

Does Financial Trading Smooth Non-Convex Markets?

Nicolas Stevens^{*†}, Peter Cramton[‡] and Martial Toniotti[§]

July 6, 2026

Abstract

In non-convex markets, a competitive equilibrium may fail to exist. This turns out to be an important issue in real-world non-convex auction markets, such as electricity markets, as it complicates pricing and requires the auctioneer to resort to out-of-market discriminatory side payments to sustain an equilibrium. We investigate whether the introduction of *convex* financial trading induces a *smoothing effect*, mitigating the issues arising from non-convexities. We develop a two-stage non-convex market model (a forward market followed by a spot market) in which convex financial traders participate in the forward market. Our model predicts that financial trading reduces the magnitude of side payments required to support the cleared allocation. To test the prediction of our model, we examine the introduction of a transaction fee on financial traders in 2020 by PJM, the US's largest electricity market. We show that the substantial decline in financial trading volume caused by this policy coincided with a significant increase in side payments, in line with our theoretical predictions.

Keywords: Nonconvex Markets, Financial Trading, Equilibrium Existence, Electricity Auctions

*Corresponding author.

†Max Planck Institute for Research on Collective Goods (Bonn, Germany) & Center for Operations Research and Econometrics, UCLouvain (Louvain-la-Neuve, Belgium). Email: stevens@coll.mpg.de

‡Max Planck Institute for Research on Collective Goods (Bonn, Germany) & University of Maryland, College Park, USA. Email: pcramton@gmail.com

§Center for Operations Research and Econometrics, UCLouvain, Louvain-la-Neuve, Belgium. Email: martial.toniotti@uclouvain.be

1 Introduction

In a non-convex market, a competitive equilibrium may fail to exist. Several classical results show that this failure might not be too severe, at least *on the large scale*: Starr (1969) shows that when an equilibrium does not exist, the distance to an equilibrium (as measured by total excess supply or demand) is invariant with the number of market participants, while Aumann (1964) shows that when there is a *continuum* of traders, an equilibrium exists in spite of the non-convexities. Yet, despite these “asymptotic” positive results, non-convexity remains problematic in real-world auctions where the auctioneer cannot simply overlook these complexities but must determine market-clearing prices (Milgrom, 2017). There is an extensive literature which approaches this problem by designing mechanisms that accommodate non-convexities (see, for instance, Milgrom and Watt (2025), and the literature covered hereafter). Instead of designing new mechanisms, we investigate whether the introduction of financial traders induces a *smoothing effect* which helps solve the issues arising from non-convexities. Two important features of financial traders are that (i) they use *convex* bids and (ii) they are *arbitragers*, thus they are expected to bid close to the market margin. We analyze how this addition of a large volume of *convex* bids *close to the margin* smooths the non-convexities present in the market and helps the auctioneer with price formation. We study this question in the context of electricity markets, which combine non-convex auction structures with active financial trading.

Electricity wholesale markets are typically organized as sealed-bid uniform-price auctions (Wilson, 2002; Cramton, 2017). Most of these auctions involve non-convexities in the bidding language that enable the market participants, in particular the electricity producers, to represent the costs and constraints of operating a power plant directly into the auction. Focusing on the United States¹, each power plant bids separately in the market, using multi-parts bids that represent their technical characteristics, and where non-convexities arise from features such as start-up costs, minimum up and down times requirements, minimum production limits, and so on. The Walrasian auctioneer collects the bids of all the market participants and solves the so-called unit commitment model (Knueven et al., 2020) in order to find the allocation that maximizes the market surplus under all technical constraints (grid constraints and production constraints represented by the bids of the market participants). The auctioneer then computes the market-clearing prices.

The presence of non-convexities has several consequences, such as on the algorithmic complexity of the auction (Knueven et al., 2020), or on the bidding behaviour and the nature of competition in the market (Reguant, 2014; Jha and Leslie, 2025). Yet, the main implication is that an equilibrium is not guaranteed to exist, although it might exist in some cases (Bikhchandani and Mamer, 1997). In concrete terms, given the surplus-maximizing allocation, the auctioneer may not be able to find a uniform price that “clears the market”. This is illustrated in the following example. Figure 1 presents an elementary electricity market with two demand bids, D_1 and D_2 , as well as two supply offers S_1 and S_2 . Assuming first that all the bids are convex, a Walrasian auctioneer maximizing the total surplus would simply intersect supply and demand: he would clear D_1 , clear 90MWh of S_1 and reject the other bids. In this configuration, the price $\pi = 30\$/\text{MWh}$ clears the market. If now S_1 is a non-convex “all-or-nothing” bid (either it is entirely cleared or entirely rejected), then the surplus-maximizing allocation clears S_1 , D_1 , as well as 10MWh of D_2 . The key issue arising from this non-convexity is that there is no uniform price that supports the cleared allocation. Indeed, if the auctioneer chooses $\pi = 20\$/\text{MWh}$, then S_1 loses money. More generally, this occurs whenever $\pi < 30\$/\text{MWh}$. If instead he chooses $\pi = 30\$/\text{MWh}$, the demand bid D_2 loses money since his willingness-to-pay for electricity is below the market price. This occurs whenever $\pi > 20\$/\text{MWh}$. Thus there is no competitive equilibrium (no price satisfies both *market-clearing* and *envy-freeness*).

Practically, the electricity auctioneers address the issue by complementing the uniform

¹Although our work is focused on the US markets, non-convexities are also present in other electricity markets. For instance, non-convexities arise from the “block orders” (with features such as all-or-nothing dispatch, minimum acceptance ratio, exclusive groups, etc.) in European or Indian markets.




Figure 1: An hourly two-sided electricity market with four bids.

price with discriminatory *side payments* (or “markups”, as called by Milgrom and Watt (2025)) paid to the market participants to support the allocation. In our example, this could mean setting the price to $\pi = 20\$/\text{MWh}$ while providing a side payment of $(30 - 20) \times 100 = 1000\$$ to S_1 in order for S_1 to break even. It can be shown more formally, and for a general auction model, that a competitive equilibrium can indeed be recovered by extending the uniform price vector with side payments (O’Neill et al., 2005). These side payments are generally regarded as undesirable for they are discriminatory, non transparent (they do not have temporal nor locational granularity), they can induce gaming behaviour (Byers and Eldridge, 2023) and they also deter investment signal (Byers and Hug, 2023). The key implication is that *some relevant cost information is not reflected in the uniform price signal but is instead handled by “out-of-market” side payments.*

There is an extensive literature that aims at addressing these issues by designing pricing mechanisms for non-convex auctions (Milgrom and Watt, 2025), and in particular for electricity auctions (O’Neill et al., 2005; Gribik et al., 2007; Chao, 2019; Stevens et al., 2024; Ahunbay et al., 2025; Stevens et al., 2026). These market design proposals depart from the classical marginal pricing rule by incorporating some non-convex costs into the price signal, thereby reducing the side payments required to sustain equilibrium. For example, one prominent proposal is to select the uniform price that *minimizes* the required side payments (Gribik et al., 2007).

In this paper, we approach the problem from a different perspective by asking whether the participation of financial traders can itself help mitigate these issues. Electricity auctions are not pure combinatorial auctions, as they include a significant share of *convex* bids, among which are the financial trading bids, often referred to as “virtual bids”. Electricity markets are organized as sequential auctions, the most important being the day-ahead *forward* market and the real-time *spot* market. The day-ahead auction is a financial market that involves both “physical” (non-convex) bidders, tied to physical assets like power plants, and “virtual” (convex) bidders that engage in arbitrage between day-ahead and real-time prices. The benefits of financial trading in electricity markets have been extensively studied in the literature (Hogan, 2016). These benefits include mitigating market power (Ito and Reguant, 2016; Mercadal, 2022), enhancing price convergence and reducing costs (Jha and Wolak, 2023), as well as providing risk hedging opportunities (Longstaff and Wang, 2004). Hogan further conjectured the additional benefit that “including virtual transactions [...] would have the effect of smoothing the day-ahead commitment and dispatch problem and reducing the required uplift [side] payments.” (Hogan, 2016) In this paper, we investigate this claim formally and empirically: by making the non-convex market more convex, financial trading exerts a *smoothing effect* that improves price formation and reduces the side payments required to sustain an equilibrium. In particular, our paper makes two main contributions.

First, we develop a two-stage stochastic market model (a forward market followed by a spot market) in which financial traders participate in the forward market. Our model predicts that financial trading improves price convergence, reduces total cost, and—most important for our analysis—mitigates the magnitude of side payments required from the auctioneer to sustain an equilibrium. More specifically, we show that virtual trading induces two distinct effects. (i) First, by behaving as arbitragers in the day-ahead market, virtual traders internalize both expected real-time conditions and the non-convex operating costs embedded in day-ahead commitment decisions. In our model, virtual traders set the price and eliminate the side payments in day-ahead. (ii) Second, by anticipating real-time conditions, financial trading implies better commitment decisions in day-ahead, leading to less lumpy activations of “fast-start resources” in real-time. This helps reduce the side payments, as the activation of fast-start resources is typically associated with side payments needed to cover their startup costs.

This *smoothing effect* is distinct from the classic *market size effect* as analyzed by Starr (1969) in the context of general equilibrium theory and recently reinterpreted by Milgrom and Watt (2025) in the context of auction theory (see also Chao (2019), Stevens et al. (2024) and Hübner (2025) for transposition of these results to the context of electricity auctions). The market size effect shows that, in a non-convex economy where equilibrium may fail to exist, the distance to equilibrium—as measured by either total excess demand or side payments—is bounded by a constant independent of the number of market participants. Hence, as the market becomes larger (increasing the number of participants), the absolute distance to equilibrium remains unchanged while its relative importance with respect to market size decreases.² The smoothing effect we analyze is not merely related to the market size, but to the specific type of *convex arbitrage bids* introduced by the financial traders.

Second, we test the prediction of our model by analyzing the introduction in 2020 of a transaction fee on financial traders in PJM, the largest electricity market in the United States. This transaction fee led to a substantial decline in financial trading volume. We provide evidence that this reduction was followed by a significant increase of side payments needed to sustain equilibrium. We find that the real-time side payments increased by 80% following the introduction of the transaction fee, while the likelihood of non-zero day-ahead side payments increased by 10%.

Our empirical analysis contributes to the literature on financial trading, with findings in line with prior empirical evidence. In particular, Long and Giacomoni (2020) simulate the PJM auction model with and without financial bids and show that the inclusion of financial bids mitigates the side payments. The main difference with our analysis is that their work relies on simulations that assume the removal of financial bids does not affect the bidding behavior more broadly. While our empirical analysis relies on a policy change in the market. Besides these simulation exercises, the two closest empirical works are Hübner (2025) and Jha and Wolak (2023). Hübner (2025) provides empirical evidence, in the context of the European electricity auction, that when the market is large, with a high share of convex bids, the probability of reaching an equilibrium becomes higher. The main difference with our work, besides the fact that we focus on PJM market instead of the European market³, is that the analysis of Hübner (2025) is focused on the *market-size effect*—considering the effect of convex bids in general—while our work focuses on the *smoothing effect* induced by a specific type of convex bids, namely the financial trading bids. Finally, Jha and Wolak (2023) provide empirical evidence of price convergence and cost-saving entailed by the participation of financial traders in CAISO market. While they

²This important result is a direct application of Shapley-Folkman lemma. See also (Bertsekas and Sandell, 1982; Bertsekas et al., 1983) for the perspective of linear programming and duality gaps.

³There are two main benefits of focusing on PJM market over the European market, for our analysis: (i) non-convexities are more prevalent in PJM market (as PJM market is nodal, the market tends to be more fragmented, which exacerbates the impact of non-convexities ; and the European market bidding language, although it includes non-convex bids, tends to be somewhat less non-convex than the unit commitment model of PJM) and (ii) financial traders’ bids can be distinguished from power plants’ bids.

consider a non-convex market, they do not specifically look at the implications for side payments and price formation.

The rest of the paper is organized as follows. Section 2 provides an overview of PJM Interconnection electricity market. In particular, it discusses the sequential nature of the market (forward and spot markets) and the role of financial traders, and it provides descriptive statistics of the share of convex and non-convex bids in the market as well as the magnitude of the side payments. It also helps motivating some modeling choices we later make. Section 3 goes on with our model of a non-convex sequential electricity market with financial trading participation. This section formalizes the intuition of the smoothing effect described above. Sections 4 and 5 cover the empirical analysis with the data from PJM. Section 6 puts our main findings in perspective of policy discussions and concludes.

