Price and Quantity Discovery without Commitment*

Stefan Bergheimer[†], Estelle Cantillon[‡] and Mar Reguant[§]

June 2023

Abstract

Wholesale electricity markets solve a complex allocation problem: electricity is not storable, demand is uncertain, and production involves dynamic cost considerations and indivisibilities. The New Zealand wholesale electricity market attempts to solve this complex allocation problem by using an indicative price and quantity discovery mechanism that ends at dispatch. Can such a market mechanism without commitment provide useful information? We document that indicative prices and quantities are increasingly informative of the final prices and quantities and that bid revisions are consistent with information-based updating. We argue that the reason why the predispatch market is informative despite the lack of commitment is that it generates private benefits in terms of improved intertemporal optimization of production plans.

KEYWORDS: electricity markets, price discovery, pre-play communication, non-trading mechanisms, coordination, intertemporal optimization.

JEL Codes: D47, D83, G14

*We are grateful to Phil Bishop and the entire team at the New Zealand Electricity Authority for giving us access to the data and sharing their insights into the market. Aislinn Carty, Tianyu Luo, Almudena Valle, and Dorian Weber provided excellent research assistance. We thank the editor, David Genesove, an anonymous referee, seminar audiences and especially Gani Aldashev, Bruno Biais, Giacomo Calzolari, Julien Combe, Peter Cramton, Michele Fioretti, Chiara Fumagalli, David Goldreich, Robin Lee, Patrick Legros, Andras Niedermayer, Ariel Pakes, Mathias Reynaert, Marzena Rostek, Otto Toivanen, and Iivo Vehviläinen for their feedback and questions that helped us crystalize our thinking. Reguant acknowledges the support of NSF grant SES-1455084. This article has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 101001732-ENECML).

[†]Compass Lexecon. Email: stefanbergheimer@gmail.com.

[‡]Université Libre de Bruxelles, FNRS, and CEPR. Email: Estelle.Cantillon@ulb.be (Corresponding author).

[§]Northwestern University, Barcelona School of Economics, NBER, and CEPR. Email: mar.reguant@northwestern.edu.

1 Introduction

Markets and prices play an important coordinating role in our economies. They direct producers where there is demand, help consumers optimize, and, more generally, foster efficient allocation of resources. This coordination role is especially crucial in wholesale electricity markets where demand and supply are uncertain and largely inelastic in the short run, and yet, because electricity is difficult to store, demand and supply must be balanced at all times to avoid system outages.

Existing electricity markets typically solve this problem by organising a sequence of markets (typically week-ahead, day-ahead, and real-time) to coordinate supply and demand. The markets gradually lock in demand and supply and reduce uncertainty for market participants. New Zealand is an exception. The market is only called once, and all physical allocation decisions are based on bids submitted one hour prior to dispatch. To nevertheless help market participants coordinate supply and demand, the New Zealand electricity system operator organises a series of indicative markets (called predispatch) starting 36 hours before dispatch where participants can submit and update bids freely, and indicative prices and quantities are produced on a regular basis. Only the last bids submitted are used for the final allocation.

Can markets without commitment, such as the New Zealand electricity market, foster efficient price and quantity discovery? We examine bidding behaviour and bid revisions in the predispatch market. Our data spans 4 years and over 80,000 trading periods, each with a predispatch market called 24 times before actual dispatch. We observe individual market participants' bids and revisions, indicative prices and quantities during predispatch, and final allocations.

We provide evidence that indicative prices and quantities are increasingly informative of final prices and quantities and that bid revisions are consistent with information-based updating. On average, 7% of generation is reallocated during predispatch as a result of these bid revisions. Prices increase very slightly (less than 1%) and become less volatile over the course of the predispatch. We provide suggestive evidence that the predispatch market facilitates generation coordination across trading periods and therefore acts as a complement to the otherwise static (single period) allocations produced by the New Zealand electricity market model.

The New Zealand predispatch market is an example of what is called an iterative mechanism, a market organisation that allows participants to update their bids based on feedback about the ongoing price before allocations are finalized. Iterative mechanisms are credited with at least three advantages. First, iterative mechanisms facilitate and support participants' decision-making. Decisions typically take the form of whether to stay in or drop-out or adjust a bid at the margin, and participants receive direct feedback on how their choices impact their allocation.

Second, iterative mechanisms elicit and aggregate private information, which can foster competition. This is the famous "linkage principle" first identified by Milgrom and Weber (1982) and generalized to multi-unit auctions by Ausubel (2004). The insight here is that iterative mechanisms generate information that helps participants update their estimates of costs or value and protect them from the winner's curse, a phenomenon that typically holds back aggressive bidding.

A third advantage of iterative mechanisms is that they help market participants optimize their allocation when the bidding language is not rich enough to capture underlying costs and preferences. Nisan and Segal (2006) have characterized the communication requirements of efficient allocations in the presence of nonconvex preferences and indivisible goods. In electricity markets, these correspond to fixed start-up costs, ramp-up and ramp-down production constraints, and unit commitment (see e.g. Reguant (2014) for evidence). Nisan and Segal (2006) show that the number of prices needed, and therefore the complexity of the required bidding language, grows exponentially with the relevant states of the world. Iterative mechanisms, which run parallel markets for commodities that are related from the participants' perspective, partially overcome this curse of dimensionality (Ausubel and Cramton, 2004). They have been used, for example, by EDF to auction generation capacity in France and by the US Federal Communication Commission to auction spectrum.

Participants in electricity markets are typically sophisticated and well-informed. Moreover, the level of market transparency in New Zealand is particularly high. So the first two advantages of iterative mechanisms we have described are unlikely to be first order in the context of the New Zealand electricity market. In contrast, Nisan and Segal (2006)'s findings imply, in particular, that market participants in electricity markets should be able to condition their allocation in one trading period on their allocation in some other trading periods, something that the New Zealand market model does not allow. This provides a rationale for an iterative mechanism such as the predispatch.

To support genuine price discovery, iterative mechanisms often include an activity rule designed to curb manipulative bids. Participants can revise their bids but cannot make a "worse" offer (where what "worse" means depends on the specific context). Alternatively, some mechanisms include a random end-time that ensures that bidders are committed to their bids. What's remarkable about the New Zealand predispatch market is that it does not contain any such form of commitment.[1] This means that bids during the predispatch can be seen as "cheap talk".

Another example of iterative mechanisms without commitment are preopening periods at stock exchanges. During preopening, traders submit and freely revise their offers during a certain period, until the market is called and the produced price serves as the opening price for the regular market. The existing literature documents that such markets *are* informative despite the lack of commitment (see e.g. Biais et al., 1999, Cao et al., 2000, and Barclay and Hendershott, 2008). The reasons proposed all include a reduction in adverse selection due to either getting access to a larger pool of liquidity at the opening of the regular trading day (Biais et al., 1999) or information-sharing (Hong and Pouget, 2021).[2]

We too document that prices and quantities are increasingly informative of final prices and quantities despite the lack of commitment. However, our proposed explanation for why this happens has nothing to do with adverse selection, which is nonexistent in the New Zealand

[1] We will argue in Section 2 that forward markets in electricity do not provide a substitute for commitment during the predispatch.

[2] This does not mean that commitment may not be valuable nevertheless (see e.g. the experimental evidence reported in Biais et al., 2014).

electricity market, but with the ability of iterative mechanisms to coordinate allocations across several markets, in our case, across several trading periods. We document that virtually all bid revisions involve several trading periods and provide examples illustrating how market participants use the predispatch market to reorganize their generation dispatch across time.

Finally, a natural concern about the New Zealand predispatch market is that it may foster tacit collusion, or at least facilitate the exercise of market power. Tacit collusion arises when market participants coordinate on a less competitive equilibrium without explicit communication or enforcement mechanisms. Markets with multiple equilibria are prone to tacit collision. Bolle (1992) and Ausubel et al. (2014) show that markets, such as the New Zealand electricity market, where participants submit supply functions typically have multiple equilibria. In such markets, high levels of market transparency can help participants coordinate on the least competitive equilibrium (see von der Fehr (2013) for a general argument, and Brown and Eckert (2022) for evidence in the Alberta wholesale electricity market).

