Preliminary Draft – Please do not cite

Learning and Market Efficiency: Evidence from the Opening of California's Electricity Markets

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

November 7, 2000

Abstract

We study price convergence between the different markets within the deregulated California wholesale electricity industry during its first 28 months of operation. Profit-maximizing traders should eliminate persistent differences in the price of power across the different markets. Institutional impediments and incomplete understanding of the markets, however, may delay or prevent 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 – have gravitated toward one another during this time period, but significant, albeit risky, profitable trading opportunities persist. A number of other tests also indicate that market inefficiencies declined over our sample period, but have not disappeared.

---

[*] *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 Finance and Economics, School of Management, Boston University and University of California Energy Institute. Email: knittel@bu.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. 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. Still, 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. We examine the relationship between the day-ahead and spot California wholesale electricity prices over a 28-month period beginning on the day these markets opened in 1998. We find that significant market inefficiencies are identifiable in our sample, but that they diminish over time, indicating that "learning" – either increased sophistication of incumbent market participants or entry of sophisticated market participants – may have taken place.

Besides the focus on the efficiency of financial markets, this application also has important implications for the movement towards deregulation of wholesale electricity markets. One of the central questions in deregulating the electric power industry is the efficiency with which competitive wholesale markets will allocate production and set prices. There is also concern about the size and competitiveness of these markets. Will the existence of many different submarkets mean that each market is fairly thin, possibly too thin to operate efficiently? Or, are the many markets actually integrated to the point that analysis of competition and efficiency can be done on an industry-wide basis, rather than through separate examination of each sub-market?

When California opened its electricity industry to competition on March 31, 1998, it was the first state to implement a restructured wholesale industry. There are, in fact, several markets within the state in which parties trade wholesale electricity. The two biggest are the California Power Exchange (PX) day-ahead forward market and the California Independent System Operator's (ISO) real-time spot market, which is known as the imbalance energy market.[1] Since both the forward and spot markets allow trade for delivery of the same product at the same locations and are open to most of the same traders, we would expect that, absent institutional trading barriers, there would be no significant, persistent expected price differences across the two markets.

Futures or forward markets generally allow participants to hedge against volatility in spot markets and provide signals of the expected spot prices. By providing a signal of the spot price, the PX day-ahead market facilitates efficient planning and optimal dispatch of electricity generating plants. If the forward prices are imperfect predictors of spot prices, the market is less efficient at serving those functions. We use ISO–PX (spot–future) price differences as a lens through which to view the markets.

Numerous papers have considered efficiency across various spot and futures markets. This is one of the first that considers efficiency in the newly restructured electricity

---

[1] The California Power Exchange opened on March 31, 1998 for trades of power scheduled to be delivered on April 1, 1998, which was officially the first day of grid operation by the California ISO.

markets.[2] By considering a new market in an industry with little experience with organized exchange trading, we are able to document the speed and extent of participant learning over time.

We find that differences between the forward and spot markets were initially quite substantial and significant. Over time, however, they appear to have declined, though there have been month-long periods of significant price differences as recently as July 2000 (the last month in our sample). We also propose a number of efficiency tests that take advantage of the idiosyncratic timing of operations in these electricity markets. These test also indicate that there were substantial departures from market efficiency in the early months of the California electricity markets, but that these have declined since then.

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 affect 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. Section 5 discusses several possible explanations for the inefficiencies documented in the previous section.

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.[3]

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 must incorporate all information available at $t-j$ about the expected spot price of electricity at t. That is,

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

---

[2] Bessembinder and Lemmon (1999) develop a model of the risk premium in electricity forward contracts and test it by analyzing trades for California-Oregon Border NYMEX futures contracts and prices for day-ahead trades reported to a market research firm. Their analysis covered wholesale transactions both before and after deregulation took effect and focused on the hedging function of the futures contracts.

[3] 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.

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 $t-j$.

These propositions reflect simply the equilibrium outcome when clever businesspersons exploit available profit opportunities. If the spot price were systematically higher or lower than the forward price at some given earlier time, then buyers would move to make their purchases in the market with lower prices and sellers would move to sell their output in the market with higher prices. These changes would push up the price in the market with the lower price and would depress the price in the market with the higher price until the forward and spot price returned to being equal to one another (in expectation).

Similarly, if any information available at time $t-j$ can be used to predict the price difference between ${}_{t-j}P_t$ and ${}_tP_t$, then traders would change their behavior to take advantage of that and, in the process, would eliminate the relationship between that information and the price difference. For instance, consider what would happen if the weather on June 8 could be used to predict the difference between the day-ahead forward price for June 10 (contracts signed June 9) and the spot price that obtains on June 10. If hot weather two days in advance systematically indicated that the spot price would be higher than the day-ahead forward price, then on hot days sellers would choose to avoid the day-ahead forward market the following day and instead sell their output in the relevant spot market on the next day. Similarly, buyers would attempt to buy all of their output in the day-ahead spot market rather than waiting for the spot market where they expect the price to be higher. These actions would tend to push up the day-ahead price and push down the spot price until the two-day ahead weather information is no longer useful in predicting the forward/spot price difference.

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. Note that this implies, as has been the case in California and elsewhere, the variance of the spot price will be larger than the variance of the forward price.

3 The California Electricity Market

Unlike more centralized electricity markets such as that which existed in the United Kingdom through the 1990s, in California there are many avenues through which agents can sell or purchase electrical energy.[4] Without a fully centralized market, it is possible that a high-cost generator could be operating and selling its output while a lower-cost generator stands idle.[5] Others have argued that such inefficiencies would eventually be mitigated by market forces, such as trade among the existing markets. In this section, we outline the California electricity market structure and discuss how trade could occur to profit from price differences across the various markets within this structure.

3.1 Forward Markets

Currently, most of the trading activity in California occurs on a day-ahead basis for hourly transactions. The California Power Exchange (PX) runs the largest of these day-ahead markets. The PX accepts bids for the hourly supply and demand of electricity for the 24 hours of the following day. This day-ahead market closes at 7:00 AM on the previous day. Day-ahead transactions may also be reached through other scheduling coordinators (SCs) operating in parallel to the PX. Many of these daily transactions submitted by SCs may in fact reflect longer term transactions that are nonetheless still required to be resubmitted to the Independent System Operator (ISO) on a daily basis.[6]

For purposes of transmission pricing, the ISO system is now divided into 24 zones.[7] Two zones, comprising northern California (NP15), and southern California (SP15), contain the overwhelming share of ISO system demand.[8] Most of the other "zones" are actually interface points between the ISO and surrounding utility systems. All SCs, including the PX, submit their preferred energy schedules, including the location of all supply and demand sources, to the ISO by 10:00 AM on the previous day. The ISO verifies the feasibility of these aggregated schedules in light of transmission and other operating constraints. If these preferred schedules are infeasible because they would result in flows on transmission lines that exceed the capacity of those lines, then the ISO runs an auction for the use of constrained transmission interfaces by utilizing schedule "adjustment bids" submitted by SCs.[9]

The schedule adjustment bids effectively establish each SCs willingness-to-pay for the use of a congested transmission interface. The preferred schedules are adjusted according

---

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

[5] See, for example, Stoft (1996).

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

[7] The 24th zone, ZP26, was added during 1999. The ISO is currently redesigning its congestion management process and the addition of several new pricing zones is likely to be an outcome of that process.

[8] Two areas of northern California, the City of San Francisco and Humboldt County, comprise independent zones for reliability purposes, but power is priced in these zones as if it were part of the surrounding NP15 zone.

[9] Before the ISO turns to its auction process, there is a prior iteration in which SCs, once notified of potential congestion, are allowed to voluntarily revise their schedules before submitting to the ISO transmission auction. Importantly, the PX, which is by far the largest scheduling coordinator, does not perform these iterative revisions to its energy schedule or to its adjustment bids.

to these bids, and a uniform price for the use of a congested interface is set at the usage value bid by the last SC whose schedule is adjusted. In this way all SCs that have scheduled transactions over a congested interface pay the same unit price for the use of that interface.[10] The PX takes these transmission prices and uses 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. The final, adjusted schedules are determined by 1:00 PM of the previous day, although further adjustments may sometimes be performed by the ISO for reliability purposes up to 5:00 PM the previous day.

