

Who benefits from ratepayer-funded auctions of transmission congestion contracts? Evidence from New York

Gordon Leslie*

December 19, 2018

Abstract

Transmission congestion contracts are derivative products that electricity retailers can use to change their future wholesale electricity price exposure to a different location. U.S. Congress is concerned by financial trader profits in auctions for these derivatives because the payouts are funded by ratepayers, not willing counterparties. I study firm-level positions in the New York Wholesale Electricity Market to investigate the causes of this concern. I find a small set of financial traders earn large, systematic profits on products that electricity retailers tend to avoid. However, trader participation can improve price signals on these and related products. Policy implications are discussed.

*Department of Economics, Monash University. 900 Dandenong Rd, Caulfield East, VIC 3145, Australia. *gordon.leslie@monash.edu*. I thank my PhD advisors: Frank Wolak, Peter Reiss, Charles Kolstad and Tim Bresnahan for their insightful and helpful comments. The article also benefited from comments by Chris Bruegge, Isa Chaves, Liran Einav, Charles Hodgson, Bowen Hua, Will Parker, Nicola Pierri, Mark Thurber, Caio Waisman and participants in seminars at Stanford University, The University of North Carolina, The University of Maryland, Virginia Tech, University of Technology Sydney, Monash University, the UC Berkeley Energy Institute Camp, the USAEE/IAEE North American Conference and the Asia-Pacific Industrial Organisation Conference. I gratefully acknowledge financial support I received during this research project from the Alfred P. Sloan Foundation Pre-doctoral Fellowship on Energy Economics (awarded by the NBER), the Gale and Steve Kohlhagen Fellowship in Economics (awarded by the Stanford Institute for Economic Policy Research) and Stanford's Program on Energy and Sustainable Development. All errors are my own.

“Across the nation, investment funds and major banks are wagering billions on [transmission congestion contracts], as they chase profits in an arcane arena that rarely attracts attention... The utilities and power companies suggest they cannot win against trading outfits that employ math specialists, often called ‘quants,’ to spot lucrative opportunities. With transmission contracts, there are tens of thousands of tradable combinations.”

The New York Times, Creswell and Gebeloff (August 14, 2014)

Regulators have long wondered whether financial trader profits in commodity markets are purely transfers from producers and consumers of the underlying product.¹ Ideally, financial traders can improve market liquidity, future price signals, and ultimately, overall market efficiency. The restructuring of wholesale electricity markets has introduced opportunities for purely financial participants, accompanied by concerns that their profits represent costly wealth transfers away from the physical participants that buy and sell wholesale electricity. In this article, I study the sources of persistent profits that traders have earned in New York’s transmission congestion contract market and discuss how ratepayers might benefit from the actions of rent-seeking traders.

Transmission congestion contracts, or TCCs, are derivative contracts that pay or collect the difference between wholesale electricity prices at two locations for a specified future time period. Physical firms (electricity retailers and generators in this case) can benefit from the availability of such contracts. An electricity retailer, which must buy electricity at a fixed location to serve its customers, can buy a TCC to change its future spot price exposure to that of a different location. For example, in the New York Independent System Operator (NYISO) wholesale market, there are 450 locations where electricity can be purchased. This means that 449 TCCs are available that pay a price difference between a given location and that of the retailer. The retailer can search among the 449 other locations for where it believes it can source its electricity most cheaply and buy the corresponding TCC to effectively pay the electricity price at that location. Such behavior can potentially lower the wholesale energy costs of a retailer. Like retailers, electricity generating firms can derive benefits from TCCs by using them to effectively sell their output at a price of a different location to their own. Finally, financial traders participate in the markets, with the motive

¹See Chapter 10, Baer and Woodruff (1929) for an early list of concerns regarding trader behavior in commodity exchanges.

to acquire derivatives at prices less than their eventual payout. Competition among traders can cause price signals for derivatives to converge on the expected payouts of the products, and these improved price signals can aid physical firms when planning their long-term energy procurement process.²

TCCs (or financial transmission rights)³ are auctioned in all formal electricity markets in the United States. Electricity customers, not willing counterparties, effectively fund the payouts of the issued derivatives. In New York, periodic, multi-product auctions offer every bilateral combination of the 450 locations – over 100,000 products.⁴ As the opening New York Times quote highlights, financial traders have consistently earned large trading profits in these notoriously complex auctions, totaling \$600m annually across four major US markets.⁵ Market monitors are concerned by these large trading profits earned by participants in TCC auctions because TCC profits result in transfers from ratepayers (CAISO Department of Market Monitoring, 2016). In November 2017, the U.S House of Representatives Subcommittee on Energy convened with the aim to, “take a hard look at whether [TCC] trading makes sense and answer this question: Does financial trading make the electricity markets more efficient, and in turn, result in benefits to consumers?”⁶

We can learn where the required efficiency gains must occur by studying the auction positions taken by the physical and financial participants in TCC auctions. The objectives of this article are to examine the sources of trading profits in TCC auctions, the persistence of the trading profits, and to understand whether financial trader participation is likely to improve market performance in this setting. Understanding the sources of trading profits will identify why the auctions are resulting in large transfers from ratepayers to TCC holders. Further, if potential barriers to eroding these profits can be identified, their removal would ease concerns related to these wealth transfers.

²Introducing financial trader participation to day-ahead electricity markets has been shown to improve day-ahead price convergence to realized real-time prices (Saravia, 2003; Jha and Wolak, 2015). See Jha and Wolak for a demonstration of how financial traders have improved the production efficiency of the physical underlying market.

³Other markets in refer to these instruments as financial transmission rights, FTRs; or congestion revenue rights, CRRs.

⁴450 locations allows for $450 \times 449 = 202,050$ directional location pairs or 101,025 unique location pairs.

⁵Sum of the yearly averages of the following: New York: Paid out \$3,760m (to all firms) and received \$2,905m from 1999-2015 (author calculation). California: Payments of \$970m to non-physical participants (banks and energy traders) and auction payments of \$450m from 2012-2015 (CAISO Department of Market Monitoring, 2016). Mid-continent (MISO): Paid out \$3,453m (to all firms) and received \$3,037m from 2013-2015 (MISO, 2015, and various issues). Pennsylvania and surrounds (PJM): Profits to non-physical participants (banks and energy traders) of \$904m from 2013-2015 (PJM, 2015, and various issues).

⁶Passage from the Opening Statement of the Honorable Fred Upton, United States House of Representatives Subcommittee on Energy (November 29, 2017).

To accomplish these objectives, I first present a stylized model of an electricity network and compute electricity market prices and TCC auction outcomes. In standard exchange settings, financial traders can improve the available quantity of a derivative product by offering counterpositions to bids and offers by physical firms. Under the TCC auction mechanism, equilibrium prices and quantities for each product are interdependent and determined simultaneously. I show that when traders buy products that are not typically purchased by physical firms they can improve the available quantity and price signals on *other* products and facilitate matching buyers and sellers for contracts in different locations throughout the network.

The results from the theoretical examples are then used to guide the empirical portion of the article where I compile microdata on 16 years of derivative prices, payouts and firm-level trading positions in the New York TCC market to examine the different types of products firms purchase and the persistence of trading profits. To organize the analysis of the rich variety of products, I classify each derivative into groups based on the two locations specified in the price difference payout, and the time horizon of the payouts. In the time dimension, electricity markets are hourly and TCCs are available covering payments for every hour over 1-, 6- or 12- months. In the location dimension, products are either: 1) *nodal* products that pay the difference between two locations where a power plant is located or 2) *zone-indexed* products that pay holders the difference between regional price indexes, which are the prices that retailers face in the spot market and are ultimately passed through to ratepayers.

I find that retailers, generators and traders purchase zone-indexed derivatives, but only generators and traders purchase nodal derivatives. Retailers usually bid on less than 1% of the products that generators and traders bid on in each auction but account for 16% of total derivative purchases. Retailers purchase their products in large quantities, for long terms, and at actuarially fair prices that on average equal derivative payouts. Generator owners, who account for 33% of derivative expenditures, earn trading profits on nodal products, but not zone-indexed products. A large portion of their derivative purchases do not appear related to their physical operations. Financial traders account for the remaining 51% of derivative expenditures, purchase a wide variety of products, and receive most of the trading profits in this market. Like generators, traders only earn systematic profits on nodal, but not zone-indexed products.

To investigate why competition between financial traders is not sufficient to erode trading profits,

I study whether trading profits persist on the same products over time. Specifically, I measure derivative price responses across the auctions that take place at regular intervals.

The main empirical finding is that 88% of the financial trader profits are earned from being the first firm to purchase previously illiquid products. Following the public revelation of a purchase of a derivative by a profitable firm, the price for that same product appreciates by approximately 10% in the subsequent auction and the profitable opportunity is eroded. This quick adjustment of prices on the same products across auctions suggests that payout premiums are not solely due to the presence of a risk premium, an opportunity cost of capital or some other fixed cost to participation. Based on these findings, I argue that profitable traders can improve price signals and liquidity, but also that they are unable to persistently profit on the same derivative products. Traders must consistently identify profitable opportunities from illiquid derivative products if they are to consistently earn profits.⁷

A major barrier to eroding overall trading profits could be the cost for new entrants to develop a technology that can identify successful trading strategies in TCC auctions. These multi-product auctions are complex, where TCC payouts and the auction allocations are determined in part by physical transmission constraints in the electricity network. Anecdotes describe successful firms consistently updating their models and aggressively enforcing non-disclosure agreements with ex-employees. Alternate explanations do not appear to explain the majority of trading profits, such as profits being derived from exploiting market power in the energy market (theorized in Bushnell, 1999; Joskow and Tirole, 2000), or from manipulative actions by traders (demonstrated in a case study for the MISO electricity market in Birge, Hortaçsu, Mercadal, and Pavlin, 2018).

For TCC auctions to benefit ratepayers, physical efficiency gains derived from the auction process must exceed the systematic trading profits that are earned by TCC holders. Trader profits are largely found on purchases of previously illiquid, nodal products that retailers tend to avoid. I use these findings to highlight some of the tradeoffs that accompany proposals to eliminate or reduce the set of derivative products offered at ratepayer-funded auctions. Greater product offerings allow greater flexibility in the procurement strategies of physical firms and greater information revelation, but given that most of these products are not utilized by electricity retailers, a broader product set

⁷These findings extend prior observations that TCCs are not priced actuarially fairly (Bartholomew, Siddiqui, Marnay, and Oren, 2003; Hadsell and Shawky, 2009; Adamson, Noe, and Parker, 2010; Olmstead, 2018).

also offers more opportunities for trader profits and associated transfers from ratepayers.

The article is organized as follows: The product, the auction mechanism and the role of financial traders are described in section 1, followed by a description of the New York setting in section 2. Data sources are described in section 3, followed by the empirical analysis. Section 4 describes the positions taken by firms and the trading profits they earn across different product types. Section 5 investigates why trading profits have not eroded over time by describing the persistence of trading profits. Section 6 then discusses the policy relevance of the findings.

1 TCC prices, TCC payouts and the role of financial traders

A theoretical platform for subsequent empirical analysis is presented in this section, describing how electricity prices and transmission congestion contract (TCC) prices are derived in a network model. Understanding the relationship between wholesale electricity prices and transmission constraints is required to both understand the TCC auction mechanism, and how financial traders can earn profits and improve market performance. The primary result is that financial traders that purchase the TCCs that physical firms do not take up can improve both the quantity available of *other* TCC products, and price signals from the auction.

A transmission congestion contract between location i and location j in hour h pays the holder:

$$LMP_{j,h} - LMP_{i,h}$$

where $LMP_{i,h}$ is the electricity price at location i in hour h . This payout is a price swap, where if the value is negative the holder must pay money. Earning realized trading profits in TCC markets requires a firm to buy (sell) the derivative for less (more) than its eventual payout.

Although in practice TCCs cover 1-, 6- or 12-months of hourly payouts and can be purchased between any of the 450 locations in the New York market, this section will consider a one-period setting with three locations to introduce the fundamental concepts behind TCC markets.

1.1 Determining wholesale electricity prices

All formal wholesale electricity markets in the US set *locational marginal prices* (LMPs) at different locations in the network each hour of the day. The prices that determine TCC payouts are LMPs in the day ahead electricity market. The LMPs in the day-ahead market can be considered the spot market in this study.⁸ Electricity market operators collect offers to supply electricity from generator owners. They then set LMPs at every location (node) in the electricity grid to minimize the system-wide as-offered cost of supplying electricity, subject to network constraints and supply meeting demand. This can mean that a cheap offer of electricity at a generating location might not be taken up if extra supply at that location will violate a line capacity constraint somewhere in the network. In such cases, prices between regions affected by this congestion will diverge and a higher cost source will be called upon in the congested regions.

To demonstrate how congestion influences prices in electricity markets, consider the network configuration, supply offers and demand in the market specified in figure 1. This example builds on Oren (2013) and will be used throughout the section. There are three locations in the example electricity market, connected by a transmission loop. All locations have generators, but only location k has consumers. The transmission line between i and j is able to accommodate flow up to a maximum capacity of 100MW, and the line between i and k has a capacity of 400MW. For strictly illustrative reasons, the remaining j to k line is unconstrained and there are no line losses from transmission.⁹ 1500MW of electricity is demanded inelastically at k , with the following offers to supply electricity:

- Generators at i : 2000MW at \$80/MWh
- Generators at j : 2000MW at \$100/MWh
- Generators at k : 2000MW at \$200/MWh

Solving for the optimal market supply is trivial in the absence of transmission constraints – generators at i produce all of the electricity because it is the cheapest source. However, the transmission

⁸Electricity markets have a day-ahead market and a real-time market. Day-ahead markets are run one day in advance to a given delivery hour. When production or demand varies from the day-ahead production allocations during the delivery hour, the real-time market determines which power plants will increase or decrease their production to balance supply and demand in the system.

⁹Resistance on each line is assumed equal and there are no transmission losses built into the solutions.

Figure 1: A three-node electricity network and example equilibrium

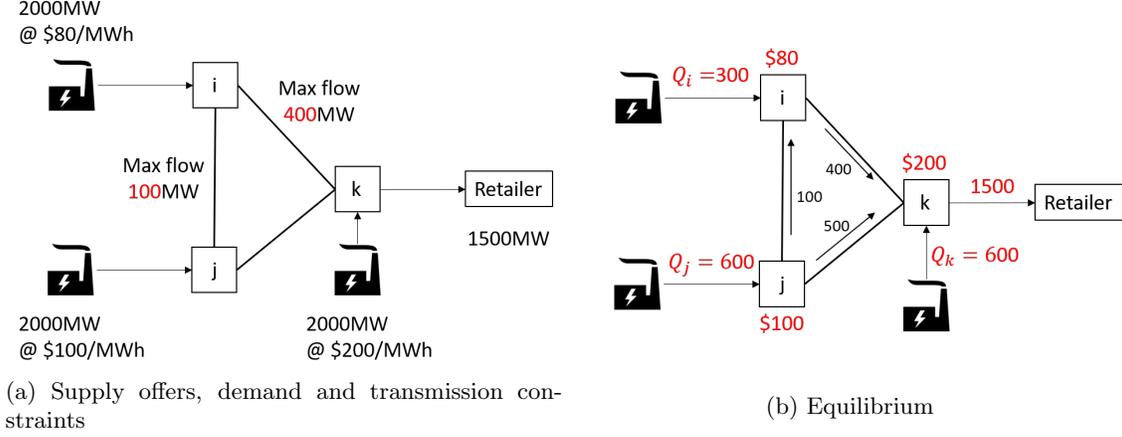


Figure (b) displays the solution to the program described in equation (1). To calculate flows on each line (the numbers inside the transmission lines), Kirchhoff's circuit laws are applied to this stylized network with no transmission losses. The formula is described in text, with the implication being that $\frac{1}{3}$ of supply at j flows via i to k , with the remaining $\frac{2}{3}$ flowing from j to k . $\frac{1}{3}$ of supply at i flows via j to k , with the remaining $\frac{2}{3}$ flowing from i to k . The body of section 1.1 describes how equilibrium prices (LMPs) are determined.

limits and the loop flow that occurs in electric circuits constrain the cost minimizing solution. The market operator solves the optimization problem described in (1) to minimize system-wide as-offered costs, with a description of the constraints to follow.

$$\text{Objective: } \min_Q 80 \cdot Q_i + 100 \cdot Q_j + 200 \cdot Q_k \quad (1)$$

$$\text{Supply = Demand: } Q_i + Q_j + Q_k = 1500$$

$$\text{Transmission constraint } i \text{ to } k: \frac{2}{3}Q_i + \frac{1}{3}Q_j \leq 400$$

$$\text{Transmission constraint } i \text{ to } j: -100 \leq \frac{1}{3}(Q_i - Q_j) \leq 100$$

$$\text{Generator constraints: } Q_i \leq 2000, Q_j \leq 2000, Q_k \leq 2000$$

$$\text{Solution: } Q_i = 300, Q_j = 600, Q_k = 600$$

The first constraint is that supply equals demand in the network. The second constraint is the transmission constraint on the i, k line, where flow can not exceed 400MW, with the $\frac{2}{3}$ multiplier on Q_i and the $\frac{1}{3}$ multiplier on Q_j due to assumed equal line resistance and Kirchhoff's circuit

laws.¹⁰ The third constraint is the transmission constraint on the i, j line, where flow can not exceed 100MW in either direction. Q_i and Q_j variables offset each other as counterflows in this constraint due to Kirchhoff's circuit laws, demonstrating that it is possible for more electricity being injected at i if more electricity is injected at j .¹¹ The final constraints are the capacities offered by the generators at each location node. The solution is displayed in figure 1. Both transmission line constraints are binding, limiting the generation that can occur at i and j .

Locational marginal prices are equal to the increase in the optimized value of the objective function in (1) from withdrawing an extra unit of electricity from the node. The prices for this example are $LMP_i = \$80/\text{MWh}$, $LMP_j = \$100/\text{MWh}$ and $LMP_k = \$200/\text{MWh}$.¹²

This three-node example highlights the interdependency of the network problem. Despite the line between j and k not having a maximum flow rating, the constraints on the other lines lead to the LMP_j and the LMP_k prices separating, in this case by $\$100/\text{MWh}$.¹³

1.2 Relating network congestion to the policy problem, transmission congestion contracts and ratepayer cash flows

Participants in electricity markets face the LMPs at the location where they generate (produce) or withdraw (consume) electricity. Even though generators at i in figure 1 receive $\$80/\text{MWh}$, the retailer at k pays $\$200/\text{MWh}$. Therefore, the cash flows from the market in figure 1 are the following:

¹⁰Given equal resistance on each line and the implication from Kirchhoff's laws that the sum of flows entering and exiting any given node must equal zero, injection of electricity at j and withdrawal at k will have $\frac{1}{3}$ flow via i and the remaining $\frac{2}{3}$ flow directly to k . This is because the i route encounters twice the number of lines, therefore twice the resistance, so the $\frac{1}{3}$ and $\frac{2}{3}$ split equates marginal losses meaning that electricity flows take the path of least resistance.

¹¹This is because electricity injected at i and j and withdrawn at k each have $\frac{1}{3}$ of the electricity flow via the i, j line.

