Trading Inefficiencies in California's Electricity Markets

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

July 7, 2003

Abstract

We study price convergence between the two major markets for wholesale electricity in California from their deregulation in April 1998 through November 2000, nearly the end of trading in one market. We would expect profit-maximizing traders to have eliminated persistent price differences between the markets. Institutional impediments and traders' incomplete understanding of the markets, however, could have delayed or prevented price convergence. We find that the two benchmark electricity prices in California – the Power Exchange's day-ahead price and the Independent System Operator's real-time price – differed substantially after the markets opened but then appeared to be converging by the beginning of 2000. Starting in May 2000, however, price levels and price differences increased dramatically. We consider several explanations for the significant price differences and conclude that rapidly changing market rules and market fundamentals, including one buyer's attempt to exercise a form of monopsony power and institutional barriers to arbitrage trades, made it difficult for traders to take advantage of opportunities that *ex post* appear to have been profitable.

* *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: jimb@IEOR.Berkeley.edu. *Knittel:* Department of Economics, University of California, Davis and University of California Energy Institute. 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 Steve Puller, Erin Mansur and Celeste Saravia for outstanding research assistance and to Jonathan Berk and Rich Lyons for useful discussions. This research was partially supported by the California Independent System Operator.

1 Introduction

While efficient markets are the cornerstone of much of financial economics, there is anecdotal evidence that transitory profit opportunities frequently exist. Some *apparent* opportunities are in fact due to institutional barriers or transaction costs that prevent a set of trades that would otherwise have positive expected return. There may also be cases in which the presence of a significant number of imperfectly informed traders creates opportunities that are not completely eliminated by the actions of better-informed market participants. In new markets, even the most experienced traders may lack sufficient information to anticipate *ex ante* what turn out to be profitable trading opportunities *ex post*. Similarly, a significant regime shift in the pricing relationship between markets can produce apparent inefficiencies since, again, speculators lack sufficient experience with what is a new pricing regime.

Over the last decade, several new and necessarily complex markets for wholesale electricity have been established in the United States. Many of these markets featured both day-ahead and 'real-time' markets for electric energy at various locations. Although these sequential markets are analagous to traditional forward and spot commodity markets, they were often not viewed as such by the institutions that created them. The real-time markets were envisioned as small markets intended to accommodate last-minute adjustments to changes in market conditions.

Complicating the challenge to market efficiency is the lack of a widely available economic means to store electricity. While firms may not broadly utilize inventories to engage in risk-free arbitrage, this does not exclude the possibility, or even expectation, that day-ahead prices for electricity would reflect the expectation of real-time spot prices.

California became the first state to restructure its electricity industry when it opened its deregulated day-ahead market on March 31, 1998. The real-time market opened on April 1, 1998. In the last half of 2000, high and extremely volatile prices for electricity in California led many to conclude that the deregulated market was dysfunctional. By the beginning of 2001, it was characterized by periodic shortages, continuing high prices, the bankruptcy of the market's largest buyer, Pacific Gas & Electric (PG&E), and the involvement of the State of California as a major trader.

We analyze the relationship between wholesale electricity prices in the day-ahead market run by the California Power Exchange (PX) and the real-time spot market overseen by the California Independent System Operator (ISO) over a 32-month period beginning on the day the markets opened in 1998 and ending just before the demise of the PX. During several parts of our sample, we find large price differences, though, before the last half of 2000, none of the price differences persisted for more than a couple months.

Disruptions to the market beginning in May 2000, however, caused the price differences to reach previously unseen levels. 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. As the PX day-ahead price became an increasingly poor predictor of spot prices, its value as a hedging and planning tool was correspondingly diminished. Eventually, major rule changes imposed by the Federal Energy Regulatory Commission (FERC) in December of 2000 led to a steep dropoff of trading volume in the PX, and effectively ended its ability to operate as a viable market. The PX day-ahead market ceased operations on January 31, 2001, and on March 9, 2001, the PX officially filed for Chapter 11 bankruptcy protection.

Having documented the significant differences between ISO and PX prices, we go on to analyze why they failed to converge. Our analysis suggests that neither risk aversion nor differing transaction costs are likely to explain the ISO-PX (spot-future) price differences. Instead, we attempt to demonstrate that the differences were attributable to four factors: imperfect information of participants that improved over time (learning), regime shifts in the price formation process that disrupted this learning, an attempt by at least one buyer to exercise a form of monopsony power and institutional barriers to financial trading that limited the number of participants.

To analyze why the price differences persisted, we present statistical analyses, but we also have traders' descriptions of their motivations and beliefs about the markets. Enron was a major trader in the markets we are studying. Investigations into its collapse and its role in the California Energy Crisis in the last half of 2000 have yielded internal emails, memoranda and presentations relevant to our analysis. For instance, in May 2002, Enron's lawyers uncovered memoranda describing trading strategies for the California electricity markets.[1] One of those strategies, labeled "Fat Boy" by the traders, took advantage of the price differences we document.

Following the release of the memos, two Enron traders pled guilty to wire fraud for their involvement in the California markets. The pleadings describe strategies other than "Fat Boy" as illegal, although some of the strategies to which they pled guilty appear to be simple arbitrage plays of one form or another. While some ISO rules declare that

[1]The memos are available of the FERC's website at http://www.ferc.gov/electric/bulkpower/pa02-2/pa02-2.htm#memo.

offers must represent genuine interest in physically selling or consuming electricity, there were no explicit penalties for arbitrage trades during the period we study and no serious attempts to discourage financial trades. Nonetheless, Enron appears to have reaped the majority of the profits from the arbitrage trades and only Enron traders are involved in criminal proceedings. Apparently, both the complexity of the markets and the ambiguous legal and regulatory treatment of these practices limited the entry of trading firms and other potential arbitrageurs, and Enron was able to capitalize on this market power in the arbitrage markets.

A number of other papers, both theoretical and empirical, consider prices in situations where the same good is traded on multiple markets. While most of these papers look for equilibrium explanations for price differences (see, for example, Shleifer and Vishny (1997)), we believe that the ISO-PX (spot-future) price differences reflected out-of-equilibrium behavior that would have eventually disappeared had the PX survived longer. Other papers have considered how markets incorporate new information (Biais, Hillion and Spatt, 1999) or how prices may reflect investors' expectations of a regime shift in the price formation process.[2] The demand in most financial markets is thought to be highly elastic,[3] so the idea that a participant could exercise market power clearly differentiates the ISO and PX from most other financial markets.

Two recent papers have analyzed differences between forward and spot prices in electricity markets and have focused on risk-aversion of the market participants. Bessembinder and Lemmon (2002) develop a model of the risk premium in electricity forward contracts and test it by analyzing one-month ahead forward trades for the California and PJM (Pennsylvania, New Jersey, Maryland) markets compared to their respective day-ahead prices. Their analysis covers wholesale transactions both before and after deregulation took effect and focuses on the hedging function of the futures contracts. Longstaff and Wang (2002) study the day-ahead to real-time price relationship in the PJM electricity market. Over an 18 month period, they find that mean day-ahead prices were lower than mean spot prices, but the median day-ahead prices was higher than the median real-time prices. They attribute these differences to the risk aversion of electricity suppliers.

Risk aversion, however, does not seem to be a very satisfactory explanation for the price differences we find. Although day-ahead and spot prices frequently do not converge, the sign of the forward "premium" is not consistent over time or location.[4] At different

[2] See, for example, Engel and Hamilton (1990), Kaminsky (1991) and Lewis (1989, 1991).

[3] See, however, Kaul, Mehrotra and Morck (2000) for evidence that this is not always the case.

[4] There also is no obvious seasonal pattern in the forward premium as Bessembinder and Lemmon's theory would suggest.

stages of our sample, for example, we find periods in which forward prices were higher than spot prices as well as periods where the opposite relationship maintained.[5] Further, although many producers were likely risk averse, these markets also included a diverse set of merchant and trading firms. As we describe below, trades were available that, albeit not risk-free, appear to have consistently positive expected value and virtually no non-diversifiable risk. Recent revelations indicate that at least one firm, Enron, was taking advantage of these opportunities in California.

Another commonly cited cause of bias in estimating spot/forward price relationships is the "peso problem." The peso problem arises when, during the sample period, investors place a positive probability on a discrete event occurring, such as a currency devaluation, and the event does not occur. This drives a wedge between the forward rate, which is affected by this probability, and the spot rate, which is unaffected. This theme is present in Krasker (1980), Lewis (1988), Kaminsky and Peruga (1990 and 1991), Bachman (1992) and Kaminsky(1993).[6] It is extremely unlikely that a peso-type problem can explain the price differences that we observe in the California market. First, the existence of both price floors and price ceilings truncates the distribution of spot/forward price differences. Second, conditional on the price floors and price ceilings, in our data we often observe the *most catastrophic* event possible, driving spot prices to one extreme or the other. Therefore, it is not the case that players in the market are placing positive probability on an event that *does not* occur within the sample.

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 explain some of the price differences. We conclude that rapidly changing market rules and market fundamentals, including PG&E's attempt to exercise monopsony power, made it difficult for traders to take advantage of opportunities that *ex post* appear profitable.

[5]While the forward premium in PJM was negative over the 18 month period studied by Longstaff and Wang, they do not report the premium for shorter time periods within the sample. The premia in the adjacent New York electricity markets were positive over the same time period (see Saravia, 2002) despite the fact that these markets shared many of the same participants and engaged in substantial cross-border trade.