In addition to the day-ahead markets operated by the PX and other SCs, schedules may also be submitted or revised up to an hour ahead of the actual delivery time. The PX operates a “day-of” market (originally called an “hour-ahead” market) that allows trades at a time closer to, but still many hours before the hour of operation.[11]

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 consumers do not yet have access to, let alone respond to, real-time prices, the ISO must ensure that supply and demand remain in continuous balance, 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 sets 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 is 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 deviates 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 are effectively financial forward positions, and the ISO imbalance energy market is the underlying spot market in which positions in these forward markets are resolved. Since August 19, 1999, however, this situation has 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 July 2000.[12] These figures give the percentage

---

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

[11] Until January 1999, the “hour-ahead” market operated on a rolling basis, with each market closing three hours before the hour of operation began. Since then, this has been operated as a “day-of” market, which is open three times per day, each time covering different blocks of hours that are about 5 to 12 hours in the future. In neither configuration has this market been particularly successful: trading volumes have been low and no transactions have taken place for in more than 25% of all hours.

[12] Data were not available for April-June 1998.

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.[13] For much of the study period, real-time volume average around 3% of total volume, but this figure has climbed to as high as 6% during the spring and summer of 2000. During high demand periods in last few months of our sample period, the real-time imbalance energy market handled as much as 25% of total volume. This high level of real-time volume has raised concerns about system reliability and has led to consideration of further efforts to discourage real-time transactions.

Figure 1 gives a diagram of the interaction of the various electricity product markets in the California ISO system. 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 do not actively bid into this market. They cannot explicitly bid demand-side adjustments into the imbalance energy 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. Table 2 gives the average prices over our sample period, April 1998 to July 2000, in the PX day-ahead and ISO real-time for the North (NP15) and South (SP15) zones.

Suppliers can 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 generate more than they were committed to provide are implicitly agreeing to take whatever price obtains in the imbalance energy market. Producers that bid into the imbalance energy market can choose to offer supply at a given price up to 45 minutes prior to the hour of production. Most suppliers of reserve capacity are also eligible to earn imbalance energy revenues. These ($/MWh) energy revenues are in addition to ($/MW) capacity payments earned by suppliers that commit to being available with varying response times. Suppliers to the ancillary services markets submit two-part bids: a “stand-by” capacity price for a given reserve service and an energy price to be paid in the event that the unit is 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 is acquired by the ISO on a day-ahead basis, following the PX auction. Suppliers who wish to sell imbalance energy through the ancillary service channel must therefore submit offers shortly after the PX market closes. Of course, they have the opportunity to earn revenues for their stand-by capacity, as well as energy production. Suppliers opting for the imbalance energy channel can wait until 45 minutes prior to the hour of delivery before finalizing their offers. A supplier that simply generates in excess of its scheduled supply makes that decision on a real-time basis, with no advance commitment.

---

[13] 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 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 have sometimes skipped over low-cost energy bids from certain reserve sources due to concerns about depleting available reserves.[14] Consequently, suppliers of some reserves earn no imbalance energy revenues even when their energy bid is below the imbalance energy price.[15]

Since August 19, 1999, however, the ISO has allocated the costs of replacment reserve capacity, a form of operating reserve, disproportionately to suppliers that produce less than their scheduled quantity and demanders that consume more than their scheduled demand. This has produced an an additional penalty on transactions that end up as net purchases from the imbalance energy market. The impact of this change in replacement reserve cost allocation is discussed further in section 4.

A supplier that has been scheduled to provide energy in one of the forward markets can also take a short position in the spot market by either offering to decrement its output through a imbalance energy bid or by simply generating less than its advance commitment. In the latter case, the supplier must make up its production short-fall through a purchase on the imbalance energy market and is effectively a consumer in this market. A decremental supply bid in the imbalance energy market is an offer to buy out of an advance supply commitment. A supplier pays the ISO an amount equal to the imbalance energy price in exchange for not having to provide the energy that it has scheduled. By bidding a decremental energy bid, a supplier has the opportunity to set the imbalance energy price, and reserves the right to generate energy in the event that the imbalance energy price is set above its decremental bid.[16]

It is important to note that the ISO calls upon decremental supply bids only when there is oversupply in real-time, and calls upon incremental supply bids only when there is undersupply. The ISO does not attempt to arbitrage price gaps when there are some incremental bids that are lower than decremental bids. In other words, even though there may be suppliers in the real-time market that are willing to pay more to buy out of their supply commitment than other suppliers may be asking for in order to fill that commitment, there is 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.

---

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

[15] This practice most likely impacts the capacity prices of ancillary services more than it does 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.

[16] 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 are certainly possible, current institutional barriers largely restrict such activity to the actual "physical" market participants.[17] Since the market has opened, further restrictions and institutional barriers have been applied in an effort to limit speculative trades. These efforts have been 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 is to maintain system reliability. All SCs that deal with the ISO are 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 will occur. There is no way to verify that a given level of consumption is completely realistic, but supply resources must be specifically identified. Bids to provide ancillary services and imbalance energy from within the ISO system must be linked to specific generation facilities. Since the ISO's ability to verify the availability of specific production sources is largely limited to its own control area, bids to supply production and reserves from outside the ISO system face much less stringent verification requirements.

The PX allows more flexibility in both the eligibility of traders and the form of trades. PX market participants must meet financial credit requirements, but need not control actual supply resources. Offers to supply through the PX take the form of "portfolio bids" that are required simply to be strictly upward sloping, piece-wise linear curves. In setting its unconstrained market price (*i.e.*, ignoring transmission constraints), the PX does not require the identity or location of specific production or consumption sources. Once the unconstrained price is set, suppliers to the PX must 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 is no specific verification of the availability of import supply.[18]

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) - are required to meet the demand needs of their distribution systems through purchases from the PX.[19] This requirement was intended to help ensure sufficient liquidity in the PX day-ahead market

---

[17] In fact, for most of the time period we study, the three investor-owned distribution utilities, which are responsible for supplying 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 August 1999, their forward purchasing could be no more than one day before actual consumption. There is now a more active longer-term block forward market run through the PX.

[18] The ISO is currently considering revising its congestion management practices.

[19] Although SDG&E completed its transition period in 1999, its "must-buy" requirement with the PX was continued into 2000.

and to establish a transparent day-ahead price. Other market participants are free to participate in other day-ahead markets, or sign direct bilateral arrangements. Although there are roughly 60 firms trading in the PX, the three IOUs account for about 90% of the energy purchases. The PX itself accounts for about 87% of the total trading volume in the ISO system during the sample period.[20]

Although the IOUs are technically required to purchase all their supply needs from the PX markets, the market process makes rigid enforcement of this requirement both impractical and undesirable. It has been well documented that demand bids into the PX are downward sloping and in fact quite elastic over some price ranges.[21] This is despite the fact that the overwhelming majority of end-use demand is incapable of receiving, let alone responding to, hourly price signals. Price-elastic demand bids in the PX clearly reflect strategic decisions by buyers to purchase in the ISO real-time imbalance energy market if the PX day-ahead price is too high. This is in part driven by the fact that the ISO imbalance energy market was subject to a price cap that was frequently 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 reflect the fact that no firms were willing to pay more than 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 is also committed to bidding into the PX day-ahead market. This energy is supplied by generation sources producing under regulatory or commercial arrangements that predated the restructuring of the California market. The price earned by these producers is set by the terms of their pre-existing "must-take" arrangements. This must-take supply is bid into the PX day-ahead market at a zero price.

In addition to the institutional and regulatory constraints on market participants, there are also differences in the transaction costs of dealing with the various markets. Two notable costs appear to have initially favored trading in the imbalance energy market over trading in the PX. Both the ISO and PX assess trading charges on all volume in their markets. However, the PX charge, which is currently $0.30/MWh applies only to volume traded by firms that use the PX as their scheduling coordinator,[23] while the ISO charge applies 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 PX as one's SC, but there is 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

---

[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 in the fall of 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 unprecented price levels experienced during May and June of that year.