¹²See Bohn, Caramanis, and Schweppe (1984) for a detailed explanation of locational marginal pricing and how the prices reflect Lagrange multipliers on the transmission constraints and shift factors. At node i , only 300MW of the 2000MW offered at $\$80$ is generated in equilibrium, therefore the marginal cost of withdrawing a unit of electricity at i is $LMP_i = 80$. However, due to the transmission constraints being binding, it is infeasible to inject an extra MW of electricity at i to be withdrawn at either j or k . Only 600MW of electricity offered at node j is utilized in the solution, therefore the marginal cost of withdrawing a unit of electricity at i is $LMP_i = 100$. Again, it is infeasible to inject an extra unit of electricity at node j to be withdrawn at node k , therefore the marginal cost of withdrawing a unit of electricity at k is $LMP_k = 200$.

¹³Further, electricity does not necessarily flow from low-cost to high-cost nodes. Although cheap electricity flows to k , the net flow on the i, j transmission line is in the j to i direction, from a higher to a lower cost location.

Entity	Cash flow	Realized cash flow (figure 1)
Retailer at k	$-LMP_k \cdot (Q_i + Q_j + Q_k)$	$-200 \cdot 1500$
Generators at i	$LMP_i \cdot Q_i$	$80 \cdot 300$
Generators at j	$LMP_j \cdot Q_j$	$100 \cdot 600$
Generators at k	$LMP_k \cdot Q_k$	$200 \cdot 600$
Market operator	$(LMP_k - LMP_i) \cdot Q_i +$ $(LMP_k - LMP_j) \cdot Q_j$	$120 \cdot 300 +$ $100 \cdot 600 = \$96,000$

The revenue in the final line is the *merchandising surplus*. Market operators collect a merchandising surplus because congestion in the system results in the retailer paying more for their energy consumption than the generators are paid for their energy production. In this example, the merchandising surplus is equal to the payouts of 300 TCCs between i and k and 600 TCCs between j and k . A policy decision must be made for how to distribute this revenue. In formal electricity markets throughout the United States, market operators securitize the merchandising surplus into TCCs in advance of the short-term energy market (described in the next section).

For the retailer at k to source 1MWh at the location i price, they would need to purchase a TCC between locations i and k that pays $LMP_k - LMP_i$. Combining the TCC payout and their spot price LMP_k means that they effectively pay LMP_i , the spot price at i :

$$\text{Retailer spot payment: } -LMP_k$$

$$\text{TCC payout: } LMP_k - LMP_i$$

$$\text{Net cashflow from spot + TCC: } -LMP_i$$

Further, if the retailer at k were to have a forward contract with a generator at i , then the combination of the forward contract with a TCC removes all price uncertainty.¹⁴ Therefore, retailers with a full set of forward price offers from suppliers at each location and the full set of TCC prices can more efficiently source electricity by picking the supplier that offers the lowest price when combined

¹⁴Consider a retailer at node A entering a forward contract to source x MWh of power from node B . In the spot market, the firm purchases x MWh at A to meet its consumption needs but owns x MWh at B from its forward position, therefore its cash flows are now exposed to a basis of $\$(LMP_B - LMP_A) \cdot x$. Notice that an x unit transmission congestion contract position exactly matches this basis differential, therefore an x unit forward contract at B combined with an x unit TCC between A and B removes all price uncertainty for the firm sourcing x MWh of electricity from node B .

with the corresponding TCC. This is one potential mechanism for TCCs to improve competition between suppliers, economic efficiency or lower the costs of procurement for retailers.¹⁵

To disentangle who loses when an entity profits from their TCC holdings, consider the sequence of events and cash flows in TCC auctions and wholesale electricity markets:

1. TCC auction

- Contracts issued, auction revenues collected by market operator

2. Day-ahead electricity market

- LMPs determined
- Market operator collects merchandising surplus from transmission congestion

3. Cash flows

- TCC holders receive payout based on realized LMPs
- Merchandising surplus + (auction revenues - TCC holder payouts) distributed to lower the transmission service charge paid by transmission ratepayers
 - All else equal, ratepayers benefit from higher auction revenues. The zero sum nature of TCC holder profits (the bracketed term) means that TCC holder profits are effectively funded by ratepayers, and TCC holder losses benefit ratepayers

TCC holder profits are transfers from ratepayers because of the design of the transmission service charge. Transmission forms a natural monopoly, with transmission owners regulated to earn a fixed rate of return in exchange for open access to their transmission lines. Wholesale market consumers collectively pay this fixed rate of return less the merchandising surplus and the difference between auction revenues and TCC holder payouts, with individual payments determined by a cost-splitting rule. This fee is a transmission service charge (TSC), where lower TCC holder profits means a bigger reduction in this charge and ultimately, lower bills to customers.¹⁶

¹⁵Alsac, Bright, Brignone, Prais, Silva, Stott, and Vempati (2004) argue that TCCs provide hedging benefits to firms. Formalizing hedging benefits is not the focus of this article. See Jha (2017) for a recent empirical investigation into the risk aversion of electricity retailers.

¹⁶Transmission owners must meet further operational and reliability targets to earn its return. See section 14.1.2 of NYISO (2010) for a detailed breakdown of the transmission service charge.

1.3 The transmission congestion contract auction and a role for financial traders

The merchandising surplus collected by market operators is stochastic and depends on equilibrium flows and prices in the network (see figure B1). TCC auctions have been designed to allocate a set of TCCs, where the collective payout to TCC holders does not exceed the merchandising surplus. Hogan (1992) proves that a given allocation of TCCs can be funded from the merchandising surplus if the set of contracts are *simultaneously feasible*. Simultaneous feasibility means that if each i, j TCC of size q resulted in q MW being injected at i and withdrawn at j in the physical electricity network, no transmission constraint in the network would be violated.¹⁷ Consequently, the volume of the TCCs that can be issued between any two locations is dependent on all other TCCs issued in the network and the transmission capacities in the electricity network. This section outlines the simultaneous feasibility constraint, the auction equilibrium and, through a series of examples, will highlight a potential role for financial traders.

The market operator collects offers to buy and sell each possible combination of TCC. A bid to buy the i, j TCC means the holder wishes to receive the future cash flow $LMP_j - LMP_i$ from the electricity market. An offer to sell the i, j derivative is equivalent to a bid to buy the j, i derivative, with the holder of such a contract receiving $LMP_i - LMP_j$. I consider only three products existing, the i, j , the j, k and the i, k , where selling a product is equivalent to buying a negative quantity. The network configuration and constraints match the running example in figure 1. For this 3-node network, the auction problem solves the following program for the vector q containing the quantity of each TCC bid that is issued:

$$\text{Objective: } \max_q b \cdot q \quad (2)$$

$$\text{Simultaneous feasibility } i, k \text{ line: } \frac{2}{3}q_{i,k} + \frac{1}{3}q_{j,k} + \frac{1}{3}q_{i,j} \leq 400 \quad (3)$$

$$\text{Simultaneous feasibility } i, j \text{ line: } -100 \leq \frac{1}{3}q_{i,k} - \frac{1}{3}q_{j,k} + \frac{2}{3}q_{i,j} \leq 100 \quad (4)$$

$$\text{Bid quantity constraints: } \bar{q} \cdot 1(\bar{q} \leq 0) \leq q \leq \bar{q} \cdot 1(\bar{q} \geq 0) \quad (5)$$

¹⁷For example, a 10 unit contract from i to j implies a 10MW injection of electricity at i and a 10MW withdrawal of electricity at j . If the implied injections and withdrawals of all contracts is not feasible given the assumed transmission capacities of the grid, then payouts to the set of TCC holders may exceed the merchandising surplus, a funding shortfall. See Hogan (1992) or appendix A for more technical details.

where b is the bid price vector for each bid in the \bar{q} vector and $q_{a,b}$ is the sum of all allocated TCCs issued between a and b .¹⁸ The auction equilibrium maximizes the as-bid valuations for the TCC allocations, subject to the simultaneous feasibility constraint. Notice the tradeoffs between the quantities of i, k , j, k or i, j contracts that can be issued. The simultaneous feasibility constraint in (3) has each additional unit of a contract type reducing the amount of other types that can be issued. However, the simultaneous feasibility constraint in (4) dictates that if more i, k or i, j TCCs are issued, it allows extra j, k TCCs to be issued.¹⁹ Therefore, depending on which constraints are binding, bidding on a particular product can increase or decrease the quantity available of another product. Derivative prices, $p_{i,k}, p_{j,k}, p_{i,j}$ are set such that all bids above (and offers below) the price are cleared and that they are transitive. Transitivity is imposed so that there is no within auction arbitrage opportunities, so $p_{i,k} = p_{i,j} + p_{j,k}$, given that the payouts for the i, k derivative is equal to the sum of the payouts of the i, j and j, k derivatives.

To demonstrate how financial traders may profit and influence auction performance, equilibrium outcomes will be described for five examples of bids, displayed in table 1. These cases are:

1. Ideal allocation: TCC prices and quantities match realized flows in the electricity market. Merchandising surplus is fully securitized into TCCs.
2. Under allocation: Low demand for one TCC reduces the available quantity of another TCC. Merchandising surplus is partially securitized into TCCs.
3. Trader liquidity and signaling 1: Traders buying a TCC with low demand can earn a profit and improve the available quantity of other TCCs.
4. Trader liquidity and signaling 2: Traders buying a TCC that is never used in the procurement strategies of physical firms can earn a profit and improve the available quantity of other TCCs.
5. Trader competition: Competition among traders on a TCC not used by physical firms can restore price efficiency on all contracts in the network.

¹⁸Formally, NYISO lists quantity of contracts in megawatts (MW). To avoid a confusion regarding the stock or flow nature of quantity, this article will not refer to the quantity in MW units, because one TCC pays the per MWh price difference between two locations over the duration of the contract.

¹⁹The constraint includes $\frac{2}{3}(q_{i,j})$ because a 1MW injection at i and a 1MW withdrawal at j means adding $\frac{1}{3}$ MW flow on the i, j line and removing $\frac{1}{3}$ counterflow from the i, j line. See Deng, Oren, and Meliopoulos (2004) for the generalized auction problem.

Table 1: Example TCC auction bids, allocations, prices and cash flows

	Example 1	Example 2	Example 3	Example 4	Example 5
	Bids				
i, k TCC:	2000 @ \$120	30 @ \$120	30 @ \$120 2000 @ \$110	30 @ \$120	30 @ \$120
j, k TCC:	600 @ \$100	600 @ \$100	600 @ \$100	600 @ \$100	600 @ \$100
i, j TCC:	No bids	No bids	No bids	2000 @ \$10	2000 @ \$10 2000 @ \$20
	Equilibrium				
i, k TCC:	$q_{i,k} = 300$ $p_{i,k} = \$120$	$q_{i,k} = 30$ $p_{i,k} = \$120$	$q_{i,k} = 300$ $p_{i,k} = \$110$	$q_{i,k} = 30$ $p_{i,k} = \$110$	$q_{i,k} = 30$ $p_{i,k} = \$120$
j, k TCC:	$q_{j,k} = 600$ $p_{j,k} = \$100$	$q_{j,k} = 330$ $p_{j,k} = \$100$	$q_{j,k} = 600$ $p_{j,k} = \$100$	$q_{j,k} = 600$ $p_{j,k} = \$100$	$q_{j,k} = 600$ $p_{j,k} = \$100$
i, j TCC:	$q_{i,j} = 0$ $p_{i,j} = \$20$	$q_{i,j} = 0$ $p_{i,j} = \$20$	$q_{i,j} = 0$ $p_{i,j} = \$10$	$q_{i,j} = 435$ $p_{i,j} = \$10$	$q_{i,j} = 435$ $p_{i,j} = \$20$
	Cash flows				
Auction revenues ^a	\$96,000	\$36,600	\$93,000	\$67,650	\$72,300
Merchandising surplus ^b	\$96,000	\$96,000	\$96,000	\$96,000	\$96,000
TCC holder payouts ^c	\$96,000	\$36,600	\$96,000	\$72,300	\$72,300
TSC rebate ^d	\$96,000	\$96,000	\$93,000	\$91,350	\$96,000
$a + b - c$					
	Simultaneous feasibility (Implied transmission flows) ^e				
i, k line	400	130	400	365	365
i, j line	-100	-100	-100	100	100

(a): The sum of the quantities of each TCC type issued multiplied by the price. (b): From the example day-ahead market in figure 1, the difference between the prices retailers pay and generators get paid in that market. (c): From the equilibrium auction quantities of each TCC type in the auction and the realized prices the example day-ahead market in figure 1, with $LMP_i = 80$, $LMP_j = 200$ and $LMP_k = 200$. (d): As explained in the cash flow description, the transmission service charge (TSC) reduction is the amount that transmission ratepayers effectively gain under the given auction and day-ahead market scenario. (e): The simultaneous feasibility constraints are shown in equations (3) and (4).

Example 1 - an “ideal” solution

Example 1 displays the TCC auction solution for the program described in equations (2)-(5) with bids for 2000 i, k derivatives at \$120 per unit, 600 j, k derivatives at \$100 per unit, and no bids on the i, j product.²⁰ These bids could reflect the physical suppliers of energy at nodes i and j

²⁰The objective function is $\max_q 120 \cdot q_{i,k} + 100 \cdot q_{j,k}$, where $q_{i,k}$ is the allocation to the i, k bidder and $q_{i,j}$ is the allocation to the i, k bidder.

wanting to use TCCs to sell at node k prices. The solution to the auction problem has 300 i, k TCCs and 600 j, k TCCs being issued. Assuming that the subsequent electricity market outcomes are as described in figure 1, the TCC quantities are equal to the quantities of generation at i and j . The equilibrium TCC prices are $p_{i,k} = \$120, p_{i,j} = \$20, p_{j,k} = \$100$, exactly equal to the realized LMP price differences between these locations.²¹

Example 1 implications: In situations where there are many bids on TCCs between generation and consumption locations, the duality between the TCC auction and the physical market simultaneous feasibility constraints results in equilibrium quantities of contracts that match the realized net flows in the market. This includes a zero quantity being issued on the i, j product, with the i, j price pinned down by the bids on the other products. When the issued contracts match the realized net flows in the market, the merchandising surplus is fully securitized. Finally, when TCC prices equal realized TCC payouts, transmission ratepayers are not transferring wealth to TCC holders.

Example 2 - an “under allocation” solution

Example 2 replicates Example 1 with an adjustment that only 30 i, k TCCs are demanded in the auction. TCC prices do not change, however, the simultaneous feasibility constraint in equation (4) would be violated if 600 i, k TCCs were to be issued, resulting in equilibrium quantities of the j, k TCC falling to $q_{j,k} = 330$, with $q_{i,k} = 30$ and $q_{i,j} = 0$.

Example 2 implications: Reduced demand on a given TCC product can reduce the quantity available for other TCC products, due to the simultaneous feasibility constraints imposed by the auction mechanism. The implied transmission flows from the quantities of issued contracts uses less transmission than in the first example, and would be suboptimal if realized in the subsequent physical market.²² Therefore, contracts are under-allocated and the merchandising surplus is not fully distributed to TCC holders. However, given that TCC prices match the realized LMP differences, transmission ratepayers still have their transmission service fee reduced by the same amount as in Example 1.²³

²¹The price solution is not unique in this case, where $p_{i,k} = \$120, p_{i,j} = \$20 + x, p_{j,k} = \$100 - x$ would also be feasible. The solution in the stylized examples in this section chooses prices among the feasible price sets to also maximize auction revenues.

²²If these flows were actually the realized quantities in the day-ahead electricity market depicted in figure 1, production would be inefficient because it would require substitution away from cheaper sources of generation to more expensive sources.

²³This is under the solution rule that chooses prices among the feasible price sets to also maximize auction revenues. Under the pricing formula outlined in Hogan (2002), $p_{i,k} = -100, p_{j,k} = 100, p_{i,j} = -200$, but the allocation is unchanged. The price rule was chosen for expositional reasons, with the purpose of the example to demonstrate the

Example 3 - trader profits from increasing available quantity and signaling optimal power flows

Example 3 replicates Example 2, but adds a financial trader that is willing to buy 2000 i, k products at a price of \$110 (italicized in table 1). The equilibrium allocation returns to that in example 1, with 300 i, k TCCs and 600 j, k TCCs being issued, so the trader participation on the i, k product increased the available quantity of the j, k product.²⁴ However, prices change to $p_{i,k} = \$110, p_{i,j} = \$10, p_{j,k} = \$100$. Assuming the realized payouts are derived from the electricity market in figure 1, the i, k derivative holders are buying the products for \$10 less than the realized contract payout. Therefore, TCC holders earn trading profits of $\$10 \cdot 300$, and consequently, transmission ratepayers receive $\$10 \cdot 300$ less than what they received in Example 1.

Example 3 implications: When demand by physical firms is low for a given product, traders that submit low priced bids for this product can profit and increase the available quantity of other products by doing so. Traders in this example have expanded the transmission capacity of the contract network. This equilibrium signals the optimal power flow configuration that can occur in the physical market and could provide benefits to physical firms by signaling the likely production levels throughout the network which could lead to more efficient forward contracting or production decisions. Formally modeling the benefits from traders increasing the available quantity of TCCs to physical efficiency gains is difficult without imposing further theoretical structure on the model.²⁵ Empirically, Jha and Wolak (2015) demonstrate the plausibility that trader participation in electricity markets can lead to better production efficiency in the context of virtual bidding.²⁶

Example 4 - trader profits from increasing available quantity by purchasing a different product

Example 4 modifies Example 3 by moving the trader bid on the i, k product to the i, j product, bidding \$10 for 2000 units (italicized in table 1). The i, j product pays differences between generator nodes in this example, so it is unlikely to form a role in any physical firm's energy procurement strategy. The auction solution allocates 435 of the i, j TCCs to the financial trader and fully al-

quantity impact of low demand for one product, with price impacts demonstrated in later examples.

²⁴The objective function changes to $\max_q 120 \cdot q_{i,k}^1 + 110 \cdot q_{i,k}^2 + 100 \cdot q_{j,k}$, where $q_{i,k}^1$ is the allocation to the 30 unit bidder, and $q_{i,k}^2$ is the allocation to the financial trader.

²⁵As a non-rigorous illustration, strategic generators at j may be more competitive in their supply if they expect more competitive generation at i , which could be signaled by this new trader-assisted TCC auction equilibrium.

²⁶The removal of barriers for financial traders to submit virtual bids to day ahead electricity markets is shown in Jha and Wolak (2015) to have lowered total generation costs in the Californian market.

locates the 600 j, k TCCs demanded.²⁷ The TCC composition of this solution differs to that in Example 3, but traders still profit at the expense of ratepayers.

Example 4 implications: Financial trader participation on products that do not match the injections and withdrawals of electricity in the physical electricity market can still improve contract allocations to physical firms and expand the set of contracts that can be issued. This is because of the simultaneous feasibility constraint (equation 4), where implied flows on one transmission line can free up congestion and allow more flows on different transmission lines, improving the quantity of TCCs available in the market and signaling the potential for 600MW of energy to be transmitted between j and k . Traders can profit in such a scenario, resulting in a smaller reduction in the transmission service charge.

Example 5 - trader competition increases available quantity and erodes trading profits

Example 5 adds extra competition to example 4. Suppose competition amongst traders to purchase the potentially mispriced i, j product induces an extra bid for 2000 i, j products at a price of \$20.²⁸ Now, the total allocations for each product are equal to those in example 4, but the extra competition from the trader bid on the i, j product has resulted in TCC prices adjusting back to be equal to the realized LMP price differences in the electricity market, leaving the collection of TCC holders with zero trading profits.

Example 5 implications: Trader competition on TCCs that are not used by physical firms as part of their procurement strategy can both improve the quantity of TCCs available in the contract market and restore all TCC prices in the network to actuarially fair prices. Therefore, trader competition can reduce trading profits and the consequent transfers from transmission ratepayers to TCC holders.

