Pivotal Suppliers and Market Power in Experimental Supply Function Competition

Jordi Brandts, Stanley S. Reynolds and Arthur Schram

January 2010

Abstract

In the process of regulatory reform in the electric power industry, the mitigation of market power is one of the basic problems regulators have to deal with. We use experimental data to study the sources of market power with supply function competition, akin to the competition in wholesale electricity markets. An acute form of market power may arise if a supplier is pivotal; that is, if the supplier's capacity is required in order to meet demand. To be able to isolate the impact of demand and capacity conditions on market power, our treatments vary the distribution of demand levels as well as the amount and symmetry of the allocation of production capacity between different firms. We relate our results to a descriptive power index and to the predictions of two alternative models: a supply function equilibrium (SFE) model and a multi-unit auction (MUA) model. We find that pivotal suppliers do indeed exercise their market power in the experiments. We also find that the SFE model better captures the behavior of pivotal suppliers than the MUA model.

Keywords : Market Power, Electric Power Markets, Pivotal Suppliers, Experiments

JEL Classification Codes : C92, D43, L11, L94

Authors

Jordi Brandts	Stanley S. Reynolds	Arthur Schram
Department of Business Economics Universitat Autònoma de Barcelona and Institut d'Anàlisi Econòmica, CSIC Campus UAB 08193 Bellaterra (Barcelona) Spain	Department of Economics McClelland Hall Eller College of Management University of Arizona Tucson, AZ 85721-0108 U.S.A.	CREED Department of Economics and Econometrics Roeterstraat 11 University of Amsterdam 1019 WB Amsterdam The Netherlands
phone +34-93-5814300 Jordi.Brandts@uab.es	Reynolds@eller.arizona.edu	phone +31-20-525.4293 Schram@uva.nl

1. Introduction

In the worldwide process of regulatory reform in the electricity industry, the possible existence of market power is one of the basic problems analysts and policy makers have to deal with. Field data document the existence of reduced competition due to market power in some electric power markets (Wolfram 1999; Borenstein et al. 2002). The severe welfare losses this may cause are a major concern that needs to be addressed to fully assess the success of the reforms. If non-competitive prices can easily persist in these markets, this creates the need to find measures to mitigate market power.

Among the features of markets that need to be taken into account in relation to market power is the presence of one or more *pivotal suppliers*. In a general sense, a producer can be considered to be pivotal if, without his capacity, the supply cannot serve the whole demand. The issue is not one of insufficient total capacity to serve the market demand, but one of particular producers controlling large enough parts of the capacity. We will refer to market power due to pivotal suppliers as *pivotal power*.

Concerns about pivotal power are the basis for some energy policy provisions. For instance, the U.S. Federal Energy Regulatory Commission (FERC) may block a generation company from charging market-based rates for energy if the company fails either of two screening tests for market power. One of the tests used by FERC is the pivotal supplier screen; a generation supplier is deemed pivotal, and therefore fails the test, if peak demand cannot be met in the relevant market without production from the supplier's capacity.[1]

In this paper we present results from laboratory experiments in which we study the effects of pivotal power. To our knowledge this is the first experimental study focusing on this specific issue. A potentially important distinction for policy-makers is the presence of a pivotal supplier vs. the supplier's incentive to exercise market power. We examine in the laboratory the extent to which pivotal suppliers actually exercise market power under varying market conditions. We study both the cases where pivotal power is evenly spread among producers and where it is concentrated in a subset of producers. The first case corresponds more to a situation of tight market capacity and the second more to a case in which particular firms may have strong influence on market outcomes, even though market capacity is large.

---

[1] The following quote is from FERC Order No. 697 (2007), pp. 18-19: "The second screen is the pivotal supplier screen, which evaluates the potential of a seller to exercise market power based on uncommitted capacity at the time of the balancing authority area's annual peak demand. This screen focuses on the seller's ability to exercise market power unilaterally. It examines whether the market demand can be met absent the seller during peak times. A seller is pivotal if demand cannot be met without some contribution of supply by the seller or its affiliates."

To study more closely the circumstances under which pivotal power may matter, we also analyze the impact on market power of a variation in the extent of demand uncertainty.

The use of laboratory experiments makes it possible to implement the desired variations in capacity distributions with a high degree of control, to study their effects under conditions that are strongly ceteris paribus. This control makes the experimental method a useful tool for studying electric power markets (Rassenti et al. 2002; Staropoli and Jullien 2006). In addition the possibility of replication allows for a very systematic study of the relevant issues. (See Falk and Heckman 2009 for a recent methodological discussion of laboratory experiments).

Previous laboratory experiments show that market power is easily exerted in environments that mirror the wholesale electricity market. Moreover, experiments have been useful for studying how certain market features can increase or limit market power; demand side bidding (Rassenti, Smith and Wilson 2003) and forward trading (Brandts, Pezanis-Christou and Schram 2008) have been shown to enhance competition.

Outside of the laboratory the notion of pivotal power has been investigated before, both with descriptive measures of suppliers' positions in the market and with theoretical models. The *Residual Supply Index* (hereafter, RSI) provides a measure of the degree of pivotal power based on fundamental economic intuition. The RSI measures the aggregate capacity of all suppliers except the largest as a fraction of total demand. The largest supplier is pivotal when this index is less than one and the lower the RSI the higher pivotal power. Field data suggest that the RSI is a useful indicator. The higher the index' value –i.e., the lower the weight of the largest supplier–, the lower were price-cost margins in the summer peak hours of the year 2000 in the California wholesale market (Rahimi and Sheffrin 2003). Wolak (2009) develops a measure of a supplier's ability to exercise unilateral market power in each half-hour period in his study of the New Zealand wholesale electricity market,. He notes that this ability to exercise market power is strengthened the greater the probability that the supplier is pivotal during the period.[2] Wolak finds a positive correlation between the average half-hourly firm-level ability to exercise unilateral market power and half-hourly market prices.

More formal analyses of pivotal power have been based on either a multi-unit auction model (Anwar 1999; Fabra et al. 2006), hereafter MUA, or the supply function equilibrium model (Klemperer and Meyer 1989), hereafter SFE. Both are models of one-shot strategic interaction. The MUA is a discrete unit model in which each supplier submits price offers for

---

[2] Wolak discusses pivotal suppliers and their significance in the New Zealand wholesale market in Section 3.4 of part 2 of his report. He defines a pivotal supplier as follows: "A supplier that faces a residual demand curve that is positive for all possible positive prices is said to be a pivotal because some of its supply is necessary to serve the market demand regardless of the offer price." [Wolak (2009), p. 115]

units of capacity under their control. Fabra, von der Fehr and Harbord (2006) utilize the MUA to compare uniform and discriminatory auction formats under distinct environments concerning costs, capacity, and demand characteristics. Our focus in the present paper is on uniform price auctions. For uniform price auctions, the MUA predicts Bertrand-like outcomes with price equal to marginal cost for cases with no pivotal firms. If pivotal firms are present then the MUA predicts equilibrium prices above marginal cost, with the extent of price-cost markups and the character of equilibrium prices depending on the nature of demand uncertainty and the allocation of capacity across firms.

The SFE assumes a completely divisible good and has been used to study a variety of issues related to electric power markets (Green 1999; Newbery 1998; Baldick et al. 2004; Bolle 2001). In the standard model each seller submits a supply function that specifies the quantities supplied at different prices (Klemperer and Meyer 1989). At the time the sellers submit their supply functions, demand is typically uncertain. The supply functions are aggregated and intersected with realized demand to obtain a uniform market price. If the range of demand variation is bounded then there is a continuum of equilibria. The equilibrium price at the upper bound of demand realizations ranges between the competitive price and the Cournot price. The assumption of divisible output and the use of supply function strategies essentially expands the set of equilibria relative to that of the MUA. If there are no pivotal suppliers then the MUA equilibrium involves marginal cost pricing; the SFE model has that equilibrium as well as additional equilibria with positive price-cost markups. If pivotal suppliers are present then the SFE model may yield equilibria with prices below those of a MUA equilibrium. Neither model of pivotal power has to date been applied to either field or laboratory data. In our data analysis we will use the input of both the descriptive index and the two theoretical models. As will be seen, the RSI proposes certain shifts in prices in response to changes in total capacity and its distribution, but is silent about what price levels should be expected. The two theoretical models allow us to go beyond that. The MUA prescribes unique pure strategy equilibrium prices for four of the five environments we consider and mixed strategy equilibria for the fifth. The SFE yields pure strategy equilibrium price sets for each parameter constellation that we study. In each case for which the MUA has a pure strategy equilibrium, this equilibrium is also an equilibrium for the SFE.

Our results show that observed market prices change with capacity levels and their distribution. The existence of overall excess capacity is not enough to guarantee competitive prices. The way in which they change is intuitive and consistent with the qualitative

predictions of the RSI. In contrast to the effects of capacities, the variations in the demand range have smaller effects on observed market prices. Overall, the SFE appears to organize the data better than the MUA, though there is a tendency for observed prices to move towards the competitive MUA prediction in treatments without pivotal power. The data are not consistent with MUA equilibrium predictions for treatments with market power; for these treatments, the SFE model is better at organizing the data. In particular, the SFE prescribes that market power caused by a symmetric reduction in capacity has a stronger effect on prices than market power caused by an asymmetric distribution of overall high capacity.

The remainder of this paper is organized as follows. In the following section, we present our experimental design and procedures. Specific theoretical predictions are provided in section 3. The results follow in section 4. Section 5 concludes and discusses the implications of our findings.

2. Design and Procedures

In the experiment there are 25 number of rounds each consisting of five periods. The demand is simulated using a simple box-design (Davis and Holt, 1993). In each of the five periods t in round r, a perfectly inelastic demand d_t^r is randomly chosen from the set $d_t^r \in \{d^{min}, \ldots, d^{max}\}$ with equal probabilities for each element in the set. In all of our treatments, $d^{max}=35$. We define $l \equiv d^{min}/d^{max}$ as the *load ratio,* which is our first treatment variable. All sessions have either $l=4/7$ (i.e., $d^{min}=20$) or $l=6/7$ ($d^{min}=30$).[3] There is a price cap given by $p^{max}=25$, i.e. no units can be traded above this price.

On the supply side there are four firms in each market. Each subject represents one firm. Each firm j offers a discrete number of units in round r, which will apply to each period t in r. Any units sold are produced at constant marginal costs $c=5$.[4] Individual supply is limited by an exogenously enforced maximum capacity s_j^{max}. These determine industry capacity, which is given by $S^{max} \equiv \sum_{j=1}^{4} s_j^{max}$.

Our second treatment variable is this industry capacity. This is given by either $S^{max}=48$ or $S^{max}=36$. Note that in both cases $d^{max}<S^{max}$, i.e., in all cases industry capacity suffices to satisfy the maximum demand. Our third treatment variable pertains to the $S^{max}=48$ case. We distinguish between the case where this capacity is distributed evenly across firms

---

[3] The load ratio affects the theoretically predicted outcomes (cf. section 3).
[4] Given our focus on pivotal power, the assumption of constant marginal costs is not restrictive.

$(s_j^{max} = 12, j = 1,..,4)$ and the case where there is asymmetric capacity $(s_j^{max} = 5, j = 1,2, s_j^{max} = 19, j = 3,4)$. For the $S^{max}=36$ treatment we only consider the symmetric case $(s_j^{max} = 9, j = 1,..,4)$.