[6]Additional evidence of the peso-problem in option pricing is uncovered in Bates (1996).

2 Price Relationships in Electricity Markets

The California electricity market is 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 10MW 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.[7]

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,

$$ {}_{t-j}P_t = E\left[{}_tP_t|\Omega_{t-j}\right] \tag{1} $$

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, ${}_{t-j}P_t -{}_t P_t$, 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

$$ {}_tP_t ={}_{t-j} P_t + \varepsilon_t, \tag{2} $$

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

[7]Note 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.

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 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.[8] 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 bids for the hourly supply and demand of electricity for the 24 hours of the following day. Bids were submitted for the day-ahead market by 7a.m. 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.[9]

Though in a first-round calculation each day, the PX calculated day-ahead prices as if all bids and offers were in a common market, limits on the capacity of electricity transmission lines often necessitated further price adjustment. Electricity transmission is extremely

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

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

capital intensive and marginal transportation costs are usually negligible. Consequently, the shadow value of the transmission capacity constraints determine the marginal cost of transporting electricity from one area to another in the network, and is zero if the capacity constraint is not binding. This shadow price is in turn based upon the relative costs of electricity supply on either side of a network constraint.

In California, for purposes of transmission pricing, the ISO system was divided into 24 zones.[10] 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.[11] 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.[12]

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

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

[12]Until 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

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 most customers did not have access to, let alone an incentive to respond to, real-time prices, the ISO had to ensure that supply and demand remained in continuous balance (by adjusting production), despite the random fluctuations of production and consumption. An imbalance energy market, run by the ISO, was created 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. After August 19, 1999, however, this situation changed slightly, as discussed below.

Table 1 gives the relative volumes of the day-ahead, hour-ahead, and imbalance energy markets for the months of July 1998 to January 2001.[13] These figures give the percentage of energy that was scheduled day-ahead from all SCs, the percentage scheduled hour-ahead, and the percentage provided in the real-time imbalance energy market.[14] For much of the study period, real-time volume averaged around 3% of total volume, but this figure climbed as high as 6% during the spring and summer of 2000. 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.

---

low and in more than 25% of all hours no transactions took place.

[13]Data were not available for April-June 1998. Trading on the PX ended in January 2001, so the data series ends there.

[14]These volumes reflect the absolute differences between aggregate final day-ahead schedules (DA), hour-ahead schedules (HA) and actual real-time load (RT). Total volume, from all three sets of markets is measured as $DA + |HA - DA| + |RT - HA|$. 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. 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. 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.

Each of these three avenues of supply – ancillary services, imbalance market bids, and excess 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,

[15]Suppliers 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.

suppliers of some reserves earned no imbalance energy revenues even when their energy bid was below the imbalance energy price.[16]

On August 19, 1999, however, the ISO began allocating the costs of replacement reserve capacity, a form of operating reserve, disproportionately to suppliers that produced less than their scheduled quantity and demanders that consumed more than their scheduled demand. This produced an additional penalty on transactions that end up as net purchases from the imbalance energy market.[17]

A supplier that was scheduled to provide energy in one of the forward markets could also take a short position in the spot market by either 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.[18]

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.

[16]This 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.

[17]As discussed later, we have examined the impact of this change.

[18]In 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.

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.[19] After the market opened, further restrictions and institutional barriers were 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 deal 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 production 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.

[19] In fact, for most of the time period we study, the three investor-owned distribution utilities, which were responsible for serving the bulk of demand within the ISO system, were not allowed to enter into longer term supply contracts or any supply contracts outside of the PX. During the summer of 1999, they won approval to purchase the PX's long-term futures product, the block forward. Before that time, their forward purchasing could be no more than one day before actual consumption.

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 from 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.[20]

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.[21] 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.[22]

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

[20]See Bohn, Klevorick, and Stalon (1999), page 13.

[21]Bohn, Klevorick, and Stalon (1999).

[22]The 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 at the beginning of August – in response to the unprecedented price levels experienced during May and June of that year.

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,[23] 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 to avoid forward trades in favor of transacting on the imbalance energy market. During that time period, 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 a firm that scheduled no supply or demand, but instead simply produced or consumed in real time without any notification, prior to August 19, 1999 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 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

[23] The 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.

at time $t-j$ then:

$$_{t-j}PX_t = E\left[ISO_t|\Omega_{t-j}\right] \tag{3}$$

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.*, $ISO_t = E\left[ISO_t|\Omega_{t-j}\right] + \varepsilon_t$), we have:[24]

$$ISO_t = PX_t + \varepsilon_t \tag{4}$$

This implication can be tested by estimating the model:

$$ISO_t - PX_t = \alpha + \varepsilon_t \tag{5}$$

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 indeed 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-1am 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 will also be serially correlated, implying the standard errors obtained from ordinary least squares will be biased.[25] 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. We describe this below and then discuss two different estimation approaches.

[24] We now suppress the $t-j$ presubscript on PX_t.

[25] For 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.

Let the information set at hour t be represented by Ω_t. Let $t = 1$ represent the beginning of an arbitrary day (*i.e.* the 12:00 midnight–1 AM hour). The PX prices for $t = 1, ..., 24$ 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 $t = -18$, which would be Ω_{-18}.[26] 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 $t = -18, ..., -1$ timeframe may continue to impact this difference for hours $t = 1, ..., 24$ (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 $t = -18, ..., -1, 0, 1, ...6$ 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:

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

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

$$
\begin{aligned}
ISO_{24} - PX_{24} &= \alpha + \varepsilon_{24} + \sum_{i=1}^{41} \theta_{24,i}\varepsilon_{24-i} \\
ISO_{25} - PX_{25} &= \alpha + \varepsilon_{25} + \sum_{i=1}^{18} \theta_{25,i}\varepsilon_{25-i} \\
&\vdots
\end{aligned}
$$

Unfortunately, our attempts to estimate a model with varying serial correlation lengths 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.

So, we have taken a simplified alternative approach. Instead of estimating a single regression 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, though 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. So, instead, we have averaged the price differences for the early and later parts of the day, using one observation per day for each. An "early" 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 that 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 alternatively, the average ISO-PX price difference in hours 1-6 and hours 8-24. 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.[27] 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 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.[28]

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, 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.

[27]We 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.

[28]As 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.

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 make money in the hands of a pure speculative trader, who, unconstrained by institutional barriers, could buy 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 use 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 make 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.[29] Table 5 summarizes the coefficients and t-statistics from including this variable in a specification of equation (5) without any month dummies. The first two columns report results from specifications that included all thirty-two months, while the remaining columns report tests during four separate time periods. Considering the entire time period, the t-statistic 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, 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.

[29]We assume that the trader trades an equal quantity each hour.

5 Possible Explanations for the Observed Price Differences

The results thus far suggest that significant price differences have persisted between the PX and ISO, and that several simple trading strategies would have made money. This section considers the extent to which transaction costs, risk aversion, or participant learning about the price formation process may explain the significant price differences.

5.1 Risk Aversion

Persistent price differences could reflect risk aversion on the part of the 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.

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 on speculating on the ISO–PX price difference have essentially no correlation with any other investments, so the risk associated with them could be diversified away by those with claims on their returns. We calculated the coefficient on the return on the S&P 500 (the β) from a CAPM model of the ISO–PX price difference and could not reject that the coefficient was 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.[30] We compared the risk-return properties of speculation on the ISO–PX price differences to the S&P 500 by computing the Sharpe ratio for the trading rules discussed in the previous section.

Calculating the Sharpe ratio requires defining the time period over which returns are computed. Typically, Sharpe ratios are computed for yearly returns. For example, Sharpe ratios are often calculated for yearly returns on the S&P 500. With 32 months of electricity price data, this makes little sense. Therefore, 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

[30] See, for example, Chapter 7 in Lyons (2001).

("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 1/18th 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:

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

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

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

Monthly returns are computed in an analogous manner. The Sharpe ratio is based on the mean and standard deviation of these returns.[31]

As a comparison, we also calculated the Sharpe ratio for someone trading in the S&P 500. 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 and then sells the shares at the end of the period.[32]

Table 6 lists the Sharpe ratios for the weekly trading rules.[33] 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.

5.2 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, one would expect all volume to move 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

[31] During 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.

[32] We 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.

[33] Sharpe ratios based on the monthly trading rules were very similar.

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 who bears 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.[34] Sellers must be induced to trade in the forward market, so the 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.[35]

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 costs in either market changed substantially at the beginning of summer 2000.

[34] This is *not* intended to be a characterization of the California market. The actual incentives in the California market are much more complex.

[35] It 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.

5.3 Buyer Market Power

Another element that distinguishes electricity markets from those for most other commodities is the pervasiveness of market power. Much has been written about the market power of suppliers in the California electricity market.[36] However, the concentration of *buyers* is also quite substantial in California, as in other electricity markets. The demand bid by Pacific Gas & Electric (PG&E) constituted 30% of overall ISO demand and 85% of demand in the NP15 zone. It is important to note that such large purchasing responsibility did not convey upon the California utilities monopsony power in the traditional sense, since the utilities in their role as distributor had no control over the aggregate levels of end-use consumption. Although the utilities could not control how much power was consumed they did have some discretion over the market in which the power was purchased. In particular, the utilities clearly understood the implications of moving their purchases between the PX day-ahead and ISO real-time markets.

