

This is the pre-peer reviewed version of the following article: *"Inefficiencies and Market Power in Financial Arbitrage: A Study of California's Electricity Markets."* Borenstein, Severin, James Bushnell, Christopher R. Knittel, and Catherine D. Wolfram. *The Journal of Industrial Economics*, Vol. 56, No. 2 (2008): 347-378, which has been published in final form at <https://www.jstor.org/stable/25483412>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions.

Inefficiencies and Market Power in Financial Arbitrage: A Study of California's Electricity Markets

Severin Borenstein, James Bushnell,
Christopher R. Knittel, and Catherine Wolfram*

December 12, 2004

Abstract

In the three years following the restructuring of the California electricity industry, 1998 to 2000, power trading occurred in both a day-ahead market and a real-time market. Despite the fact that the power traded in these two major markets was for delivery at the same times and locations, prices differed significantly in many months. We consider several explanations for persistent price differences between the markets. We conclude that uncertainty about regulatory penalties for trading in the real-time market caused most firms to eschew arbitrage between the two markets. The few firms that did carry out this (risky) arbitrage did not find it profit-maximizing to eliminate the price differences. Due to California's electricity restructuring plan, the investor-owned utilities, which were the primary buyers of electricity, had little incentive to respond to the price differences. In the summer of 2000, however, when prices in both markets skyrocketed, we argue that the utilities' incentives changed in a way that was consistent with one utility's subsequent attempts to move demand between markets to minimize their purchase costs.

**Borenstein*: Haas School of Business, University of California, Berkeley, University of California Energy Institute, and NBER. Email: borenste@haas.berkeley.edu. *Bushnell*: University of California Energy Institute. Email: bushnell@haas.berkeley.edu. *Knittel*: Department of Economics, University of California, Davis, University of California Energy Institute and NBER. Email: crknittel@ucdavis.edu. *Wolfram*: Haas School of Business, University of California, Berkeley, University of California Energy Institute, and NBER. Email: wolfram@haas.berkeley.edu. We are grateful to Meredith Fowlie, Ryan Kellogg, Steve Puller, Erin Mansur and Celeste Saravia for outstanding research assistance and to Jonathan Berk, Al Klevorick, Rich Lyons, Frank Wolak and participants in NBER 2003 Summer Institute for useful discussions. This research was partially supported by the California Independent System Operator. A much earlier version of this paper was released in October 2001 as POWER Working Paper PWP-086, "Trading Inefficiencies in California's Electricity Markets"

1 Introduction

In product markets, it is well-understood that a firm that discovers a profitable market opportunity will generally maximize profits by producing less than the quantity that would drive price to the firm's marginal cost. The parallel analysis for financial markets suggests that if one firm sees a profitable trading opportunity, its trading will tend to reduce the profitability of the strategy, but it will not trade to the point that the marginal trade by itself breaks even. Put differently, the firm will have market power in the trading opportunity, though perhaps only briefly, and will take into account its effect on the strategy's profitability when it decides how much trading to do.

For two reasons, market power in trading opportunities has seldom been analyzed.¹ First, most opportunities are open to a large enough set of potential traders that the resulting equilibrium eliminates profits on the marginal trade. Second, even if only one firm can execute the trade, in most financial markets the firm can trade sequentially at different prices. By making sequential small trades, it can effectively price discriminate, trading until the profit on the marginal trade is zero. If either of these conditions hold, persistent profitable trading opportunities will not be observed in equilibrium.

In some cases, however, neither condition may hold. Institutional or legal constraints, or asymmetric information, may limit the number of agents that recognize a trading opportunity and are in a position to exploit it. Market rules or design may make it difficult for a strategic trader to sequentially price discriminate in its trading.² As a result, persistent price differences may be observed.

We argue that the California electricity market, which operated from 1998 through 2000, presented such a case. Two major markets accommodated trading of power for delivery at a specific location in a specific hour. Trading in the Power Exchange (PX) took place the day before delivery while trading in the Independent System Operator's (ISO) real-time market took place at the time of delivery. The products traded were identical, but we show that prices exhibited systematic and *ex ante* predictable differences that presented profitable trading opportunities. A variety of simple trading rules would have yielded positive returns that appear to more than compensate for the associated trading risk.

¹The role of corporate raiders in takeover battles is one exception (see Grossman and Hart, 1980 and Kyle and Vila, 1991).

²Zitzewitz (2003) considers the case of open-end mutual funds, where funds' decisions about how to price transactions can lead to profitable trading opportunities. He argues that agency issues allow arbitrage opportunities to persist in this case.

Once we establish the existence of significant price differences between the markets, we address the plausible explanations for this phenomenon. In financial markets, the most common explanation is risk aversion. Indeed, some research refers to such price differences as “risk premia” without addressing the alternative explanations we raise in this paper.³ We demonstrate that risk aversion is not a plausible explanation (a) because the direction of the premium shifts between buyers and sellers from month to month, (b) because the risk from trading on these expected price differences is highly diversifiable, and (c) because the magnitude of the gains are very large relative to the variance of returns (yielding very high Sharpe ratios). Also, a “peso problem” explanation – extreme outcomes that are possible, but not observed in the data set – is not applicable here.⁴ Regulatory constraints on prices – both floors and ceilings – limited the risk associated with such trading, and the most extreme prices permitted actually occur in the dataset.

Transaction costs are the other common explanation for persistent price differences. We show that the direct transaction costs of trading are too small to plausibly explain the persistence of predictable price differences of the magnitude we observe.

Transaction costs considered more broadly, however, could explain why more market participants did not take advantage of the apparent trading opportunity, and thus eliminate its profitability. We document that restrictions on speculative trading in these markets, and penalties for breaching those restrictions, were unclear. It is clear that some believed that the California ISO and PX had given tacit approval to such activities while others believed that it constituted a violation of ISO and PX rules and would eventually lead to punishment. In fact, some individuals who engaged in these trades have faced no repercussions. Conversely, these trades were included in the list of activities that were the basis for punishing traders at Enron.⁵

Even with a limited number of firms arbitraging the price differences, it appears that prices in the two markets were probably converging prior to the sudden change in market performance that occurred in summer 2000. The data before May 2000 suggest that the inefficiency dissipated as more traders came to understand the opportunity, but the statistical evidence of convergence is not strong.

As was widely reported, beginning in the summer of 2000, the California markets experienced drastic price increases. At the same time, the price differences between the

³See Longstaff and Wang (2004), who study the Pennsylvania-New Jersey-Maryland electricity market.

⁴Kaminsky (1993) studies a situation where the peso problem appears to be relevant.

⁵Unfortunately, none of the available data permit us to distinguish between fear of punishment and lack of understanding as possible reasons that other market participants did not exploit this market inefficiency.

ISO and PX widened, as the ISO price came to persistently and dramatically exceed the PX price. Average day-ahead prices in the Power Exchange were more than 15% below prices for the same product in the real-time market of the ISO and, by September 2000, prices in the ISO were higher than prices in the Power Exchange for over 70 percent of the hours.

We offer an explanation for the timing of this change that is consistent with a limited number of market participants that could individually influence the ISO-PX price difference. By summer 2000, the incentives of the major buyers – utilities in California – had changed. Due to the structure of regulation, utilities were much more motivated to reduce energy purchase costs than they had been in the previous two years. We present both data and documentary evidence showing that the largest buyer in the market, Pacific Gas and Electric (PG&E), attempted to reduce its purchasing costs by exacerbating the ISO-PX price difference in a way that reduced the price in the PX where PG&E carried out most of its purchasing.

While the story of California’s electricity debacle is itself interesting, the implications of our analysis extend beyond this particular market. Our analysis suggests that impediments that reduce the number of firms that can take advantage of profitable arbitrage trades can give market power to those that do engage in such trades and, thus, result in persistent price differences across markets. This weakens the ability of the forward market to provide an accurate signal of market conditions in the spot market. Our analysis also suggests a problem associated with using uniform-price auctions: they prevent firms from sequentially trading away inefficient price differences. It is also related to the ongoing policy debates about whether traders without physical positions in a market should be allowed to trade (see Saravia 2003).

The paper proceeds as follows. In the next section, we discuss the role and feasibility of arbitrage in electricity markets. In section 3, we describe the California forward and spot markets and some of the institutional rules that affected trading in them. Section 4 begins by laying out some simple statistics on the extent of market integration, and then presents results from more complete tests for market efficiency. In Section 5, we discuss several factors that could explain the price differences identified in Section 4. We present evidence suggesting that the price differences cannot be attributed to risk aversion, transaction costs or traders’ inability to learn about the profitable trading strategies in real time. We then go on to describe both statistical and documentary evidence on the behavior of individual firms and their beliefs about the markets. Our analysis suggests that the price difference persisted because some firms had market power in the trading required to push the prices

together.

2 Price Relationships in Electricity Markets

The California electricity market was one of many markets in which transactions occur on both a forward and a spot basis. In an efficient commodity market with risk-neutral traders, all contracts – forward and spot – for delivery of the good at the same time and location will, on average, transact at the same price. For instance, a contract signed on June 9 for delivery of 10 megawatt-hours (MWh) of power at 4pm on June 10 should bear a price that is an unbiased forecast of the spot price for electricity at 4pm on June 10. If the forward price differs systematically from the spot price, this can be due either to risk aversion on the part of some traders in the market or some impediment or cost that prevents full integration of the markets. In this section, we explain how one would expect the market to operate in the absence of risk aversion or impediments to integration.⁶

If there are no transaction costs and all traders are risk neutral, then the price at time $t - j$ for delivery of power at time t incorporates all information available at $t - j$ about the expected spot price of electricity at t . That is,

$${}_{t-j}P_t = E[{}_tP_t | \Omega_{t-j}] \quad (1)$$

where Ω_{t-j} is the information set available at $t - j$, the left subscript on price is the time at which the contract is traded, and the right subscript indicates the designated time for delivery of the power.

Equation (1) says that the forward price must be an unbiased predictor of the spot price. It also implies that the forward price incorporates all information available at the time it is in effect. The spot price can, and in most cases will, differ from this forward price, but the deviation, ${}_tP_t - {}_{t-j}P_t$, will have a distribution with a mean of zero and will be orthogonal to all information available at time $t - j$.

