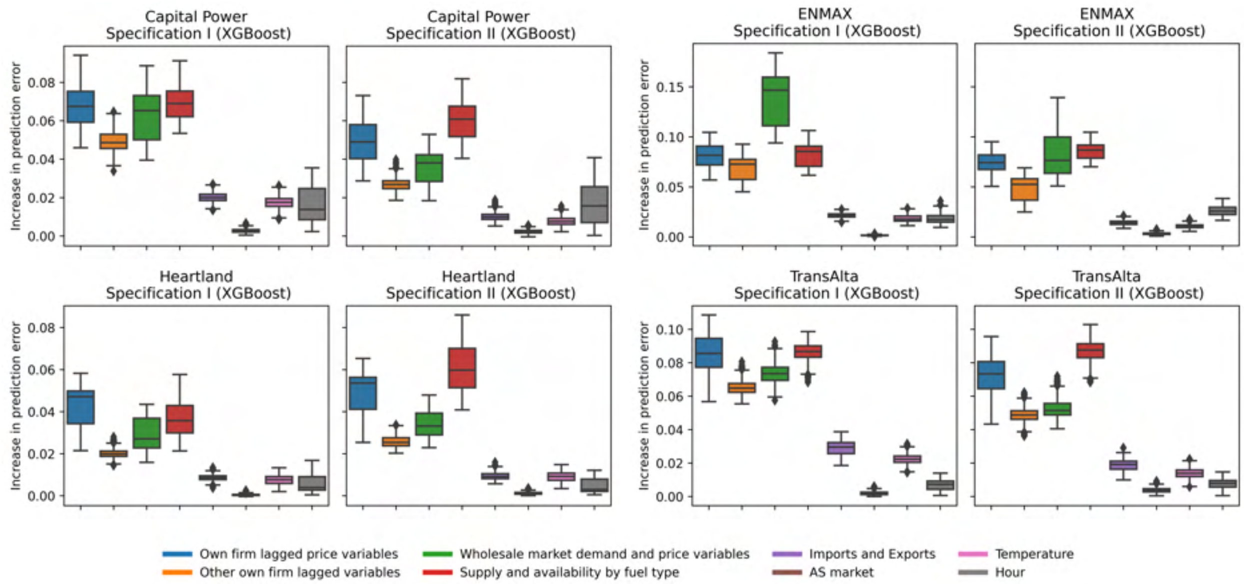

Figure B.4: Box Plot of Grouped Permutation Importance - Ridge

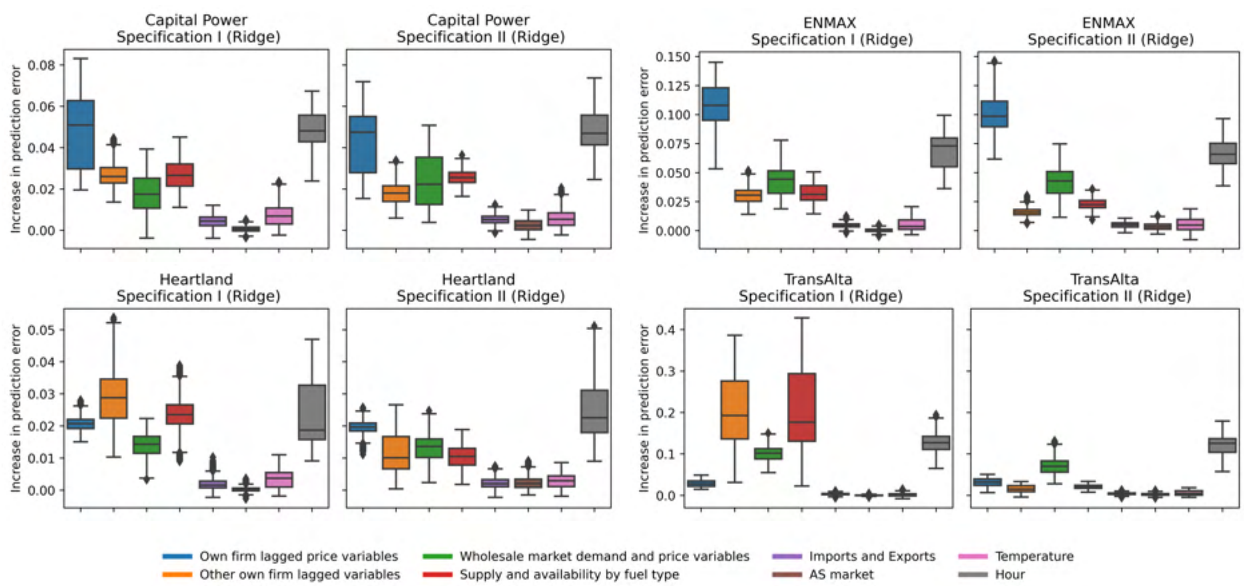

Figure B.5: Box Plot of Top 10 Explanatory Variables - XGBoost

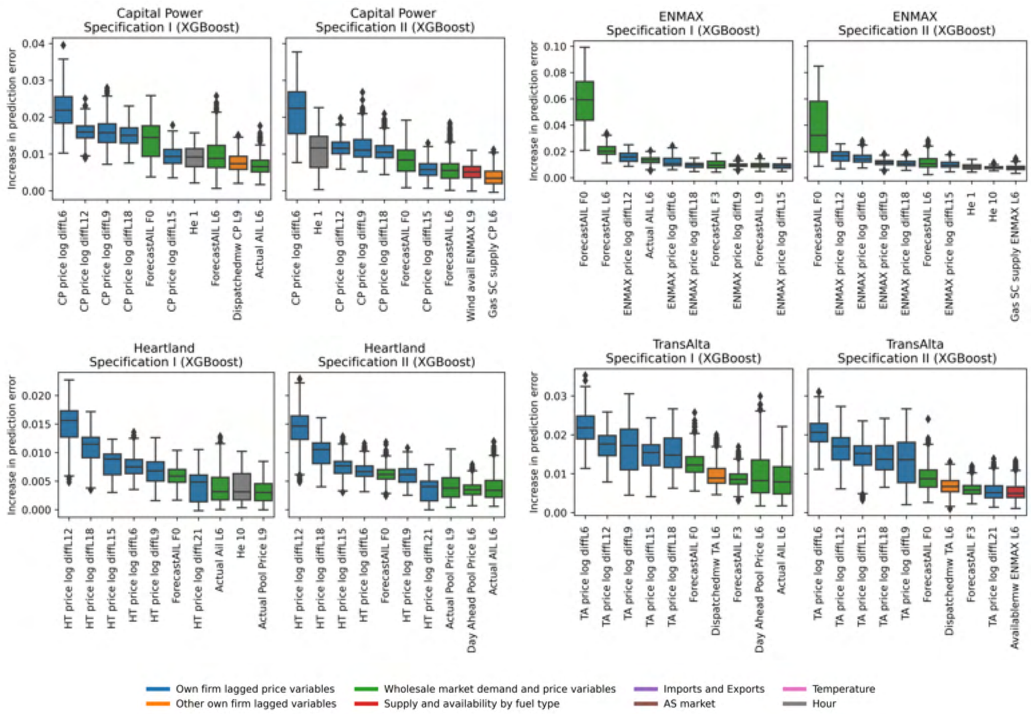

Figure B.6: Box Plot of Top 10 Explanatory Variables - Ridge

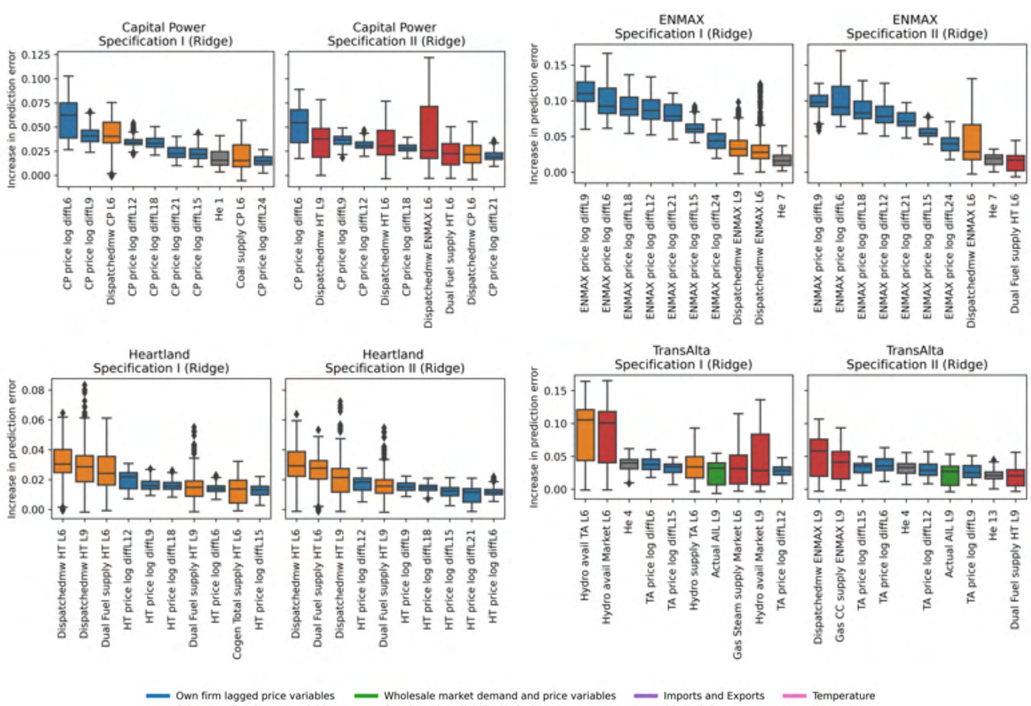

Figure B.7: Box Plot of Finer Group Analysis

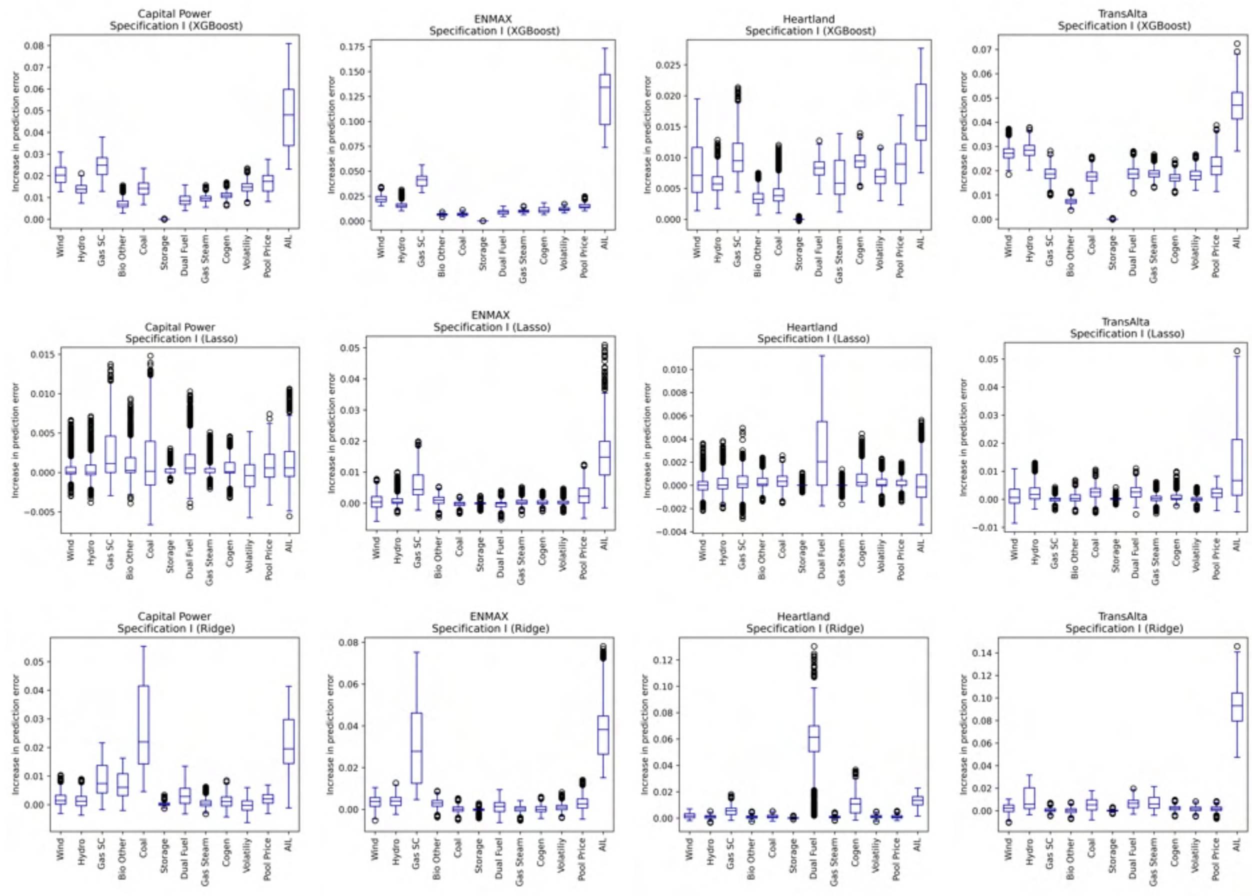

Appendix C Marginal Cost Estimates and Assets Used to Construct Price Variables

To construct our price variable to be predicted, we identify for each firm those generating assets whose capacity is offered into the market at prices above short-run marginal cost. To estimate unit-level marginal cost functions, we use daily natural gas price data from Alberta's Natural Gas Exchange (NGX) and weekly Powder River Basin coal price data provided by the U.S. Energy Information Administration to serve as our fuel input prices.