We make no claims as to the nature - collusive or not - of the equilibrium in the New Zealand wholesale electricity market. After all, like most other electricity markets, the New Zealand electricity market is characterized by a small number of participants, repeated interactions, inelastic demand, and limited informational asymmetry, all conditions that are favorable to the emergence of collusion, or at least tacit collusion.

However, our results do not provide any indication that the predispatch market facilitates participants' coordination on a high price equilibrium: prices barely rise over the course of the predispatch and bid revisions are largely driven by new information arrival. Moreover, the exact mechanism through which predispatch could incrementally support such coordination is unclear. In a recent paper, Kamada and Kandori (2020) study what they call "revision games" where players can update their actions repeatedly until the market is called. The predispatch market can be seen as a revision game. They show that even a small probability of not being able to upgrade their action (in our context, market participants missing the deadline for submitting their bids) can help support coordination on a less competitive equilibrium. However, in their setting, play evolves over time to an increasingly competitive outcome as players update their actions. This is not borne out in our data: if anything prices increase very slightly (by less than one percent) over the course of the predispatch.

In a separate paper (Bergheimer et al., 2023), we explore the impact of the reduced uncertainty produced by the predispatch market on participants' ability to *unilaterally* exert market power. We find that the reduced uncertainty about residual demand facilitates the exertion of market power by market participants and, in particular, by hydro-based generators, who can easily reallocate their intraday production to take advantage of price differences.

2 The New Zealand predispatch market for electricity

The current organization of the electricity market in New Zealand is the result of the economy-wide pro-market reforms that swept the country in the 1980s and 1990s. Before that, electricity production, transmission, distribution, and retail were all under public ownership and vertically integrated. Transmission was separated early on through the creation of Transpower, which today acts as the system operator. Between 1996 and 1999, the generating assets of the monopoly generation company were progressively split to make way for five independently-operated firms: Contact, Genesis, Meridian, Mighty River Power (now Mercury), and Trust Power (now Manawa Energy). Distribution and retailing were also separated at that time, with the five incumbent generation companies inheriting the retail business of the former electricity monopoly.

The foundations for the electricity wholesale market were laid out in 1996. The design relies on a single settlement (dispatch) market where energy and reserves are co-optimized on the basis of the bids received by electricity producers and industrial consumers. The dispatch algorithm maximizes the area between the demand and supply curves, taking physical constraints into account.[3] Co-optimization means that dispatch may occasionally deviate from the cost-minimizing dispatch when doing so reduces the cost of reserves. Prices are nodal, i.e. location-specific, reflecting the geography of the country and the large transmission losses that go with it.

Participation in the wholesale market is compulsory for all electricity producers and industrial consumers, including vertically-integrated firms. Electricity producers are asked to submit energy and reserve bid schedules, i.e. step functions that describe quantities offered at each price, as well as maximal capacity and ramp-up and ramp-down constraints for each half-hour (trading period).[4] Bid schedules are specific to units or stations. Industrial consumers are also requested to submit bid schedules. (For simplicity, we will use the term "bid" as short-hand notation for bid schedule in the rest of the paper.)

The market is cleared for each trading period sequentially. This means that the only dynamic consideration that the market model takes into account is ramp-up and ramp-down constraints from the previous trading period. It is preceded by a predispatch market, that opens 36 hours before dispatch, where indicative prices and quantities are generated using exactly the same inputs and optimization model as for the final dispatch. Specifically, every 2 hours, Transpower runs the model over a 72-period horizon (the so-called long schedule) and, every half hour, it runs it over an 8-period horizon (short schedule). This means that, for every trading period, 24 indicative predispatch markets are run (initially at the frequency of once every two hours, then once every half an hour) prior to final dispatch.

During the predispatch, market participants can update these bids at their will until "gate closure" which happens one hour before dispatch (two hours before dispatch until June 28, 2017). After gate closure, restrictions apply. Intermittent generators (wind) and industrial consumers are expected to update their estimates of their generation and load after gate closure. Other market participants - generators - can only update their bids in exceptional circumstances (e.g.

[3] Retail demand does not participate actively in the wholesale market. Forecast load is instead used when producing dispatch instructions. Alvey et al. (1998) describe the model used for scheduling, pricing, and dispatch (SPD) in detail.

[4] A bid schedule can have up to 5 price bands for generation bids, and up to 10 price bands for demand-side bids.

an unplanned outage) and can only change the quantity offered, not the price, starting from the highest price band. Commitment is limited. Wind generators and industrial consumers are never bound by their bids. Other generators are only committed by their last bids.

Real-time dispatch uses the latest load forecast, current generation, and the last submitted bids during predispatch as inputs to the market model and generates dispatch instructions at the frequency of once every 5 minutes.

The wholesale market is complemented by a voluntary hedge market, where market participants take positions either on the Australian Security Exchange (ASX) (mostly) or in the over-the-counter market. This hedge market operates on a very different time horizon. Around 97% of traded contracts are monthly or quarterly contracts that cover all trading periods, or all peak trading periods, in a given month or quarter. This means that this hedge market does not provide a substitute for the lack of commitment in the predispatch market.

Market transparency is promoted at all stages. After each predispatch, price forecasts, load forecasts and the aggregate supply curve at reference nodes are published. In addition, individual market participants are informed about their cleared bids. The entire history of bids and offers during predispatch and all inputs to the final dispatch and pricing are published within two or three days. A website, WITS (which stands for Wholesale Information Trading System), provides real-time information about the state of the market and operational constraints, including prices, load, generation, outages, transmission constraints and flows between the North and South islands. Separate websites provide anonymized data on hedging positions and information about hydro reserves.

3 Data

Our data span the period between 1 January 2014 and 30 September 2018. For each half hour (trading period) and each node,[5] we observe the bidding behavior of market participants during the predispatch market, indicative prices and quantities generated by each predispatch, and final prices and quantities.[6,7] We additionally observe all public market-relevant information such as installed capacity at all nodes, load forecasts, planned and unplanned outages for each node and trading period, hourly regional weather realizations, and daily levels of hydro reservoirs.

Table 1 provides an overview of our bidding data and outcomes at different stages of the market. By default, each generator and each industrial consumer must submit a bid schedule at the time of the first predispatch round for all the nodes at which they are active. The top panel of Table

[5]With some slight abuse of language, we call a node, not only the physical injection or exit point on the grid but every unique "node x bidding unit" observation. When two participants are active at a physical node, they face the same price, but their behavior may still differ. Likewise, several generating units owned by the same participant may be connected to the grid at the same physical node while submitting different bids.

[6]Transpower solves two versions of its program, one in which bids from industrial consumers are taken into account as submitted (the so-called price-responsive schedule), and one where they are replaced by a vertical demand at their maximum demanded quantity (non-responsive schedule). For the purpose of our paper, we use the price-responsive schedules for the predispatch data.

[7]Predispatch data are missing for 2,083 trading periods (2.5%) so our final dataset covers 81,146 trading periods. There is no indication of systematic bias in this censoring.

Table 1: Bidding behavior and market outcomes during predispatch

	# nodes	5%	25%	50%	75%	95%
Generation bids per trading period	77	1	1	3	5	11
Demand bids per trading period	34	1	1	1	2	6
Prices (NZ$/MWh)						
First predispatch price	89	9.5	43.5	57.8	79.0	158.6
Last predispatch price	89	23.2	47.4	58.3	75.1	123.2
Change from first to last predispatch	89	-58.3	-13.5	0.1	14.1	44.2
Generation (MWh)						
First predispatch quantity	89	3,392	3,966	4,797	5,307	6,061
Last predispatch quantity	89	3,386	3,964	4,781	5,289	6,043
Reshuffling from first to last predispatch	89	0.04	0.06	0.07	0.10	0.14
Industrial consumers (MWh)						
First predispatch quantity	34	943	989	1,017	1,047	1,094
Last predispatch quantity	34	885	935	970	1,004	1,052
Change from first to last predispatch	34	-132	-81	-47	-15	27
Reshuffling from first to last predispatch	34	0.01	0.02	0.04	0.05	0.08

Notes: The unit of observation for bids is a trading period x node. Generation bids exclude bids from wind units. Time stamps uniquely define bids. The unit of observation for prices and quantities is a trading period. Nodal prices are quantity-weighted to produce an average price for the trading period. There are 81,146 trading periods with complete coverage for predispatch prices and quantities between 1 January 2014 and 30 September 2018. Last predispatch prices and quantities use the last bids submitted and updated load and wind forecast at the beginning of the trading period. Reshuffling refers to the sum of node-level absolute changes in quantities, divided by two, and normalized by the total generation (resp. total industrial load) for that trading period.