In subsequent regulatory investigations, PG&E has acknowledged their recognition that "every additional MW cleared in the PX market increased the price for every MW purchased" and that "the same is true for the ISO market." [37] Given their size, it is possible that PG&E could, *ceteris paribus*, reduce its overall purchasing costs by shifting purchases from the PX to the ISO real-time market. This could raise prices in the real-time market, but PG&E recognized that "paying a higher price in the ISO market for the incremental portion of total load 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] Given that PG&E's net purchase volume in the PX was much larger than in the ISO, it could make sense for PG&E to tolerate a higher ISO price if doing so lowered costs on its inframarginal purchases in the PX significantly.

Such practices became increasingly prominent during the summer of 2000, prompting serious concern at the ISO. The ISO had not anticipated the levels of volume then experienced in the real-time market and felt that such increases threatened system reliability. On June 13, the ISO issued a market notice warning of extensive day-ahead "under-scheduling" of demand and pointing out the substantial difference between the incremental costs of load in the two markets. As table 1 illustrates, however, the volume handled by the real-time market steadily grew throughout the summer of 2000.

As PG&E pursued this strategy, we would expect that eventually sellers would respond

[36] See, for instance, Borenstein, Bushnell and Wolak (2002) and Joskow and Kahn (2002).
[37] PG&E (2002) p.271
[38] PG&E (2002), p.009.

by trying to sell more of their output in the ISO market. This might not be the case for the generation still owned by the three major utilities, since their incentives under continued regulation were complex, but it would clearly be the case for the deregulated suppliers, known as exempt wholesale generators (EWGs), who were only on the sell side of the wholesale market. There is indeed strong evidence that by September 2000 the EWGs in the North were trying to move most of their volume out of the PX and into the ISO.[39] Figure 3 shows that the share of EWG production sold through the ISO was just a few percent through June 2000.[40] Beginning with slight increases in July and August, the share of EWG production in the North rose to nearly 50% by September. In the South, there was a substantial, but smaller, increase in sales through the ISO. One might ask why EWGs were selling any output through the PX by September. One explanation is that moving more volume out of the PX and into the ISO would depress the ISO price for the output that they were already selling there.[41]

5.4 Learning about the Price Formation Processes

While "physical" market participants appear to have slowly responded to the persistent price differences in northern California during 2000, the question remains as to why other firms with no physical assets were not also taking advantage of an apparently profitable trading opportunity. If enough firms were engaged in such a form of arbitrage, prices in the two markets would have again converged in expectation. We have already considered the role of transactions costs and risk aversion, and neither provides a satisfactory explanation. In fact, we know of one firm, Enron, that was trading on these opportunities. The "Fat Boy" strategy described in internal Enron memos that have subsequently become public consisted of making purchases in day-ahead markets and then implicitly reselling that purchased supply onto the real-time market. It appears that few other firms were either

[39] 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 eneable it to sell power into the Real Time market. Thus Williams does not have the incentives that were apparently driving Enron to schedule load 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."

[40] These figures are calculated from ISO-instructed real-time production (in the numerator) and total unit generation (in the denominator). As of this writing, data for the numerator are available only back to February 1999, so the figure starts at that date.

[41] Also, any output sold through the PX's block forward contracts was required to be delivered through the PX day-ahead market. The EWGs did sell such contracts, but we have not been able to obtain data on the amount of power they sold through block forwards and were thus required to deliver through the PX.

aware of these opportunities or in a position to take advantage of them.[42] If a small number of firms are aware of arbitrage opportunities, they will not want to trade at such volumes as to eliminate those opportunities. It appears that such factors were at play in California during 2000.

Why were only a few firms trading on these price differences? In any new market, it may take participants time to learn about how market rules, market fundamentals and their own behavior affect prices. Learning could take two forms: either more sophisticated participants enter a market over time, or the existing participants become more proficient at taking advantage of price differences. Of course, if market rules and market fundamentals are changing rapidly enough, even the most sophisticated traders may not be able to keep up, and any accumulated experience with the market may be irrelevant. There were significant changes in the ownership shares of sellers into the California markets at the beginning of our sample period when the investor-owned utilities sold their plants to private wholesale generators, although the entry (or exit) of market participants seems unlikely to have created the huge price differences we see at the end of our sample. However, institutional restrictions meant to deter purely financial transactions along with the unstable political and market environment during 2000 most likely discouraged the entry of other sophisticated trading firms, thus contributing to the persistence of those price differences.

The second form of learning seems more likely to be a factor in explaining our results. While it is easy to talk abstractly about the distributions of prices and price differences, in reality a trader in these markets is constantly updating her beliefs about these distributions, and must recognize that her knowledge of the underlying distribution of prices is very imperfect. Furthermore, in dynamic and new markets, such as in the California electricity market, the distribution that a trader faces *is* constantly changing as market rules are modified and as other firms modify their behavior.[43] The result is that traders do not *know* the true distribution of prices they face on any given day.

To be somewhat more concrete, assume that a seller is trying to decide whether to sell in the PX or ISO. He is risk neutral and would like to sell where the price will be higher. He has a limited history of data from which he gleans that the mean price in the PX has been

[42] At least two other firms, El Paso Electric and Dynegy, have admitted in filings to FERC of engaging in some transactions profiting from the day-ahead/real-time price relationship in California, although they deny that these trades were analogous to Enron's "Fat Boy" strategy. See Dynegy (2002) and El Paso Electric (2002).

[43] Indeed, Knittel and Roberts (2000) find that the parameters associated with a variety of financial models of California electricity prices have changed since the inception of the market.

lower than in the ISO, but a statistical test indicates that the means are not statistically different. Furthermore, he knows that rule changes have taken place recently (after some of the data on which he is making his comparison) that could affect the ISO–PX price relationship. The seller will still form some guess about the expected price differential, but it is easy to see how there could be an underlying systematic price difference that the seller does not uncover for a fairly long period of time. There are a number of events that could have altered the underlying price formation process in the California markets. We have already noted the rule change that affected payments for replacement reserves after August 19, 1999, but the results in tables 3 and 4 indicate that this rule change did not disrupt a trend towards price convergence. More of the α coefficients were significant in the six months preceding August 1999 than in the six months following that month. A number of other factors, including weather, time-of-year, plant ownership, and other market rules (including the market price caps) also changed during our sample period, and these all could affect the market participants' expectations about prices. As a result, it is difficult to pinpoint the effect of a particular event that changed traders' expectations.

The Collapse of the Market

As is now well known, the California electricity market went through dramatic changes beginning in summer 2000 as average wholesale prices more than doubled from their highest previous level. The relationship between ISO and PX prices also changed: 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). Factors discussed in the previous sections, such as participants' relative risk aversion or transaction costs are unlikely to explain these results. There is no evidence that risk aversion or transaction costs changed, and the changes would have had to be extremely large and isolated to the North. It is more likely that there was a shock to the price formation process that either carried over to the ISO-PX price relationship or elicited a response that affected the ISO-PX price relationship. As we discuss below, the latter seems most likely. The overall price shock that began in mid-2000 was due in part to unprecedented increases in the prices of two inputs to electricity generation. As Figure 4 documents, the prices for natural gas and NOx pollution permits increased by factors of around 3 and 50 respectively, beginning in the spring of 2000. These increases had the greatest effect on the costs of the units that are marginal at peak times, 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, increased the volatility of prices.

While volatility in itself would not change the ISO-PX equilibrium price relationship, it could have impacts while players sorted out the new incentives and likely outcomes. In particular, it appears that along with the regime shift in late May and early June, extremely high ISO prices became more common during periods of very tight supply. If market participants were slow to recognize (or believe) this price distribution change, then the PX prices would underforecast ISO prices because the PX prices would not incorporate, or would underweight, the potential for extremely high prices in the ISO. Of course, we wouldn't expect such forecast errors to persist indefinitely, but the unprecedented price levels could cause adjustment to be rather slow.

This explanation, however, does not square with the very different price relationship changes in the North and South. Since PG&E was the major buyer in the North, their attempts to lower their purchase price by shifting some of their demand to the ISO would have affect prices in the North more than in the South. As discussed above and documented in Figure 3, suppliers evenutally followed the load to the ISO, but not until the price differences had persisted for several months.

The migration of the non-utility suppliers out of the PX undermined the market as it reduced the number of players on the sell side of the market. 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

One of the dominant questions surrounding the reorganization of the electricity industry is the role of markets in coordinating the short-run operation of electricity systems. Debates over the proper role of system operators, and the degree to which market incentives can eliminate operational inefficiencies, continue.

In this paper, we have studied one aspect of the California market performance: the interaction between the day-ahead market of the PX and the real-time market of the ISO. Although these markets play very different institutional roles, and operate under quite different market rules, they are 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 has established that significant price differences existed between the PX and the ISO during several periods in the first 32 months of operation, particularly during the last seven months. Also, it appears that some risky trading strategies with positive expected return existed.

From our analysis, two potential explanations emerge for the persistent and unprecedented differences between day-ahead and real-time prices during the summer of 2000. One is that the many regulatory and structural changes introduced into the market during this period overwhelmed the ability of most firms to accurately predict price differences. It is important to note that the presence of a single, or even several, firms sophisticated enough to detect profitable trades is not a guarantee that prices will converge. Each firm individually would not want to eliminate price differences completely. Rather the markets need a sufficient number of sophisticated firms that the combined trading activity of all of them causes prices to coverage. Electricity markets are extremely complex, and most active participants have had limited experience trading commodities in unregulated markets.

