Covariance Matrix Estimation using an Errors-in-Variables Factor Model with Applications to Portfolio Selection and a Deregulated Electricity Market

Warren R. Scott, Warren B. Powell

Sherrerd Hall, Charlton Street, Princeton, NJ 08544

Abstract

We propose an errors-in-variables factor model which extends the classical capital asset pricing model (CAPM) to the case where the market returns contain additive noise. Using the model, we propose a method for choosing portfolios of assets, such as U.S. stocks in the S&P 100 and virtual electricity contracts in a regional transmission organization. Virtual electricity contracts relate real-time and day-ahead electircity prices and can be used for hedging or speculation. Using the errors-in-variables factor model, we first estimate the covariance matrix of electricity prices at different locations and then construct efficient portfolios which balance risk and return. We present applications using US equities and virtual electricity contracts to show the new covariance matrix estimation technique can effectively manage risk.

Keywords:
Portfolio Selection, Factor Models, Errors-in-Variables, Electricity Market, Covariance Matrix

---

1. Introduction

Covariance matrices have been fundamental to choosing diversified portfolios of assets dating back to Harry Markowitz (see Markowitz (1999)). However, the sample covariance matrix is often ill-conditioned and typically not appropriate for Markowitz portfolio theory. Additionally, when calculating

Email addresses: `wscott@alumni.princeton.edu` (Warren R. Scott), `powell@princeton.edu` (Warren B. Powell)

the covariance matrix of non-stationary random variables, estimation with data that is too old may not effectively capture the current market conditions. The use of factor models to estimate large covariance matrices of asset returns dates back to William Sharpe (see Sharpe (1963)). The most well-known factor models for capital assets are the capital asset pricing model, which uses excess market returns as the only factor (see Sharpe (1964)), and the Fama-French 3-factor model (see Fama and French (1992)), which uses three factors: excess market returns, small minus big market capitalization, and high minus low book-to-market ratio.

In the United States, regional transmission organizations manage the generation and transmission of electricity. For example, PJM is an independent and profit-neutral regional transmission organization that is responsible for the grid operation and reliability for most of the northeast United States (see PJM (2011)). PJM sets the locational marginal prices (LMP's) taking into account the sizes and locations of the electricity generation plants and loads, along with transmission constraints. The day-ahead PJM market determines the day-ahead electricity price at each node for each hour of the day. All bids and offers for electricity (including virtual bids) are submitted by noon the day-ahead. At 4 pm PJM determines the next day's day-ahead LMP's. Virtual bids that clear in the day-ahead market are purely financial contracts, typically referred to as "incs" (increments) or "decs" (decrements) for selling or buying electricity in the day-ahead market, respectively. The payoff of these contracts is the difference between the real-time and day-ahead price, minus transaction fees. These contracts can be used by generators and load serving entities to hedge positions or speculate.

This paper makes the following contributions: (1) We propose an errors-in-variables extension of the CAPM that accounts for errors in the observations of the market returns. (2) We show how the errors-in-variables version of CAPM can be used to estimate the covariance matrix of the returns of assets, especially when we use relatively short histories so that we can capture short-term changes in volatility. (3) We evaluate portfolios optimized using the covariance matrices produced by CAPM, errors-in-variables CAPM, and the Fama-French three-factor model on the S&P 100. (4) We evaluate portfolios optimized using covariance matrices produced by CAPM and errors-in-variables CAPM in the PJM electricity market.

This paper is organized as follows. Section 2 summarizes how factor models are used to estimate covariance matrices. We present an errors-in-variables factor model which assumes there are errors in the observations of the market returns, which can be used to construct a covariance matrix. Section 3 summarizes the typical Markowitz portfolio which requires a covariance matrix for the returns of the assets. We present an extension of the Markowitz portfolio which can be used for choosing portfolios of electricity contracts. In Section 4, we show that the covariance matrix constructed with the errors-in-variables factor model has benefits for choosing portfolios of U.S. equities as well as portfolios of electricity contracts.

2. Estimating Covariance Matrices

In this section, we show how CAPM and factor models can be used to estimate covariance matrices. We then propose an extension to CAPM which we call the errors-in-variables CAPM, which allows for errors in the observations of the market returns. We show how the errors-in-variables CAPM can be used to estimate the covariance matrix of asset returns, even if we use a relatively short price history.

2.1. The Capital Asset Pricing Model

The capital asset pricing model (CAPM) described in Sharpe (1964), Lintner (1965), and Black (1972) relates the excess return of an asset to the excess return of the market portfolio under certain market equilibrium conditions. Assume we have N assets. Let Y_j be the excess return of asset j where $j = 1, ..., N$, and let Y^m be the excess return of the market. Assuming each investor invests in a mean-variance optimal portfolio, CAPM states

$$Y_j = \alpha_j + \beta_j Y^m + \epsilon_j, \tag{1}$$

where $\mathbb{E}[\epsilon_j] = 0$, and is independent from Y^m. CAPM concludes that $\alpha_j = 0$, for $j = 1, ...N$. Assuming α_j and β_j are deterministic, it follows from Equation (1) that

$$\text{Cov}[Y_j, Y^m] = \beta_j \text{Var}[Y^m]. \tag{2}$$

Replacing the true covariance with the sample covariance, this gives an equation for estimating β_j which is equivalent to the least-squares estimate of β_j,

$$\hat{\beta}_j = \frac{\text{Cov}[Y_j, Y^m]}{\text{Var}[Y^m]}. \tag{3}$$

One powerful application of CAPM is building covariance matrices, where we can write

$$
\begin{aligned}
\mathrm{Cov}[Y_i, Y_j] &= \mathrm{Cov}[\alpha_i + \beta_i Y^m + \epsilon_i, \alpha_j + \beta_j Y^m + \epsilon_j] \\
&= \mathrm{Cov}[\beta_i Y^m + \epsilon_i, \beta_j Y^m + \epsilon_j] \\
&= \mathrm{Cov}[\beta_i Y^m, \beta_j Y^m] + \mathrm{Cov}[\beta_i Y^m, \epsilon_j] + \mathrm{Cov}[\beta_j Y^m, \epsilon_i] + \mathrm{Cov}[\epsilon_i, \epsilon_j] \\
&= \beta_i \beta_j \mathrm{Var}[Y^m] + \beta_i \underbrace{\mathrm{Cov}[Y^m, \epsilon_j]}_{0} + \beta_j \underbrace{\mathrm{Cov}[Y^m, \epsilon_i]}_{0} + \mathrm{Cov}[\epsilon_i, \epsilon_j] \\
&= \beta_i \beta_j \mathrm{Var}[Y^m] + \mathrm{Cov}[\epsilon_i, \epsilon_j]. \qquad (4)
\end{aligned}
$$

