JOURNAL OF APPLIED ECONOMETRICS
J. Appl. Econ. **22**: 1127–1144 (2007)
Published online in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/jae.982

EFFICIENCY OF THE CALIFORNIA ELECTRICITY RESERVES MARKET

KONSTANTINOS METAXOGLOU[a] AND AARON SMITH[b]*

[a] *Department of Economics, University of California Davis, California, USA*
[b] *Department of Agricultural and Resource Economics, University of California, Davis, California, USA*

SUMMARY

We test the efficiency of the California electricity reserves market by examining systematic differences between its day- and hour-ahead prices. We uncover significant day-ahead premia, which we attribute to market design characteristics. On the demand side, the market design established a principal–agent relationship between the markets' buyers (principal) and their supervisory authority (agent). The agent had very limited incentives to shift reserve purchases to the lower priced hour-ahead markets. On the supply side, the market design raised substantial entry barriers by precluding purely speculative trading and by introducing a complicated code of conduct that induced uncertainty about which actions were subject to regulatory scrutiny. We use a high-dimensional vector moving average model to estimate the premia and conduct correct inferences. To obtain exact maximum likelihood estimates of the model, we develop a new EM algorithm that seamlessly incorporates missing data and applies directly to general moving average time series models. Our algorithm uses only analytical expressions: the Kalman filter and a fixed interval smoother in the E step and least squares-type regressions in the M step. Copyright © 2007 John Wiley & Sons, Ltd.

1. INTRODUCTION

During the last 15 years, a massive wave of restructuring has transformed the wholesale electricity sector in the United States. This reform created a myriad of interrelated markets that determine prices and ensure a continuous balance between supply and demand as the physics of electricity requires. In California, the new setup appeared to work well from its implementation in 1998 until the summer of 2000. However, on June 14, 2000, the state faced its first rolling blackout since World War II. More supply disruptions were to follow, with 38 blackouts and service interruptions occurring between November 2000 and May 2001 (California Department of Justice Office of the Attorney General, 2004).

During this period, the state's energy markets operated inefficiently by exhibiting spot prices that substantially exceeded forward prices on average (Borenstein *et al.*, 2004). In this article, we study California's electricity reserves markets. These markets were designed to ensure the availability of sufficient standby supply. Therefore, they played a critical role in the reliability of the electricity grid. Unlike their forward energy counterparts, the reserves markets have continued to operate since the crisis, so their efficiency is of ongoing policy relevance. We show that the reserves markets performed inefficiently not only before and during the crisis, but also after the crisis. In our 3-year sample, the total cost of this inefficiency is of the order of $1 billion.

* Correspondence to: Aaron Smith, Department of Agricultural and Resource Economics, University of California Davis, One Shields Avenue, Davis, CA 95616, USA. E-mail: adsmith@ucdavis.edu

For the period that we analyze, the reserves markets enabled trade in five types of products intended to respond to contingencies, such as demand surges and supply deficiencies. They operated on a daily basis, along with the state's forward and spot energy markets, with transactions over two different time horizons: a day-ahead and an hour-ahead. Purchasing day- and hour-ahead reserves was equivalent to obtaining the right, but not the obligation, to call energy up to the amount of reserve bought. Therefore, a transaction over day- and hour-ahead reserves was equivalent to signing a European-style call option, where the underlier was energy (see Chao and Wilson, 2002). The California Independent System Operator (CAISO), a non-profit corporation that had been already assigned the role of running the state's wholesale power grid, undertook the task to manage the day- and hour-ahead reserves sales and purchases.

We define a reserves market as efficient if day-ahead prices are unbiased predictors of the corresponding hour-ahead prices for the same type of reserve during the same hour. We test this hypothesis and find that day-ahead prices were considerably higher than hour-ahead prices for all the reserve types we analyze. This day-ahead premium existed before, during, and after the crisis. If the markets were efficient, then risk-neutral buyers would have transferred their demand to the lower-priced hour-ahead markets, or risk-neutral sellers would have shifted their offers to the higher-priced day-ahead markets, so that the equilibrium relation was restored. The demand side had weak incentives to shift its purchases from the day- to the hour-ahead markets because of the principal–agent relationship that existed between the CAISO (agent) and the buyer (principal) of the reserves. On the supply side, the market structure raised entry barriers by precluding purely speculative trading and creating a complicated code of conduct. The vague descriptions of practices subject to regulatory scrutiny and sanctions deterred sellers from bidding down the premium.

Several alternative supply-side explanations exist for the persistence of the day-ahead premium. First, if the day-ahead market were more costly for sellers to transact in, then we would expect a day-ahead premium to persist. However, given that the day- and hour-ahead markets were operated by the same entity and cleared in the same way, there is little evidence of such differential transaction costs. Second, the premium would persist if all sellers were risk averse and were receiving a risk premium to sell in the day- rather than the hour-ahead market. For example, Bessembinder and Lemon (2002) explain forward premia in energy (rather than reserves) markets as compensation to risk-averse power companies for bearing risk (see also Longstaff and Wang, 2004). For reserves markets, we find day-ahead premia between 25% and 50% in most cases. These premia are much larger than those in energy markets, both in absolute terms and when measured relative to the standard deviation. It is difficult to imagine why a rational risk-averse seller would require such a large premium as daily compensation for committing generating capacity one day in advance.

A limited number of incumbents did profit from the day-ahead premium. Among those who took advantage of systematically higher day-ahead prices were Enron traders, by means of their so-called 'Get Shorty' strategy. Violating market rules and acting as pure speculators, they sold day-ahead at a high price expecting to buy back hour-ahead for a low price, thus gaining the difference. According to taped discussions available from the web site of the Snohomish Public Utility District, Enron traders were so allured by the margin in their buybacks that they did not hesitate to engage in illegal business practices, having learned well how to game the market rules. Enron traders repeatedly submitted false information regarding the physical availability of reserve resources that they did not have (see Section 4.3). The extent of any illegal behavior by other companies is unknown.

DOI: 10.1002/jae

In Section 2 of the paper, we discuss the mechanics of the California electricity reserves markets. Accounting for the market structure, we build a reduced form model of the difference (spread) between day- and hour-ahead prices as a high-dimensional vector moving-average process. In Section 3, we develop a new state-space representation of this vector moving average that enables exact maximum likelihood estimation by the EM algorithm using analytical expressions. In contrast, existing methods require numerical differentiation, which is computationally infeasible in our high-dimensional framework. In Section 4, we elaborate further our arguments regarding the effects of the market design on the size and persistence of the day-ahead premia, and we conclude in Section 5.

2. THE CAISO RESERVES MARKET

The CAISO is responsible for the reliability of the high-voltage transmission grid within its control area, which covers almost the entire state of California, by maintaining a continuous balance between electricity demand and supply. Among other things, the CAISO ensures that sufficient reserves of electricity are maintained on an hourly basis every day, as required by criteria and standards with which its operations conform (CAISO Tariff, Section 2.5.2.1). These reserves include the following five products procured in day- and hour-ahead markets: regulation up, regulation down, spinning, non-spinning and replacement.

Flexible suppliers that can increase and decrease their output instantly under automatic control provide regulation up and regulation down, respectively. A supplier's unit, already operating at its minimum level, can provide spinning reserve if it can convert reserves to energy within 10 minutes and maintain that output for at least 2 hours. The same response time and availability requirements hold for curtailable demand and offline suppliers that offer non-spinning reserve. Units that are capable of starting up, if not already operating, and increasing their output to a specified level within 60 minutes may provide replacement reserves. Curtailable demand with a 60-minute response time is also a potential source of replacement reserves (CAISO Tariff, sections 2.5.14–2.5.17).

Between 1999 and 2002, participants in the California wholesale electricity market were required to submit daily schedules of their forward energy commitments to the CAISO via entities known as scheduling coordinators. Each scheduling coordinator aggregated a set of forward commitments to form a balanced schedule, i.e., a schedule such that supply was equal to demand. The list of scheduling coordinators included municipal utilities (e.g., Los Angeles Department of Water and Power), in-state public utilities (e.g., Pacific Gas and Electric), private entities (e.g., Duke, Reliant, and Enron), public utilities from neighboring states (e.g., Puget Sound Energy), and the Bonneville power administration of the US Department of Energy.