We can summarize this discussion by rewriting (1) slightly differently as,

$${}_tP_t = {}_{t-j}P_t + \varepsilon_t, \quad (2)$$

⁶Note that the discussion in this section relies on there being a sufficient number of competitive entities able to take advantage of any spot/forward price differences. It does not rely on perfect competition in the production of electricity. Even if considerable market power exists in the electricity supply, we would still expect no systematic price difference between forward and expected spot prices if both markets continue to support significant volume.

where ε_t is a random variable that has mean zero and is uncorrelated with Ω_{t-j} . That is, ε_t incorporates all of the shocks to the market that occur between $t - j$ and t . This implies, as has been the case in California and elsewhere, the variance of the spot price will be larger than the variance of the forward price.

It is worth noting that we do not assume any particular relationship with regard to the intertemporal patterns of electricity spot prices. Intertemporal arbitrage through storage is extremely costly in electricity markets, because electricity is not storable. While there are technologies to store potential energy, for instance by charging a battery or pumping water uphill, these methods are quite expensive and inefficient, losing more than 50% of the energy stored. For these reasons, it is common for electricity prices to fluctuate by as much as 300% or more within a day without creating profitable intertemporal arbitrage opportunities.

3 The California Electricity Market

During the first several years following electricity restructuring in California there were many avenues through which agents could sell or purchase wholesale electrical energy.⁷ In this section, we outline the California electricity market structure and discuss how traders could have profited from price differences across the various markets within this structure.

3.1 Forward Markets

Until December of 2000, most of the trading activity in California occurred on a day-ahead basis for hourly transactions. The California Power Exchange (PX) ran the largest of these day-ahead markets. The PX accepted supply and demand bids for each hour of the following day. Bids were submitted for the day-ahead market by 7am on the day before delivery. Day-ahead transactions were also reached through other scheduling coordinators (SCs) operating in parallel to the PX. Many of these daily transactions submitted by SCs in fact reflected longer term transactions that were nonetheless still required to be resubmitted to the Independent System Operator (ISO) on a daily basis.⁸

In a first-round calculation each day, the PX calculated day-ahead prices as if all bids and offers were in a common California-wide market. Limits on the capacity of electricity

⁷For more detailed descriptions of the various markets and their timing, see Bohn, Klevorick, and Stalon (1999) and Wolak, Nordhaus, and Shapiro (1998).

⁸The Automated Power Exchange (APX), for example, operated a 168 hour energy market on a rolling horizon.

transmission lines within California often necessitated further price adjustments. For purposes of transmission pricing, the California ISO system was divided into 24 zones.⁹ Two zones, comprising northern California (NP15) and southern California (SP15), contained the overwhelming share of ISO system demand. Most of the other “zones” were actually interface points between the ISO and surrounding utility systems. All SCs, including the PX, submitted their preferred energy schedules, including the location of all supply and demand sources, to the ISO by 10:00 am on the day before delivery. The ISO verified the feasibility of these aggregated schedules in light of transmission and other operating constraints. If these preferred schedules were infeasible because they would result in flows on transmission lines that exceeded the capacity of those lines, then the ISO ran an auction for the use of constrained transmission interfaces by utilizing schedule “adjustment bids” submitted by SCs.

The schedule adjustment bids effectively established each SCs willingness-to-pay for the use of a congested transmission interface. The preferred schedules were adjusted according to these bids, and a uniform price for the use of a congested interface was set at the usage value bid by the last SC whose schedule was adjusted. In this way all SCs that had scheduled transactions over a congested interface paid the same unit price for the use of that interface.¹⁰ The PX took these transmission prices and used them to determine zonal energy prices for all power traded in the PX. The difference between the PX price of two zones is equal to the ISO transmission charge for power shipped in the congested direction between those two zones.

In addition to the day-ahead markets operated by the PX and other SCs, schedule changes or revisions were permitted up to an hour ahead of the actual delivery time. The PX operated a “day-of” market (originally called an “hour-ahead” market) that allowed trades at a time closer to, but still many hours before the hour of operation.¹¹

⁹There were 23 zones when the ISO began operations in April 1998. The 24th zone, ZP26, was added during 1999.

¹⁰See Bushnell and Oren (1997) for a more detailed description of transmission pricing in the California market.

¹¹Until January 1999, the “hour-ahead” market operated on a rolling basis, with each market closing three hours before the hour of operation began. After that, it was operated as a “day-of” market, which was open three times per day, each time covering different blocks of hours that were from 5 to 12 hours in the future. This market was not particularly successful in either configuration: trading volumes were low and in more than 25% of all hours no transactions took place.

3.2 The Spot Market

The designers of the California market envisioned that the bulk of all transactions would be scheduled in one of the day-ahead or hour-ahead markets. However, since electricity is very costly to store and demand is inelastic, the ISO had to ensure that supply and demand remained in continuous balance by adjusting production. The ISO ran an “imbalance” energy market to handle these deviations. Like the PX, the imbalance energy market set a uniform price based upon the offer price of the marginal supplier.

The forward markets have often been described as “physical” power markets, in the sense that delivery of power was technically required to fulfill a transaction. During the first part of our sample period, there were no penalties explicitly associated with this delivery requirement. A market participant whose delivery or consumption of power deviated from its final schedule was simply charged, or paid, the ISO imbalance energy price for the hour in question depending on whether the SC turned out to be in a short or long position in real time. In this sense, the day-ahead and hour-ahead schedules were effectively financial forward positions, and the ISO imbalance energy market was the underlying spot market in which positions in these forward markets were resolved.¹²

Table 1 gives the relative volumes of the day-ahead, hour-ahead, and imbalance energy markets for the ten quarters between July 1998 and November 2000.¹³ The total volume across all three markets, which is summarized in the fourth column of the table and which serves as the denominator for the first three columns, reflects the total volume of power traded across all three markets.¹⁴ For much of the study period, day-ahead volume accounted for more than 92% of total volume, but this figure drops down below 90% during the fall of 2000. We discuss this pattern further in section 5. Most of the balance was taken up in the real-time market. As the last column of the table indicates, during the summer and fall of 2000, $RT > DA$, indicating that the power scheduled to be sold in the day-ahead market was less than the power that was eventually sold in real time, in nearly 90% of the hours. During high demand periods in the last few months of our sample period, the real-time imbalance energy market handled as much as 33% of total volume. This high level of real-time volume raised concerns about system reliability and prompted debates over the merits of further efforts to discourage real-time transactions.

¹²After August 19, 1999, this situation changed slightly, as discussed below.

¹³Data were not available for April-June 1998. Our analysis ends in November 2000, so 4Q 2000 reflects only October and November 2000.

¹⁴It is measured as $DA + |HA - DA| + |RT - HA|$, where DA reflects the power scheduled for delivery at the time of the day ahead market, HA reflects the power scheduled as of the hour ahead and RT is the actual real-time demand. Absolute values are used to reflect trading that reverses earlier positions.

The ISO imbalance energy market was not intended to be a full market, but instead to maintain reliability in the face of randomly fluctuating supply and demand. As such, consumers did not actively bid demand adjustments into this market. However, since there was, until August 1999, no explicit penalty for deviating from scheduled consumption, demand could passively take a position in the imbalance market simply by consuming more or less than it was scheduled to consume.

Suppliers could sell power in the imbalance energy market in three ways: by actively bidding into an imbalance energy market, by passively supplying more than was scheduled, or in conjunction with the supply of ancillary service, or reserve, capacity. Suppliers to the ancillary services markets submitted two-part bids: a “stand-by” capacity price for a given reserve service and an energy price paid in the event that the unit was actually called upon to generate. Producers that simply generated more than they were committed to provide were implicitly agreeing to take whatever price obtained in the imbalance energy market. Producers that bid into the imbalance energy market could choose to offer supply at a given price up to 45 minutes prior to the hour of production. Most suppliers of reserve capacity were also eligible to earn imbalance energy revenues. These (\$/MWh) energy revenues were in addition to (\$/MW) capacity payments earned by suppliers that commit to being available with varying response times. Each of these three avenues of supply – ancillary services, imbalance market bids, and excess or insufficient generation – involves a different degree of advance commitment. The bulk of ancillary service capacity was acquired by the ISO on a day-ahead basis, after the PX auction had closed. Suppliers who wished to sell imbalance energy through the ancillary service channel therefore had to submit offers a day before the service was actually used. Of course, they had the opportunity to earn revenues for their stand-by capacity, as well as any energy production. Suppliers opting for the imbalance energy channel could wait until 45 minutes prior to the hour of delivery before finalizing their offers. A supplier that simply generated in excess of its scheduled supply made that decision on a real-time basis, with no advance commitment.

The original ISO tariff specified that imbalance energy bids from all sources, reserve and imbalance energy providers, be treated equally and combined into a single supply offer curve. In practice, ISO operators sometimes skipped over low-cost energy bids from certain reserve sources due to concerns about depleting available reserves.¹⁵ Consequently, suppliers of some reserves earned no imbalance energy revenues even when their energy

¹⁵Suppliers of the most responsive form of reserve, regulation – which was used by the ISO for automated second-by-second adjustment to respond to imbalances at particular places on the grid – were never eligible to earn imbalance energy revenues. The capacity prices for this form of reserve consequently were significantly higher than those for other reserves, reflecting this lost revenue opportunity.

bid was below the imbalance energy price.¹⁶

A supplier that was scheduled to provide energy in one of the forward markets could also take a short position in the spot market either by offering to decrement its output through an imbalance energy bid or by simply generating less than its advance commitment. In the latter case, the supplier had to make up its production short-fall through a purchase on the imbalance energy market and was effectively a consumer in this market. A decremental supply bid in the imbalance energy market was an offer to buy out of an advance supply commitment. A supplier paid the ISO an amount equal to the imbalance energy price in exchange for not having to provide the energy that it had scheduled. By bidding a decremental energy bid, a supplier had the opportunity to set the imbalance energy price, and reserved the right to generate energy in the event that the imbalance energy price was set above its decremental bid.¹⁷

It is important to note that the ISO called upon decremental supply bids only when there was oversupply in real-time, and called upon incremental supply bids only when there was undersupply. The ISO did not attempt to arbitrage price gaps when there were some incremental bids that were lower than decremental bids. In other words, even though there may have been suppliers in the real-time market that were willing to pay more to buy out of their supply commitment than other suppliers required in order to fill that commitment, there was no mechanism for instituting these Pareto improving trades through the ISO. The actual magnitude of these inefficiencies has not been measured, but is an important empirical question.