2 PJM Interconnection market

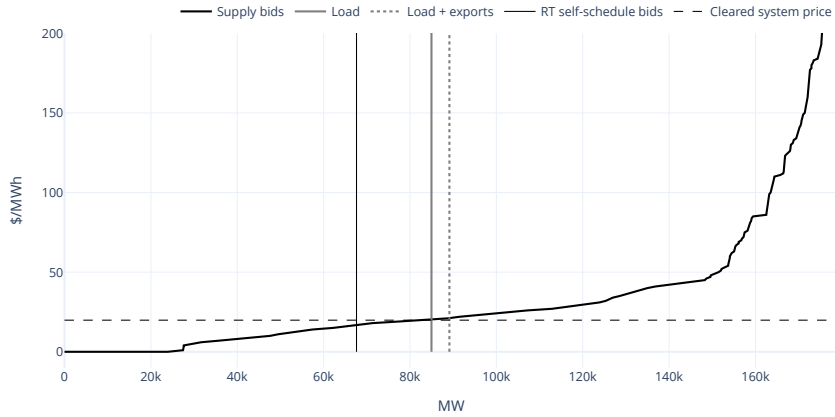
PJM Interconnection operates the largest wholesale electricity market in the United States, serving 67 million people across 13 states, and monitoring more than 1,400 power plants accounting for roughly 180GW of generation capacity.⁴ In 2024, the hourly average demand including exports was 95 GWh (with a peak load plus exports of 154 GWh), served by production coming mainly from gas (44%), coal (15%) and nuclear (32%) (Monitoring Analytics, 2025a).

As it is common in the US, the wholesale electricity market works as a sequence of sealed-bids uniform price auctions organised by PJM, which follow the principles of bid-based, security-constrained, economic dispatch with locational marginal prices (Hogan, 2021). Two cornerstones of the wholesale market are the real-time *spot* market and the day-ahead *forward* market. The real-time market is the *physical* market when a physical commodity is sold “on the spot” and where the trades correspond to physical exchanges of energy. The day-ahead market is a *financial* market: there is no physical delivery but only trades of financial contracts, which are derivatives of the real-time price of electricity. Yet, it is a financial market which has significant physical implications as its outcome is used for coordination and commitment of power plants for the next day (Stoft, 2002). Generation companies and consumers’ electricity providers participate in these auctions by submitting multi-parts bids representing their willingness to sell or to buy electricity. PJM market also includes a detailed model of the power grid representing the constraints of trade between the locations (thousands of nodes are represented in PJM auction model). The real-time market clears every five minutes while the day-ahead market clears once per day (it is a multi-periods market that clears at once the 24 hourly periods of the next day). These markets also co-optimize both energy and ancillary services products.⁵

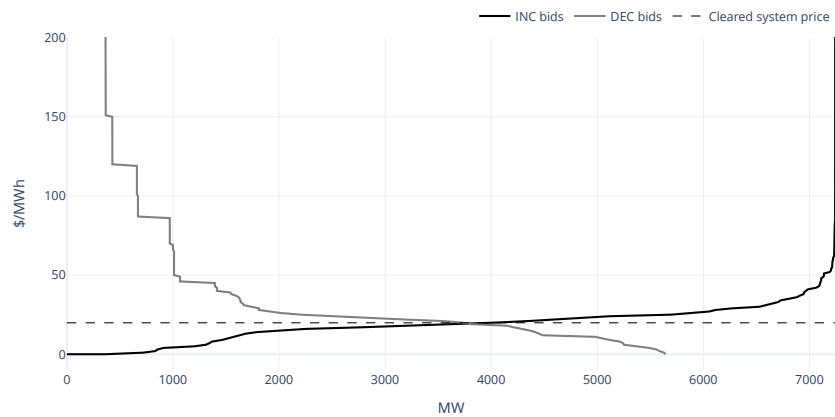
Bids and non-convexities. The day-ahead electricity auction in PJM mainly includes two types of bids: the “physical” bids and the “virtual” bids. Physical bids are tied to physical assets. They further split into two types: dispatchable generation bids and self-schedule bids. The dispatchable generation are *non-convex* (or partly non-convex) bids that enable the market participants to express the true economics of their power plants: they are multi-part “unit commitment” bids which include explicit information about costs and constraints of production. The main sources of non-convexities relate to production constraints such as minimum production limits or minimum up and down times requirements of power plants as well as to fixed costs such as start-up costs and no-load costs (let us emphasize that whenever we talk about “fixed costs” in this paper, we mean the operational avoidable fixed costs, such as start-up costs, as opposed to the capital costs which are sunk and have no implication for our discussion). The self-schedule bids are convex and correspond to “offer to supply a fixed block of MW, as a price taker” (Monitoring Analytics, 2018), which

⁴The installed capacities in 2024 were roughly 90GW of gas, 35GW of coal, 34GW of nuclear, 5GW of oil, 6GW of hydropower, >10GW of wind and >10GW of solar.

⁵See PJM manual for a detailed timeline and specific rules of these market (PJM, 2025b).



(a) PJM merit order curve (physical bids)



(b) PJM virtual bids (INC and DEC)

Figure 2: PJM bidding curves by category on June 12th, 2019, at 9am.

in PJM means a continuous quantity at zero price. Power plants have the obligation to participate into the day-ahead market of PJM, but they are allowed to either participate as dispatchable generation or to self-schedule (PJM, 2025b).⁶ Figure 2a shows the merit order curve of physical bids for a sample hour (12/06/2019 at 9am): the day-ahead self-schedule bids are the flat zero-price segment while the dispatchable generation bids are the remaining part of the curve. This aggregated merit order curve is of course a simplification, as it ignores the network constraints (although the selected period corresponds to an hour with little congestion), the fixed costs and other constraints of production. Yet, we see on Figure 2a that supply and demand (the inelastic load) cross closely to the actual PJM system price at that time. The figure also shows the *real-time* self-scheduling, which corresponds to the amount of inflexible supply in real-time. As we observe, this coincides to most of the supply committed in day-ahead.

Besides these physical bids, PJM day-ahead market also includes virtual bids which are pure financial trading: market participants playing as arbitrageur between the day-ahead financial market and the real-time spot market. Concretely, a virtual trader who sells a quantity q at the day-ahead price p_{DA} have to buy it back at the real-time price p_{RT} , thus he makes a profit $q \times (p_{DA} - p_{RT})$. There are three types of virtual bids in PJM: INC, DEC and UTC orders, which are all *convex* bids. An INC order is a sell, or “increment”, order (an injection) in one location of the grid. A DEC is a buy, or “decrement”, order (a withdrawal) in one location of the grid. A UTC is Up To Congestion Transactions order

⁶The main restriction to self-scheduling is that these bids are not compensated for their losses in the market: to incentivize units to be flexible, self-schedule units are not eligible to receive side payments.




Figure 3: PJM convex bid volume as a percentage of total load.

which is a combination of an INC and a DEC order in two different locations: an arbitrage between two locations of the grid, similar to a FTR contract (Hogan, 2016). Figure 2b also shows the INC and DEC bid curves on the same sample hour as Figure 2a. We observe these curves intersect closely around the market price and are fairly flat around the market price, thus *these virtual bids add a large volume of convex bids close to the margin*. That is: unlike self-schedule bids, which are convex *but at zero price*—thus not at the margin—, virtual bids are convex *and close to the margin*, which means more likely to set the price. This is a crucial observation for the smoothing effect analysed in our paper, which will be formalized later in section 3. Note that there is an explicit distinction in PJM market data between physical and financial bids. This is a major asset for our analysis, and it contrasts with European electricity markets, where bids often aggregate multiple assets and do not separate physical from financial offers. There, a high share of convex bids does not necessarily reflect active financial trading, but might relate to a more convex physics: a smoother operating day with fewer binding non-convex constraints. PJM data allows us to cleanly disentangle these two effects.

From the description of these different bids, it emerges that electricity auctions such as PJM are not pure combinatorial auctions: although they include non-convexities, they also include a fair amount of convex bids. Figure 3 shows the daily *cleared* volumes of INC, DEC and UTC as well as day-ahead self-schedule bids as a percentage of the total day-ahead load. Table 1 also provides more figures of the hourly average volume of the different category of bids, focusing on the period 2018 to 2021 which will be used in our empirical analysis.⁷ We observe that the volume these convex bids is a significant share of the market.

Side payments. Because an equilibrium is not guaranteed to exist in a non-convex market, PJM relies on discriminatory side payments (sometimes called “uplift payments”⁸) in

⁷We report here the *cleared* volume of virtual bids (and not the bided volume). Self-scheduling are bids at a price of zero, thus bided and cleared volume should be the same. The non-convex dispatchable bids are computed using the data of dispatchable bids and keeping only the ones with positive fixed costs.

⁸“Uplifts” payments are paid for various reasons which are not all related to the non-convexities *per se*, such as ancillary services uplifts and reliability services uplifts (cf. PJM manual 28 on Operating Agreement Accounting for rules on how these uplifts are computed (PJM, 2025a)). The side payments directly related to the non-convexities, which are shown in Figure 4 are the “operating reserve uplifts” paid both in the day-ahead and in the real-time markets, which correspond to: “The total resource offer amount for generation, including startup and no-load costs as applicable, is compared to its total energy market value for specified operating period segments during the day [...] If the total value is less than the offer amount, the difference

Average hourly volumes	GW
Day-ahead load	88
<i>Physical bids:</i>	
– Available non-convex dispatchable bids	139
– Day-ahead self-schedule bids	25
<i>Virtual bids:</i>	
– Cleared INC bids	2.5
– Cleared DEC bids	3.8
– Cleared UTC bids	15.4

Table 1: Average hourly volume of bids per category, over the period 2018 to 2021 in PJM.

order to restore the incentives of market participants to follow the market allocation. We use these side payments as a measure of how far the price-allocation pair stands from an equilibrium. It can be shown more formally that when an equilibrium does not exist, the distance to an equilibrium is measured by the side payments (more formally, by the “lost opportunity costs”⁹, cf. Stevens et al. (2024)). Small side payments are an indication that the cleared quantities and the uniform price are close to be an equilibrium, while large side payments are an indication the market outcome is further away from an equilibrium. Large side payments indicate that the price signal does not reflect some relevant cost information. Figure 4 shows the monthly side payments related to the non-convexities. We observe that they stand for 7.9M\$/month and that most of these side payments are paid in real-time rather than in day-ahead.

3 Model

Model environment. Let us consider a sequence of two markets: the day-ahead (DA) forward market and the real-time (RT) spot market. This sequence of markets are represented on Figure 5. There are two types of agents: the virtual bidders, active in the DA market, who play as arbitrageur; and the physical bidders, active in both the DA and RT markets, who are producers of electricity, tied to some generation assets. The physical bidders incur both a fixed start-up cost s and a variable cost of production. These physical producers further split into two categories: the fast-start resources (FSR) and the slow-start resources (SSR). The FSR are flexible assets, which can be started-up in real-time even though they are not committed in the DA market. On the contrary, the SSR that are not committed in DA cannot participate in the RT market (this assumption aligns with Figure

is credited to the PJM Member” (PJM, 2025a) A unit whose day-ahead revenue, at the day-ahead price, does not cover his cost receives a DA operating reserve uplift credit. Then, in real-time, this unit might have a schedule that deviates from its day-ahead plan. She then receives balancing (real-time) operating reserve uplift credits for loss exposure in the spot market due to following PJM dispatch. We discuss more the different types of uplifts in Appendix C.

⁹Starr (1969) and Arrow and Hahn (1971) use total excess supply or demand to measure the distance to equilibrium: given a price, each participant maximizes her profit but the market clearing constraint might be violated as an equilibrium does not exist. This market clearing constraint violation (excess supply or demand) is used as a measure of distance to equilibrium. Another way to measure the distance to equilibrium, which is natural in context of auctions where the auctioneer determine both the price and the cleared quantities, is with the concept of *lost opportunity costs* (cf. Stevens et al. (2024) for a formal argument and a detailed discussion). By definition, the quantities cleared by the auctioneer satisfy the market clearing constraint, but the cleared quantities might not maximize the profit of the market participants given the cleared price: they violate envy-freeness. The lost opportunity costs measure the monetary incentives that market participants have to deviate from the cleared quantities. In particular, some market participants might be losing money given the cleared price and quantities (cf. the example of Figure 1). Therefore, the auctioneer complements the uniform price with side payments to make sure that exposed power plants break even. These side payments typically only stand for a share of the lost opportunity costs, as some “lost opportunities” are not compensated (e.g. a market participant that is not producing, but could have made a positive profit at the cleared price is not compensated for his lost opportunity). Thus the side payments is only approximative measure of the distance to equilibrium.