1.4 Summary of TCC auction examples

TCC auctions provide physical firms (retailers and generators) the opportunity to purchase a derivative that can change their price exposure to that of a different location. Financial traders can also

²⁷The objective function changes to $\max_q 120 \cdot q_{i,k} + 100 \cdot q_{j,k} + 10 \cdot q_{i,j}$.

²⁸The objective function changes to $\max_q 120 \cdot q_{i,k} + 100 \cdot q_{j,k} + 10 \cdot q_{i,j}^1 + 20 \cdot q_{i,j}^2$, where $q_{i,j}^1$ and $q_{i,j}^2$ are the allocations to the financial traders.

participate, and with the numerous products available, there may be opportunities to profit. Even when financial firms purchase products that are not counterpositions to physical firms, under this auction mechanism they can still increase the available quantity of TCCs that physical firms purchase, and signal the optimal configuration of the network that may allow firms to more efficiently enter forward contracts with each other. A further feature of this mechanism is that price transitivity holds throughout the network. Therefore, all products will be priced even if the equilibrium quantity is zero. The empirical portion of the article will study the sources of profits in TCC auctions, their persistence, and their relationship to the liquidity of the product (whether it had a zero equilibrium quantity) in prior auctions.

2 Setting: The New York TCC market

2.1 Defining a *derivative* and a *contract*

The average monthly payouts of the derivatives studied in this article take the following form:

$$r_{i,j,T_1,T_2} = \frac{1}{m(T_1, T_2)} \sum_{h=T_1}^{T_2} \underbrace{(LMP_{j,h} - LMP_{i,h})}_{\text{Price swap}} \quad (6)$$

where r is the average monthly revenue (or payout) to the derivative holder, i and j index nodes, T_1 and T_2 denote the first and last hour of payments the derivative covers and $LMP_{x,h}$ denotes the electricity price per MWh at location x in hour h .²⁹ $m(T_1, T_2)$ is the duration of the derivative payouts in months, either being 1-, 6- or 12- months and all derivatives start and end on the first and last hour of a calendar month. In finance terminology, $LMP_{j,h} - LMP_{i,h}$ is a locational price swap; in electricity market terminology, $LMP_{j,h} - LMP_{i,h}$ is the congestion price difference between a point of injection (POI) i and a point of withdrawal (POW) j , with the price being that of the day-ahead market. The price for this derivative is also standardized to a monthly average, denoted $p_{i,j,T_1,T_2,t}$, where t indexes the auction it was sold in.

Throughout, a *derivative* will refer to the (i, j, T_1, T_2) financial product with payouts defined by

²⁹ $LMP_{x,h}$ consists of three components, the price at a reference node plus a component that captures line losses and a congestion component. Line losses tend to be small and transmission congestion contracts pay the difference in the congestion component of the prices, where $LMP_{i,h} - LMP_{j,h} \approx CP_{i,h} - CP_{j,h}$ where CP is the congestion component of the price.

equation (6). A *contract* will refer to a q unit position purchased on the (i, j, T_1, T_2) derivative by a firm f . The payout of the contract is $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot r_{i,j,T_1,T_2}$. An example contract follows:

- Transmission congestion contract from Linden Cogen (POI) to N.Y.C. (POW) for each hour between May 1 2008 - April 30 2009, for 3 units
 - Nodes/locations: $i = \text{Linden Cogen}, j = \text{N.Y.C.}$
 - Start and end hour: $T_1 = 12\text{am May 1 2008}, T_2 = 11\text{pm April 30 2009}$
 - Length: $m(T_2, T_1) = 12 \text{ months}$
 - Quantity: $q_{i,j,T_1,T_2,f} = 3$
- Purchased at auction for \$90,110.07 by J. P. Morgan Ventures Energy Corporation
 - Total contract expenditure: $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot p_{i,j,T_1,T_2} = \$90,110.07$
 - Derivative average monthly price: $p_{i,j,T_1,T_2} = \$\frac{90,110.07}{3*12} = \$2,503.06$
 - Firm: $f = \text{J. P. Morgan Ventures Energy Corporation}$
- Locational price differences ($LMP_{POW} - LMP_{POI}$) accrue hourly
 - Total contract payout: $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot r_{i,j,T_1,T_2} = \$132,045.15$
 - Derivative average monthly payout: $r_{i,j,T_1,T_2} = \$\frac{132,045.15}{3*12} = \$3,667.92$
 - Derivative average monthly realized profit: $r_{i,j,T_1,T_2} - p_{i,j,T_1,T_2} = \$1,164.86$
 - Total contract realized profit: $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot (r_{i,j,T_1,T_2} - p_{i,j,T_1,T_2}) = \$132,045.15 - \$90,110.07 = \$41,935.08$

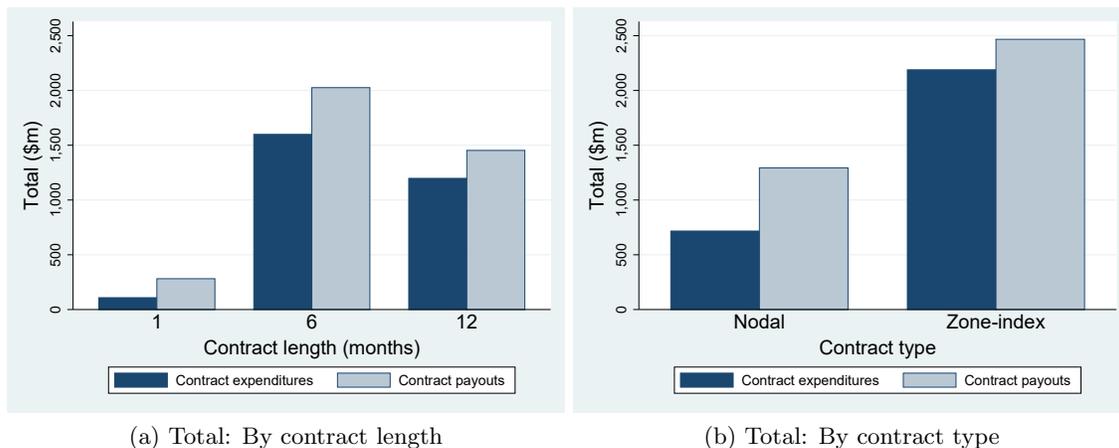
The remainder of this section outlines the product specifications available for purchase, the firm types that participate in this market and the timing of the auctions and the public release of auction outcomes.

2.2 Derivative specifications available for purchase

A wide variety of transmission congestion contract specifications can be purchased at auction. In the T_1, T_2 time horizon dimension, all products studied are of 1-, 6- or 12- months duration. 6- and

12- month contracts attract the greatest expenditure by firms (figure 2 (a)). Collectively, holders of all derivative durations earned revenues greater than expenditures from their contract positions in the NYISO TCC auctions from 1999-2015.

Figure 2: Contract expenditures and payouts by contract specification



Figures (a) and (b) display the sum of $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot p_{i,j,T_1,T_2}$ and $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot r_{i,j,T_1,T_2}$. Sample sizes (number of contracts entered between 1999 and 2015 for each contract grouping) equal to 38,822 for 1- month contracts, 24,412, for 6- month contracts, 14,238 for 12- month contracts, 68,125 for nodal contracts and 9,347 for zone-indexed contracts.

In the location dimension, there are 450 price nodes in the New York grid, resulting in approximately 100,000 i, j derivative specifications available.³⁰ A map of the transmission network and these nodes is shown in figure 3.

In addition to the price nodes, figure 3 displays 11 price zones. Nodal derivatives pay the difference in the electricity prices at the two nodes. Zone-indexed derivatives pay the difference between two zonal prices (z_1, z_2) , which are a quantity weighted average of the nodal prices where

³⁰450 locations allows for $450 \cdot 449 = 202,050$ directional location pairs. Given that $r_{i,j} = -r_{j,i}$ and all other variables share this transitive property, this number is halved to give 101,025 observations. The number of locations is not constant across all auctions, with some nodes being added and removed over the sample window.

electricity is withdrawn in a given zone, with payouts equal to:

$$\begin{aligned}
 r_{z1,z2,T_1,T_2} &= \sum_{h=T_1}^{T_2} (LMP_{z2,h} - LMP_{z1,h}) \\
 &= \sum_{h=T_1}^{T_2} \left(\sum_{j \in z2} w_{j,h} LMP_{j,h} - \sum_{i \in z1} w_{i,h} LMP_{i,h} \right)
 \end{aligned} \tag{7}$$

Mixed derivatives that pay the price difference between a node and a zone-index are classified as nodal.

The distinction between nodal and zonal products is important because in the NYISO market, producers of electricity receive nodal prices whereas consumers of electricity pay the zonal prices, described in equation (7). Therefore, different firms may demand different products depending on their operations in the wholesale market (See Tangeras and Wolak, 2017, for more detail on the competitive impacts of nodal versus zonal prices on demand).

Zonal contracts attract the greatest expenditure (figure 2 (b)), despite having far fewer potential specifications available and many less overall contracts issued. Collectively, holders of both derivative types earn revenues greater than expenditures from their contract positions, however, nodal contract holders receive proportionally larger revenues than their expenditures compared to holders of zonal contracts.

2.3 Participants in the derivative market

133 firms were awarded a TCC in the New York market between 1999-2015, 117 of which purchased a TCC at auction.³¹ This subsection describes the three firm types (retailers, generators and traders)³² that participate in this market and the motives for their participation. Descriptive statistics on the expenditures and payouts of the positions entered in the TCC market for each group are in figure 4, and a full list firms and their firm type classifications is in appendix A.3.

³¹Some firms were directly awarded TCCs under a grandfathering arrangement. These contracts are not studied in this article - only derivatives that were bought or sold for a price set at auction are considered.

³²Toole (2014) also classifies firms into three firm types that he labels as hedgers, speculators and unknown.

Figure 3: Map of the NYISO network

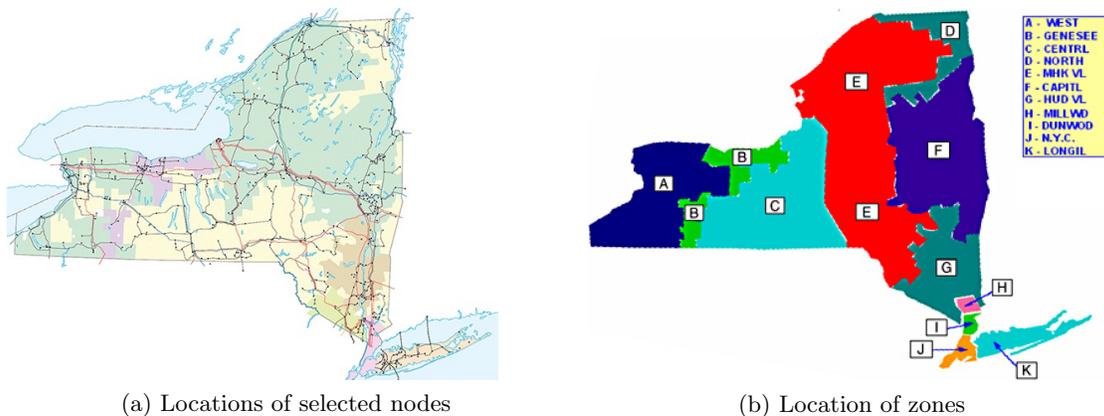
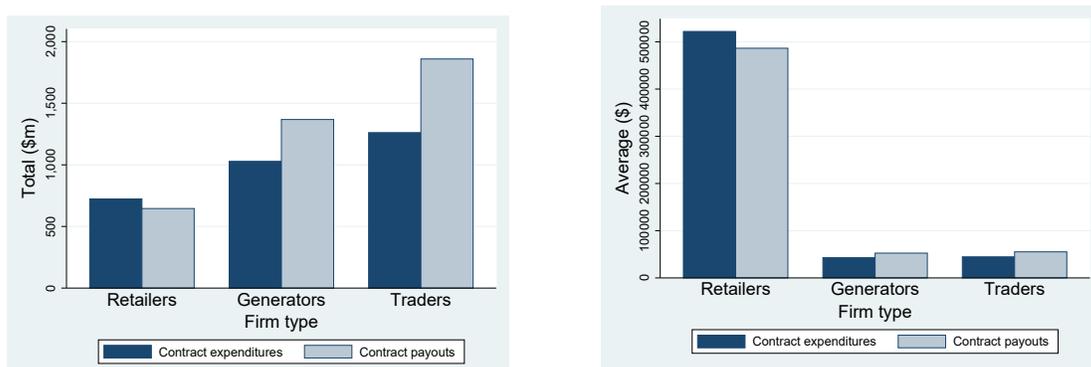


Figure (a) displays major transmission lines and ownership regions, with a selection of nodes. Figure (b) displays the zones which use the contained nodes to form a weighted price index. External electricity markets that have import/export price nodes are ISO-NE (East), Hydro Quebec (North), Ontario Hydro (North West) and PJM (South).

Figure 4: Contract expenditures and payouts of participants



(a) Total expenditures and payouts

(b) Average expenditures and payouts (restricted to purchases)

Figure (a) displays the sum of $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot p_{i,j,T_1,T_2}$ and $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot r_{i,j,T_1,T_2}$. Figure (b) displays the average of $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot p_{i,j,T_1,T_2}$ and $q_{i,j,T_1,T_2,f} \cdot m(T_2, T_1) \cdot r_{i,j,T_1,T_2}$. Sample sizes (number of contracts entered between 1999 and 2015 for each firm type) are 3,295 for the retailers, 59,425 for the generators and 76,905 for the traders. Average expenditures and payouts are constructed only using contracts purchased for a positive price due to compositional differences in the amount of long and short positions entered by each firm group.

2.3.1 Retailers

Firms that purchase electricity from the New York wholesale electricity market to meet the consumption demands of their customers are classified as retailers. Retailers that are the sole provider for a geographic region are regulated in the prices they can charge their retail customers. In New York, retailers face spatially aggregated zonal prices for the electricity they withdraw from a given node.

Overall, retailers pay slightly more for their TCCs than the TCCs pay out, and retailers are the smallest participant group in terms of total derivative expenditure (figure 4a). However, retailers spend approximately 10 times more per contract than other firm groups (figure 4b), meaning that retailers buy contracts with larger quantities and durations than other firms.

2.3.2 Generators

Firms that own electricity generating plants in New York that are not retailers are classified as generators. These firms supply electricity and may have local market power at the price nodes where their generating units are located. Any market power diminishes at other price nodes.³³

Generator participation in the TCC market has been theoretically scrutinized in Bushnell (1999) and Joskow and Tirole (2000) because generators have the ability to influence electricity spot prices (and therefore TCC payouts) via their production decisions. If a generator can influence the payout of a particular derivative, the derivative is worth more in their hands than anyone else – generators may exercise market power to increase their TCC payout. Such a situation would not be economically efficient if the TCC holder is a low cost generator and its production is replaced by a higher cost source. As summarized by Bushnell, auctions of TCCs could result in contracts “flowing to those that can abuse them the most.” An implication from the theory will be examined in section 4.2.2.

Overall, generating firms received net payouts on their derivative positions of \$1,367m, exceeding their net expenditure by \$340m from 1999-2015 (figure 4a).

³³Market power is related to competitor locations and the capacities of the transmission grid. If transmission capacity was infinite, the ability of a generator to influence prices would be the same across all nodes (and there would be a uniform price).

2.3.3 Traders

All remaining firms with no physical interests in the New York electricity market are classified as financial traders. These firms are largely investment banks or energy traders. These firms are motivated to make a profit in this market, by purchasing underpriced products and selling overpriced products. TCC profits solely determine the success of a trading firm or the TCC division of the firm, whereas TCC profits are only a small portion of total revenues for retailing and generating firms. Section 1.3 showed how trader activity in TCC auctions can benefit the physical players in the market.

Overall, trading firms received net payouts on their derivative positions of \$1,859m, exceeding their net expenditure by \$598m from 1999-2015 (figure 4a).

2.4 Information release and the sequence auctions in New York

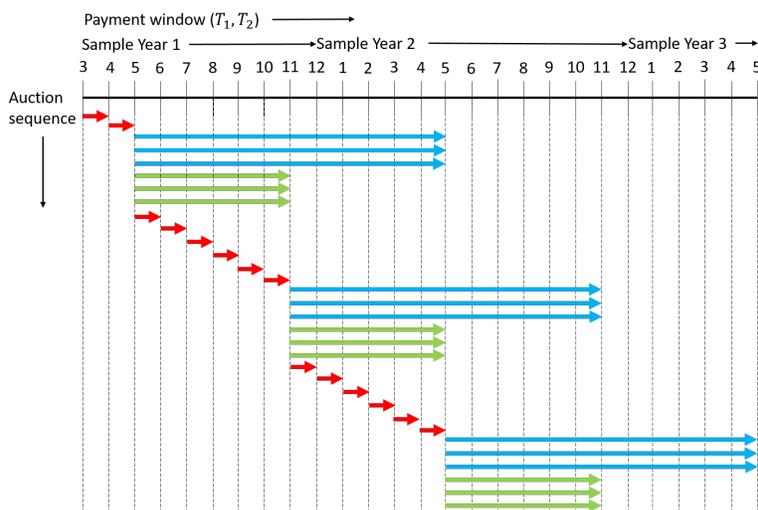
A single auction (indexed by t) allocates a set of TCCs that have a common time horizon, defined by T_1 and T_2 in equation (6). Firms (f) can bid to buy, or offer to sell, any of the $\approx 100,000$ possible i, j location pairs with this time horizon. The auction process is extended from that described in section 1.3 to match the network configuration of New York (more detail in Appendix A). The size of the position awarded to a firm on a derivative product is denoted $q_{i,j,t,f}$. Auction prices are transitive in the location nodes ($p_{i,k,T_1,T_2,t} = p_{i,j,T_1,T_2,t} + p_{j,k,T_1,T_2,t}$) and the issued contracts (the collection of $q_{i,j,t,f}$) are simultaneously feasible.³⁴

There are two crucial features of the allocation process that will be utilized in the analysis. First, prices are observed for every derivative. Even if a firm is not allocated a contract on a given derivative, a price is set and represents the price at which the market operator would have sold or bought a derivative had bids above or offers below that price been placed. For example, in a three node system, the i, k derivative and the j, k derivative might have had bids placed on them, and given the constraints on the auctioneers problem, this is enough to set a price for the i, j derivative that did not receive a bid.

Second, the auctions are sequential, with restrictions on information flows to participants. Figure 5 displays a representation of the auction structure, with the duration of the derivative specified in

³⁴The simultaneous feasibility constraint is discussed in further detail in Appendix A, where the rules to account for overlapping time horizons of different products sold at different auctions are explained.

Figure 5: Order of auction vintages and their payout windows



Derivatives of 1 month duration are red, 6 month duration are green and 12 month are blue. The length of the arrow covers the payout period for a derivative. The auctions for each vintage occur in order from the top of the diagram to the bottom.

the horizontal dimension and the order in which auctions occur in the vertical dimension. The 6- and 12- month derivatives either begin in May or November, with each vintage auctioned in three to five tranches, one week apart. The 1- month derivatives are available for each month of the year, sold in a single auction.³⁵ Entering an auction, firms have access to public information on prior auction prices for every possible TCC and the payout of the TCC if the payout period has been realized. Further, all issued contracts are reported, containing complete information on the TCC specification, the size of the contract and the firm that purchased it. However, the bids in TCC auctions are released 3 months after the fact and with anonymized identifiers placed on the location nodes and the firm identities.

