Published in final edited form as:

J Comput Graph Stat. 2010 ; 19(4): 947–962. doi:10.1198/jcgs.2010.09188.

Sparse Multivariate Regression With Covariance Estimation

Adam J. Rothman, Ph.D. [Student], Elizaveta Levina [Associate Professor], and Ji Zhu [Associate Professor]

Department of Statistics, University of Michigan, Ann Arbor, MI 48109-1107

Elizaveta Levina: elevina@umich.edu

Abstract

We propose a procedure for constructing a sparse estimator of a multivariate regression coefficient matrix that accounts for correlation of the response variables. This method, which we call multivariate regression with covariance estimation (MRCE), involves penalized likelihood with simultaneous estimation of the regression coefficients and the covariance structure. An efficient optimization algorithm and a fast approximation are developed for computing MRCE. Using simulation studies, we show that the proposed method outperforms relevant competitors when the responses are highly correlated. We also apply the new method to a finance example on predicting asset returns. An R-package containing this dataset and code for computing MRCE and its approximation are available online.

Keywords

High dimension low sample size; Lasso; Multiple output regression; Sparsity

1. Introduction

Multivariate regression is a generalization of the classical regression model of regressing a single response on p predictors to regressing $q > 1$ responses on p predictors. Applications of this general model arise in chemometrics, econometrics, psychometrics, and other quantitative disciplines where one predicts multiple responses with a single set of prediction variables. For example, predicting several measures of quality of paper with a set of variables relating to its production, or predicting asset returns for several companies using the vector autoregressive model (Reinsel 1997), both result in multivariate regression problems.

Let $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^T$ denote the predictors, let $\mathbf{y}_i = (y_{i1}, \dots, y_{iq})^T$ denote the responses, and let $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \dots, \varepsilon_{iq})^T$ denote the errors, all for the i th sample. The multivariate regression model is given by

$$\mathbf{y}_i = \mathbf{B}^T \mathbf{x}_i + \varepsilon_i \text{ for } i=1, \dots, n,$$

where $\mathbf{B}$ is a $p \times q$ regression coefficient matrix and n is the sample size. Column k of $\mathbf{B}$ is the regression coefficient vector from regressing the k th response on the predictors. We make the standard assumption that $\varepsilon_1, \dots, \varepsilon_n$ are iid $N_q(0, \Sigma)$. Thus, given a realization of the predictor variables, the covariance matrix of the response variables is Σ .

The model can be expressed in matrix notation. Let $\mathbf{X}$ denote the $n \times p$ predictor matrix where its i th row is $\mathbf{x}_i^T$, let $\mathbf{Y}$ denote the $n \times q$ random response matrix where its i th row is $\mathbf{y}_i^T$, and let $\mathbf{E}$ denote the $n \times q$ random error matrix where its i th row is ε_i^T ; then the model is

$$\mathbf{Y} = \mathbf{XB} + \mathbf{E}.$$

Note that if $q = 1$, the model simplifies to the classical regression problem where $\mathbf{B}$ is a p -dimensional regression coefficient vector. For simplicity of notation we assume that columns of $\mathbf{X}$ and $\mathbf{Y}$ have been centered and thus the intercept terms are omitted.

The negative log-likelihood function of $(\mathbf{B}, \Omega)$, where $\Omega = \Sigma^{-1}$, can be expressed up to a constant as

$$g(\mathbf{B}, \Omega) = \text{tr} \left[\frac{1}{n} (\mathbf{Y} - \mathbf{XB})^T (\mathbf{Y} - \mathbf{XB}) \Omega \right] - \log |\Omega|. \quad (1.1)$$

The maximum likelihood estimator for $\mathbf{B}$ is simply $\hat{\mathbf{B}}^{\text{OLS}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$, which amounts to performing separate ordinary least squares estimates for each of the q response variables and does not depend on Ω .

Prediction with the multivariate regression model requires one to estimate pq parameters which becomes challenging when there are many predictors and responses. Criterion-based model selection has been extended to multivariate regression by (Bedrick and Tsai 1994) and (Fujikoshi and Satoh 1997). For a review of Bayesian approaches for model selection and prediction with the multivariate regression model see the article by Brown, Vannucci, and Fearn 2002) and references therein. A dimensionality reduction approach called reduced-rank regression (Anderson 1951; Izenman 1975; Reinsel and Velu 1998) minimizes (1.1) subject to $\text{rank}(\mathbf{B}) = r$ for some $r \leq \min(p, q)$. The solution involves canonical correlation analysis, and combines information from all of the q response variables into r canonical response variates that have the highest canonical correlation with the corresponding predictor canonical variates. As in the case of principal components regression, the interpretation of the reduced-rank model is typically impossible in terms of the original predictors and responses.