The other possible explanation for at least the initial emergence of the price-differences during the summer of 2000 was an attempt by the largest buyers of electricity to exercise a form of monoposony power over these markets. We presented data and internal company memos consistent with this explanation of initial price differences in 2000. Although such attempts may constitute a significant regime shift in market patterns, and lead to significant short term price differences, they do not by themselves explain the persistence of price differences throughout the summer of 2000. A sophisticated trader, in hindsight, should have been able to profit from these price differences. With enough traders, any attempts to induce and profit from persistent price differences between markets should be eventually undermined by the forces of arbitrage. In the complex and frequently-disrupted California electricity market, it appears that such learning did not occur fast enough for prices to converge before the eventual demise of the day-ahead market.

References

[1] Bachman, D. (1992). "The Effect of Political Risk on the Forward Exchange Bias: The Case of Elections," *Journal of International Money and Finance* vol. 11, pp. 208-219.

[2] Bates, D. S. (1996). "Dollar Jump Fears, 1984-1992: Distributional Abnormalities Implicit in Foreign Currency Futures Options," *Journal of International Money and Finance* vol. 15, pp. 65-93.

[3] Bessembinder, H. and M. Lemmon (2002). "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets," *Journal of Finance* vol. 57, June, pp. 1347-1382.

[4] Biais, B., P. Hillion, and C. Spatt (1999). "Price Discovery and Learning during the Preopening Period in the Paris Bourse," *Journal of Political Economy* vol. 107, December, pp. 1218-48.

[5] Bohn, R., A. Klevorik, and C. Stalon (1999). "Second Report on Market Issues in the California Power Exchange Energy Markets." Prepared for the Federal Energy Regulatory Commission.

[6] Borenstein, S., J.B. Bushnell, and F. Wolak (2002). "Diagnosing Market Power in California's Deregulated Wholesale Electricity Market." *A*merican Economic Review, vol. 92, December, pp. 1376-1405.

[7] Bushnell, J. and S. Oren (1997). "Transmission Pricing in California's Proposed Electricity Market." *Utilities Policy* vol. 6, no. 3.

[8] 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.fed.us/Electric/bulkpower/PA02-2/pa02-2.htm. May 22, 2002.

[9] 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.fed.us/Electric/bulkpower/PA02-2/pa02-2.html. May 22, 2002.

[10] Engel, C. and J.D. Hamilton (1990). "Long Swings in the Dollar: Are they in the Data and Do Markets Know it?," *American Economic Review* vol. 80, September, pp. 689-713.

[11] Enron (2002). "Enron Memo Dated December 6, 2000. 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.fed.us/Electric/bulkpower/PA02-2/pa02-2.htm#memo.

[12] Joskow, P.L. and E. Kahn (2002). "A Quantitative Analysis of Pricing Behavior in California's Wholesale Electricity Market During Summer 2000." *The Energy Journal* vol. 23, December, pp. 1-35.

[13] 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.

[14] Kaminsky, G. and R Peruga (1990). "Can a Time-Varying Risk Premium Explain Excess Returns in the Forward Market for Foreign Exchange?," *Journal of International Economics* 28, 47-70.

[15] Kaminsky, G. and R. Peruga (1991). "Credible Crisis: The Dollar in the Early 1980s," *Journal of International Money and Finance* vol. 10, pp. 170-192.

[16] Kaul, A., V. Mehrotra, and R. Morck (2000). "Demand Curves for Stocks Do Slope Down: New Evidence from an Index Weights Adjustment," *Journal of Finance* vol. 55, April, pp. 893-912.

[17] Knittel, C. R. and M. R. Roberts (2001). "Financial Models of Deregulated Electricity Prices: An Application to the California Market," mimeo, Boston University.

[18] Krasker, W. (1980). "The 'Peso-Problem' in Testing the Efficiency of Forward Exchange Markets," *Journal of Monetary Economics* vol. 6, pp. 269-276.

[19] Lewis, K. K. (1988). "The Persistence of the 'Peso-Problem' When Policy is Noisy," *Journal of International Money and Finance* vol. 8, pp. 1-11.

[20] Lewis, K.K. (1989). "Changing Beliefs and Systematic Forecast Errors," *American Economics Review* vol. 79, August, pp. 621-636.

[21] Lewis, K.K. (1991). "Was There a 'Peso Problem' in the U.S. Term Stucture of Interest Rates: 1979-1982?," *International Economic Review* vol. 31, February, pp. 159-173.

[22] Longstaff, F.A. and A.W. Wang (2002). "Electricity Forward Prices: A High-Frequency Empirical Analysis," mimeo, UCLA.

[23] Lyons, R. K. *The Microstructure Approach to Exchange Rates.* Forthcoming, MIT Press: Cambridge, MA.

[24] 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.

[25] Peroni, E. and R. McNown (1998). "Noninformative and Informative Test of Efficiency in Three Energy Futures Markets," *The Journal of Futures Markets* vol. 18, pp. 939-964.

[26] 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.fed.us/Electric/bulkpower/PA02-2/pa02-2.htm. May 22, 2002.

[27] Saravia, C. (2002). "Efficient Futures Markets and Production Efficiency: The Case of the New York Electricity Market," mimeo, UC Berkeley.

[28] Shleifer, A. and R.W. Vishny (1997). "The Limits of Arbitrage," *Journal of Finance* vol. 52, March, pp. 35-55.

[29] Stoft, S. (1996). "California's ISO: Why Not Clear the Market?" *The Electricity Journal* vol. 9, pp. 38-43.

[30] 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.fed.us/Electric/bulkpower/PA02-2/pa02-2.htm. May 22, 2002.

[31] 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.

Table 1

Monthly Average Trading Volumes in the Forward and Spot Markets

Month	% Day Ahead	% Hour Ahead	% Real Time	Volume Traded (MWh)	% HA > DA	% RT > HA	DA / RT	DA / RT if < 1
Jul-98	91.6%	3.6%	4.7%	29,646	76.0%	70.6%	0.933	0.924
Aug-98	91.6%	4.6%	3.8%	31,057	75.2%	68.8%	0.927	0.917
Sep-98	92.6%	4.1%	3.3%	28,274	66.4%	52.1%	0.951	0.924
Oct-98	94.3%	3.5%	2.3%	24,601	70.1%	55.6%	0.957	0.943
Nov-98	94.3%	3.4%	2.3%	23,927	66.3%	52.8%	0.960	0.939
Dec-98	93.1%	4.0%	2.9%	25,045	73.3%	63.7%	0.947	0.930
Jan-99	92.1%	4.7%	3.2%	24,399	74.7%	66.6%	0.936	0.920
Feb-99	91.2%	4.8%	3.9%	24,322	75.9%	83.3%	0.916	0.909
Mar-99	91.6%	5.5%	2.9%	24,728	75.9%	71.4%	0.926	0.916
Apr-99	92.2%	4.9%	3.0%	24,747	76.1%	53.1%	0.944	0.927
May-99	91.9%	4.6%	3.4%	24,743	70.5%	62.5%	0.941	0.918
Jun-99	91.2%	5.8%	3.0%	27,227	74.8%	52.3%	0.934	0.912
Jul-99	91.2%	5.0%	3.8%	30,150	78.2%	37.5%	0.953	0.922
Aug-99	91.4%	5.0%	3.7%	30,063	72.3%	48.9%	0.945	0.915
Sep-99	91.9%	5.2%	2.9%	28,517	77.5%	52.1%	0.939	0.916
Oct-99	92.3%	4.2%	3.5%	27,826	72.6%	43.5%	0.958	0.930
Nov-99	92.9%	3.9%	3.2%	26,324	72.9%	35.2%	0.973	0.931
Dec-99	92.6%	4.9%	2.5%	26,460	82.2%	52.9%	0.946	0.921
Jan-00	92.6%	4.7%	2.7%	26,296	79.6%	42.7%	0.954	0.922
Feb-00	90.9%	5.1%	4.0%	25,725	82.5%	80.7%	0.914	0.906
Mar-00	88.2%	5.4%	6.3%	25,623	83.7%	86.7%	0.886	0.876
Apr-00	91.0%	4.5%	4.5%	25,692	83.1%	71.3%	0.923	0.908
May-00	88.9%	5.1%	5.9%	26,968	88.8%	87.9%	0.892	0.887
Jun-00	89.0%	5.7%	5.3%	30,424	89.1%	67.9%	0.903	0.885
Jul-00	90.3%	5.1%	4.7%	29,927	92.7%	71.8%	0.916	0.900
Aug-00	85.8%	5.5%	8.7%	31,331	90.7%	80.9%	0.864	0.859
Sep-00	89.6%	4.7%	5.7%	29,032	86.9%	68.1%	0.908	0.897
Oct-00	92.1%	4.4%	3.5%	27,294	71.2%	42.8%	0.962	0.931
Nov-00	90.4%	4.1%	5.5%	26,254	82.8%	74.1%	0.916	0.901
Dec-00	83.8%	3.8%	12.3%	26,116	85.3%	95.2%	0.839	0.839
Jan-01	86.8%	3.9%	9.3%	25,330	79.2%	90.2%	0.871	0.866

Table 2

Price Summary Statistics April 1998-November 2000 ($/MWh)

Variable	Mean	Std Dev	Skewness	Kurtosis	Min	Max
PX North	46.89	56.86	5.53	58.53	0.00	1099.99
PX South	44.30	58.83	5.22	41.84	0.00	750.00
ISO North	54.80	77.67	3.67	23.26	-325.60	750.00
ISO South	45.20	71.77	4.64	34.99	-428.15	750.00
ISO-PX North	7.92	52.57	1.68	44.44	-709.01	689.85
ISO-PX South	0.91	50.85	1.43	50.88	-709.01	688.93

Table 3

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

	Early Hours 1-6					Late Hours 8-24				
Month	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.