It is typically assumed $\mathrm{Cov}[\epsilon_i, \epsilon_j] = 0, \quad \forall i \neq j$ (see Sharpe (1963)). In matrix notation with N assets,

$$\underbrace{Y}_{Nx1} = \underbrace{\alpha}_{Nx1} + \underbrace{\beta}_{Nx1} \underbrace{Y^m}_{1x1} + \underbrace{\epsilon}_{Nx1}. \qquad (5)$$

and

$$\mathrm{Cov}[Y] = \beta \beta^T \mathrm{Var}[Y^m] + \mathrm{Var}[\epsilon], \qquad (6)$$

where $\mathrm{Var}[\epsilon]$ is diagonal by assumption. The estimator given by Equation (6) is always positive semidefinite and guaranteed to be positive definite when $\mathrm{Var}[\epsilon]$ has full rank. This is a major advantage over the sample covariance matrix which can easily be rank deficient. However, it is well known that difficulties may arise when using Equation (3) to estimate β. Barra (2000) says that, in practice, fundamental changes in an asset as well as specific events that affect an asset may distort historical values of β. For this reason we would like to create a more robust version of CAPM that allows for errors in the market returns, which also allows us to use more recent histories that capture short term changes in volatility.

2.2. Errors-in-Variables CAPM

We now formulate our model differently than Equation (1). We first assume the true excess return of asset j has a linear relationship with the true excess market return, where $j = 1, ...N$,

$$Y_j = \alpha_j + \beta_j Y^m. \qquad (7)$$

We next assume we do not observe Y_j, the true excess return of asset j; we only observe Y_j', which is equal to the true excess return of asset j plus the noise Y_j'',

$$Y_j' = Y_j + Y_j''. \qquad (8)$$

Analogously, we next assume we do not observe Y^m, the true excess return of the market; we only observe $Y^{m'}$, which is equal to the true excess return of the market plus the noise $Y^{m''}$,

$$Y^{m'} = Y^m + Y^{m''}. \tag{9}$$

To summarize, at each time period, we are only able to observe $Y^{m'}, Y'_1, Y'_2, ..., Y'_N$. In Section 2.1, we assumed we could observe the true excess return of the market; now we assume we can only observe a noisy version of the excess return of the market.

Using Equations (7), (8), and (9) we are able to derive an analogous estimator to Equation (6) for estimating the covariance matrix of assets,

$$\begin{aligned}
\mathrm{Cov}[Y'_i, Y'_j] &= \mathrm{Cov}[\alpha_i + \beta_i Y^m + Y_i^{''}, \alpha_j + \beta_j Y^m + Y_j^{''}] \\
&= \mathrm{Cov}[\beta_i Y^m + Y_i^{''}, \beta_j Y^m + Y_j^{''}] \\
&= \mathrm{Cov}[\beta_i Y^m, \beta_j Y^m] + \mathrm{Cov}[\beta_i Y^m, Y_j^{''}] + \mathrm{Cov}[\beta_j Y^m, Y_i^{''}] + \mathrm{Cov}[Y_i^{''}, Y_j^{''}] \\
&= \beta_i \beta_j \mathrm{Var}[Y^m] + \beta_i \underbrace{\mathrm{Cov}[Y^m, Y_j^{''}]}_{0} + \beta_j \underbrace{\mathrm{Cov}[Y^m, Y_i^{''}]}_{0} + \mathrm{Cov}[Y_i^{''}, Y_j^{''}] \\
&= \beta_i \beta_j \mathrm{Var}[Y^m] + \mathrm{Cov}[Y_i^{''}, Y_j^{''}],
\end{aligned}$$

where we assumed Y^m and $Y_j^{''}$ were uncorrelated for each asset $j = 1, ..., N$. This can be written in matrix form as

$$\mathrm{Cov}[Y'] = \beta\beta^T \mathrm{Var}[Y^m] + \mathrm{Var}[Y^{''}]. \tag{10}$$

In CAPM described in Section 2.1, estimating $\mathrm{Var}[Y^m]$ is straightforward using the sample variance, because we assumed we could observe Y^m. In this section, we assume we only observe $Y^{m'}$ and hence we must first estimate Y^m before we can estimate $\mathrm{Var}[Y^m]$.

2.2.1. Estimating β

The method of instrumental variables can be used to estimate β in the errors-in-variables model given in Equations (7), (8), and (9) (see Durbin (1954); Kendall and Stuart (1961); Söderström and Stoica (1983)). An instrumental variable, Z, should be correlated with the market return, Y^m, but uncorrelated with the errors in the observations of Y^m and Y_j. We assume we have M observations of the returns of N different assets. The following equation gives the instrumental variables estimate of β_j for each asset $j = 1, ..., N$:

$$Z^T Y^{m'} \hat{\beta}_j = Z^T Y_j^{'}. \tag{11}$$

If $\lim_{M\to\infty} \frac{1}{M} Z^T Y_j^{m''} = 0$, $\lim_{M\to\infty} \frac{1}{M} Z^T Y_j^{''} = 0$, and $\lim_{M\to\infty} \frac{1}{M} Z^T Y^m = \Sigma$, then Equation (11) yields a consistent estimate for β_j,

$$\begin{aligned} Z^T(Y^m + Y^{m''})\hat{\beta}_j &= Z^T(Y^m \beta_j + Y_j^{''}), \\ (Z^T Y^m + Z^T Y^{m''})\hat{\beta}_j &= Z^T Y^m \beta_j + Z^T Y_j^{''}. \end{aligned}$$

Now, taking the limit as M goes to infinity,

$$\begin{aligned} \lim_{M\to\infty} \frac{1}{M} \left(Z^T Y^m + Z^T Y^{m''} \right) \hat{\beta}_j &= \lim_{M\to\infty} \frac{1}{M} \left(Z^T Y^m \beta_j + Z^T Y_j^{''} \right), \\ \lim_{M\to\infty} \frac{1}{M} Z^T Y^m \hat{\beta}_j &= \lim_{M\to\infty} \frac{1}{M} Z^T Y^m \beta_j, \\ \Sigma \left(\lim_{M\to\infty} \hat{\beta}_j \right) &= \Sigma \beta_j, \\ \lim_{M\to\infty} \hat{\beta}_j &= \beta_j. \end{aligned}$$

For the numerical work in this paper, we use the returns of an equal-weighted portfolio as a reasonable instrumental variable.