Figure 4: PJM side payments due to non-convexities in day-ahead and real-time.




Figure 5: Model environment: market timing.

2a, where we saw that a large share of the units committed in day-ahead self-schedule in real-time, thus are inflexible). An SSR unit who is committed in DA can still be turned off in RT but the start-up cost s is sunk. There is an exogenous inelastic demand which is made of a predictable load d and a random shock $\epsilon \in [-\bar{D}, \bar{D}]$ on demand, which is uniformly distributed: in DA, the demand is d ($> \bar{D}$), while in RT, the demand is $d + \epsilon$. This exogenous shock should be viewed as representing all sorts of uncertainties that arise between day-ahead and real-time and which affect the net load served by SSR and FSR: uncertainties on the load, on wind production, on the availability of some base-load units (for instance due to forced outages), or on imports, exports and grid constraints.

We assume physical bidders bid truthfully in the sequential markets and that they are myopic: they do not anticipate the real-time conditions when bidding in day-ahead. To qualify this seemingly restrictive assumption, let us stress that physical assets could also participate as virtual bidders. Thus the myopic assumption should be viewed as an assumption that the arbitrage is made entirely through the virtual bids. In the model, we consider a continuum of slow start-units with a total cost equal to $C_S(q) = \frac{c_S}{2}q^2 + s_Sq$ and a continuum of fast start-units with a total cost equal to $C_F(q) = c_Fq + s_Fq$. There is a base-load slow-start unit of capacity \bar{d} with marginal cost equal to 0 and a fixed cost of $s_S\bar{d}$. We assume $c_F > c_S\bar{D} + s_S$ and $s_F > s_S$. That is, fast-start units are always more expensive than slow-start units. These settings are represented on Figure 6.

Note that our modelling assumptions (a two-stage market with perfectly inelastic demand, random real-time shocks, and higher-cost resources in real time than in the day-ahead market) provide a reasonable approximation of PJM market operations (cf. section 2). They are also consistent with standard frameworks of financial trading encountered in

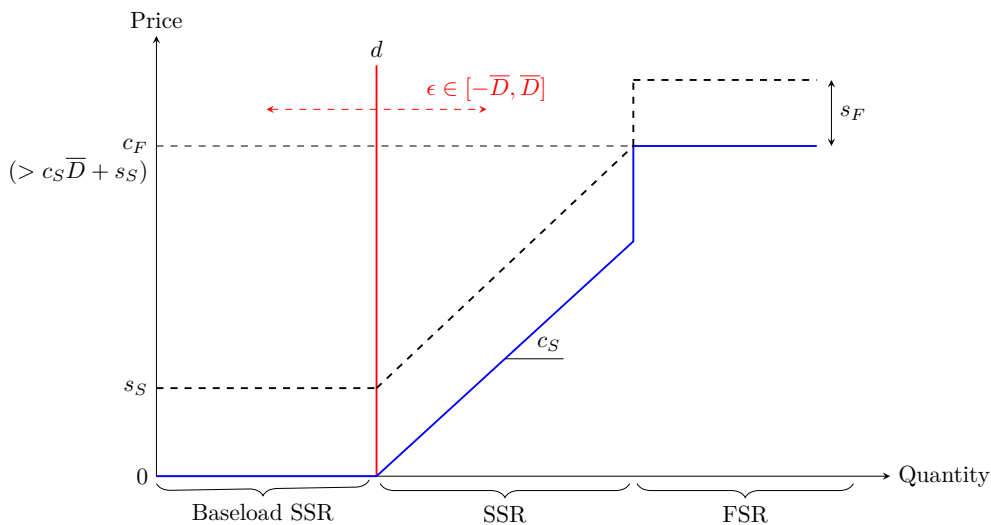


Figure 6: Model environment: market participants.

the literature (Tang et al., 2016; Mather et al., 2017). In particular, the distinction between slow- and fast-start resources is well established in the electricity market literature (Eldridge et al., 2023; Kazempour and Hobbs, 2017a,b).

Mechanism description. Our model of a *non-convex* market with arbitragers is slightly artificial, for it is in fact *convex*. This modeling approach is justified by the fact that including non-convexities renders derivation of analytical results intractable. Yet, by making some assumptions on the allocation rule and settlement rule of the mechanism, we can mimic the behaviour of a non-convex auction. We shall solve numerically at the end of this section an instance of a truly non-convex market to illustrate and confirm our analytical findings. Meanwhile, we assume the following mechanism, illustrated in Figure 7:

- *Allocation rule:* the auctioneer clears the allocation that maximizes social welfare. That is: demand is cleared as long as its willingness-to-pay is higher than average cost of the next infinitesimal producers. Figure 7 shows the two main scenarios. In scenario VB_1 , the demand curve from the virtual bids intersects horizontally with the average cost curve of the SSR. These bids are thus partially cleared by the auctioneer until they cross with the average cost of SSR. Indeed, clearing VB_1 entirely would lower social welfare as their willingness-to-pay is lower than the average cost of the SSR above their curve. In scenario VB_2 , the demand curve intersects vertically with the average cost curve of the SSR. These bids are fully cleared by the auctioneer.
- *Pricing and settlement rule:* the auctioneer uses marginal pricing rule, thus the marginal agent sets the price. In scenario VB_1 , the virtual bids are partially cleared and thus set the price: $p_1 = VB_1$. In scenario VB_2 , the virtual bids are fully cleared and the price is set by the marginal SSR: $p_2 = c_S VB_2$. Producers who do not break even receive side payments to offset their losses.

An important observation is that the fixed costs s_F and s_S enter the allocation rule's computations, whereas they are not necessarily accounted for in the pricing rule. This mimics the typical two-stage procedures used by electricity auctioneers who first solve a comprehensive unit commitment model to compute the allocation and then solve a convex restriction of the original problem in which the commitment decisions are fixed, and thus the startup costs omitted from the pricing problem (O'Neill et al., 2005).




Figure 7: Model environment: allocation and pricing rules.

Equilibrium without virtual trading. Without virtual trading, the equilibrium prices and side payments in day-ahead and real-time are:

$$p_{DA} = 0, \quad SP_{DA} = s_S d \quad (1)$$

$$\mathbb{E}_\epsilon(p_{RT,\epsilon}) = \frac{c_F}{2}, \quad \mathbb{E}_\epsilon(SP_{RT,\epsilon}) = \frac{s_F \bar{D}}{4} \quad (2)$$

In day-ahead, the lumpy baseload unit covers the entire production, and there is no other activation of slow-start units. This is followed in real-time by lumpy activations of fast-start units, depending on the shock ϵ . The baseload unit sets the price to zero in day-ahead, as well as in real-time when there is a negative shock. When the shock is positive, the fast-start units set the price in real-time. There is a price spread between DA and RT markets—as we shall see, virtual traders will arbitrage away this spread. As the price does not cover the fixed costs, side payments are paid in DA to cover the fixed costs of the base-load SSR. Side payments are also paid in RT to cover the fixed costs related to the activation of FSR. The total expected cost corresponds to both the baseload unit cost and the expected cost of activating fast-start units:

$$Cost = s_S d + \frac{\bar{D}}{4}(c_F + s_F) \quad (3)$$

Equilibrium with virtual trading. Let us consider the case of “perfect” virtual trading, where virtual traders compete *à la Bertrand*: they compete by outbidding each other until the no-arbitrage condition holds. This corresponds to arbitragers who bid an infinite quantity at a price P such that $p_{DA} = \mathbb{E}_\epsilon(p_{RT,\epsilon})$. As shown in Figure 2b, although actual virtual traders in PJM do not submit bids at the exact same price P , their bids remain relatively flat around the market price.

Virtual bidders anticipate the risk of real-time activations of fast-start units and their impact on prices. As a consequence, they buy more quantities in day-ahead, leading to the activation of k slow-start units. The expected real-time price depends on the number of slow-start units k cleared in day-ahead: this is expressed by equation (4). Furthermore, in order for unit k to be cleared, the demand from virtual bidder should exceed its average cost (cf. the allocation and pricing rules above): this is expressed by equation (5). Thus,

the equilibrium prices as a function of k are:

$$p_{DA} = P = c_S k + s_S \quad (4)$$

$$\mathbb{E}_\epsilon(p_{RT,\epsilon}) = \frac{1}{2\bar{D}} \left(\int_0^k c_S \epsilon d\epsilon + \int_k^{\bar{D}} c_F d\epsilon \right) = \frac{1}{2\bar{D}} \left(\frac{c_S}{2} k^2 - c_F k + c_F \bar{D} \right) \quad (5)$$

Using the no-arbitrage condition, we get:

$$0 = \frac{c_S}{2} k^2 + k(-c_F - 2c_S \bar{D}) + \bar{D}(c_F - 2s_S) \quad (6)$$

Solving the second-order equation, given $0 \leq k \leq \bar{D}$ (notice that one need $c_F \geq 2s_S$ to have $k^* \geq 0$), leads to

$$k^* = \frac{c_F + 2c_S \bar{D} - \sqrt{c_F^2 + (2c_S \bar{D})^2 + 2c_F c_S \bar{D} + 4c_S s_S \bar{D}}}{c_S} \quad (7)$$

Inclusion of virtual trading leads to activation of k^* SSR and reduces the price spread between day-ahead and real-time from $\frac{c_F}{2}$ to zero. Simple comparative statics on k^* leads to the following results: the lower the competitive advantage of SSR with respect to FSR, the fewer SSR are committed in day-ahead following virtual bidders' strategy.

Proposition 1. *Under perfect virtual trading, k^* slow-start resources are committed in day-ahead, such that (focusing on the case $k^* > 0$):*

1. $\partial k^* / \partial s_S < 0$: the higher the fixed cost of slow-start units, the lower k^* .
2. $\partial k^* / \partial c_F > 0$: the higher the variable cost of fast-start units, the higher k^* .
3. $\partial k^* / \partial c_S < 0$: the higher the variable cost of slow-start units, the lower k^* .
4. $\partial k^* / \partial s_F = 0$: k^* is insensitive to the fixed cost of fast-start units.

The equilibrium expected side payments in the market that includes virtual trading are then as follows:

$$SP_{DA} = 0 \quad (8)$$

$$\mathbb{E}_\epsilon(SP_{RT,\epsilon}) = \frac{1}{2\bar{D}} \int_k^{\bar{D}} s_F (\epsilon - k) d\epsilon = \frac{s_F \bar{D}}{4} - \frac{s_F k}{2} \left(1 - \frac{k}{2\bar{D}} \right) \quad (9)$$

As far as side payments in day-ahead are concerned, we see that there is a discrete effect: introducing virtual trading, in our model, fully eliminates side payments. Regarding the side payments in real-time, the two main observations are: they are reduced but not eliminated, and they are linearly increasing with s_F , the fixed costs of FSR.

Proposition 2 (Side payments 1). *Perfect virtual trading decreases expected side payments both in day-ahead and in real-time.*

Proposition 3 (Side payments 2). *Real-time side payments are linearly increasing with s_F .*

This is illustrated in Figure 8. There, we see k^* reduces side payments with respect to the case of no virtual trading (cf. Proposition 2). The figure also shows that the side payments scale linearly with s_F (cf. Proposition 3): given k^* , whose value does not change with s_F , going from $s_F = 0.5$ to $s_F = 1$ implies the same increase in side payments as going from $s_F = 1$ to $s_F = 1.5$.

These results follow virtual trading bidding behaviour, which has two main features: (i) they anticipate real-time conditions and internalise the fixed costs of the physical resources




Figure 8: Side payments as a function of k (with $d = 2$, $\bar{D} = 1$, $c_S = 1$, $s_S = 0.5$, $c_F = 2$, $t = 0.25$) and three values of s_F : $s_F = 0.5$, $s_F = 1$ and $s_F = 1.5$.

(cf. equation (4)) and (ii) they set the price in DA. Together, this leads to a day-ahead price that reflects the fixed costs, thus side payments drop to zero ; and it also entails better commitment decisions in day-ahead that lead to fewer activations of fast start resources in real-time, thus less side payments in real-time.