Firm j offers units for sale in round r by bidding a discrete 'supply function', s_j^r. This is a vector of up to s_j^{max} supply prices, p_{jk}^{sr}, ordered from low to high at which firm j is willing to sell units: $s_j^r = (p_{j1}^{sr}, p_{j2}^{sr},..., p_{js_j^{max}}^{sr})$, with $p_{jk}^{sr} \geq p_{jk-1}^{sr}, k \geq 2$.[5] Subjects can offer fewer than s_j^{max} units by not entering prices for them. Equivalently, they can offer $m < s_j^{max}$ units by setting $p_{jk}^{sr} = 26, k = m+1,..., s_j^{max}$.[6] The individual supply functions are combined and supply prices are ordered from low to high to obtain the market supply function for round r, denoted by $s^r = (s^r(1),...., s^r(S^{max}))$, which is a vector of the S^{max} submitted supply prices (if necessary, supplemented with infinite prices for units not supplied), ordered from low to high. Finally, a uniform transaction price, p_t^r, is determined in each period t of r by comparing d_t^r to s^r: $p_t^r \equiv \min\{s^r(d_t^r), p^{max}\}$. Note that if $s^r(d_t^r) > p^{max}$, then supply cannot satisfy demand at a price below p^{max}, and $k < d_t^r$ units are sold (where k is uniquely determined by $s^r(k) \leq 25$ and $s^r(k+1) > 25$).

Finally, the payoffs of firm j in round r, π_j^r, are determined by the uniform prices in each of the five periods of a round and the marginal costs:

$$\pi_j^r = \sum_{t=1}^5 (p_t^r - c) q_{tj}^r, \quad j = 1,..., 4$$

where q_{tj}^r denotes the number of units sold by j in period t of round r.

In the experiment, subjects submit supply functions by entering a price for each possible unit in a table. To ease the task, the software fills gaps between units priced. For example, if a subject enters a price of 5 for unit 1 and then 7 for unit 5, then units 2-4 are automatically priced at 7, though the subject can subsequently change them. In addition, the software does not allow decreasing prices across units. The subject is free to withhold units from the market by leaving them unpriced, as long as all subsequent units remain unpriced as well. No supply price is submitted until the subject finalizes and confirms the complete set. There is no time limit for submission of the supply functions.

---

[5] The "s" superscript on the price variable indicates that the variable is a price offer made by a seller.
[6] Units offered at a price above the price cap, $p^{max}=25$ will not be sold.

After all four subjects have submitted a supply function, they are aggregated and the result is confronted with 5 subsequent demand realizations – the 5 periods of a round - yielding 5 prices. Each realization appears on the subjects' monitors for 5 seconds. After the 5 periods, the subject can page back and forth between the periods until satisfied. After everyone has indicated that they are ready the next round commences.

The results of a period appear on the screen graphically and in numbers. Figure 1 shows an example of the graph a subject could see – the text is in Dutch. This was a treatment where each subject had a maximum of 9 units to supply. The graph shows the results for period 3 (shown at the bottom on the right), where $d_3^t = 30$ and $p_t^r = 13$ were realized (bottom left) and this subject sold 7 units (bottom center). The graph shows the realized demand function (including the price cap), the price (light colored line) and the 9 supply prices submitted by this subject for this round (rising from 6 to 22). The 7 units sold by the subject are shown in dark gray, the 2 that were not sold are light gray.[7] Notice that this subject was willing to sell her 8[th] unit at price 13 as well, but other subjects had also submitted units at this price. In this case the computer randomly appoints subjects to slots on the aggregate supply function.[8]

Figure 1: Screenshot of Period Results
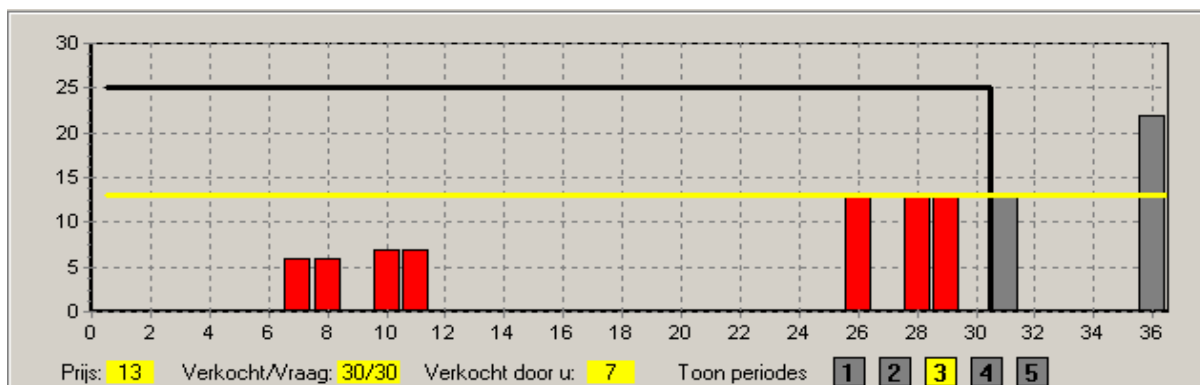
Notes. Translation from Dutch: Prijs = Price; Verkocht/Vraag = Sold/Demand; Verkocht door u = Sold by you; Toon periodes = Show periods.

The three treatment variables presented above (load ratio, industry capacity and symmetry of capacity) are varied between subjects. Table 1 presents an overview of these treatments and the number of markets we ran for each of them. In addition, it gives the average subject earnings for each treatment. Note that this simple design allows us to investigate the role of

---

[7] Subjects did not receive information about the offer price of units sold by other suppliers; we considered this to be closest to what happens in the field.
[8] We use pro-rata on the margin rationing in the experiments. This scheme for rationing supply is commonly used in wholesale electricity auctions. The way in which excess supply is rationed may have an impact on bidding; see Kremer and Nyborg (2004) for a theoretical analysis.

pivotal suppliers in a straightforward manner. First, in treatments *hsl* and *hsh* no single firm is pivotal because any three firms can cover maximum demand (d^{max}=35). Second, *lsl* and *lsh* deal with the case where every firm is pivotal for at least some possible demand quantities: together the four firms have sufficient capacity for maximum demand, but three firms can only supply 27 units. Third, *hah* covers the situation where two of the four firms (i.e., firms 3 and 4) are pivotal. Any combination of three firms including these two suffices for maximum demand but a combination of firms 1 and 2 with either 3 or 4 can only supply 29 units. Finally, note the large variation in earnings. This is a first indication that market power matters.

Table 1: Treatment Overview

		Low Capacity: $S^{max}=36$	High Capacity: $S^{max}=48$	
		Symmetric: $s_j^{\max} = 9,\ j = 1,..,4$	Symmetric: $s_j^{\max} = 12,\ j = 1,..,4$	Asymmetric: $s_j^{max} = 5,\ j = 1,2$ $s_j^{max} = 19,\ j = 3,4$
Load Ratio	**Low**: *l*=4/7	*lsl*; n=6; €45.14	*hsl*; n=6; €15.89	--
	High: *l*=6/7	*lsh*; n=5; €67.13	*hsh*; n=5; €27.28	*hah*; n=6; €43.60

Notes. The entries in the cell show the acronym, *xyz* (*x*=capacity, *y*=symmetry and *z*=load ratio); the number of markets we have data for, n; and the average earnings in euro (€).

The experiment was conducted in seven sessions at the CREED laboratory for experimental economics of the University of Amsterdam. 132 subjects were recruited by public advertisement on campus and were mostly undergraduate students in economics, business and law. They were allowed to participate in only one experimental session. Each session lasted for about 2-3 hours. Earnings in the experiment were denoted in experimental francs. We used an exchange rate of 250 francs to 1 euro. All subjects received a starting capital of 1250 francs, which was part of their earnings. There was no show-up fee. Subjects earned between €12.40 and €112.20 with an average of €42.71.

At the outset of each session, subjects were randomly allocated to the laboratory terminals and were asked to read the instructions displayed on their screens.[9] Then they were introduced to the computer software and given five trial rounds to practice with the software's features. Subjects were told that during these trial rounds other subjects' decisions would be simulated by the computer, which was programmed to make random decisions, and

---

[9] A transcript of the instructions (translated from Dutch) is included in Appendix 1.

that gains or losses made during those rounds would not count for their final earnings from participation. Once the five trial rounds were over, the pool of subjects was divided into independent groups (markets) of 4 subjects.

Each session then consisted of 25 repetitions (rounds), each round taking approximately 3-4 minutes to be completed. In each session only one treatment was run with 5 or 6 groups per session. As mentioned at the beginning of the section, each round consisted of five periods, meant to correspond to different times of the day, with possibly different demand realizations. Subjects made supply decisions for each round; decisions were held fixed across the five periods within the round.

Subjects stayed in the same market for the whole session and did not know who of the other subjects were in the same market as themselves. The interaction in fixed groups approximates best actual circumstances in the kind of electric power markets that we are interested in. Fixed interaction is used in all previously cited experiments on electric power markets. The procedure has also the advantage that the observations from the different groups are statistically independent from each other.

3. Theoretical Predictions and Hypotheses

As mentioned in the introduction, we center our theoretical analysis on a descriptive index, the RSI (Residual Supply Index) and two theoretical models, the MUA (multi-unit auction model) and SFE (supply function equilibrium model). In this section we formulate specific hypotheses derived from these benchmarks. Our hypotheses will pertain to prices. In particular, we will use the volume weighted average price (VWAP, hereafter), which is defined as the monetary value of trades divided by the number of units traded per round. The VWAP provides a useful way to compare observed prices in experiments to theoretical predictions. Formally, let $p(d)$ be the expected price when d units of output are demanded. The expected VWAP is defined as,

$$\overline{P}^e = \sum_{d=20}^{35} \frac{\delta d\, p(d)}{E[d]},$$

where superscript e denotes expectation and δ is the probability of each possible demand level.

The RSI is an indicator of market power given by the following expression:

$$RSI = \frac{total\ capacity - largest\ seller's\ capacity}{demand\ quantity} = \frac{S^{\max} - s_3^{\max}}{d},$$

where we have used the fact that in our notation seller *j*=3 has the highest capacity in all treatments (as does seller *j*=4). Table 2 shows the range of RSI for each of our treatments, as well as the midpoint of the interval.

Table 2: RSI

| | | Low Capacity | High Capacity: | |
		Symmetric:	Symmetric:	Asymmetric:
Load	**Low**	[0.71,1.25], 0.98	[1.03,1.80], 1.415	--
Ratio	**High**	[0.71,0.83], 0.77	[1.03,1,20], 1.115	[0.83,0,97], 0.90

Notes. The first entry in the cell shows the range within which the RSI falls for the various possible demand realizations in the treatment concerned. The second entry shows the midpoint of that interval.

If no firms are pivotal for any demand quantity, as in high-capacity treatments *hsh* and *hsl*, then the RSI exceeds one in all periods of all rounds; in this case any group of three firms has enough capacity to meet demand. If the largest firm is pivotal for all demand quantities, as in treatments *lsh* and *hah*, then the RSI is less than one for all periods of all rounds. For treatment *lsl* firmsº are pivotal for low demand quantities but not for high quantities. RSI is less than one for some periods and greater than one for other periods in treatment *lsl*. We expect treatments with higher average values of RSI to have lower market prices.

The descriptive index is useful but has clear limitations, as it does not propose specific price levels for the different parameter configurations we analyze. The MUA and the SFE take us a bit further in this regard. Both models analyze the interaction between firms as a one shot strategic game, where the strategies consist of supply functions. Of course, the subjects in our experiment are engaged in a 25-round repeated game, so that the equilibrium prescriptions do not exactly pertain to the environment we study. However, as in many other studies the equilibria of the one shot game are relevant benchmarks, particularly given the known and finite time horizon we use. The central difference between the two models is that the MUA pertains to discrete production units while the SFE specifies continuously divisible output.

MUA Predictions

The MUA considers the game as an auction in which each firm *j* submits a vector of offer prices, selected from non-negative real numbers, for discrete units of output, s_j^r. This game is analyzed in Anwar (1999) and Fabra, et al. (2006). For our parameters, this formulation yields pure strategy equilibria for treatments *hsh*, *hsl*, *hah*, and *lsh*; the equilibrium is in mixed strategies for treatment *lsl*. Details are given in Appendix 2.