3.3 Market Participants

Unlike more established commodity futures or forward markets, trading in the California electricity market was intended to be restricted to the actual producers and purchasers of electricity. As such, it was thought that trading would be restricted to hedging, and not speculative, activity. Although, in reality, speculative trades were certainly possible, institutional barriers largely restricted such activity to the actual “physical” market participants. After the market opened, further restrictions and institutional barriers were

¹⁶This practice most likely impacted the capacity prices of ancillary services more than it did the market-clearing imbalance energy price. In this paper we restrict our analysis to the relationship between the forward (PX day-ahead) and spot (ISO real-time imbalance) energy prices. The relationship of prices for reserves to these energy prices is an important topic for future study.

¹⁷In a perfectly competitive market with minimal transactions costs, we would expect that imbalance energy bids, in both the upward (incremental) and downward (decremental) directions, would be equal to marginal production costs.

applied in an effort to limit speculative trades. These efforts were motivated by a concern that such trades might destabilize the system and negatively impact the reliability of the network.

The primary function of the ISO was to maintain system reliability. All SCs that dealt with the ISO were supposed to present credible evidence of the ability to physically deliver and consume all power scheduled through the ISO system, as well as the specific locations where this activity would occur. There is no way to verify that a given level of consumption is completely realistic, but supply resources had to be specifically identified. Bids to provide ancillary services and imbalance energy from within the ISO system had to be linked to specific generation facilities. Since the ISO's ability to verify the availability of specific production sources was largely limited to its own control area, bids to supply energy and reserves from outside the ISO system faced much less stringent verification requirements.

The PX allowed more flexibility in both the eligibility of traders and the form of trades. PX market participants had to meet financial credit requirements, but did not need to control actual supply resources. Offers to supply through the PX took the form of "portfolio bids" that were required simply to be strictly upward sloping, piece-wise linear curves. In setting its unconstrained market price (*i.e.*, ignoring transmission constraints), the PX did not require the identity or location of specific production or consumption sources. Once the unconstrained price was set, suppliers to the PX had to identify their production source, either the specific generator within the ISO system or the transmission interface over which the supply will be imported. As with the ISO, there was no specific verification of the availability of import supply.

During a four year transition period starting in 1998, the three large investor-owned utilities (IOUs) in the California ISO system – Pacific Gas & Electric (PG&E), Southern California Edison (SCE), and San Diego Gas & Electric (SDG&E) – were required to meet the demand needs of their distribution systems through purchases in the PX. This requirement was intended to help ensure sufficient liquidity in the PX day-ahead market and to establish a transparent day-ahead price. Other market participants were free to participate in other day-ahead markets, or sign direct bilateral arrangements. Although there were roughly 60 firms trading in the PX, the three IOUs accounted for about 90% of the energy purchases. The PX itself accounted for about 87% of the total trading volume in the ISO system during the sample period.¹⁸

¹⁸See Bohn, Klevorick, and Stalton (1999), page 13.

Although the IOUs were technically required to purchase all their supply needs from the PX markets, the market process made rigid enforcement of this requirement both impractical and undesirable. It has been well documented that demand bids into the PX were downward sloping and in fact quite elastic over some price ranges.¹⁹ This is despite the fact that nearly all of end-use demand was incapable of receiving, let alone responding to, hourly price signals. Price-elastic demand bids in the PX clearly reflected strategic decisions by buyers to purchase in the ISO real-time imbalance energy market if the PX day-ahead price was too high. This was in part driven by the fact that the ISO imbalance energy market was subject to a price cap that was at times binding during our sample period, while PX prices were capped at a much higher level that was never binding. A large part of the elastic portion of PX demand bid curves reflected the fact that no firms were willing to pay more than the ISO energy price cap for power in a forward market, since that was the maximum allowable price in the spot market.²⁰

A large amount of energy supply in the California market was also committed to bidding into the PX day-ahead market. This energy was supplied by generation sources producing under regulatory or commercial arrangements that predated the restructuring of the California market. The price earned by these producers was set by the terms of their pre-existing “must-take” arrangements. This must-take supply was bid into the PX day-ahead market at a zero price.

In addition to the institutional and regulatory constraints on market participants, there were also differences in the transaction costs of dealing with the various markets. Two notable costs appear to have initially favored trading in the imbalance energy market over trading in the PX. Both the ISO and PX assessed trading charges on all volume in their markets. However, the PX charge, which was close to \$0.30/MWh, applied only to volume traded by firms that used the PX as their scheduling coordinator,²¹ while the ISO charge applied to all energy actually consumed in the ISO system *including that traded in the PX*. Thus, one could avoid the PX trading charge by not transacting using the PX as one’s SC, but there was no way to avoid the ISO trading charge.

The allocation of ancillary services costs until August 19, 1999 also provided incentives

¹⁹Bohn, Klevorick, and Stalton (1999).

²⁰The ISO imbalance energy price was capped at \$250/MWh when the market opened in April of 1998. The cap was raised to \$750/MWh on October 1, 1999, and subsequently lowered again in 2000 – to \$500/MWh at the beginning of July and \$250/MWh on August 8 – in response to the unprecedented price levels experienced during May and June of that year.

²¹The PX administrative charge applied to all volume traded by entities that used the PX as their scheduling coordinator, including volume in the ISO market. Thus, in order to avoid the PX administrative charge an entity had to use some SC other than the PX.

to avoid forward trades in favor of transacting on the imbalance energy market. Prior to August 19, 1999, all ancillary services costs were allocated based upon scheduled volume, rather than actual consumption. In August 1999, this was changed so that ancillary service costs from replacement reserve were assessed disproportionately to real-time transactions, rather than exempting such transactions. Thus, prior to August 1999, a firm that scheduled no supply or demand, but instead simply produced or consumed in real time without any notification, could have avoided paying for the reliability benefits provided by system reserves. These costs ranged between roughly 6-10% of the cost of energy. Despite these costs, the bulk of energy was still traded a day ahead, indicating that the institutional barriers, both real and perceived, and underlying benefits of forward trading outweighed the transaction cost differential.

4 Tests for Market Efficiency

In this section, we analyze the convergence of the ISO and PX energy prices. Monthly averages of these prices for the NP15 (North) and SP15 (South) zones are plotted in Figures 1a and 1b, and Table 2 provides summary statistics over the entire sample. Our sample period begins with the opening of the markets on April 1, 1998 and ends on November 30, 2000, the last month in which the PX could be considered fully functional.

We begin our analysis of price convergence by testing for systematic differences between the ISO and PX prices. Market efficiency implies that if agents are risk neutral and transaction costs are absent then, at the time the PX prices are determined, they should represent unbiased estimates of ISO prices. Formally, this implies that if PX prices are set at time $t - j$ then:

$${}_{t-j}PX_t = E[ISO_t | \Omega_{t-j}] \quad (3)$$

where Ω_{t-j} is the information set available at time $t - j$. Defining the realization of the ISO price at time t to be its expectation, conditional on the information set Ω_{t-j} , plus a random component ε_t , (*i.e.*, $ISO_t = E[ISO_t | \Omega_{t-j}] + \varepsilon_t$), we have:²²

$$ISO_t = PX_t + \varepsilon_t \quad (4)$$

This implication can be tested by estimating the model:

$$ISO_t - PX_t = \alpha + \varepsilon_t \quad (5)$$

²²We now suppress the $t - j$ presubscript on PX_t .

If the PX price is an unbiased forecast of the ISO price then $\alpha = 0$. We begin by estimating equation (5), allowing each month to have a different intercept, for zones NP15 and SP15.

There is good reason to think that shocks to the price differences between the PX and ISO prices were serially correlated, and empirical tests confirm that they were. Because the PX prices in a given day were all set at the same time, the errors in (5) are almost certain to be correlated across the hours in a day.

At 7:00 am each day PX participants submitted supply and demand bids for the 24 hour period beginning with the midnight-1:00 am hour of the following day. Because PX prices were determined in 24-hour “blocks,” shocks to either supply or demand (such as weather changes) that take place after PX prices were determined can have an impact on each ISO–PX price difference within a “block.” Since these shocks are serially correlated, the ISO–PX price differences also will be serially correlated, implying the standard errors obtained from ordinary least squares will be biased.²³ It is important to note that this institutional environment implies that *even in an efficient market* ISO–PX price differences are likely to be serially correlated.

Because of the timing of the PX market, the exact serial correlation structure that one would expect is quite complex. In the appendix, we describe the full correlation structure and two methods we used in attempting to estimate it. Unfortunately, neither approach proved tractable.

We have taken a simplified alternative approach. Instead of estimating a single version of (5) with all 24 hours of each day, one could estimate 24 separate hourly regressions. In this approach one regression would include all of the 848 hour-1 observations in our sample period, another all of the 848 hour-2 observations, and so on. By our discussion above, the regressions for the first 7 hours of the day would, in a fully efficient market, exhibit no serial correlation, while the regressions for hours 8-24 would have errors that follow an MA(1) process. This approach would yield consistent estimates of both the parameters and the standard errors; however it would be less efficient than a regression that pools the hours and takes into account the cross-hour correlations.

One drawback of this approach is that it yields 24 different sets of regression results, which would be difficult to interpret jointly. We have instead averaged the price differences for the early and later parts of the day, using one observation per day for each. An “early”

²³For example, if a summer day turns out to be hotter than was forecasted when PX prices were determined, the ISO–PX errors are all likely to be positive and therefore correlated.

observation is the average ISO-PX price difference for hours 1-6, while a late observation is the average ISO-PX price difference for hours 8-24. We drop hour 7, because it is the hour in which market participants generally submit bids. It is unclear whether the ISO-PX price difference during hour 7 would be correlated across days in an efficient market.

Thus, for each of the zones, sample periods, and specifications we analyze, we estimate an “early” regression and a “late” regression where the dependent variable is the average ISO-PX price difference in hours 1-6 and hours 8-24, respectively. In a fully efficient market, the early regressions would exhibit no serial correlation and the residuals from the late regressions would follow an MA(1) process. We estimate these equations using separate constant terms for each month, which indicate the average price ISO-PX differences for that month during the hours examined. Tables 3 and 4 present the results of this analysis for the North and South, respectively, including the Newey-West standard errors of the estimates, and the estimated price difference as a proportion of the average PX price during the same hours.²⁴ The shaded areas highlight p-values that indicate the estimates are significant at the 5% level. The coefficients demonstrate that PX prices were significantly different from ISO prices during the majority of months during 1998, except in the South during the later hours. After that, until May 2000, prices were less likely to differ consistently over a month and appeared to be converging. Beginning in May 2000, particularly in the North, price started to be consistently higher in the ISO. The magnitudes of the differences were also substantial, both overall and as a fraction of the ISO price levels.²⁵

4.1 Trading Rules Based Only On Prior Information

While the results presented thus far suggest that there have been significant differences between the PX and ISO prices in certain months, no distinct pattern emerges. For instance,

²⁴We estimate by OLS and report Newey-West standard errors (assuming an MA(1) error process for both early and late regressions), rather than using a GLS procedure that corrects for an MA(1) error process, because there is also substantial heteroskedasticity. The error variance is much greater during months of high average prices. Although one would expect there to be no serial correlation in the early period, this is rejected by the data. This is another indication that the market is not fully efficient.