1 shows that generation is the active side of the market during predispatch: bidding on the demand side mostly sticks to the minimum level of activity, whereas the median producer submits three different bid schedules for the same node and trading period over the course of the predispatch.

The second panel of the table shows that prices go up by less than 1% over the course of the predispatch and that their dispersion goes down.

The third panel documents generation. Aggregate scheduled generation barely changes over the course of the predispatch (median change of 12 MWh, less than 0.25% of aggregate generation) but, nevertheless, around 7% of generation is reshuffled, where reshuffling in a trading period is defined as the sum of absolute node-level changes in scheduled generation between the first and last predispatch, divided by twice total generation that trading period.

The last panel shows that industrial consumers account for approximately 20% of electricity consumption. Their scheduled consumption does not change much over the course of the predispatch, and when it does, it mostly goes down, in line with the observation that submitted demand schedules by industrial consumers adjust quantities downward for very high prices, but are essentially vertical otherwise. About 4% of industrial load gets reshuffled over the course of the predispatch.

Table 2 provides summary statistics for realized generation for each technology. Schedulable hydro accounts for 55% of electricity generation on average. It is followed by geothermal and gas (combined cycle), with 18% and 10% generation share, respectively.

Table 2: Characteristics of realized generation, by technology

	#	% active nodes		Generation (MWh)				Gen. share
	nodes	off-peak	peak	Mean	SD	Min	Max	Mean
Coal	4	0.25	0.33	117.0	130.0	0.0	500.0	0.02
Cogeneration	8	0.92	0.94	200.0	42.0	40.0	374.0	0.04
Diesel	1	0.00	0.01	0.0	4.0	0.0	156.0	0.00
Gas, combined cycle	3	0.62	0.63	460.0	195.0	0.0	1,139.0	0.10
Gas, open cycle	5	0.14	0.38	87.0	98.0	0.0	391.0	0.02
Geothermal	12	0.94	0.94	798.0	64.0	0.0	888.0	0.18
Hydro, run of river	8	0.79	0.91	176.0	66.0	26.0	319.0	0.04
Hydro, schedulable	36	0.77	0.89	2,604.0	658.0	917.0	4,438.0	0.55
Wind	11	0.89	0.89	247.0	138.0	0.0	589.0	0.05

Notes: The unit of observation is a trading period (N = 81,146). A trading period is considered to be a peak trading period when generation in that trading period belongs to the top 10 percentile of generation observed in the sample. One generation node consisting of a battery is excluded.

Technologies differ in their production profiles. Given its importance in the New Zealand electricity mix, hydro generation is active both off-peak and on-peak (with a smaller proportion of nodes active during off-peak time). Cogeneration and geothermal are two baseload technologies with little variation in generation levels and stable production patterns independent of the time of the day, as witnessed by the stable fraction of nodes active both during peak and off-peak times and the low standard deviation relative to mean generation. Wind generation is highly variable, but its production profile is independent of the state of demand. Finally, thermal production varies considerably and, except for combined cycle, increases during peak times, reflecting their cost structure and the role that these technologies play in the electricity generation mix.[8]

4 The role for a market

Economists since Hayek have valued markets, and the prices they generate, for their ability to enable "rapid adaptation to changes in the circumstances of time and place" facing decentralized economic agents (Hayek, 1945, p. 524). In this section, we quantify the residual uncertainty about final allocations that prevails in the system 36 hours before dispatch. After all, if this is minimal, there is little role for a market, be it with or without commitment. Central planning would do.

One source of residual uncertainty stems from wind and load forecasts. The system operator Transpower produces load forecasts for all nodes not participating in the wholesale market (essentially retail nodes). These are used as inputs to the predispatch and dispatch models. In addition, windmill operators have a *bona fide* obligation to submit accurate generation forecasts and must update those at least once every 30 minutes within 2 hours of dispatch. Table 3 shows that wind generation and load tend to be overestimated at the time of the first predispatch.

[8]The main economic difference between a combined cycle gas turbine and an open cycle gas turbine is their cost structures, with open cycle gas turbines being more expensive to operate.

Table 3: Load and wind forecast errors

	5%	25%	Median	75%	95%
Demand forecast errors (MWh)	-230	-86	-11	60	218
Demand forecast errors relative to market size	-0.05	-0.02	0.00	0.01	0.05
Wind forecast errors (MWh)	-142	-61	-16	30	112
Wind forecast errors relative to market size	-0.03	-0.01	0.00	0.01	0.02
Net forecast errors (MWh)	-248	-93	3	96	274
Net forecast errors relative to market size	-0.05	-0.02	0.00	0.02	0.06

Notes: The unit of observation is a trading period ($N = 81,146$). Forecast errors for load are computed as the difference between forecast load in the last and first predispatch. Wind forecast errors are computed as the difference between the final forecast at dispatch time and the wind forecast for the first predispatch. The first predispatch is used as reference to compute the relative numbers.

As these effects go in opposite directions, their net effect is symmetric around zero, with most observations falling within 6% of the actual market size.

Other sources of uncertainty include short-run changes in production and transmission circumstances. To systematically explore how these, together with uncertainty about load and wind generation, impact actual allocation of generation, we apply machine learning techniques to predict prices and generation at each production node based on information available before bidding starts. Any discrepancy between our best prediction and the observed allocation provides a measure of the residual uncertainty about final allocations that prevails 36 hours before dispatch.

We consider two broad sets of models: LASSO penalized regressions and random forests. The models include as predictors very much the same kind of information that Transpower uses to predict load (weather, seasonal, week and hour-of-the-day variables, lagged dependent variables) as well as node-specific information and system-level information about generation, such as hydraulic information about reservoirs and outage status.[9] Importantly, the model relies only on information available 36 hours before dispatch. Generation and prices are predicted at the node level. We focus on average predictions for prices due to higher volatility in prices. We split the sample into training and testing observations and provide summary statistics on the performance of predictions on the testing set.[10]

Table 4 summarizes the results from these prediction models. We compare the predictive power of three models: forest (F), LASSO (L), and the first predispatch (P). The predictive power of the models is evaluated based on the mean absolute errors (MAE) between the prediction and the outcome of the final predispatch, root mean squared error (RMSE), and R^2. To make the

[9] It is worthwhile to note that the exercise we carry out is different from the one that Transpower solves. Transpower seeks to predict load, which it uses, *together with the bids submitted by market participants*, as an input to the New Zealand electricity market model to pin down generation dispatch. What we are doing is exploring to what extent we can bypass the market (i.e., participants' bids) and predict final generation allocations based on information available 36 hours before dispatch.

[10] Our training sample uses 20% random days. Lagged variables are generated before taking the training sample, so that the random sample already contains all necessary variables. One aspect that departs from what one would do in practice is that we train the model using a random draw of the entire period, i.e., from 2014-2018. As an alternative approach, we could have more explicitly considered a model that only uses past data. In previous iterations of the prediction model, we estimated some models for 2017 with only past data and we obtained similar results.

MAE and RMSE more comparable across technologies, we normalize the MAE and the RMSE by the average generation of each technology.

The top panel describes the results for the quantity predictions for each technology. The results indicate large differences across technologies in terms of predictability. Hydro, combined cycle gas, and geothermal generation are highly predictable, as implied by the small MAE and RMSE, and high R^2. While hydro and combined cycle units adjust significantly during the predispatch process, a large share of their generation can be predicted. Open cycle gas and wind generation are less predictable, which is intuitive given that open cycle gas units tend to manage last-minute changes in dispatch as a much larger share of their output and the fact that wind is intermittent. For example, for open gas units the MAE can be as high as 45%. For wind, errors are estimated to be around 20-30% of average production in the first predispatch, which are not uncommon forecast errors for wind generation. The specification of random forests performs better than LASSO for all technologies, with lower MAE and RMSE, and higher R^2. Our predictions also tend to outperform predictions produced by the first predispatch, particularly the forest models.