[23] The PX administrative charge applies to all volume traded by entities that use the PX as their scheduling coordinator, including volume in the ISO market. Thus, in order to avoid the PX administrative charge an entity would have to use some SC other than the PX.

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 are now charged 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 range 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.

3.4 Efficiency in the California Electricity Markets

Given the institutional peculiarities of the California electricity marketplace, the nature of speculative trading in these markets is somewhat different from that in most commodity markets. Indeed, as mentioned above, the market design in many ways was intended to discourage, rather than facilitate, speculative trades. Despite this, we present evidence below that the prices in the two major markets for electrical energy have tended to converge over time. In this section we discuss the mechanics of a potential speculative trade.

Since the price of energy in the spot, or imbalance energy, market was in fact the only "penalty" applied for not meeting advance supply commitments (prior to August 19, 1999), speculative trades simply took the form of deviations from advance schedules. Consider, for example, a firm that expects the PX price will be higher than the ISO imbalance energy price in a given hour. This firm would bid to supply power into the PX, and would in fact produce nothing. The firm's position would be settled at the PX minus the ISO price multiplied by the quantity committed in the PX. Since August 1999, a firm executing such a trade would also be charged for some share of the cost of replacement reserves in addition to the real-time energy price.

Conversely, if a firm expected that the PX price would fall below the ISO price in a given hour, it would bid demand into the PX and not consume in real-time. This is equivalent to buying the power a day-ahead and selling in the imbalance energy market. It is important to note that this latter trade is subject to less ex-ante verification than one involving taking a long position in the forward market. This is because most successful supply bids with the exception of imports, must provide evidence of a credible source of production. Traders on the demand side are not required to provide evidence that they are capable of taking delivery. In addition, the entity carrying out these trades is not subject to disproportionate charges for replacement reserve.

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 given in table

3 for our sample period. Our sample period begins with the opening of the markets on April 1, 1998 and ends on July 31, 2000.

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

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

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. We estimate (5) by ordinary least squares using each set of hourly prices in each zone as a separate observation. The dependent variable is the ISO price in a given zone minus the PX price for the same zone. The results reported in table (4) suggest that, in a number of months, the PX price systematically differed from the ISO price.

Estimation of equation (5) by ordinary least squares, however, makes the assumption that the shocks to the price differences between the PX and ISO prices are not serially correlated. There is good reason to think that this would not be the case, and empirical tests indeed confirm that it is not. Because the PX prices in a given day are 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 submit supply and demand bids for the 24 hour period beginning with the midnight-1am hour of the following day. Because PX prices are determined in 24-hour "blocks," shocks to either supply or demand (such as weather changes) that take place after PX prices are 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.

The implication of this serial correlation is that, while the parameter estimates are unbiased, the OLS standard errors in table (4) are likely to be understated. Because of

---

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

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

the timing of the PX market, the exact serial correlation structure that one would expect in even an efficient market is quite complex. We describe this below and then discuss two different estimation approaches.

Let the information set at hour t be represented by Ω_t. Let $t = 1$ represent the beginning of an arbitrary day (*i.e.* 12:00 midnight). 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 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 degree 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:

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

---

[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 $t = 6$.

$$ISO_{25} - PX_{25} = \alpha + \varepsilon_{25} + \sum_{i=1}^{18} \theta_{25,i}\varepsilon_{25-i}$$

$$\vdots$$

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-one observations in our sample period, another all of the 848 hour-two 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. Table (5) presents the results of this analysis, 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 not significant at the 5% level.

---

[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.

Examination of table (5) suggests that the statistical significance of the ISO-PX price differences has declined over time. A slightly more systematic approach for the most part confirms this. Table (7) presents the average monthly p-values for three time periods: April-December 1998, Januray-August 1999, and September 1999-July 2000. The periods are 9, 8, and 11 months respectively. We break between the first two periods rather arbitrarily at the end of 1998. We break after August 1999, because it was late in that month that the ISO changed the way it charged reserve capacity costs to real-time transactions. Table (7) indicates that in the North-early, North-late, and South-early, there was a marked pattern of declining significance over time. In the South-late, the difference has not been significantly different from zero on average in any period.[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, they were instead 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 presented in table (6). The results are largely consistent with those in table (5). The monthly constant terms appear to become less statistically significant over time. Table (7) again confirms this. The reasons that these results are quite similar to those in table (5) is that for most periods replacement reserve capacity payments are very close to or equal to zero.

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 twenty-eight months of the markets.

---

[28] Despite this, the joint test against all constant terms in the period being zero rejects the null hypothesis in each zone, both early and late, for each of the three time periods, at at least the 5% level.

The first simple rule we evaluate stipulates that a trader always makes sales or purchases in the market that would have been the most advantageous in the previous month. The first test we implement assesses 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 July 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 construct 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 (8) summarizes the coefficients and t-statistics from including this variable in a specification of equation (5) without any month dummies. The first column reports results from specifications that included all twenty-eight months, while the second, third and fourth columns report tests on the beginning, middle and end of our data set. Considering the entire time period, the t-statistic are greater than 2 in all specifications except the late hours in the South (where it is nearly 2), suggesting that the 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 \$2.50 per MWh traded in the North during early hours. The results on the three separate time periods, however, suggest that there were fewer opportunities to profit from the simple rule more recently.

We next considered whether the trading rule would make money in the hands of a more representative market participant – one who owned supply and decided whether to sell it through the PX or the ISO based on the previous month's price differences. In this case, we need to evaluate the trading rule in comparison to some benchmark behavior. We assumed that, absent the trading rule, the supplier would have sold 100% of his production through the PX in all months. We compare this decision rule to one under which the supplier sells all his power through the market that had the higher average price in the previous month. Again we evaluate whether the supplier would make money by using a variable that is equal to one if the ISO price was higher in the previous month and zero otherwise, since we are assuming the supplier would have sold through the PX anyway. For example, in May 1998, our trading rule correctly predicts that the ISO prices will be lower, but since the supplier would have sold all of his volume through the PX anyway, the rule does not help him re-allocate his volume. The coefficient and t-statistics from this variable are reported in the bottom of Table (8). Again, for the specifications including all 28 months, in all but the later hours in the South, the t-statistic are greater than 2. Using this rule, and assuming that the company sold 1 MWh every day during the early period in every

---

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

month, a supplier in the North would have increased his profit by \$2.06 per day.[30] Unlike the speculative trader, our representative trader earns statistically significant profits in 2000 for three out of the four specifications. The trading rule leads to significant losses during two months (November 1999 and May 2000) when the ISO price was higher than the PX, but since we assume that the representative trader would have sold through the PX anyway, these losses do not count against the profitability of the rule.

The bottom of Table (8) 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 28 months of the market. Unlike the monthly tests, even the speculative trader makes significant profits in the early hours of 2000, although these are largely driven by returns to trading during the last three months of the sample when the weekly rule is able to adapt to extremely high ISO prices more quickly than the monthly rule.

4.2 Additional Tests of Efficiency

In this section, we analyze whether a trader could have used additional variables to predict systematic differences between the ISO and PX prices. Specifically, we test whether the ISO–PX price differences are explained by the price difference from the same time period one or two days before. If traders are taking advantage of all of the relevant information in the market, then previous ISO–PX price differences (except for the previous day price difference in the case of the later part of each day) should not help predict the current price difference.

Table (9) reports the results for the 1-day and 2-day lagged price differences. The results for the previous-day price difference suggest that the difference has clear predictive power for the late periods (hours 8-24), which they indeed should as explained earlier. This should not, however, be the case for the early periods (hours 1-6). In the North zone, this test of market efficiency fails not just in the early part of our sample, but even in the most recent period. In the South zone, however, there is indication of learning, with both the lag coefficient and the degree of statistical significance declining over time.

The results for the 2-day lagged price are largely consistent with market efficiency and show some evidence of learning. In the early periods, both zones show statistically significant correlations during 1998, but these effects decline in both size and statistical significance over time. In the late periods, the two-day lag price difference show no sign of predictive power in the North and none in the South until the most recent time period in our data, when there is a surprising negative and significant correlation with the two-day lag.