2.2.2. Estimating Y^m

We first rewrite Equations (7), (8), (9) using an additional index for each of the M observations. Letting $i = 1, ..., M$ be the index for the observation and $j = 1, ..., N$ be the index for the asset, we can write

$$Y_{ij} = \alpha_j + \beta_j Y_i^m, \tag{12}$$

$$Y_{ij}' = Y_{ij} + Y_{ij}^{''}, \tag{13}$$

$$Y_i^{m'} = Y_i^m + Y_i^{m''}. \tag{14}$$

In order to obtain estimates for Y^m in closed form, we make the following assumptions:

Assumption 1. $\{Y_i^{m''}\}_{i=1}^M$ *are i.i.d. and*

$$Y_i^{m''} \sim \mathcal{N}(0, (\sigma^m)^2), \qquad i = 1, ..., M.$$

Assumption 2. $\{Y_{ij}^{''}\}_{i=1}^{M}$ *are i.i.d. and*

$$Y_{ij}^{''} \sim \mathcal{N}(0, (\sigma_j)^2), \qquad i = 1, ..., M, j = 1, ..., N.$$

Assumption 3.

$$\sigma^m = \sigma_1 = \cdots = \sigma_N.$$

Assumption 4.

$$\alpha_1 = \cdots = \alpha_N = 0.$$

In order to estimate the market returns, $Y_1^m, \cdots, Y_M^m$, we first write the log-likelihood:

$$\begin{aligned}
&l(Y_1^m, \cdots, Y_M^m) \\
&= \sum_{i=1}^{M} \left[-\frac{1}{2}\ln(2\pi(\sigma^m)^2) - \frac{(Y_i^{m'} - Y_i^m)^2}{2(\sigma^m)^2} \right] + \sum_{j=1}^{N}\sum_{i=1}^{M} \left[-\frac{1}{2}\ln(2\pi(\sigma_j)^2) - \frac{(Y_{ij}^{'} - Y_{ij})^2}{2(\sigma_j)^2} \right] \\
&= \sum_{i=1}^{M} \left[-\frac{1}{2}\ln(2\pi(\sigma^m)^2) - \frac{(Y_i^{m'} - Y_i^m)^2}{2(\sigma^m)^2} \right] + \sum_{j=1}^{N}\sum_{i=1}^{M} \left[-\frac{1}{2}\ln(2\pi(\sigma_j)^2) - \frac{(Y_{ij}^{'} - \beta_j Y_i^m)^2}{2(\sigma_j)^2} \right] \\
&= \sum_{i=1}^{M} \left[-\frac{1}{2}\ln(2\pi(\sigma^m)^2) - \frac{(Y_i^{m'} - Y_i^m)^2}{2(\sigma^m)^2} \right] + \sum_{j=1}^{N}\sum_{i=1}^{M} \left[-\frac{1}{2}\ln(2\pi(\sigma_j)^2) - \frac{(Y_{ij}^{'} - \beta_j Y_i^m)^2}{2(\sigma^m)^2} \right].
\end{aligned}$$

Now, maximizing this expression with respect to $Y_1^m, \cdots, Y_M^m$ we get,

$$\begin{aligned}
&\underset{Y_1^m, \cdots, Y_M^m}{\arg\max} \left[\sum_{i=1}^{M} \left[-\frac{(Y_i^{m'} - Y_i^m)^2}{2(\sigma^m)^2} \right] + \sum_{j=1}^{N}\sum_{i=1}^{M} \left[-\frac{(Y_{ij}^{'} - \beta_j Y_i^m)^2}{2(\sigma^m)^2} \right] \right] \\
= \;&\underset{Y_1^m, \cdots, Y_M^m}{\arg\max} \left[\sum_{i=1}^{M} -(Y_i^{m'} - Y_i^m)^2 + \sum_{j=1}^{N}\sum_{i=1}^{M} -(Y_{ij}^{'} - \beta_j Y_i^m)^2 \right]
\end{aligned}$$

Now taking the derivative with respect to Y_i^m and setting it to zero, we get

$$
\begin{aligned}
2(Y_i^{m'} - Y_i^m) + \sum_{j=1}^{N} 2(Y_{ij}' - \beta_j Y_i^m) &= 0, \\
Y_i^{m'} - Y_i^m + \sum_{j=1}^{N} Y_{ij}' - \sum_{j=1}^{N} \beta_j Y_i^m &= 0, \\
Y_i^m + Y_i^m \sum_{j=1}^{N} \beta_j &= Y_i^{m'} + \sum_{j=1}^{N} Y_{ij}', \\
Y_i^m &= \frac{Y_i^{m'} + \sum_{j=1}^{N} Y_{ij}'}{1 + \sum_{j=1}^{N} \beta_j}. \qquad (15)
\end{aligned}
$$

This equation gives an estimator for the true market excess return Y_i^m at each period $i = 1, ..., M$. Note that the estimate of the scalar Y_i^m depends on the β for each asset as well as the observed excess return for each asset in period i, $Y_{i1}', ..., Y_{iN}'$. This is fundamentally very different from classical CAPM, because now the regression for each asset depends on the regression of all the other assets as well.

2.2.3. Visualizing the Difference

CAPM treats the market returns as noise-free, but our errors-in-variables CAPM explicitly assumes the market returns have noise. When there are errors in the explanatory variable, Durbin (1954) and Kendall and Stuart (1961) show that least-squares linear regression typically yields an estimate of the slope which is too small in magnitude. This may correspond to the empirical observation that the CAPM estimate of β is typically too small in magnitude (see Barra (2000)). In Figure 1(a) and 1(b) we show the CAPM estimates of β for Chevron and ExxonMobil using daily data over 50 days. The excess returns of the market are shown on the horizontal axis and the excess returns of Chevron and ExxonMobil are shown on the vertical axis, respectively. The residuals are drawn vertically because there is no noise in the market returns.

In Figure 2(a) and 2(b) we show the errors-in-variables CAPM estimates of β for Chevron and ExxonMobil. The excess returns of the market are shown on the horizontal axis and the excess returns of Chevron and Exxon-Mobil are shown on the y-axis, respectively. The residuals are no longer drawn vertically because there is noise in the market returns. The residuals are now estimated using the equations described in Sections 2.2.1 and 2.2.2. The estimates of β are now larger as predicted by the theory. In particular, if the estimates of β are too small in magnitude, we should expect the residuals to be correlated across assets. Figure 3 shows the residuals for Chevron and ExxonMobil using CAPM, and they appear to be positively correlated. Figure 3(b) shows the residuals for Chevron and ExxonMobil using the errors-in-variables CAPM. Visually, the residuals appear to be smaller and less correlated, as desired.