The bottom panel summarizes the results for quantity-weighted prices at the island level. Prices are less predictable than generation for most technologies, and the first predispatch can be substantially noisier than a prediction model. This is in part due to the fact that the first predispatch provides a first pass at the market dispatch, and firms re-adjust their bids, reducing the volatility in prices, as already shown in Table 1.

While the predictions models can get an average prediction of market outcomes, there is still substantial uncertainty. Furthermore, the first predispatch appears to be noisier. This suggests that there is room for updating between the first and the last predispatch. While some of this reallocation can be predicted, not all of it can. There remains a good amount of uncertainty about final allocations 36 hours before dispatch. This provides a role for informative price signals to further coordinate supply and demand.[11] Comparing the numbers in Table 3 and Table 4, this uncertainty is not limited to the aggregate level of generation needed (load and wind uncertainty) but also to the allocation of generation across units.

5 Price and quantity discovery during the predispatch

In this section, we document the process of price and quantity discovery during predispatch. We first show that predispatch prices and quantities are increasingly informative of final prices and quantities. We then explore the extent to which predispatch prices reflect developments inside the predispatch market rather than developments in the contemporaneous spot prices. The findings support the hypothesis that information arises from the predispath market itself. We then zoom in on bid revisions and document that (1) bid revisions are more frequent during the short schedule part of the predispatch, and that (2) everything else equal, bid revisions are more frequent when new information arises (wind and load forecast revisions, new outage announcements).

[11]For another perspective on the value of markets to allocate electricity production, see Cicala (2022).

Table 4: Residual uncertainty about generation allocation and prices

	MAE			RMSE			R2		
	F	L	P	F	L	P	F	L	P
Technology-level generation									
- Hydro, run of river	0.07	0.10	0.12	0.10	0.13	0.18	0.95	0.92	0.86
- Hydro, schedulable	0.03	0.04	0.05	0.04	0.05	0.07	0.98	0.97	0.95
- Gas, combined cycle	0.07	0.14	0.12	0.10	0.20	0.21	0.97	0.89	0.87
- Gas, open cycle	0.29	0.45	0.45	0.43	0.61	0.74	0.87	0.74	0.65
- Geothermal	0.01	0.03	0.02	0.02	0.04	0.03	0.99	0.95	0.97
- Coal	0.21	0.42	0.34	0.34	0.58	0.66	0.92	0.77	0.72
- Wind	0.16	0.25	0.26	0.21	0.31	0.34	0.87	0.73	0.73
Average prices									
- North Island	0.14	0.19	0.32	0.25	0.30	0.52	0.73	0.60	0.35
- South Island	0.15	0.20	0.33	0.24	0.30	0.63	0.78	0.66	0.33

Notes: F = forest predictions, L = LASSO predictions, P = first predispatch. The results only include testing observations. MAE stands for mean absolute errors between the prediction and last predispatch, normalized by mean quantities at the technology level, RMSE stands for root mean squared error, also normalized by mean quantities. The top 0.1% observations with high prices (either final price or predispatch) are censored due to the importance of outliers driving the R^2 measure.

5.1 Convergence and increasing informativeness of predispatch prices and quantities

The first part of our empirical exploration into price and quantity discovery in the predispatch market builds on an approach first implemented by Biais et al. (1999) to study price discovery in the preopening period of the Paris stock exchange.

Let q^r_{nt} denote the indicative quantity for node n at time period t produced during the r^{th} round of the predispatch, with the convention that q^0_{nt} is the best forecast based on information available before the start of the predispatch.

For every predispatch round r, we regress the quantity revision over the entire predispatch on quantity revisions up to round r:

$$q^{24}_{nt} - q^0_{nt} = \alpha_r + \beta_r(q^r_{nt} - q^0_{nt}) + \varepsilon_{rnt}. \tag{1}$$

If predispatch quantities are uninformative, then β_r should be equal to zero. Conversely, if predispatch quantities are informative, β_r should be positive. If predictions are unbiased but noisy, we would expect β_r to be between zero and one. As predictions become more accurate, β_r should converge to one.[12]

Regressions are carried out separately for each predispatch round to account for the non-stationary process of learning during the predispatch. We are thus comparing the same predispatch round, for different production nodes and trading periods. We run the equivalent

[12] Another way to see this is to rewrite (1) in the mathematically equivalent, but statistically less convenient, equation $q^{24}_{nt} = \alpha_r + \beta_r q^r_{nt} + (1 - \beta_r)q^0_{nt} + \varepsilon_{rnt}$. β_r can then be interpreted as the weight of round r's indicative quantity in predicting final quantity.

Figure 1: Evidence for quantity and price discovery

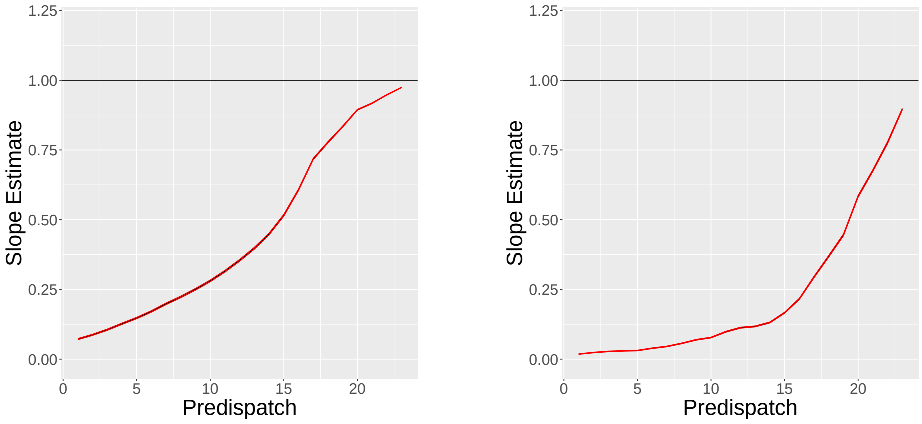

(a) Quantity convergence β (b) Price convergence β

Notes: The figure displays the estimated β_r coefficient and its 5% confidence interval for equation (1), and its equivalent for prices, as a function of the predispatch round.

regressions for prices. As best forecasts based on information before the predispatch, q_{nt}^0 and p_{nt}^0, we use the forecasts produced by the random forest predictor from Section 4.

The top panel of Figure 1 reports the slope estimates for the quantity equation (1) and for the North Island price equation (the results for the South Island are qualitatively similar). Consistent with the hypothesis that the predispatch market generates information, the coefficient β_r increases over the course of the predispatch and reaches one by the end of the predispatch. The estimated slope coefficient for the price equation remains low until the beginning of the short schedule, suggesting that information aggregation is picking up only then. The slope for the quantity equation increases steadily over the course of the predispatch.

Section 4 showed that residual uncertainty at the time the predispatch market opens differed across production technologies. Production at combined cycle gas power plants and geothermal stations could be predicted with little uncertainty, whereas residual uncertainty remained high for production at open cycle gas turbines nodes. Figure 2 reports the result of running equation (1) separately for four different technologies: schedulable hydro, run-of-river hydro, combined cycle gas and open cycle gas. The results reflect the combination of the quality of the random forests forecast, and the flexibility and exposure to last minute events of the specific technologies. Specifically, the main difference between schedulable hydro and run-of-river hydro is that run-of-river hydro is more dependent on short-run variations in river flows, leaving some generation uncertainty until close to dispatch. Comparing Figure 2 (a) and Figure 2 (b) confirms that quantity discovery indeed picks up later for run-of-river hydro than for schedulable hydro. In contrast, there is little difference in the evolution of the estimated β_r for combined and open cycle gas turbines, except for the slightly later start of quantity discovery for open cycle gas turbines. This may be due to both technologies being subject to similar operational constraints. We will see later (Table 7) that the time horizon of bid revisions is similar for both technologies.