Equilibria for high-capacity treatments *hsh* and *hsl* involve all firms offering their capacity at a price equal to marginal cost ($c = 5$), i.e., $p_{jk}^{sr} = 5, \forall j, r, k,$ yielding an equilibrium price $p_t^r = 5, \forall r, t$ and profits $\pi_j^r = 0, \forall j, r.$ In other words, when none of the firms are pivotal, the MUA model predicts that competition will work perfectly with the equilibrium price equal to marginal cost. Our experimental setup involves discrete price units $\{5, 6, \ldots, 25, 26\}$, rather than prices chosen from a continuum. For high-capacity treatments our experiments have equilibrium prices equal to either 5 or 6.

Pure strategy equilibria of MUA for *hah* and *lsh* involve asymmetric strategies, in which 3 firms offer all units at a low price, and the 4[th] firm (one of the two high-capacity firms for *hah*) offers all of its units at 25 (the price cap). The equilibrium price is equal to 25 for all demand realizations. The firm submitting high-price offers earns lower expected profit than its rivals; the low-price offers of rivals leave the high-price firm with no incentive to reduce its offers. These equilibria embody maximal exercise of market power; firms extract the maximum possible surplus in equilibrium.[10]

In treatment *lsl* firms are pivotal for high demand quantities, but not for low demand quantities. There are no pure strategy equilibria of the MUA for *lsl*. Fabra, et al (2006) derive a mixed strategy equilibrium for the case in which firms are restricted to make a single price offer for their capacity. However, Anwar (1999) shows that mixing over a single offer is not an equilibrium when each firm can make distinct offers for multiple units. We are not aware of analytical results for mixed strategy equilibria of MUA in which each firm can submit multiple offers. However, we can provide bounds on mixed strategy prices. Our experiment requires firms to submit offers in discrete price units from the set $\{5, 6, \ldots, 25, 26\}$; a unit offered at 26 will not be accepted and is equivalent to withholding the unit. Each firm submits offers for 9 units. A firm's strategy is a non-decreasing offer schedule for 9 units; each firm has a finite set of strategies to choose from.[11] It is well known that any finite *n*-person non-cooperative game has at least one mixed strategy Nash equilibrium. In Appendix 2 we show that expected equilibrium profit for a seller has a positive lower bound. This profit bound permits us to bound the equilibrium VWAP for treatment *lsl*: $\overline{P}^e \geq 11.55.$

---

[10] This effect of the asymmetric distribution of capacities is reminiscent of the price competition environment with capacity constraints studied experimentally by Davis and Holt (1994). With a symmetric distribution of capacity the (pure strategy) equilibrium price is equal to marginal cost, while for a certain asymmetric distribution of the same capacity the equilibrium prices are above marginal cost.

[11] The strategy set is finite, but very large. There are 14,307,150 strategies to choose from.

SFE Predictions

The second theoretical approach permits firms to submit continuous supply functions to an auctioneer. Klemperer and Meyer (1989) formulate and analyze game-theoretic models in which demand is uncertain and strategies are continuous, non-decreasing supply functions for infinitely divisible output. A Nash equilibrium for such a game is termed a Supply Function Equilibrium (SFE). Genc and Reynolds (2008) extend the SFE analysis of Klemperer and Meyer to permit capacity constraints and supply functions with discontinuities (e.g., step functions).

The SFE formulation has been used in a number of studies to predict behavior in naturally occurring wholesale electricity markets (Green 1999; Newbery 1998; Baldick et al. 2004; Bolle 2001) in which suppliers submit offers for discrete units of output. Output is not infinitely divisible in our experiments. Each firm (subject) submits offers for between 5 and 19 discrete units of output, depending on the treatment. This permits us to explore whether or not the SFE model provides useful predictions of behavior in an environment with discrete units. In addition, our experiments permit us to compare the predictive power of the SFE model to that of the MUA model in a particular setting.[12]

Details of the SFE method applied to our parameters are presented in Appendix 3. Here, we present the main results derived from this theory. The first point to make is that Nash equilibrium pure strategies of the MUA model are also equilibrium strategies of the SFE model. Second, the SFE model admits additional pure strategy equilibria compared to the MUA model.

Consider our high-capacity *hsh* and *hsl* treatments. The only equilibrium for the MUA model has price equal to marginal cost (or one tick above marginal cost, for discrete prices); a firm has an incentive to undercut any rival offers that are above marginal cost. However, with infinitely divisible output, if a firm's rivals submit smooth upward sloping supply curves then the firm's best response is to offer its supply at prices above marginal cost. Klemperer and Meyer (1989) show that in general there are multiple supply function equilibria and these equilibria involve non-negative price-cost markups; a SFE with positive markups is sometimes referred to as an implicitly collusive equilibrium. Supply function equilibria for some of the treatments are illustrated in Figure 2. For *hsh* and *hsl* any aggregate supply

---

[12] Under some market rules, one theory may be much more suitable than the other. If market rules limit firms to submitting offers with one or two steps, then the MUA model is appropriate and SFE is not. Some market rules allow firms to submit upward sloping supply functions. For example, the Southwest Power Pool RTO runs an energy balancing market in which each firm submits multiple price-quantity pairs. This RTO interpolates linearly between pairs to yield a piece-wise linear, upward sloping supply curve for the firm. See: http://www.spp.org/section.asp?group=328&pageID=27. A SFE model would be more appropriate than a MUA model for such market rules.

function between (and including) the two bold curves indicated by A and B is consistent with a SFE.[13]

One way to characterize supply function equilibria is by the equilibrium price they generate when d^{max} is realized, i.e., $p_t^r|_{d^{max}} = s^r(35)$. As illustrated in Figure 2, the set of SFE for *hsh* and *hsl* is characterized by $s^r(35) \in [5,25]$, i.e., the price at maximum demand can lie anywhere between the competitive price and the Cournot price, which in this case is the same as the monopoly price.

Consider now treatments *lsh* and *lsl* for which each firm is pivotal; other firms cannot fully compensate if one firm withholds units. The market power induced in these treatments has as a consequence that supply functions at or near marginal cost for all units for all firms

Figure 2: Aggregate Supply Functions for SFE
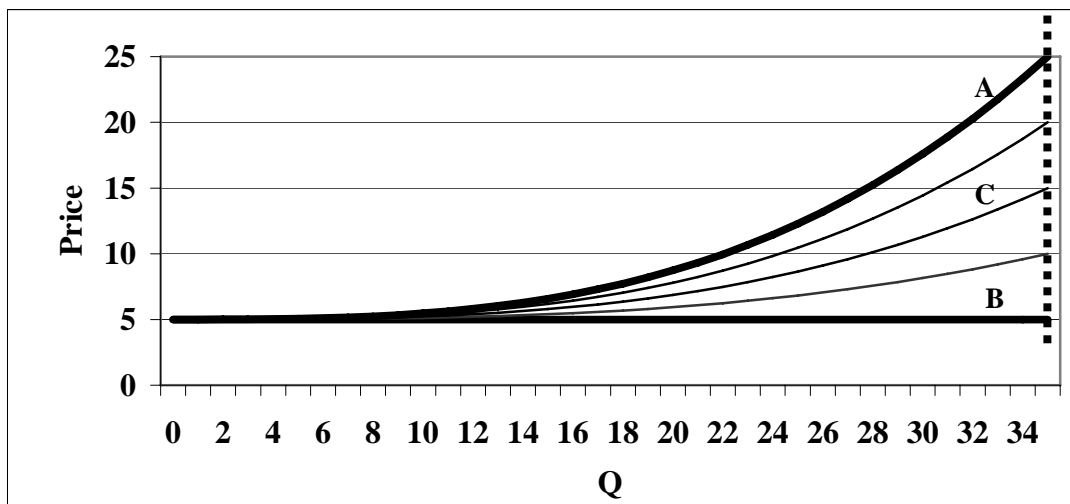
Notes. The curves show possible aggregate supply curves for smooth supply function equilibria. The vertical dashed line at Q=35 indicates maximum demand d^{max}. For *hsh* and *hsl*, any curve between A and B constitutes a SFE. For *lsh* and *lsl* the set of aggregate supply curves for smooth SFE is reduced to the curves between A and some curve C, above curve B.

are not equilibrium strategies. In fact only the aggregate supply functions between some function C, above B, and A are SFE for these low capacity treatments (cf. figure 2). More specifically, for *lsh*, functions characterized by $s^r(35) \in [21,25]$ are SFE and for *lsl* this holds for functions with $s^r(35) \in [18,25]$.

---

[13] We refer to a SFE in which the aggregate supply function is differentiable over the range of possible demand quantities as a *smooth SFE*. For example, supply function equilibria associated with aggregate supply functions labeled A and C in figure 2 are smooth.

For *lsh* we have to also consider non-smooth SFE. If one allows firms to submit non-smooth step-function supply functions (formally, right-continuous functions of price) then the asymmetric equilibria of the MUA with price equal to 25, in which 3 firms offer all units at a low price and the 4[th] firm offers all of its units at 25 are also SFE for treatment *lsh* (but not for any of the other treatments with symmetric capacity distribution).

For treatment *hah* the asymmetric equilibria of the MUA, with one of the two high-capacity firms bidding in all units at a price of 25 (yielding price equal to 25 for all demand realizations), are also supply function equilibria. There are additional supply function equilibria for treatment *hah* with prices below the price cap. In these equilibria, low-capacity firms offer all units at a low price and high-capacity firms have upward sloping (in fact, linear) supply functions, which are between supply curves A and B in Figure 3; $s^r(35) \in [11.9, 25]$ for these equilibria.

Table 3 summarizes the theoretical predictions from the two approaches in terms of volume weighted average prices ($\overline{P}^e$). In summary, the MUA yields very precise predictions for four out of five cases, but only mixed strategy equilibria for treatment *lsl*. SFE predictions

Figure 3: SFE for Asymmetric Capacity Treatment
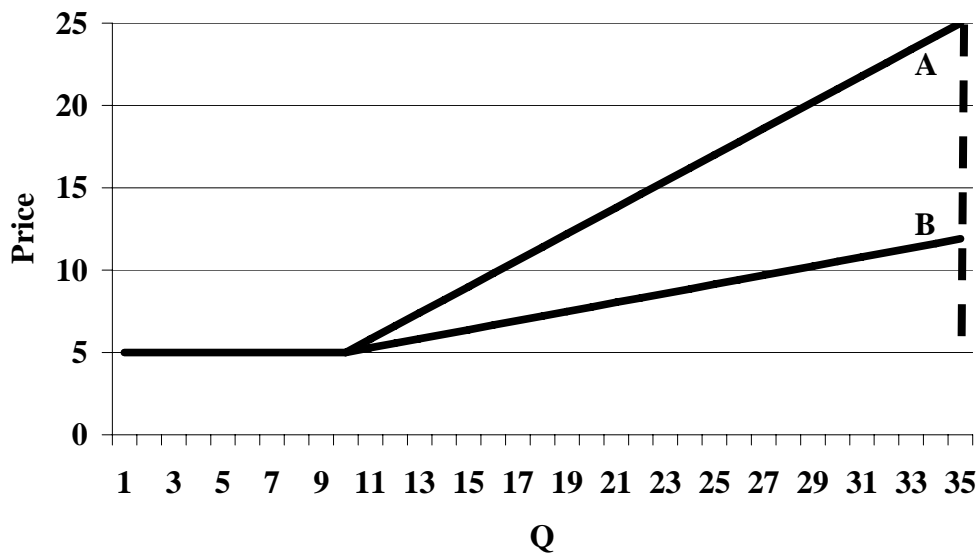
Notes. The small firms submit low price offers for their entire capacity, jointly 10 units (in the graph, these low offers are set equal to 5, but they may be larger than the marginal costs). Each large firm submits a linear, increasing supply function. The aggregate supply function is horizontal for units 1 – 10 and increases after unit 10. The most competitive of these equilibria reaches a price of 11.9 at $d^{max} = 35$; the aggregate supply is labeled B. The least competitive of these equilibria reaches the price cap of 25 at $d^{max} = 35$; the aggregate supply is labeled A.