³⁵There are occasional auctions for TCCs that cover 24 months of payments, but only 1-, 6- and 12- month auctions have occurred on a consistent schedule each year.

3 Data sources

Data on derivatives and contracts are available to the public at the NYISO TCC website,³⁶ with mechanical details of the data construction found in Appendix A.

3.1 Contract data

Contract observations are defined by i, j, T_1, T_2, f , the locations (i, j) and time horizon (T_1, T_2) specified in the derivative contract purchased, and the firm (f) that purchased the contract. The key variables are the prices, payouts and quantities of the contract.³⁷ Data for all contracts are available since the market began in 1999. There are 139,625 contracts in the contract dataset.

3.2 Derivative data

Derivative observations are defined by i, j, T_1, T_2, t , the locations and time horizon specified in the derivative, and the auction t that it was sold in. Each auction t has attached a common duration window T_1, T_2 for all i, j derivatives (T_1, T_2 will be dropped in later notation). Derivative data are available for 235 auctions from November 2006 to December 2015. There are approximately 450 nodes available to be used in a derivative specification each auction³⁸, giving approximately 100,000 i, j location pair observations per auction t . This gives approximately 23,500,000 i, j, t observations.

The number of derivative observations greatly exceeds the number of contract observations. The auction mechanism sets prices for each derivative in every auction regardless of whether a firm purchased any given derivative (refer to section 1.3 and 1.4). The derivatives studied are restricted to types purchased by firms over the sample window. There are 304,039 unique $(i, j, m(T_1, T_2))$ derivative types, where $m(T_1, T_2)$ is the number of months the derivative spans. The sample is restricted to the 14,969 of 304,039 unique $(i, j, m(T_1, T_2))$ types where a contract was ever issued, leaving 1,151,374 i, j, t derivative observations. Attached to each observation are price and payout per month duration variables $p_{i,j,t}$ and $r_{i,j,t}$. Both directions of a derivative are not included in the data because it is a duplication with $p_{i,j,t} = -p_{j,i,t}$ and $r_{i,j,t} = -r_{j,i,t}$.³⁹

³⁶http://www.nyiso.com/public/markets_operations/market_data/tcc/index.jsp

³⁷In the raw data, i, j, T_1, T_2, f does not uniquely identify each observation. This is because a firm that bids a step function will get an issued contract for each step that clears at auction. Given that the price per unit is the same, I aggregate these into one observation and add the size of each contract into the single, unique observation.

³⁸The number of nodes increased over the sample window.

³⁹The empirical analysis does not tend to be sensitive to the direction of the derivative when listed either in

3.3 Auction bid data

Bidding data is released three months after each auction and lists anonymized identifiers rather than the names of the POI and POW locations and the identity of the firm. With these anonymized identities, the auction data contain all bids as defined by the firm, the product and the auction, with information on the quantity of units demanded and the bid price. I have compiled bid data from 2006-2015. I describe an algorithm for decoding a subset of these identifiers in Appendix A. When comparing the decoded subset of auction dataset to the contract dataset, only 10% of the expenditures and profits are earned on locations that were not decoded or by firms that were not decoded.⁴⁰ The auction dataset is used in some descriptions of firm bidding behavior, with the sample outlined at a case-by-case basis in the analysis.

4 Firm participation and trading profits in TCC auctions

Section 4.1 describes firms' participation and purchases. Section 4.2 investigates which firms earn systematic trading profits on which products. Profit sources are also investigated for generators that purchase TCCs at locations where they own generating units.

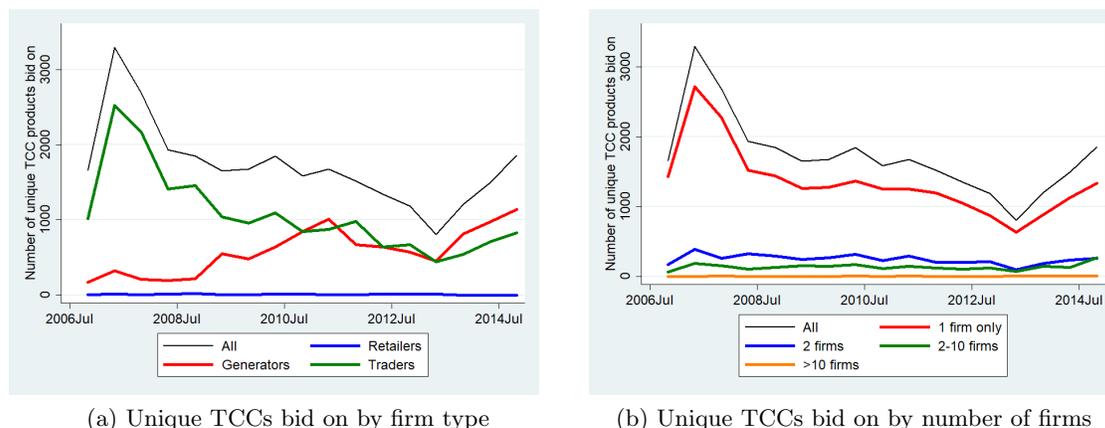
4.1 Participation of firms in TCC auctions

Figure 6a displays the number of unique TCC products bid on by retailers, generator owners and financial traders in every 12-month auction since 2006. Retailers bid on a tiny portion (less than 1%) of the products that generators and traders bid on. On average, all retailers collectively bid on 7 different TCC specifications for each vintage of 12 month auctions, whereas generators and traders bid on 581 and 1,069 different products. Figure 6b displays the number of firms that place a bid on each TCC that received at least one bid at auction. Very few of the $\approx 100,000$ permutations of location-pair derivative specifications are bid on in each auction, with even less products receiving

the direction that faces a positive price or if assigned arbitrarily in the direction from the location with the larger identification number to the lower identification number.

⁴⁰See appendix A. For the decoded locations, market clearing prices can be applied to allocate clearing quantities to participants, and realized revenues can be applied to recover ex-post contract profits. The total profits when split across classes and firms in the auction data are proportional to the total profits from the corresponding period in the awards data. Enough identities are recovered to cover 90% of the contract expenditures and profits from the contract data but only 45% of total contracts. The decoded data is more likely to contain locations that are more frequently specified in issued contracts.

Figure 6: Unique TCC bids in 12-month auctions, 2006-2015



Figures display the number of unique TCC products (defined by two locations) bid on for each vintage of 12 month TCC products. Vintages are either November to November or May to May. The firm type counts in figure (a) include the firms that were decoded from the auction data, described in appendix A. Otherwise, all firms and all TCC locations are included, regardless of whether the true location was decoded.

bids from multiple firms.

Despite the small set of products retailers bid on, they are not insignificant in their participation. Retailers account for 16% of derivative expenditures, with 84% of retailer expenditures on zone-indexed contracts and 96% on 6 or 12 month duration contracts. Generator owners account for 33% of derivative expenditures, and financial traders account for the remaining 51%.⁴¹ Retailers on average enter much larger and longer positions (figure 4), consistent with their larger exposure to the spot market and potentially a hedging motive. Individually, the majority of contracts held by generators and traders are for small positions that are tiny relative to the aggregate price exposure faced by major retailers in the procurement of electricity or relative to the sale of electricity by generators.⁴²

Overall, retailers restrict their participation to large purchases of zonal products, whereas generators and traders buy a mix of both zonal and nodal products, often in small quantities. The radically different purchase behavior of retailers to generators and traders could be explained by

⁴¹TCC expenditures are displayed in table B1 in the data appendix.

⁴²The median contract size of 6- and 12-month TCCs is 5 units for generators and 3 for traders. In 2015 Orange and Rockland Utilities, Inc. purchased an average of 655MWh of electricity from wholesale markets each hour and received approximately \$75,000 each hour from its customers. In 2015, Consolidated Edison's New York City retailer averaged approximately 10 times those figures (Consolidated Edison Inc. (2015), pages 20 and 24.).

regulatory incentives. Retailers face zonal prices in the wholesale market and the prices retailers can charge their retail customers are determined via public utility commission rate-setting processes. There may be some risk to retailers that losses from trading activity not linked to the procurement of energy will be disallowed and profits from such activities could be used to lower retail prices. With generator owners and traders able to keep any trading profits they earn, we see that they are much more likely to bid on and purchase a wide range of products, even in small quantities.

4.2 Systematic trading profits across firm and product types

Prior work on TCC auctions identified that contract prices were not equal to expected contract payouts (Bartholomew, Siddiqui, Marnay, and Oren, 2003; Hadsell and Shawky, 2009; Adamson, Noe, and Parker, 2010; Olmstead, 2018). I extend this work in two dimensions. First, instead of studying issued contracts, I study derivatives in order to use the information contained from products that had a zero equilibrium quantity in the auction.⁴³ Second, I explore the link between derivative product design, the types of firms that profit on each product and whether profits might be linked to downstream payout manipulation.

A derivative's price should equal the expected payout of the derivative under free entry of risk neutral firms and no private information. That is, if each derivative auctioned has some expected payout $E(r) = \mu$, then its price p should equal $E(r)$. Consider a single derivative that is auctioned, and denote $I = 1$ when it is purchased by firm I and $I = 0$ otherwise. Then,

$$\underbrace{E(r|I, p)}_{\text{Deriv. payout}} = \underbrace{p}_{E(r|I=0)} + \underbrace{E(\mu - \beta(t_I)|I = 1)}_{\text{Rent to firm I}} \cdot I \quad (8)$$

In equation (8), $\beta(t_I)$ is the bid firm I places when it receives some signal t_I . Under the assumptions of free entry of risk neutral firms and complete information, conditional on firm I being awarded the object, $\beta(t_I) = \mu$. However, if the assumptions of risk neutrality or complete information are violated then it could be that $\beta(t_I) \neq \mu$ when firm I is awarded the object. If firm I values the derivative at more than its expected value, or it persistently overestimates the payout, then $\beta(t_I) > \mu$ when $I = 1$. If firm I has the ability to purchase derivatives for less than its expected

⁴³A similar analysis to these previous papers that uses the contract data is found in section A.4 of Leslie (2018), which highlights that tests for prices being equal to expected payouts can not be rejected for zonal products, but is rejected for nodal products.

payout, or that all firms value the derivative at less than its expected value, then $\beta(t_I) < \mu$ when $I = 1$.

Define $I_{i,j,t,f}^q$ as an indicator set to 1 if firm type f was issued an i, j TCC in this auction, set to -1 if firm type f was issued an j, i TCC in this auction (sold the i, j derivative) and set to 0 otherwise. The model to be estimated is a statistical analogue to equation (8) and has the following specification:

$$r_{i,j,t} - p_{i,j,t} = \alpha + \sum_{f \in F} \delta_f I_{i,j,t,f}^q + \epsilon_{i,j,t} \quad (9)$$

$p_{i,j,t}$, $r_{i,j,t}$ are the average monthly prices and payouts. Derivative payouts exclude any discount factor.⁴⁴ The direction of the i, j derivative is chosen such that $p_{i,j,t} > 0$, with the j, i derivative excluded to avoid double counting.⁴⁵ Under this organization of the data, if α is non-zero then there is a different risk premium attached to buying versus selling the derivative. Therefore, α is equal to the expected payout premium for derivatives not purchased but with a positive price, which should be zero with free entry of risk neutral firms. If $\alpha = 0$, δ_f is equal to the expected difference between the payout of the derivative and the market clearing bid when firm f is awarded the derivative. $\delta_f = 0$ implies that when the firm purchases the object, it on average receives a payout equal to the price it paid for the object. If $\delta_f > 0$ the firm on average enters profitable contracts, either receiving a payout greater than the price it pays for the derivative or that it pays a payout less than the price it was paid. If $\delta_f < 0$ the firm on average enters unprofitable contracts.

To emphasize the nature of the derivative data, whereby prices and payouts exist for each derivative regardless of whether a firm was actually issued a contract on that derivative, note that a retailing firm is issued a contract in 0.1% of observations ($\overline{|I_{i,j,t,RET}^q|} = 0.001$), with generators and traders each issued contracts for 3% of the derivative observations ($\overline{|I_{i,j,t,GEN}^q|} = 0.031$, $\overline{|I_{i,j,t,TRA}^q|} = 0.034$).

⁴⁴The small payout lengths and monthly payouts mean that applying a discount rate correction has a negligible impact on the results.

⁴⁵The results are not sensitive to arbitrarily listing the derivatives in the i, j direction based on their identifying code.

4.2.1 Estimates of derivative prices and payouts

Columns I-III of table 2 report the estimates of the parameters in equation (9) for all derivatives and for partitioned samples of the nodal and zone-indexed derivatives. The majority of products available in this market are nodal, but we earlier saw that zonal contracts attract greater total expenditure. The unit of observation is a location pair derivative available in auction t . Estimates are obtained using ordinary least squares and the standard errors are clustered at a vintage T_1, T_2 level given the transitivity property of prices and payouts.

Table 2: Estimates of average monthly derivative payouts

	I	II	III	IV
	All	Zonal	Nodal	Nodal
	$\overline{p_{i,j,t}} = 1686$	$\overline{p_{i,j,t}} = 3821$	$\overline{p_{i,j,t}} = 1667$	$\overline{p_{i,j,t}} = 1667$
α [Constant]	32.99 (91.84)	206.50 (255.21)	31.49 (90.78)	
δ_{RET} [$I_{RET,t}^q$]	-40.67 (125.80)	-105.53 (130.76)	90.26 (174.16)	76.14 (196.43)
δ_{GEN} [$I_{GEN,t}^q$]	93.58 (43.74)	14.41 (115.25)	95.58 (44.85)	58.29 (46.02)
δ_{TRA} [$I_{TRA,t}^q$]	162.25 (38.50)	-84.46 (126.01)	174.53 (37.71)	112.30 (39.22)
i, j pair fixed effect	N	N	N	Y
N	1,151,374	10,506	1,140,868	1,140,868
N_A	235	235	235	235

Each column reports estimates of equation (9) for the denoted sample using ordinary least squares. Standard errors clustered at a vintage level T_1, T_2 reported in parentheses. All contract prices and payouts are divided by the number of months a contract covers. Summary statistics for the variables used in estimation are found in table B2. The direction of the price swap derivative is chosen such that $p_{i,j} > 0$. The estimated coefficients do not substantially differ when the direction is chosen arbitrarily.

First, examine the estimates in column I that pools all derivatives. Products offered but not sold are priced on average \$32.99 less than their payout (α), but this coefficient is not detected to be statistically different from zero at a size 5% hypothesis test. Failing to reject $\alpha = 0$ is consistent with free entry of risk neutral firms. Retailers are predicted to receive an average payout of $\hat{\delta}_{RET} = \$40.67$ less per month of contract payments than the price they pay. Generators and traders are predicted to receive an extra $\hat{\delta}_{GEN} = \$93.58$ and $\hat{\delta}_{TRA} = \$162.25$. However, only for generators and traders are these estimates detected to be statistically different from zero.

Systematic profits are not detected for any firm type on zone-indexed derivatives (column II).

However, systematic profits are detected for generators and traders on nodal derivatives (column III). The average price faced by generators for their nodal derivatives is \$804.53, and using the estimated value of δ_{GEN} implies an average payout premium to generators of $\frac{95.58}{804.53} \cdot 100 = 11.9c$ per dollar value of the position. The equivalent calculation for financial trading firms estimates an average payout premium of $\frac{174.53}{1113.09} \cdot 100 = 15.7c$ per dollar.

To summarize, the results add statistical robustness to the observations in figures 2 and 4, that systematic profits are only earned by generators and traders on nodal products. Retailers are not detected to earn trading profits and were earlier shown to largely confine their participation to the zonal products (that are priced actuarially fair).

Column IV reports estimates for equation (9) when replacing α with fixed effects for each i, j pair. This allows for each location-pair to have a systematic premium or discount, meaning that the δ_f coefficients measure whether different firm types earn systematic profits in excess of any location-pair premium. For this specification, only traders are detected to earn a non-zero payout premium, equal to \$112.30 beyond the average payout for the i, j location products they purchase, suggesting that financial traders are adept at buying (selling) the products that are underpriced (overpriced) relative to their usual premium.

4.2.2 Extension: Are generator profits tied to generator operations?

Existing theories for why TCC auctions may result in systematic trading profits to some firms have predicted that generating firms can earn systematic profits from TCC positions tied to their generator operations due to their ability to influence downstream electricity prices (Bushnell, 1999; Joskow and Tirole, 2000). If generators can influence the payout of a TCC by exercising market power, the TCC is worth more in their hands than in the hands a firm that does not have this ability. The results in table 2 emphasize trader profits, however, generating firms also earn systematic profits in this market. This extension examines the possibility that generators might be profiting in the TCC market due to their downstream physical market power.

Monthly average profit ($r_{i,j,f,t} - p_{i,j,f,t}$) for all contracts ever purchased by generators are regressed on indicator variables SZ - denoting the firm owns a power plant in the same zone as one of the i, j locations specified in the derivative, and SN - denoting the firm owns a power plant at the exact node as one of the i, j locations specified in the derivative. The data for power plant

locations is described in appendix A. Only 1,219 of the 23,951 generator held contracts included in the estimates have a location specified in the payout where the holder owns a power plant. 3,832 contracts have a location specified in the payout which is in the same zone as a power plant owned by the holder. The estimates are:

$$\widehat{r_{i,j,f,t} - p_{i,j,f,t}} = 187.2 - 22.6 \quad SZ_{i,j,f,t} + 37.2 \quad SN_{i,j,f,t}$$

(56.7) (82.5) (182.0)

The estimates show that for the 23,951 derivatives purchased at a positive price by generating firms at generating nodes, there is no average profit differential associated with a firm’s power plant ownership at a node specified in the derivative contract.⁴⁶

An implication from the theories in Bushnell (1999) and Joskow and Tirole (2000) is that derivative payouts are increasing in the size of the position held by a generator. If a firm can exercise market power and influence LMPs at certain locations, then their incentive to do so is increasing in their exposure to contracts that are linked to that location. Further, if a firm had some other mechanism available to manipulate the payouts of TCCs at the margin (as demonstrated in a case study of a financial trader in Birge, Hortaçsu, Mercadal, and Pavlin (2018)), we might expect derivative payouts to be increasing in the size of their open position.⁴⁷

Equation (9) is extended to investigate whether derivative payouts are related to the size of firm derivative positions:

$$r_{i,j,t} - p_{i,j,t} = \alpha + \sum_{f \in F} \delta_f I_{i,j,t,f}^q + \sum_{f \in F} \rho_f Q_{i,j,t,f}^{POST} + \epsilon_{i,j,t} \quad (10)$$

where $Q_{i,j,t,f}^{POST}$ is the number of contracts firm type f holds on the i, j derivative.⁴⁸ The estimates

⁴⁶The sample is the 23,951 contracts issued to generating firms at generating nodes for a positive price, with prices and payouts standardized by the length of the contract. Standard errors are clustered at a vintage level (all contracts with the same T_1 and T_2).

⁴⁷Birge et al. (2018) study the positions of a firm that was investigated by the Federal Energy Regulatory Commission. The manipulation under investigation was in a virtual market, where financial traders can offer supply in the day ahead market and close out their position in the real-time market and influence FTR (TCC) payouts. The virtual market trades in question totaled \$390,000, compared to \$1b of positions taken annually in the MISO market for financial transmission rights. Birge et al. examine whether similar behavior is widespread but are impeded by the anonymity of firm identities.