Let us make two remarks. First, real-time side payments are not eliminated by virtual trading. Because the physical real-time market involves non-convexities in the production, the uniform price alone might not clear the market and still requires side payments. One way—which we will not discuss further in our paper—to improve these issues is by amending the auction pricing rule itself, deviating from marginal pricing, to reflect some of the fixed costs into the price signal. In our model, this would for instance mean setting the price to $c_F + s_F$ instead of c_F when FSR are activated in real-time. This pricing rule would eliminate the need for side payments in real-time, in our model. This is the spirit of the so-called “fast-start pricing approaches” discussed in the literature and implemented by some US ISOs. Our goal is thus not to argue that these discussions on the right pricing rule are irrelevant—it remains important and it is compatible with virtual trading—but to highlight the implications of virtual trading itself. Second, virtual trading fully eliminates day-ahead side payments. One caveat though is that our model assumes a competition *à la Bertrand* with an infinite quantity of virtual traders, bidding at a price P . Following our pricing rule definition, it implies the day-ahead price is set by P . If instead there was a limited volume of virtual traders, the demand curve could intersect the supply curve vertically (cf. Figure 7). As a consequence, SSR k would set the price to $c_S k$, which would entail some side payments to cover the fixed costs. Besides, our model also neglects the multi-period nature of the market, inter-temporal constraints and most unit commitment constraints that are encountered in actual markets. Thus, although our model predicts that virtual trading fully eliminates DA side payments, this result should be viewed as a theoretical benchmark, and observed outcomes may depart from it for several reasons (we come back to this in section 4).

Let us turn to the analysis of the expected costs. With virtual trading, the total expected

cost as a function of k is:

$$Cost(k) = s_S d + s_S k + \frac{1}{2\bar{D}} \left(\int_0^k \frac{c_S}{2} \epsilon^2 d\epsilon + \int_k^{\bar{D}} \frac{c_S k^2}{2} + (c_F + s_F)(\epsilon - k) d\epsilon \right) \quad (10)$$

$$= \underbrace{s_S d + \frac{\bar{D}}{4}(c_F + s_F)}_{\text{cost without VB}} + \underbrace{s_S k}_{\text{sunk cost}} - \underbrace{\left(\frac{k^3 c_S}{6\bar{D}} + \frac{k^2(c_F + s_F)}{4\bar{D}} + \frac{k^2 c_S}{4} - k \frac{(c_F + s_F)}{2} \right)}_{\text{expected cost savings}} \quad (11)$$

The cost is a sum of three terms: (i) the cost without virtual virtual trading, (ii) the sunk cost implied by committing k SSR and (iii) the cost savings entailed by committing k SSR. Comparing expression (11) with expression (3), direct mathematical derivations lead to the following observation:

Proposition 4 (Efficiency 1). *Perfect virtual trading improves cost efficiency compare to no virtual trading.*

Although virtual trading improves efficiency, it does not reproduce the first-best commitment decisions. The reason is that the virtual traders only react to the day-ahead and real-time price signals. With non-convexities, the price does not convey all the relevant cost information—mainly, it overlooks the fixed costs. For example, k^* does not depend on s_F : $\partial k^*/\partial s_F = 0$. While clearly, if $s_F \rightarrow \infty$, assuming s_S remains bounded, it would be optimal to commit all slow-start units, thus having $k = \bar{D}$.

Proposition 5 (Efficiency 2). *Perfect virtual trading does not reproduce the first-best efficient day-ahead commitment plan k^{**} : it entails lower commitments of SSR than optimal, thus $k^* < k^{**}$.*

Propositions 4 and 5 are illustrated on Figure 9. Virtual trading entails committing k^* SSR in day-ahead which reduces the overall costs compared to the situation without virtual trading (Proposition 4). Yet, it does not lead to the first-best commitment plan k^{**} (the optimum of social planner’s two-stage stochastic optimization problem) but leads to under-commitment, thus $k^* < k^{**}$ (Propositions 5).¹⁰

Let us finally remark that, in *some* scenarios, virtual trading may reduce efficiency. For instance, for $\epsilon = -\bar{D}$, the cost would have been lower without financial trading. However, what matters is that *in expectation* financial trading improves efficiency.¹¹

Equilibrium with costly virtual trading. Let us consider the case of “costly” virtual trading where a transaction cost t is introduced. Virtual traders competing *à la Bertrand* now leads to the following arbitrage condition: $\mathbb{E}_\epsilon(p_{RT,\epsilon}) - p_{DA} = t$, thus the price spread must exactly refund the transaction fee. Using a similar reasoning as previously, one can compute the SSR commitments k^t resulting from virtual trading with a transaction cost t . This corresponds to the $k^t \in [0, \bar{D}]$ solving:

$$0 = \frac{c_S}{2} k^2 + k(-c_F - 2c_S \bar{D}) + \bar{D}(c_F - 2s_S - 2t) \quad (12)$$

Condition (12) is the counterpart of (6) when transaction cost t is included. Comparing (12) and (6) directly leads to $k^t \leq k^*$ (notice that one need $t \leq \frac{c_F - 2s_S}{2}$ to have $k_t \geq 0$), and

¹⁰Notice that $s_F = 0$ does not imply $k^* = k^{**}$: the fixed costs s_S also distort incentives as they are reflected in the day-ahead price through the virtual traders but not in the real-time prices. More surprisingly, $s_F = s_S = 0$ would not imply $k^* = k^{**}$ either, in our model. If $s_S = 0$ then $k^{**} = \bar{D}$: there is no stranded cost involved in committing SSR thus it is optimal to commit them all. While $k^* < \bar{D}$ for $s_S = 0$, as $k^* = \bar{D}$ would imply the day-ahead price is always greater than the real-time price, thus the virtual traders would not achieve this full commitment outcome. The reason of this failure is that the commitment decisions in day-ahead entail some inflexibility in real-time. Even with $s_F = s_S = 0$, the commitment leads to virtual trading that fails to be fully efficient.

¹¹PJM (2015) argued that virtual trading can deter efficiency based on some deterministic examples that assume an exogenous bidding behavior for financial traders, and as we argue, what matters is what happens *in expectation* and *at equilibrium*, and we further discuss this in section 6.




Figure 9: Total expected costs as a function of k (with $d = 2$, $\bar{D} = 1$, $c_S = 1$, $s_S = 0.5$, $c_F = 2$, $s_F = 1$, $t = 0.25$).

$\frac{\partial k^t}{\partial t} < 0$, thus the higher the fee the lower k^t . A transaction fee implies that less financial trading volume gets cleared by the auction, and the higher the fee t the lower the cleared volume.

All the equilibrium results derived above as functions of k still hold, substituting k with k^t (the prices in equations (4)–(5), the side payments in equations (8)–(9) and the cost in equation (11)). Concretely, the side payments are reduced compared to no-virtual trading but this improvement shrinks as t increases. This is illustrated on Figure 8. Similarly, the cost is improved compared to no-virtual trading but this improvement shrinks as t increases. This is illustrated on Figure 9.

Proposition 6 (VB with transaction cost t). *Virtual trading with transaction cost t leads to $k^t \leq k^*$ such that:*

- $Cost(k^*) \leq Cost(k^t) \leq Cost(0)$
- $SP(k^*) \leq SP(k^t) \leq SP(0)$

Main takeaways. To sum-up, there are three main takeaways from our model. (i) First, virtual trading improves price convergence between DA and RT, although a transaction fee t would entail a persistent price difference of t . (ii) Second, more virtual trading leads to better commitment decisions in DA (more activations of slow-start units), which lowers the expected costs. Both the improvement of price convergence and the reduction of total cost have been analysed in the literature (see, for instance, Mather et al. (2017) for an analytical model, Kazempour and Hobbs (2017a,b) for numerical simulations relying on a comprehensive electricity auction model and Jha and Wolak (2023) for empirical evidences from CAISO market). (iii) The third effect—and most important for our analysis—is that higher virtual trading mitigates the need for side payments, thus it leads to a uniform pricing outcome closer to an equilibrium. Large side payments means some cost information is not well reflected in the market price and is instead handled by out-of-market payments, while low side payments means the uniform price better signals the overall cost information. Here, there are two distinct effects. In DA, virtual trading set the price leading to prices that endogenise the fixed costs of SSR which eliminates the day-ahead side payments. In RT, as virtual traders have anticipated the risk of FSR activation, the likelihood of lumpy FSR activations is reduced which mitigates the associated side payments.

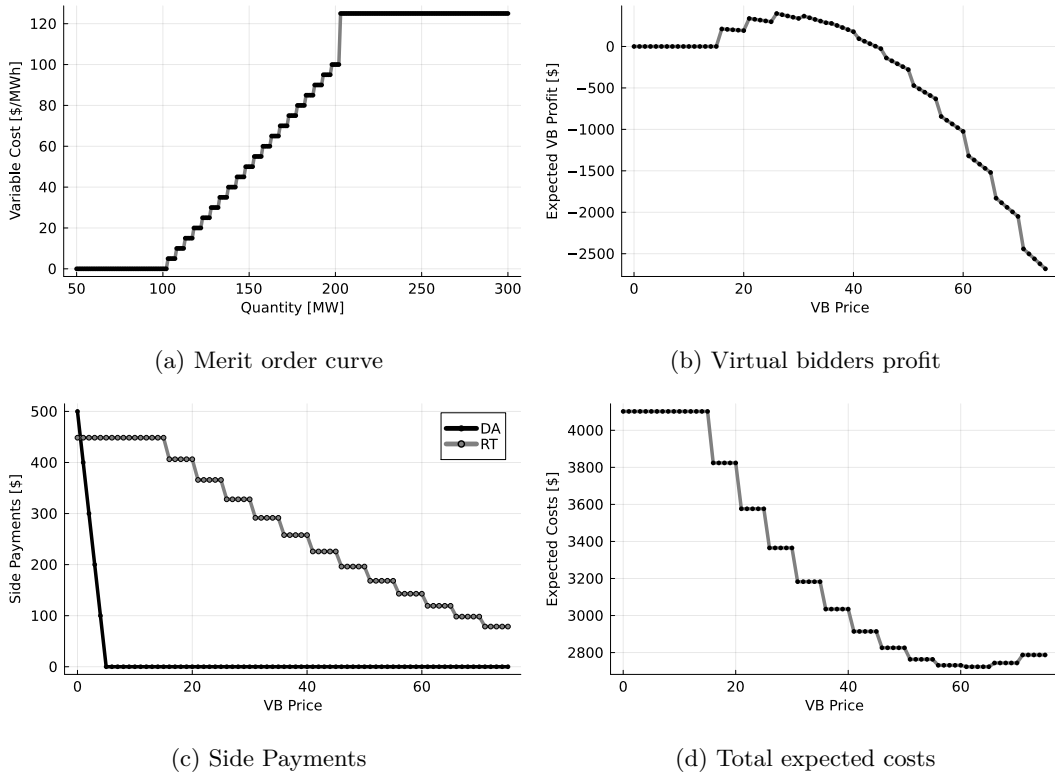


Figure 10: Results of the discrete example.

A discrete example Let us consider the following non-convex electricity market. There are 20 SSR and 20 FSR, with a capacity of 5MW each and a minimum production limit of 2 MW each. The fixed cost of SSR are \$50, while it is \$85 for the FSR. The variable cost of production is $i \times 5\$/\text{MWh}$ for $i \in \{1, 2\}$ for the SSR and 125\$/MWh for the FSR. There is also a baseload SSR unit of 100MW with variable cost 0 and fixed cost equal to \$500. The merit order curve of this example is shown on Figure 10a. $d = 100\text{MW}$ and the demand shocks ϵ are uniformly distributed in the set $\{2i | i \in -50 \dots 50\}$. We consider the case of perfect virtual trading with a large volume of VB who choose a bid-price strategy P . As in our model, the auctioneer clears the surplus-maximizing allocation and uses marginal pricing complemented by side payments, both for the day-ahead and real-time auctions. This instance is fairly close to the settings of our model, except that the previous continuum of agents is replaced by a finite number of lumpy producers.