Table 3: Theoretical Predictions of Volume Weighted Average Price

Capacity:	low load ratio		high load ratio	
	MUA	SFE	MUA	SFE
High symmetric	{5, 6}	[5, 16.4]	{5, 6}	[5, 21.3]
Low symmetric	≥ 11.6	[12.4, 16.4]	{25}	[18, 21.3] & {25}
High asymmetric	n.a	n.a	{25}	[11.2, 23.1] & {25}

include the pure strategy equilibrium predictions of MUA, as well as additional predictions of intervals of average prices that are based on smooth supply function equilibria.

We will use these theoretical predictions to organize our data on volume-weighted-average prices in several ways. We will study whether observed prices remain within the interval prescribed by the SFE and, if so, whether they are well approximated by the more extreme predictions of the MUA model, shown in table 3.

In addition, we will also take a more qualitative look at the data and test a set of formal hypotheses about the comparative static effects that the results in table 3 predict for our treatment variables. The null hypothesis we use as a benchmark stems from the naïve view that prices should not be expected to differ across treatments, since in all our treatments total capacity is sufficient to serve the maximum demand,. The distinct alternative hypotheses are based on both the midpoint values of the RSI (table 2) and the predictions that have been derived using MUA and SFE (table 3). For this comparative static analysis, when considering the predictions of the SFE we focus on the mid-point of the equilibrium price interval. The comparisons we perform pertain to the distinguished treatment variations of total capacity, capacity distribution and demand load factor and to the direct comparison of the two ways in which pivotal power is present in our design.

In the following hypotheses $\overline{P}_x$ stands for the volume-adjusted average price in treatment x. The first two hypotheses refer to the symmetric reduction of capacity, with a high and low demand load factor respectively:

1. With a high load ratio, the presence of pivotal firms, due to symmetrically distributed low total capacity, increases average prices (predicted by RSI, MUA and SFE). Formally:

$$H1_0: \quad \overline{P}_{hsh} = \overline{P}_{lsh} \qquad \text{vs.} \qquad H1_1: \quad \overline{P}_{hsh} < \overline{P}_{lsh}$$

2. With a low load ratio, market power caused by a symmetric reduction in capacity causes an increase in average prices (predicted by RSI, MUA and SFE). Formally:

$$\text{H2}_0: \quad \overline{P}_{hsl} = \overline{P}_{lsl} \qquad \text{vs.} \qquad \text{H2}_1: \quad \overline{P}_{hsl} < \overline{P}_{lsl}$$

The next hypothesis refers to the change in distribution of the high total capacity level for the high demand load factor.

3. With a high load ratio, the presence of pivotal firms, due to asymetrically distributed high total capacity, increases average prices (predicted by RSI, MUA and SFE). Formally:

$$\text{H3}_0: \quad \overline{P}_{hsh} = \overline{P}_{hah} \qquad \text{vs.} \qquad \text{H3}_1: \quad \overline{P}_{hsh} < \overline{P}_{hah}$$

Observe that for the first three hypotheses RSI, SFE and MUA predict the same shifts in price levels. This does not hold for our last three hypotheses. The next hypothesis refers to the two ways in which pivotal power can appear.

4. With a high load ratio, market power caused by a symmetric reduction in capacity has a stronger effect on average prices than market power caused by asymmetry (predicted by RSI and SFE).[14] Formally:

$$\text{H4}_0: \quad \overline{P}_{lsh} = \overline{P}_{hah} \qquad \text{vs.} \qquad \text{H4}_1: \quad \overline{P}_{lsh} > \overline{P}_{hah}$$

The RSI predicts a shift simply because the aggregate capacity that is left after a pivotal supplier withdraws his capacity from the total is smaller under lsh than under hah. The SFE picks this up; some of the lower prices that are equilibrium for the hah treatment are not part of the equilibria for lsh. In contrast, the MUA model does not suggest a difference between these two cases; both ways of introducing pivotal power lead to the same asymmetric equilibrium with the highest possible price.

The two remaining pair-wise comparisons pertain to the impact of changing the load factor. The RSI is mute here, since it does not take demand conditions into account.

5. With a high symmetric capacity, the change from a low to a high load factor yields higher prices (predicted by RSI and SFE). Formally:

$$\text{H5}_0: \quad \overline{P}_{hsh} = \overline{P}_{hsl} \qquad \text{vs.} \qquad \text{H5}_1: \quad \overline{P}_{hsh} > \overline{P}_{hsl}$$

---

[14] The midpoint of the price range [18, 21.3] is higher than the one for [11.2, 23.1].

6. With a low symmetric capacity, the change from a low to a high load factor leads to higher prices (predicted by RSI, MUA and SFE). Formally:

$$H6_0: \quad \overline{P}_{lsh} = \overline{P}_{lsl} \qquad \text{vs.} \qquad H6_1: \quad \overline{P}_{lsh} > \overline{P}_{lsl}$$

Observe that both the RSI and the SFE prescribe a directional shift for all the cases we consider. The prescriptions of the MUA do not change for two of the parameter changes. For the parameters of our design this model suggests that (hypotheses 4) it makes no difference whether pivotal power results from symmetrically distributed lower capacity than from an asymmetric redistribution of a given level of capacity. In addition, the MUA suggests that (hypothesis 6) changes in demand conditions only affect behavior under low total capacity.

4. Results

We start with a general qualitative overview of the supply functions submitted by our subjects. This is followed by an analysis of the aggregate supply functions. We then present data on average volume weighted average prices, compare them with the equilibrium predictions and formally test our hypotheses H_1-H_6. In this way we gradually move from a complete picture of the supply functions to more specific hypotheses.

When submitting their individual supply functions, subjects typically submitted all units that they had available. In the low capacity treatments (where each subject had a capacity of 9 units) on average 8.9 units were offered at a price lower than or equal to p^{max} (=25). In the symmetric high capacity cases (12 units each) on average 11.7 units were offered. In the asymmetric treatment *hah* (two firms with 5 units and two with 19) the low capacity firms always offered all units whereas the firms with high capacity on average offered 18.6 out of 19 units at a price lower than or equal to 25. This is an indication that attempts to exert market power were done by offering units at high prices, not by withholding them altogether.[15]

Figure 4 gives the average aggregate supply function per treatment, distinguishing between the low load ratio and high load ratio cases. In both panels the ranking of the functions, in relation to the load factor, is the same. The highest prices are asked for the low capacity treatments *lsh* and *lsl* and the lowest for the symmetric high capacity cases *hsh* and *hsl*. The supply function for the asymmetric capacity case lies somewhere in between, in the

---

[15] Across all treatments, in 82.2% of the rounds the aggregate supply function offered the maximum total capacity at prices lower than or equal to 25.

top panel. This simple presentation of aggregate supply functions already suggests that price shifts will be like those proposed by SFE.

Figure 4: Aggregate Supply Functions

High Load Ratio
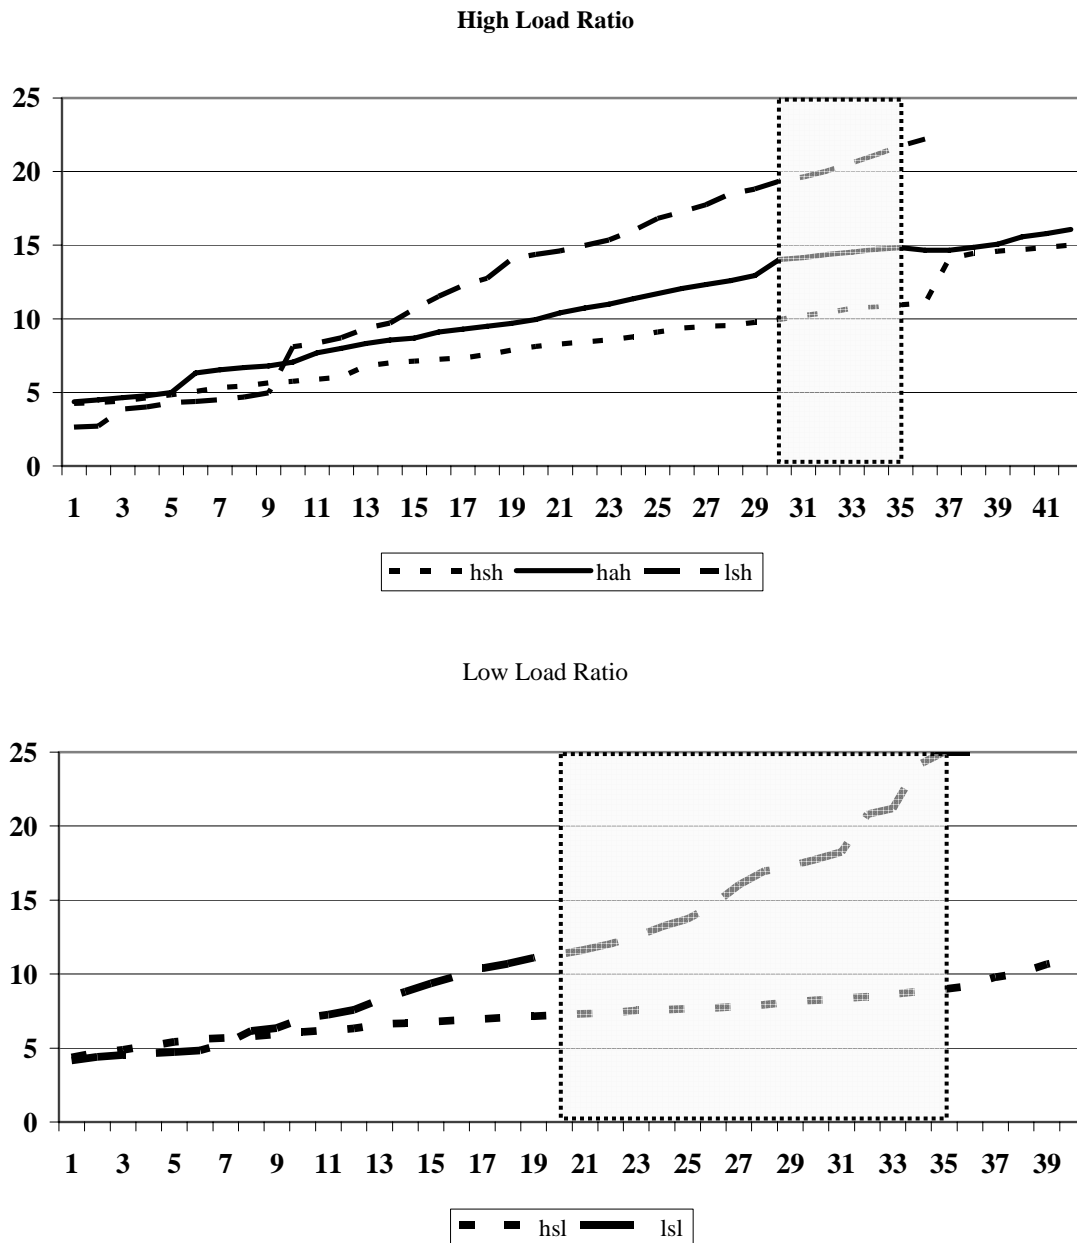
Low Load Ratio

Notes: The lines show the average supply function across all rounds for a high load ratio (top panel) and low load ratio (bottom panel). The window with dashed contour shows the area $[d^{min},d^{max}]$ within which demand may vary. Rounds are indicated on the horizontal axes, prices on the vertical axes.