⁴⁸Given the overlapping auction structure shown in figure 5, this value totals all contracts on the i, j derivative with overlapping payouts to the product sold in auction t . $Q_{i,j,t,f}^{POST}$ is negative if the firm type has a positive j, i position.

of this model are displayed in table B3, with ρ_f not detected to be different from zero for any firm type on nodal contracts and the $\hat{\delta}_f$ estimates similar to those in table 2. Regardless, the point estimate $\hat{\rho}_{GEN} = -0.01$ is small and negative.

To summarize, I find no evidence that generator trading profits systematically differ with either the product being tied to locations related to their power plant operation or with the size of their open positions. It is plausible that generators are deterred from exploiting their downstream physical market power in TCC auctions due to regulatory rules that allow market operators to withhold TCC payouts if they determine that a firm exploited their contract position via market power.⁴⁹ Alternatively, this estimation technique might not be statistically powerful enough to detect such actions.⁵⁰

5 Price updating and the persistence of trading profits in TCC auctions

Systematic profits should erode over time in the absence of entry costs as other firms mimic the successful firms. However, profits have not eroded over time (figure 7). For each of the 16 years of auctions, profits from nodal contracts have been positive, whereas zonal contract profits are centered around zero. This is despite a steady year-to-year increase in the number of firms that were observed to purchase at least one contract over the sample window.

If traders are managing to earn systematic profits by purchasing a distinct set of products to the physical firms, there may be a barrier preventing other traders from competing for these opportunities. The unique information revelation structure across sequential auctions can be exploited to examine how trading profits are being earned and whether price discovery occurs following purchases by profitable firms. If firms have constant profit margins over the same products across auctions, this could indicate the presence of a risk premium, an opportunity cost of capital or some

⁴⁹Regulations exist to deter such activity, with the Federal Energy Regulatory Commission having jurisdiction to investigate and potentially withhold payments to TCC or virtual traders if they were found to have altered their downstream strategies because of their TCC positions (Alderete, 2013).

⁵⁰An ideal test for firms possessing the ability or incentive to perform downstream actions to influence asset payouts would be to estimate the impact TCC positions have on their electricity bidding strategies. Given that such data is not available for this market, the indirect test that $\rho_f = 0$ could only identify marginal changes in derivative payouts with contract holdings, whereas a structural model of electricity bidding strategies may be able to identify inframarginal changes in derivative payouts to firm holdings.

other cost to participation. However, if a firm earns a profit on a particular product and the next time that same product is auctioned the profitable opportunity is removed, some other barrier may exist. This could represent well informed firms earning payouts greater than the prices paid (see Wilson, 1967)⁵¹ in the first round of an auction, with their information advantage diminished in subsequent rounds after it has been revealed to the market. Such information diffusion could benefit physical participants.

Figure 7: TCC holder profits and participants, 1999-2015

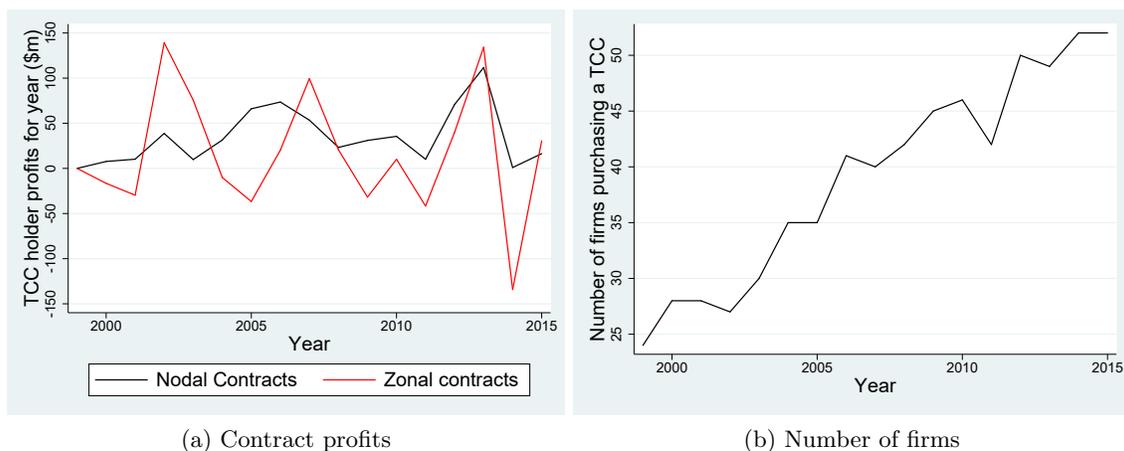


Figure (a) aggregates profits from all TCCs with a start hour in the calendar year. Figure (b) counts the number of firms that were observed to buy at least one TCC in the calendar year.

5.1 Price updating across auction rounds of the same products

The TCC auctions are sequential for contracts of 6- and 12- months duration (figure 5). Each derivative d is defined by (i, j, T_1, T_2) , and denote a given subset of these derivatives as D . For each auction round (ar), the following statistic can be constructed:

$$\frac{1}{|D|} \sum_{d \in D} 100 \cdot \frac{p_{i,j,T_1,T_2,ar} - p_{i,j,T_1,T_2,ar=1}}{p_{i,j,T_1,T_2,ar=1}} \quad (11)$$

⁵¹Examples of studies of private information advantages with similar empirical consequences are found in oil drainage tract auctions (Hendricks and Porter, 1988) and insurance (Chiappori and Salanie, 2000).

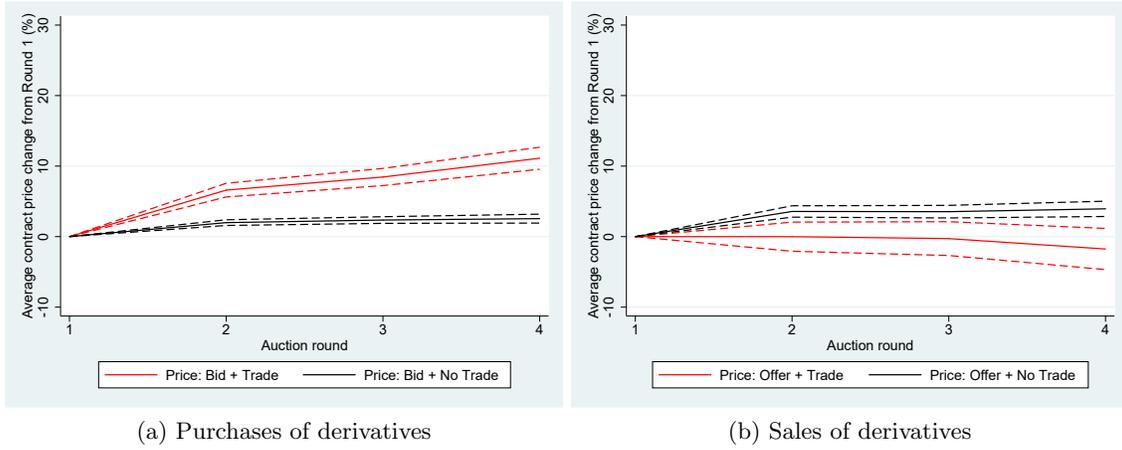
The statistic is the mean of the percentage derivative price change in auction round ar relative to the round one derivative price for products in the set D . The information structure for the sequential auctions (held one week apart) is as follows: Immediately after each auction, the prices for every derivative and the contract awards (including the identity of the firm) are made public. Bids by each firm are not made available to the public in time for the next auction.

The empirical strategy is to estimate the price response of derivatives in subsequent auction rounds following the revelation that a firm was awarded that derivative in the first round of the auction. A comparison group is formed with derivatives that were bid on but not awarded a contract. This comparison group is relevant because both sets of products were in demand, but this information is only released for the set that were purchased. Therefore, if there is information content attached to the award of a derivative, the price of a derivative should rise after it is revealed that a well-informed firm is awarded that derivative, whereas we may not expect to see such a response after a bid that was below the market clearing price. This is because, for 6- and 12- month derivative auctions, the same set of products with the exact same payout specifications are offered across each round. For comparison reasons, the sample used in this section is restricted to the products observed in the auction dataset, described in Appendix A.

Figure 8 plots the average price process for bids and offers as specified in equation (11). Derivative prices appreciate an average 7 to 11% following a purchase in round one award, consistent with some form of information revelation. Prices only appreciate 2% for contracts receiving a bid but without an award (and therefore no revelation that a bid was placed). Equal but opposite responses are not seen for offers to sell. Derivatives with an offer that does not result in an issued contract look similar to bids with no trade, but when a contract is sold, prices do not fall. A potential explanation for the bid/offer asymmetry is that bids to buy TCCs might clear without a willing counterparty. For a sell offer to clear a non-zero quantity, other firms must buy a combination of contracts that form a counterposition. Therefore, we might expect these products to be priced fairly if firms are willing to take opposite positions on related contracts in round 1 of the auction.

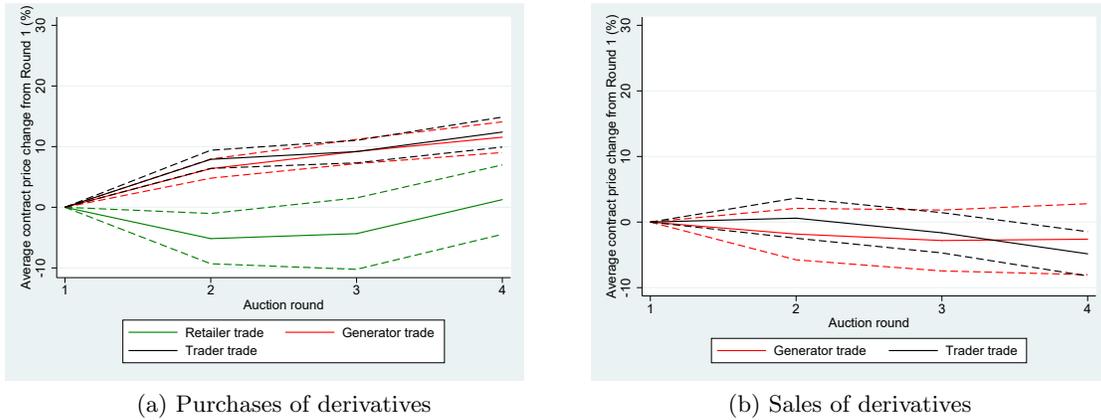
Figure 9 splits the price dynamics by the type of firm awarded a contract in round 1. In the first chart, the sets of derivatives included are all derivatives that were purchased at a positive price, split by the firm type that purchased that derivative. We see that prices do not respond to retailer awards but do respond to generator and trader awards. The second chart is analogous to the first

Figure 8: Price updating following purchases, bids, sales and offers



All series plot equation (11). The purchases chart compares two sets D of derivatives, those that were purchased by any firm at a positive price in round one to those that were not awarded to any firm but receive a round one bid. The sales chart is analogous to the purchases chart but for negatively priced products. Sample restricted to derivatives with $p_{i,j,T_1,T_2,ar=1} > \$1,000$ and where price changes are within a threefold increase or decrease. Means and pointwise 95% confidence intervals plotted. Sample sizes: (a) 2,980 and 9,850. (b) 1,009 and 4,059.

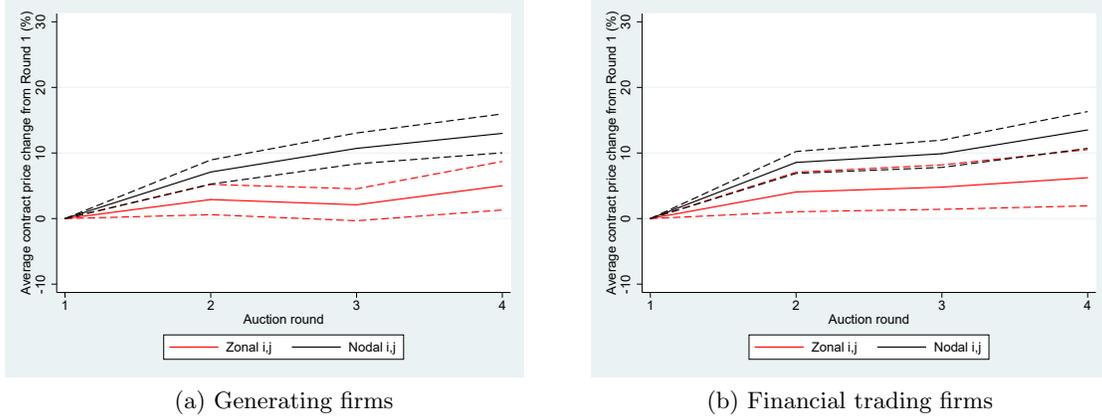
Figure 9: Price updating following purchases and sales, by firm type



All series plot equation (11). The purchases chart compares three sets D of derivatives, those that were purchased by any firm at a positive price across the three firm groupings, retailers, generators and traders. The sales chart is analogous to the purchases chart but for negatively priced products, with retailers excluded for sample size reasons. Sample restricted to derivatives with $p_{i,j,T_1,T_2,ar=1} > \$1,000$ and where price changes are within a threefold increase or decrease. Means and pointwise 95% confidence intervals plotted. Sample sizes: (a) 61, 1,151 and 1,211. (b) 296 and 465.

chart but for derivative sales.⁵² Similar patterns are seen to figure 9 when splitting the sample into profitable and unprofitable firms in figure B2.

Figure 10: Price updating following purchases of zone-indexed and nodal contracts, by firm type



All series plot equation (11). Both charts compare two sets D of derivatives, sets of round 1, positive price purchases split by nodal and zonal derivatives. The first chart plots these sets for generating firm purchases, the second chart plots the sets for trading firm purchases. Sample restricted to derivatives with $p_{i,j,T_1,T_2,ar=1} > \$1,000$ and where price changes are within a threefold increase or decrease. Means and pointwise 95% confidence intervals plotted. Sample sizes: (a) 193 and 958. (b) 173 and 1,038.

Finally, I compare the price responses to generator and trader bids on zone-indexed and nodal contract specifications in figure 10. The market responds more to a nodal contract award than a zonal contract award for both generating and trading firms.

To summarize, the market updates derivative prices following the revelation of purchases by generating and trading firms, particularly for nodal products. Given the auctions studied in this section sell the same sets of products one week apart, it is difficult to attribute the systematic trading profits earned in these auctions solely to risk premiums, the opportunity cost of capital or a fixed, per auction participation cost. Firms that purchase a derivative could be revealing some private information to other participants about the value of the derivative. In the context of the examples in section 1.3, it may be that some of the first round purchases are improving the supply of other TCCs to the market and earning a trading profit, where trader competition in subsequent auctions on related products removes this opportunity.

⁵²The responses to retailer offers are omitted for sample size reasons. The 95% confidence interval covers -40% to 20% for the second round.

5.2 Price updating across auctions with different vintages

We have seen that generating firms and trading firms earn systematic trading profits, and that after they buy a derivative its price appreciates, diminishing the potential to earn profits on that exact same product in the next auction. This poses a puzzle – how have firms been able to systematically earn profits year after year? This section describes the persistence of trading profits on products with the same (i, j) location-pairs over different vintages. The motivation is to identify whether profitable firms have some persistent forecasting advantage tied to a single location in the network.

To examine the persistence of profits, each awarded contract is classified into a quartile based off of the profitability of the underlying (i, j) derivative to the contract for the previous (T_1, T_2) vintage.⁵³ Within this vintage of contract, each contract location-pair is then classified as being in one of the following five categories:

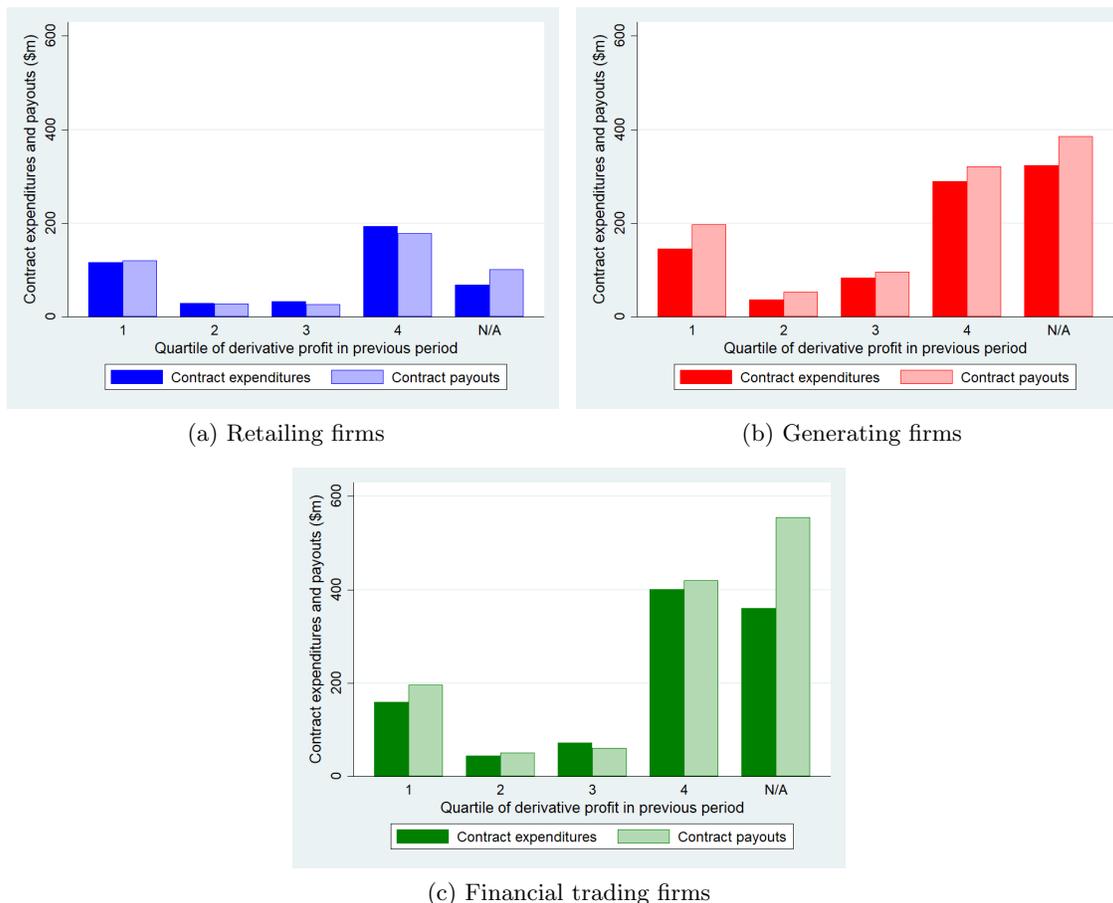
- 1-4: Quartiles 1-4 in derivative profit in previous auction of the (i, j) derivatives that were purchased
- N/A: Zero quantity of the (i, j) location-pair derivative was purchased by any firm in the previous auction

The N/A category is substantial. We observed in the three node auction example in section 1.3 that not every i, j derivative has non-zero TCC allocations at auction. Unsurprisingly, given the 100,000s of potential specifications available, many of them are not purchased each auction. Figure 11 displays the contract costs and payouts by firm type for contracts in each of the five categories as defined by the outcome for the contract in the previous vintage. A large portion of TCC purchases by generating and trading firms were for (i, j) derivatives that were not purchased by any firm in the previous vintage. These previously untraded, or low liquidity, contracts make up 88% of financial trader profits.