The CAISO required scheduling coordinators to maintain sufficient reserves to absorb any real-time deviations from their balanced forward energy schedules. The scheduling coordinators' share of total scheduled energy determined their share of the CAISO total reserve requirements. For example, a scheduling coordinator with 200 megawatt-hours of scheduled energy was required to maintain twice the reserves of a scheduling coordinator with 100 megawatt-hours scheduled energy. In many cases, the scheduling coordinators met some of their reserve obligations by self-providing with a mix of the generation and curtailable demand they represented.

The scheduling coordinators purchased in the reserves markets the difference between their total reserve obligations and their self-provisions. In these market transactions, the CAISO acted as an

DOI: 10.1002/jae

agent between a scheduling coordinator that was required to buy and a scheduling coordinator that was willing to sell reserves. The former received a bill and the latter received a payment. The CAISO had full discretion over the proportion of reserves transacted in the day-ahead market and the proportion transacted hour-ahead. Historically, around 80% of these transactions took place day-ahead with the rest deferred hour-ahead (CAISO, various monthly reports). This principal–agent relationship between a purchasing scheduling coordinator (the principal) and the CAISO (the agent) partially explains the day-ahead premia we find in Section 4.

For the day-ahead market, sellers had to bid their reserves by 12:00 noon on the day prior to the relevant trading day for all of its 24 hours. The CAISO cleared each hour-ahead reserve market one hour prior to the beginning of each trading interval (e.g., by 11:00 for the interval 12:00–13:00). In both the day- and hour-ahead markets, bids had to identify fully the name, location and technical characteristics of their reserve resources, if the resources were within the CAISO control area. In the case of reserves coming from outside the CAISO control area (imports), the bids had only to identify their point of entry to the CAISO control area (CAISO Tariff, sections 2.5.14–2.5.17). Therefore, purely speculative trading was precluded by the rules in the day- and hour-ahead reserves markets.

The day- and hour-ahead reserve bids had capacity and energy components. The capacity component was a single quantity and price pair (megawatt, $/megawatt) that indicated willingness to supply reserves. It was used in the reserves markets to produce the market clearing prices ignoring the energy component. The energy component was a step supply curve with up to 11 quantity and price pairs (megawatt-hour, $/megawatt-hour), and it determined the position of a reserve supplier in the bid stack of the spot energy market should it be called to supply energy. The quantity of the last pair in the energy component was the same as the quantity of the single pair in the capacity component. Ultimately, a reserve supplier could receive two payments. The first payment was made for having amounts of electricity production capacity on standby. The second payment was for the fraction of the available electricity used in the spot energy market. Hence, in a call-option context, the price of the capacity component was the option's price, calling energy from reserves was equivalent to exercising the option and the spot energy price was the option's strike price.

For the day- and hour-ahead reserve markets, the CAISO software created a step supply curve by stacking the capacity components of reserve offers in ascending order of their prices. To produce the market clearing price, this market supply curve was crossed with a vertical demand representing the purchase decision of the CAISO. However, in about 30% of the hours in our sample, market transactions did not occur for three main reasons. The first reason was sufficient self-provision outside the market, such that residual demand for reserves was zero. Second, reserve products that could respond faster to a dispatch instruction could be procured in place of a slower responding reserve. For example, procurement of regulation in place of spinning could cause zero demand in the spinning market. Third, missing hour-ahead prices arose due to the fact that CAISO sometimes procured all of its reserve requirements in the day-ahead markets only. Hours with zero procurement gave rise to missing prices.

3. VECTOR MOVING-AVERAGE MODEL AND ESTIMATION METHOD

In this section, we develop the moving-average model implied by the reserves market structure, and we introduce a new EM algorithm for exact maximum likelihood estimation of the model's

DOI: 10.1002/jae

parameters. The CAISO cleared the day-ahead markets by 12:00 the day before the trading day and it cleared the hour-ahead markets one hour prior to the beginning of the relevant trading intervals (e.g., 12:00 for the interval 13:00–14:00). For the first interval of the trading day (24:00–1:00) there was a window of 11 hours from the time that the sellers submitted their day-ahead bids (12:00) to the time that the CAISO cleared the corresponding hour-ahead market (23:00). Similarly, for the last interval of the trading day (23:00–24:00), there was a 34-hour time span. It follows that the difference between day- and hour-ahead prices for each hour depends on accumulated information in the intervening 11–34 hours, even in an efficient market.

We define an observation in our sample as $y_{hd} = \text{PHA}_{hd} - \text{PDA}_{hd}$, where PHA_{hd} and PDA_{hd} are the hour- and day-ahead prices for hour ending $h = 1, \ldots, 24$ and day d. If we denote the information arriving in hour h and day d by w_{hd}, then we can express the difference between day- and hour-ahead prices as a moving average process of 11 to 34 w_{hd} terms (Borenstein *et al.* 2004). Because the hour-ahead market clears one hour before the beginning of the relevant trading interval (and therefore two hours before the end of the interval), we write:

$$
\begin{aligned}
y_{1d} &= \beta_1 + w_{23,d-1} + \theta_{1,1} w_{22,d-1} + \cdots + \theta_{1,11} w_{12,d-1} \\
y_{2d} &= \beta_2 + w_{24,d-1} + \theta_{2,1} w_{23,d-1} + \theta_{2,2} w_{22,d-1} + \cdots + \theta_{24,34} w_{12,d-1} \\
&\vdots \\
y_{24d} &= \beta_{24} + w_{22d} + \cdots + \theta_{24,22} w_{24,d-1} + \theta_{24,23} w_{23,d-1} + \theta_{24,24} w_{22,d-1} + \cdots + \theta_{24,34} w_{12,d-1}
\end{aligned}
$$

with an equivalent 24-dimensional vector moving average representation of order 1:

$$
\begin{aligned}
Y_d &= B + A_0 W_d + A_1 W_{d-1} \\
&= B + U_d + \Theta U_{d-1} \qquad (1)
\end{aligned}
$$

where $Y_d = \text{PHA}_d - \text{PDA}_d$, $B = (\beta_1, \ldots, \beta_{24})^{\mathsf{T}}$, $U_d = A_0 W_d$, and $\Theta_1 = A_1 A_0^{-1}$. The serially uncorrelated zero mean Gaussian error vector U_d has covariance matrix Σ_u. Because our hypothesis is that the day-ahead prices are unbiased estimates of the hour-ahead prices, we test whether B is statistically significantly different from zero. Taking into account the moving average structure is important for correct inference about B. Rather than reporting the sample mean price spreads together with robust standard errors, we estimate (1) by maximum likelihood for two reasons. First, maximum likelihood improves efficiency by enabling joint estimation of (B, Θ, Σ_u). Second, estimation by maximum likelihood allows us to explicitly account for the missing data that we have in our sample.

The available methods for exact maximum likelihood estimation of Gaussian vector moving-average (VMA) models employ numerical techniques, which are numerically unstable and computationally infeasible for high-dimensional vector processes. For the 24-dimensional VMA(1) model in (1), we performed one iteration of the numerical gradient procedure of GAUSS on a PC with 1 GB of RAM and an Intel Pentium 4 processor (2.8 GHz). This operation took almost 37 minutes to complete, and to maximize the likelihood we would require many more such iterations. We make exact maximum likelihood estimation feasible by proposing a new estimation method that completes one iteration every 9 seconds, which is almost 200 times faster than the numerical gradient procedure. Specifically, we propose a state-space representation that allows the

DOI: 10.1002/jae

EM algorithm to produce exact maximum-likelihood estimates using analytical expressions. Using the fact that a VMA process has the same Wold representation as a VMA plus white noise, we add noise to the observation equation and employ the EM algorithm as in Shumway and Stoffer (1982) and Engle and Watson (1983). The E step requires a pass of the Kalman filter and a fixed interval smoother that also handles the abundant missing observations in our sample; the M step collapses to least-squares type regressions.