Table 4 shows volume weighted average prices both for all rounds and for the last 5 rounds, averaged over all groups of each treatment, together with the equilibrium predictions for the two models we consider.

Table 4: Predicted and Actual Volume Weighted Average Prices

| | Equilibrium Predictions | | Actual | |
	MUA	SFE	All Rounds	Last 5 Rounds
High Load Ratio				
High Sym Cap (hsh)	{5,6}	[5, 21.3]	10.48	9.14
Low Sym Cap (lsh)	{25}	[18, 21.3] & {25}	20.34	19.95
High Asym Cap (hah)	{25}	[11.2, 23.1] & {25}	14.48	13.48
Low Load Ratio				
High Sym Cap (hsl)	{5,6}	[5, 16.4]	8.09	6.98
Low Sym Cap (lsl)	≥ 11.6	[12.4, 16.4]	16.50	17.91

Focus first on the two cases with high symmetric capacity, *hsh* and *hsl*. For both load factors prices are above the prediction of the MUA and within the range of the SFE predictions. Note also that prices are lower in the last five rounds than in earlier rounds. For low symmetric capacity with the high load factor – *lsh* - prices are below the MUA prediction and again within the interval of the SFE. For *lsl,* observe that prices are somewhat above the upper limit of the SFE interval and above the lower limit of the mixed-strategy equilibrium support. Finally, for the high asymmetric case, *hah*, prices are below the MUA prediction and within the SFE interval.[16] Formal tests of differences in average VWAP are presented below, when we discuss the results for our hypotheses testing.

Given that average prices are different in the final five rounds than across all rounds, the dynamics of the VWAP may be important. Therefore, we now examine how these prices changed over time during the experiments. Figure 5 presents their development across rounds, separately for each treatment. Starting with the symmetric treatments, figure 5 shows that prices for both low capacity treatments are substantially and consistently above those for the high capacity treatments. The differences increase over rounds: the primary reason is that average prices for high symmetric capacity treatments decrease steadily.

Comparing *hsh* to *hsl* one can see that prices for the former are above those for *hsl* in all rounds. In addition, average prices in these high symmetric capacity experiments are above the (highest) MUA prediction of 6 in all rounds. Thus, aggregate behavior in these high symmetric capacity experiments appears to be inconsistent with MUA predictions, although prices appear to be moving toward the MUA prediction over time. We will explore this issue further when we examine data from individual markets. The ordering of average prices in *hsh*

---

[16] The fact that prices stay away from the extreme predictions of the MUA could be attributed to a behavioral tendency not to choose prices at the edges of the choice space. However, it is worth pointing out here that in experiments with the double auction and the box demand design prices often do go all the way to the extremes (Davis and Holt, 1993).

Figure 5: Development of Volume Weighted Average Prices
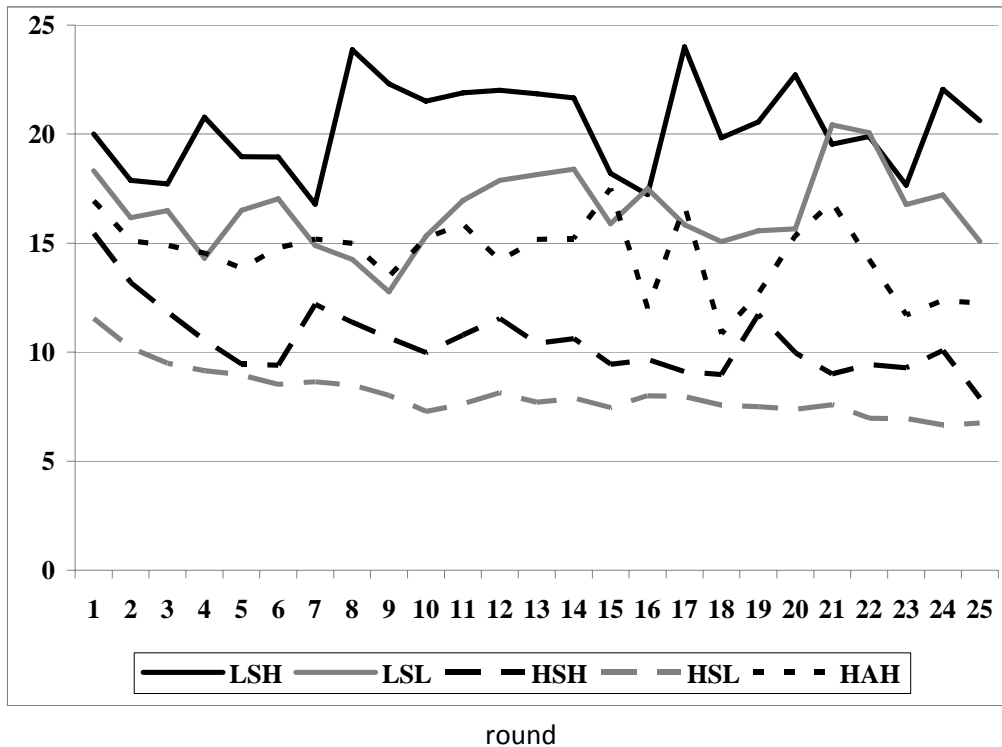
round

Notes. For each treatment the graph shows the volume weighted average price at each round.

and *hsl* would be consistent with a single upward sloping aggregate supply function in a SFE. This is true because the low demand realizations in *hsl* would cut the aggregate supply curve at prices below the clearing prices for *hsh*.

Average prices for the low symmetric capacity treatments (solid lines in Figure 5) vary across rounds, but tend to stay within to slightly above the intervals predicted for smooth SFE (see Table 4). Note that average prices for *lsh* are inconsistent with the MUA prediction of 25. Their dynamics show no tendency toward this prediction. Average prices for *lsl* are consistent with the MUA prediction, in the sense that they are above the lower bound prediction for the mixed strategy MUA equilibrium.

Finally, average prices for *hah* vary over rounds but tend to lie within the interval of equilibrium prices for smooth SFE. In later rounds, average prices are in the lower portion of this predicted interval. Average prices for *hah* are clearly inconsistent with the pure strategy equilibrium prediction of 25 for MUA. If anything, they are converging away from this predicted level.

We conclude that the time trends in our data could be interpreted as convergence in the direction of the MUA prediction only in the symmetric high capacity cases. Even in these cases, this predicted level has not been reached, even after 25 rounds.

Figure 5 aggregates observations across markets but disaggregates across rounds. We do the reverse in figure 6, which shows average prices across rounds separately for each market. To highlight the effects of learning we distinguish between the average across all rounds and the average across the final five rounds. This figure confirms that just like the aggregate prices (figure 5), the average prices per market lie largely within the bounds predicted by SFE. In the absence of market power, the observations for *hsh* and *hsl* appear to be drawn towards the competitive prices predicted by MUA. For each of these two treatments, all groups but one had average prices near the MUA prediction for the last five rounds. For the treatments with market power, the observations appear to be more or less uniformly spread over the predicted interval of smooth SFE prices.[17]

Figure 6: Average Price per Market
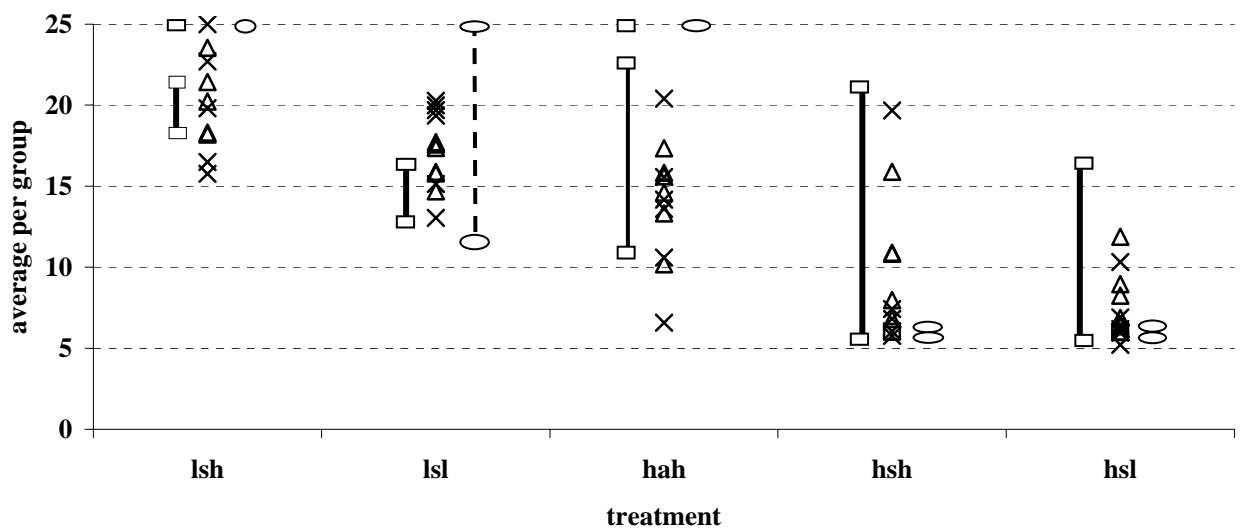
Notes. Markers denote the volume weighted average price per market across all rounds (triangles) or last five rounds (crosses). Rectangles and lines connecting rectangles denote SFE predictions. Ovals and the dashed line between the two ovals for lsl denote MUA predictions.

---

[17] Note that for *lsh* there is one group that has VWAP at the MUA prediction of 25 in the last 5 rounds. In the MUA equilibrium, three firms offer all units at a low price and one firm offers all units at a price of 25. The data for this group reveal that there are two firms offering all units at a low price and two firms offering at the price cap.

We now move to the tests of the hypotheses 1 to 6 about differences in prices across treatments presented in section 3.[18] Table 5 presents the results of Mann-Whitney tests for all pairwise differences in means across the five treatments. It takes the (volume weighted) average price per market (across all rounds) as the unit of observation. The p-values pertaining to our six hypotheses are shown in italics. For the hypotheses, we only need to consider these results. Observe that five out of six of the differences in italics are statistically significant at the 10%-level or better (four are significant at the 1%-level) and therefore support the alternative hypotheses against the null of no differences in average prices.

Table 5: Pairwise Mann-Whitney Tests for Volume Weighted Average Prices

	hsl	*hah*	*lsh*	*lsl*
hsh	*0.137*	*0.089*	*0.004*	0.009
hsl	-	0.002	0.002	*0.001*
hah	-	-	*0.002*	0.032
lsh	-	-	-	*0.002*

Notes. Cell entries give the *p*-value for the Mann-Whitney test for the null hypothesis that the difference in means between the treatments in the row and column concerned are equal to zero. Results in italics are relevant for the hypotheses developed in section 3, as explained in the main text.

We summarize the results of our hypotheses testing in the following way:

(1) Symetrically decreasing capacity (with either load ratio) and of asymmetrically redistributing a given total capacity (with the high load ratio) all have the positive effects on prices predicted by the RSI and both theoretical models. In other words, when the RSI, MUA and SFE all yield the same comparative static prediction, this is confirmed by our data.

(2) With a high load ratio, market power caused by a symmetric reduction in capacity has a stronger effect than market power caused by asymmetry; this is in accordance with the *RSI* and SFE, while MUA is mute on this particular comparison.

(3) With high symmetric capacity the change from a low to a high load ratio does not significantly affect prices, an effect predicted both by RSI and SFE. In contrast, for low symmetric capacity, the change in load ratio does lead to higher prices, an effect predicted both by RSI, MUA and SFE.[19]

Finally, we consider the matter of *who* is exerting market power. We focus on the firm that determines the supply price of the 35[th] unit. If this is the same firm in every round, then in some sense behavior within a market is stable. However, if market power is being exerted,

---

[18] As mentioned in section 3 we only evaluate separate variations of the different treatment variables.
[19] Observe in table 3 that for the case of high capacity the range of SFE equilibrium prices only changes with the load factor in terms of an increase in the upper limit of the interval. In contrast, for low capacity the change in the price predictions consists in increases of both the lower and upper limit of the interval and in the addition of the highest possible price.

the firm supplying this unit will typically have lower earnings than other firms in the round concerned.[20] Table 6 shows the extent to which the price of the 35[th] unit was determined by one or two firms in the market.