Statistical support for the figures is found in table B4, where tests for prices being equal to expected payouts is rejected for lagged quartile 1 and the not previously traded contract groups. Therefore, if a firm takes a profitable position on a derivative between locations i and j , the

⁵³For example, all contracts covering November 1 2008 to April 30 2009 have their derivative profits from November 1 2007 to April 30 2008 calculated $(r_{i,j,t-1} - p_{i,j,t-1})$. For six and twelve month derivatives, the $t - 1$ values are for the same i, j pair for the derivative beginning 12 months earlier. For one month auctions, this is for the derivative beginning one month earlier.

Figure 11: Contract costs and payouts by past derivative performance



Figures plot the total contract costs and contract payouts for derivatives purchased by the specified firm group. Sample is restricted to derivatives traded since 2007, where derivative prices is available. Quartile groupings are determined by the quartile ranking of issued contracts with a common time horizon, for the per unit of derivative profits in the previous vintage, as defined in the paper body.

opportunity to profit on the i, j product in the next auction disappears. Similarly, figure B3 shows that after payouts are revealed, derivatives that are poor performing attract less bidders and lower prices in the subsequent auction for that derivative, whereas better performing derivatives attract more bidders and higher prices.⁵⁴

⁵⁴The mechanism behind these patterns is not definitive, it may be that when an asset performs poorly it is because of more bidders or higher prices. Olmstead (2018) observes in Ontario that financial transmission rights are underpriced when less bidders participate and more likely to be actuarially fairly priced when there are many bidders.

To summarize, financial traders are compensated for being the first firm to purchase a contract on a derivative that was not purchased in previous auctions. After a contract has been purchased and revealed to be profitable, the market learns the given product was underpriced and accurately adjusts their payout expectations in the subsequent period. The market is able to close the profit margins on products that firms recently purchased, but profit margins exist when a firm is the first to buy a product that was not purchased in the previous auction.

To identify profitable opportunities, firms may need to possess a forecasting technology for illiquid derivatives that did not have a contract issued in the prior period. Therefore, a regulator's objective of designing the auction to facilitate price discovery might be working, where markets respond to some form of information revealed by some firms purchasing a contract. However, a transfer cost of these profits is imposed on ratepayers, narrowing the policy questions to whether the benefits are worth it for ratepayers, or whether improved competition could achieve a similar outcome but with a lower transfer.

5.3 What barriers prevent competition from eroding systematic trading profits?

Traders must consistently identify a new set of mispriced derivative products each auction if they are to systematically profit. Although this article does not uncover how these mispriced derivatives are identified, this section relates the empirical findings to anecdotes regarding financial trader operations.

Arce (2013) describes the existence of both sunk and ongoing resources being devoted to active trading in TCCs. The mechanism to set electricity prices and TCC auctions are nonlinear, constrained optimization problems. Therefore, a microfounded forecasting strategy requires an understanding of the physics behind electricity networks. Some traders build proprietary electricity network models that can generate prices from different inputs of demand, supply and transmission capacities. The forecast inputs are consistently updated as private information is acquired or public data is released from past electricity markets and TCC markets, along with planned transmission and generator outages. Price forecasts are then used to form a bidding strategy. Arce claims that TCC traders must be competent in each of physics, computing and economics, and also require a

high tolerance for tedium. Given the difficulties to find, train and retain traders, it usually takes between 12 and 24 months for an FTR desk to successfully consolidate.

Given the costs involved in developing and maintaining a proprietary black box to trade in TCCs, trading firms must earn some trading profits to continue participating. These costs could be representing a barrier to TCC profits eroding. Furthermore, Creswell and Gebeloff (August 14, 2014) describe an additional difficulty of being able to enter the market, with the most profitable trading firm in New York, DC Energy, requiring non-disclosure and non-compete agreements with their employees. DC Energy has demonstrated their preparedness to enforce these agreements.⁵⁵

Figure B5 reports estimates of systematic trading profits at a firm level. Given that the number of issued contracts is more sparse at a firm level, the estimates have low power but 4 major firms are detected as earning systematic trading profits in this market. The firms are financial traders Boston Energy Trading and Marketing, DC Energy, DC Energy New York and DC Energy New England, along with two generator owners, Hydro Quebec and EDF Trading North America. Together, these firms account for 17% of contract expenditures and 50% of contract profits (figure B4). Given 117 firms have ever participated in the market, this concentration of profits suggests these are firms more adept to identifying profitable opportunities. Further investigation into these firms reveals that Hydro Quebec almost exclusively purchased contracts with a point of injection at the import/export node between Quebec and New York, whereas EDF⁵⁶ and the profitable trading firms buy products across all price nodes in the network, consistent with the predictions and the earlier findings that traders profit from buying the products physical firms do not purchase.⁵⁷

Taking the empirical results and the trader anecdotes together, it appears that profitable financial traders must have some technology to identify profitable trading opportunities among products

⁵⁵Refer to Creswell and Gebeloff (August 14, 2014) for a description of a lawsuit filed by DC Energy against an ex-employee that moved to a company that began to trade in TCCs soon after.

⁵⁶EDF is classified as a generator owner, even though their trading subsidiary is listed as the TCC trader. They earn substantially less profits than DC Energy and Boston Energy Trading and Marketing, but bid on products that are not tied to their power plant ownership, therefore they might be better considered a financial trader.

⁵⁷Hydro Quebec provides an interesting case study as the only firm with systematic profits in the TCC markets that limits their participation to a single local node. In *DC Energy, LLC v. HQ Energy Services*, DC Energy (DC) took a counterposition to Hydro Quebec (HQ). DC unsuccessfully accused HQ of manipulating prices at the Quebec export node, where the day-ahead electricity price frequently dropped below long term averages to \$0/MWh for periods when HQ held TCCs with payouts decreasing in the Quebec price (Cramton, 2007). An observer might speculate that the otherwise information-rich DC Energy and their subsidiaries (accounting for \$212m of the \$860m TCC profits observed in this dataset) took a position based on a model of TCC payout forecasts, where it might not have taken the position if it had known that HQ, endowed with an operational information advantage, would take the opposite position.

that were not purchased in previous auctions. Once they act on these opportunities, there are enough participants in the market who update their expectations for the payout of that derivative to erode any further profits that can be made on that product. Therefore, to continue to earn trading profits these firms must update their models of future electricity prices to uncover new opportunities for trading profits without other firms replicating their trading strategy.

6 Policy discussion: Who benefits from ratepayer-funded auctions for transmission congestion contracts?

Three firm groups participate in TCC auctions, with electricity ratepayers the fourth, non-participating stakeholder group. Retailers were shown to have purchased predominantly zonal products in large quantities and due to regulatory incentives might prefer to abstain from taking speculative positions on contracts that are not linked to their procurement strategies. On average, retailers pay actuarially fair prices for their derivatives.

Generators were shown to mostly purchase derivatives unrelated to their physical operations. Unlike retailers, generators purchase both zonal and nodal contracts that are offered at auction. On average they earn systematic profits from their trading positions. Therefore, generators may benefit from some of the derivatives that allow them to sell electricity to different locational prices to their own, but they also receive benefits simply by profiting from their positions.

Financial traders have no physical interests that can be enhanced by holding a TCC. Like generators, traders purchase both zonal and nodal contracts that are offered at auction and do not always purchase large quantities. Traders have no reason to participate in these markets if they are unable to earn trading profits, which I have shown they are able to do systematically. Under the TCC auction mechanism, trader purchases on products with low demand can increase the quantity of other TCC products available for others to buy and potentially improve price signals.

Although the physical and financial firms appear to benefit from the existence of TCC auctions, transmission ratepayers effectively fund the trading profits earned by generating and trading firms. TCC auctions allocate the merchandising surplus market operators receive from transmission congestion in the spot market to TCC holders, with the auction proceeds used to lower ratepayer bills. Concerns from U.S. Congress and market monitors have focused on the distributional aspect of the

auctions that appear to be transferring wealth away from ratepayers to TCC holders. Therefore, these policymakers would want to see a more efficient electricity market and consumer benefits attached to the systematic profits being earned.

It is difficult to claim that transmission ratepayers benefit or lose out from trader participation in the TCC auctions without a formal welfare analysis (which is not able to be performed without generator production and cost data). However, we have developed a picture of where any benefits to ratepayers from the current auction regime must come from. Ratepayers would need to have their expenses fall by approximately \$50m per year to prefer funding TCC auctions that deliver systematic profits to traders. The results present a case that traders buy many of the products that physical firms do not purchase, and improve price signals on previously illiquid products. These trader actions must result in procurement cost savings (be it from more efficient generation/transmission planning, or retailer contract positions) for ratepayers to benefit from the auction framework. When evaluating the auction regime from a broader welfare perspective, planners might also consider the resources financial traders use when obtaining their forecasts and trading strategies, and the administrative costs of running the auctions.

The magnitude of the regulator's problem is substantial, with clear distributional consequences. TCC profits earned by financial trading firms totaled \$855m from 1999-2015 in New York, \$420m in California from 2012-2015 and \$904m in the PJM market from 2013-2015.⁵⁸ This study has shown that in New York, TCC profits are systematic and have not diminished over time. It is unclear that future entry of traders will occur to increase the auction revenues and consequently lower electricity customer bills. To this end, policy modifications have been suggested or implemented, each of which would likely reduce trading profits but may also restrict the benefits physical firms derive from TCC markets.

First, there is the option for market operators to disband the auctions and distribute the merchandising surplus it collects from transmission congestion in the short-term energy markets in another manner. Eliminating the auctions would of course eliminate derivative trading profits, the consequent transfers of wealth and any costly investment in information traders incur via their participation. However, as shown by the participation of retailers in New York's TCC market, the

⁵⁸New York, author calculation, California, see CAISO Department of Market Monitoring (2016) and PJM see PJM (2015) and various issues.

benefits to physical firms from having products available to source or sell electricity to different locations would be lost by disbanding the auctions, along with any benefits tied to the power flow and price guidance provided by the auctions.

Second, this proposal is extended by Bushnell and Wolak (2005) who propose directly allocating the merchandising surplus to retailers as a collection of derivatives.⁵⁹ If retailers hold a collection of TCCs, it may facilitate greater competition among suppliers – retailers that hold a TCC between their location and that of the supplier and enter a forward contract with a supplier have certainty regarding their procurement costs and pick from the cheapest option. The revenues they collect from their remaining TCC holdings could be used to lower the revenues they can recover from their retail customers. CAISO Department of Market Monitoring (2016, 2017) propose that market operators could still facilitate derivative markets for locational price swaps, but have them set up such that contracts are formed only when willing counterparties take opposite positions.⁶⁰

A third policy modification has been implemented in New Zealand. There, following a stakeholder process, a single TCC between two locations was made available, with the remaining merchandising surplus distributed via direct allocation (see Energy Market Services, 2012). Although this necessarily reduces the ability for firms to source or sell their electricity to different locations via these particular auctions, it could increase liquidity at these locations by concentrating participation into a smaller set of products and also remove the complexity of the auction.⁶¹ In New York, the set of 11 zone prices (55 TCC combinations) received greater expenditure on TCCs than the 100,000 TCC combinations available between price nodes, with retailers restricting their participation to zonal products. Further, zonal products were consistently purchased, priced actuarially fairly and were not subject to large TCC holder profits. It is left for further work to evaluate a proposal that restricts the set of products offered in New York to zonal products. Considerations include the lower participation costs from a simpler auction, the loss of product choice for firms

⁵⁹This position is also suggested in CAISO Department of Market Monitoring (2016). The article contends that transmission ratepayers effectively take counterpositions on TCCs so the auction should be updated such that contracts are only entered by willing counterparties.

⁶⁰Whether such derivative markets would be liquid and provide valuable price signals is uncertain. Black (1986) summarizes a large literature discussing why markets for some derivative products fail to exist, whereas markets for other derivative products are liquid. One feature of liquid derivative markets is that both physical and financial firms participate, with financial traders needing to collect some rents in order to participate.

⁶¹The potential benefits of reducing choice sets in a variety of settings are explored in Levin and Milgrom (2010). Black (1986) gives a summary regarding derivative product design and qualitative predictors for the success or failure of derivatives to be liquid. Derivative payout structures that are at centralized locations or at an index level tend to be more liquid.

to manage locational price differentials and the impact removing some profitable opportunities for financial traders will have on their participation.

7 Conclusion

To fund their participation in derivative markets, financial traders must earn trading profits. In markets for transmission congestion contracts, trader profits have attracted regulatory attention because TCCs are auctioned and TCC holder profits are effectively funded by transmission ratepayers. I have described, using simple models of TCC auctions, the potential for financial traders to improve auction outcomes by purchasing the derivative products retailers and generators do not purchase. I showed that retailers bid on a tiny proportion of products relative to financial traders. 88% of trader profits are earned by firms that are the first to purchase a previously illiquid product, but profitable opportunities are quickly competed away in subsequent auctions. This pattern has persisted for 16 years in the New York market, suggesting that there is a barrier to more trading firms being able to spot the initial profitable opportunity and in turn erode the trading profits earned in this market.

Regulators need to decide how to distribute the merchandising surplus collected by operators of formal wholesale electricity markets. These revenues accrue when transmission lines get congested, where consumers of electricity in importing regions pay more than the payments suppliers of electricity in exporting regions receive. Every formal electricity market in the United States distributes these revenues as transmission congestion contracts that are sold at auction. These contracts pay the holder future locational price differences in electricity prices and the auction revenues are used to lower transmission ratepayer bills. The merchandising surplus could be used for other purposes than to fund TCC holder payouts. The results of this article highlight the tradeoffs that regulators need to weigh up when considering the modifications to the distribution rule. The current rule of securitizing the merchandising surplus into TCCs via a network auction has resulted in financial traders earning large trading profits, but they earn them from performing actions policy makers often want to see from financial participants – buying a large variety of previously illiquid products (and providing price discovery). It is an open question as to how these actions benefits electricity retailers in this market, given that they bid on less than 1% of the products that financial traders

bid on.

If regulators wish to revise their policy to reduce large wealth transfers from electricity ratepayers to derivative holders, they could consider a direct allocation policy for the merchandising surplus from transmission congestion, or a restriction on the products offered at auction. The Californian Department of Market Monitoring has proposed that the merchandising surplus not be used to securitize TCCs and for a centralized price swap clearing pool be established (CAISO Department of Market Monitoring, 2017). Here, price swap contracts with the same payout structures of TCCs will only be created when willing counterparties take opposite positions. Under this proposal it is still possible for traders to profit and provide price discovery, but only by taking counterpositions to combinations of contracts that another firm (or collection of firms) willingly enters. It is left as further work to investigate whether modifications to the derivative product set offered at auction will improve economic outcomes. There may be opportunities to compare policies if ongoing reviews of these auctions result in policy changes. A pre-post study that can measure the realized physical costs from electricity generation may build upon the description of profit sources in this article to provide further insight into the physical efficiency impacts from the policy changes.

References

- ADAMSON, S., T. NOE, AND G. PARKER (2010): “Efficiency of financial transmission rights markets in centrally coordinated periodic auctions,” *Energy Economics*, 32(4), 771–778.
- ALDERETE, G. B. (2013): “FTRs and Revenue Adequacy,” in *Financial Transmission Rights*, pp. 253–270.
- ALSAC, O., J. BRIGHT, S. BRIGNONE, M. PRAIS, C. SILVA, B. STOTT, AND N. VEMPATI (2004): “The rights to fight price volatility,” *IEEE Power and Energy Magazine*, 2(4), 47–57.
- ARCE, J. (2013): “Trading FTRs: Real Life Challenges,” in *Financial Transmission Rights*, pp. 271–288. Springer.
- BAER, J. B., AND G. P. WOODRUFF (1929): *Commodity Exchanges*. New York and London: Harper & Brothers.
- BARTHOLOMEW, E. S., A. S. SIDDIQUI, C. MARNAY, AND S. S. OREN (2003): “The New York transmission congestion contract market: Is it truly working efficiently?,” *The Electricity Journal*, 16(9), 14–24.
- BIRGE, J. R., A. HORTAÇSU, I. MERCADAL, AND J. M. PAVLIN (2018): “Limits to Arbitrage in Electricity Markets: A case study of MISO,” *Energy Economics*, 75, 518–533.
- BLACK, D. G. (1986): *Success and failure of futures contracts: theory and empirical evidence*. Salomon Brothers Center for the Study of Financial Institutions, Graduate School of Business Administration, New York University.
- BOHN, R. E., M. C. CARAMANIS, AND F. C. SCHWEPPE (1984): “Optimal pricing in electrical networks over space and time,” *The Rand Journal of Economics*, pp. 360–376.
- BUSHNELL, J. (1999): “Transmission rights and market power,” *The Electricity Journal*, 12(8), 77–85.
- BUSHNELL, J. B., AND F. A. WOLAK (2005): “A Proposal for Pro-rata Congestion Revenue Right Allocation,” Discussion paper, Market Surveillance Committee of the California ISO.
- CAISO DEPARTMENT OF MARKET MONITORING (2016): “Shortcomings in the congestion revenue right auction design,” Discussion paper.
- (2017): “Market alternatives to the congestion revenue rights auction,” Discussion paper.
- CHIAPPORI, P.-A., AND B. SALANIE (2000): “Testing for asymmetric information in insurance markets,” *Journal of Political Economy*, 108(1), 56–78.
- CONSOLIDATED EDISON INC. (2015): “Con Edison 2015 Annual Report,” Discussion paper.
- CRAMTON, P. (2007): “Affidavit of Professor Peter Cramton on behalf of DC Energy, LLC,” *DC Energy, LLC v. HQ Energy Services (US) Inc.*, Federal Energy Regulatory Commission EL07-67-000.

- CRESWELL, J., AND R. GEBELOFF (August 14, 2014): “Traders Profit as Power Grid Is Overworked,” *The New York Times*.
- DENG, S.-J., S. OREN, AND S. MELIOPOULOS (2004): “The inherent inefficiency of the point-to-point congestion revenue right auction,” in *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*. IEEE.
- ENERGY MARKET SERVICES (2012): “An Introduction to the New Zealand FTR Market,” Discussion paper.
- HADSELL, L., AND H. A. SHAWKY (2009): “Efficiency and profit in the NYISO transmission congestion contract market,” *The Electricity Journal*, 22(9), 47–57.
- HENDRICKS, K., AND R. H. PORTER (1988): “An empirical study of an auction with asymmetric information,” *The American Economic Review*, pp. 865–883.
- HOGAN, W. (2002): “Financial transmission right formulations,” *Report, Center for Business and Government, John F. Kennedy School of Government, Harvard University, Cambridge, MA*.
- HOGAN, W. W. (1992): “Contract networks for electric power transmission,” *Journal of Regulatory Economics*, 4(3), 211–242.
- JHA, A. (2017): “Regulatory Induced Risk Aversion: Coal Procurement at US Power Plants,” Discussion paper, Working Paper.
- JHA, A., AND F. A. WOLAK (2015): “Testing for market efficiency with transaction costs: An application to financial trading in wholesale electricity markets,” Working Paper.
- JOSKOW, P. L., AND J. TIROLE (2000): “Transmission rights and market power on electric power networks,” *The Rand Journal of Economics*, pp. 450–487.
- LESLIE, G. W. (2018): “Essays in energy and environmental markets,” Ph.D. thesis, Stanford University.
- LEVIN, J., AND P. MILGROM (2010): “Online advertising: Heterogeneity and conflation in market design,” *The American Economic Review: Papers & Proceedings*, 100(2), 603–607.
- MISO (2015): “2014 State of the Market Report for the MISO Electricity Markets,” Discussion paper.
- NYISO (2005): “Transmission Services Manual,” Discussion paper.
- (2010): “Open Access Transmission Tariff,” Discussion paper.
- (2015): “Transmission Congestion Contracts Manual,” Discussion paper.
- (2016): “2016 Load and Capacity Data Gold Book,” Discussion paper.
- OLMSTEAD, D. E. H. (2018): “Ontario’s auction market for financial transmission rights: An analysis of its efficiency,” *The Energy Journal*, 39(1), 235–251.
- OREN, S. S. (2013): “Point to Point and Flow-Based Financial Transmission Rights: Revenue Adequacy and Performance Incentives,” in *Financial Transmission Rights*, pp. 77–94. Springer.