We write (1) as

$$Y - B = \overline{\mathbf{\Theta}}\mathbf{u} \tag{2}$$

$$\mathbf{Y} = \begin{bmatrix} Y_1 \\ \vdots \\ Y_n \end{bmatrix}, \mathbf{u} = \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_n \end{bmatrix}, \overline{\mathbf{\Theta}} = \begin{bmatrix} \Theta & I_d & 0 & \cdots & 0 & 0 \\ 0 & \Theta & I_d & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \Theta & I_d \end{bmatrix}$$

with I_{n+1} and I_d being $(n+1) \times (n+1)$ and $d \times d$ identity matrices. The size of matrix $\overline{\mathbf{\Theta}}$ is $(dn \times d(n+1))$. Ignoring the constant term, we write the exact Gaussian log-likelihood for (1) as

$$l(\theta|\mathbf{Y}) = -\frac{1}{2}\ln|\overline{\mathbf{\Theta}}(I_{n+1} \otimes \Sigma_u)\overline{\mathbf{\Theta}}^{\mathsf{T}}| - \frac{1}{2}(\mathbf{Y} - B)^{\mathsf{T}}(\overline{\mathbf{\Theta}}(I_{n+1} \otimes \Sigma_u)\overline{\mathbf{\Theta}}^{\mathsf{T}})^{-1}(\mathbf{Y} - B) \tag{3}$$

where $\theta = \text{vec}(B, \Theta, \Sigma_u)$, and $^{\mathsf{T}}$, $|\cdot|$ and $\otimes$ denote the transpose, determinant and Kronecker product operators, respectively. The task of evaluating the expression in (3) is complicated because it is difficult to compute the inverse and the determinant of the $dn \times dn$ covariance matrix (Lütkepohl, 1993, section 7.2.1). The inverse and the determinant of the covariance matrix are also highly nonlinear in θ and, therefore, to obtain the maximum likelihood estimate via numerical maximization is extremely demanding computationally, especially if d is large.

Using the result that a VMA of some order q plus white noise remains a VMA of the same order q (Pieris, 1988, theorem 2), we write $u_t + \Theta u_{t-1} \equiv \upsilon_t + \Gamma\upsilon_{t-1} + \varepsilon_t$, where υ_t and ε_t denote vector white noise processes. This setup allows us to treat the lag of υ_t as observable in the complete-data log-likelihood that underlies the EM algorithm. Therefore, we write (1) in state-space form as

$$\begin{aligned} Y_t &= B + Za_t + \varepsilon_t \;\; \varepsilon_t \sim N(0, \Sigma_\varepsilon) \\ a_t &= Ta_{t-1} + \eta_t \;\; \eta_t \sim N(0, \Sigma_\eta) \\ Z &= [I_d \quad \mathbf{\Gamma}], \; \begin{matrix} a_t^{\mathsf{T}} = [\upsilon_t^T, \upsilon_{t-1}^T]^{\mathsf{T}} \\ \eta_t^T = [\upsilon_t^T, 0]^T \end{matrix}, \; T = \begin{bmatrix} 0 & 0 \\ I_d & 0 \end{bmatrix}, \Sigma_\eta = \begin{bmatrix} \Sigma_\upsilon & 0 \\ 0 & 0 \end{bmatrix} \end{aligned} \tag{4}$$

We label the equation for Y_t the observation equation and we refer to the equation for a_t as the state equation. The Gaussian independently distributed disturbance vectors ε_t and η_t are mutually uncorrelated and independent of a_0. We maintain diagonality for Σ_ε for the remainder of this section.

If $\Sigma_\varepsilon = 0$, then (4) is identical to (1) and $\Gamma \equiv \Theta$. However, if $\Sigma_\varepsilon \neq 0$, then there exists a unique mapping from $(Z, \Sigma_\varepsilon, \Sigma_\upsilon)$ to (Θ, Σ_u) by matching the moving average parameters Θ to the coefficients in the Wold representation of (4). This matching yields $\Theta = ZK$ (Hamilton, 1994,

DOI: 10.1002/jae

section 13.5), where K denotes the steady-state Kalman gain matrix. Moreover, Σ_u equals the steady-state value of the covariance of the one-step prediction error, which we also obtain from the Kalman filter.

The Wold representation provides the unique mapping from $(Z, \Sigma_\varepsilon, \Sigma_\upsilon)$ to (Θ, Σ_u). However, the reverse mapping is not unique, which implies that the parameters $(Z, \Sigma_\varepsilon, \Sigma_\upsilon)$ are not separately identified. Therefore, to identify the parameters in (4), we set Σ_ε to a constant, which allows us to obtain maximum likelihood estimates of (Z, Σ_υ). We then calculate the maximum likelihood estimates of the parameters of interest (Θ, Σ_u). The maximum likelihood estimates of (Θ, Σ_u) are invariant to the value we choose for Σ_ε. This identification strategy becomes clear using a simple univariate MA(1) example: $y_t = \eta_t + \theta\eta_{t-1} + \varepsilon_t$, where η_t and ε_t are white noise processes with variances σ^2 and ω^2. The autocovariances of y_t are $\gamma_0 = E[y_t^2] = (1+\theta^2)\sigma^2 + \omega^2$, $\gamma_1 = E[y_t y_{t-1}] = \theta\sigma^2$, and $\gamma_j = E[y_t y_{t-j}] = 0$ for all $j > 1$. Hence, we have only two moments from which we need to retrieve θ, σ^2 and ω^2. This task is impossible unless a fixed value is assumed for ω^2.

Omitting constants, the exact *complete-data log-likelihood* for (4) is

$$l(\theta|\mathbf{Y}, a) = \frac{n}{2}\ln|\Sigma_\varepsilon^{-1}| + \left(\frac{n+1}{2}\right)\ln|\Sigma_\upsilon^{-1}| - \frac{1}{2}\text{trace}\left(\sum_{t=1}^{n}\Sigma_\varepsilon^{-1}\varepsilon_t\varepsilon_t^\mathsf{T} + \sum_{t=0}^{n}\Sigma_\upsilon^{-1}\upsilon_t\upsilon_t^\mathsf{T}\right) \tag{5}$$

where $\varepsilon_t = Y_t - B - Za_t$, $\upsilon_t = a_t - Ta_{t-1}$. To maximize the *incomplete (observed)-data* log-likelihood $l(\theta|\mathbf{Y})$, we apply the EM algorithm of Dempster *et al.* (1977). In the E step of the ith iteration of the algorithm we calculate the expectation of the *complete-data* log-likelihood $l(\theta|\mathbf{Y}, a)$ conditional on $\mathbf{Y}$ and the previous iterate $\theta^{(i)}$, which enables us to produce the new iterate $\theta^{(i+1)}$ in the M step. Alternating repeatedly between the E and the M steps of the algorithm until convergence, we obtain the maximum likelihood estimates of Z, B, and Σ_υ, which in turn allows us to retrieve the maximum-likelihood estimates of Θ and Σ_u.

The E step of the algorithm, which produces the expected complete-data log-likelihood, requires passes of the Kalman filter and a fixed-interval smoother to obtain the conditional mean and variance of the state and disturbance vectors. With Σ_ε being diagonal, we apply the univariate filtering and smoothing algorithm of Koopman and Durbin (2000) directly.

In the smoother pass, we use the iterations summarized in Durbin and Koopman (2001, section 4.3.1), which avoid the computationally burdensome inversion of the contemporaneous covariance matrix $P_{t|t} = \text{var}(a_t|Y_t)$. This univariate algorithm saves computation time and simplifies the calculation of the exact log-likelihood because it does not require matrix inversion in either the filter or the smoother pass. The same algorithm also easily handles missing observations in Y with the appropriate modifications of the Kalman gain in the filter pass and the implied adjustments in the smoother pass (Durbin and Koopman, 2001, section 4.8).