²⁵As explained in section 3, since the beginning of the market, sellers in the real-time market could potentially earn not just energy, but also capacity reserve payments. The risk associated with being formally in the replacement reserve market, however, was that the unit would be called to generate only if the ISO needed to increment generation, so capacity reserve payments came with some risk. Until late August 1999, the buyers in real-time faced none of these costs because they were spread across all day-ahead scheduled transactions. Since August 1999, the costs of these reserve payments have been borne disproportionately by real-time buyers. An extreme interpretation of these rules would be to consider replacement reserve payments to be part of the full ISO price, so that the test of market efficiency would be to compare the ISO price plus replacement reserve price to the PX, ISO+R-PX. The results using ISO+R-PX as the price spread are largely consistent with those in tables 3 and 4, in large part because replacement reserve capacity payments were very close to zero during most periods.

in the first four months of trading, ISO prices were lower in both the North and South during both the early hours (1-6) and late hours (8-24), although the negative coefficients were only statistically significant in three out of the eight late-hour specifications. In the next four months of trading, most coefficients are positive, though there are several months when this is not true in the South during early hours. It is unclear from the results presented so far whether a trader would have been able to capitalize on the significant price differences we find. To gain insight on that question, we consider some simple trading rules and evaluate whether they would have made money in the first thirty-two months of the markets.

The first simple rule we evaluate assumes that a trader always makes sales or purchases in the market that would have been the most advantageous in the previous month. We assess whether our simple trading rule would have made money in the hands of a pure speculative trader, who, unconstrained by institutional barriers, could have bought in the market he believed would be less expensive and sell in the more expensive market. For instance, a trader following our rule in either zone would have used the estimates from April 1998, suggesting that the ISO prices were lower (both early and late), to sell in the PX and buy in the ISO during May 1998. We considered whether this strategy, implemented from May 1998 (we start here since there is no previous month's prediction for April 1998) through November 2000, would have made money.

We consider a very simple form of the test that uses the prediction from the previous month regardless of the statistical significance of the price difference. We test this by constructing a variable that is equal to one if the ISO price was higher in the previous month, so that the trading rule indicates that the trader should buy in the PX and sell in the ISO and negative one if the trading rule indicates purchases should be made in the ISO and sales in the PX.²⁶ Table 5 summarizes the coefficients and t-statistics from including this variable in a specification of equation (5) without any month dummies. The first row reports results from specifications that included all thirty-two months, while the remaining rows report tests during four separate time periods. Considering the entire time period, the t-statistics are greater than 2 in all specifications except the late hours in the South, suggesting that the simple trading rule produces positive and statistically significant profits for three out of four hour-zone combinations. For instance, the trader would have made an average profit of \$7.54 per MWh traded in the North during early hours. The results on the four separate time periods, however, suggest that most of the significant profit opportunities occurred at the beginning and the end of our sample period,

²⁶We assume that the trader trades an equal quantity each hour.

and that the market seemed to be converging before May 2000.

The bottom of Table 5 reports coefficients and t-statistics from tests of our trading rule at a weekly periodicity, where the trader commits to a trading strategy each week based on the price differences observed over the previous week. The results are similar to the monthly results, confirming that real profit opportunities existed in the first 32 months of the market, particularly during the first eight and last seven months of our sample period. Figure 2 plots the cumulative daily profits from our trading rules. The results suggest that a trader would have made considerable profits and would *never* have negative cumulative profits.

5 Explaining Forward-Spot Price Differences

The results thus far suggest that significant price differences persisted between the PX and the expected ISO prices, and that several simple trading strategies would have made money. This section considers several possible explanations for the differences. We find that two common explanations for the existence of forward-spot price differences even in completely competitive markets – risk aversion and differential trading costs across markets – are not consistent with the data. We then examine explanations in which some firms exercise market power in the arbitraging function.

5.1 Risk Aversion

Persistent differences between a forward and spot price could reflect risk aversion on the part of market participants. The conditions under which this will occur, however, are actually rather restrictive and the direction in which this would change the ISO–PX price relationship is ambiguous. So long as there are a significant number of competitive risk neutral buyers or sellers, these players would cause the forward and expected spot prices to converge, regardless of the degree of risk aversion among other participants.

In fact, risk neutrality, or near risk neutrality, may be a fairly accurate description of many of the players in the PX and ISO. The returns to speculative trades on the ISO–PX price difference had essentially no correlation with an investment in the market portfolio, so the risk associated with them could be diversified away. A regression of the ISO–PX price difference on a constant and the same-day return on the S&P 500 index cannot reject that the price difference has a β of zero.

Even if the risk associated with betting on the ISO–PX price difference is diversifiable, however, behavioral models of investor decisions suggest that some positive net-present-value investments will be passed over if the variance of the returns, relative to their mean, is high compared to alternative investments.²⁷ We compared the risk-return properties of speculation on the ISO–PX price differences to investing in an S&P 500 index fund by computing the Sharpe ratio for the trading rules discussed in the previous section.²⁸

Calculating the Sharpe ratio requires defining the time period over which returns are computed. We calculate the Sharpe ratio of the weekly trading rule using weekly returns and the monthly trading rule using monthly returns. In addition, we assume that the trader trades a total of one megawatt during each period (“early” or “late”) equally weighted across hours of the period. For example, a trader using the trading rule for Northern California ISO and PX prices in hours 8 to 24 would trade 1/17th of a megawatt each hour. Therefore, the weekly return is calculated as follows:

During periods where the trader buys in the PX and sells in the ISO:

$$\frac{\sum_{day=1}^{day=7} (\bar{P}_{ISO} - \bar{P}_{PX})}{\sum_{day=1}^{day=7} \bar{P}_{PX}} - \text{Weekly Prime Rate}$$

During periods where the trader buys in the ISO and sells in the PX:

$$\frac{\sum_{day=1}^{day=7} (\bar{P}_{PX} - \bar{P}_{ISO})}{\sum_{day=1}^{day=7} \bar{P}_{ISO}} - \text{Weekly Prime Rate}$$

Monthly returns are computed in an analogous manner. The Sharpe ratio is based on the mean and standard deviation of these returns.²⁹ As a comparison, we also calculated the Sharpe ratio for someone trading in the S&P 500 over the same time horizons. To calculate the earnings, we assume that a trader invests the same amount of money in the S&P as she would have invested in the California electricity market following our simple trading rule. For instance, during periods when the trader buys in the ISO and sells in the PX, she invests an amount equal to the average price in the ISO in the S&P 500 and then sells the shares at the end of the period.³⁰

²⁷See, for example, Chapter 7 in Lyons (2001).

²⁸The Sharpe ratio measures the ratio of the excess return relative to a benchmark security divided by the standard deviation of the excess return. See Sharpe (1994).

²⁹During two weeks in the south, during the early hours and one week in the north, during early hours, the average ISO price was negative at a time that the rule implied purchase from the ISO, so the trading rule would imply a negative investment. We drop these weeks from the Sharpe ratio calculation, since they imply in effect infinite positive returns. Dropping these observations biases downward the ratios.

³⁰We used the trading rules and prices for the late hours in the North to determine the amount invested in the S&P. The results are virtually the same if we use a different zone/period or just equal investments in all weeks.

Table 6 lists the Sharpe ratios for the weekly trading rules.³¹ The table illustrates that the returns from the trading rule were not the result of excess risk. In each period, the Sharpe ratios are considerably larger than those in the S&P 500. Speculating on the ISO–PX price difference had a much better return/risk ratio than investing in an S&P 500 index.

5.2 Estimation Risk

In demonstrating both that there were systematic patterns of ISO–PX price differences and that simple trading rules would have been profitable, we carried out our analysis using data from April 1998 to November 2000. With the limited data that market participants had, particularly early in the sample period, it could be difficult to infer the most profitable behavior. In any new market, it may take participants time to learn about how market rules, market fundamentals and their own behavior affect prices. One might then ask how rapidly a trader could learn of the profitability of a trading rule during the sample period.

To investigate this issue, we re-ran the tests for the profitability of trading rules on a rolling basis using only the data available at different points in the sample. For example, using the “last week” trading rule, we could ask how certain a trader could be of the profitability of the rule after, for example, five weeks of market operation. In that case the trader would have five weeks of data, of which the first week does not contribute observations because there is no prior week outcome on which to base trades. Running the regression for the 28 days in this sample (days 8 through 35), we would find a p-value of 0.14 on the test of the profitability of this rule. The level of certainty, however, increases (p-value drops) rapidly with a few more weeks of data. Figure 3a shows the p-value of the “last week” trading rule for the four zone/time combinations. In all four cases, it is clear that a trader considering this rule would have been more than 95% certain of its profitability by week 10, and would have been virtually certain of its profitability by week 20. With the monthly trading rule, inference of profitability is only slightly slower, as demonstrated by figure 3b, and the rule’s performance becomes less reliable for late-South near the end of the dataset.

5.3 Transaction Costs Within and Between Markets

Efficient price convergence between forward and spot markets can fail to occur if there are differential costs associated with contracting in either market. Absent other incentives,

³¹Sharpe ratios based on the monthly trading rules were very similar.

one would expect all volume to be traded in the lower cost market.

This may not occur, however, because either legal or political considerations constrain one or both parties, or because one or both parties receive other benefits from trading in the higher cost market, such as faster or easier settlements or more user-friendly bidding or dispatch rules. In that case, the price difference between the markets will depend on the incidence of the trading cost.