Table 6: Firms Determining $s^r(35)$

	$\tau=hsh$	$\tau=hsl$	$\tau=hah$	$\tau=lsh$	$\tau=lsl$
One firm	0.48	0.56	0.70	0.56	0.66
Two firms	0.82	0.84	0.91	0.87	0.9

Notes. Numbers give the average (across markets in a treatment) fraction of the 25 rounds that the price of unit 35 was determined by a single firm (2[nd] row) or two (out of the four) firms (3[rd] row).

Note that if all firms are equally likely to determine $s^r(35)$ the fractions in the 2[nd] row should all be approximately 0.25 and those in the third row approximately 0.5. This is obviously not the case. For the asymmetric treatment, *hah*, one may expect the two large firms to alternate, yielding fractions 0.5 and 1, respectively. On average there appear to be only small differences across the symmetric treatments. 50-60% of these prices are determined by a single firm in any market and two firms account for more than 80%. We conclude that there is a strong asymmetry in the bidding by distinct firms in a market. Even when all four firms have (equal) market power in *lsh* and *lsl*, (almost) 90% of the prices at unit 35 are determined by only 2 of the 4 firms. While we clearly observe heterogeneous bidding by subjects, it is not consistent with the MUA predictions of asymmetric equilibria for treatments *lsh* and *hah*. These equilibria involve low offers for all units by three subjects and all capacity of the fourth subject offered at the price cap. We did not observe this in *lsh* or *hah* markets.

We finish with a direct statistical comparison between the two theoretical models. It is not possible to directly compare the the MUA and SFE VWAP predictions, because MUA predicts one or two discrete points for four out of five treatments whereas SFE predicts intervals. Obviously, SFE has a better shot at predicting correctly because it predicts a larger subset of the possible VWAP's (with the exception of *lsl*). To circumvent this problem, we use the mean value of the prices predicted by SFE as the SFE point predictions.[21] For the range of possible MUA predictions in *lsl*, we also take the mean value. Table 7 shows the resulting predictions.

---

[20] If other firms do not offer at high prices (as predicted for *lsh* by MUA, for example) then it may easily happen that the same firm provides the 35[th] unit, round after round. When market power is exerted in our experiments, such asymmetry typically did occur. Here we analyze whether firms alternated in exerting such power or that it was typically the same firm that did so.

[21] For lsh and hah, we did not use the prediction of 25 for SFE. Doing so does not qualitatively change our conclusions.

Table 7: Mean Values of Predicted Prices

	hsh	*hsl*	*hah*	*lsh*	*lsl*
MUA	5.5	5.5	25	25	18.3
SFE	13.2	10.7	12.2	19.6	14.4
Mean observation	10.5	8.1	14.5	20.3	16.5

One approach would be to use the observation per market i in treatment j (either across all rounds or for only the last 5 rounds) and compare to the prediction: $\theta_j \equiv VWAP_i - P^*_j$, where P^*j is the prediction from table 7 (either for SFE or for MUA). Because the MUA predictions are at the corners of the outcome space we cannot assume that θ_j is normally distributed). Instead, we compare $|VWAP_i - P^*_j|$ across treatments. Our results show that, across all rounds, the SFE absolute prediction errors are significantly smaller than those for the MUA (p=0.02). For the last 5 rounds they are smaller but not significantly so (p=0.61). The change in the last five rounds is probably due to the decrease in prices, in the direction of the MUA prediction, in treatments hsl and hah.

5. Concluding Discussion

We set out to experimentally study the effects of pivotal power, motivated both by the results of empirical field data studies and by the predictions of recent relevant theoretical models. A first conclusion from our experiments is that the more fundamental intuitions about the impact of pivotal power are supported by our data. Prices are higher when (some) firms are pivotal. The existence of aggregate excess capacity is not enough to guarantee competitive prices. This general result is accordance both with the predictions of the intuitive RSI and of both the SFE and the MUA model.

Our experiments also permit us to compare the predictive power of two alternative theoretical models: the MUA model based on discrete output units and the SFE model based on divisible output. The MUA model provides sharp equilibrium predictions for 4 out of our 5 treatments; for these treatments the MUA predicts either competitive pricing or monopoly pricing. The SFE model predicts a larger set of equilibria that includes the sharp MUA equilibria as well as additional equilibria based on upward sloping supply functions.

These additional equilibria of the SFE are useful for explaining observed behavior in two principle ways. First, for treatments with no market power, observed supply functions were upward sloping and prices tended to remain above the competitive price prediction of the MUA. The implicitly collusive equilibria of the SFE capture this behavior. The movement of

prices toward marginal cost for most markets, however, did appear to show dynamics in the direction of the MUA prediction. Only in this weak way does the basic insight, that the "implicitly collusive" equilibria that arise when offer prices and quantities are infinitely divisible are eliminated once some discreteness is introduced, obtain support in our data.

Second, the MUA model predicts monopoly pricing (at the price cap) for two of our market power treatments. Observed behavior in these two treatments was inconsistent with this sharp MUA prediction; behavior was more consistent with additional equilibria from the SFE that involve upward sloping supply functions and lower market prices. As mentioned above, we do not think that the fact that prices stay away from the extreme predictions of the MUA can be simply explained by a behavioral tendency not to choose prices at the edges of the choice space. In double-auction experiments with the box demand design prices often do go all the way to the extremes. An additional consideration here is that our design involves a repeated game, albeit a finite one. Given that repeated interaction can facilitate tacit collusion in experimental oligopoly settings (Abbink and Brandts, 2008), it is noteworthy that in our case it does not lead to the attainment of a monopoly price one-shot equilibrium.

As noted in the introduction, market power due to the presence of pivotal suppliers has been documented to contribute to high prices and inefficiency in wholesale electricity markets and is a significant concern for public policy toward the electric industry. Our experimental results are consistent with evidence from naturally occurring electricity markets that pivotal power contributes to higher market prices. An important finding however, is that the exercise of market power by pivotal suppliers in our experiments was not as severe as equilibrium predictions of the MUA model. These predictions require that agents adopt strategies that support an asymmetric equilibrium with payoffs that differ substantially across agents (even for agents with identical costs and capacities). Experimental results for treatments with pivotal suppliers were more consistent with SFE predictions involving lower prices.

Our data show that – for our parameter choices – pivotal power due to symmetric low capacities has a stronger effect on prices than high total but asymmetrically distributed capacity. This is predicted both by the RSI and SFE. However, the MUA predicts no change, with monopoly prices for both cases. In practice, this suggests that it is not as bad to have a few large companies and a few small ones in a wholesale electricity market as to have the same total number of average level firms. In a context of repeated interaction this could be driven not only by the logic underlying static models, but also by the fact that collusion may be easier under symmetric conditions.

A final result that we wish to highlight here pertains to the effects of increasing the load factors. We find that when both models suggest that it will affect prices it does have this effect. We interpret this as indicating that the models do identify a potential influence factor, but that it only shows up in the data when it is a strong force. Stronger variations in demand reduce market power, in a situation where this power is otherwise strongest.

Appendix 1

This appendix gives the English translation of the original Dutch instructions for the sessions with symmetrich high capacity (12 units per producer) and low load factor. The instructions were programmed as html pages. Horizontal lines indicate page separations.

INSTRUCTIONS

You are about to participate in an economic experiment. The instructions are simple. if you follow them carefully, you can make a substantial amount of money. Your earnings will be paid to you in euro's at the end of the experiment.

In the experiment, we use the currency **'franc'**. At the end of the experiment, we will exchange the francs for guilders. The exchange rate to be used is 1 euro for 250 francs. **For each 1000 francs, you will therefore receive €4.**

We will use numerical examples in these instructions. These are only meant to be an illustration and are irrelevant for the experiment itself.

In these instructions, you may click on the links at the bottom of each page to move forward or backward. Sometimes, there will be more text on a page than can fit onto your screen. When that is the case, you can use the scroll bar on the right to move down.

next page

ROUNDS AND PERIODS

The experiment will consist of **25 rounds** today, preceded by **5 practice rounds**.

In the 25 rounds, you will be a member of a group. Aside from you, the group will consist of **3 other people**. The composition of the group is anonymous. You will not know who is in the group with you. Others will not know that you are in their group. The composition of your group is the same for the whole experiment. You will have nothing to do with people in other groups.

In the experiment, you will participate in a market, in which fictitious goods will be produced and sold. The final consumers of the good will be simulated by the computer. All participants will be producers of the good. There are 4 producers in each group.

In the **practice rounds, you will not be in a group** with other participants. The computer will simulate the choices of other group members. It does so in a completely random manner. You cannot learn anything about others' behavior from these simulated choices.

Each round will consist of **5 periods**. In each period, the computer will decide how many of the goods to buy. You do not need to do anything between periods. At the beginning of the round, you will decide how many units you are willing to produce and sell.

This choice will be valid in each of the 5 periods in that round. The remainder of these instructions will explain the market and the rules you must abide by.

SIMULATED BUYERS

In this experiment, the decisions to buy (fictitious) goods are not made by participants but by the computer. This will be done as follows.

In **each period**, the computer will buy **between 20 and 35 units** of the good. Each number between 20 and 35 (inclusive) is equally likely. Because there are 16 integer numbers between 20 and 35, in each period there is a probability of 1/16 that any one of these number will be drawn. After a number has been drawn, it may be drawn again in a next period of a round.

To determine the **price** that the computer will pay per unit bought in a period, it is determined at what (minimum) prices the group members are willing to sell each unit. Below, we will explain how this determines the price paid by the computer.

The computer will never pay more than **25 francs per unit**, however. If not enough sellers are willing to sell for a price lower than 25, the computer will buy as many as it can for 25 francs.

PRODUCTION AND COSTS

At the beginning of the experiment, each participant will receive 1250 francs as a starting capital. You will see this amount on your screen when the experiment starts.

In each round, each participant is a producer who must decide how many units of the good he or she wishes to produce. No producer is allowed to produce more than **12 units**.

For each unit a producer is willing to produce, he or she must determine the minimum price that he or she wishes to receive for that unit. We will call this the '**ask price**'. How this is reported, will be explained shortly.

There are **costs** related to producing goods. For each unit you produce, you must pay **5 francs**.

ASK PRICES

For each unit you would like to offer, you need to indicate at what price (the 'ask price') you are willing to sell it. You may ask different prices for distinct units. For this, the following rules apply.

If you offer a unit for sale, you must also offer all preceding units. For example, if you indicate a minimum price asked for unit 3, you must also offer units 1 and 2.

Your price asked for a unit must always be **higher than or equal to the price asked for the preceding unit**. So: your ask price for the second unit may not be lower than for the first unit. Your ask price for the third unit may again not be lower than your ask price for the second unit, etc. Each producer can produce at most 12 units.

Your ask price may be lower than the costs. Note that you may make a loss on that unit in that case. For example. assume that your price asked for the first three units is 3. Assume that the three units are bought by the computer at a price of 4. For each unit, your production costs are 5 and your revenue is 4, so you make a loss of 1. For the three units together, your loss is 3.

All units for which you ask a positive price are offered on the market. However, a unit is only sold if the computer is willing to pay your ask price for that unit. How this is determined will be discussed shortly.

SUBMITTING YOUR ASK PRICES

nr	cost	cumm.	ask price
1	5	5	
2	5	10	
3	5	15	
4	5	20	

Confirmation

To enter your ask prices, you will use a window that looks like this. Note that you can only see the first 4 units. In the experiment, you will be able to scroll down to the other units. This is not possible in these instructions.