Omitting constants, the *expected complete-data log-likelihood* is

$$\begin{aligned} Q(\theta|\theta^{(i)}) &= \int l(\theta|\mathbf{Y}, a) f(a|\mathbf{Y}, \theta^{(i)}) da \\ &= \frac{n}{2}\ln|\Sigma_\varepsilon^{-1}| - \frac{1}{2}\text{trace}\left(\sum_{t=1}^{n}\Sigma_\varepsilon^{-1}(\varepsilon_{t|n}\varepsilon_{t/n}^\mathsf{T} + \text{var}(\varepsilon_t|\mathbf{Y}_n))\right) \end{aligned}$$

DOI: 10.1002/jae

$$+\left(\frac{n+1}{2}\right)\ln|\Sigma_{\upsilon}^{-1}| - \frac{1}{2}\text{trace}\left(\sum_{t=0}^{n}\Sigma_{\upsilon}^{-1}(\upsilon_{t|n}\upsilon_{t|n}^{\mathsf{T}} + \text{var}(\upsilon_t|\mathbf{Y}_n))\right) \tag{6}$$

where $\varepsilon_{t|n} = E[\varepsilon_t|\mathbf{Y}_n]$ and $\upsilon_{t|n} = E[\upsilon_t|\mathbf{Y}_n]$. Using standard matrix calculus results, the M step of the algorithm, which solves for $\theta^{(i+1)}$ by maximizing $Q(\theta|\theta^{(i)})$ with respect to θ, implies the following analytical expressions:

$$\begin{bmatrix} Z^{(i+1)} \\ B^{(i+1)} \end{bmatrix} = \begin{bmatrix} \sum_{t=1}^{n}(a_{t|n}a_{t|n}^{\mathsf{T}} + P_{t|n}) & \sum_{t=1}^{n} a_{t|n} \\ \sum_{t=1}^{n} a_{t|n}^{\mathsf{T}} & n \end{bmatrix}^{-1} \begin{bmatrix} \sum_{t=1}^{n}(a_{t|n}(Y_t - \upsilon_{t|n})^{\mathsf{T}} - P_{t|n}^{a\upsilon}) \\ \sum_{t=1}^{n}(Y_t - \upsilon_{t|n})^{\mathsf{T}} \end{bmatrix} \tag{7}$$

$$\Sigma_{\upsilon}^{(i+1)} = \left(\frac{1}{n+1}\right)\sum_{t=0}^{n}((\upsilon_{t|n})(\upsilon_{t|n})^{\mathsf{T}} + \text{var}(\upsilon_t|\mathbf{Y}_n)) \tag{8}$$

where $a_{t|n} = E[a_t|\mathbf{Y}_n]$, $P_{t|n} = \text{var}(a_t|\mathbf{Y}_n)$, and $P_{t|n}^{a\upsilon}$ is the conditional covariance between $a_{t|n}$ and $\upsilon_{t|n}$. Additionally, because Y_t is of dimension $d \times 1$ and Σ_ε is diagonal, (7) is equivalent to d univariate regressions.

The chosen value of Σ_ε does not affect the moving average parameter estimates (θ, Σ_u), but it does affect the algorithm's speed of convergence. We performed numerous simulations for a wide range of values of Σ_ε to assess the impact of noise in the observation equation on the global convergence rate of the EM algorithm, which we denote by $r = \lim_{i\to\infty}(\|\theta^{(i+1)} - \theta^*\|/\|\theta^{(i)} - \theta^*\|)$, where $\theta^{(i)} \to \theta^*$. We used $r = \max r_j$, $r_j = \lim_{i\to\infty}(|\theta_j^{(i+1)} - \theta_j^{(i)}|/|\theta_j^{(i)} - \theta_j^{(i-1)}|)$, where $\theta_j^{(i)}$ is the jth element of θ in the ith iteration (see Meng and Rubin, 1994). Our simulations showed that r is decreasing in the diagonal elements of Σ_ε, which implies a faster rate of convergence for larger Σ_ε. Moreover, because a larger value of Σ_ε implies a smaller value of Σ_υ, these simulation findings are consistent with the theoretical result that the convergence rate is decreasing in the fraction of information contained in the unobserved state variable (Krishnan and McLachlan, 1997, section 3.9.3).

Finally, several methods exist for estimating the covariance of maximum likelihood estimates obtained from the EM algorithm. A closed-form expression for the asymptotic covariance of Θ (Lütkepohl, 1993, section 6.3; Reinsel, 1993, section 4.3.1) is the following:

$$\begin{aligned} \text{var}(\text{vec}(\Theta)) &= (1/n)\Sigma_U \otimes \Sigma_Y^{-1} \\ \text{vec}(\Sigma_Y) &= (I - (\Theta \otimes \Theta))^{-1}\text{vec}(\Sigma_U) \end{aligned} \tag{9}$$

This expression is particularly useful for high-dimensional models where numerical differentiation is computationally demanding. For low-dimensional models, Meilijson (1989) shows that numerical computation of the empirical observed information matrix consistently estimates the observed information matrix. We use Meilijson's method to calculate the standard errors for the mean hourly spreads in B, and we calculate the standard errors associated with the moving average parameters in θ using the analytical expression of equation (9). Another option for consistent standard error estimation is to use the bootstrapping method for state-space models of Stoffer and Wall (1991).

DOI: 10.1002/jae

4. EMPIRICAL ANALYSIS

We use day- and hour-ahead market clearing prices publicly available from the Open Access Same Time Information System of the CAISO, between August 1999 and August 2002. Although the reserve markets started their operations on April 1, 1998, the CAISO did not record hour-ahead prices until June 1, 1999 and did not distinguish regulation up from regulation down prices before August 18, 1999. It is also widely accepted that the entire restructured wholesale electricity sector in California performed without any major problems from its inception up to May 2000. It then entered an almost year-long period of severe crisis that ended in late June 2001. Because we are also interested in the reserve markets' performance in the aftermath of the crisis, we analyze three distinct periods: (1) pre-crisis (August 1999–April 2000); (2) crisis (May 2000–June 2001); and (3) post-crisis (July 2001–August 2002).

All the resources providing reserves in the CAISO control area came from two major zones (areas within the state connected by Path 15, a major transmission line). The first zone was north of Path 15 (NP15) and the second zone was south of Path 15 (SP15). We focus our efficiency tests on NP15 prices for the four products because NP15 and SP15 prices are almost always equal. During the post-crisis period, all four products have identical NP15 and SP15 prices both day- and hour-ahead (except for less than 0.1% of non-spinning observations). Only 8% of the NP15 and SP15 pre-crisis and crisis day- or hour-ahead prices were different for the four highest-quality products on average.

We define an observation as the spread between the hour- and day-ahead price. Therefore, we treat a spread as missing if either the day- or the hour-ahead price is missing. Missing day- and hour-ahead prices arise from self-provision or from substitution of lower-quality reserves with higher-quality reserves (e.g., substituting regulation in place of spinning).

Furthermore, missing hour-ahead prices arise due to the fact that CAISO sometimes procured all of its reserve requirements in the day-ahead markets only. As a result of these missing observations, the number of hourly observations in each of the three periods lies between 24 (post-crisis, spinning, hour 1) and 400 (crisis, regulation up, hour 16), with an average of 263. Consequently, we expect our estimates for some hours to be more precise than others. Overall, about 30% of the data are missing, which underlies the importance of accounting for the missing data using our maximum likelihood estimation procedure. Due to the large number of hours with missing observations for replacement, we exclude this product from our efficiency tests.

4.1. Efficiency Tests

The model in equation (1) contains 540 moving average parameters, so maximum likelihood estimation via numerical differentiation is computationally infeasible, even by gradient-based methods. For this reason, and to incorporate the missing data, we estimate the model by applying our EM algorithm of Section 3. Figure 1 shows that maximum likelihood estimates of mean spreads are negative for most hours in all three periods. For the pre-crisis period, most of the 24 hourly spreads do not differ significantly from zero for regulation down. Regulation up is the product with the largest number of significant hourly negative spreads (21). For the same period, we see significant spreads for 15 hours of spinning. For non-spinning, there are only four significant spreads and all are positive. The overwhelming majority of the hourly spreads for all products are statistically significant and negative for both the crisis and the post-crisis periods. For the crisis period, we see additional costs for day- versus hour-ahead purchases that are as high as

DOI: 10.1002/jae

$26.9/megawatt (regulation up, hour 12, crisis). In the post-crisis periods, we see more moderate spreads compared to the crisis period (e.g., regulation down, hour 6, $8.2/megawatt), but they are still large in percentage terms.