- PATTON, D., P. LEEVANSCHAICK, AND J. CHEN (2016): “2015 State of the Market Report for the New York ISO Markets,” Discussion paper, Potomac Economics.
- PJM (2015): “State of the Market Report for PJM,” Discussion paper.
- SARAVIA, C. (2003): “Speculative trading and market performance: the effect of arbitrageurs on efficiency and market power in the New York electricity market,” Center for the Study of Energy Markets Working Paper.
- TANGERAS, T. P., AND F. A. WOLAK (2017): “Testing for market efficiency with transactions costs: An application to convergence bidding in wholesale electricity markets,” *Department of Economics, Stanford University*.
- TOOLE, C. J. (2014): “An empirical analysis of the New York Independent System Operator’s Transmission Congestion Contract market: Speculator and hedger transaction characteristics, competition, and profit,” Master’s thesis, The Pennsylvania State University.
- UNITED STATES HOUSE OF REPRESENTATIVES SUBCOMMITTEE ON ENERGY (November 29, 2017): “Hearing on “Powering America: Examining the Role of Financial Trading in the Electricity Markets,”” *115th Congress*.
- WILSON, R. B. (1967): “Competitive bidding with asymmetric information,” *Management Science*, 13(11), 816–820.

Appendix A: The NYISO transmission congestion contract auction and data sources

A.1 Additional auction details

NYISO administers transactions in the New York wholesale electricity market. This appendix extends the description of electricity market payments and the TCC auction provided in section 1 with some additional auction nuances and references. Information on the operation of the New York wholesale electricity market and transmission congestion contract market is available in the market rules (NYISO, 2015). A less technical, yet succinct overview can also be found in Toole (2014). For general explanations not specific to NYISO, Alsac, Bright, Brignone, Prais, Silva, Stott, and Vempati (2004) contains a terrific high level summary and Hogan (1992, 2002) a more complete explanation. For the specific New York auction, refer to NYISO (2010).

NYISO pays the generators their nodal price for what they inject and NYISO receives from loads (firms that buy wholesale electricity) the zonal price where they withdraw. This is the source of NYISO's merchandising surplus. Hogan (1992) shows that a set of financial transmission rights (FTRs) that is simultaneously feasible in the electric grid satisfies revenue adequacy. This means, that if the set of injections and withdrawals implied by a set of FTRs could feasibly occur given the transmission constraints of the electric grid, the merchandising surplus the market operator collects will be greater than or equal to the payouts the holders of the FTRs will collectively receive.

Each market has idiosyncratic auction rules for FTRs, with NYISO choosing to perform a simultaneous auction for every combination of price swaps in the network. NYISO collects price and quantity bids for locational price difference derivatives from auction participants. Then, it solves a non-linear optimization problem that:

- Sets auction shadow prices at each node to maximize the as-bid value of allocated TCCs.
 - Denote node shadow prices as $P_{Auction,i}$ for node i . Therefore, the equilibrium price of the i,j derivative in the notation of the article is $p_{i,j} = P_{Auction,j} - P_{Auction,i}$. This is the practical mechanism that enforces the transitivity of derivative prices.
 - Firms bid on a POI/POW pair. Bids to buy clear if it is greater than the difference in the node shadow prices.

- Constraint is that all implied injections and withdrawals from the derivatives are feasible in the physical transmission grid, with assumed transmission capacities for the problem released to participants prior to the auction. Further, for zonal bids, fixed injections and withdrawals at specific nodes are assumed, as described in section 19.9.7 of NYISO (2010).
 - A bid for a derivative that pays $20*(LMP_j - LMP_i)$ implies that 20 MW is injected at A and is withdrawn at B.
 - If all injections and withdrawals from a set of contracts that would be issued at a given set of auction shadow prices are not feasible given the assumed transmission capacities throughout the electric grid (derived from Kirchhoff's Law) then the prices and allocation are not a solution to the auctioneer's problem.
- All bids that are above the auction shadow prices are allocated the contract. So a bid for a derivative that pays $20*(LMP_j - LMP_i)$ will be awarded a contract if the bid price is greater than $p_{i,j} = P_{Auction,j} - P_{Auction,i}$.
 - Supplying this contract is the equivalent of bidding on the contract that has the opposite payment, $20*(LMP_i - LMP_j)$. Therefore, this auction is not simply a sale of goods, it can indirectly match other buyers and sellers.

The feasibility constraint of the TCC auction is modified to allow for contract periods to overlap. Before each auction which may cover 1, 6 or 12 months of derivative payments, the existing contracts and the proportion of the NYISO grid to be auctioned are known. Therefore, existing contracts are factored in to the implied injections and withdrawals from the contracts and the available transmission capacity is scaled to reflect the amount of transmission capacity being released. If 12 month contracts are auctioned off in 4 tranches, these scale factors will be 25%, 50%, 75% and then 100%.

Other practical matters include that transmission capacities are stochastic, they can vary with weather and can have unexpected outages. Therefore, when allocating FTRs, market operators must decide how much capacity to release - release too much and they might have a revenue shortfall, too little and they will maintain a surplus. Over a period of time, NYISO on average is revenue adequate (see Patton, LeeVanSchaick, and Chen, 2016, for a recent annual report covering the wholesale and

TCC markets, demonstrating the revenue adequacy of the TCC contract positions for the NYISO), with rules that transmission owners make up or receive any differences from merchandising surplus and FTR/TCC payouts.

The revenues from the TCC auctions are split amongst transmission owners. Transmission owners are regulated to earn a fixed rate of return, given that they form natural monopolies and it is inefficient to have them participate in markets as strategic players. The total revenues they are entitled to receive under the regulated return is calculated, then the TCC auction payments are taken away from that figure, with the remainder paid by transmission ratepayers via a cost-sharing formula outlined in NYISO (2005) and NYISO (2010). Therefore, in effect, the higher the TCC auction payments, the less ratepayers ultimately have to pay transmission owners.

A.2 Auction bids and results data

All data are available to the public at the NYISO TCC website, http://www.nyiso.com/public/markets_operations/market_data/tcc/index.jsp. This section details the construction of the derivative, contract datasets, which are closely related and have common information merged on to each other. Then the decoding of the anonymized auction data is described and compared to the derivative and contract data.

The main data used in this analysis is at a derivative level. The auction prices for these derivatives were collected from the “View nodal prices” link on the NYISO webpage, that lists the shadow prices generated from every auction. These files are appended, with a unit of observation constructed as being a derivative start date (T_1), end date (T_2), auction round (ar), POI (i), POW (j).

The derivative payouts are sourced from the “DAM marginal losses and congestion” link. The unit of observation is constructed as being month-of-sample, POI, POW and the relevant variable is the payout to an i, j derivative for the sample month. For each observation in the auction prices data, the payouts for the T_1, T_2 window are calculated and merged onto the dataset. Although data for derivative payouts is available since the introduction of the auctions in 1999, the auction prices are only available from late 2006, therefore the derivative dataset is restricted to derivatives issued at auction between 2006-2015.

A separate but related dataset, containing all contracts issued from 1999 is found at the “Sum-

mary Of Transmission Contracts” tab. Each observation contains start date (T_1), end date (T_2), POI (i), POW (j), firm (f), purchase price per MW (p) and quantity in MW (q). Again, payouts are merged on to each observation to give r .

The quantity variables on the contract dataset are transformed and merged onto the derivative dataset. These variables are derived from the derivative holdings of each firm in the data entering and following each auction. To generate the $q_{i,j,t,f}$, $Q_{i,j,t,f}^{POST}$ and $I_{i,j,t,f}^q$ variables in the derivative dataset, each variable is created for each firm, giving each derivative $3*117$ extra variables. The values for these variables are described in the body of the text.

For both the contract and derivative datasets, power plant ownership information from NYISO (2016) is attached to each node. For the contract dataset, an observation is marked if the contract holder holds a power plant at a node specified in the contract or in the same zone as a node in the contract.

To summarize, the derivative dataset contains prices and payouts for every derivative available at auction with a unit of observation being derivative start date (T_1), end date (T_2), auction round (ar), POI (i), POW (j). Information attached to each observation includes the price and payout of each derivative (scaled by the length of time the derivative payout covers), the $3*117$ variables relating to the holdings entering and leaving each auction for each firm, and indicator variables that list the type of nodes the contract contains (generating/non-generating). The contract dataset only contains issued contracts, with a unit of observation defined as the start date (T_1), end date (T_2), POI (i), POW (j) and firm f . The information contained in the contract dataset include the prices, payouts and quantities of derivatives issued, along with the type of nodes in the contract.

Each firm is also classified as a retailer, generator or trader. These classifications are listed in the appendix A.3.

Decoding the anonymized identities of locations and firms in NYISO’s Transmission Congestion Contract auction data

NYISO publicly releases all bids and offers entered into TCC auctions at <http://mis.nyiso.com/public/P-27list.htm>. Each auction is for a given start date and end date, with each bid a price/quantity pair. Unlike the contract dataset, each bid/offer has an anonymized identifier in place of the firm that places the bid/offer and anonymized identifiers in place of the POI and POW. These anonymized identifiers are stable across auctions.

To analyze auction behavior, a large set of the anonymized identifiers have been decoded by combining the information across the publicly available auction and contract datasets. The underlying principle behind the algorithm is to utilize the equilibrium contracts data that contains a market clearing price and quantities sold to each firm for a given location-pair to find bids and offers in the auction data that could generate the same quantity allocations for the given market clearing price.

1. For a given start date, end date and location-pair that has a non-zero equilibrium contract quantity, calculate the number of firms that bought this contract, sold this contract and store the sizes of these contracts and the clearing price p
2. In the auction data for that given start date and end date, take a given location-pair (these are anonymized identifiers)
 - (a) Calculate the clearing parcels and quantities that are implied by a clearing price of p
 - (b) Mark the pair as a potential match if the clearing parcels and quantities implied by this price match the equilibrium data
 - (c) If one of the bids/offers is equal to the market clearing price, it is a potential marginal bid. Allow the parcel quantity for that bid/offer to be less than the size of the bid/offer when determining if the location-pair is a potential match.
 - (d) Iterate to the next location pair in the auction data and continue until all location pairs have been marked as a potential match or otherwise.
 - (e) If there is only one potential match, assign the POI and POW listed in the equilibrium contract data to the anonymized identifiers.
3. Iterate to the next location pair in the equilibrium contract data and stop after all observed contract location pairs have had this procedure performed.

The algorithm is restricted to marginal bids. The algorithm matches 94 of the anonymized location identifiers to actual locations. Although less than half of the locations are decoded, they cover almost all of contract expenditures.

The next step of the algorithm recovers firm identities in the auction data.

1. For a given start date, end date and location-pair that both have matches to the anonymized location identifiers, calculate the number of firms that bought this contract, sold this contract and store the sizes of these contracts and the clearing price p
2. In the auction data, match the parcel sizes bid/offered that clear at p to clearing quantities observed.
3. If there the parcels are uniquely matched, assign the firm name to the anonymized firm identifier.

The algorithm matches 49 of the anonymized firm identifiers to the 117 firms that ever won a contract. Although less than half of the firms are decoded, they cover almost all of contract expenditures and profits made.

Table 3 compares the auction data to the awards data to examine the selection of the auction data. When defining a bid as a step function between a unique pair of locations (with a positive price a bid to buy between a POW and POI, a negative price an offer to sell between that POW and POI), the top panel of table 3 shows there are 489,409 bids in the data, 136,798 of which both the POW and POI location identifiers are decoded. Only on the decoded locations can the auction clearing prices and realized revenue information be mapped to each bid. Using this information, the value of the contracts generated between the decoded locations is \$2.7 billion, just less than the \$3.1 billion total observed in the awarded contract data covering the same period in the second panel of table 3. Comparing the top to the bottom panel gives insight into the selection of the auction data. The auction data only covers 43% ($38,370/89,124$) of the awarded contracts, but 90% of the expenditure and profit values.

The top panel of table 3 shows that the proportion of bids and offers on locations that were decoded that were successful in winning a contract was 28% ($38,370/136,768$). Given there were 489,409 bids in total and 89,124 contracts generated, this means that the remaining locations had a 14% $(89,124-38,370)/(489,409-136,798)*100$ of bids and offers that won a contract. Overall, the data selection for the auction data appears to cover higher value contracts with higher clearing rates. Given the algorithm to decode the auction data relies on matching award data to the auction data, it is vacuously true that locations that do not have awarded contracts can not be decoded and will result in the auction data covering the more liquid locations.

Coverage of the zone-indexed contracts is better than the nodal contracts, with a greater proportion of the retailer awards also seen in the auction data than the generators and retailers. Overall, the returns by contract class are similar in both datasets, but the returns by firm type differ in that retailer returns are higher using the auction data and generator returns are lower. The unknown firm types in the auction data are firms whose identities were not decoded. To reconcile the retailer and generator return differences, the collective return for the unknown firms of 2% could be explained by having the unknown category contain some of the losing retailers and winning generators.

The patterns in the awards data are broadly seen in the shorter sample of restricted locations observed in the auction data, summarized in table 3. The value of zone-indexed and nodal contracts are roughly equal, but the quantity of nodal contracts are much greater. Shorter duration contracts are more profitable, with traders realizing the greatest profits, followed by generators and then retailers. Retailers have far fewer bids (defined as a step function on a node pair) at 945, than the 55,000+ of the generators and traders, but have a higher conversion rate of bids to contracts of 50% compared with approximately 27%. For the purposes of the analysis in section 5.1 and the change in the number of bidders in section 5 across auctions, the derivatives included are less likely to contain illiquid, low price products.

A.3: Classification of participating firms into firm types

Each firm that participates in these auctions has been classified into three distinct categories based on their core business. Footnotes describe discretionary categorization decisions. First, any firm that purchases wholesale electricity in New York is classified as a retailer. Second, any firm that operates an electric generating facility that is not a retailer is classified as a generator. These two firm types are physical players in the electricity market and may have a hedging motive to participate in auctions for transmission congestion contracts. Third, all remaining firms that have no physical interests in the New York electricity market are classified as traders, who are assumed to speculate with the motive to make profits from trading. The motives of the participants are not definitive, physical players can speculate, and non-physical players may have positions to hedge.

All classifications were decided by the author, based on web searches of the firm, FERC listings of retailers and NYISO lists of generating plants and their ownership. In many cases, the listed

Table 3: Comparing implied awards from auction data with the award data: Costs and returns by contract class, 2006-2015

Sample	N bids	N decoded	N contracts	Expenditures	Profits	ROI
	Auction Data					
All	489409	136798	38370	\$ 2692.2 m	\$ 454.7 m	16.9 %
Zone-indexed	202957	21825	5955	\$ 1291.8 m	\$ 139.5 m	10.8 %
Nodal	286452	114973	32415	\$ 1400.4 m	\$ 315.1 m	22.5 %
1 month	269531	71441	21428	\$ 346.3 m	\$ 75.3 m	21.7 %
6 month	117091	35896	9939	\$ 1113.1 m	\$ 241.9 m	21.7 %
12 month	102786	29461	7003	\$ 1232.9 m	\$ 137.4 m	11.1 %
Retailers	1254	945	471	\$ 325.3 m	\$ 16 m	4.9 %
Generators	193880	56491	16309	\$ 859.9 m	\$ 162.9 m	18.9 %
Traders	218139	59951	15569	\$ 1030.3 m	\$ 266 m	25.8 %
Unknown	76136	19411	6021	\$ 476.7 m	\$ 9.8 m	2 %
	Awarded contracts data					
All	.	.	89124	\$ 3056.8 m	\$ 502.5 m	16.4 %
Zone-indexed	.	.	8959	\$ 1712.5 m	\$ 142.5 m	8.3 %
Nodal	.	.	80165	\$ 1344.3 m	\$ 360 m	26.8 %
1 month	.	.	51781	\$ 380.1 m	\$ 102 m	26.8 %
6 month	.	.	21775	\$ 1277.2 m	\$ 245.4 m	19.2 %
12 month	.	.	15568	\$ 1399.5 m	\$ 155.1 m	11.1 %
Retailers	.	.	1088	\$ 469.5 m	\$ -32.8 m	-7 %
Generators	.	.	39267	\$ 1077.7 m	\$ 207.9 m	19.3 %
Traders	.	.	48769	\$ 1509.6 m	\$ 327.5 m	21.7 %

A bid is a step function between a unique point of injection (POI), point of withdrawal (POW), start date, end date and firm (with a positive price a bid to buy between a POW and POI, a negative price an offer to sell between that POW and POI). All contract data from the auction dataset (the top panel) is for the location identifiers that were decoded and assume that bids less than or equal to the market clearing price are fully cleared. The bottom panel contains the full set of awarded contracts over the period. Contract expenditures sum the absolute value from the initial contract price across the class of contract defined by the row - buying and selling a \$1m contract are both listed as a \$1m contract. Profits are the sum of the profits for all contract positions. ROI is a modified return on investment for the asset class, equal to the total profits divided by the absolute value of contract expenditures, listed in the preceding two columns.

owner of the generator is a subsidiary or parent of a firm listed as the trading entity in the TCC data. In such cases, the classification rule applies to any and all businesses in the conglomerate, so a conglomerate will not have some subsidiaries listed across the different classifications of firms, they will all be contained in one classification.⁶²

Retailers: Allegheny Energy Supply Company, LLC; CECONY-LSE; Central Hudson Enter-

⁶²Toole (2014) classifies firms into “speculator”, “hedger” and “unknown” categories, analyzing the types of derivatives these groups are more likely to purchase. The main difference between my list and that of Toole is that generating firms tend to fall into the hedging category in Toole’s analysis.

prises Corp.; Central Hudson Gas & Electric Corp.; Con Edison Solutions, Inc.; Consolidated Edison Co. of New York, Inc.; Consolidated Edison Energy, Inc.; Constellation Energy Commodities Group, Inc.; Constellation Energy Services of New York, Inc.; Constellation NewEnergy, Inc.; Direct Energy Business Marketing, LLC; El Paso Merchant Energy, L.P.; Energy Services Providers, Inc.; Freeport Electric; Green Island Power Authority; Indeck-Corinth LP; Jamestown Board of Public Utilities; Long Island Power Authority; Macquarie Energy LLC; New York Municipal Power Agency; New York State Electric & Gas Corp.; Niagara Mohawk Power Corp.; Orange & Rockland Utilities, Inc.; Pepco Energy Services; PPL Utilities; Public Power, LLC; Public Service Electric & Gas Co.; Rochester Gas & Electric Corp.; Village of Rockville Centre.