Figure 10 shows the market outcomes, as a function of the bid price strategy of the virtual traders. No virtual trader corresponds to $P^* = 0$, as no virtual trader would get cleared at this price. Perfect virtual trading corresponds to full arbitrage between day-ahead and real-time prices, leading to zero profit for the virtual traders. Here, this corresponds to the strategy $P^* \approx 44$ (cf. Figure 10b). This strategy entails dispatching 30MW of SSR (thus 6 SSR commitments) in day-ahead. As in our model, this leads to cost reduction compared to no virtual trading (cf. Proposition 4), as observed on Figure 10d, although it does not quite minimize the expected costs (cf. Proposition 5). As observed on Figure 10c, the day-ahead side payments are null and the real-time side payments decrease—almost cut by two—compared to no virtual trading (cf. Proposition 2).

The outcome in presence of a transaction cost t can also be deduced from the figures. Instead of zero profit, a transaction cost t implies the virtual traders should make a per-MW profit of t , thus shifting their bidding strategy to the left on Figure 10b, increasing both costs and side payments compared to the perfect virtual trading case (cf. Figures 10d and




Figure 11: PJM cleared virtual bids (INC, DEC and UTC). The vertical line, on November 1, 2020, corresponds to the introduction of transaction fees on UTC bids.

10c), as expected from Proposition 6. These simulations also hint on what would happen in case of imperfect competition among virtual bidders. The price strategy maximizing virtual bidders’ profit in Figure 10b ($P^* = 26$) should be read as the outcome if virtual bidders behave as a monopsony instead of behaving *à la Bertrand*. We observe that imperfect virtual trading—either monopsony virtual bidder or virtual trading with transaction cost t —improves both costs and side payments compared to no virtual bidders, although it does not perform as well as perfect virtual trading.

We ran some more simulations introducing heterogeneity in the portfolio of power plants. The results confirm the previous observations and are reported in the appendix B.

4 Empirical strategy and data

Methodology. In order to test the predictions of our model, we analyze the effect of a policy that impacted the transaction fees paid by some of the virtual traders in PJM Interconnection market. On November 1, 2020, a transaction fee was introduced on the UTC virtual bids (Monitoring Analytics, 2021; FERC, 2020). INC and DEC virtual bids were already subject to transaction fees while UTC bids were not. This was judged “unjust, unreasonable, and unduly preferential”¹² by the FERC, the US energy regulator, who decided that the UTC bids should pay day-ahead and balancing operating reserve charges equivalent to a DEC at the UTC sink point (FERC, 2020). The transaction fee applies to cleared quantities (as in our model of section 3).

Following the implementation of this policy, the cleared volume of UTC virtual bids dropped significantly. This is shown on Figure 11 where the vertical red line corresponds to the implementation of the new transaction fee. We see this transaction fee did not substantially affect the cleared volumes of other virtual bids (INC and DEC) but led to a significant decrease of cleared UTC bids (it is worth noting that, at the time, UTC represented the largest share of total financial trading volume, cf. Table 2). Overall,

¹²“We find that PJM’s current uplift allocation rules are unjust, unreasonable, and unduly preferential because they do not allocate uplift to UTCs. Accordingly, we direct PJM to submit, in a compliance filing within 45 days from the date of this order, a replacement rate that revises PJM’s current uplift allocation rules to allocate uplift to UTCs in a manner that treats a UTC, for uplift allocation purposes, as if the UTC were equivalent to a DEC at the sink point of the UTC. As a result, under the replacement rate, UTCs will be allocated both real-time uplift and day-ahead uplift.” (FERC, 2020)

	INC		DEC		UTC	
	Pre-Treat.	Post-Treat.	Pre-Treat.	Post-Treat.	Pre-Treat.	Post-Treat.
Mean	64.3	55.9	87.2	102.2	456.6	180.3
Median	62.5	55.3	83.3	99.8	462.1	175.3
Stand. Dev.	16.5	15.7	28.6	27.8	101.9	46.5
Max	126.7	104.5	181.0	228.9	749.8	341.1
# observations	978	298	978	298	978	298

Table 2: Daily cleared volume of virtual bids [GWh/day], pre- and post-treatment.

the total hourly average cleared volume of virtual bids (INC, DEC and UTC) went from 25GW/h to 14GW/h following the introduction of the transaction fee.

This drop in the cleared volume is aligned with economic expectations and with our model predictions: a transaction fee impacts the profit-maximizing problem of the arbitrageurs, who offset their bid to cover the fee, leading to a reduction of the cleared volume. Our model predicts that the decline in financial trading volume should coincide with an increase in side payments (cf. Proposition 6). Did the inclusion of a transaction fee on November 1, 2020, and the associated sharp drop in virtual trading volume, indeed lead to an increase in side payments, as predicted by our model?

To analyze the effect of this policy, we leverage the public market data of PJM¹³. PJM publishes comprehensive market data including the detailed physical bids, the virtual bids, the load, the self-scheduling volume, data on side payments, data related to grid congestion and outages, etc. We focus on the period from end of February 2018 to September 2021. This period is selected for it is stable in the market rules, aside from the policy change on UTC transaction fee. On February 20, 2018, FERC issued an order which limit the number of bidding points at which virtual bids may be submitted by market participants (Docket No. ER18-88, FERC (2018)). This changed the scope and nature of virtual bids, thus we start our dataset after February 20, 2018. Besides, on September 1, 2021, PJM implemented a new “fast-start pricing” rule on both the day-ahead and real-time markets.¹⁴ Up to September 2021, PJM relied on the classic marginal cost pricing rule, as in our model. Concretely, this new pricing rule deviates from marginal pricing for it allows the fixed non-convex costs of fast-start resources to be reflected in the price signal. This is a disruptive change in the pricing rule, thus we end our dataset on September 1, 2021.¹⁵ Overall, our database includes 1,276 days, excluding 12 days between February 20, 2018, and September 1, 2021, because of missing data.

Our main variables of interest are the day-ahead and real-time side payments paid by PJM, i.e. the counterpart of SP_{DA} and SP_{RT} in our model of section 3. As it should be

¹³except the Henry hub gas price time series which is sourced from EIA, the US Energy Information Administration.

¹⁴Extensive discussions about “fast-start” pricing at PJM started in 2017 (PJM, 2017), following the trend initiated by the US Regulatory Commission which launched a consultation about price formation in power auctions in 2014 (FERC, 2014). This new pricing approach, sometimes referred to as “extended LMP”, consists of taking the linear programming (LP) relaxation of the unit commitment problem. The main difference between fast-start pricing and extended LMP discussed in the scientific literature, is that fast-start pricing only consider the LP relaxation for a limited set of resources, namely those that are “fast start”. Roughly said, fast-start pricing deviates from marginal pricing for it partly reflects the start-up costs, and other operational fixed costs, in the price signal, cf. the discussion in section 3 about including s_F in the price signal.

¹⁵Note that, the resources qualifying as “fast-start” for this pricing rule stand for a small share of all resources. Although this pricing policy is disruptive in theory, because of the limited number of resources which are effectively treated as “fast-start”, the effect of fast-start pricing was expected to be moderate, according to simulations conducted by PJM in 2019 (PJM, 2019). Ex post, it appears the magnitude of its impact has been almost nonexistent in day-ahead and moderate in real-time, although some months exhibit several percentage point difference in prices between fast-start pricing and marginal pricing (Monitoring Analytics, 2022, 2025b) (with the caveat that these *simulated* marginal prices would have been different if they were actually used for settlement, since the bidding behaviour would likely have been different as well).

	Day-ahead			Real-time		
	All Days	Pre-Treat.	Post-Treat.	All Days	Pre-Treat.	Post-Treat.
Mean	47,620	51,925	33,494	215,705	167,351	374,400
Median	7582	7471	7595	109,672	90,588	220,052
Stand. Dev.	86,722	92,403	62,746	296,569	210,532	445,874
Max	958,693	958,693	646,474	2,604,937	1,789,976	2,604,937
% of 0	22%	24%	15.7%	0.08%	0.1%	0%
# observations	1276	978	298	1276	978	298

Table 3: Day-ahead and real-time side payments [\$/day], pre- and post-treatment.




Figure 12: Days per month with zero side payments in day-ahead in PJM market.

clear by now, side payments are used as a measure of distance to equilibrium: they are the discriminatory payment made to market participants to compensate for their losses at the market price. PJM provides data of the daily amount of side payments (these payments indeed do not have a finer, hourly, granularity). Table 3 provides descriptive statistics of day-ahead and real-time daily side payments. Let us highlight two main observations. First, most of the side payments due to non-convexities are paid in real-time, as already observed on Figure 4: average real-time side payments are about four times larger in real-time than in day-ahead. This should mostly be read as side payments related to the activation of fast-start resources in real-time, as in our model. Second, 22% of the days exhibit zero day-ahead side payments, which means about 6 days per month. Figure 12 shows the number of days per month with zero day-ahead side payments.

To estimate the impact of the policy on the side payments, we estimate two different models: one for day-ahead side payments and one for real-time side payments.

Real-time model. For the real-time side payments, we estimate the following ordinary least squares (OLS) regression:

$$SP_{RT,t} = \alpha + \beta_1 \text{Treatment}_t + \beta_2 s_{F,t} + \beta_3 \text{loadShock}_t + \eta \mathbf{X}_t + u_t \quad (13)$$

where $SP_{RT,t}$ corresponds to the real-time side payments for each day t , Treatment_t is the indicator variable representing the introduction of a transaction fee on virtual bids on November 1, 2020, $s_{F,t}$ are the fixed costs of FSR, loadShock_t represents the load shocks and \mathbf{X}_t are the control variables. The main parameter of interest is β_1 which measures the treatment effect, thus by how much did the transaction fee increased the amount of




Figure 13: Henry Hub gas price and PJM fixed costs of fast-start resources.

real-time side payments. From our model we expect $\beta_1 > 0$: a transaction fee increases the amount of side payments needed to sustain an equilibrium.

In addition to the impact of financial trading, our model also points to a direct effect of both FSR fixed costs (i.e. s_F in our model) and load shocks (i.e. $\epsilon \in [-\bar{D}, \bar{D}]$ in our model) on side payments. These effects are captured by the parameters β_2 , which is the effect of the fixed costs of FSR, and β_3 , which is the effect of the load shock between day-ahead and real-time. We briefly comment on the FSR data and load shock data before turning to the control variables.

FSR BIDS DATA—PJM publishes comprehensive bid data of dispatchable resources, including, for each resource, the hourly offered curve, start-up and no-load costs, maximum and minimum production limit, start time and minimum up and down time. We use these bid data to construct a time series of the fast-start resource fixed costs (the counterpart of parameter s_F in our model of section 3). In PJM, eligible fast-start resources should have a total time to start and minimum run time less than or equal to one hour (PJM, 2025b). Using these criteria and focusing on non-convex resources (with positive fixed costs), leads to about 180 units that are “fast-start”.¹⁶ These resources have a capacity mostly between 0 and 100MW, with a median at 43MW (cf. the distribution in Figure C.4 in appendix). Each one of these resources has a fixed cost which corresponds to the sum of the start-up cost and the no-load cost, and which is mostly between 0 and 10,000\$ (cf. the distribution in Figure C.4 in appendix). We aggregate the fixed costs of these resources by taking a weighted average, where the plant capacity is used as weights: the weighted average FSR fixed cost is around 4000\$. Figure 13 shows the output time series as well as the Henry hub gas price, which explains some but not all the FSR fixed costs variations. This time series corresponds to $s_{F,t}$ in equation (13). Parameter β_2 measures by how much the real-time side payments increase when the fixed costs of FSR increases. From our model (cf. Proposition 3), we expect $\beta_2 > 0$: more fixed costs, typically refunded by side payments, naturally increases the amount of side payments.

LOAD SHOCKS—We compute load shocks (or load forecast error) as the hourly square difference between DA and RT loads. This time series corresponds to $loadShock_t$ in equation (13). We interpret it as the counterpart to the load shock $\epsilon \in [-\bar{D}, \bar{D}]$ in our model of section 3. Figure 14 shows the monthly average of the load shocks and compare it with the daily load and daily load standard deviation. The average daily load error is 12GW, thus 500MW/h, meaning 0.5% of the hourly load. We observe that load shocks tend to be

¹⁶The public figures available speak about 200 to 300 resources (Monitoring Analytics, 2025b).




Figure 14: Real-time daily load, daily load standard deviation and daily load error.

higher during the summer months, when both the load and load variation are higher. We expect $\beta_3 > 0$: larger shocks implies more FSR activations thus larger side payments.