To assess the efficiency gain from maximum likelihood estimation, we calculated the 24 sample mean spreads and constructed confidence intervals using Newey–West standard errors (1 lag). These sample means closely approximate our maximum likelihood estimates, but the associated confidence intervals are wider than for our maximum likelihood estimates. This improvement in efficiency from maximum likelihood estimation translates into larger number of significant spreads, as shown in Table I.

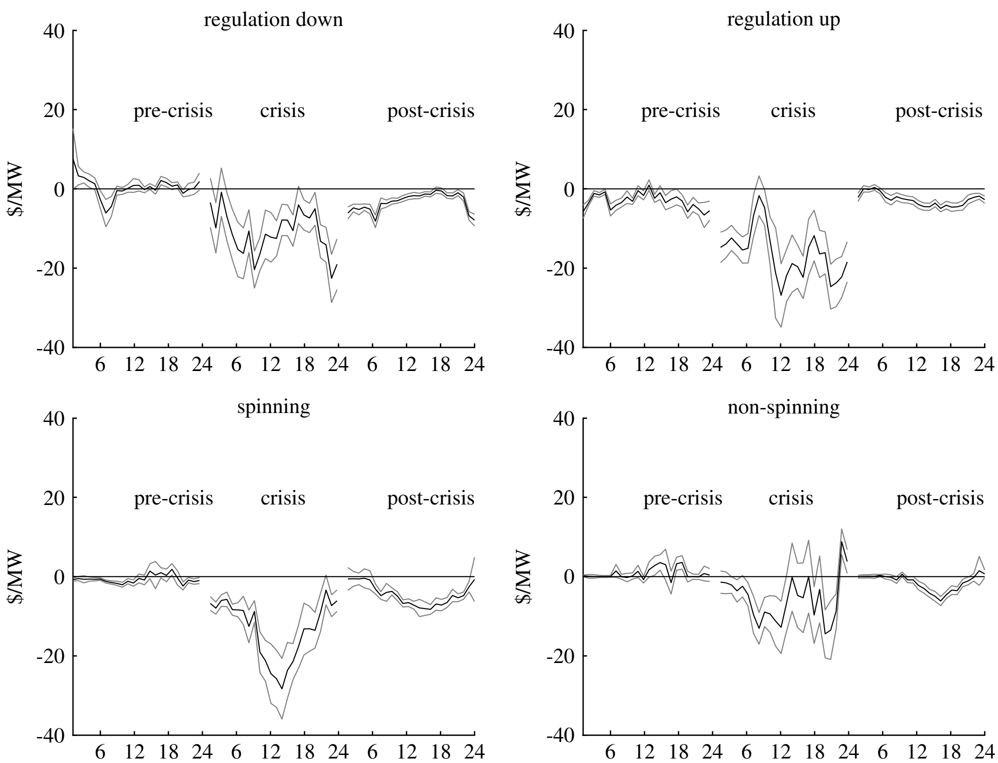

Figure 1. Spread estimates with 95% confidence intervals

Table I. Statistically significant negative price spreads at 0.05 level using Newey–West (1 lag) and VMA standard errors

	Pre-crisis		Crisis		Post-crisis	
Reserve	N–W	VMA	N–W	VMA	N–W	VMA
Regulation down	4	10	11	21	9	22
Regulation up	7	21	20	23	11	20
Spinning	15	15	17	23	18	18
Non-spinning	0	4	8	15	12	13

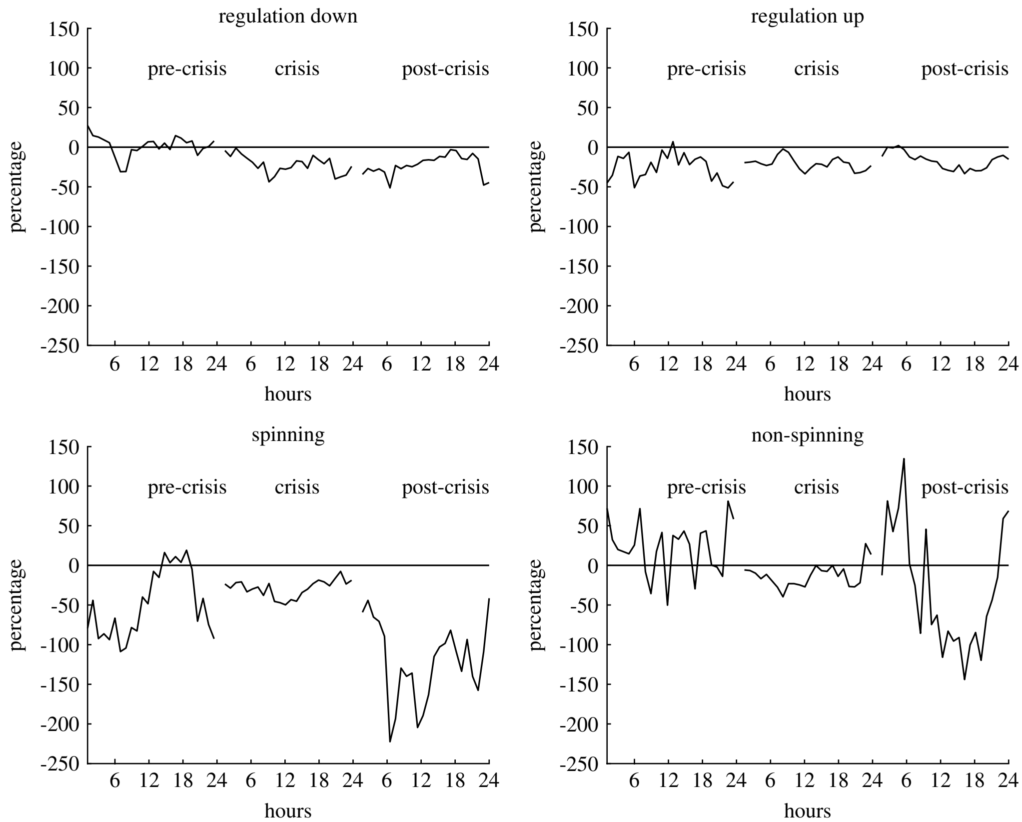

Figure 2. Percentage spread estimates

In Figure 2, we provide the percentage day-ahead premia, which are evaluated as the ratio of the mean price spread to the mean hour-ahead price for the same hour. For all three periods, the largest percentage day-ahead premia are observed for spinning: 109%, hour 7, pre-crisis; 50%, hour 12, crisis; 222%, hour 6, post-crisis. The corresponding spinning mean hour-ahead prices are: \$1.12/megawatt, \$49.31/megawatt and \$1.41/megawatt. These premia estimates well exceed those of earlier studies for energy markets both in California and other parts of the country. In Borenstein *et al.*, the largest spot premium observed (40%) is that for NP15 hours between 1 : 00 and 6 : 00 in September 2000. The largest significant unconditional hourly forward premium in Longstaff and Wang for the Pennsylvania–New Jersey–Maryland eastern hub between June 2000 and November 2002 is 14% for hour 20 : 00. In Saravia (2003), the implied day-ahead transmission premia for western and central New York are about 12% and 3% before the introduction of speculative trading in November 2001. After the introduction of speculative trading, the same premia reduce to 4.9% and 2.8%.