3. Portfolio Selection

In the presence of a risk-free asset and shorting, it is well known that the optimal Markowitz portfolio can be calculated in closed form. Below we briefly summarize the main points (see Campbell et al. (1997)).

3.1. Markowitz with a Risk-free Asset

Markowitz (1952) and Sharpe (1963) describe a method for choosing a portfolio of assets in the presence of a risk-free asset. Let $\bar{r}$ be an $N \times 1$ column vector of the expected returns over one period of the N risky assets, and assume the risk-free return over one period is r_0. Let $\bar{Y}$ be an $N \times 1$ column vector of the expected excess returns of the N risky assets, $\bar{Y}_j = \bar{r}_j - r_0$, $j = 1, ..., N$. Let Σ be the $N \times N$ covariance matrix of the returns of the N risky assets over the next time period. Let α be an $N \times 1$ column vector where α_j represents the proportion of total wealth invested in asset j, $j = 1, ..., N$. In addition, let α_0 be the proportion of total wealth invested in the risk-free asset, and assume $\alpha_0 + \alpha_1 + \cdots + \alpha_N = 1$.

The expected return of the portfolio α after one period can be written

$$\begin{aligned}
\mu(\alpha) &= \alpha_0 r_0 + \alpha_1 \bar{r}_1 + \cdots + \alpha_N \bar{r}_N \\
&= r_0 - r_0(\alpha_0 + \alpha_1 + \cdots + \alpha_N) + \alpha_0 r_0 + \alpha_1 \bar{r}_1 + \cdots + \alpha_N \bar{r}_N \\
&= r_0 + \alpha_1(\bar{r}_1 - r_0) + \cdots + \alpha_N(\bar{r}_N - r_0) \\
&= r_0 + \alpha^T \bar{Y}.
\end{aligned}$$

The variance of the portfolio α can be calculated using

$$\sigma^2(\alpha) = \alpha^T \Sigma \alpha.$$

Letting A be a measure of risk aversion, the Markowitz portfolio is chosen by maximizing the expected return of the portfolio minus a penalty for risk,

$$\max_{\alpha}\left[\mu(\alpha)-\frac{A}{2}\sigma^2(\alpha)\right]=\max_{\alpha}\left[r_0+\alpha^T\bar{Y}-\frac{A}{2}\alpha^T\Sigma\alpha\right].$$

Now, setting the gradient with respect to α equal to zero, we obtain

$$\begin{aligned}\bar{Y}-A\Sigma\alpha &= \vec{0},\\ \alpha &= \frac{1}{A}\Sigma^{-1}\bar{Y},\end{aligned} \tag{16}$$

and $\alpha_0 = 1-\alpha_1-\cdots-\alpha_N$.

3.2. Without a Risk-free Asset

Without a risk-free asset, the expected return of the portfolio α over one period can now be written as

$$\begin{aligned}\mu(\alpha) &= \alpha_1\bar{r}_1+\cdots+\alpha_N\bar{r}_N\\ &= \alpha^T\bar{r}.\end{aligned}$$

We now require $\alpha_1+\cdots+\alpha_N = 1$. The equation for the variance of the return of portfolio α is still

$$\sigma^2(\alpha)=\alpha^T\Sigma\alpha.$$

We can now formulate the problem of finding α as the following optimization problem:

$$\max_{\alpha}\left[\mu(\alpha)-\frac{A}{2}\sigma^2(\alpha)\right]$$

subject to $\alpha_1+\cdots+\alpha_N=1$.

3.3. Assets in Electricity Markets

The relationship between the day-ahead and real-time electricity prices at the PJM hubs varies greatly by the hub and the hour of the day. To illustrate this, we plot the cumulative sum of the real-time prices minus the day-ahead prices ($ per MWh) for a fixed location and a fixed hour of the day. In Figures 4 and 5, we show the cumulative sum of these price differences in order to determine whether electricity is more expensive in the day-ahead or real-time market on average. In Figure 4 we see that, for the off-peak hour of 1 a.m., the relationship between the day-ahead and real-time prices vary greatly by the hub, and no general conclusion can be made about whether electricity is more expensive in the day-ahead or real-time market. In Figure 5, for the on-peak hour of 5 p.m., the dynamics are slightly different than the off-peak hour, but the conclusions are similar. Each hub has different dynamics, and the relationship between the day-ahead market and real-time market is complicated. For example, Figure 4 shows that for the New Jersey Hub at 1 a.m., electricity is cheaper in the day-ahead market on average. However, Figure 5 shows that for the New Jersey Hub at 5 p.m., on average, electricity is cheaper in the real-time market.

Next, we examine a single hub and plot the cumulative sum of the real-time price minus the day-ahead price by the hour of day. In Figure 6, we show the cumulative sum of the differences in day-ahead and real-time prices for each hour of the day, showing a systematic discrepancy by hour of day. This is even more evident at the New Jersey Hub shown in Figure 7. From 1 a.m. to 4 p.m., electricity is cheapest in the day-ahead market on average. However, from 5 p.m. to midnight, electricity is cheapest in the real-time market on average.

We will now attempt to model the difference in the day-ahead and real-time electricity price by the hour of the day and the hub. The number of assets, N, is now equal to the number of hub locations multiplied by 24. We will treat each hour of day at each location as a separate asset. Let P^{DA} be an $N \times 1$ vector of the day-ahead prices at each location and each hour of day. Let α by an $N \times 1$ vector where α_j represents the amount of electricity we buy at location and time j on a fixed day, $j = 1, ...N$. If α_j is negative, then we sell electricity in the day-ahead. We assume we start with one dollar and set the cost of buying electricity in the day ahead market to zero (self-financing),

$$\alpha_1 P_1^{DA} + \cdots + \alpha_N P_N^{DA} = \alpha^T P^{DA} = 0.$$

We enter these contracts the day before the electricity is actually transmitted. The next day we can liquidate our positions in the real-time market. Let P^{RT} be the $N \times 1$ vector of the real-time prices at each location and each hour of day. P^{RT} corresponds to P^{DA}, although P^{RT} is still unknown when we choose our portfolio α. We will first assume P^{DA} is known when we choose our portfolio α.