Figure 2: Quantity discovery differences across technologies

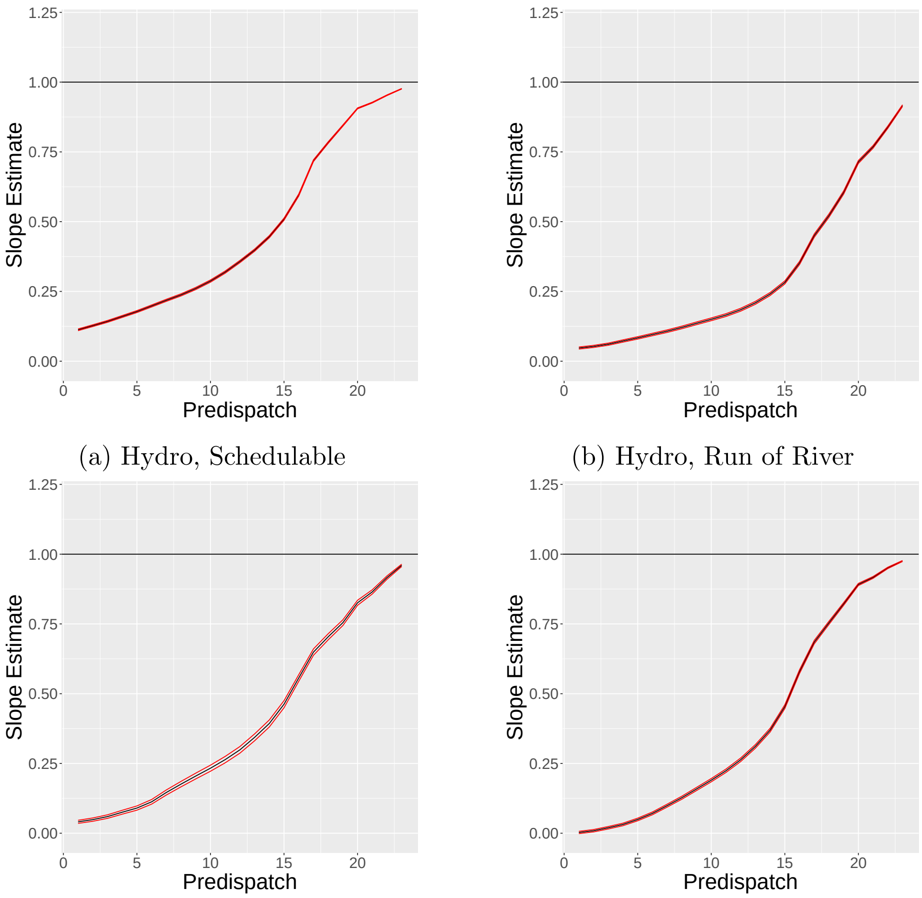

(a) Hydro, Schedulable

(b) Hydro, Run of River

(c) Gas, Combined Cycle

(d) Gas, Open Cycle

Notes: The figure displays the estimated β_r coefficient and its 5% confidence interval for equation (1) for a range of technologies.

5.2 Parallel markets and contribution to price and quantity discovery

A challenge when studying markets where identical or similar assets are traded is the potentially confounding effect of contemporaneous transactions in related markets. In our case, this takes the form of spot market transactions, which are based on the last predispatch, happening at the same time as earlier predispatch rounds for future trading periods. How can we know that the informational role that we have documented is performed by the predispatch, rather than the contemporaneous spot market, which involves commitment?

We examine the relationship between predispatch prices over time by regressing the quantity-weighted predispatch price for a given round r and trading period t, p_t^r, on previous lagged predispatch prices for the same trading period, and the most recent spot market price, i.e., the price of the trading period that has cleared right before the current predispatch round (and lags thereof):

$$p_t^r = \alpha + \beta_{11} p_t^{r-1} + \beta_{12} p_t^{r-2} + \cdots + \beta_{21} p_{\tau(t,r)}^{spot} + \beta_{22} p_{\tau(t-1,r)}^{spot} + \cdots + \epsilon_t^r,$$

where p_t^{spot} refers to the quantity-weighted spot price at trading period t and $\tau(t,r)$ gives the time period for which the spot market just cleared when round r of trading period t is happening. This means that such spot price is already known to market participants.[13] It is important to note that, by design, this spot price applies to a different hour of the day. Its defining feature is that it will be financially settled, as opposed to just being indicative.

Due to the high degree of correlation between predispatch rounds, we also consider a differenced model, in which the variables are first differenced in the regression.

Table 5 shows the results from the regression above. One can see in columns (1) and (2) that predispatch prices are very well predicted by the prices in previous rounds while the spot price is not nearly as important. This should not be surprising as we have already established that predispatch prices converge. Columns (3) and (4) focus on the differenced model, which more directly examines price *updates* in the predispatch market. Changes in spot prices are only weakly correlated with changes in predispatch prices and the coefficients are not significant, while lagged changes in predispatch prices are strongly correlated with current changes in predispatch prices. This suggests that price discovery is indeed happening in the predispatch market, independently of developments in the spot market.

The negative coefficients on lagged changes in predispatch prices shown in Table 5 are interesting on their own. They suggest some kind of reversion to the mean where an increase in price in one round is followed by a correction downwards, in the spirit of *tâtonnement*. Indeed, in the data, we find that almost 50% of the time price changes occur in the opposite direction from one round to the other. This suggests that prices in the predispatch market oscillate as information is revealed.

[13] For example, at the 21^{st} predispatch round for 4 pm in the afternoon, the last dispatch price at 2:00 pm has just been revealed. This is what we call the spot price.

Table 5: Predispatch price updates based on lagged predispatch and spot prices

	Levels		First Differences	
	(1)	(2)	(3)	(4)
L.Weighted price	0.679	0.493	-0.338	-0.431
	(0.022)	(0.024)	(0.019)	(0.025)
L2.Weighted price		0.176		-0.226
		(0.020)		(0.016)
L3.Weighted price		0.100		-0.118
		(0.010)		(0.010)
Spot price	0.015	0.017	0.008	0.013
	(0.014)	(0.016)	(0.008)	(0.013)
L.Spot price		0.000		0.009
		(0.000)		(0.008)
L2.Spot price		0.000		0.004
		(0.000)		(0.004)
Constant	20.411	14.010	-0.303	-0.263
	(1.713)	(1.240)	(0.069)	(0.081)
Observations	1,841,258	1,672,969	1,756,867	1,589,409

Notes: The unit of observation is a trading period x predispatch round. The price is a quantity-weighted average across nodes using final generation quantities. The spot price is a quantity-weighted average following the same approach and represents the price of the last trading period that has cleared at the time market participants bid in the market and will be *financially* settled. The goal of the regression is to separate "indicative" settlements vs. settlements that involve commitment. Clustered standard errors at the trading period reported in parenthesis.

5.3 New information arrival and bid updates

The previous results have established that the indicative prices and quantities produced by the predispatch market are increasingly informative and that this information is largely produced by the predispatch market itself, rather than the contemporaneous spot market. In this section, we zoom in on the bid revisions submitted by market participants since they drive the price and quantity revisions that we observe during the predispatch market.

Market participants can submit a bid revision any time during the predispatch, up to gate closure (conditions apply afterwards). A bid revision is characterized by a time stamp and the identity of the market participant who submitted it. We focus on generation-side bid revisions. We find that virtually all bid revisions involve several trading periods. Two thirds of bid revisions involve several nodes.