---

[30] Note that the test statistics in the previous paragraph also reflect the probability that our simple trading rule would dominate a rule where the supplier sold half of his volume in the ISO and half in the PX.

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. There is some evidence that the price differences have eroded over time, suggesting that they were an artifact of the market's infancy. Under certain circumstances, significant price differences will persist between markets. This section considers the extent to which risk aversion, learning or transaction costs explain our results.

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.

To understand the effect of risk aversion on the part of some traders, consider two simple examples. First, consider if one small buyer of power is risk averse. That buyer will prefer to buy in the PX to reduce his risk exposure. He is small relative to the market, however, and will not affect the relationship between the ISO and PX price. In essence, he will be able to purchase price insurance for free.

Now consider what would happen if a large block of buyers were risk averse. They would all want to purchase in the PX to lock in a price. At first, this would tend to increase the PX price relative to the ISO price. But if there was also a large competitive block of buyers who were risk neutral, they would now see an opportunity since the PX would have a systematically higher price than the ISO. They would decrease their purchases in the PX and move as much of their demand as possible to the ISO, where the expected price would be lower. If there were enough of these risk-neutral buyers, they would offset the risk averse ones and push the prices back into equality (in expectation).

Risk-neutral buyers are not the only ones who would have an incentive to offset the effect of the risk averse buyers. Even if all buyers were risk averse, if there were a large number of competitive risk-neutral sellers, they would have an incentive to move into the PX, seeking the higher expected price, and by doing so, eventually eliminating the forward/spot price differential.

The point here is that risk aversion by some traders does not imply that the PX and expected ISO prices will deviate. So long as there are also significant numbers of risk-neutral traders, the competitive force of those traders will 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.

For one, any one position a trader takes is small relative to their entire business. If, for instance, a trader thought that the PX price overstated the expected ISO price, she could sell in the PX and buy back power in the ISO. There would be risk associated with doing this since the ISO price is a random variable, but this bet would be for just one day (or one hour) of power. If the trader made a business of such risky arbitrage, she would be taking hundreds or thousands of such risks each year and these risks would be nearly completely uncorrelated with one another. The law of large numbers would result in the trader being fairly indifferent to the risk exposure on any one day. In such situations, the trader is unlikely to exhibit much risk aversion on any one bet.

In addition, of course, returns on bets 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. In fact, we calculated the coefficient on the return on the S&P500 from a CAPM model of the ISO–PX price difference and could not reject that the coefficient was zero.

Finally, it is important to recognize that if risk aversion is driving a gap between the PX and ISO, it is ambiguous what the sign of α would be. If buyers are systematically more risk averse than sellers, one would expect the PX price to exceed the ISO price as buyers pay a premium for price certainty. If sellers are systematically more risk averse, then the PX price would be discounted relative to the ISO price as sellers are willing to accept a lower price in the forward market in exchange for price certainty. It seems unlikely, however, that the relative risk aversion between buyers and sellers would periodically flip. The fact that α changes sign throughout our sample casts further doubt on the hypothesis that our results are driven by risk aversion.

5.2 Learning

The results in Tables 5, 6 and 7 indicate that significant price differences have diminished over time, suggesting that they have been eroded by learning. Learning could take two forms: either more sophisticated participants have entered the market over time, or the existing participants are better able to take advantage of price differences. Although there have been a number of plant ownership changes over our sample period, it is unlikely that learning has been driven by the entry of new market participants. Even in the presence of a number of unsophisticated traders, a significant price difference will persist only if the volume traded by the sophisticated trader is too small to drive the market prices together. For instance, if PX prices were consistently higher than the ISO prices, a sophisticated buyer would have to be small enough that even moving all of his demand to the ISO would not drive the ISO price up to the PX price.

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.[31] 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 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. Considering the results in Tables 5 and 6, however, it is not obvious that this rule change disrupted convergence. Considering all four specifications in Table 5, for instance, more of the α coefficients were significant in the five months preceeding August 1999 than in the five 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) have also been changing, 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 which changed traders' expectations. Given the general trend toward price convergence, documented in Table 7, it seems likely that either the effect of shocks or the number of events which significantly altered traders' expectations about the price formation process have diminished over time.

5.3 Transaction Costs Within and Between Markets

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

---

[31] In the exchange-rate literature, this is known as the "Peso Problem."

markets.[32] 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.[33]

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. Given the importance of these questions, empirical analysis of market performance will be a crucial line of research in the coming years.

In this paper, we have studied one aspect of the California market's performance, the interaction between the two dominant energy markets, the day-ahead market of the PX and the real-time market of the ISO. Although these markets operate under very different institutional roles, and according to quite different sets of 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 much of 1998. PX prices were higher in the first four months and lower in the following four months. There have still been months since 1998 in which prices have differed significantly between the ISO and PX, but there seems to be a trend towards convergence. Nonetheless, it appears that some risky trading strategies with positive expected return still exist.

---

[32] This is *not* intended to be a characterization of the California market. It is just a very simple example to illustrate the point. The actual incentives in the California market are much more complex.

[33] It is possible that traders on one side all 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.

Our results are consistent with a gradual learning process among players in the market. Complete price convergence is a very high standard during a time in which not only the players are learning, but the market institutions are undergoing frequent changes. Still, it is surprising that rather simple trading rules based upon the average price-differential of the previous month or week earned a positive return during even the most recent the period of our study.[34]

Lastly, it is important to note that several changes to the ISO markets have taken place during our sample period and numerous changes are now under consideration that could make the kind of trading analyzed here more difficult. In particular, concern about grid instability have led some participants to suggest that the ISO should place additional charges on transactions that take place in the real-time market. The ISO is also considering redesign of its congestion management system. All of these changes will diminish the value of any learning by the market about the way prices interact under the current rules.

---

[34] Although our results indicate that the markets are becoming more efficient with regards to their interaction *with each other*, this does not necessarily indicate that the markets themselves are fully efficient. For example, other work has indicated that significant market power has affected both markets during some time periods.

References

[1] Bessembinder, H. and M. Lemmon (1999). "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets." Mimeo, Arizona State University.

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

[3] Borenstein, S. Bushnell, J.B. and F. Wolak (1999). "Diagnosing Market Power in California's Deregulated Wholesale Electricity Market." POWER Working Paper PWP-064, University of California Energy Institute.

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

[5] Newey, Whitney K., and Kenneth D. West (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, vol. 55, no. 3, pp. 703-708.

[6] Peroni, E. and R. McNown (1998). "Noninformative and Informative Test of Efficiency in Three Energy Futures Markets." *The Journal of Futures Markets*, vol. 18, no. 8, pp. 939-964.

[7] Stoft, S. (1996). "California's ISO: Why Not Clear the Market?" *The Electricity Journal*, vol. 9, no.10, pp.38-43.

[8] Wolak, F.A., Nordhaus, R. 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
Jul-98	91.9%	3.5%	4.6%	29646	79.2%	73.5%
Aug-98	92.0%	4.4%	3.6%	31057	78.4%	71.6%
Sep-98	93.1%	3.8%	3.1%	28274	69.2%	54.3%
Oct-98	94.4%	3.4%	2.2%	24601	72.9%	58.0%
Nov-98	94.5%	3.2%	2.3%	23927	69.0%	55.0%
Dec-98	93.3%	3.9%	2.8%	25045	76.3%	66.4%
Jan-99	92.4%	4.5%	3.1%	24399	77.8%	69.4%
Feb-99	91.5%	4.7%	3.8%	24322	79.0%	86.8%
Mar-99	91.8%	5.4%	2.8%	24728	79.0%	74.3%
Apr-99	92.3%	4.7%	2.9%	24747	79.4%	55.4%
May-99	92.4%	4.4%	3.2%	24743	73.4%	65.1%
Jun-99	91.5%	5.6%	2.9%	27227	77.9%	54.4%
Jul-99	91.5%	4.8%	3.7%	30150	81.5%	39.1%
Aug-99	91.7%	4.8%	3.4%	30063	75.2%	51.1%
Sep-99	92.3%	4.9%	2.7%	28518	80.7%	54.3%
Oct-99	92.5%	4.0%	3.4%	27826	75.7%	45.3%
Nov-99	93.0%	3.8%	3.3%	26324	76.0%	36.7%
Dec-99	92.8%	4.7%	2.5%	26460	85.6%	55.1%
Jan-00	92.8%	4.5%	2.7%	26296	82.9%	44.5%
Feb-00	91.2%	4.9%	3.9%	25725	85.9%	84.1%
Mar-00	88.7%	5.3%	6.0%	25623	87.2%	90.3%
Apr-00	91.4%	4.3%	4.3%	25692	86.6%	74.4%
May-00	89.4%	5.0%	5.6%	26968	92.5%	91.5%
Jun-00	89.8%	5.5%	4.7%	30424	92.8%	70.7%
Jul-00	90.8%	4.8%	4.3%	29956	92.7%	71.6%