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

	Early Hours 1-6					Late Hours 8-24				
Month	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)

Monthly				
Epoch	North Early	North Late	South Early	South Late
All Months	7.54	9.28	1.82	1.18
	(7.30)	(5.85)	(2.29)	(0.85)
May-Dec, 1998	2.64	2.61	1.56	1.72
	(4.58)	(2.27)	(2.91)	(1.53)
Jan-Aug, 1999	2.90	0.46	2.20	0.77
	(3.90)	(0.40)	(3.32)	(0.68)
Sept 1999-April 2000	-0.59	3.41	0.39	3.21
	(-0.77)	(2.47)	(0.49)	(2.33)
May-Nov, 2000	27.64	33.55	3.34	-1.28
	(7.40)	(5.60)	(1.03)	(-0.23)

Weekly				
Epoch	North Early	North Late	South Early	South Late
All Months	7.99	8.54	3.53	1.72
	(7.96)	(5.52)	(4.69)	(1.29)
April 8-Dec 31, 1998	3.09	3.21	1.37	1.90
	(5.73)	(3.19)	(2.60)	(1.91)
Jan-Aug, 1999	1.59	0.71	0.60	0.77
	(2.05)	(0.61)	(0.88)	(0.69)
Sept 1999-April 2000	0.54	1.93	0.68	3.38
	(0.72)	(1.38)	(0.88)	(2.46)
May-Nov, 2000	29.87	31.68	12.79	0.68
	(8.30)	(5.21)	(4.31)	(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

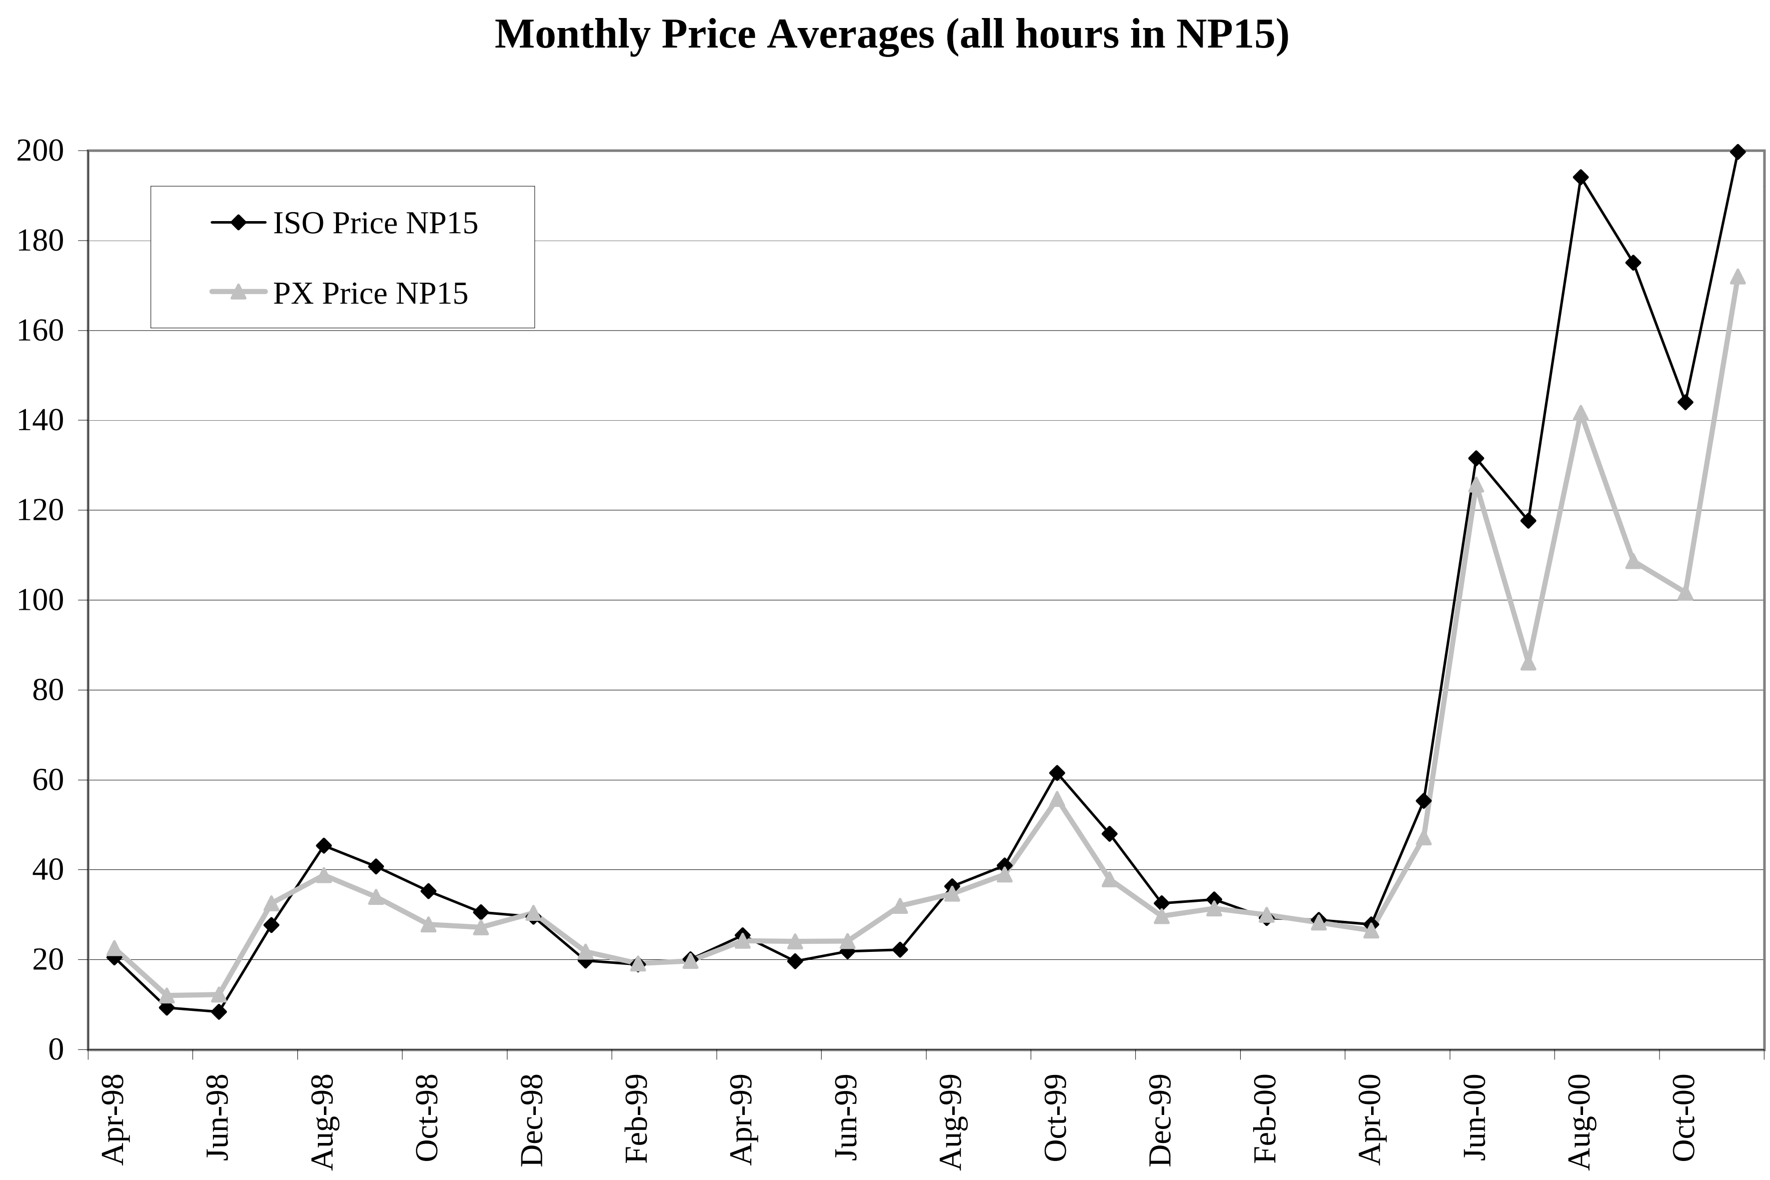
Monthly Price Averages (all hours in NP15)
ISO Price NP15
PX Price NP15
200
180
160
140
120
100
80
60
40
20
0
Apr-98
Jun-98
Aug-98
Oct-98
Dec-98
Feb-99
Apr-99
Jun-99
Aug-99
Oct-99
Dec-99
Feb-00
Apr-00
Jun-00
Aug-00
Oct-00