Figure 3 shows the evolution of the number of bid revisions during the predispatch. Since predispatch rounds last for different lengths, the numbers are normalized and expressed as the percentage of nodes subject to revision per hour. The figure shows that the number of revisions increases significantly after predispatch switches from the long schedule, when predispatch rounds last two hours, to the short schedule, when predispatch rounds last half an hour (vertical line). At the beginning of the predispatch, about 3% of node-level bids are revised every hour. This increases to 11% when the predispatch switches to the short schedule. Bid revision activity drops at gate closure, as constraints apply to bid revisions. Acceleration of trading activity close

Figure 3: Bid revision activity during the predispatch

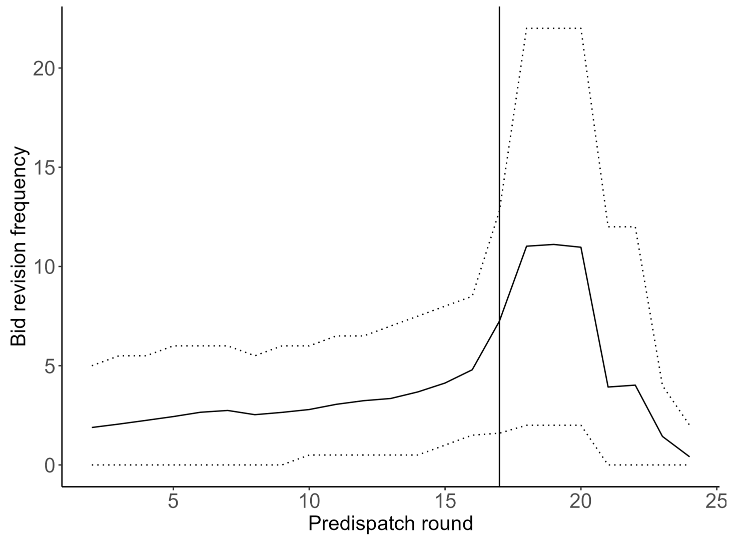

Notes: The figure shows the frequency of node-level bid revisions, for a given predispatch round (mean in solid line; 10 and 90 percentiles in dotted lines). The unit of observation is a trading period x predispatch round x generation node. Wind generation nodes are excluded. Bid revision frequency is defined as the percentage of nodes subject to a bid revision per hour. Vertical line at round 17, the round at which predispatch switches from the long to the short schedule. Gate closure starts at round 20 until June 28, 2017 and at round 22 afterwards.

to the market end-time is also documented for preopening periods at stock exchanges (Biais et al., 1999).

We next explore the determinants of bid revisions. Let $y_{intr} \in \{0, 1\}$ denote whether firm i revises their bid for node n and trading period t during the r^{th} round of the predispatch.[14] Similarly, let $y_{itr} \in \{0, 1\}$ denote whether firm i revises any bid for trading period t in the r^{th} round of the predispatch.

We construct several measures of exogenous information arrival at the market and node level. First, we measure changes, during the predispatch, in available generation capacity. For each node, we define $\Delta\text{Capacity}_{nrt}$ (Δ own capacity, in the table) as the absolute value of the change in available capacity between round $r-1$ and round r at that node (firm-level change in own capacity is defined as the mean of the node-level absolute value changes). This information is reported by participants alongside their bids. We use it as a proxy for unplanned outages.[15] Likewise, we define $\Delta\text{Capacity}_{(-i)rt}$ (Δ others' capacity) as the changes between round $r-1$ and round r in available capacity at nodes owned by other market participants, again measured as the sum of nodel-level absolute changes. Second, we measure market-level revisions in net load forecasts defined as the net change in load and wind forecasts between round $r-2$ and $r-1$ (Lagged Δ net load). This information is available to participants when they bid in round r.

Finally, we construct two measures for the market feedback received in the previous round. Let $\Delta q_{ntr} = |q_{nt}^{r-1} - q_{nt}^{r-2}|$ (Lagged Δ quantity) describe the change in indicative quantity at node n

[14] By convention, we say that a market participant revised their bids in the r^{th} round if they submitted a bid revision between the time of the $(r-1)^{th}$ and the r^{th} predispatch.

[15] News about new and unplanned outages is available in real-time to all participants on the WITS system.

Table 6: Determinants of bid revisions (linear probability model)

	Node-level				Firm-level		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ own capacity		0.00461	0.00464	0.00317	0.03994	0.04052	0.02971
		(0.00113)	(0.00115)	(0.00081)	(0.00804)	(0.00825)	(0.00545)
Δ others' capacity				0.00092			0.00128
				0.00005			0.00012
Lagged Δ net load				0.00004			0.00008
				(0.00001)			(0.00001)
Lagged Δ quantity			0.00076	0.00064		0.00068	0.00058
			(0.00013)	(0.00010)		(0.00061)	(0.00047)
Lagged Δ price			0.00011	0.00004		0.00015	0.00005
			(0.00001)	(0.00001)		(0.00002)	(0.00001)
Round-node FE	X	X	X	X			
Round-firm FE					X	X	X
Obs. (million)	132.7	132.0	126.3	126.3	20.7	19.8	19.8
Adjusted R2	0.02	0.05	0.05	0.26	0.15	0.15	0.38

Notes: The unit of observation is a trading period x generation node x predispatch round for node-level regressions and a trading period x generation firm x predispatch round for firm-level regressions. At the firm level, a revision is defined as taking a value of one as long as the bid schedule of one of its plants is revised. Wind excluded. Standard errors clustered at the node level (specifications (1)-(4)), and at the firm level (specifications (5)-(7)) in parenthesis.

and trading period t from predispatch round $r-2$ to predispatch round $r-1$ (Δp_{ntr} is defined analogously). A positive value for Δq_{ntr} can be the result of a previous change in bids (between round $r-2$ and $r-1$) or a change in market circumstances that leads the market model to select another point on the market participant's bid schedule at node n.[16] The interpretation for Δp_{ntr} is similar and the two variables will tend to be correlated, except that, because bid schedules are step functions, indicative quantities can change without indicative prices changing, and vice versa. Firm-level changes in indicative price and quantity are defined as the average of the node-level changes in price and quantity.

We run two sets of linear probability regressions, one at the node level and, because two-thirds of bid revisions involve several nodes, one at the market participant level. Figure 3 shows that bid revisions are more frequent in later stages of the predispatch. Additionally, some technologies might be more prone to frequent revisions than others. Therefore, we control for round-node (in node-level regressions) and round-market participant fixed effects (in firm-level regressions).

Table 6 summarizes the results. All coefficients are positive, as expected, and significant at the 1% level, except lagged own quantity change in specifications 6 and 7. Changes in the available capacity of a market participant increases the likelihood that they revise their bids. Changes in the indicative price and quantity at a node during the previous round are also associated with an increased probability of submitting a bid revision for that node, but their explanatory power

[16]Brown et al. (2018) find that participants in the Alberta's wholesale electricity market respond to rival offer changes, as revealed by the local market authority's historical trading reports.

Table 7: Number of trading periods involved in a bid revision (node level)

	# revs.	5%	25%	50%	75%	95%
Overall	5.76	1.0	3.0	6.0	15.0	48.0
- Coal	4.76	1.0	2.0	6.0	15.0	34.0
- Cogeneration	3.77	2.0	11.0	29.0	48.0	69.0
- Diesel	2.02	1.0	3.0	8.0	24.0	44.0
- Gas, combined cycle	4.66	1.0	3.0	9.0	23.0	45.0
- Gas, open cycle	2.74	1.0	3.0	8.0	21.0	48.0
- Geothermal	3.03	1.0	5.0	13.0	23.0	48.0
- Hydro, run of river	5.50	1.0	2.0	5.0	14.0	49.0
- Hydro, schedulable	6.29	1.0	3.0	5.0	12.0	48.0

Notes: The unit of observation is a node-level bid revision (timestamp x node). The first column reports the average number of bid revisions per node and trading period. Note that we only include bid revisions that happen during the predispatch, i.e., revisions submitted after the first predispatch. $N = 980,777$.

is small. Finally, changes in competitors' available capacity and changes in net load increase the probability of a bid revision and these variables have a strong explanatory power as evidenced by the increase in the adjusted R^2. These results confirm that the predispatch market reacts to new information.

6 Why does price and quantity discovery happen without commitment?

So far, we have shown that market participants actively participate in the predispatch market and that their revised bids contribute to making indicative predispatch prices and quantities increasingly informative. In this section, we explore the possible *private* incentives for bid revisions given the absence of commitment. After all, there is no reason for a market participant to submit a bid revision if they do not privately benefit from it.

We already noted that essentially all bid revisions involve several trading periods. Table 7 provides a detailed breakdown of the number of trading periods involved in a bid revision, by technology. The time span of bid revisions reflects the technical attributes of each technology. For example, more than 25% of bid revisions for gas-powered thermal units involve at least 21 trading periods according to Table 7. This is close to the typical warm-up time for these units, which is evaluated to be at around 10-12 hours. Coal-powered plants have faster start-up times and bid revisions tend to involve fewer trading periods for this technology. Half of the bid revisions for cogeneration involve at least 29 trading periods, presumably reflecting the production constraints of the industrial process (paper and pulp, dairy, ...) paired with electricity generation. At the other extreme, 75% of bids for schedulable hydro involve less than 12 trading periods (6 hours).