Our data exhibit substantial excess kurtosis such that, for all series, we reject normality with a p-value of 0.000 using a Jarque–Bera test. To account for the effect of the excess kurtosis on our estimates, we re-estimated our VMA(1) model and the 24 sample means treating any observation below the 1st and above the 99th percentiles as missing. These 'outliers' are due to sharp increases

DOI: 10.1002/jae

in the slope of the supply curves at high quantity levels, especially under tight supply conditions (see Knittel and Roberts, 2005). In most hours, the market clears on the flat portion of the supply curves, leading to a low price. However, the market occasionally clears on the steep portion of the supply curves and the price increases dramatically. The magnitude of the outliers was determined by the different price caps imposed by the Federal Energy Regulatory Commission (FERC). The price caps declined steadily throughout the sample, and generated kurtosis values greater than 50 for most pre-crisis hours, greater than 20 for most crisis hours, and around 10 for most post-crisis hours.

Excluding these outliers, the largest kurtosis was 16.8 (non-spinning, post-crisis), which is a dramatic reduction from the raw data and of a similar magnitude to that of many financial returns time series. We summarize the number of statistically significant spreads for each product by period in Table II. Because the results change little, we conclude that our findings are not driven by the excess mass in the tails of the distribution. With outliers treated as missing, the largest percentage day-ahead premia in all three periods still arise for spinning: 127%, hour 7, pre-crisis; 49%, hour 12, crisis; 240%, hour 6, post-crisis. The corresponding spinning mean hour-ahead prices are \$1.12/megawatt, \$46.36/megawatt and \$1.68/megawatt.

To indicate the nature of the estimated moving average coefficients, we illustrate the regulation down case (interval 17:00–18:00, all periods) in Figure 3, which shows that most of the information arrives in the market during business hours. The spikes between 13:00 and 18:00 are present for all products in all periods and are compatible with the comments of a member of the CAISO staff: 'The phones ring all the time up to late afternoon, then it is quiet.' Two daily events accentuate this pattern. The day-ahead procurement results and the market clearing prices were finalized by 13:00. Changes to final day-ahead schedules due to unit failures or other requirements with a potential impact on the hour-ahead markets were published by about 18:00.

4.2. Why did the Large Premia Persist?

On the demand side of the market, the CAISO did not react to the day-ahead premia because it had no incentive to shift to the lower-priced hour-ahead markets. As the agent for the purchasing scheduling coordinators, the CAISO had full discretion over the proportion of the reserves purchased in each market, but it was the scheduling coordinator that paid the bill. The reliability of the state's electricity grid was the primary objective of the CAISO, and ensuring the availability of sufficient reserves was an important tool in achieving this objective. By purchasing reserves a day ahead, the CAISO reduces the probability that insufficient reserves will be available in the

Table II. Statistically significant negative price spreads at 0.05 level treating outliers as missing using Newey–West (1 lag) and VMA standard errors

	Pre-crisis		Crisis		Post-crisis	
Reserve	N–W	VMA	N–W	VMA	N–W	VMA
Regulation down	7	8	12	22	15	24
Regulation up	16	24	19	23	13	21
Spinning	22	21	21	24	18	18
Non-spinning	13	1	15	19	14	12

DOI: 10.1002/jae

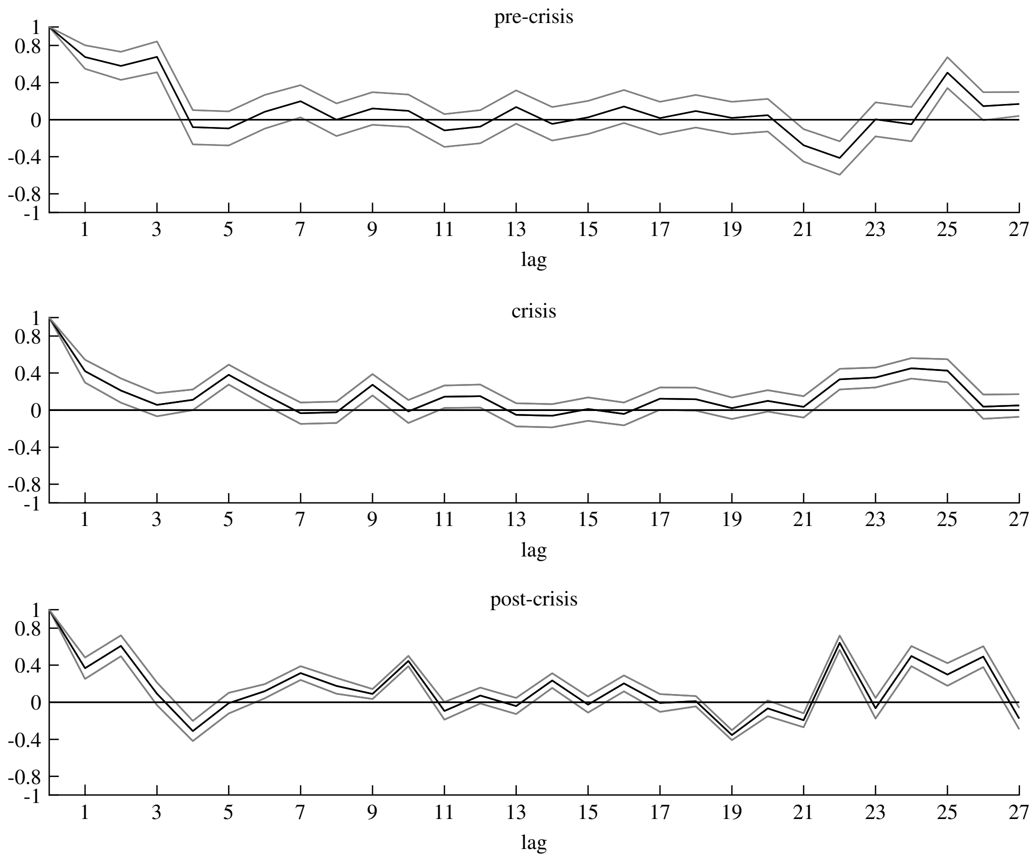

Figure 3. MA coefficients with 95% confidence intervals: regulation down 17 : 00–18 : 00

hour-ahead market, thereby reducing the probability of stage 2 or 3 emergencies.[1] To provide further support for the lack of the agent's incentives to shift purchases toward hour-ahead, we calculated the mean percentage of purchases made day-ahead on an hourly basis for the three periods. Figure 4 shows that the CAISO makes over 80% of procurements a day ahead for almost all products in all three periods.

Buyers of reserves could not shift away from day-ahead reserve prices that came with a high premium. However, they could avoid this premium by self-providing reserves and thereby skipping the market. Figure 5 shows a clear upward trend in the percentage of megawatts self-provided by scheduling coordinators for all types of reserves over the three periods. This upward trend is most pronounced for spinning and non-spinning: about 20% in the pre-crisis period, well above 50% during crisis, and close to 80% in the post-crisis. Therefore, the scheduling coordinators are increasingly seeking to trade with themselves by self-providing, rather than entering the market

[1] The CAISO issues emergency notices when the percentage of actual or anticipated reserves falls below those dictated by regional operating reliability criteria. Under the last (worst) stage of emergency, stage 3, when reserves are between 1.5% and 3% of the demand served, involuntary curtailment of service to consumers ('rotating outages') is required. For example, rotating load interruptions of 500 MW were declared during stage 3 emergency on March 19, 2001.