Figure 1b

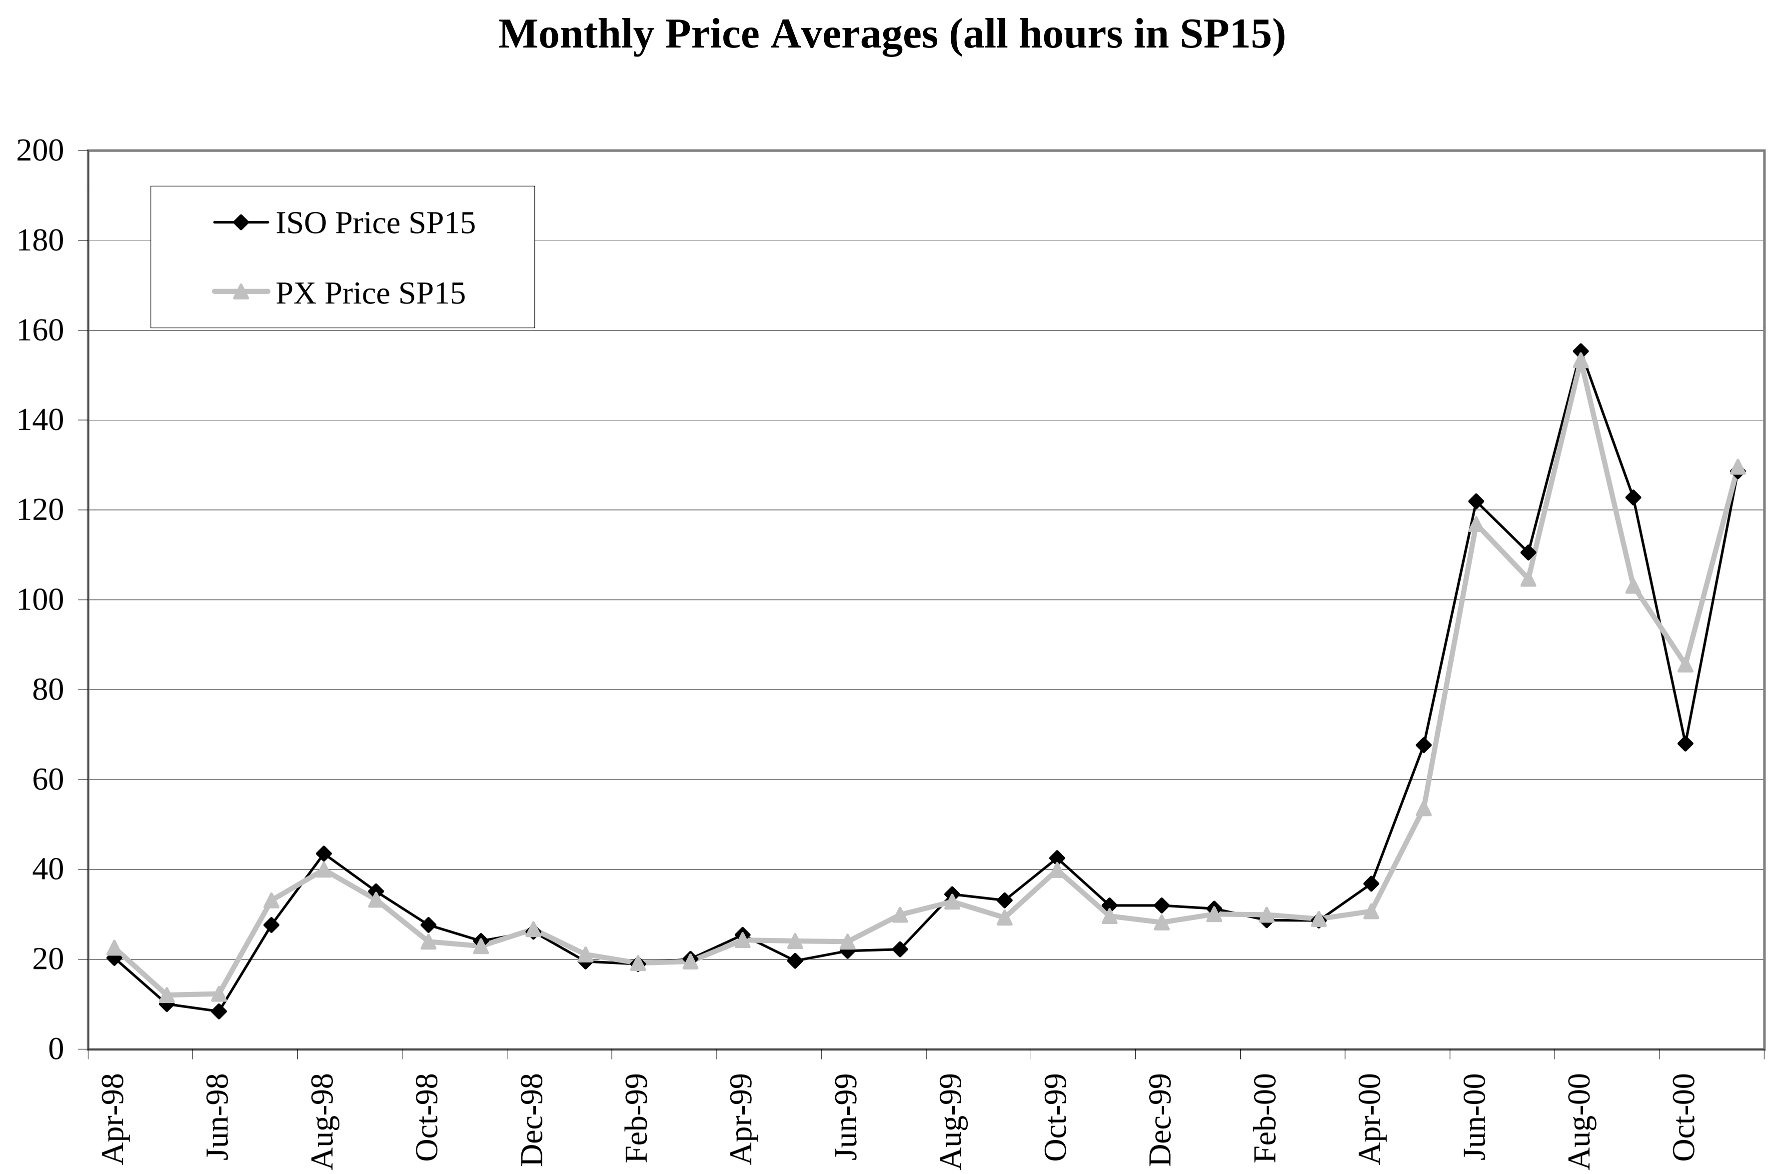

Monthly Price Averages (all hours in SP15)
ISO Price SP15
PX Price SP15
200
180
160
140
120
100
80
60
40
20
0
Apr-98
Jun-98
Aug-98
Oct-98
Dec-98
Feb-99
Apr-99
Jun-99
Aug-99
Oct-99
Dec-99
Feb-00
Apr-00
Jun-00
Aug-00
Oct-00