[29] Generation unit-level efficiency parameters were acquired from documents published by the Alberta Market Surveillance Administrator, Alberta Utilities Commission, and CASA (2004).

We estimate the marginal cost of natural gas and coal units using the summation of fuel costs, costs of environmental compliance, and variable operating and maintenance costs. The fuel input costs are determined by the fuel price (natural gas or coal) multiplied by the unit-specific efficiency (heat rate) parameters.[30] Wind and solar units are assumed to have zero marginal cost. Further, a sizable portion of Alberta's market consists of cogeneration facilities that generate electricity as a by-product of an on-site industrial process (e.g., oil sands). We set the marginal cost of generation at their bid price, which is systematically at a price of \$0/MWh. There are a small number of biomass units in the province that are operated by small fringe producers and have an on-site industrial process (e.g., forestry, pulp mill). We assume that the small fringe of producers behave as perfectly competitive producers and their bids reflect their marginal costs. Similar to the cogen facilities, these units often bid a price of \$0/MWh. Finally, there are several hydroelectric facilities that represent a small portion of the market's output (approximately 2% in 2021). The marginal cost of these units reflects the opportunity cost of the stored energy. Hydro units often bid their assets in at a price of \$0/MWh or at high prices in excess of \$900/MWh. These high bids reflect in part regulatory and ecological constraints. We set the marginal cost of these units to be equal to their bids. We anticipate the possible biases from this assumption to be minimal given these units reflect a small portion of the market.

A descriptive analysis of asset-level price-cost markups indicates that the following assets for each large firm are included as "strategic" assets:

29. Powder River Basin coal price data can be accessed at: https://data.nasdaq.com/data/EIA/COAL-us-coal-prices-by-region. NGX gas price data were provided to the authors by the Alberta Market Surveillance Administrator.

30. For a more detailed discussion of the estimation of unit-level marginal cost, see Brown, Eckert, and Shaffer (2023).

- TransAlta: Keephills 1 & 2, Sundance 4 & 6, Sheerness 1 & 2
- Heartland: Battle River 4 & 5, Poplar Hill 1, Sheerness 1 & 2, Valleyview 1 & 2
- Capital Power: Shepard, Cloverbar 1-3, Genesee 1 - 3
- ENMAX: Shepard, Summit-Crossfield 1-3, Cavalier, Calgary Energy Centre

Notice that a subset of the assets appear for multiple firms. This arises because these units are co-owned with two firms having offer control over a certain percentage of the asset's capacity.