Table 2: Price Summary Statistics April 1998-July 2000 ($/MWh)

Variable	**Mean**	**Std Dev**	**Min**	**Max**
PX North	34.87	45.38	0.00	1099.99
PX South	33.76	47.18	0.00	750.00
ISO North	37.16	59.01	-325.60	750.00
ISO South	34.72	60.61	-428.15	750.00

Table 3: Average North and South PX and ISO Prices by Month

Month	Region	ISO Price	PX Price
April, 1998	North	20.49	22.64
	South	20.30	22.64
May	North	9.30	12.06
	South	10.08	12.06
June	North	8.38	12.25
	South	8.38	12.34
July	North	27.73	32.52
	South	27.62	33.14
August	North	45.40	38.80
	South	43.53	39.96
September	North	40.77	33.97
	South	35.13	33.25
October	North	35.26	27.85
	South	27.81	23.94
November	North	30.58	27.24
	South	24.08	22.92
December	North	29.59	30.43
	South	26.13	26.74
January, 1999	North	19.79	21.79
	South	19.50	21.09
February	North	18.98	19.19
	South	18.98	19.19
March	North	20.09	19.74
	South	20.09	19.48
April	North	25.42	24.25
	South	25.42	24.32
May	North	19.66	24.07
	South	19.66	24.06
June	North	21.87	24.15
	South	21.87	23.93
July	North	22.21	32.01
	South	22.21	29.91
August	North	36.34	34.65
	South	34.47	32.80
September	North	40.97	38.98
	South	33.09	29.28
October	North	61.60	55.77
	South	42.54	39.88
November	North	48.03	37.90
	South	31.95	29.64
December	North	32.58	29.70
	South	32.01	28.19
January, 2000	North	33.37	31.38
	South	31.30	30.05
February	North	29.26	29.97
	South	28.72	29.93
March	North	28.81	28.25
	South	28.63	29.02
April	North	27.87	26.48
	South	36.79	30.67
May	North	55.37	47.20
	South	67.67	53.64
June	North	131.55	125.73
	South	121.95	116.85
July	North	117.70	86.08
	South	110.54	104.67

TABLE 4

Dependent Var: ISO-PX

Year	Month	Zone	OLS Coef Early Hrs 1-6	OLS SE Early Hrs 1-6	OLS P-value Early Hrs 1-6	OLS Coef Late Hrs 8-24	OLS SE Late Hrs 8-24	OLS P-value Late Hrs 8-24
1998	April	NP15	-3.571	2.706	0.187	-1.556	6.511	0.811
1998	May	NP15	-1.876	2.662	0.481	-2.860	6.405	0.655
1998	June	NP15	-1.153	2.706	0.670	-4.856	6.511	0.456
1998	July	NP15	-6.133	2.662	0.022	-4.203	6.405	0.512
1998	August	NP15	0.280	2.662	0.916	9.206	6.405	0.151
1998	September	NP15	3.517	2.706	0.194	8.255	6.511	0.205
1998	October	NP15	8.922	2.706	0.001	6.776	6.511	0.298
1998	November	NP15	3.717	2.706	0.170	3.108	6.511	0.633
1998	December	NP15	-3.681	2.662	0.167	0.432	6.405	0.946
1999	January	NP15	-1.321	2.662	0.620	-2.194	6.405	0.732
1999	February	NP15	-1.052	2.801	0.707	0.178	6.740	0.979
1999	March	NP15	-1.934	2.662	0.468	1.218	6.405	0.849
1999	April	NP15	-0.341	2.706	0.900	1.787	6.511	0.784
1999	May	NP15	-2.364	2.662	0.375	-4.793	6.405	0.455
1999	June	NP15	-2.706	2.706	0.318	-2.007	6.511	0.758
1999	July	NP15	-11.289	2.662	0.000	-9.278	6.405	0.148
1999	August	NP15	-2.021	2.662	0.448	3.382	6.405	0.598
1999	September	NP15	0.764	2.706	0.778	2.464	6.511	0.705
1999	October	NP15	-0.968	2.662	0.716	7.758	6.405	0.226
1999	November	NP15	6.637	2.706	0.014	11.420	6.511	0.080
1999	December	NP15	1.506	2.662	0.572	3.481	6.405	0.587
2000	January	NP15	1.364	2.662	0.609	1.968	6.405	0.759
2000	February	NP15	1.080	2.753	0.695	-1.203	6.622	0.856
2000	March	NP15	-2.039	2.662	0.444	1.785	6.405	0.781
2000	April	NP15	-1.714	2.706	0.527	3.100	6.511	0.634
2000	May	NP15	14.348	2.662	0.000	6.546	6.405	0.307
2000	June	NP15	11.805	2.706	0.000	3.966	6.511	0.543
2000	July	NP15	22.663	2.662	0.000	36.134	6.405	0.000
1998	April	SP15	-4.249	2.294	0.064	-1.578	6.345	0.804
1998	May	SP15	-1.876	2.257	0.406	-1.767	6.242	0.777
1998	June	SP15	-1.114	2.294	0.627	-4.994	6.345	0.431
1998	July	SP15	-5.794	2.257	0.010	-5.354	6.242	0.391
1998	August	SP15	-3.398	2.257	0.132	6.389	6.242	0.306
1998	September	SP15	-1.475	2.294	0.520	3.310	6.345	0.602
1998	October	SP15	2.406	2.294	0.295	4.381	6.345	0.490
1998	November	SP15	2.489	2.294	0.278	0.815	6.345	0.898
1998	December	SP15	-1.397	2.257	0.536	-0.275	6.242	0.965
1999	January	SP15	-0.300	2.257	0.894	-2.009	6.242	0.748
1999	February	SP15	-1.030	2.374	0.664	0.171	6.567	0.979
1999	March	SP15	-1.110	2.257	0.623	1.274	6.242	0.838
1999	April	SP15	-0.341	2.294	0.882	1.679	6.345	0.791
1999	May	SP15	-2.330	2.257	0.302	-4.793	6.242	0.443
1999	June	SP15	-1.960	2.294	0.393	-1.965	6.345	0.757
1999	July	SP15	-7.089	2.257	0.002	-7.857	6.242	0.208
1999	August	SP15	-3.300	2.257	0.144	3.677	6.242	0.556
1999	September	SP15	-0.136	2.294	0.953	5.340	6.345	0.400
1999	October	SP15	-2.649	2.257	0.241	4.465	6.242	0.475
1999	November	SP15	-2.558	2.294	0.265	4.299	6.345	0.498
1999	December	SP15	5.007	2.257	0.027	3.445	6.242	0.581
2000	January	SP15	1.788	2.257	0.428	0.788	6.242	0.900
2000	February	SP15	1.529	2.333	0.513	-2.035	6.453	0.753
2000	March	SP15	-1.736	2.257	0.442	0.268	6.242	0.966
2000	April	SP15	-1.356	2.294	0.555	9.648	6.345	0.129
2000	May	SP15	10.849	2.257	0.000	16.162	6.242	0.010
2000	June	SP15	16.686	2.294	0.000	0.081	6.345	0.990
2000	July	SP15	3.567	2.257	0.114	7.747	6.242	0.215