To illustrate this with a simple example, assume that the trading cost in the spot market is $C_s = 1$ and the trading cost in the forward market is $C_f = 2.50$. Absent other considerations, we would expect traders to abandon the forward market and make all transactions in the spot market. Now assume that buyers are constrained to buy the bulk of their power in the forward market, while sellers are completely indifferent between the markets.³² Sellers must be induced to trade in the forward market, so the net price they receive must be as high as in the spot market. If the buyer paid the trading charge in each market, then the price in the spot market would have to equal the price in the forward market in order to induce sellers to do business in the forward market. The buyers, however, would pay that price plus C_f . If the charge were assessed on sellers, then the price in the forward market would have to exceed the price in the spot market by 1.50, so that the sellers would be indifferent between the markets.

In reality, if both markets survive even though they have different direct trading costs, it is likely because both parties get some additional benefits from the higher direct-cost market. The difference in the direct trading costs is likely to then be a bound on the extent to which the prices in the two markets can differ. The incidence of the difference between the trading charges will be shared between the buyers and sellers depending on which side, on the margin, gets greater value from trading in the higher cost market.³³

The ISO-PX price differences that we've demonstrated are difficult to square with an explanation of differential trading costs for two reasons. First, the direction of the price difference changes numerous times during the period we study while there is little evidence that the relative cost of transactions in the two markets changed significantly, and no evidence that changes in the forward premium or discount is associated with changes in relative transaction costs. Second, the price differences that began in May 2000 are far in excess of the magnitudes of transaction costs. We know of no evidence that transaction

³²This is *not* intended to be a characterization of the California market. The actual incentives in the California market were much more complex.

³³It is possible that traders on one side will strictly prefer the market with the lower direct trading costs, even before accounting for the trading costs, in which case the equilibrium price spread between the markets could be greater than the difference in trading costs.

costs in either market changed substantially at the beginning of summer 2000, and, as we point out in Section 4 (footnote 25), the trading costs are so small that our results remain largely unchanged when we adjust the ISO price to reflect the trading costs imposed on buyers in that market.

5.4 Market Power in Arbitrage and Barriers to Entry

We have established that (1) there were profitable (in expectation) risky arbitrage opportunities between the ISO and PX power markets using simple trading rules, (2) that the risk associated with these trades was not great compared to the potential return and was diversifiable, (3) that it should have been apparent to traders early in the life of the market that these arbitrage opportunities existed, and (4) that transaction costs do not seem to be a viable explanation for the persistence of these price differences. Thus, it seems unlikely that outcomes we observed could be explained as part of a competitive financial market for power. In this section, we discuss evidence on the market power and incentives of three types of parties that could have profited from the ISO–PX price differences: electricity buyers, electricity sellers and arbitrageurs.

Electricity Buyer Market Power

Among the “physical” players in a position to take advantage of ISO–PX price differences were the three utilities that accounted for most of the demand in the market. The utilities were expected to purchase the bulk of their demand (as forecasted a day ahead of time) through the PX and use the ISO to cover imbalances caused by last-minute demand shocks. Though no attempt was ever made to penalize the utilities for using the ISO market, there was a common perception that they should not make significant purchases of forecastable demand in real time.

Prior to the spring of 2000, the utilities also had little incentive to attempt to reduce their aggregate purchase costs by moving purchases between the markets, but that changed around May 2000. To understand why, one needs to understand the Competition Transition Charge (CTC). The CTC was a surcharge on all power that was designed to allow the utilities to recover losses incurred when their capital stock of generation plants was effectively devalued by the deregulation process. Each utility was assigned an aggregate amount that it was allowed to recover through the CTC. Furthermore, each utility was allowed to collect a CTC surcharge on power sold to all customers in its service area until either it recovered its stranded costs or until March 2002, whichever came first. The CTC surcharge, however, was not a fixed amount. Instead, the law fixed the retail price utilities

charged for energy (at about 6 cents per kWh equal to \$60/MWh). The difference between the retail revenue earned at the fixed retail price and the wholesale cost of electricity was the CTC payment to the utility. Thus, when wholesale prices were low, the CTC recovery payment was high. This was the situation for most of the first two years of the market, and most observers believed that the utilities would collect their stranded cost prior to the March 2002 cutoff. In fact, SDG&E, the smallest of the utilities, did complete its CTC collection in June 1999, after which retail customers, not SDG&E, were responsible for changes in wholesale purchase costs.³⁴ So long as the utilities believed that the March 2002 cutoff would not be binding, they had little incentive to try to minimize their purchase cost.³⁵

All that changed in May/June 2000. As is now well known, the California electricity market went through dramatic changes beginning in May 2000 as average wholesale prices more than doubled from their highest previous level. The overall price shock was due in part to unprecedented increases in the prices of two inputs to electricity generation: natural gas tripled in price and NOx pollution permits increased by a factor of 50 beginning in the spring of 2000 (see Figure 4). These increases had the greatest effect on the costs of the units that were marginal at peak times – which in California are low-efficiency, high-polluting, gas-fired plants – so they increased the steepness of the industry supply curve particularly at the higher output levels. The steeper supply curve, in turn, increased the incentives of some sellers to exercise market power at peak times. Thus, the cost changes, and accompanying changes in sellers’ production incentives led to drastic price increases.

When wholesale prices increased to well above \$60/MWh in June 2000, PG&E and SCE began collecting “negative CTC payments.” In other words, they were losing money on each kilowatt-hour sold, which depleted their CTC balance and made it much more likely that the March 2002 cutoff for CTC collection would have been binding. If the cutoff had become binding, then utility shareholders would have been the residual claimants on any reduction of power procurement costs prior to March 2002. Thus, the increase in price levels gave the utilities stronger incentives to lower their procurement costs.

Though the three utilities were major buyers in the power market, their market shares did not give them monopsony power in the traditional sense, since the utilities in their

³⁴Actually, in late August 2000, the State passed legislation reimposing a fixed rate on SDG&E, but also made it clear that SDG&E would be made whole for any losses it suffered as a result of this change. See Bushnell & Mansur (forthcoming) for further details.

³⁵Reductions in the wholesale price would only have sped up collection of their CTC and would not have increased the total amount collected. The only benefit from reducing the wholesale price, therefore, was the forgone interest from collecting this money sooner. Given that interest rates were low, this was a weak incentive.

role as distributor had no control over the aggregate quantities of end-use consumption. They did, however, have discretion over the market in which the power was purchased.

Because the supply curve in the PX was upward sloping (see figure 5 for an example³⁶), if a utility shifted some of its purchases from the forward market to the spot market, and this shift was not anticipated by suppliers, it would lower the forward price.³⁷ In a very simple model, the move would not change the ISO real-time price because the ultimate level of demand would not be altered; the intersection of the market level (or “physical”) supply curve and the demand would be unchanged.

This logic is depicted graphically in figures 6 and 7. In figure 6, the expected total retail electricity demand is represented by the inelastic demand curve $\bar{Q}$ and the market level supply curve is represented by the upward-sloping supply curve S . The market is in equilibrium with the forward price equal to the expected spot price and no net transactions occurring in the real-time market. Deviations between the forward and spot prices occur only when the inelastic demand differs from its forecasted level. For example, if the real-time demand level is lower than forecasted, then the net quantity transacted in the spot market will be negative and the market will move down the market supply curve resulting in a spot price that is lower than the forward price. Conversely, if there is a positive shock to demand, the spot price will be greater than the forward price.

In figure 7, there is an unanticipated decrease in the forward market demand representing the decision of a buyer (such as one of the utilities) to shift γ units of demand from the forward market to the spot market. This is accompanied by an unanticipated increase in the spot market demand. The forward price is reduced. Because final demand and supply remain unchanged, the spot market price is unchanged. Alternatively, if some generation is available only at higher cost in real time – for instance, because there is a (possibly implicit) penalty for large sales in the real-time market – then this strategy could increase the ISO price. Still, the net impact could be to reduce procurement costs if the

³⁶Note that nearly 20,000 MWs of supply was “must-offer” meaning that the owners had to bid their power through the PX at a price of \$0/MWh.

³⁷Why the PX supply curve was upward sloping is a question we don’t attempt to answer here. If all bidders had symmetric expectations about the spot price, were risk neutral and faced no penalty for using the spot market, the PX would effectively be a financial forward market and participants would stand ready to buy or sell at their expected spot price with infinite elasticity. Risk aversion on the part of some buyers and sellers would lead to upward-sloping supply and downward-sloping demand in the PX - even though the presence of other risk-neutral firms could be expected to eliminate price differences. Similarly, a penalty or tax on real time transactions would do so as it would in a sense move “physical” transactions into the forward market. In order to have this effect the penalty would have to be non-linear in the size of the real-time transaction. This includes policies that ignore a modest reliance on the real-time market but react to significant real-time volumes.

savings from the price reduction on a large purchase quantity in the PX were greater than the increased cost on the price increase on a comparatively small purchase quantity in the ISO.

There is strong documentary and empirical evidence that PG&E attempted just such a strategy by moving demand out of the PX. For instance, in a subsequent regulatory filing, they described this strategy and explained how, “paying a higher price in the ISO market for the incremental portion of total load [demand] was more economical than bidding higher prices into the PX market and paying a much higher price in the PX for every MW purchased” in that market.³⁸

Figure 8 helps identify the timing of PG&E’s attempt to move demand out of the forward market. It plots the fraction of each of the three utilities’ total end use demand that they bid into the forward market at the eventual ISO price.³⁹ SCE and SDG&E both consistently bid 70%-80% of their demand into the PX, while the fraction that PG&E bid in began declining in May 2000 and fell from averaging about 80% in January-April 2000 to about 50% in August through November 2000. Figure 9 highlights differences among the demand curves the utilities bid in the PX. PG&E and SDG&E both bid downward sloping demand curves into the PX, while SCE bid nearly completely inelastic demand curves. As market prices rose through summer 2000 for the reasons discussed above, the market equilibrium shifted along PG&E’s and SDG&E’s demand curves. SDG&E offset this by shifting their demand curve out between June and August 2000. PG&E did not do this; in fact it shifted its demand slightly inward. As a result, PG&E purchased less and less through the PX market.⁴⁰

The relationship between ISO and PX prices changed markedly in May 2000, consistent with a change in PG&E’s buying strategy. Beginning in May 2000, PX prices in the North averaged substantially below ISO prices (see figure 1a), with the difference becoming still larger in July 2000. Prices in the South exhibited much less change; the PX prices averaged only slightly lower than the ISO in SP15 (see figure 1b).