Generators: AES Creative Resources, L.P.; American Electric Power Service Corp.; Bayonne Energy Center, LLC; Brookfield Energy Marketing LP; Bruce Power Inc.; Castleon Commodities Merchant Trading L.P.; Dynegy Marketing and Trade, LLC (DMT); Dynegy Power Marketing, LLC; EDF Trading North America, LLC; EDP Renewables North America LLC; Exelon Generation Company LLC; GDF Suez Energy Resources NA, Inc; GenOn Energy Management, LLC; Hess Corporation; HQ Energy Services (US);⁶³ Integrys Energy Services, Inc.; Mercuria Energy America, Inc; Mirant Americas Energy Marketing LP; New York Power Authority;⁶⁴ NextEra Energy Power Marketing, LLC; Noble Americas Energy Solutions LLC; Noble Americas Gas & Power Corp; NRG Power Marketing LLC; Ontario Power Generation Inc.; PSEG Energy Resource & Trade, LLC; Select Energy New York, Inc.; Select Energy, Inc.; Selkirk Cogen Partners, L.P.; Shell Energy North America (US), L.P.; Sithe Energy Marketing, L.P.; Talen Energy Marketing, LLC; TransCanada Power Marketing, Ltd.

Traders: 330 Fund I LP; AC Energy, LLC; Amber Power, LLC; Appian Way Energy Partners East, LLC; Aquila Energy Marketing Corp.; BJ Energy LLC; Black Oak Capital LLC; BNP Paribas Energy Trading GP; Boston Energy Trading and Marketing LLC; BP Energy Company; Cargill Power Markets, LLC; Centaurus Energy Master Fund, LP; Citadel Energy Products LLC; Citadel Energy Strategies LLC; Citigroup Energy Inc.; Credit Suisse Energy LLC; DB Energy Trading LLC; DC Energy LLC; DC Energy New England, LLC; DC Energy New York, LLC; DTE Energy Trading Inc; E.ON Global Commodities North America LLC; Emera Energy Services, Inc; Enron Power Marketing; ENTEGRA CAPITAL MANAGEMENT LP; Entergy-Koch Trading, LP; EPIC Merchant Energy L.P.; EPIC Merchant Energy NY LP; Franklin Power LLC; Galt Power Inc.; GRG Energy LLC; J Aron and Company; J. P. Morgan Ventures Energy Corporation; KFW Energy Trading, LLC; Lighthouse Energy Trading Co., Inc.; MAG Energy Solutions Inc.; Merchant Energy Group (MEGA); Merrill Lynch Capital Services, Inc.; Merrill Lynch Commodities, Inc.; Midwest Energy Trading East LLC; Morgan Stanley Capital Group, Inc.; Nalcor Energy Marketing Corporation; Northern States Power Company; Ocean Power LLC; Old Lane Commodities, LP; OPD Energy LLC; Orthogonal Energy, LLC; Petra Technical Consultant Group, LLC; PG&E Energy Trading; Powerex Corporation; Pythagoras Global Investors LP; Quark Power LLC; RAM Energy Products LLC; RBC Energy Services LP; Royal Bank of Canada; Saracen Energy East LP; Saracen Energy West LP; Saracen Energy, LP; Saracen Power LP; Sempra Energy Trading LLC; SESCO Enterprises LLC; SIG Energy, LLLP; Silverhill Ltd., GP for Power Fund LPs.; Solios Power LLC; Split Rock Energy LLC; TransAlta Energy Marketing (U.S.) Inc.; Twin Cities Power, LLC;

⁶³Hydro Quebec is a peculiar case that has been classified as a generator for two reasons. First, it can purchase electricity for consumption, with the retail operation outside the NYISO. Second, it is a major net exporter to the NYISO.

⁶⁴New York Power Authority (NYPA) is a publicly owned generator owner but does not have a standard profit maximization objective function. For the analysis of auction positions in this paper, NYPA's classification is irrelevant as they never purchased a TCC at auction, with their only positions existing from grandfathered TCCs.

TXU Energy Services; Viridian Energy NY, LLC; Vitol Inc.; Williams Power Company Inc.

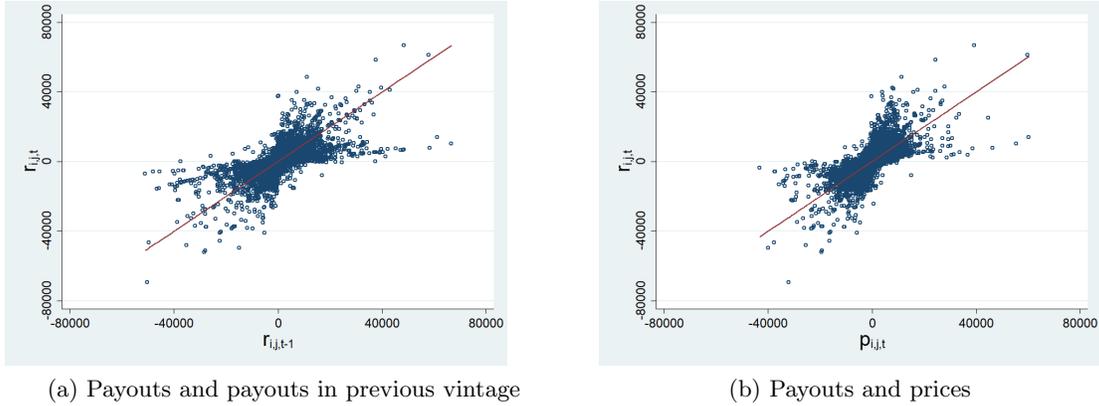
Appendix B: Additional figures and tables

Table B1: Expenditures on TCC contract positions

	Retailers					
	Zonal	Nodal	1 month	>1 month	Round 1	Round > 1
Total (\$m)	685	133	34	785	165	620
	Generators					
	Zonal	Nodal	1 month	>1 month	Round 1	Round > 1
Total (\$m)	901	844	225	1,520	249	1,271
	Traders					
	Zonal	Nodal	1 month	>1 month	Round 1	Round > 1
Total (\$m)	1,373	1,300	358	2,315	309	2,007

Contracts are classified into groups based on the zonal, nodal, 1- month or >1- month characteristics, and whether for the >1- month products they were sold in the first round or a later round. Given positions can be short or long, the absolute value of expenditures is the variable underlying the statistics in the table ($|q_{i,j,T_1,T_2,f} \cdot m(T_2,T_1) \cdot p_{i,j,T_1,T_2}|$). Sample restricted to the purchases in 2006-2015 where auction round information is available.

Figure B1: Derivative payouts, lagged payouts and prices



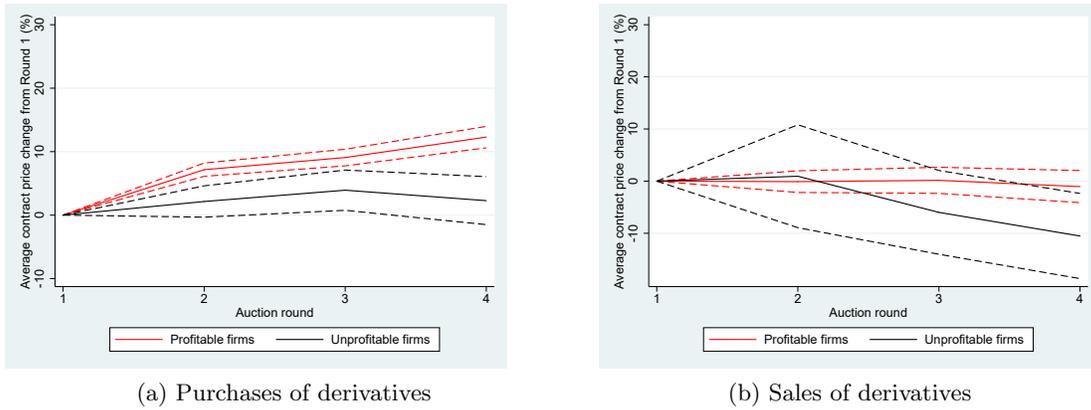
Figures plot average monthly payouts for a MW of each issued contract ($r_{i,j,t}$) against the payout for that contract in the prior vintage ($r_{i,j,t-1}$) and the price ($p_{i,j,t}$) for all contracts entered from 1999-2015. For six and twelve month derivatives, the $t - 1$ values are for the same i, j pair for the derivative beginning 12 months earlier. For one month auctions, this is for the derivative beginning one month earlier. The line in red plots the values where the y-axis is equal to the x-axis.

Table B2: Summary statistics of the location-pair-auction (i, j, t) derivatives studied

		Variable name			Description
		$p_{i,j,t}$	$r_{i,j,t}$		Price and payout of derivative
Mean		1686	1842		
Std. dev.		3215	4004		
		$q_{i,j,t,RET}$	$q_{i,j,t,GEN}$	$q_{i,j,t,TRA}$	Number of derivative units at auction
Mean		0.02	0.18	0.23	
Std. dev.		1.29	3.16	3.62	
		$I_{i,j,t,RET}^q$	$I_{i,j,t,GEN}^q$	$I_{i,j,t,TRA}^q$	Indicator = 1 if allocated contract at auction
Mean		0.001	0.031	0.034	
Std. dev.		0.031	0.174	0.182	
		$Q_{i,j,t,RET}^{POST}$	$Q_{i,j,t,GEN}^{POST}$	$Q_{i,j,t,TRA}^{POST}$	Size of open position after auction
Mean		0.85	2.06	2.03	
Std. dev.		20.46	24.74	18.61	

1,151,374 i, j, t observations in each cell. The absolute value of each variable is reported because the location direction a derivative enters the model is arbitrary. p and r , the derivative price and payout, are divided by the length of the contract. RET , GEN and TRA aggregate all allocations to retailing, generating and trading firms into a single firm grouping. Open position refers to derivatives held on an i, j derivative that has a payout window that covers T_1 .

Figure B2: Price updating following purchases or sales, by profitable and unprofitable firms



All series plot equation (11). The purchases chart compares two sets D of derivatives, those that were purchased by any firm at a positive price across, split by firms that earned positive and negative profits over the sample window. The sales chart is analogous to the purchases chart but for negatively priced products. Means and pointwise 95% confidence intervals plotted.

Sample sizes: (a) 2,638 and 342. (b) 984 and 85.

Table B3: Estimates of average monthly derivative payouts - extended model

	All $\overline{p_{i,j,t}} = 1686$	Nodal $\overline{p_{i,j,t}} = 1667$	Zonal $\overline{p_{i,j,t}} = 3821$
$\delta_{RET} [I_{RET,t}^q]$	7.00 (118.61)	126.72 (173.74)	-32.97 (146.65)
$\delta_{GEN} [I_{GEN,t}^q]$	90.03 (44.81)	96.38 (45.68)	-26.62 (118.73)
$\delta_{TRA} [I_{TRA,t}^q]$	162.02 (41.80)	176.66 (41.05)	-104.48 (139.72)
$\rho_{RET} [Q_{RET,t}^{POST}]$	-0.30 (0.17)	-0.28 (0.15)	-0.55 (0.30)
$\rho_{GEN} [Q_{GEN,t}^{POST}]$	0.19 (0.16)	-0.01 (0.14)	0.81 (0.44)
$\rho_{TRA} [Q_{TRA,t}^{POST}]$	-0.02 (0.30)	-0.08 (0.39)	0.17 (0.56)
α [Constant]	32.98 (91.87)	31.63 (90.87)	195.56 (261.93)
N	1,151,374	1,140,868	10,506
N_A	235	235	235

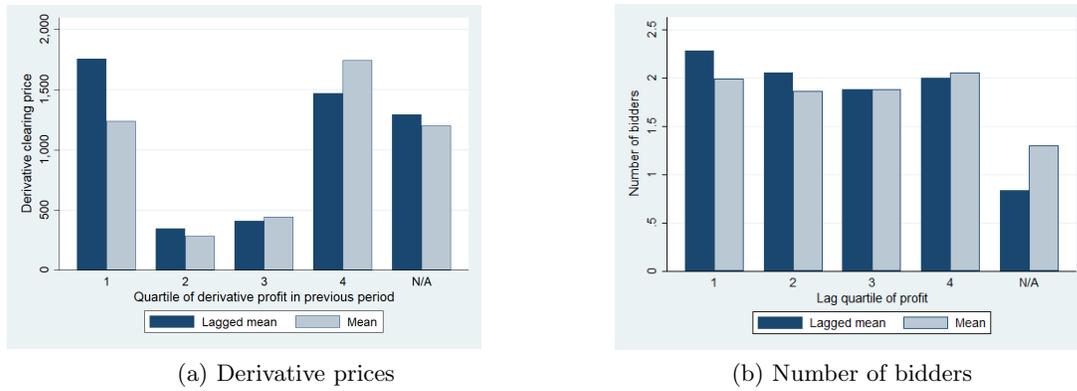
Estimates of equation (9), using ordinary least squares. Standard errors clustered at a vintage level T_1, T_2 reported in parentheses. All contract prices and payouts are divided by the number of months a contract covers. Summary statistics for the variables used in estimation are found in table B2. The direction of the price swap derivative is chosen such that $p_{i,j} > 0$. The estimated coefficients do not substantially differ when the direction is chosen arbitrarily. Comparing the estimates of table B3 to table 2, first note that the common δ_f coefficients estimated in both specifications are not sensitive to relaxing the restrictions on the ρ_f , size of position parameters. For nodal contracts, a test that $\rho_f = 0$ for any firm group with a test size of 5% fails to reject that the average marginal effect of increasing a firm's open position on derivative payouts is zero. There is no evidence to support claims that firms are performing downstream actions to influence derivative payouts in this market (as shown in a case study in Birge, Hortaçsu, Mercadal, and Pavlin (2018)).

Table B4: Coefficients for systematic trading profit tests, by previous auction performance

Prev. auction performance	Q1	Q2	Q3	Q4	N/A
$\alpha [q_{i,j,t,f} \cdot (T_2 - T_1) \cdot (r_{i,j,t} - p_{i,j,t})]$	15,263 (6,807)	3,386 (2,614)	-859 (3,612)	5,008 (16,035)	9,023 (4,656)
$\overline{q_{i,j,t,f} p_{i,j,t}}$	70,343	16,071	29,586	121,523	23,536
N	6002	6913	6422	7276	32016

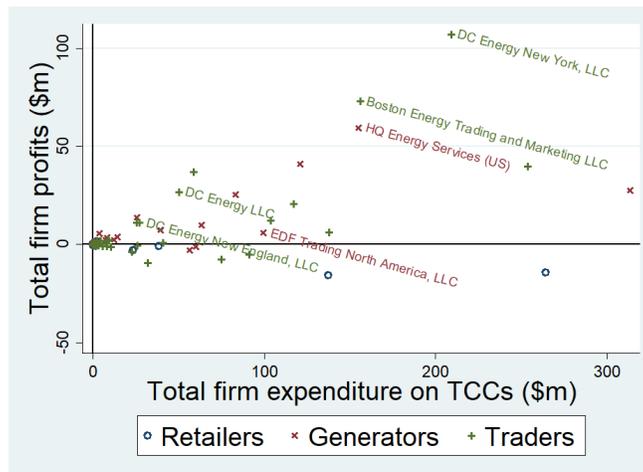
Standard errors are clustered at a vintage level (all contracts with the same T_1 and T_2) in parentheses. The null hypothesis for efficient markets is equivalent to $\alpha = 0$. The unit of observation is a unique contract, defined by location pair i, j , auction it was purchased in t (t defines the payout window T_1 and T_2), and the firm holder f . The sample contains all 1, 6 and 12 month contracts issued from 1999 to 2015. Q1-Q4 refer to contracts on i, j derivatives that were in the first to fourth quartiles of profits in the previous auction. N/A refers to contracts on i, j derivatives that were not issued in the previous auction.

Figure B3: Dynamics of prices, payouts and bidder numbers by past derivative performance



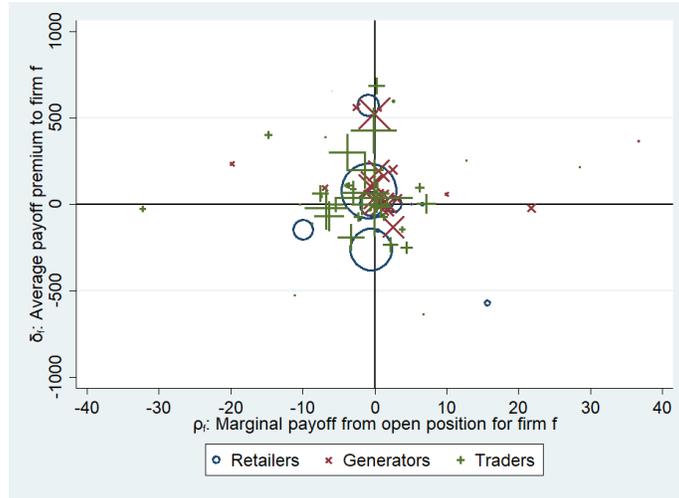
Samples are restricted to derivatives issued at a positive price since 2007. Quartile groupings are determined by the quartile ranking of issued contracts with a common time horizon, for the per unit of derivative profits in the previous vintage, as defined in the paper body. “N/A” denotes a derivative with a contract issued, but no contracts were issued for that POI/POW location pair in the previous vintage. Prices and payouts are scaled by the length of the contract. The number of bidders sample is restricted to the derivatives that were decoded by the algorithm discussed in Appendix A.2. Hypothesis tests with equality of means under the null are rejected at a 5% level of significance for all variables and groups, with the exception of lagged quartile 3, number of bidders.

Figure B4: Total contract costs and profits by firm

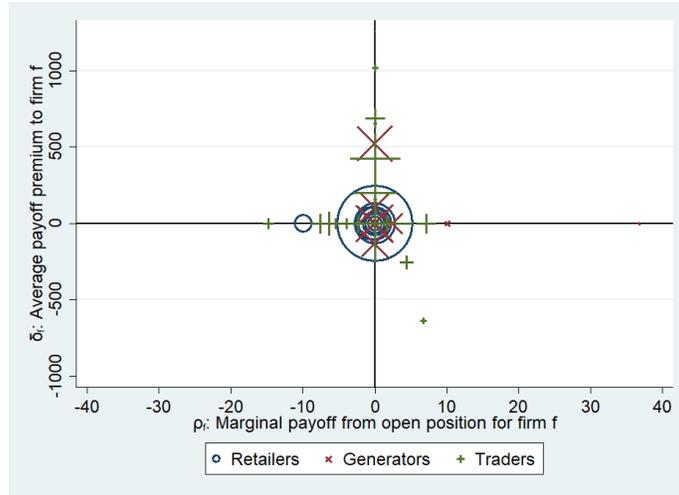


Aggregate firm values for all 1, 6 and 12 month awarded contracts since 2006. Total contract costs is the sum of the absolute value of each contract position taken by a firm.

Figure B5: Estimates of payout premiums by firm



(a) All point estimates



(b) All point estimates with 95% confidence intervals not covering zero.

Top figure plots firm level estimates of δ_f (the coefficient on the firm contract indicator variable) and ρ_f (the coefficient on the firm open position variable) as specified in equation (9) for the 117 firms ever observed to purchase TCCs over the sample window. Second figure replaces $\delta_f = 0$ or $\rho_f = 0$, if that hypothesis test at a 10% level of significance is not rejected. All markers are weighted by the sum of the total costs a firm incurred when purchasing TCCs over the sample window for the firms included on the chart. The six major firms observed to have statistically detectable values of δ_f greater than zero, are EDF Trading North America, Boston Energy Trading and Marketing, Hydro Quebec, DC Energy, DC Energy New York and DC Energy New England. These firms account for 17% of contract expenditures and 50% of contract profits. The figures show that a handful of extra firms have also earned positive aggregate profits, but the testing technique did not detect them as systematically earning trading profits.