TABLE 5

Dependent Variable:		ISO-PX	OLS Coef	Percent of ISO	N-W SE	N-W P-value	OLS Coef	Percent of ISO	N-W SE	N-W P-value
			Early	Early	Early	Early	Late	Late	Late	Late
Year	Month	Zone	Hrs 1-6	Hrs 1-6	Hrs 1-6	Hrs 1-6	Hrs 8-24	Hrs 8-24	Hrs 8-24	Hrs 8-24
1998	April	NP15	-3.571	0.322	1.831	0.051	-1.556	0.065	1.127	0.168
1998	May	NP15	-1.876	0.857	0.821	0.023	-2.860	0.234	1.428	0.045
1998	June	NP15	-1.153	0.766	0.461	0.013	-4.856	0.431	1.905	0.011
1998	July	NP15	-6.133	0.524	1.554	0.000	-4.203	0.122	4.555	0.357
1998	August	NP15	0.280	0.012	1.215	0.818	9.206	0.169	4.519	0.042
1998	September	NP15	3.517	0.128	1.040	0.001	8.255	0.178	4.301	0.055
1998	October	NP15	8.922	0.276	1.208	0.000	6.776	0.187	1.264	0.000
1998	November	NP15	3.717	0.134	1.180	0.002	3.108	0.098	0.833	0.000
1998	December	NP15	-3.681	0.155	2.444	0.132	0.432	0.014	2.266	0.849
1999	January	NP15	-1.321	0.092	1.034	0.202	-2.194	0.101	0.689	0.002
1999	February	NP15	-1.052	0.086	0.568	0.064	0.178	0.008	0.478	0.710
1999	March	NP15	-1.934	0.163	0.931	0.038	1.218	0.053	1.033	0.238
1999	April	NP15	-0.341	0.020	0.860	0.692	1.787	0.063	2.637	0.498
1999	May	NP15	-2.364	0.205	1.190	0.047	-4.793	0.207	1.356	0.000
1999	June	NP15	-2.706	0.364	1.113	0.015	-2.007	0.072	3.607	0.578
1999	July	NP15	-11.289	1.409	4.662	0.016	-9.278	0.329	4.848	0.056
1999	August	NP15	-2.021	0.104	1.454	0.165	3.382	0.078	5.718	0.554
1999	September	NP15	0.764	0.025	1.730	0.659	2.464	0.055	5.123	0.631
1999	October	NP15	-0.968	0.027	3.094	0.754	7.758	0.110	8.045	0.335
1999	November	NP15	6.637	0.195	3.128	0.034	11.420	0.215	4.768	0.017
1999	December	NP15	1.506	0.059	1.678	0.370	3.481	0.099	1.100	0.002
2000	January	NP15	1.364	0.051	1.616	0.399	1.968	0.056	1.411	0.163
2000	February	NP15	1.080	0.040	1.334	0.418	-1.203	0.040	1.437	0.403
2000	March	NP15	-2.039	0.102	1.157	0.079	1.785	0.056	1.372	0.194
2000	April	NP15	-1.714	0.130	2.062	0.406	3.100	0.092	3.769	0.411
2000	May	NP15	14.348	0.365	3.820	0.000	6.546	0.105	11.081	0.555
2000	June	NP15	11.805	0.186	7.157	0.099	3.966	0.025	31.250	0.899
2000	July	NP15	22.663	0.314	9.693	0.020	36.134	0.263	12.073	0.003
1998	April	SP15	-4.249	0.409	1.707	0.013	-1.578	0.066	1.127	0.162
1998	May	SP15	-1.876	0.857	0.821	0.023	-1.767	0.133	2.059	0.391
1998	June	SP15	-1.114	0.741	0.454	0.014	-4.994	0.443	1.799	0.006
1998	July	SP15	-5.794	0.497	1.595	0.000	-5.354	0.156	4.789	0.264
1998	August	SP15	-3.398	0.187	1.580	0.032	6.389	0.119	4.794	0.183
1998	September	SP15	-1.475	0.076	1.741	0.397	3.310	0.080	3.814	0.386
1998	October	SP15	2.406	0.152	1.666	0.149	4.381	0.137	1.389	0.002
1998	November	SP15	2.489	0.183	1.254	0.048	0.815	0.029	0.668	0.223
1998	December	SP15	-1.397	0.087	1.569	0.373	-0.275	0.009	2.150	0.898
1999	January	SP15	-0.300	0.022	1.138	0.792	-2.009	0.093	0.694	0.004
1999	February	SP15	-1.030	0.084	0.566	0.069	0.171	0.008	0.478	0.721
1999	March	SP15	-1.110	0.094	0.946	0.241	1.274	0.055	1.016	0.210
1999	April	SP15	-0.341	0.020	0.860	0.692	1.679	0.059	2.647	0.526
1999	May	SP15	-2.330	0.202	1.181	0.049	-4.793	0.207	1.356	0.000
1999	June	SP15	-1.960	0.264	1.063	0.066	-1.965	0.071	3.637	0.589
1999	July	SP15	-7.089	0.885	4.083	0.083	-7.857	0.279	5.438	0.149
1999	August	SP15	-3.300	0.205	1.735	0.058	3.677	0.088	4.894	0.453
1999	September	SP15	-0.136	0.008	2.162	0.950	5.340	0.136	5.827	0.360
1999	October	SP15	-2.649	0.100	2.082	0.204	4.465	0.092	4.090	0.275
1999	November	SP15	-2.558	0.175	3.686	0.488	4.299	0.111	2.450	0.080
1999	December	SP15	5.007	0.199	2.230	0.025	3.445	0.100	1.097	0.002
2000	January	SP15	1.788	0.072	1.720	0.299	0.788	0.024	1.363	0.563
2000	February	SP15	1.529	0.058	1.629	0.348	-2.035	0.069	1.508	0.177
2000	March	SP15	-1.736	0.090	1.283	0.176	0.268	0.008	1.279	0.834
2000	April	SP15	-1.356	0.103	1.958	0.489	9.648	0.208	7.601	0.205
2000	May	SP15	10.849	0.313	2.969	0.000	16.162	0.198	14.170	0.254
2000	June	SP15	16.686	0.319	5.467	0.002	0.081	0.001	28.929	0.998
2000	July	SP15	3.567	0.076	5.565	0.522	7.747	0.056	12.147	0.524