OTHER MARKET CONTROLS—Besides the drivers of side payments captured in the model of section 3, the actual PJM market involves additional complexities that can affect the magnitude of side payments, such as the multi-period nature of the market or the presence of transmission grid constraints. PJM provides a list of what they consider to be the main drivers of side payments which, besides generator market offers and load error between day-ahead and real-time, also includes (i) the load, (ii) the grid congestions, (iii) generation availability, (iv) the amount of self-scheduling as well as (v) emergency procedure events (PJM, 2026). These data are included as control variables \mathbf{X} in equation (13). We briefly describe these data below (more in depth description of the data is provided in Appendix C). (i) As far as the load is concerned, higher load is associated with more commitment of resources, which can give rises to more side payments. High load variation can also put the system under stress and lead to side payments. PJM load has a clear seasonality, with both a summer and a winter peak (cf. Figure 14). Yet, the daily load profile exhibits significantly more variations during the summer peak rather than the winter peak. (ii) Grid congestions also affect the prevalence of non-convexity and the magnitude of side payments: congestions tend to make the market more fragmented, which exacerbates the non-convexities of the market. (iii) Unavailability of generation resources can also result into activations of costlier fast-start resources, at the right of the merit order curve, associated with more side payments. (iv) The volume of self-scheduling impact the amount of flexibility available in the system, which also affect side payments. (v) Finally, emergency procedure events can indicate a high level of stress in the system. Altogether, these elements can influence the side payments. For example, inspecting the data, the relatively low side payments observed from December 2019 to May 2020 on Figure 4 seem to be driven by favourable system conditions: a relatively mild winter (low load with low variability), fewer congestions, few emergency events, relatively low load error between DA and RT.

The control variables \mathbf{X} in the model of equation (13) are split in three categories which are *market controls* (Henry hub gas price, day-ahead price and standard deviation, real-time load and standard deviation), *date controls* (winter-summer and weekend-weekdays indicators) and *system stress controls* (DA and RT self-scheduling, DA congestion, generation availability—both active units, generation capacity—and emergency events).

Day-ahead model. As noted earlier in this section, day-ahead side payments data exhibit two main features relative to real-time side payments (Table 3): (i) they are substantially smaller in magnitude and (ii) they are equal to zero about 20% of the time. Note that this share of days with zero day-ahead side payments is consistent with our model’s prediction that virtual trading participation may eventually eliminate day-ahead side payments. Instead of an OLS model, we estimate the following binary response logit model:

$$\mathbb{P}(SP_{DA,t} > 0|x) = G(\alpha + \beta \text{Treatment}_t + \eta \mathbf{X}_t) \quad (14)$$

where G is the logistic function ($G(z) = e^z/(1 + e^z)$), $SP_{DA,t}$ corresponds to the day-ahead side payments at date t and $\mathbb{P}(SP_{DA,t} > 0|x)$ the probability of having positive side payments given the full set of explanatory variables x (here, the Treatment_t and the controls \mathbf{X}_t). Treatment_t is, as previously, the indicator variable representing the introduction of the transaction fee and \mathbf{X}_t are the control variables. The main parameter of interest is β which measures the treatment effect. We expect $\beta > 0$: there are higher chances of non-zero side payments when a transaction fee is introduced.

We use similar types of controls as for the real-time model, except that we obviously remove variables related to the real-time (we use DA load and not RT load, DA self-scheduling and not RT self-scheduling, etc.). Similarly to the real-time model, the control variables \mathbf{X} are split in three categories which are market controls (Henry hub gas price, day-ahead load and standard deviation), date controls (winter-summer and weekend-weekdays indicators) and system stress controls (DA self-scheduling and generation availability—both active units, generation capacity).

5 Main results

Tables 4 and 5 report the main results of both the real-time and day-ahead models. In both tables, column (I) reports the results of the full model, including all controls, for the period starting on February 20, 2018, to September 1, 2021. Columns (II) to (IV) correspond to models in which some controls are removed: columns (II) removes system stress controls, column (III) further removes date controls and column (IV) removes all controls. Columns (V) and (VI) are robustness tests on the time window. On February 20, 2018, the definition of UTC products changed. As we observe on Figure 11, this is followed by a period of roughly two months of steep increase of the UTC volume before it stabilizes, which likely corresponds to the time it took for traders to adapt their trading strategy. Column (V) filters these two first months (thus starting the model on April 20, 2018). On the other end of our time window, on September 1, 2021, PJM implemented a new pricing rule. Column (VI) adds four more months to the regression (thus terminating on December 31, 2021 instead of August 31, 2021).

As far as the real-time model is concerned, let us stress three main observations. First, we find an average treatment effect of 135k\$. Thus our model estimates that the introduction of the transaction fee on virtual bids increased the daily real-time side payments by 135k\$ (with a 95% confidence interval ranging from 103 to 168k\$). This represents an economically significant increase as, from Table 3, the pre-treatment magnitude of the daily side payments was 167k\$. Thus the transaction fee roughly increased by 80% the real-time side payments. Second, from columns (I) to (VI), these findings are statistically significant and they are robust to sensitivities on both the control variables and on the time window. Note that the slightly smaller effect in column (VI) is aligned with expectations since the implementation of the fast-start pricing policy on September 1st, 2021, is expected to have mitigated the side payments. Third, both the FSR fixed costs and the load shocks positively affect the side payments. All else equal, higher FSR fixed costs—typically recovered through side payments—naturally lead to higher side payments.¹⁷ Similarly, higher load shocks, which are expected to trigger FSR activations, also lead to higher side payments.

¹⁷The lower significance of FSR fixed cost compared to the other variables can be explained by the way

	(I)	(II)	(III)	(IV)	(V)	(VI)
Treatment Effect (in K dollars)	135.09*** (16.55)	159.79*** (16.24)	160.67*** (16.44)	207.78*** (17.93)	127.44*** (17.13)	108.75*** (16.77)
FSR fixed cost	33.97** (12.64)	11.4 (8.63)	6.51 (8.71)	31.59*** (9.57)	42.38** (13.29)	13.3 (10.49)
DA-RT load error	0.0003*** (3.7e-05)	0.0003*** (3.7e-05)	0.0003*** (3.7e-05)	0.0005*** (4e-05)	0.0003*** (3.9e-05)	0.0003*** (3.7e-05)
System stress controls	X				X	X
Date controls	X	X			X	X
Market controls	X	X	X		X	X
Sample size	1276	1276	1276	1276	1216	1395
R squared	0.46	0.42	0.41	0.21	0.46	0.41

Table 4: Real-time side payments results (OLS model).

	(I)	(II)	(III)	(IV)	(V)	(VI)
Logit Coefficient	0.6954*** (0.1996)	0.7156*** (0.1923)	0.5416** (0.1870)	0.5189** (0.1757)	0.7840*** (0.2005)	0.7232*** (0.1880)
System stress controls	X				X	X
Date controls	X	X			X	X
Market controls	X	X	X		X	X
Sample size	1276	1276	1276	1276	1216	1395

Table 5: Day-Ahead side payments results (logit model).

Regarding the day-ahead side payments, our model estimates a logit coefficient of 0.7. Thus, as expected, the introduction of the transaction fee increased the probability of having non-zero side payments in day-ahead. As for the real-time model, we find this estimate is statistically significant and robust to our sensitivities on both the control variables and on the time window. A logit coefficient of 0.7 means the treatment double the odds ($e^{0.695} \approx 2.0$) of having non-zero side payments in day-ahead. Given an average probability of non-zero day-ahead side payments before treatment was 76%, the logit coefficient means the treatment increased on average by +10% the chances of having non-zero side payments in day-ahead.¹⁸

6 Discussion and conclusions

The effect of financial trading in a non-convex electricity market has been subject to debate among stakeholders and in the scientific literature. In a note from 2015, PJM argued that virtual bids could in principle give rise to side payments by distorting the commitments made in day-ahead, leading to under or over commitment, associated with side payments

we measure them. They are the average of the FSR fixed cost, whereas we expect only the least costly FSR units to have an impact. This is an artefact of our model where the fixed cost of FSR are linear, whereas it is of course not the case in reality.

¹⁸ $p_1 = 0.76 = \frac{e^z}{1+e^z}$ from which we compute $e^z = 3.16$. Then $p_2 = \frac{e^z e^\beta}{1+e^z e^\beta}$ with $\beta = 0.695$ from which we compute $p_2 = 0.86$ thus $p_2 - p_1 = 0.1$.

(PJM, 2015). PJM’s argument was not based on an equilibrium analysis though, but on stylized examples that assume exogenously some bidding behaviour for virtual traders (while, as we argue in our model, what matters is what happens *at the equilibrium* and *in expectation*). From this analysis, PJM recommended charging the virtual bids a transaction fee to compensate for the side payments they may cause.

In 2016, Hogan argued against the idea of allocating uplift charges to virtual traders as this increase in transaction cost would deter participation and have material consequences on the efficiency of the market (Hogan, 2016).

Later in 2020, the FERC order that led to the implementation of the transaction fee on UTCs (which we analyzed in section 4) gave rise to contrasted viewpoints on the benefits of virtual trading. The utilities were in favor of this policy: “PJM Utilities Coalition conclude that because UTCs are transacting in the markets, they should receive their respective share of the costs that result from market operations.” (FERC, 2020) In contrast, some trading firms active in PJM market argued that “because UTCs promote price convergence and help with price formation, allocation of uplift to UTCs disincentivizes UTCs’ efficiency enhancing characteristics on high priced days when there is risk of high uplift.” (FERC, 2020)

Our model and empirical analysis contribute to shed light on this debate. Our article argues that the introduction of financial trading in an electricity market induces a *smoothing effect*. Financial traders are arbitrageurs that bid close to the market price, resulting in the addition of a large volume of convex bids close to the market margin. Our model shows that, by internalizing the fixed costs in their bids and anticipating the real-time system conditions, the financial traders mitigate the side payments needed to sustain an equilibrium. In day-ahead, virtual traders set the prices and as a consequence, fixed costs are reflected in the uniform price signal instead of being refunded by side payments. This eliminates day-ahead side payments (in the scope of our model). In real-time, higher price convergence and better commitment decisions in day-ahead lead to less lumpy activations of fast-start resources in real-time usually associated with side payments. Our empirical analysis confirm the predictions of our model. According to our analysis, the introduction of a transaction fee in PJM market on the financial traders, on November 1, 2020, led to a substantial drop of financial trading volume which in turn led to an increase of side payments. In day-ahead, the likelihood of non-zero side payments increased by 10% following the introduction of the transaction fee. In real-time, the amount of side payments increased by 80%.

FERC justified the introduction of transaction fees on virtual UTC bids on the ground that other virtual bids (INC and DEC) were already subject to similar fees, which was judged “unjust, unreasonable, and unduly preferential”. A natural alternative proposition would have been to remove transaction fees on INC and DEC, instead of introducing a fee on UTC. From our results, virtual trading reduces side payments and a transaction fee is counter-productive. As stressed by Hogan (2016), the inexistence of an equilibrium and associated side payments have their source in non-convexities, not in financial trading.

References

- Mete Şeref Ahunbay, Martin Bichler, and Johannes Knörr. Pricing optimal outcomes in coupled and non-convex markets: Theory and applications to electricity markets. *Operations Research*, 73(1):178–193, 2025.
- Kenneth Joseph Arrow and Frank Hahn. *General competitive analysis*. Holden-Day, 1971.
- Robert J Aumann. Markets with a continuum of traders. *Econometrica*, pages 39–50, 1964.
- Dimitri Bertsekas and Nils Sandell. Estimates of the duality gap for large-scale separable nonconvex optimization problems. In *1982 21st IEEE conference on decision and control*, pages 782–785. IEEE, 1982.