Let r be the $N \times 1$ vector of returns. The return r_j can now be defined as simply $r_j = P_j^{RT} - P_j^{DA}$, $j = 1, ..., N$. Let $\bar{r}$ be the expected value of the return vector (P^{DA} is known but P^{RT} is random). The expected return of our portfolio can now be written as,

$$\mu(\alpha) \quad = \quad \alpha_1 \bar{r}_1 + \cdots + \alpha_N \bar{r}_N. \tag{17}$$

Let Σ be the covariance matrix of the return vector r. The variance of the portfolio is

$$\sigma^2(\alpha) = \alpha^T \Sigma \alpha,$$

and we can choose our portfolio by maximizing the following objective,

$$\max_{\alpha} \left[\mu(\alpha) - \frac{A}{2} \sigma^2(\alpha) \right]$$

subject to $\alpha^T P^{DA} = 0$. Using the definition of $\mu(\alpha)$ in Equation (17), this problem becomes

$$\max_{\alpha} \left[\alpha^T \bar{r} - \frac{A}{2} \alpha^T \Sigma \alpha \right]$$

subject to $\alpha^T P^{DA} = 0$.

4. Numerical Results

We compare our covariance matrix estimation technique described in Section 2.2 with the CAPM covariance matrix estimation in Section 2.1. One of the typical ways to evaluate a portfolio is to look at the Sharpe ratio, the annualized excess return divided by the annualized standard deviation. Additionally, the maximum draw-down gives an idea of how much a portfolio could potentially lose. The portfolio selection methods described in Section 3.1 require covariance matrices. We show the results of both covariance matrix estimation techniques for investing in U.S. equities and electricity in the PJM grid, in terms of how well they choose portfolios. In both cases we ignore transaction fees.

4.1. S&P 100

To see if the errors-in-variables version of CAPM (Equation (10)) has benefits over traditional CAPM (Equation (6)) for estimating covariance matrices, we evaluate their performance in terms of the quality of Markowitz portfolios produced by each method. We use four-week treasury bills as a proxy for the risky-free asset, ignoring transaction costs and allowing shorting.

Table 1 shows statistics of various S&P 100 portfolios over the period of 2004 to 2011. The market portfolio is the S&P 500 index. The equal-weighted portfolio puts equal weight on each of the S&P 100 stocks, re-balancing daily. CAPM refers to a Markowitz portfolio of the S&P 100 stocks re-balanced daily where the covariance matrix is estimated using Equation (6) with a calibration length of 5 days. The Fama-French method refers to a Markowitz portfolio of the S&P 100 stocks re-balanced daily where the covariance matrix is estimated using the Fama-French 3-factor model with a calibration length of 5 days. The errors-in-variables CAPM method refers to a Markowitz portfolio of the S&P 100 stocks re-balanced daily where the covariance matrix is estimated using Equation (10) with a calibration length of 5 days. In this case, the errors-in-variables CAPM Markowitz portfolio has a higher Sharpe ratio than the traditional CAPM Markowitz portfolio.

Portfolio	μ	σ	Sharpe Ratio	Max Drawdown
Riskfree	.019	.001	.000	.000
Market	-.004	.218	-.106	.565
Equal Weight	.038	.233	.080	.539
CAPM	.302	.258	1.099	.249
Fama-French 3-Factor	.140	.264	.460	.341
Errors-in-Variables CAPM	.377	.256	1.399	.341

Table 1: The portfolio statistics of the S&P 100 portfolios. μ is the annualized geometric return, and σ is the annualized standard deviation. Transaction fees are ignored.

To perform a fair comparison of the two covariance matrix estimation techniques, we next vary the risk aversion, A, in Equation (16), and create an efficient frontier. In Figure 8, we plot the efficient frontier for each of the portfolio selection methods. For a fixed value of A, we back-test each of the portfolio selection methods based on CAPM, the Fama-French 3-factor model, and errors-in-variables CAPM. The portfolio called market portfolio with risk-free is a portfolio which only uses the risk-free asset and the market portfolio. Each portfolio is updated daily, and transaction costs are ignored. We then calculate the annualized return and standard deviation of each portfolio. We then repeat this for many values of A in order to plot an efficient frontier. Figure 8 shows that the portfolio with errors-in-variables CAPM appears to outperform the portfolio with CAPM for each calibration length. For the calibration lengths 20 and 100 days, the portfolio using the Fama-French covariance matrix performs best. However, for the short calibration of length 5 days and the long calibration of length 500 days, the portfolio using the errors-in-variables CAPM has the highest Sharpe ratio. The highest Sharpe ratios are obtained when calibrating over 5 days, likely due to the fact that financial markets are non-stationary.

4.2. PJM

Next, we compare both covariance matrix estimation techniques, CAPM and errors-in-variables CAPM, for estimating the covariance matrix of the electricity market returns at each of the PJM hubs as described in Section 3.3. The self-financing constraint, $\alpha_1 P_1^{DA} + \cdots + \alpha_N P_N^{DA} = 0$, described in Section 3.3 requires knowing the day-ahead prices, which are not known when choosing the portfolio. For the purpose of numerical exercise, we assume the day-ahead prices at each location and time are equivalent. This is a reasonable approximation to get a portfolio that is almost self-financing. In Figure 9, for many values of the risk aversion parameter A, we show the performance of the portfolios constructed using CAPM and errors-in-variables CAPM using a calibration length of 5 days. As before, when estimating covariance matrices over only 5 days of data, portfolios using errors-in-variables CAPM produce higher Sharpe ratios than those with CAPM. For calibration lengths of 20, 100, and 500 days, the portfolios had very similar performances for each covariance estimation techniques. In the non-stationary settings of choosing PJM and S&P 100 portfolios, the errors-in-variables CAPM appears to do a better job estimating covariance matrices with a small amount of calibration data, resulting in portfolios with improved Sharpe ratios, ignoring transaction costs.

5. Conclusion

We began with the intent to estimate covariance matrices of the returns on assets in order to choose portfolios and manage risk. Factor models provide a simple and numerically efficient way to estimate full rank covariance matrices.

However, the classical way to estimate factor loadings defines the residuals as the vertical distance between the asset returns and the regression line. A natural extension is to define the residuals as the Euclidean distance between the asset returns and the regression line (total least squares). One downside of using the Euclidean distance to measure the size of the residuals is that it gives equal weights to each dimension, which may not be appropriate. In addition, the regression for each asset can no longer be estimated individually, because the fitted factor values should be the same across assets.

Allowing for errors in the factors (the market returns), we wrote down an errors-in-variables extension to CAPM and made assumptions, which allowed us to estimate the covariance matrix of the asset returns in closed form. We used the method of instrumental variables to calculate the β's and maximum likelihood estimation to estimate the true market returns. We visually showed that the residuals of the assets appeared to be smaller and less correlated across assets, compared with the residuals in the traditional CAPM.