Unlike most other electricity markets, including the Continental West European (CWE), the Iberian and Nordic markets in Europe, and PJM, California, and Colombia in the Americas, the New Zealand wholesale electricity market does not allow market participants to express

Figure 4: Evidence for intertemporal optimization

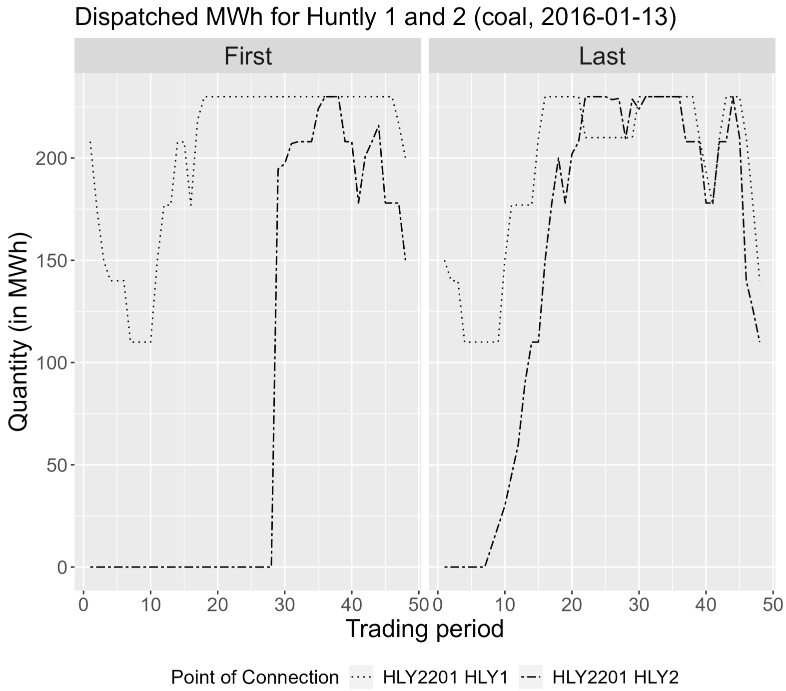

Notes: The figure shows an example of scheduled generation over the course of the day (48 trading periods) for the Huntly coal power station. One can observe that Unit 2 changes its schedule to improve its ramping profile over the course of the predispatch market.

preferences over intertemporal production profiles.[17] Prices and quantities are set for each trading period sequentially. The only dynamic consideration taken into account is ramp up and ramp down constraints from the previous trading period.

This means that the predispatch market, and the indicative prices and quantities it produces, is the only mechanism through which generation units are able to optimize their production profile over time. Specifically, whenever Transpower calls the market, it produces indicative prices and quantities for 8 (during the short schedule) or 72 (during the long schedule) *consecutive* trading periods. Market participants can then get a good sense of the forecast generation schedule of their plants over this time horizon, in the absence of further changes. The prices and aggregate supply curves generated in the process also give them an indication of how to adjust their bids to change their forecast generation schedule, if desirable.[18] This will of course not guarantee that they will secure the desired schedule since prices and quantities are not final, but the probability that it does steadily increases over the course of the predispatch, given the convergence we have documented. Such benefit of iterative mechanisms has been emphasized by Ausubel and Cramton (2004) among others.

Figure 4 provides one example of intertemporal reallocation of production over the course of

[17]These can take the form of a minimum revenue requirement as in the Iberian market (Reguant, 2014), explicit fixed costs bids in addition to "simple" bids as in the Colombian market (Balat et al., 2022), block and linked orders and other multi-period contingent bids (Tirez et al., 2012) as used in the CWE area.

[18]For example, a market participant who is allocated a time-varying schedule of production, based on the current predispatch results, might want to adjust their bids to smooth out their production level and avoid costly short-run changes in generation.

the predispatch. It shows the evolution of scheduled production on January 13, 2016, for the first and second coal units at Huntly, a station located on the North Island and operated by Genesis. At the first round of the predispatch, the second unit was on only for 10 hours of the day, starting at trading period 28 (1:30 pm). By the end of the predispatch, it is committed for a longer period, a useful property given the ramp-up costs of such units.

To provide a more systematic assessment of the presence of improved intertemporal allocation, we construct two measures of dynamic efficiency at the day and technology level. First, we define the number of daily startups as follows:

$$Startup^r_{d,tech} = \sum_{n \in tech} 1_{\{q^r_{ndt}>0 \ \& \ q^r_{nd,t-1}=0\}},$$

where r refers to the predispatch round, n is a node belonging to a technology $tech$, d is a day, and t is a trading period. Everything else equal, a day with fewer startups will tend to be dynamically more efficient. Second, we define a measure of ramping, i.e., the change in quantities from one trading period to the next. Thermal power plants have additional costs from quickly changing their output, and therefore, we consider our ramping to be quadratic in output changes:

$$Ramp^r_{d,tech} = \sum_{n \in tech} \sum_{t \in TP} (q^r_{ndt} - q^r_{nd,t-1})^2$$

Figure 5 shows evidence consistent with dynamic efficiency improving over the course of the predispatch market. Panel (a) shows that the number of coal start-ups decreases from the first to the last predispatch. Due to the large startup costs of coal power plants (Wolak, 2007; Reguant, 2014; Gowrisankaran et al., 2023), the reduced number of startups can reduce operational costs. Panel (b) shows that the output in the last predispatch is also less subject to ramping costs.

Panels (c)-(d) show the same patterns, albeit less stark, for combined cycle gas units. As Table 2 indicated, these units were largely used as baseload during our sample period, explaining the low number of start-ups and the little change thereof.

As a point of comparison, panels (e)-(f) show the value of our dynamic efficiency measures for schedulable hydro. Schedulable hydro adjusts production during the day, with several nodes going from zero to positive production on a daily basis, as shown by the number of startups.[19] Contrary to the thermal plants, there is no discernible difference between the first and last predispatch. This is intuitive as ramping constraints and dynamic costs are not as relevant for schedulable hydro.

Improved intertemporal allocation of production is clearly a private benefit. To materialize it requires informative prices and quantities, so market participants have a collective interest in the quality of the predispatch. When a market participant revises their bids, they are not only optimizing their own production but also improving information for other market participants. In turn, more informative prices provide an incentive for market participants to use them to update their bids, thereby creating a virtuous cycle. This provides a rationale

[19]Note that there are 36 schedulable hydro nodes and, therefore, the number of daily startups can be quite high.

Figure 5: Measures of dynamic efficiency over the course of the predispatch market