Figure 2

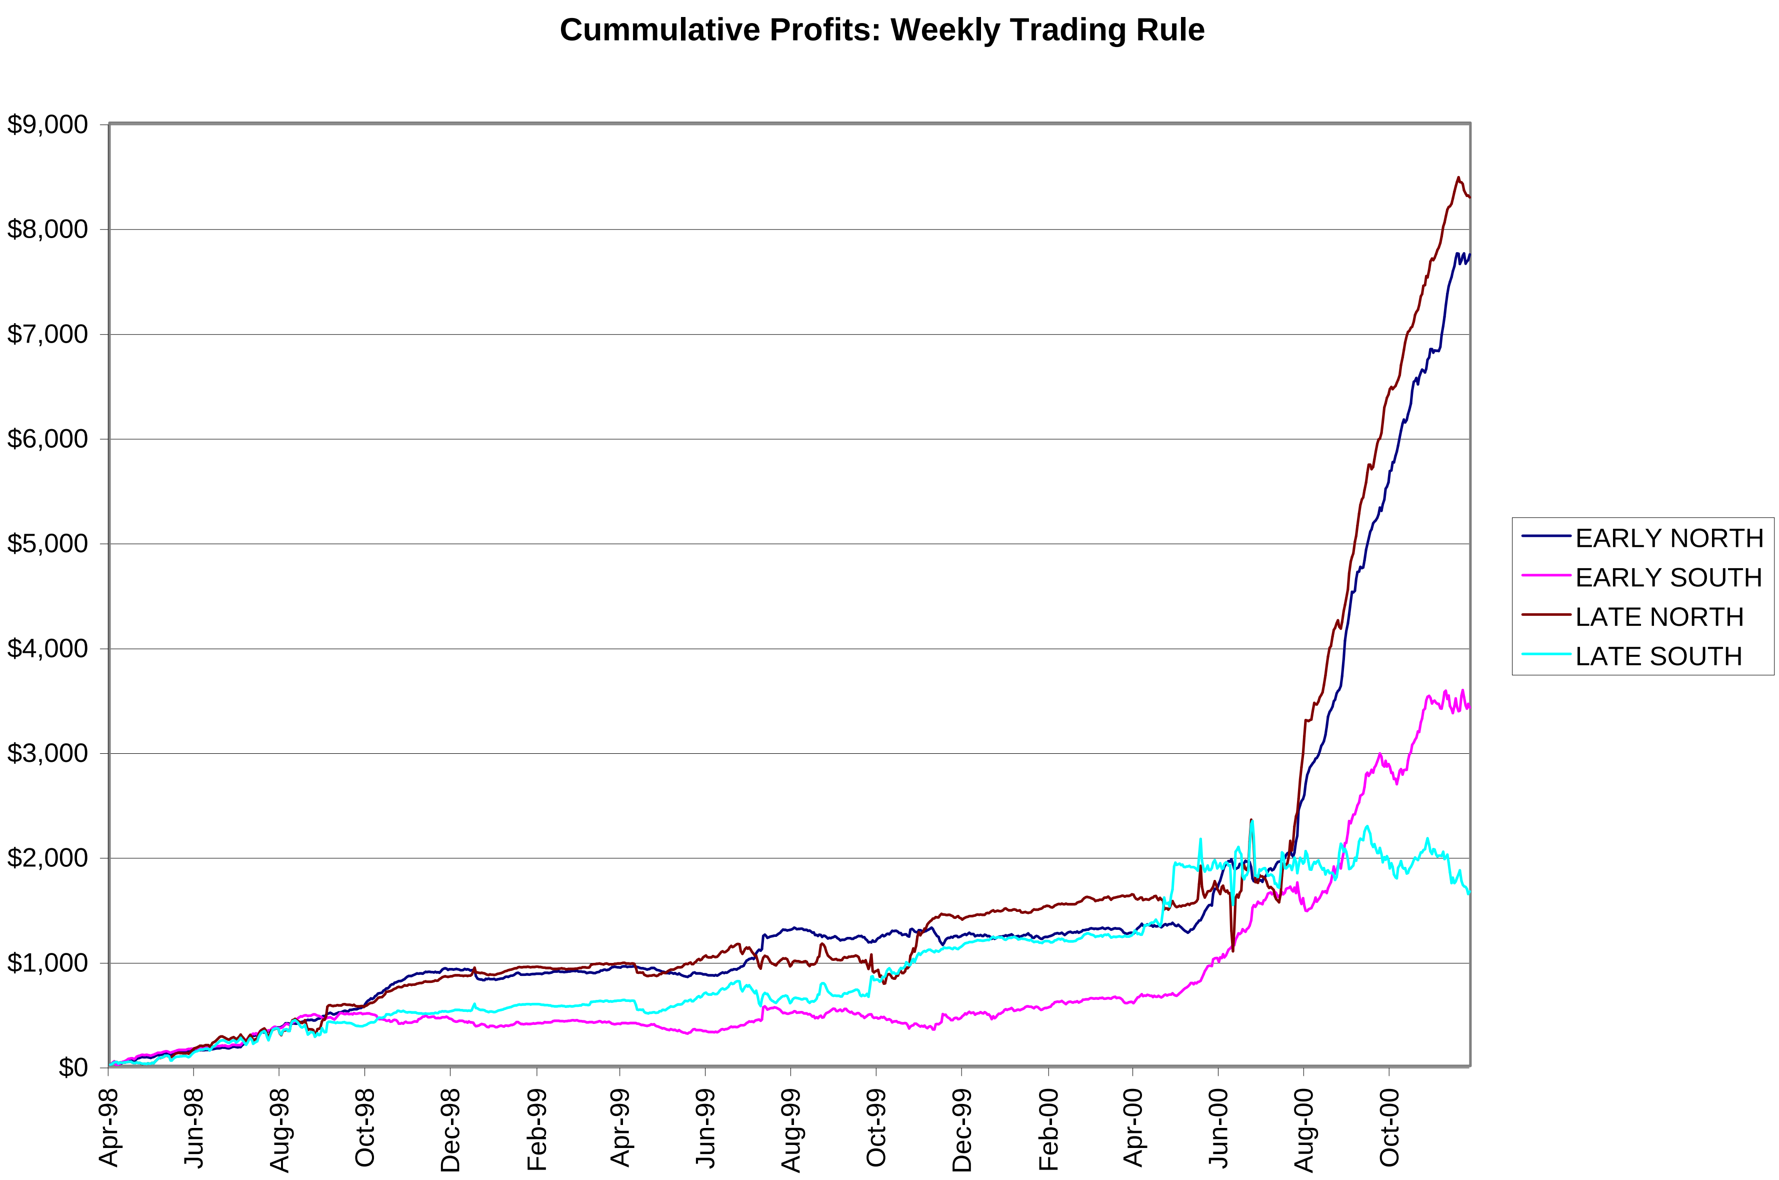
Cummulative Profits: Weekly Trading Rule
$9,000
$8,000
$7,000
$6,000
$5,000
$4,000
$3,000
$2,000
$1,000
$0
Apr-98
Jun-98
Aug-98
Oct-98
Dec-98
Feb-99
Apr-99
Jun-99
Aug-99
Oct-99
Dec-99
Feb-00
Apr-00
Jun-00
Aug-00
Oct-00
EARLY NORTH
EARLY SOUTH
LATE NORTH
LATE SOUTH