All of our evidence suggests that PG&E pursued the monopsony strategy but SCE did not. There is no record indicating why they did not, but it is possible that the presence of

³⁸See PG&E (2002), p.009.

³⁹We use the ISO price in order to control for changes in cost and supply conditions that affect the “relevant” part of the utilities’ demand curves. We use the ISO price in the north for PG&E and the ISO price in the south for the other two utilities.

⁴⁰The abrupt flattening of the utilities’ August demand curves at \$250, most notable for PG&E, reflects their rational response to the lower ISO price cap of \$250, which became effective August 8th, 2000. Even though PG&E’s PX demand flattens more than the other utilities’ near the ISO price cap, and the price cap was sometimes binding, the pattern in figure 9 is very similar if we drop those hours.

SDG&E as an additional buyer in the south made it harder for them to move the PX price. Also, SCE could free ride off of PG&E's strategy as it benefitted from the lower PX prices without having to pay the higher ISO price on any of its own purchases. Because SDG&E had completed its collection of the CTC by 2000 and thereafter passed its purchase costs through to retail customers, SDG&E faced much less incentive to minimize its purchase costs.

Market Power of Arbitrageurs

The strategy discussed in the previous section relies on the shift in demand across markets being unanticipated. It is clear that at least some firms operating in the market knew of the predictable forward/spot price differences and devised strategies to arbitrage the price differences. Since firms that had no physical supply nor served any end-use demand were not supposed to trade in the ISO market, pure arbitrageurs were technically not allowed.⁴¹ However, several parties, most famously Enron, traded large amounts beyond their physical positions in the markets.

Enron's activities illustrate the possible strategies. Enron's physical presence in California included the power from the generation assets that their subsidiary in Oregon, Portland General Electric, regularly exported to California and the obligations that their subsidiary Enron Energy Services (EES) had to meet the demand of several large buyers, including the University of California, who had opted to leave the utilities and buy power from EES. To take advantage of the fact that ISO prices were consistently higher than PX prices after May 2000, EES could overstate their demand in the PX market. They could then sell to the ISO the difference between what they bought in the PX and what they actually needed to meet demand of customers like the University of California. Enron internal memos released to the FERC described this as the "Fat Boy" trading strategy (see Yoder and Hall, 2000). Other documents describe the reverse strategy—a party could schedule more generation than it intended to provide through the PX and then buy it back through the ISO when the PX price was expected to be higher than the ISO—as "Thin Man."

Restricting the market to physical parties created one barrier to entering the ISO–PX arbitrage business. In addition, there was ambiguity about whether arbitrage trades violated ISO and PX rules. Among the parties that were allowed to trade in both the ISO and the PX there could well have been either differences of opinion about how to interpret

⁴¹Traders could get around the restrictions to only make physical trades by scheduling power to be supplied or consumed at 'import' interfaces with neighboring regions. The California ISO had limited ability to monitor the production or consumption activity outside its own control area.

the rules or different valuations of the risks associated with skirting the rules. Rules for traders in the PX and ISO were collected in their Market Monitoring and Information Protocols (MMIP), and included general prohibitions against “gaming” and “anomalous market behavior,” which it defines to include, “bidding patterns that are inconsistent with prevailing supply and demand conditions” (California ISO, MMIP, 2.1.1.4). Enron was aware of these provisions, as they are described in Yoder and Hall (2000). By June 2000, the parties had reason to believe that the ISO would not penalize overscheduling since they were concerned about PG&E’s underscheduling. One party, Reliant, claimed that the ISO took actions to assist Reliant in overscheduling demand through the PX (see FERC, 2003, p. VI-24).

After the release of the Enron memos, FERC initiated an investigation of trading strategies in the California markets. Two interesting facts have come out of this investigation: Arbitrage profits were concentrated between Enron and one other large firm with possible ties to Enron, and Enron took steps to coordinate arbitrage trades among market participants. This suggests that the ambiguity in the rules governing arbitrage and the restriction of physical players may have been sufficient to give Enron market power in the arbitrage market.

Information on the concentration of arbitrage trades comes from an analysis that the ISO staff did of which parties benefited from the “Fat Boy” strategy (California ISO, 2003). The report identifies hours in which firms scheduled substantially more than their actual demand (specifically, when forecast exceeded actual by more than 13% or by more than 25MWs) and calculated the profits they earned on the excess. Though 33 parties earned more than \$20,000 through overscheduling between January 1 and October 1, 2000, Enron trades accounted for about 28% of the volume of “Fat Boy” activity, and Powerex, the marketing arm of BC Hydro, accounted for 15%. The Enron memos and other internal documents claim that Enron had assisted Powerex in making Fat Boy transactions (see Yoder and Hall, 2000, p. 2), suggesting that they may have shared their strategy with them. Powerex has subsequently denied having such a relationship with Enron (Peterson, 2002).

Internal Enron documents also suggest that Enron was attempting to coordinate arbitrage trades across parties, and had implemented specific profit-sharing rules with certain parties. FERC (2003) cites sections of the Enron Services Handbook, which appears to contain instructions to traders. Under a section, “Who do you call and what action to take?” there are six parties listed under the Fat Boy transaction and instructions to the Enron traders to tell them to “fake or increase load [demand]” in the PX. The sharing

rules for four of the six parties are straight 50-50 splits of profits or losses, while the other two have more complicated sharing rules. FERC (2003) challenges the legality of these arrangements because they were wholesale power trades that had not been filed with or approved by FERC.

The presence of arbitrageurs would mitigate the success of PG&E's strategy and thus reduce the price difference, but would not drive it to zero. Given the uniform price auction in both the forward and spot market, profits from the arbitrage trades were equal to the price difference times the amount traded, implying that the profit maximizing trade would reduce the price difference, but not eliminate it. This is depicted in figure 10. We begin with the same monopsony trade as in figure 7, but allow for an arbitrageur, with no physical assets, to respond to the price difference by buying q_a MWh of power in the forward market. Because the arbitrageur cannot take delivery of the power, she is forced to reverse her position in the spot market, which effectively requires her to bid q_a MWh in the spot market at a price of zero.

By increasing the forward demand, the arbitrage trade increases the forward price, but does not influence the spot price, since the ultimate demand level is unchanged. Therefore, the price difference is reduced, but not eliminated. The optimal amount of the arbitrage trade will depend on the shape of the supply curve and the size of the monopsony trade.⁴² In addition, greater competition in arbitraging would result in greater aggregate arbitrage trades. At some point, we would expect the competitive pressures to eliminate the price difference. In the case of the California electricity markets, the evidence suggests that the arbitrage market was sufficiently concentrated such that price differences remained.

Seller Responses to Monopsony Strategy

Besides firms who took advantage of the price difference using mostly financial trades, we would expect that firms with net sell positions in California would respond to PG&E's demand reduction strategy by selling more of their output in the ISO market. The deregulated suppliers, known as exempt wholesale generators (EWGs), were primarily on the sell side of the wholesale market. We have analyzed the PX bid curves of the EWGs; while there is some indication that some of them were selling less through the PX especially by the fall of 2000, the data are too noisy to decipher an overall trend towards shifting supply

⁴²To see this, suppose the supply curve was quadratic in quantity, $S = a + bQ + cQ^2$. In this case the forward price would be $p_f = a + b(Q_f + q_a) + c(Q_f + q_a)^2$, while the spot price would be $p_s = a + bQ_s + cQ_s^2$. The monopolist arbitrageur would maximize $(p_s - p_f)q_a$, which yields a solution of $q_a^* = \frac{1}{2b+4cQ_f} (bQ_s - bQ_f + cQ_s^2 - cQ_f^2)$; this is an increasing function of c , the curvature of the supply function, and $Q_s - Q_f$.

out of the PX. For one thing, as figure 4 points out, input costs were rising substantially over the same time period, so even if the firms weren't responding to PG&E's strategy, they would have shifted their supply curves up. Also, a number of the EWGs had small demand positions they were trying to cover, so they often bid both supply and demand curves. Analyzing the net sales they bid is not always meaningful because we don't know the true underlying demand they were bidding to serve. Finally, some of the suppliers were obliged under the terms of long-term contracts to bid through the PX, so the correct analysis of their bidding in response to PG&E's strategy would have to account for these contract positions.

There is documentary evidence suggesting that some of the suppliers moved into the spot market. For example, Williams Energy in its FERC disclosure in the Enron proceeding (see Williams 2002) states that "unlike Enron, Williams has dispatch rights to generation assets in California that enable it to sell power into the Real Time market. Thus Williams does not have the incentives that were apparently driving Enron to schedule demand in the Day Ahead schedule which it could cut to sell energy in the Real Time market. Williams could simply sell its own generation in the Real Time market."

It is also clear from the forward-spot price differential that any seller response was not sufficiently large to drive the price difference to zero. As explained earlier, however, with the uniform price auctions for power, other parties *profiting* by offsetting PG&E's strategy would not want to trade until the price difference disappeared.

The migration of volume out of the PX undermined its viability. The PX suffered another blow, which was ultimately fatal, when the Federal Energy Regulatory Commission announced a preliminary ruling in November and final decision in December 2000 that required the three California utilities to stop selling their own power through the PX. Volume in the PX plummeted in December 2000 and January 2001. On January 31, 2001, the California Power Exchange ceased operations of a day-ahead electricity market.

6 Conclusion

It is sometimes taken for granted that profitable arbitrage opportunities should not persist in reasonably functioning financial markets. However, capital restrictions, limited information, and other more explicit barriers to trading can limit the number of traders. This can result in profitable trading opportunities that can be sustained for long periods.

In this paper, we have studied one such trading opportunity in the California electricity

market: the interaction between the day-ahead market of the PX and the real-time market of the ISO. Although these markets played very different institutional roles, and operated under quite different market rules, they were fundamentally markets for the same product, a unit of electrical energy to be consumed in a given hour at a given location in the network. The level of price convergence between these two markets is therefore an indicator of the ability of firms to overcome informational and institutional barriers to efficient trade.

Our work establishes that significant price differences existed between the PX and the ISO during several periods in the PX's 32 months of operation, particularly during the last seven months. Furthermore, it appears that some trading strategies with positive expected return existed. The strategies were risky, but the risk was highly diversifiable and the expected return was quite large in comparison to the risk.

From our analysis, one main explanation emerges for the persistent and unprecedented differences between day-ahead and real-time prices during the summer of 2000. Several of the traders in the market appear to have exercised market power in their trading positions. We documented an attempt by the largest buyer of electricity, PG&E, to exercise a form of monoposony power over these markets. We also presented data and internal company memos suggesting that Enron made arbitrage trades to take advantage of the price differentials, and we demonstrated that an arbitrageur like Enron with market power may not want to trade until the price differential disappears.