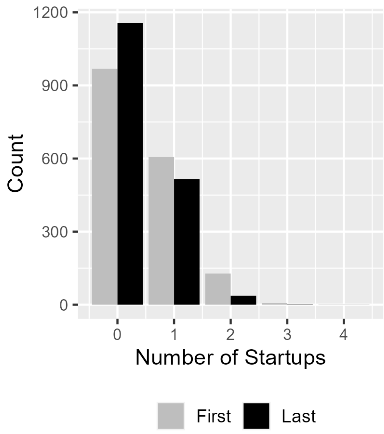

(a) Coal Startups

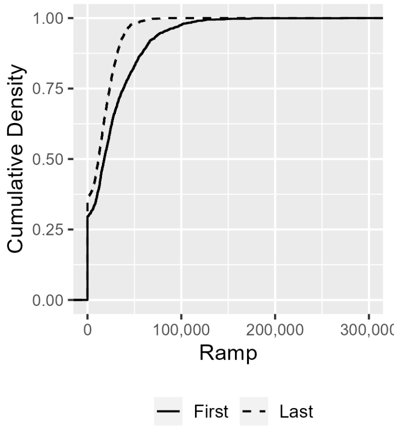

(b) Coal Ramp

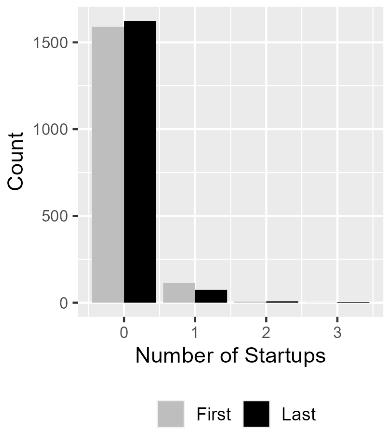

(c) CCGT Startups

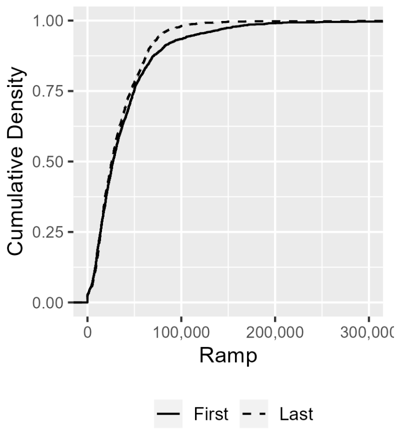

(d) CCGT Ramp

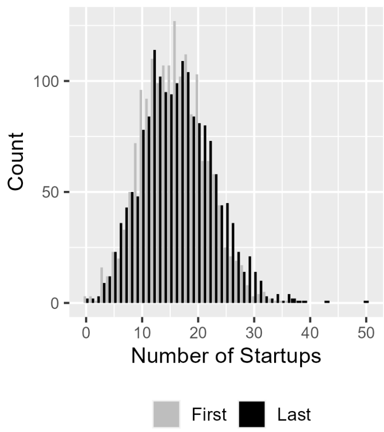

(e) Schedulable Hydro Startups

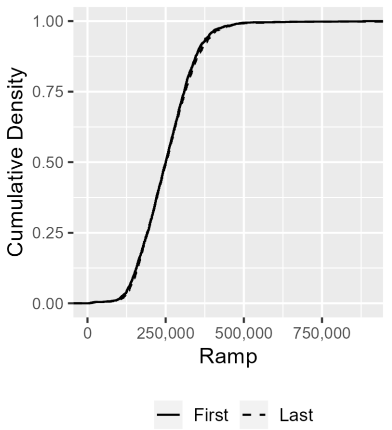

(f) Schedulable Hydro Ramp

Notes: The figure shows how alternative measures of dynamic efficiency evolve between the first and last predispatch. The unit of observation is a day x technology pair. Startups are measured as the daily number of startups for units of a given technology. Ramps for each node are measured as the sum of squared changes in output from one trading period to the next. This number is then aggregated at the technology level.

for the observed active participation in the predispatch market and its informativeness, despite the lack of commitment. Note, however, that the level of bid revisions that we observe, and therefore the level of price informativeness, may not be optimal, given that bid revisions produce a positive externality.

7 Concluding comments

Wholesale electricity markets are notoriously incomplete (Wilson, 2002). Existing market designs are all pragmatic attempts to solve the complex allocation that electricity production and dispatch entail. The New Zealand electricity market is no exception. Its distinguishing feature is the use of a non-binding indicative predispatch market, before final allocations are decided.

Non-binding iterative markets are uncommon: preopening periods at stock exchanges and initial public offerings (IPOs) seem to be the other two examples. They raise the concern that participation is uninformative at best, manipulative at worst. We show that bid revisions in the New Zealand market are motivated by new information arrival and that predispatch prices and quantities are increasingly informative of final prices and quantities.

Our explanation for the informativeness of the predispatch market is that market participants derive a private benefit from effective price discovery, in the form of improved intertemporal coordination of production plans. This contrasts with the reasons given for the informativeness of preopening periods at stock exchanges and IPOs which rely on asymmetric information and adverse selection.

Could commitment nevertheless help? Introducing some form of commitment is on the agenda of New Zealand policy-makers. It is motivated by the observed increasing frequency of peak periods when offered generation is tight relative to load (despite otherwise sufficient installed capacity), and the perception that this arises because the financial incentives for keeping thermal generation capacity "warm" for the probable event that it may be needed have decreased in the context of higher fuel and carbon prices and increased wind penetration that makes prices more difficult to predict. The proposed solutions include measures to improve the accuracy of wind generation forecasts, predispatch market feedback based both on predicted load (like today) but also on load forecast sensitivity cases, and the introduction of commitment in the form of a hours-ahead market that would lock in parts of the supply and demand 8 hours or so before dispatch (Electricity Authority, 2023).

Our results provide two insights into this question. First, our argument that there is a virtuous cycle at play in the predispatch market implies, by the same token, that any external measure to improve the informativeness of the resulting price signals will strengthen the incentives of market participants to use the predispatch market to convey information about their production plans, thereby further improving the informational value of the predispatch. Second, the intertemporal nature of market participants' optimisation problem that we have documented suggests that careful attention should be paid to the way commitment is introduced unless some form of multiperiod bidding is introduced.

Our paper has documented the efficiency benefits of information in terms of improved production scheduling. The flip side of information is that it also increases the market participants' ability to exert market power. We turn to this question in our follow-up paper (Bergheimer et al., 2023).

References

Alvey, T., D. Goodwin, X. Ma, D. Streiffert, and D. Sun (1998). A security-constrained bid-clearing system for the New Zealand wholesale electricity market. *IEEE Transactions on power systems 13*(2), 340–346.

Ausubel, L. M. (2004). An efficient ascending-bid auction for multiple objects. *American Economic Review 94*(5), 1452–1475.

Ausubel, L. M. and P. Cramton (2004). Auctioning many divisible goods. *Journal of the European Economic Association 2*(2-3), 480–493.

Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek, and M. Weretka (2014). Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies 81*(4), 1366–1400.

Balat, J., J. E. Carranza, J. D. Martin, A. Riascos, et al. (2022). The effects of changes in the regulation of the Colombian wholesale electricity market in a structural model of complex auctions. Technical report, Banco de la Republica de Colombia.

Barclay, M. J. and T. Hendershott (2008). A comparison of trading and non-trading mechanisms for price discovery. *Journal of Empirical Finance 15*(5), 839–849.

Bergheimer, S., E. Cantillon, and M. Reguant (2023). The benefits and costs of information: Evidence from the New Zealand wholesale electricity market.

Biais, B., C. Bisière, and S. Pouget (2014). Equilibrium discovery and preopening mechanisms in an experimental market. *Management Science 60*(3), 753–769.

Biais, B., P. Hillion, and C. Spatt (1999). Price discovery and learning during the preopening period in the Paris Bourse. *Journal of Political Economy 107*(6), 1218–1248.

Bolle, F. (1992). Supply function equilibria and the danger of tacit collusion: the case of spot markets for electricity. *Energy economics 14*(2), 94–102.

Brown, D. P. and A. Eckert (2022). Pricing patterns in wholesale electricity markets: Unilateral market power or coordinated behavior? *The Journal of Industrial Economics 70*(1), 168–216.

Brown, D. P., A. Eckert, and J. Lin (2018). Information and transparency in wholesale electricity markets: evidence from Alberta. *Journal of Regulatory Economics 54*(3), 292–330.

Cao, C., E. Ghysels, and F. Hatheway (2000). Price discovery without trading: Evidence from the Nasdaq preopening. *The Journal of Finance 55*(3), 1339–1365.

Cicala, S. (2022). Imperfect markets versus imperfect regulation in US electricity generation. *American Economic Review 112*(2), 409–441.

Electricity Authority (2023). Driving efficient solutions to promote consumer interests through winter 2023. *Decision.*

Gowrisankaran, G., A. Langer, and W. Zhang (2023, January). Policy uncertainty in the market for coal electricity: The case of air toxics standards. Working Paper 30297, National Bureau of Economic Research.

Hayek, F. (1945). The use of knowledge in Society. *The American Economic Review 35*(4), 519–530.

Hong, J. and S. Pouget (2021). Liquidity formation and preopening periods in financial markets. *Economica 88*(351), 697–723.

Kamada, Y. and M. Kandori (2020). Revision games. *Econometrica 88*(4), 1599–1630.

Milgrom, P. R. and R. J. Weber (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, 1089–1122.

Nisan, N. and I. Segal (2006). The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory 129*(1), 192–224.

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

Tirez, A., P. Luickx, and D. Woitrin (2012). Impact of sophisticated orders on spot power prices. In *2012 9th International Conference on the European Energy Market*, pp. 1–7. IEEE.

von der Fehr, N.-H. M. (2013). Transparency in electricity markets. *Economics of Energy & Environmental Policy 2*(2), 87–106.

Wilson, R. (2002). Architecture of power markets. *Econometrica 70*(4), 1299–1340.

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