The first column (**nr**) indicates the number of the unit. The second column (**cost**) gives the cost per unit (5). The next column (**cumm**) gives the total costs for that level of production. The last column (**ask price**) will be used to enter the minimum price you wish to receive for that unit.

You indicate your willingness to sell units by entering the amount you want to receive in the column **ask price**. It is up to you to decide how many different numbers you wish to enter, as long as no ask price is lower than the preceding one. You may enter a different number for each unit, the same for all units or anything in between. It is also up to you to decide how many units you want to offer. There is a maximum of 12, however.

To help you when entering numbers, the following happens. If you enter a price for a unit, the same number is automatically entered in all previous units for which no number had been entered yet. For example, if you start by entering a price of 12 in unit 3, 12 is also entered in units 1 and 2. If you then enter 22 for unit 5, 1-3 stay at 12 but 22 is entered for unit 4. You may practice this in the practice rounds. Units where you do not enter a number are not offered for sale.

When you are satisfied, **you must confirm your choice**. As long as you have not done so, you can still change every and any price asked. Note that your decision is not valid until you have confirmed. The experiment will not proceed until everyone has confirmed her or his production decision. You must also confirm if you wish to produce zero units. You do so by clicking the confirmation button without entering any numbers.

DETERMINING THE PRICE PAID

The price paid by the computer for any unit bought is often not equal to the ask price. **The price paid is never lower than the ask price**, however.

After the computer has determined how many units it wants to buy in a period, it considers all the ask prices in your group of producers. It **first** buys the unit offered at the **lowest** price, then the second lowest price, etc. The price it pays is the same for all units bought.

In each period of a round, the number of units the computer wants to buy is some randomly drawn quantity between 20 and 35.

Example
Let's say that in some period, the computer wants to buy 29 units. It then checks whether it can buy all 29 units at a price of 25 francs or less. If there at least 29 ask prices less than or equal to 25, the computer will buy all 29 units. If not, it will buy as many units as it can for 25   francs.
If there are at least 29 ask prices lower than or equal to 25, the computer chooses the 29 lowest ask prices. The price paid is then equal to to the 29th ask price. For example, if 12

units are offered for 10 francs, 10 units for 12 francs, and 10 units for 18 francs, the 29th price is 18 francs. All 29 units will be sold for 18 francs. Note that some units for which the ask price is 18 remain unsold, however.

The same procedure holds for any other quantity (randomly) chosen by the computer. One way to picture how the price paid is determined is as follows. Consider all of the ask prices submitted by members of your group. Order them from low to high. Then count how many units can be sold for 25 francs or less.

If this is **less** than the number chosen by the computer, then the price is 25 and all units with ask prices less than or equal to 25 are sold.

If this is **more** than or equal to the number chosen, the computer is able to buy the units it wants. It looks for the lowest price at which it can by all of these units. **This is the price the computer pays for all units.** Note that this price paid is the ask price of the last unit bought. For all other units bought the ask price is lower. For all units not bought, the ask price is at least as high as the price paid.
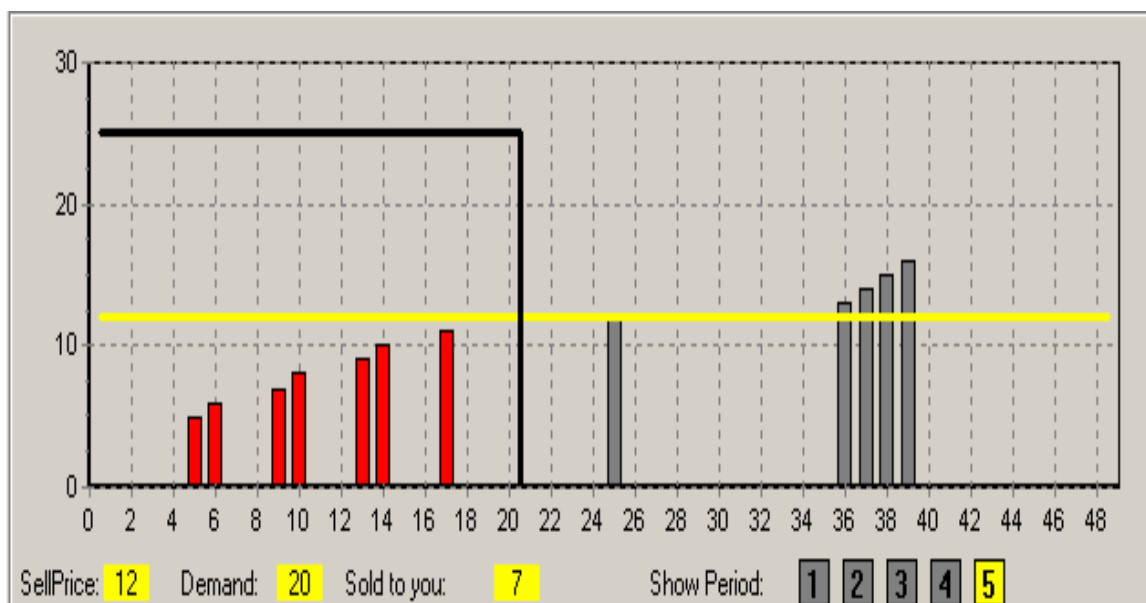
5 PERIODS

After everyone has confirmed their ask prices at the beginning of the round, the computer orders these from low to high. **If two ask prices are equal, the computer randomly determines their order.** Then, it runs the 5 periods of the round. In each period, it randomly determines the number of units it wants to buy.

In every period, it determines which units are bought and what price is to be paid. It will show you the results for a few seconds and then move on to the next period. After the 5th period, you will be able to review all of the periods of the round at your own pace. **The experiment will only proceed to the next round after all participants have indicated that they are ready.**

RESULTS OF A PERIOD

This is an example of how the results of a period will be shown to you.

These are the results of period 5, which can be seen from the yellow square with a '5' in the bottom right corner. In this period, the price paid was 12 and the computer wanted to buy 20 units. This participant sold 7 of these units. This information is given in the bottom left corner.
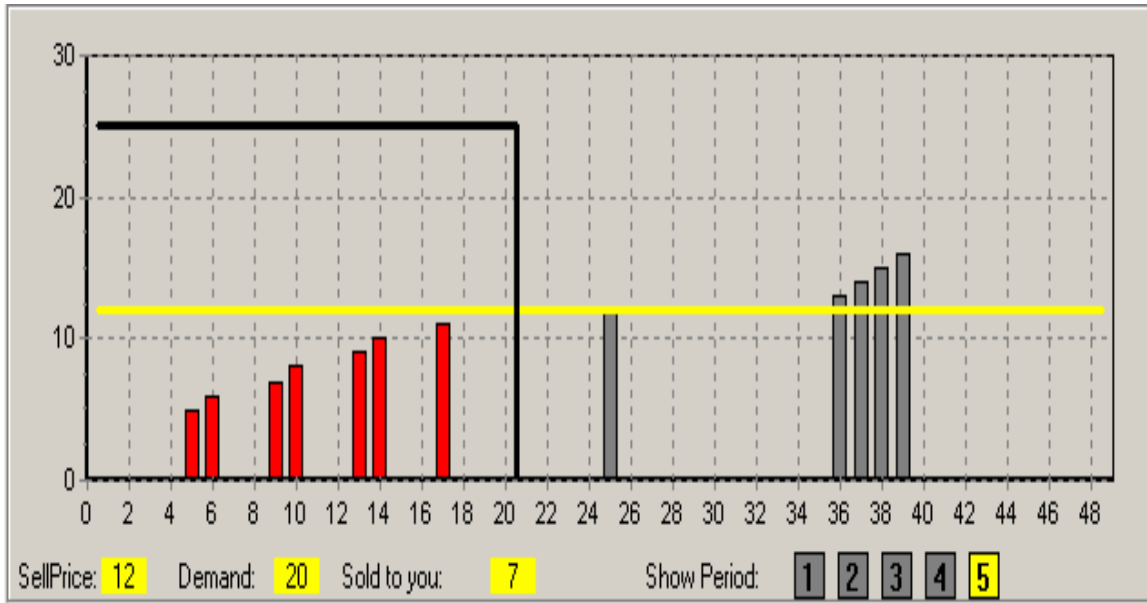
The bars in the graph are the ask prices submitted by this participant. Notice that this participant offered (all) 12 units. The ask prices ranged between 5 and 16. Red bars indicate units that were sold in this period and grey bars indicate unsold units. The location of a bar is determined by its place relative to the ask prices of all participants, ordered from low to high. For example, the bar indicating an ask price of 5 is the 5th unit. This means that there were five units offered by other participants at ask prices of 5 or lower.

The graph shows more details. The number of units the computer wanted to buy (20) is shown by the black line. This line also shows the maximum price of 25 francs. The price paid (12) is given by the horizontal yellow line in the graph. Notice that in this example the computer was able to buy all 20 units at a price lower than 25, because there are unsold units with an ask price less than 25.

Note that **you will not see what prices were asked by the other participants**, only your own. Also note that this participant offered a unit at a price of 12 that was not sold, even though the price paid is 12. This can only be the case if at least one other producer entered an ask price of 12 and the computer (randomly) put that other unit before this one.

RESULTS OF A PERIOD

After everyone has entered and confirmed their ask prices, a graph like this is shown for each of the 5 periods. Each period is shown for a few seconds. After the 5th period, you will be able to review any period by clicking on the numbers in the bottom right corner.

You will have to indicate that you have finished reviewing all of the information of a round by clicking on a 'Ready' button (not shown here). **We will not proceed to the next round until everyone has indicated that they are ready**.

previous page next page

RESULTS OF A PERIOD

period	demand	price	sold	costs	profit
1	25	12	5	25	35
2	29	12	6	30	42
3	26	12	5	25	35
4	28	12	6	30	42
5	20	12	4	20	28

total over periods: 182

Ready

Here you see the 'Ready' button where you can indicate that you have finished reviewing the periods of a round.

You also see a table summarizing the results of a round. This will appear after period 5 has been completed by the computer.

The table has one row for each period. The first column gives the **period**. The second column (**demand**) shows the quantity that the computer wants to buy in that period. If the number is in **black**, it was able to buy all units. If it is in **red**, there were not enough ask prices smaller than or equal to 25.

Under the header '**price**' you will find the price paid per unit in that period. The number of units sold by you is given in the column '**sold**'.

Your production costs (the number sold times 5) are given in the column '**costs**'. Finally, your profit in this period (the price multiplied by the number sold minus your costs) is given under '**profit**'.

Between the table and the 'Ready'-button, we give your aggregate earnings in this round. This is the sum of your profits in the five rounds.

AGGREGATE EARNINGS

During the whole experiment, a window in the op left corner will keep track of the round and period you are in. It also gives your aggregate earnings in francs. At the end of the experiment, your francs will be converted into euros.

This brings you to the end of the instructions. You may now take your time to reread parts of the instructions. When you are satisfied that you understand them, you can indicate to us that you are finished, by clicking the '**ready**' button at the bottom of this screen. After that, you may still page through the instructions. However, when everyone has indicated that they are ready, we will move on to the practice rounds.

Appendix 2 – Equilibria of the MUA

Treatments *hsh* and *hsl*:

Each seller has 12 units of capacity in these treatments. In *hsh* demand quantities have a discrete uniform distribution from 30 to 35; the distribution is from 20 to 35 for *hsl*. No seller is pivotal in these treatments. As a consequence, each seller offering all their units at marginal cost ($c = 5$) is a Nash equilibrium.

For our experimental environment with discrete price offers there are also equilibrium strategies involving some offers at 6 (one tick above marginal cost) that yield equilibrium

prices equal to 6 for at least some demand realizations. Consider treatment *hsh*. Suppose that sellers select strategies such that

    1) 29 units are offered at 5,

    2) 19 units are offered at 6, and

    3) each seller offers at least 2 units at 5.