Figure 3

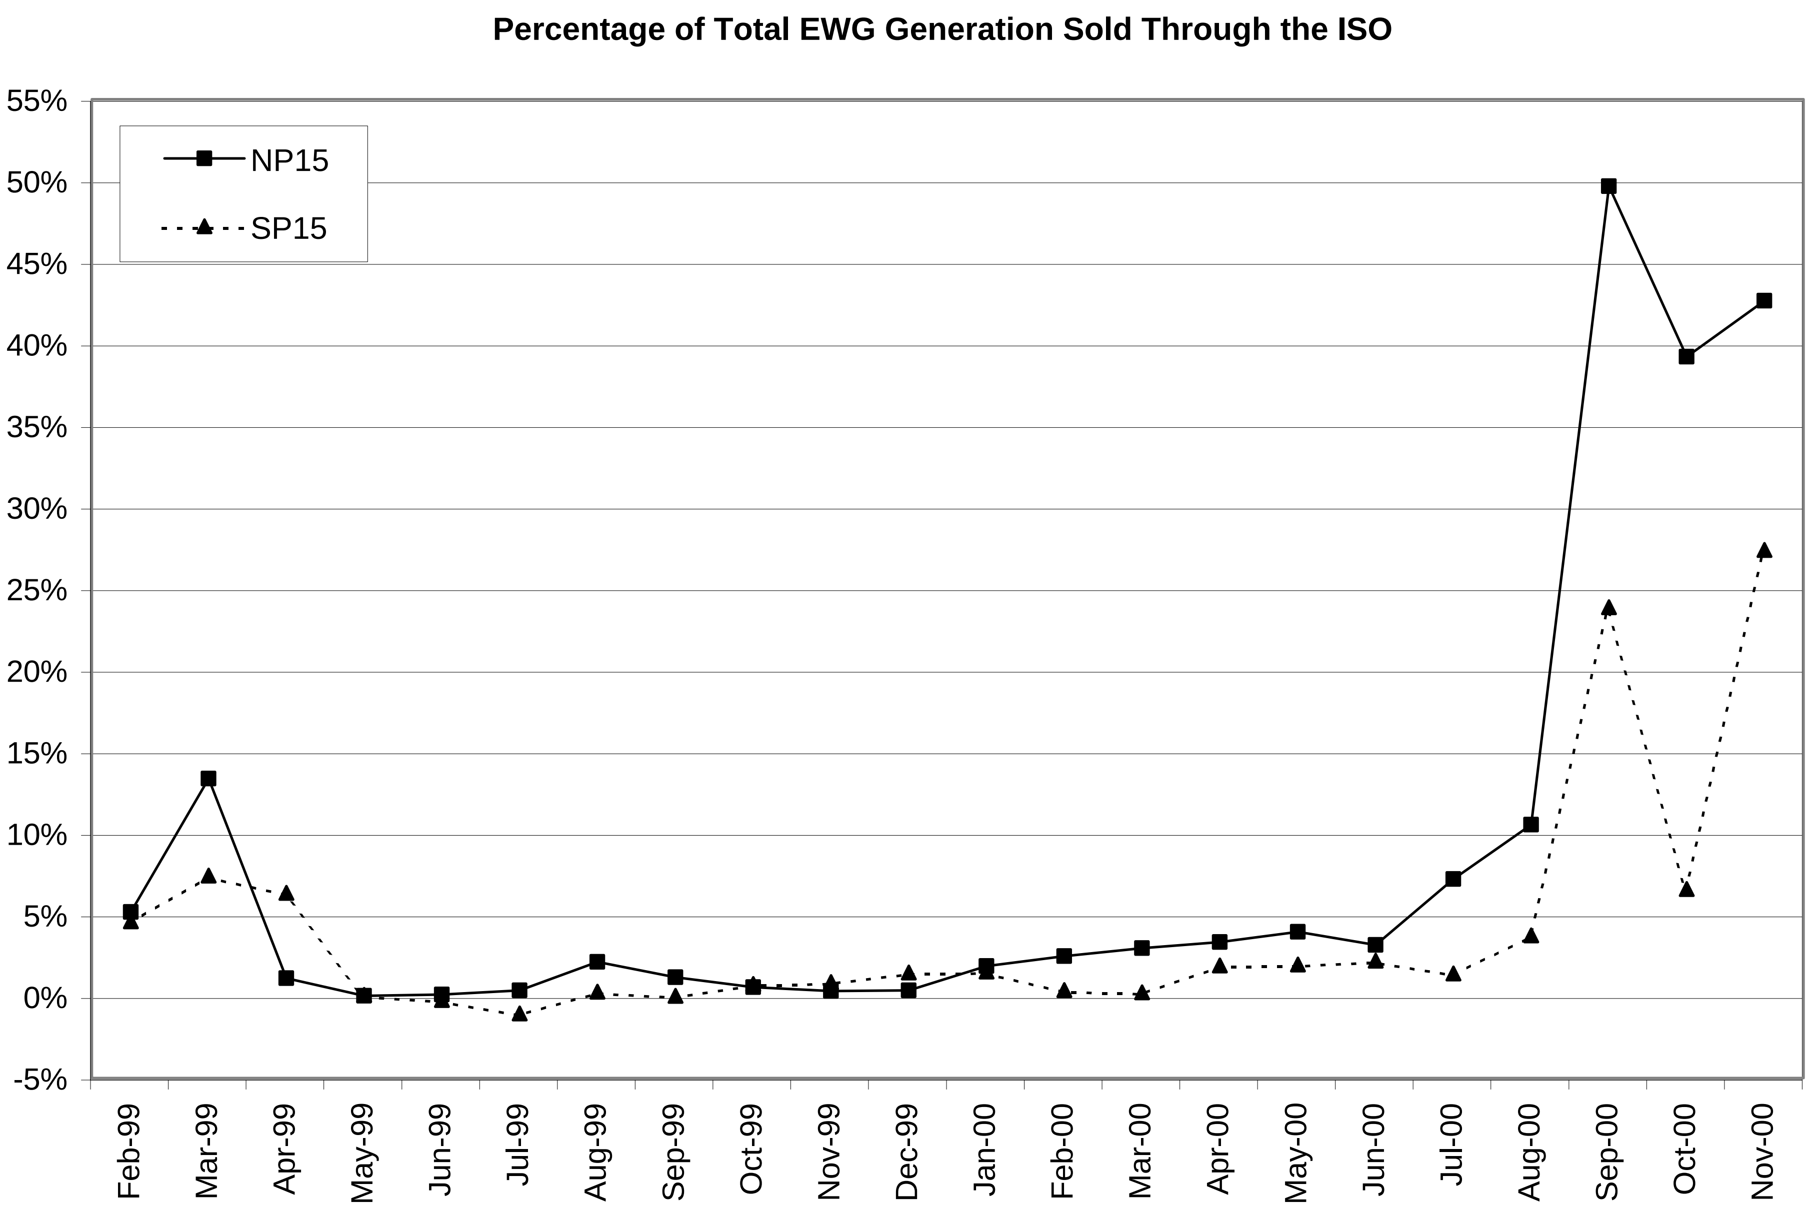

Percentage of Total EWG Generation Sold Through the ISO
NP15
SP15
55%
50%
45%
40%
35%
30%
25%
20%
15%
10%
5%
0%
-5%
Feb-99
Mar-99
Apr-99
May-99
Jun-99
Jul-99
Aug-99
Sep-99
Oct-99
Nov-99
Dec-99
Jan-00
Feb-00
Mar-00
Apr-00
May-00
Jun-00
Jul-00
Aug-00
Sep-00
Oct-00
Nov-00

Figure 4

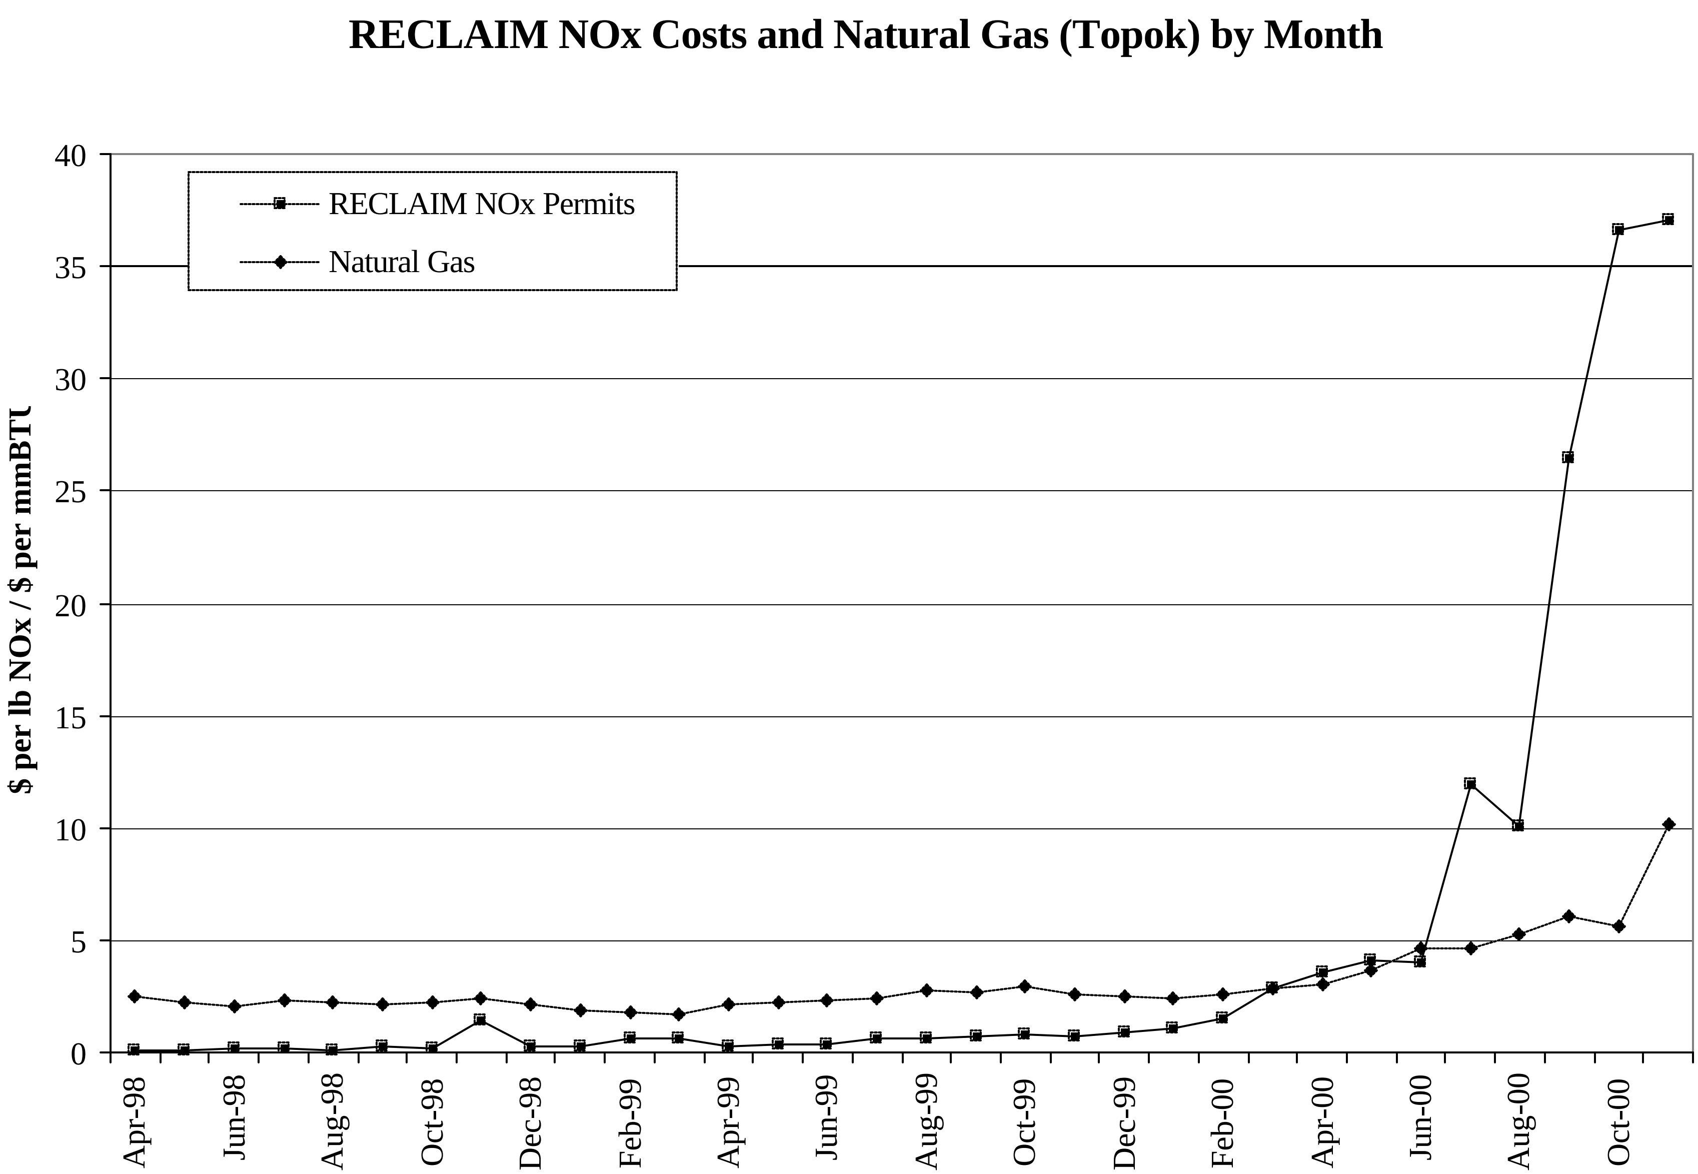
RECLAIM NOx Costs and Natural Gas (Topok) by Month
RECLAIM NOx Permits
Natural Gas
$ per lb NOx / $ per mmBTU
40
35
30
25
20
15
10
5
0
Apr-98
Jun-98
Aug-98
Oct-98
Dec-98
Feb-99
Apr-99
Jun-99
Aug-99
Oct-99
Dec-99
Feb-00
Apr-00
Jun-00
Aug-00
Oct-00