More traders, in hindsight, should have been able to profit from these price differences. With enough traders, any attempts to induce or profit from persistent price differences between markets should be eventually undermined by the forces of arbitrage, and we should observe the price differentials disappearing. Why we did not observe this in the California markets is open to speculation. One plausible explanation is that the market rules provided a barrier to entry for the types of trades that Enron used. Other potential arbitrageurs may have been deterred by the market rules prohibiting traders from misrepresenting their physical positions. It is also possible that the generators, who were earning extremely large profits as the overall market price level rose, perceived political constraints to earning even more money from the markets by flouting the rules and shifting more sales into the ISO. A second explanation is that in the complex and frequently-disrupted California electricity market, too few participants learned about the price differences fast enough for prices to converge before the eventual demise of the day-ahead market. Unfortunately, the documentary evidence provides much better insight into why some parties took certain actions than into why others *did not* take those same actions.

References

- [1] Bohn, R., A. Klevorick, and C. Stalton (1999). "Second Report on Market Issues in the California Power Exchange Energy Markets." Prepared for the Federal Energy Regulatory Commission.
- [2] Bushnell, J. and E. Mansur (forthcoming). "Consumption under Noisy Price Signals: A Study of Electricity Retail Rate Deregulation in San Diego." *Journal of Industrial Economics*.
- [3] Bushnell, J. and S. Oren (1997). "Transmission Pricing in California's Proposed Electricity Market." *Utilities Policy* vol. 6, no. 3.
- [4] California ISO, Department of Market Analysis (2003). "Supplemental Analysis of Trading and Scheduling Strategies Described in Enron Memos." Available at <http://www.caiso.com/docs/2003/06/18/2003061806053424839.pdf>, June 2003.
- [5] California ISO, Market Monitoring and Information Protocol. Available at <http://www.caiso.com/docs/2002/05/16/2002051610063028441.html>.
- [6] Dynegy (2002). "Responses of Dynegy Power Marketing Inc. to Request for Admission. Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices." Federal Energy Regulatory Commission Docket No. PA02-2-000. Available at <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2.asp>. May 22, 2002.
- [7] El Paso (2002). "El Paso Electric Company Response to Data Request. Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices." Federal Energy Regulatory Commission Docket No. PA02-2-000. Available at <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2.asp>. May 22, 2002.
- [8] Federal Energy Regulatory Commission (FERC) (2003). "Final Report on Price Manipulation in Western Markets. Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices." Federal Energy Regulatory Commission Docket No. PA02-2-000. Available at <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2.asp>. March, 2003.
- [9] Grossman, S. and O. Hart (1980). "Takeover Bids, the Free Rider Problem, and the Theory of the Corporation," *Bell Journal of Economics* vol. 11, Spring, pp. 42-64.
- [10] Kaminsky, G (1993). "Is There a Peso Problem? Evidence from the Dollar/Pound Exchange Rate, 1976-1987." *American Economic Review* vol. 83, June, pp. 450-472.

- [11] Kyle, A.S. and J.-L. Vila (1991). "Noise Trading and Takeovers," *RAND Journal of Economics* vol. 22, Spring, pp. 54-71.
- [12] Longstaff, F.A. and A.W. Wang (2004). "Electricity Forward Prices: A High-Frequency Empirical Analysis," *Journal of Finance*, vol. 59 (4), August, pp. 1877-1900.
- [13] Lyons, R. K. *The Microstructure Approach to Exchange Rates*. Forthcoming, MIT Press: Cambridge, MA.
- [14] Newey, W. K., and K. D. West (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* vol. 55, pp. 703-708.
- [15] Peterson, K. (2002). "Affidavit of Ken Peterosn on behalf of Powerex Corp. ," in Docket No. PA02-2-000 before the Federal Energy Regulatory Commission. Available at <http://www.powerex.com/pdf/KGPAffidavit2.pdf> May, 2002.
- [16] PG&E (2002). "Responses of Pacific Gas & Electric Company to Data Request. Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices." Federal Energy Regulatory Commission Docket No. PA02-2-000. Available at <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2.asp>. May 22, 2002.
- [17] Saravia, C. (2003). "Efficient Futures Markets and Production Efficiency: The Case of the New York Electricity Market," mimeo, UC Berkeley.
- [18] Sharpe, W.F. (1994). "The Sharpe Ratio," *Journal of Portfolio Management* vol. 21, Fall, pp. 49-58.
- [19] Williams (2002). "Responses to Commission Staff's Data Request of Williams Energy Marketing & Trading Company. Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices." Federal Energy Regulatory Commission Docket No. PA02-2-000. Available at <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2.asp>. May 22, 2002.
- [20] Wolak, F.A., R. Nordhaus, and C. Shapiro (1998). "Preliminary Report on the Operation of the Ancillary Services Markets of the California Independent System Operator (ISO)." Available at www.ucei.berkeley.edu/ucei/filings.html.
- [21] Yoder, C. and S. Hall (2000). "Traders' Strategies in the California Wholesale Power Markets/ISO Sanctions." Memorandum to Richard Sanders. Available

at <http://www.ferc.gov/industries/electric/indus-act/wem/pa02-2.asp>. December 6, 2000.

- [22] Zitzewitz, E. (2003). “Who Cares About Shareholders? Arbitrage-Proofing Mutual Funds,” *Journal of Law, Economics, & Organizations* vol. 19, October, pp. 245-280.

Appendix: Correlation Structure of Hourly Price Differences

The timing of PX and ISO trading is described in the text. Because of the timing of the PX market, the exact serial correlation structure that one would expect is quite complex. We describe this below and then discuss two different estimation approaches. Neither of these estimation approaches were tractable. This led us to adopt the simplified approach described in the text.

Let $t = 1$ represent the beginning of an arbitrary day (*i.e.* the 12:00 am–1 am hour). The PX prices for $t = 1, \dots, 24$ are set conditional on the information set available at the time the PX supply and demand bids were made, which is likely to be between 6:00 am and 7:00 am (hour 7) of the previous day, or at $t = -17$. Therefore the information available at the time the PX market closed is Ω_{-17} .⁴³ At time $t = 6$, PX prices are calculated for hours 25 to 48, but these prices are conditional on the information set Ω_6 . The process continues *ad infinitum*.

The consequence of this process when econometrically modeling the difference between ISO and PX prices is that the serial correlation among the error terms is expected to be of varying lengths, depending on the time of day of the observation. A shock that causes the difference between the ISO and PX prices to diverge during the $t = -16, \dots, 0$ time frame may continue to impact this difference for hours $t = 1, \dots, 24$ (likely at a decreasing rate). However, since PX prices at time $t = 25$ are set conditional on an information set that takes into account any shocks that preceded $t = 6$, an efficient market would imply that a shock at $t = -16, \dots, -1, 0, 1, \dots, 6$ should not be correlated with the difference between the PX and ISO price at $t = 25$. Also, it is likely that the level of correlation between prices set on the same day will be larger than correlations between prices on successive days. For instance, the correlation between the error in hour 1 and the previous hour (hour 24 from the previous day) is likely to be smaller than the correlation between hour 2 and hour 1, because the latter were determined under the same information set. Thus, both the number of lagged hours with which an error is likely to be correlated and the size of that correlation with each lag will vary by hour of the day.

We can write the price difference as a moving average process that explicitly recognizes the correlation with earlier hours. For each hour, we would expect correlation back to the time at which the price was set for that hour, that is, 6am–7am of the previous day. We can therefore write the process as:

⁴³Because supply and demand bids may take some time to be formulated, we make the assumption that they are made during the time period of 6:00 am to 7:00 am, and are therefore set conditional on the information set available at $t = 6$.

$$\begin{aligned}
ISO_1 - PX_1 &= \alpha + \varepsilon_1 + \sum_{i=1}^{18} \theta_{1,i} \varepsilon_{1-i} \\
ISO_2 - PX_2 &= \alpha + \varepsilon_2 + \sum_{i=1}^{19} \theta_{2,i} \varepsilon_{2-i} \\
&\vdots \\
ISO_{24} - PX_{24} &= \alpha + \varepsilon_{24} + \sum_{i=1}^{41} \theta_{24,i} \varepsilon_{24-i} \\
ISO_{25} - PX_{25} &= \alpha + \varepsilon_{25} + \sum_{i=1}^{18} \theta_{25,i} \varepsilon_{25-i} \\
&\vdots
\end{aligned} \tag{6}$$

Unfortunately, our attempts to estimate a model with varying moving average components have not led to convergence. One can obtain consistent estimates of the standard errors from OLS estimation based on the Newey-West (1987) procedure. This requires modifying the standard Newey-West estimator to account for the variable lengths of correlations. Unfortunately, the covariance matrix of the modified Newey-West estimators is not guaranteed to be positive semi-definite, and indeed yielded imaginary standard errors for some specifications.