Then the market clearing price is 6 for each demand outcome (since $d \geq 30$). First, note that no seller has an incentive to switch a unit offered at 5 to 6. When offered at 5 this unit is sold with probability one and has profit equal to one. When offered at 6 this unit is sold with probability less than one, with profit equal to one if sold. Second, one can show that no seller has an incentive to switch a unit offered at 6 to 5. While this switch increases the probability that the unit will be sold, this switch also reduces the average price and hence average profit on other units. On balance the impact of a lower average price outweighs the increased probability of selling the unit.

For treatment *hsl* there are equilibria that yield a market clearing price of 6 for high demand outcomes, but not for low demand outcomes.

Treatments *lsh* and *hah*:

In *lsh* each seller has 9 units of capacity and demand quantities are $d \in \{30,31,\ldots,35\}$, with equal probabilities. Let seller j offer its entire capacity at the price cap; that is, $p_{lk}^{sr} = p^{\max} = 25$, for $k = 1,\ldots,9$. Let sellers $l \neq j$ choose offers, $p_{lk}^{sr} \leq 17$, for $k = 1,\ldots,9$. Then the market price is 25 for each demand realization. Expected profit for seller j is,

$$E\left[\pi_j\right] = (25-5)\left(E[d] - \sum_{l \neq j}\left(s_l^{\max}\right)\right) = 110.$$

Sellers $l \neq j$ earn the maximum possible profit for a seller in this environment:

$E\left[\pi_j\right] = (25-5)s_l^{\max} = 180$ Seller j has no incentive to defect since she would have to reduce her offers to 17 or less in order to increase her quantity sold. Even if seller j sold her entire capacity at a price of 17 her payoff would be 108, which is less than the payoff of 110 associated with the high price strategy. So the asymmetric strategies described above are Nash equilibrium strategies; see Fabra, et al (2006) for more details. There are four asymmetric equilibria of this type, with a different seller acting as the high price seller in each equilibrium.

In *hah* there are two small sellers, each with 5 units of capacity, two large sellers, each with 19 units of capacity, and demand quantities are $d \in \{30,31,\ldots,35\}$, , with equal

probabilities. Suppose that one of the large sellers offers their entire capacity at the price cap. This would ensure that the market price is at the price cap (25) for each possible demand realization. If the other three sellers offer all of their units at prices less than or equal to 8, the high price seller has no incentive to change their strategy. These strategies are a Nash equilibrium. There are two asymmetric Nash equilibria for *hah*, with one of the large sellers acting as the high price seller in each equilibrium.

Treatment *lsl* :

The only equilibrium is in mixed strategies. By offering all of its capacity at the price cap, a single seller can guarantee itself expected profit of,

$$E\left[\overline{\pi}\right] = \sum_{d=20}^{35} \left[\delta\left(25-5\right)\max\left\{0, d - \sum_{l\neq j}\left(s_l^{\max}\right)\right\}\right] = 45,$$

where $\delta = 1/16$ is the probability of each possible demand level for *lsl*. This expected profit is a lower bound for a firm's mixed strategy equilibrium profit. This bound on profit permits us to bound a measure of expected equilibrium price.

Given the definition of $p(d)$ in the main text, total expected equilibrium profit for the four firms in the market is defined by:

$$Total\ profit = \sum_{d=20}^{35} \delta d\left(p\left(d\right)-c\right)$$
.

Since each firm must earn at least $\overline{\pi}$ in equilibrium, we have the following inequality:

$$Total\ profit = \sum_{d=20}^{35} \delta d\left(p\left(d\right)-c\right) \geq 4\overline{\pi}$$
.

This implies that:

$$\sum_{d=20}^{35} \delta d\left(p\left(d\right)\right) \geq 4\overline{\pi} + cE\left[d\right],$$

and this permits us to place a lower bound on the volume weighted average price:

$$\overline{P}^e = \sum_{d=20}^{35} \frac{\delta d\left(p\left(d\right)\right)}{E\left[d\right]} \geq \frac{4\overline{\pi}}{E\left[d\right]} + c \approx 11.55$$

Appendix 3 – Equilibria of the SFE Model

In the derivations for the SFE model we treat both price and quantity as continuous variables.

Treatments *hsh* and *hsl* :

There are no pivotal suppliers for these two treatments. Consider the profit for firm i in the event that demand is d, given that rival firm j chooses a differentiable supply function $s_j(p)$ for $j \neq i$. If the clearing price is p and firm i supplies the residual demand, $d - \sum_{j \neq i} s_j(p)$, then its profit is:

$$\pi_i(p,d) = (p-c)\left(d - \sum_{j \neq i}(s_j(p))\right)$$

We seek a supply function $s_i(p)$ for firm i that has the property that the clearing price p maximizes $\pi_i(p,d)$ with $s_i(p) = d - \sum_{j \neq i} s_j(p)$, for each possible $d \in \left[\underline{d}, \overline{d}\right]$. The necessary conditions for an (interior) optimal price for d for each firm i yield a system of ordinary differential equations for supply functions:

$$\sum_{j \neq i} s_j'(p) = \frac{s_i(p)}{(p-c)}$$

for $i = 1,\ldots,4$. There is a continuum of symmetric solutions to this system of the form:

(*) $s_i(p) = \tfrac{1}{4}\overline{d}\left(\frac{p-c}{p'-c}\right)^{\frac{1}{2}}$

where p' is a price parameter that can take on any value in the interval, $\left(c, p^{\max}\right]$; p' is the market clearing price associated with equilibrium supply strategies in (*) at maximum demand quantity, $\overline{d}$. Figure 2 in the main text illustrates aggregate supply functions based on the strategies in (*). Note that in the limit as p' approaches c the supply strategy in (*) converges to the Bertrand strategy of offering all units at marginal cost.

Treatments *lsh* and *lsl* :

For these treatments any one of the sellers is pivotal for some or all demand quantities. When pivotal suppliers are present a strategy of offering capacity at prices close to marginal cost will not be a symmetric equilibrium strategy. If a seller's rivals use strategies in (*) with p' close to c then the seller would prefer to offer all of their capacity at the price cap rather than use strategy (*). Genc and Reynolds (2008) show that the symmetric supply function strategies in (*) are equilibrium strategies for capacity constrained pivotal sellers for a restricted set of p' parameters. For treatment *lsh* the supply functions in (*) are equilibrium

strategies for $p' \in [21.7, 25]$;; for treatment *lsl* the supply functions in (*) are equilibrium strategies for $p' \in [17.7, 25]$. The equilibria associated with these price parameters are the basis for the (volume weighted) average equilibrium price predictions that we provide in Table 3 of the main text.

Treatment *hah* :

In this treatment there are two small sellers (each with 5 units of capacity) and two large sellers (each with 19 units of capacity). There are quasi-symmetric supply function equilibria of the following form. The small sellers each offer their capacity at a low price (e.g., at or near marginal cost). The two large sellers compete for the remaining residual demand (d – 10) by choosing supply functions that are increasing in price. By using arguments similar to those used earlier in this Appendix one can show that there is a supply function equilibrium in which each large seller i uses the linear strategy:

(A3.4) $$s_i(p) = \tfrac{1}{2}\left(\overline{d} - 10\right)\left(\frac{p-c}{p'-c}\right)$$

where the price parameter satisfies, $p' \in [11.9, 25]$. As in the other treatments there is a continuum of equilibria.

References:

Abbink, K. and J. Brandts (2008), "24. Pricing in Bertrand Competition with Increasing Marginal Costs", Games and Economic Behavior, 63, 1-31.

Anwar, A. W. (1999), "The Case for a Discriminatory Pricing Rule in Competitive Electricity Pools", Mimeo, Department of Economics, University of Edinburgh.

Baldick, R. and W. Hogan (2002), "Capacity Contrained Supply Function Equilibrium Models of Electricity Markets: Stability, Non-decreasing Constraints, and Function Space Iterations", University of California Berkeley, PWP-089 Working Paper, August 2002.

Baldick, R., R. Grant and E. Kahn (2004), "Theory and Applications of Linear Supply Function Equilibrium in Electricity Markets", Journal of Regulatory Economics, 25, 143-167.
Bolle F. (2001), "Competition with supply and demand functions", Energy Economics, 23,253-277.

Borenstein, S. J. Bushnell and F. Wolak (2002), "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market", American Economic Review, **92**, 1376-1405.

Brandts, J., P. Pezanis-Christou and A. Schram (2008), "Competition with Forward Contracts: A Laboratory Analysis Motivated by Electricity Market Design", The Economic Journal, 118, 192-214

Davis, D. and C. Holt (1993), Experimental Economics, Princeton University Press.

Davis, D. and C. Holt (1994), "Market Power and Mergers in Laboratory Experiments with Posted Prices", RAND Journal of Economics, 25(3), 467-487.

Fabra, N., N.-H. von der Fehr and D. Harbord (2006), "Designing Electricity Auctions," Rand Journal of Economics, 37, 1, 23-46.

Falk, A. and J Heckman (2009), "Lab Experiments are a Major Source of Knowledge in the Social Sciences," Science, 326, 535-538.

Federal Energy Regulatory Comission (2007), 18 CFR Part 35 (Docket No. RM04-7-000; Order No. 697), June 21, 2007.

Genc, T. and S. Reynolds (2008), "Supply Function Equilibria with Capacity Constraints and Pivotal Electricity Suppliers," working paper, University of Arizona.

Green, R. (1999), "The Electricity Contract Market in England and Wales", Journal of Industrial Economics, 47, 1, 107-124.

Holmberg, P. (2004), "Unique Supply Function Equilibrium with Capacity Constraints", Mimeo, Uppsala University, November 2004.

Klemperer, P. and M. Meyer (1989), "Supply Function Equilibria in Oligopoly under Uncertainty", Econometrica, 57, 6, 1243-1277.

Kremer, I. and K. Nyborg (2004), "Divisible Good Auctions: The Role of Allocation Rules", RAND Journal of Economics, 35, 147-159.

Newbery, D. (1998), "Competition, contracts, and entry in the electricity spot market", RAND Journal of Economics, 29, 4, 726-749.

Newbery, D. (2002), "Problems of liberalising the electricity industry", European Economic Review, 46, 919-927.

Rahimi, A. and A. Sheffrin (2003), "Effective Market Monitoring in Deregulated Electricity Markets", IEEE Transactions on Power Systems, 18, 2, pp. 486-493.

Rassenti, S., V. Smith and B. Wilson (2002), "Using experiments to Inform the Privatization/Deregulation Movement in Electricity", The Cato Journal, 22, 3, winter 2002.

Rassenti, S., V. Smith and B. Wilson (2003), "Controlling Market Power and Price Spikes in Electricity Networks: demand-side bidding", Proceedings of the National Academy of Science, 100, 5, 2998-3003.

Sheffrin, A. (2001), "Critical Actions Necessary for Effective Market Monitoring", Mimeo, Department of Market Analysis, California ISO.

Sheffrin, A. (2002), "Predicting Market Power Using the Residual Supply Index", Mimeo, Department of Market Analysis, California ISO.

Staropoli C. and C. Jullien (2006), "Using Laboratory Experiments to Design Efficient Market Institutions: The Case of Wholesale Electricity Markets," Annals of Public and Cooperative Economics, 77, 4, 555-577.

Wolak, F. "An Assessment of the Performance of the New Zealand Wholesale Electricity Market", Investigation report, May 2009, available at
http://www.comcom.govt.nz/BusinessCompetition/Publications/Electricityreport/DecisionsList.aspx

Wolfram, C. (1999), "Measuring Duopoly Power in the British Electricity Spot Market", The American Economic Review, Vol. 89, No. 4.