- Dimitri Bertsekas, Gregory Lauer, Nils Sandell, and Thomas Posbergh. Optimal short-term scheduling of large-scale power systems. *IEEE Transactions on Automatic Control*, 28(1):1–11, 1983.
- Sushil Bikhchandani and John W Mamer. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of economic theory*, 74(2):385–413, 1997.
- Conleigh Byers and Brent Eldridge. Auction designs to increase incentive compatibility and reduce self-scheduling in electricity markets. *arXiv preprint arXiv:2212.10234v3*, 2023.
- Conleigh Byers and Gabriela Hug. Long-run optimal pricing in electricity markets with non-convex costs. *European Journal of Operational Research*, 307(1):351–363, 2023.
- Hung-po Chao. Incentives for efficient pricing mechanism in markets with non-convexities. *Journal of Regulatory Economics*, 56(1):33–58, 2019.
- Peter Cramton. Electricity market design. *Oxford Review of Economic Policy*, 33(4):589–612, 2017.
- Brent Eldridge, Bernard Knueven, and Jacob Mays. Rethinking the price formation problem—part 1: Participant incentives under uncertainty. *IEEE Transactions on Energy Markets, Policy and Regulation*, 1(4):480–489, 2023.
- FERC. Price Formation in Organized Wholesale Electricity Markets, Docket No. AD14-14-000. <https://www.ferc.gov/sites/default/files/2020-05/AD14-14-operator-actions.pdf>, 2014.
- FERC. PJM Interconnection, L.L.C., 162 FERC ¶ 61,139 (Docket No. ER18-88-000), 2018.
- FERC. PJM Interconnection, L.L.C., 172 FERC ¶ 61,046 (Docket No. EL14-37-000), 2020.
- Paul R Gribik, William W Hogan, Susan L Pope, et al. Market-clearing electricity prices and energy uplift. *Cambridge, MA*, 2007.
- William W Hogan. Virtual bidding and electricity market design. *The Electricity Journal*, 29(5):33–47, 2016.
- William W Hogan. Strengths and weaknesses of the pjm market model. In Jean-Michel Glachant, Paul Joskow, and Michael Pollitt, editors, *Handbook on Electricity Markets*, pages 182–204. Edward Elgar Publishing, 2021.
- Thomas Hübner. Approximate equilibria in nonconvex markets: Theory and evidence from european electricity auctions. *arXiv preprint arXiv:2503.02464*, 2025.
- Koichiro Ito and Mar Reguant. Sequential markets, market power, and arbitrage. *American Economic Review*, 106(7):1921–1957, 2016.
- Akshaya Jha and Gordon Leslie. Start-up costs and market power: Lessons from the renewable energy transition. *American Economic Review*, 115(2):690–724, 2025.
- Akshaya Jha and Frank A Wolak. Can forward commodity markets improve spot market performance? evidence from wholesale electricity. *American Economic Journal: Economic Policy*, 15(2):292–330, 2023.
- Jalal Kazempour and Benjamin F Hobbs. Value of flexible resources, virtual bidding, and self-scheduling in two-settlement electricity markets with wind generation—part i: Principles and competitive model. *IEEE Transactions on Power Systems*, 33(1):749–759, 2017a.
- Jalal Kazempour and Benjamin F Hobbs. Value of flexible resources, virtual bidding, and self-scheduling in two-settlement electricity markets with wind generation—part ii: Iso models and application. *IEEE Transactions on Power Systems*, 33(1):760–770, 2017b.

- Bernard Knueven, James Ostrowski, and Jean-Paul Watson. On mixed-integer programming formulations for the unit commitment problem. *INFORMS Journal on Computing*, 32(4):857–876, 2020.
- Amanda Long and Anthony Giacomoni. Exploring the impacts of virtual transactions in the pjm wholesale energy market. In *2020 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2020.
- Francis A Longstaff and Ashley W Wang. Electricity forward prices: a high-frequency empirical analysis. *The journal of finance*, 59(4):1877–1900, 2004.
- Jonathan Mather, Eilyan Bitar, and Kameshwar Poola. Virtual bidding: Equilibrium, learning, and the wisdom of crowds. *IFAC-PapersOnLine*, 50(1):225–232, 2017.
- Ignacia Mercadal. Dynamic competition and arbitrage in electricity markets: The role of financial players. *American Economic Journal: Microeconomics*, 14(3):665–699, 2022.
- Paul Milgrom. *Discovering prices: auction design in markets with complex constraints*. Columbia University Press, 2017.
- Paul Milgrom and Mitchell Watt. A walrasian mechanism with markups for nonconvex markets. *The Review of Economic Studies*, 2025.
- Monitoring Analytics. State of the market report for pjm, sec. 3, 2018.
- Monitoring Analytics. State of the market report for pjm 2020, sec. 4, 2021.
- Monitoring Analytics. State of the market report for pjm 2021, sec. 3, energy market, 2022.
- Monitoring Analytics. State of the market report for pjm 2024, sec. 3, energy market, 2025a.
- Monitoring Analytics. Fast start pricing in pjm (caiso price formation enhancements: Phase 2), 2025b.
- Richard P O’Neill, Paul M Sotkiewicz, Benjamin F Hobbs, Michael H Rothkopf, and William R Stewart Jr. Efficient market-clearing prices in markets with nonconvexities. *European journal of operational research*, 164(1):269–285, 2005.
- PJM. Virtual Transactions and Uplift in PJM (FERC Technical Conference), January 2015.
- PJM. Proposed enhancements to energy price formation, 2017.
- PJM. PJM, Exploring the Impacts of Price Formation Enhancements in PJM’s Wholesale Energy Markets, *FERC Technical Conference*, June 2019.
- PJM. Pjm manual 28: Operating agreement accounting. https://www.pjm.com/pjmfiles/directory/manuals/m28/index.html#Sections/52_Credit_for_Operating_Reserve.html, 2025a.
- PJM. Pjm manual 11: Energy and ancillary services market. <https://www.pjm.com/-/media/DotCom/documents/manuals/m11.pdf>, 2025b.
- PJM. Drivers of uplift. <https://www.pjm.com/markets-and-operations/energy/drivers-of-uplift>, 2026.
- Mar Reguant. Complementary bidding mechanisms and startup costs in electricity markets. *The Review of Economic Studies*, 81(4):1708–1742, 2014.
- Ross M Starr. Quasi-equilibria in markets with non-convex preferences. *Econometrica*, pages 25–38, 1969.

- Nicolas Stevens, Anthony Papavasiliou, and Yves Smeers. On some advantages of convex hull pricing for the european electricity auction. *Energy Economics*, 134:107542, 2024.
- Nicolas Stevens, Richard O’Neill, and Anthony Papavasiliou. Average incremental cost pricing in electricity auctions. *Energy Economics*, page 109047, 2026.
- Steven Stoft. *Power system economics: designing markets for electricity*, volume 468. IEEE press Piscataway, 2002.
- Wenyuan Tang, Ram Rajagopal, Kameshwar Poolla, and Pravin Varaiya. Model and data analysis of two-settlement electricity market with virtual bidding. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 6645–6650. IEEE, 2016.
- Robert Wilson. Architecture of power markets. *Econometrica*, 70(4):1299–1340, 2002.

A Proofs of the propositions

Proof of Proposition 1. We prove the four points. Let us first define for convenience $A = c_F^2 + (2c_S\bar{D})^2 + 2c_Fc_S\bar{D} + 4c_Ss_S\bar{D}$, the expression inside the square root in (7).

1. $\partial k^*/\partial s_S < 0$ is straightforward from (7).
2. $\partial k^*/\partial c_F > 0$ can be seen from equation (6): this is a convex second-order polynomial and a higher c_F shifts the function upward ($\forall k \in [0, \bar{D}]$) thus the smaller root (k^*) moves to the right. Alternatively, using (7), one can compute:

$$\frac{\partial k^*}{\partial c_F} = \frac{1}{c_S} \left(1 - \frac{c_F + c_S\bar{D}}{\sqrt{A}} \right) > 0$$

Indeed, $\sqrt{A} = \sqrt{(c_F + c_S\bar{D})^2 + 3(c_S\bar{D})^2 + 4c_Ss_S\bar{D}} > c_F + c_S\bar{D}$

3. $\partial k^*/\partial c_S < 0$ can be seen from equation (6): clearly, $\forall k \in [0, \bar{D}]$, a higher c_S shifts the function downward thus the smallest root (k^*) moves to the left. Alternatively, using (7), one can compute:

$$\begin{aligned} \frac{\partial k^*}{\partial c_S} &= \frac{-c_F}{c_S^2} + \frac{\sqrt{A}}{c_S^2} - \frac{4c_S\bar{D}^2 + c_F\bar{D} + 2s_S\bar{D}}{c_S\sqrt{A}} < 0 \\ &\Leftrightarrow -c_F\sqrt{A} + c_F^2 + c_Fc_S\bar{D} + 2c_Ss_S\bar{D} < 0 \\ &\Leftrightarrow \sqrt{c_F^2 + (2c_S\bar{D})^2 + 2c_Fc_S\bar{D} + 4c_Ss_S\bar{D}} > c_F + c_S\bar{D} + \frac{2c_Ss_S\bar{D}}{c_F} \\ &\Leftrightarrow \frac{s_S^2}{c_F^2} + \frac{s_S}{c_F} < \frac{3}{4} \\ &\Leftrightarrow c_F > 2s_S \end{aligned}$$

which is nothing more than the condition for $k^* > 0$.

4. $\partial k^*/\partial s_F = 0$ is straightforward from (7). □

Proof of Proposition 2. The day-ahead side payments drop from $s_S d$ to zero. The real-time side payments are reduced by $\frac{s_F k}{2} \left(1 - \frac{k}{2\bar{D}} \right)$ (which is > 0 for any $k < \bar{D}$). □

Proof of Proposition 3. k^* does not depend on s_F , thus $\mathbb{E}_\epsilon(SP_{RT,\epsilon})$ is a linear function in s_F . We then compute the derivative:

$$\frac{\partial \mathbb{E}_\epsilon(SP_{RT,\epsilon})}{\partial s_F} = \frac{1}{2\bar{D}} \int_k^{\bar{D}} (\epsilon - k) d\epsilon > 0$$

□

Proof of Proposition 4. We prove that

$$s_S - \frac{k^2 c_S}{6\bar{D}} + \frac{k(c_F + s_F)}{4\bar{D}} + \frac{k c_S}{4} - \frac{(c_F + s_F)}{2} < 0 \quad (\text{A.1})$$

Injecting (6) into (A.1) to replace k^2 , and rearranging the terms, the left-hand-side leads to

$$L(k) = k \left(\frac{-c_F}{12\bar{D}} + \frac{s_F}{4\bar{D}} - \frac{5c_S}{12} \right) + \frac{s_S}{3} - \frac{s_F}{2} - \frac{c_F}{6}$$

$L(k)$ is a linear function in k , with $k^* \in [0, \bar{D}]$. A sufficient condition for $L(k^*) < 0$ is thus $L(0) < 0$ and $L(\bar{D}) < 0$. Using the model assumptions $c_S \bar{D} + s_S < c_F$ and $s_F > s_S$, we conclude:

$$L(0) = \frac{s_S}{3} - \frac{s_F}{2} - \frac{c_F}{6} < 0$$

$$L(\bar{D}) = -\frac{c_F}{4} - \frac{s_F}{4} - \frac{5c_S \bar{D}}{12} + \frac{s_S}{3} < 0$$

which proves our statement. (Alternatively, one can prove Proposition 4 by analysing the shape, the first and second derivatives and the optimum k^{**} , of the cost curve $Cost(k)$, as done in the proof of Proposition 5.) \square

Proof of Proposition 5. The first best optimal commitment in day-ahead can be computed by optimizing the cost function $Cost(k)$ in equation (11) with respect to k . Computing the first and second derivatives lead to

$$\frac{\partial Cost}{\partial k} = s_S - \frac{c_S k^2}{2\bar{D}} + k \left(\frac{c_F + s_F}{2\bar{D}} + \frac{c_S}{2} \right) - \frac{c_F + s_F}{2} \quad (\text{A.2})$$

$$\frac{\partial^2 Cost}{\partial k^2} = -\frac{c_S k}{\bar{D}} + \frac{c_F + s_F}{2\bar{D}} + \frac{c_S}{2} \quad (\text{A.3})$$

We can check that $\frac{\partial Cost}{\partial k}|_0 < 0$ (as $c_F + s_F > 2s_S$) and $\frac{\partial Cost}{\partial k}|_{\bar{D}} > 0$; as well as $\frac{\partial^2 Cost}{\partial k^2} > 0$ for $k \in [0, \bar{D}]$. Thus $Cost(k)$ is a polynomial cubic function in k , which is convex for $k \in [0, \bar{D}]$, decreasing at $k = 0$ and increasing at $k = \bar{D}$. This is illustrated on Figure 9. The optimal k^{**} satisfies $\frac{\partial Cost}{\partial k} = 0$, thus

$$c_S k^2 + k(-c_S \bar{D} - c_F - s_F) + (c_F + s_F - 2s_S)\bar{D} = 0 \quad (\text{A.4})$$

which differs from (6). Clearly, $k^{**} > k^*$ as the second-order function defined by (A.4) is above the second-order function defined by (6) (k^2 coefficient is bigger and $\bar{D}s_F \geq ks_F$ for $k \in [0, \bar{D}]$). This implies the smallest root of (6) is smaller than (A.4). \square

Proof of Proposition 6. Let us prove the inequalities stated:

- $k^t \leq k^*$ is immediate from equation (12): the second order polynomial defined in (12) is below from the one defined in (6), thus the smallest root is smaller.
- $Cost(k^*) \leq Cost(k^t) \leq Cost(0)$ is immediate from the expression of the cost function $Cost(k)$ which is decreasing at $k = 0$, convex for $k \in [0, \bar{D}]$ and minimized at k^{**} with $k^t \leq k^* \leq k^{**}$ (cf. Proposition 5).
- $SP(k^t) \leq SP(0)$ is immediate from the expression of the side payments in (9).
- $SP(k^*) \leq SP(k^t) \leq SP(0)$ can be seen by computing the derivative of the side payments: $\frac{\partial \mathbb{E}_\epsilon(SP_{RT,\epsilon})}{\partial k} = \frac{s_F}{2} \left(\frac{k}{\bar{D}} - 1 \right) < 0$ for $k \in [0, \bar{D}]$. Thus the side payments are a decreasing function of k for $k \in [0, \bar{D}]$. The conclusion follows from $0 \leq k^t \leq k^* < \bar{D}$.