DOI: 10.1002/jae

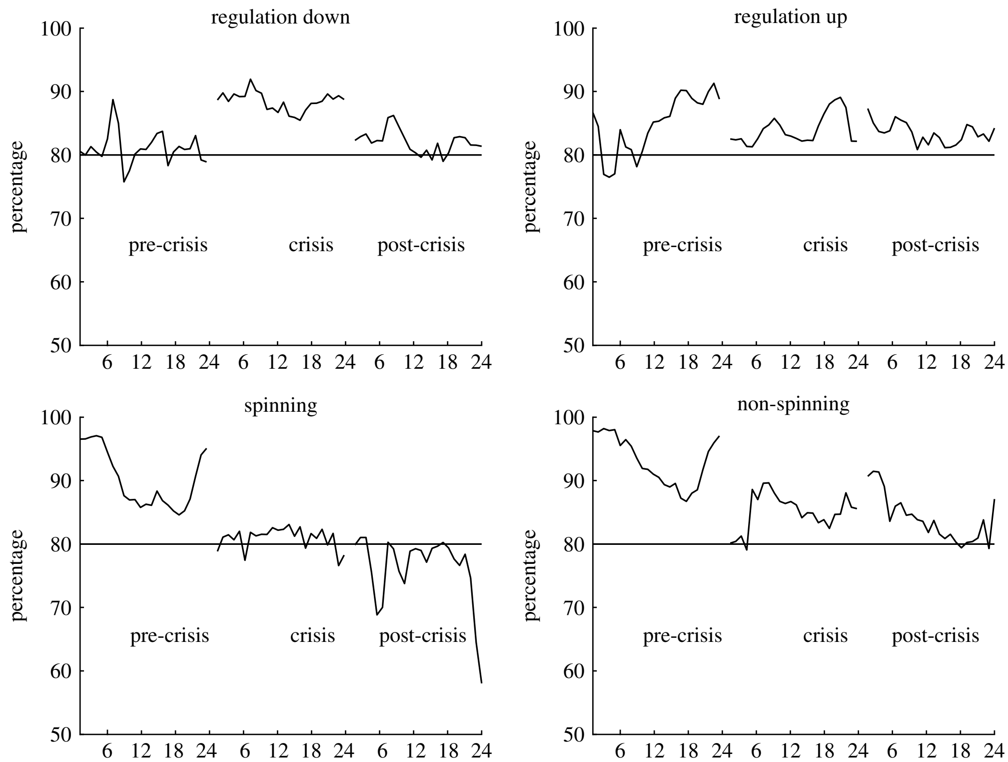

Figure 4. Percentage of megawatt procurements day-ahead

and trading with other scheduling coordinators. Given the low transaction costs to participating in the market, this pattern strongly suggests that the market is inefficient and the day-ahead cost is excessive. Moreover, by removing themselves from the market the scheduling coordinators reduce liquidity and cause the market to be less efficient.

In spite of the demand constraints, why did the suppliers not ultimately bid the day-ahead prices down? The answer lies in restrictive and vague market rules that raised barriers to entry. According to the market rules, sellers could not offer reserves from resources that they did not physically identify by name and location within the CAISO control area. The CAISO maintained the right to randomly test the availability of reserves from resources attached to their winning bids. A compliance failure to such a test implied payment rescissions and penalties (CAISO Tariff, section 2.5.26). Hence, by precluding purely speculative trading, the market rules raised barriers to entry and reduced the chances for the incumbents to shrink the price differences.

To take advantage of the day-ahead premia, a speculator would sell reserves day-ahead and buy back hour-ahead. For example, Table I shows peak-hour average spreads for regulation down equal to −1.90 (pre-crisis), −15.27 (crisis), and −2.69 (post-crisis). These spreads represent the average daily profits from selling day- and buying back hour-ahead. Dividing these average spreads by the corresponding standard deviations yields −0.05 (pre-crisis), −0.15 (crisis), and −0.14 (post-crisis), and similar calculations for other products and hours produce ratios of a similar order of magnitude. These ratios far exceed the ratio of daily returns to daily standard deviation for

DOI: 10.1002/jae

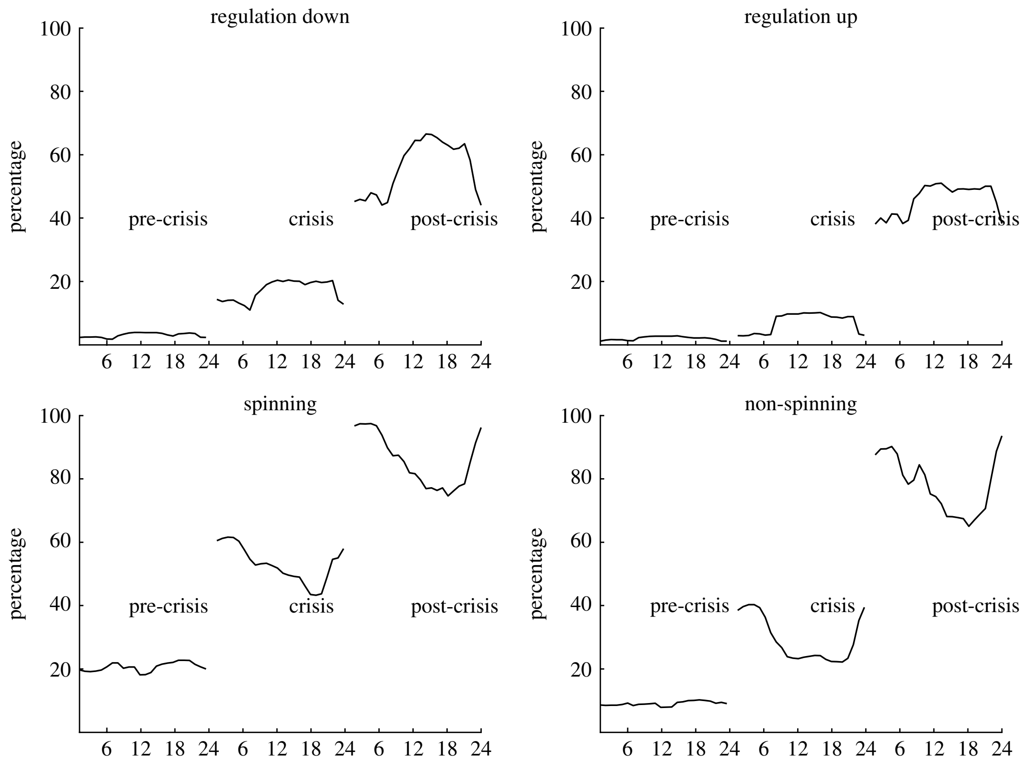

Figure 5. Percentage of megawatt self-provided

the S&P 500 (the 'Sharpe ratio'), which was only 0.03 during this sample period. Moreover, the correlation between the daily day-ahead premia and the returns on the S&P 500 is negligible. Thus, in a market with clear rules and no entry barriers, even risk-averse speculators would be willing to trade on the day-ahead premia.

From our discussions with the CAISO staff, we understand that hour-ahead buybacks of reserves previously sold to the day-ahead market were considered a legitimate practice. However, the same conclusion cannot be easily drawn from the Market Monitoring and Information Protocol (MMIP) of the CAISO Tariff. The MMIP is the code of conduct in the CAISO markets and it provides a description of what constitutes anomalous market behavior and gaming, which are termed practices subject to regulatory scrutiny. The MMIP definition of gaming is vague and is characterized as 'Taking unfair advantage of the rules and procedures set forth in the CAISO Tariff ... It may also include actions or behaviors that may render the CAISO Markets vulnerable to price manipulation to the detriment of their efficiency.' This lack of clarity in the definition of gaming likely induced uncertainty about sanctions and penalties associated with profitable hour-ahead buybacks. It also served as an extra entry barrier that protected those incumbents who had managed to learn the market rules well.

The last impediment to the elimination of the day-ahead premia due to entry forces came from a desire for semi-autarky. The provision of regulation was restricted exclusively to sellers with resources located within the CAISO control area and there was a quota (50%) on imports regarding

DOI: 10.1002/jae

spinning and non-spinning reserves from outside (CAISO, Ancillary Services Requirements Protocol, sections 4.1–5.3). These restrictions also served as barriers to entry.

According to the annual reports by the CAISO Department of Market Analysis, total reserve procurement costs over the three periods were $178 million (pre-crisis), $3952 million (crisis), and $302 million (post-crisis), for a total of $4.4 billion. Given that about 80% of this $4.4 billion was incurred at approximately a 35% premium through day-ahead purchases, market participants left a large amount of money on the table. However, a limited number of incumbents did exploit the implied profitable trading opportunities by selling reserves day-ahead and buying them back hour-ahead. Table III shows calculations by the CAISO staff, which reveal that six scheduling coordinators made net gains totaling $57 million from buybacks during 2000 and 2001 (CAISO, Department of Market Analysis, October 2002). This CAISO report was prepared as part of the 'catch-all' FERC investigation of price manipulation in electricity and natural gas markets in the West, which was initiated in February 2002.