Table 1

Quarterly Average Trading Volumes in the Forward and Spot Markets

Quarter	% Day Ahead	% Hour Ahead	% Real Time	Volume Traded (MWh)	% Hours in which $RT > DA$
3Q 1998	92.3%	3.9%	3.8%	29,674	83.8%
4Q 1998	94.1%	3.5%	2.4%	24,531	77.4%
1Q 1999	91.9%	4.9%	3.2%	24,489	87.4%
2Q 1999	92.1%	4.9%	2.9%	25,537	80.0%
3Q 1999	91.9%	4.9%	3.3%	29,573	72.9%
4Q 1999	92.8%	4.2%	3.1%	26,876	67.4%
1Q 2000	90.9%	4.9%	4.2%	25,885	83.3%
2Q 2000	90.2%	4.9%	4.9%	27,682	88.9%
3Q 2000	89.4%	4.9%	5.7%	30,108	89.5%
4Q 2000	91.5%	4.2%	4.4%	26,782	76.8%

Table 2

Price Summary Statistics April 1998-November 2000 (\$/MWh)				
Variable	Mean	Std Dev	Min	Max
PX North	46.89	56.86	0.00	1099.99
PX South	44.30	58.83	0.00	750.00
ISO North	54.80	77.67	-325.60	750.00
ISO South	45.20	71.77	-428.15	750.00
ISO-PX North	7.92	52.57	-709.01	689.85
ISO-PX South	0.91	50.85	-709.01	688.93

Table 3

Dependent Variable is ISO-PX in NP15 (Newey-West 1 day lag MA structure)

Month	Early Hours 1-6					Late Hours 8-24				
	OLS Coef	Percent PX	Percent ISO	N-W SE	N-W P-value	OLS Coef	Percent PX	Percent ISO	N-W SE	N-W P-value
April, 1998	-3.484	0.239	0.314	1.807	0.054	-1.556	0.061	0.065	1.127	0.168
May	-1.876	0.461	0.857	0.821	0.023	-2.860	0.189	0.234	1.428	0.045
June	-1.153	0.434	0.766	0.461	0.013	-4.856	0.301	0.431	1.905	0.011
July	-6.133	0.344	0.524	1.554	0.000	-4.203	0.109	0.122	4.555	0.356
August	0.280	0.012	0.012	1.215	0.818	9.206	0.204	0.169	4.519	0.042
September	3.517	0.147	0.128	1.040	0.001	8.255	0.217	0.178	4.301	0.055
October	8.922	0.381	0.276	1.208	0.000	6.776	0.230	0.187	1.263	0.000
November	3.717	0.155	0.134	1.180	0.002	3.108	0.109	0.098	0.833	0.000
December	-3.681	0.134	0.155	2.444	0.132	0.432	0.014	0.014	2.266	0.849
January, 1999	-1.321	0.084	0.092	1.034	0.202	-2.194	0.092	0.101	0.689	0.001
February	-1.052	0.079	0.086	0.568	0.064	0.178	0.008	0.008	0.478	0.710
March	-1.934	0.140	0.163	0.931	0.038	1.218	0.056	0.053	1.033	0.238
April	-0.273	0.016	0.016	0.852	0.749	1.787	0.067	0.063	2.637	0.498
May	-2.364	0.170	0.205	1.190	0.047	-4.793	0.171	0.207	1.355	0.000
June	-2.706	0.267	0.364	1.113	0.015	-2.007	0.067	0.072	3.607	0.578
July	-11.289	0.585	1.409	4.662	0.016	-9.278	0.248	0.329	4.847	0.056
August	-2.021	0.095	0.104	1.454	0.165	3.382	0.085	0.078	5.718	0.554
September	0.764	0.026	0.025	1.730	0.659	2.464	0.058	0.055	5.123	0.631
October	-0.968	0.026	0.027	3.094	0.754	7.758	0.123	0.110	8.045	0.335
November	6.637	0.242	0.195	3.128	0.034	11.420	0.274	0.215	4.768	0.017
December	1.506	0.063	0.059	1.678	0.370	3.481	0.110	0.099	1.100	0.002
January, 2000	1.364	0.053	0.051	1.616	0.399	1.968	0.059	0.056	1.411	0.163
February	1.080	0.042	0.040	1.334	0.418	-1.203	0.038	0.040	1.437	0.402
March	-2.039	0.092	0.102	1.157	0.078	1.785	0.059	0.056	1.372	0.193
April	-1.714	0.115	0.130	2.062	0.406	3.100	0.101	0.092	3.768	0.411
May	14.348	0.575	0.365	3.820	0.000	6.546	0.117	0.105	11.080	0.555
June	11.805	0.228	0.186	7.156	0.099	3.966	0.025	0.025	31.250	0.899
July	22.663	0.458	0.314	9.693	0.020	36.134	0.357	0.263	12.072	0.003
August	41.223	0.476	0.323	5.447	0.000	54.344	0.331	0.249	12.819	0.000
September	56.180	0.683	0.406	10.471	0.000	68.311	0.572	0.364	9.125	0.000
October	42.986	0.513	0.339	7.613	0.000	39.985	0.370	0.270	5.402	0.000
November	33.580	0.224	0.183	11.229	0.003	25.448	0.142	0.124	8.864	0.004

Table 4

Dependent Variable is ISO-PX in SP15 (Newey-West 1 day lag MA structure)

Month	Early Hours 1-6					Late Hours 8-24				
	OLS Coef	Percent PX	Percent ISO	N-W SE	N-W P-value	OLS Coef	Percent PX	Percent ISO	N-W SE	N-W P-value
April, 1998	-4.162	0.286	0.400	1.684	0.014	-1.578	0.062	0.066	1.126	0.162
May	-1.876	0.461	0.857	0.821	0.023	-1.767	0.117	0.133	2.059	0.391
June	-1.114	0.426	0.741	0.454	0.014	-4.994	0.307	0.443	1.799	0.006
July	-5.794	0.332	0.497	1.595	0.000	-5.354	0.135	0.156	4.789	0.264
August	-3.398	0.157	0.187	1.580	0.032	6.389	0.135	0.119	4.793	0.183
September	-1.475	0.070	0.076	1.741	0.397	3.310	0.087	0.080	3.814	0.386
October	2.406	0.177	0.152	1.666	0.149	4.381	0.157	0.137	1.389	0.002
November	2.489	0.225	0.183	1.254	0.048	0.815	0.030	0.029	0.668	0.223
December	-1.397	0.080	0.087	1.569	0.373	-0.275	0.009	0.009	2.149	0.898
January, 1999	-0.300	0.022	0.022	1.138	0.792	-2.009	0.085	0.093	0.694	0.004
February	-1.030	0.078	0.084	0.566	0.069	0.171	0.008	0.008	0.478	0.721
March	-1.110	0.086	0.094	0.946	0.241	1.274	0.058	0.055	1.016	0.210
April	-0.273	0.016	0.016	0.852	0.749	1.679	0.063	0.059	2.647	0.526
May	-2.330	0.168	0.202	1.181	0.049	-4.793	0.171	0.207	1.355	0.000
June	-1.960	0.209	0.264	1.063	0.066	-1.965	0.066	0.071	3.637	0.589
July	-7.089	0.469	0.885	4.083	0.083	-7.857	0.218	0.279	5.438	0.149
August	-3.300	0.170	0.205	1.735	0.057	3.677	0.096	0.088	4.894	0.453
September	-0.136	0.008	0.008	2.162	0.950	5.340	0.158	0.136	5.826	0.360
October	-2.649	0.091	0.100	2.082	0.204	4.465	0.101	0.092	4.090	0.275
November	-2.558	0.149	0.175	3.686	0.488	4.299	0.125	0.111	2.450	0.080
December	5.007	0.248	0.199	2.230	0.025	3.445	0.111	0.100	1.097	0.002
January, 2000	1.788	0.077	0.072	1.720	0.299	0.788	0.024	0.024	1.362	0.563
February	1.529	0.062	0.058	1.628	0.348	-2.035	0.064	0.069	1.508	0.177
March	-1.736	0.083	0.090	1.283	0.176	0.268	0.008	0.008	1.279	0.834
April	-1.356	0.093	0.103	1.958	0.489	9.648	0.263	0.208	7.601	0.205
May	10.849	0.455	0.313	2.969	0.000	16.162	0.247	0.198	14.170	0.254
June	16.686	0.469	0.319	5.467	0.002	0.081	0.001	0.001	28.928	0.998
July	3.567	0.082	0.076	5.565	0.522	7.747	0.059	0.056	12.146	0.524
August	21.395	0.399	0.285	6.072	0.000	-8.131	0.042	0.044	10.902	0.456
September	29.755	0.517	0.341	8.888	0.001	12.472	0.102	0.092	10.807	0.249
October	-29.171	0.480	0.924	6.539	0.000	-16.409	0.172	0.207	6.541	0.012
November	7.075	0.083	0.077	11.267	0.530	-3.934	0.027	0.028	10.158	0.699

Table 5

PROFITABILITY OF TRADING RULES (average profit per MWh)

<i>Monthly</i>				
Epoch	North Early	North Late	South Early	South Late
All Months	7.54 (7.30)	9.28 (5.85)	1.82 (2.29)	1.18 (0.85)
May-Dec, 1998	2.64 (4.58)	2.61 (2.27)	1.56 (2.91)	1.72 (1.53)
Jan-Aug, 1999	2.90 (3.90)	0.46 (0.40)	2.20 (3.32)	0.77 (0.68)
Sept 1999-April 2000	-0.59 (-0.77)	3.41 (2.47)	0.39 (0.49)	3.21 (2.33)
May-Nov, 2000	27.64 (7.40)	33.55 (5.60)	3.34 (1.03)	-1.28 (-0.23)
<i>Weekly</i>				
Epoch	North Early	North Late	South Early	South Late
All Months	7.99 (7.96)	8.54 (5.52)	3.53 (4.69)	1.72 (1.29)
April 8-Dec 31, 1998	3.09 (5.73)	3.21 (3.19)	1.37 (2.60)	1.90 (1.91)
Jan-Aug, 1999	1.59 (2.05)	0.71 (0.61)	0.60 (0.88)	0.77 (0.69)
Sept 1999-April 2000	0.54 (0.72)	1.93 (1.38)	0.68 (0.88)	3.38 (2.46)
May-Nov, 2000	29.87 (8.30)	31.68 (5.21)	12.79 (4.31)	0.68 (0.12)

T-statistics in parentheses.

Table 6: Sharpe Ratios for the Weekly Trading Rule

	4/98- 12/98	1/99- 8/99	9/99- 4/00	5/00- 11/00	Total Sample
North – Early	.73	.61	1.38	1.68	.71
North – Late	.86	.97	.95	1.37	.97
South – Early	.77	.92	.44	.65	.64
South – Late	.80	.94	.90	1.02	.87
S&P 500	-.09	.13	.04	-.25	-.09

Figure 1a

Monthly Price Averages (all hours in NP15 - \$/MWh)




Figure 1b

Monthly Price Averages (all hours in SP15 - \$/MWh)




Figure 2

Cummulative Profits: Weekly Rule




Figure 3a

P-Values of Weekly Trading Rule




Figure 3b

P-Values of Monthly Trading Rule




Figure 4

RECLAIM NOx Costs and Natural Gas (Topok) by Month




Figure 5

Aggregate Supply Curve for All Firms

Tuesday, 15 Aug 2000, Hour 16




Figure 6

Market Equilibria




Figure 7

Unanticipated Decrease in Forward Demand




Figure 8

Monthly Average Percentage of Total Demand Bid into PX at ISO Price




Figure 9

Utility PX Demand as a Fraction of Total Utility Demand




Figure 10

Monopoly Arbitrage with Unanticipated Decrease in Forward Demand