TABLE 6

Dependent Variable: ISO+R-PX

Year	Month	Zone	OLS Coef Early Hrs 1-6	Percent of ISO Early Hrs 1-6	N-W SE Early Hrs 1-6	N-W P-value Early Hrs 1-6	OLS Coef Late Hrs 8-24	Percent of ISO Late Hrs 1-6	N-W SE Late Hrs 8-24	N-W P-value Late Hrs 8-24
1998	April	NP15	4.293	0.388	1.837	0.020	6.366	0.266	1.110	0.000
1998	May	NP15	6.020	2.749	0.822	0.000	5.065	0.414	1.428	0.000
1998	June	NP15	1.701	1.131	0.681	0.013	-0.645	0.057	1.694	0.703
1998	July	NP15	-2.357	0.202	1.743	0.177	22.111	0.643	8.007	0.006
1998	August	NP15	2.190	0.093	1.026	0.033	57.315	1.054	14.098	0.000
1998	September	NP15	3.937	0.144	1.057	0.000	24.506	0.529	10.355	0.018
1998	October	NP15	8.922	0.276	1.208	0.000	7.315	0.202	1.254	0.000
1998	November	NP15	3.717	0.134	1.180	0.002	3.859	0.122	0.839	0.000
1998	December	NP15	-3.681	0.155	2.444	0.132	3.376	0.106	2.865	0.239
1999	January	NP15	-1.321	0.092	1.034	0.202	-1.232	0.057	0.690	0.075
1999	February	NP15	-1.052	0.086	0.568	0.064	1.242	0.058	0.474	0.009
1999	March	NP15	-1.934	0.163	0.931	0.038	2.030	0.088	1.037	0.051
1999	April	NP15	-0.341	0.020	0.860	0.692	3.915	0.137	2.469	0.113
1999	May	NP15	-2.364	0.205	1.190	0.047	-2.545	0.110	1.394	0.068
1999	June	NP15	-2.706	0.364	1.113	0.015	-0.102	0.004	4.114	0.980
1999	July	NP15	-11.289	1.409	4.662	0.016	2.372	0.084	7.246	0.743
1999	August	NP15	-2.017	0.104	1.454	0.166	9.485	0.219	8.115	0.243
1999	September	NP15	0.777	0.025	1.729	0.653	7.484	0.166	6.459	0.247
1999	October	NP15	-0.958	0.027	3.094	0.757	20.341	0.288	5.557	0.000
1999	November	NP15	6.647	0.195	3.128	0.034	13.101	0.247	4.879	0.007
1999	December	NP15	1.515	0.059	1.678	0.367	3.990	0.113	1.115	0.000
2000	January	NP15	1.374	0.051	1.616	0.396	2.283	0.064	1.399	0.103
2000	February	NP15	1.089	0.041	1.334	0.414	-0.542	0.018	1.333	0.684
2000	March	NP15	-2.029	0.101	1.157	0.080	2.698	0.084	1.484	0.069
2000	April	NP15	-1.709	0.130	2.062	0.407	3.780	0.112	3.870	0.329
2000	May	NP15	14.355	0.365	3.819	0.000	26.330	0.421	13.535	0.052
2000	June	NP15	11.933	0.188	7.155	0.096	102.675	0.643	54.967	0.062
2000	July	NP15	22.702	0.315	9.700	0.019	47.878	0.348	14.756	0.001
1998	April	SP15	3.615	0.348	1.717	0.036	6.344	0.265	1.109	0.000
1998	May	SP15	6.020	2.749	0.822	0.000	6.158	0.462	2.064	0.003
1998	June	SP15	1.738	1.155	0.671	0.010	-1.018	0.090	1.674	0.543
1998	July	SP15	-5.037	0.432	1.620	0.002	157.396	4.597	94.489	0.096
1998	August	SP15	-1.789	0.098	1.548	0.248	55.125	1.024	14.392	0.000
1998	September	SP15	-1.069	0.055	1.763	0.545	19.887	0.480	9.976	0.047
1998	October	SP15	2.406	0.152	1.666	0.149	4.920	0.153	1.387	0.000
1998	November	SP15	2.489	0.183	1.254	0.048	1.566	0.056	0.684	0.022
1998	December	SP15	-1.397	0.087	1.569	0.373	2.672	0.090	2.778	0.336
1999	January	SP15	-0.300	0.022	1.138	0.792	-1.048	0.048	0.692	0.131
1999	February	SP15	-1.030	0.084	0.566	0.069	1.235	0.058	0.474	0.009
1999	March	SP15	-1.110	0.094	0.946	0.241	2.085	0.090	1.020	0.041
1999	April	SP15	-0.341	0.020	0.860	0.692	3.807	0.134	2.472	0.124
1999	May	SP15	-2.330	0.202	1.181	0.049	-2.576	0.111	1.391	0.064
1999	June	SP15	-1.960	0.264	1.063	0.066	-0.103	0.004	4.143	0.980
1999	July	SP15	-7.089	0.885	4.083	0.083	3.794	0.135	7.812	0.627
1999	August	SP15	-3.296	0.205	1.735	0.058	9.781	0.233	7.408	0.187
1999	September	SP15	-0.126	0.007	2.162	0.954	8.260	0.211	6.433	0.200
1999	October	SP15	-2.639	0.100	2.082	0.205	10.092	0.208	4.568	0.027
1999	November	SP15	-2.548	0.174	3.686	0.490	5.679	0.147	2.594	0.029
1999	December	SP15	5.016	0.199	2.230	0.025	3.977	0.115	1.104	0.000
2000	January	SP15	1.798	0.072	1.720	0.296	1.128	0.034	1.357	0.406
2000	February	SP15	1.538	0.059	1.629	0.345	-1.374	0.046	1.416	0.332
2000	March	SP15	-1.726	0.090	1.283	0.179	1.305	0.041	1.360	0.338
2000	April	SP15	-1.351	0.103	1.958	0.490	10.409	0.225	7.562	0.169
2000	May	SP15	10.856	0.313	2.968	0.000	35.951	0.440	15.783	0.023
2000	June	SP15	16.814	0.322	5.477	0.002	98.790	0.656	53.017	0.063
2000	July	SP15	3.607	0.076	5.568	0.517	21.348	0.154	13.093	0.103

TABLE 7

Dependent Variable:	ISO-PX	Average |Percent of ISO| Early	Average P-value Early	Average |Percent of ISO| Late	Average P-Value Late
Period	Zone				
April-December 1998	NP15	0.353	0.116	0.166	0.170
January-August 1999	NP15	0.336	0.148	0.116	0.376
Sept 1999-July 2000	NP15	0.327	0.233	0.122	0.365
April-December 1998	SP15	0.354	0.117	0.130	0.279
January-August 1999	SP15	0.222	0.256	0.107	0.332
Sept 1999-July 2000	SP15	0.137	0.318	0.091	0.388

Dependent Variable:	ISO+R-PX	Average |Percent of ISO| Early	Average P-value Early	Average |Percent of ISO| Early	Average P-Value Late
Period	Zone				
April-December 1998	NP15	0.586	0.042	0.377	0.107
January-August 1999	NP15	0.305	0.155	0.095	0.285
Sept 1999-July 2000	NP15	0.136	0.293	0.228	0.141
April-December 1998	SP15	0.584	0.157	0.802	0.116
January-August 1999	SP15	0.222	0.256	0.102	0.270
Sept 1999-July 2000	SP15	0.138	0.318	0.207	0.154

TABLE 8

PROFITABILITY OF TRADING RULES

BY MONTH	All Months Avg Profit per MWh	All Months t-stat	April-Dec 1998 Avg Profit per MWh	April-Dec 1998 t-stat	Jan-Aug 1999 Avg Profit per MWh	Jan-Aug 1999 t-stat	Sept 1999-July 2000 Avg Profit per MWh	Sept 1999-July 2000 t-stat
Speculative Trader								
North Early	2.50	3.73	2.64	4.58	2.91	3.78	2.13	1.48
North Late	4.56	2.93	2.61	2.27	0.74	0.73	8.41	2.41
South Early	1.77	3.35	1.56	2.91	1.98	2.95	1.79	1.62
South Late	2.75	1.96	1.72	1.53	1.13	1.06	4.52	1.45
Representative Trader								
North Early	2.06	3.56	1.52	3.02	NA	NA	3.78	2.91
North Late	4.24	2.83	2.28	3.34	-0.54	-1.18	8.73	2.53
South Early	1.05	2.65	0.13	0.49	NA	NA	2.39	2.63
South Late	2.25	1.68	1.01	1.72	-0.25	-0.57	4.75	1.53

NA = Not Applicable. The average ISO-PX price difference was negative during all months, so the trading rule does not generate any profi

BY WEEK	All Weeks Avg Profit per MWh	All Weeks t-stat	April-Dec 1998 Avg Profit per MWh	April-Dec 1998 t-stat	Jan-Aug 1999 Avg Profit per MWh	Jan-Aug 1999 t-stat	Sept 1999-July 2000 Avg Profit per MWh	Sept 1999-July 2000 t-stat
Speculative Trader								
North Early	3.15	4.83	3.09	5.73	1.88	2.37	4.04	2.81
North Late	3.87	2.56	3.21	3.19	0.44	0.44	6.62	1.89
South Early	1.74	3.48	1.37	2.60	0.44	0.63	2.87	2.75
South Late	2.40	1.77	1.90	1.91	0.45	0.43	4.07	1.30
Representative Trader								
North Early	2.31	4.11	1.57	3.46	-0.51	-3.35	4.74	3.67
North Late	3.82	3.06	2.43	3.25	-0.69	-1.93	7.84	2.67
South Early	0.97	2.70	-0.09	-0.42	-0.77	-3.44	2.93	3.56
South Late	2.03	1.84	1.03	1.51	-0.59	-1.63	4.53	1.73