In our numerical work, we showed that Markowitz portfolios constructed with covariance matrices estimated with our errors-in-variables CAPM had higher Sharpe ratios than traditional CAPM. In particular, portfolios constructed using covariance matrices over the very short period of five days benefited greatly by the new covariance estimation technique, ignoring transaction costs. Overall, the new covariance matrix estimation technique appeared to be effective in controlling the variance of portfolios of virtual electricity contracts.

References

Barra, 2000. Predicted beta.
URL `http://www.barra.com/support/library/PredictedBeta.pdf`

Black, F., 1972. Capital market equilibrium with restricted borrowing. The Journal of Business 45 (3), 444–455.

Campbell, J., Lo, A., MacKinlay, A., 1997. The econometrics of financial markets. Vol. 1. princeton University press Princeton, NJ.

Durbin, J., 1954. Errors in variables. Revue de l'Institut international de statistique, 23–32.

Fama, E., French, K., 1992. The cross-section of expected stock returns. Journal of finance, 427–465.

Kendall, M., Stuart, A., 1961. The Advanced Theory of Statistics: Inference and Relationship. Vol. 2. Hafner Publishing Company.

Lintner, J., 1965. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. The review of economics and statistics 47 (1), 13–37.

Markowitz, H. M., March 1952. Portfolio selection. Journal of Finance 7 (1), 77–91.

Markowitz, H. M., 1999. The early history of portfolio theory: 1600-1960. Financial Analysts Journal 55 (4), pp. 5–16.
URL `http://www.jstor.org/stable/4480178`

PJM, 2011. Pjm 101.
URL `http://www.pjm.com/sitecore/content/Globals/Training/Courses/ol-pjm-101`

Sharpe, W., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. The Journal of finance 19 (3), 425–442.

Sharpe, W. F., 1963. A simplified model for portfolio analysis. Management Science 9 (2), pp. 277–293.
URL `http://www.jstor.org/stable/2627407`

Söderström, T., Stoica, P., 1983. Instrumental variable methods for system identification. Vol. 161. Springer-Verlag Berlin.

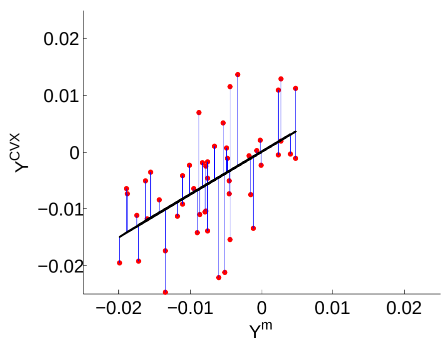

(a)

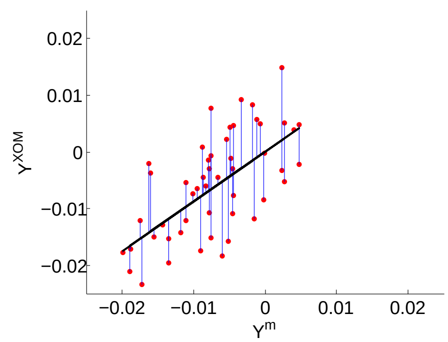

(b)

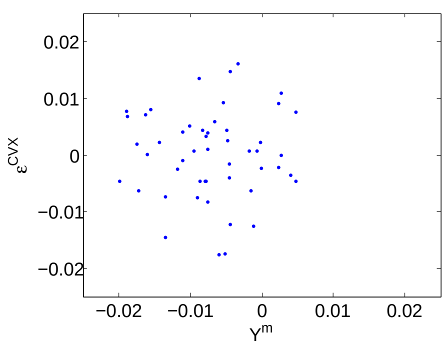

(c)

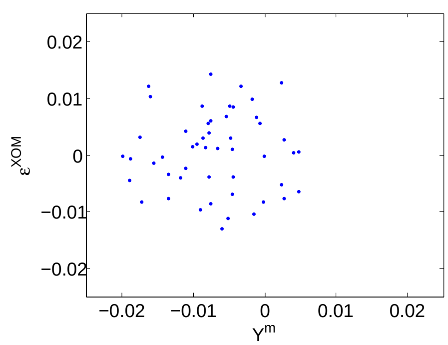

(d)

Figure 1: (a) The CAPM estimate of β for Chevron. (b) The CAPM estimate of β for ExxonMobil (c) The residuals for the fit of Chevron's β. (d) The residuals for the fit of ExxonMobil's β.

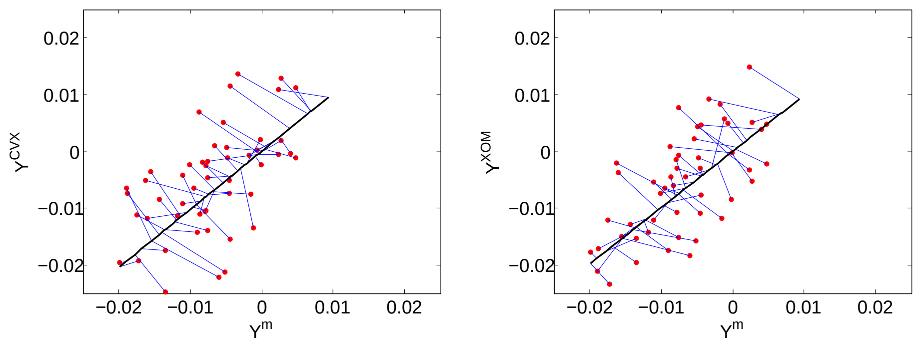

(a) (b)

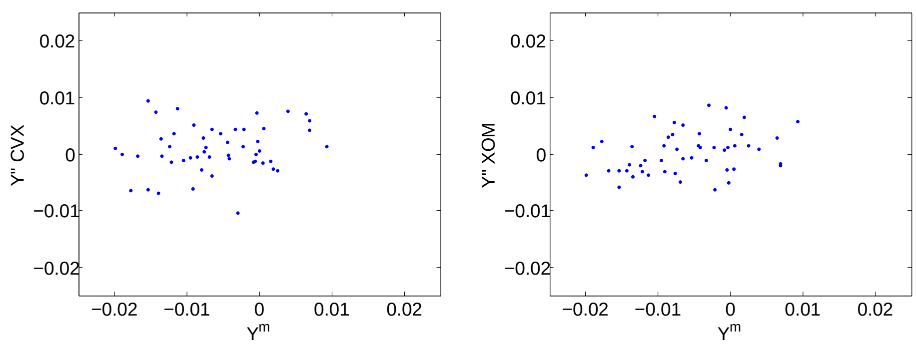

(c) (d)

Figure 2: (a) The error-in-variables CAPM estimate of β for Chevron. (b) The error-in-variables CAPM estimate of β for ExxonMobil. (c) The residuals for the fit of Chevron's β. (d) The residuals for the fit of ExxonMobil's β.