Enron's filing for bankruptcy in December 2001 and the subsequent substantial decline in spot electricity and natural gas prices triggered allegations that it had manipulated them. Five months later, in May 2002, FERC released on its web site two internal Enron memoranda from December 2000 that described various fraudulent trading schemes employed by its traders in the Western markets. Among these strategies, 'Get Shorty' took advantage of the differences between the day- and hour-ahead reserve prices by selling day-ahead and buying back hour-ahead, but in an illegal way. The recipe was based almost exclusively on the submission of false information by Enron traders who understood well the market rules.

Two chief Enron executives of its West Power Trading Division (West Power), Timothy Belden and Jeffrey Richter, fully acknowledged their repeated fraudulent business practices in CAISO reserve markets in their written guilty pleas (United States of America: plaintiff vs. Timothy N. Belden: defendant, United States of America: plaintiff vs. Jeffrey S. Richter: defendant). Acting as if market rules allowed pure speculative trading, West Power submitted numerous times day-ahead reserve bids from resources that did not have standby available supply. Most commonly, Enron traders scheduled resources outside the CAISO control area. They did so because the market rules required only the disclosure of the point of entry to the CAISO control area for reserves coming from outside resources. The strategy was successful in most cases. They bought back in the hour-ahead the entire amount of the day-ahead commitment, eliminating the chances for payment rescissions and penalties and generating millions of dollars in profits (Table III). Although FERC concluded that 'Get Shorty' violated the anti-gaming prohibition of the CAISO MMIP (FERC, 2003), CAISO recognized that it was impossible to identify whether the remainder of the scheduling coordinators in Table III also engaged in the illegal buybacks that Enron employees confessed to enacting.

Table III. Buybacks by scheduling coordinator

Scheduling coordinator	Gains ($)	Losses ($)	Net ($)
Coral Power, LLC	18,140,839	−1,026,754	17,114,085
Sempra Energy Trading Corporation	13,436,678	−376,652	13,060,026
Avista Energy Inc.	11,977,712	− 149,293	11,828,418
Modesto Irrigation District	10,583,973	−266,593	10,317,380
Enron Power Marketing Inc.	5,311,040	−256,312	5,054,728
British Columbia Power Exchange	1,351,613	−345,586	1,006,027

DOI: 10.1002/jae

5. CONCLUDING REMARKS

We estimate premia that are significant, both economically and statistically, for purchases in the day- over the hour-ahead electricity reserves markets in California between 1999 and 2002. We show that the purchase of one unit of reserves day-ahead came at a price up to three times larger than its hour-ahead counterpart, even a year after the official end of the state's energy crisis. Taking into account the markets' structure, we obtain these day-ahead premia by estimating a high-dimensional vector moving-average process via exact maximum likelihood. In doing so, we derive a new state-space EM algorithm that generates analytical expressions of the likelihood equations.

We attribute the pronounced market inefficiencies to the poor market design. A principal–agent relationship between the reserves' buyers (principal) and the markets' supervisory authority (agent) eliminated the incentives on the demand side of the market to shift purchases to the lower-priced hour-ahead market. Consequently, the buyers increasingly avoided the market by self-providing reserves. Design features on the supply side raised entry barriers by precluding speculative trading and inducing uncertainty about which practices were subject to regulatory scrutiny and sanctions. As a result, a handful of incumbents enjoyed millions of dollars in profits, but much more was left on the table.

The outcome of the intrinsically complicated set of reserves market rules was information scarcity rents, which some market participants collected by means of illegal business practices. This was exactly the case with the Enron traders and their so-called 'Get Shorty' strategy. In the absence of entry barriers, these rents would have been collected in a very short period of time, and the premia would have disappeared.

Although ambiguous electricity market rules provide the regulator with the flexibility to redefine them at will, they come at the cost of potential market failures. These costs should not be ignored, especially in light of the Standard Market Design efforts of the Federal Energy Regulatory Commission that have been in place since the late spring of 2002. Policy makers should realize that practices such as paying a premium to procure reserves a day ahead impose large costs on the system, which may well exceed any real-time reliability benefits.

ACKNOWLEDGEMENTS

Smith is a member of the Giannini Foundation of Agricultural Economics, and this research was partially funded by a Giannini Foundation Minigrant. We would like to thank Christopher Knittel for invaluable comments and suggestions on earlier drafts of this paper. We also benefited from the feedback by James Bushnell and James Prieger, and by the participants in the UC Davis Econometrics (Spring 2005) and UCEI (Fall 2005) seminar series. Jamal Batakji, from the California ISO, and Frank Wolak generously shared with us their knowledge of the markets' operations. Finally, we are particularly grateful to the comments of the co-editor John Rust and two anonymous referees, which significantly improved the final version of the draft. All remaining errors and omissions are ours.

REFERENCES

Bessembinder H, Lemmon M. 2002. Equilibrium pricing and optimal hedging in electricity forward markets. *Journal of Finance* **57**: 1347–1382.

DOI: 10.1002/jae

Borenstein S, Bushnell J, Knittel CR, Wolfram C. 2004. Inefficiencies and market power in financial arbitrage: a study of California's electricity markets. *POWER Working Paper PWP-138*, University of California Energy Institute.
California ISO Department of Market Analysis. 2002. Analysis of trading and scheduling strategies described in Enron memos.
California Department of Justice Office of the Attorney General. 2004. Attorney General's Energy White Paper, A Law Enforcement Perspective on the California Energy Crisis: Recommendations for Improving Enforcement and Protecting Consumers in Deregulated Energy Markets.
Chao HP, Wilson R. 2002. Multi-dimensional procurement auctions for power reserves: robust incentive-compatible scoring and settlement rules. *Journal of Regulatory Economics* **22**: 161–183.
Dempster A, Laird N, Rubin DB. 1977. Maximum likelihood estimation from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society (Series B)* **39**: 1–38.
Durbin J, Koopman SJ. 2001. *Time Series Analysis by State Space Models*. Oxford University Press: Oxford.
Engle RF, Watson M. 1983. Alternative algorithms for the estimation of dynamic factor, MIMIC and varying coefficient regression models. *Journal of Econometrics* **23**: 385–400.
Federal Energy Regulatory Commission. 2003. Final *Report on Price Manipulation in Western Markets, Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices*, Docket No. PA02-2-000.
Hamilton J. 1994. *Time Series Analysis*. Princeton University Press: Princeton, NJ.
Knittel CR, Roberts MR. 2005. An empirical examination of restructured electricity prices. *Energy Economics* **27**: 791–817.
Koopman SJ, Durbin J. 2000. Fast filtering and smoothing for multivariate state space models. *Journal of Time Series Analysis* **21**: 281–296.
Krishnan T, McLachlan T. 1997. *The EM Algorithm and Extensions*. Wiley: New York.
Longstaff FA, Wang AW. 2004. Electricity forward prices: a high-frequency empirical analysis. *Journal of Finance* **59**: 1877–1900.
Lütkepokhl H. 1993. *Introduction to Multiple Time Series Analysis*. Springer: New York.
Meng XL, Rubin DB. 1994. On the global and component-wise rates of convergence of the EM algorithm. *Linear Algebra and its Applications* **199**: 413–425.
Meilijson I. 1989. A fast improvement to the EM algorithm on its own terms. *Journal of the Royal Statistical Society (Series B)* **51**: 127–138.
Pieris S. 1988. On the study of some functions of multivariate ARMA processes. *Journal of Multivariate Analysis* **25**: 146–151.
Reinsel GC. 1993. *Elements of Multivariate Time Series Analysis*. Springer: New York.
Saravia C. 2003. Speculative trading and market performance: the effect of arbitrageurs on efficiency and market power in the New York electricity market. *CSEM Working Paper CSEM WP 121*, University of California Energy Institute.
Shumway RH, Stoffer DS. 1982. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis* **3**: 253–264.
Stoffer DS, Wall KD. 1991. Bootstrapping state-space models: Gaussian maximum likelihood estimation and the Kalman filter. *Journal of the American Statistical Association* **86**: 1024–1033.

DOI: 10.1002/jae