TABLE 9

EFFECT OF LAGGED PRICE DIFFERENCES

	Full Sample	Full Sample	1998	1998	Jan-Aug 1999	Jan-Aug 1999	Sep 1999- July 2000	Sep 1999- July 2000
ONE-DAY LAG	Coefficient	SE	Coefficient	SE	Coefficient	SE	Coefficient	SE
North Early	0.47	0.09	0.44	0.06	0.23	0.10	0.49	0.11
North Late	0.32	0.08	0.40	0.08	0.44	0.08	0.30	0.09
South Early	0.18	0.09	0.34	0.07	0.18	0.07	0.13	0.11
South Late	0.31	0.08	0.31	0.07	0.47	0.10	0.30	0.09
TWO-DAY LAG								
North Early	0.29	0.17	0.26	0.07	-0.10	0.07	0.34	0.23
North Late	-0.14	0.11	0.04	0.07	-0.09	0.11	-0.17	0.12
South Early	0.05	0.08	0.16	0.07	-0.16	0.08	0.04	0.10
South Late	-0.20	0.10	-0.07	0.08	-0.08	0.13	-0.22	0.11

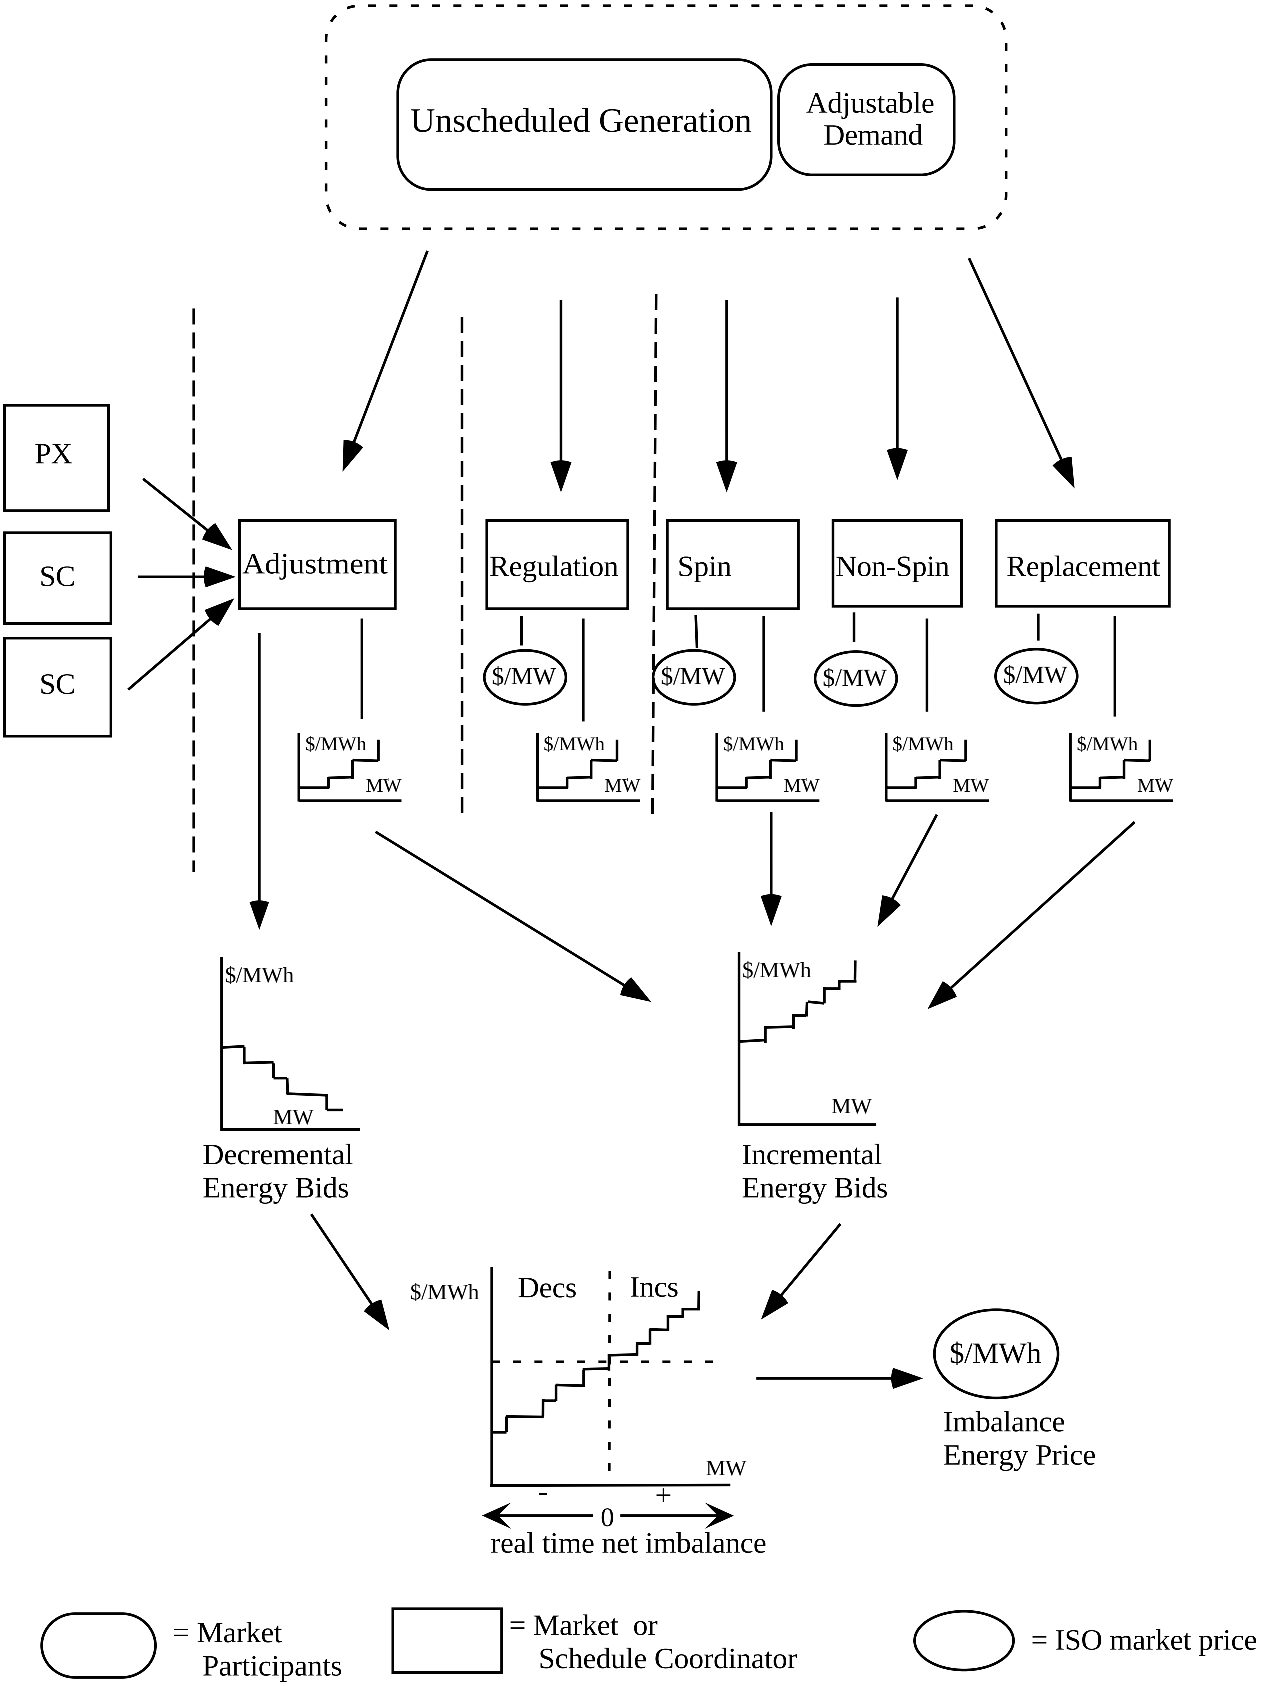

Figure 1: California ISO Product Markets