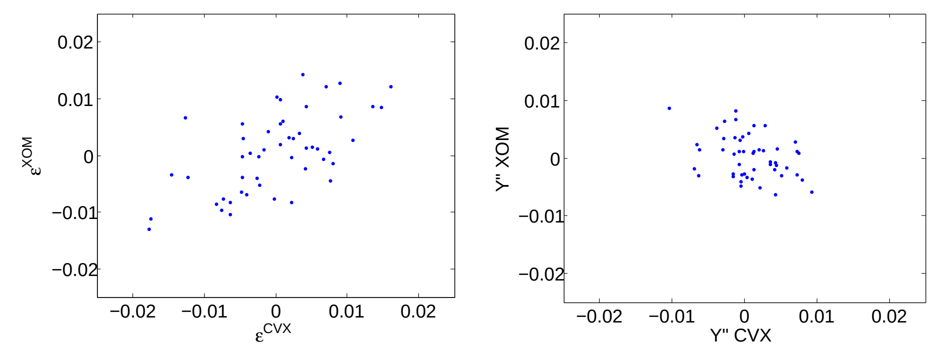

(a) (b)

Figure 3: (a) The residuals between Chevron and ExxonMobil using CAPM as shown in Figures 1(a) and 1(b). (b) The residuals between Chevron and ExxonMobil using the errors-in-variables CAPM as shown in Figures 2(a) and 2(b).

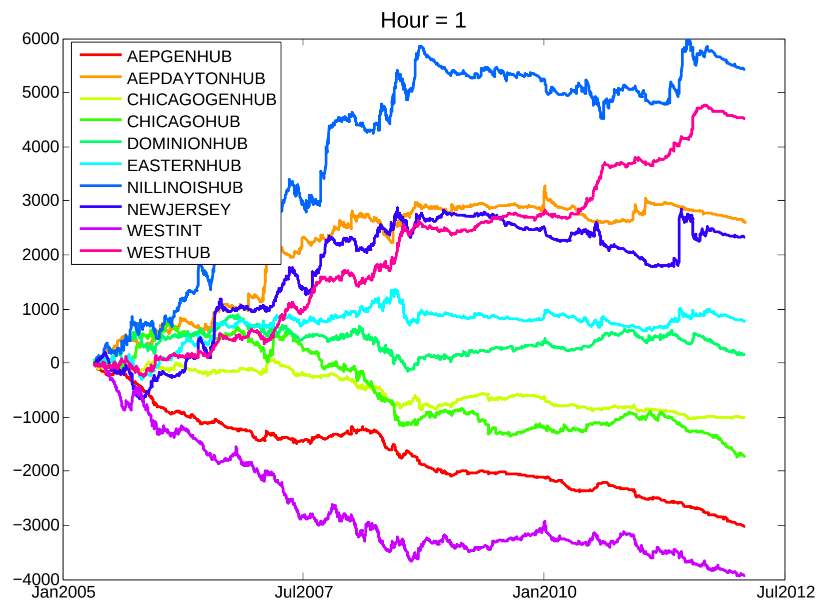

Figure 4: The cumulative sum of the 1 AM real-time price minus the 1 AM day-ahead price for the PJM hubs. When the cumulative sum slopes upwards, the day-ahead price is lower than the real-time price on average. When the cumulative sum slopes downwards, the real-time price is lower than the day-ahead price on average.

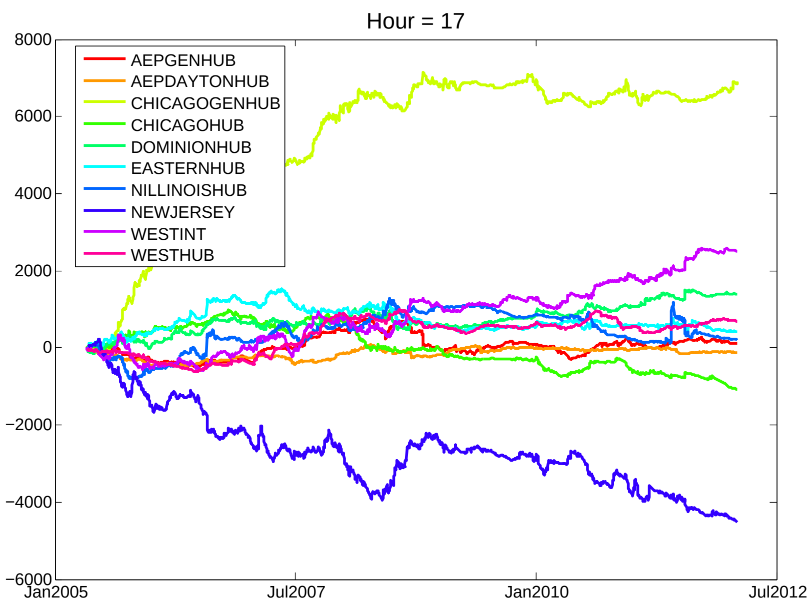

Figure 5: The cumulative sum of the 5 PM real-time price minus the 5 PM day-ahead price for the PJM hubs.

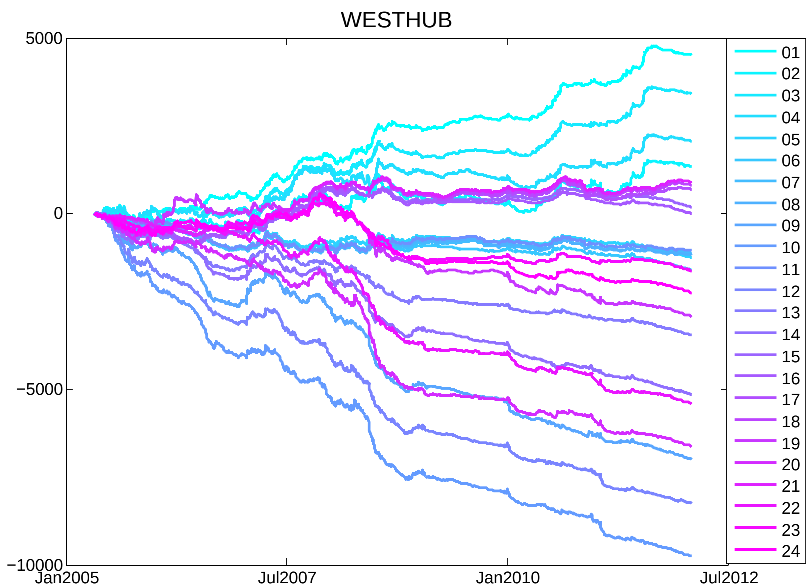

Figure 6: The cumulative sum of the real-time price minus day-ahead price for a particular hour of the day at the Western Hub. Early in the morning, electricity is cheapest in the day-ahead market on average. In the middle of the day, electricity is cheapest in the real-time market on average.