\square

B Some more numerical examples

We construct a portfolio of power plants, introducing some heterogeneity between the SSR and the FSR. We build on the portfolio from the example in section 3, perturbing randomly both the costs, the minimum and the maximum production limits of both SSR and FSR. The resulting merit order curve is shown on Figure B.1a. This resembles the merit order curve in Figure 10a, but with more heterogeneity between the power plants. One can see

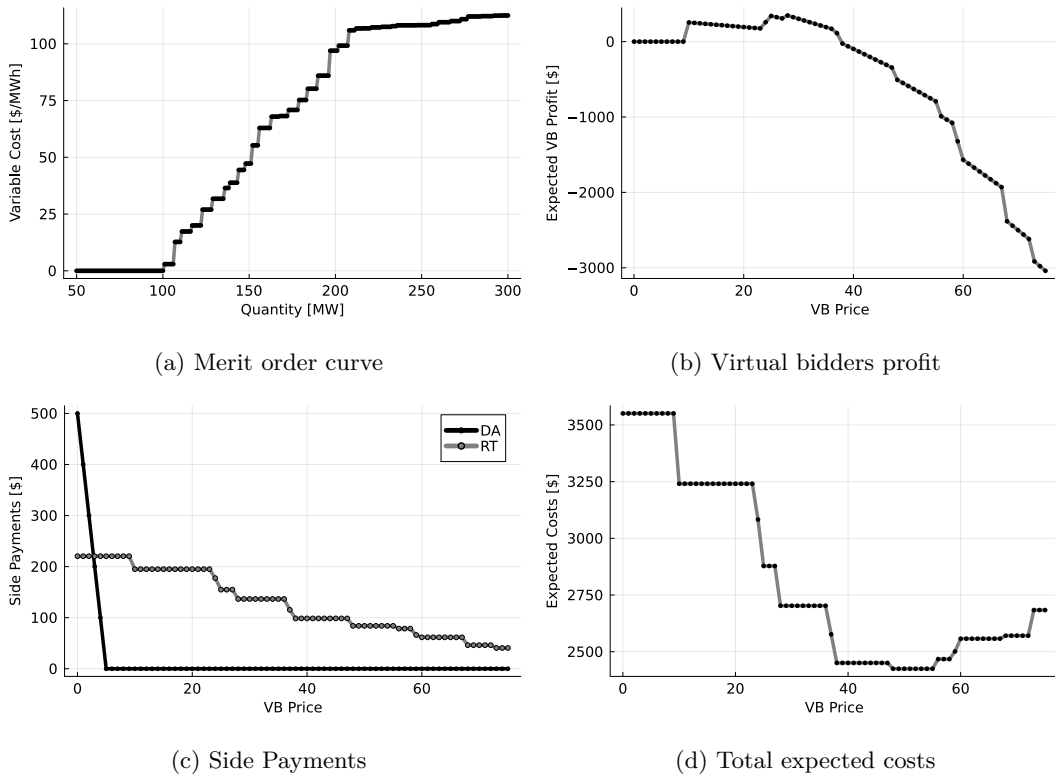


Figure B.1: Results of the perturbed discrete example.

from Figure B.1 that all the previous observations from section 3 (both the model and the simulations on the numerical example) still apply to this example.

Notice that, because of the non-convexity in the game, a pure strategy equilibrium is not guaranteed to exist. Instead, the game may admit a mixed-strategy equilibrium, in which virtual traders randomize, for example, between the two strategies that yield zero expected profit.

C PJM data

This appendix provides more in-depth information about PJM data. We comment about each data one by one.

MISSING DATA. As mentioned in section 4, our database includes 1276 dates, excluding 12 dates with missing data. The dates with missing data are: 2018-12-25, 2018-12-29, 2018-12-31, 2019-10-09, 2019-12-31, 2020-07-25, 2020-12-29, 2020-12-31, 2021-01-07, 2021-01-24, 2021-07-25, 2021-05-27. The missing data are coming from side payments data, self-schedule data and DA load data. Side payments data are missing for 2018-12-25 and 2018-12-29. Self-schedule data are missing for 2018-12-31, 2019-10-09, 2019-12-31, 2020-07-25, 2020-12-31, 2021-01-07, 2021-05-27, 2021-11-01 and 2021-12-13. DA load data are missing for 2020-12-29, 2021-01-24 and 2021-07-25.

SIDE PAYMENTS DATA. Figure C.2 shows the distribution of the side payments (related to non-convexities). Figure C.3 shows all the types of uplift payments that exist in PJM, divided per category. Other uplifts payments can be divided into energy market uplifts, ancillary services uplifts and reliability services uplifts. Energy market uplift—besides Balancing Operating Reserve—include “Balancing Operating Reserve Lost Opportunity Cost” which are uplift due to out-of-merit re-dispatch for transmission constraint control. Ancillary services uplifts include “Regulation and Frequency Response Service”, “Synchronized

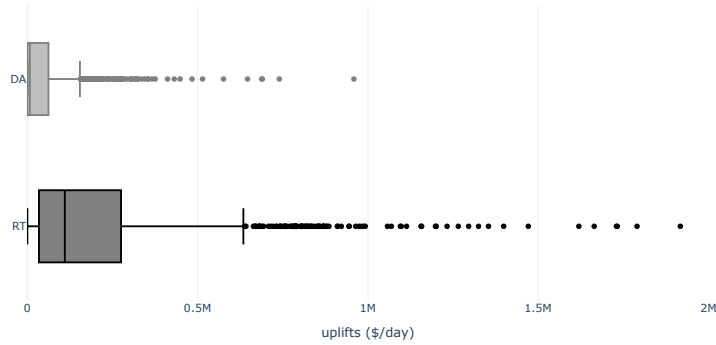


Figure C.2: PJM daily side payments distribution.




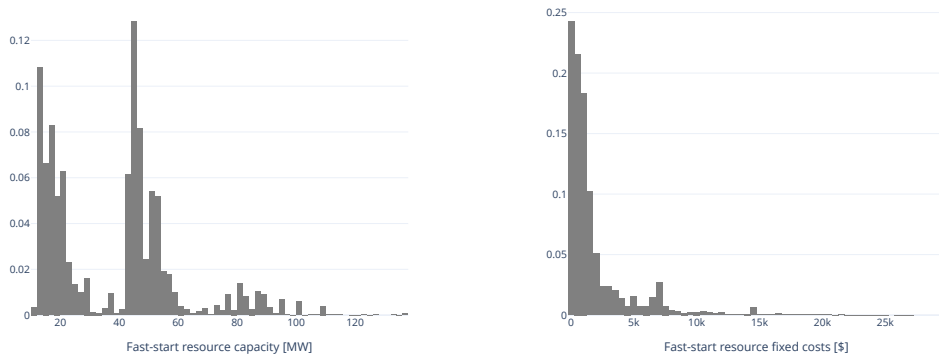
Figure C.3: Overview of all categories of PJM “uplifts”.

Reserve” and “Non-Synchronized Reserve”. Reliability services uplifts include “Black Start Service” and “Reactive Supply and Voltage Control Service”.

GENERATION BIDS. Figure C.4 provides the distribution of the capacity and fixed costs of fast-start resources. As noted in section 4, fast start resources are mostly small power plants, with a capacity between 0 and 100MW and a median at 43MW. Their start-up costs are mostly between 0 and 10,000\$.

LOAD. We aggregate PJM hourly load into daily load values by summation. We also compute the standard deviation of the hourly load within each day as a measure of load variations: a high standard deviation indicates a sharp difference between peak and offpeak loads within a day, often associated with steep load increase. Both measures are computed for day-ahead and real-time loads. Figure 14 shows the resulting time series for the real-time loads.

GRID CONGESTIONS. Direct data on grid congestions, such as the shadow prices associated to each line or the congestion rent, are not available on PJM. Yet, the severity of congestions can be inferred from prices data. We test two metrics of congestion, that are included in the controls. One is based on a load-weighted average of the congestion component of PJM total price. The other is based on the standard deviation of the locational prices at each hour. Days with high standard deviation of locational prices are indeed likely to be days where grid constraints are binding.



(a) FSR capacity distribution

(b) FSR fixed costs distribution

Figure C.4: PJM fast-start resource capacity and fixed costs distributions.

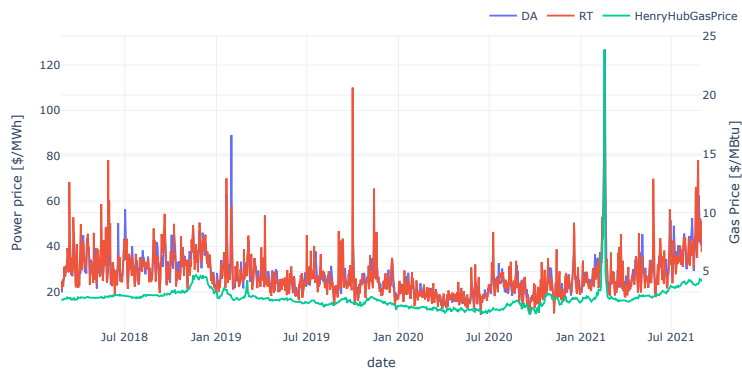


Figure C.5: Gas price and PJM prices.

MARKET PRICES. PJM provides times series of hourly locational price of electricity in day-ahead and real-time. We compute aggregate daily prices (cf. Figure C.5). We use the system price (price at the reference hub) as a control in our regression.

GENERATION AVAILABILITY. PJM bids data also enable to compute the generation capacity available each day (i.e. some plants might be unavailable due to planned outages or for other reasons). The resulting time series of both number of plants available and MW capacity available is provided in Figure C.6. There is a clear seasonality in plants availability, with most non-availabilities occurring during the fall and the spring, thus outside the winter and summer peak loads.

SELF-SCHEDULING. Self-scheduling are the physical generation bids at zero price (as opposed to the dispatchable generation bids). PJM provides data of hourly real-time self-scheduling. We compute the volume of day-ahead self-scheduling using the bid data, as the supply offers at zero price (cf. Figure 2a). Figure C.7 shows the resulting output, which fits well with the average figures of day-ahead self-scheduling reported by PJM market monitoring.

EMERGENCY PROCEDURE EVENTS. Emergency procedure are issued by PJM to ensure reliability of the system. These events relate to controlled load curtailment and load shedding, generation curtailment, geomagnetic disturbance, voltage support, hot and cold




Figure C.6: PJM daily available generation capacity and number of active units.




Figure C.7: PJM bid volume by category.




Figure C.8: Number of emergency events per month in PJM.




Figure C.9: Distribution of the emergency events per category (only the most significant categories are shown).

weather alerts, etc. We compute a time series of number of daily events. We count only the events labelled as Action, Warning and Alert. If an event spans over multiple days, it is counted as an event on each of these days. Figures C.8 shows the resulting number of emergency events per month in PJM. Figure C.9 shows the distribution of these events per category.