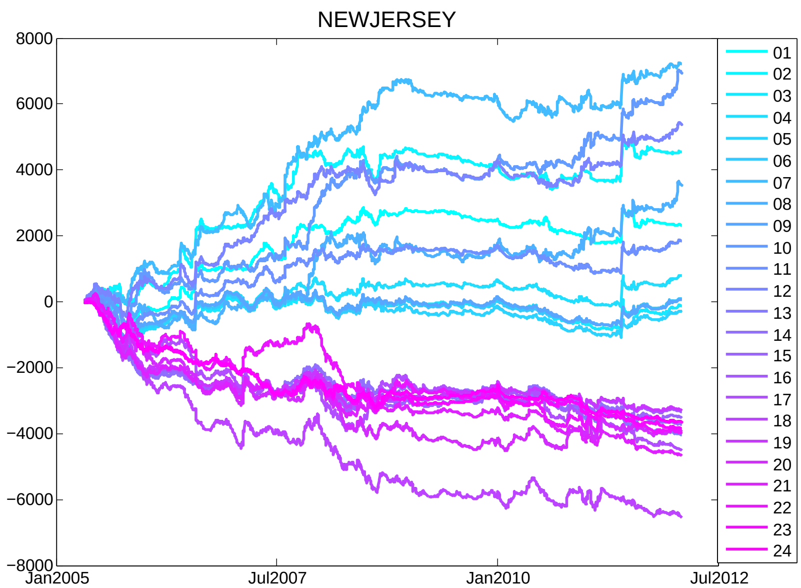

Figure 7: The cumulative sum of the real-time price minus day-ahead price for each hour of the day at the New Jersey Hub.

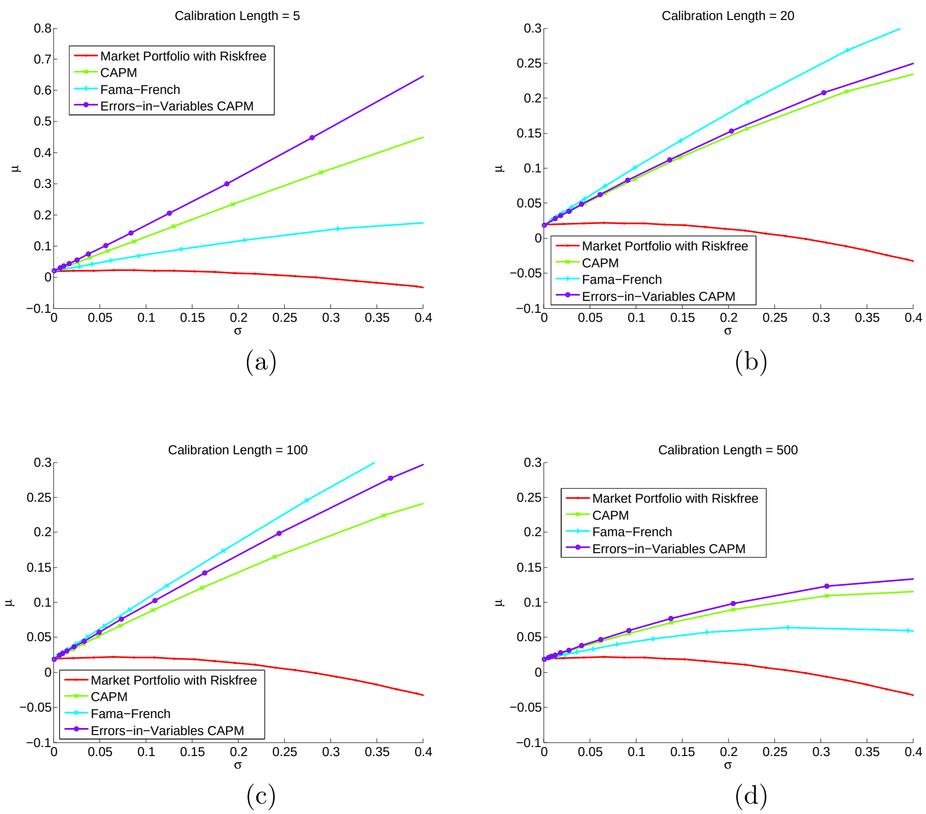

Figure 8: The efficient frontiers of the S&P 100 portfolios. μ is the annualized geometric return (ignoring transaction costs), and σ is the annualized volatility. We show the results using calibration lengths of (a) 5 days. (b) 20 days. (c) 100 days. (d) 500 days. Transaction fees are ignored.

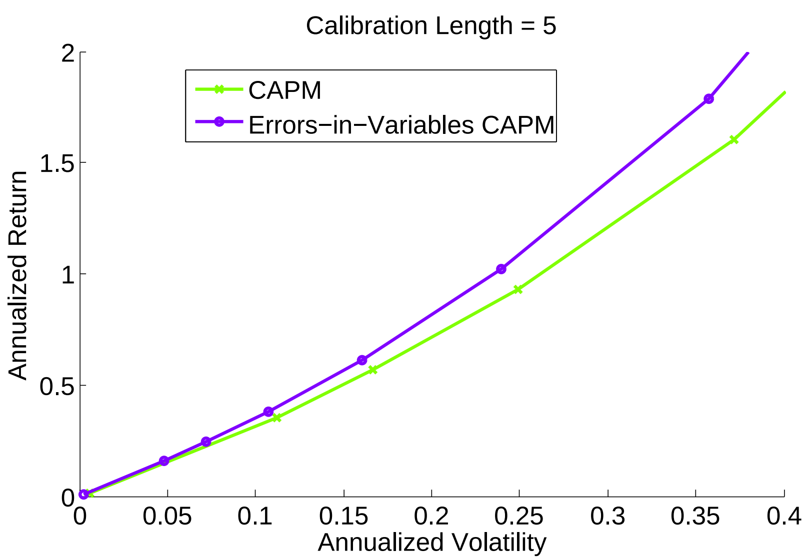

Figure 9: The efficient frontier for the PJM model. Transaction fees are ignored.