Other approaches aimed at reducing the number of parameters in the coefficient matrix $\mathbf{B}$ involve solving

$$\hat{\mathbf{B}} = \arg \min_{\mathbf{B}} \text{tr}[(\mathbf{Y} - \mathbf{XB})^T (\mathbf{Y} - \mathbf{XB})] \text{ subject to: } C(\mathbf{B}) \leq \mathbf{t}, \quad (1.2)$$

where $C(\mathbf{B})$ is some constraint function. A method called factor estimation and selection (FES) was proposed by (Yuan et al. 2007), who applied the constraint function

$C(\mathbf{B}) = \sum_{j=1}^{\min(p,q)} \sigma_j(\mathbf{B})$, where $\sigma_j(\mathbf{B})$ is the j th singular value of $\mathbf{B}$. This constraint encourages sparsity in the singular values of $\mathbf{B}$, and hence reduces the rank of $\mathbf{B}$; however, unlike reduced-rank regression, FES offers a continuous regularization path. A novel approach for imposing sparsity in the entries of $\mathbf{B}$ was taken by (Turlach, Venables, and Wright 2005), who proposed the constraint function, $C(\mathbf{B}) = \sum_{j=1}^p \max(|b_{j1}|, \dots, |b_{jq}|)$. This method was recommended for model selection (sparsity identification), and not for prediction because of the bias of the L_∞ -norm penalty. Imposing sparsity in $\mathbf{B}$ for the purposes of identifying “master predictors” was proposed by (Peng et al. 2010), who applied a combined constraint function $C(\mathbf{B}) = \lambda C_1(\mathbf{B}) + (1 - \lambda)C_2(\mathbf{B})$ for $\lambda \in [0, 1]$, where $C_1(\mathbf{B}) = \sum_{j,k} |b_{jk}|$, the lasso constraint (Tibshirani 1996) on the entries of $\mathbf{B}$ and

$C_2(\mathbf{B}) = \sum_{j=1}^p (b_{j1}^2 + \dots + b_{jq}^2)^{0.5}$, the sum of the L_2 -norms of the rows of $\mathbf{B}$ (Yuan and Lin 2006). The first constraint introduces sparsity in the entries of $\mathbf{B}$ and the second constraint introduces zeros for all entries in some rows of $\mathbf{B}$, meaning that some predictors are irrelevant for all q responses. Asymptotic properties for an estimator using this constraint with $\lambda = 0$ have also been established (Obozinski, Wainwright, and Jordan 2008). This combined constraint approach provides highly interpretable models in terms of the prediction variables. However, all of the methods above that solve (1.2) do not account for correlated errors.

To directly exploit the correlation in the response variables to improve prediction performance, a method called Curds and Whey (C&W) was proposed by (Breiman and Friedman 1997). C&W predicts the multivariate response with an optimal linear combination of the ordinary least squares predictors. The C&W linear predictor has the form $\hat{\mathbf{Y}} = \hat{\mathbf{Y}}^{\text{OLS}} \mathbf{M}$, where $\mathbf{M}$ is a $q \times q$ shrinkage matrix estimated from the data. This method exploits correlation in the responses arising from shared random predictors as well as correlated errors.

In this article, we propose a method that combines some of the strengths of the estimators discussed above to improve prediction in the multivariate regression problem while allowing for interpretable models in terms of the predictors. We reduce the number of parameters using the lasso penalty on the entries of $\mathbf{B}$ while accounting for correlated errors. We accomplish this by simultaneously optimizing (1.1) with penalties on the entries of $\mathbf{B}$ and $\mathbf{\Omega}$. We call our new method multivariate regression with covariance estimation (MRCE). The method assumes predictors are not random; however, the resulting formulas for the estimates would be the same with random predictors. Our focus is on the conditional distribution of $\mathbf{Y}$ given $\mathbf{X}$ and thus, unlike in the Curds and Whey framework, the correlation of the response variables arises only from the correlation in the errors.

We also note that the use of the lasso penalty on the entries of Ω has been considered by several authors in the context of covariance estimation (Yuan and Lin 2007; d'Aspremont, Banerjee, and El Ghaoui 2008; Friedman, Hastie, and Tibshirani 2008; Rothman et al. 2008). However, here we use it in the context of a regression problem, thus making it an example of what one could call *supervised* covariance estimation: the covariance matrix here is estimated in order to improve prediction, rather than as a stand-alone parameter. This is a natural next step from the extensive covariance estimation literature, which has been given surprisingly little attention to date; one exception is the joint regression approach of Witten and (Tibshirani 2009). Another less directly relevant example of such supervised estimation is the supervised principal components by (Bair et al. 2006).

The remainder of the article is organized as follows: Section 2 describes the MRCE method and associated computational algorithms, Section 3 presents simulation studies comparing MRCE to competing methods, Section 4 presents an application of MRCE for predicting asset returns, and Section 5 concludes with a summary and discussion.

2. Joint Estimation of $\mathbf{B}$ and Ω via Penalized Normal Likelihood

2.1 The MRCE Method

We propose a sparse estimator for $\mathbf{B}$ that accounts for correlated errors using penalized normal likelihood. We add two penalties to the negative log-likelihood function g to construct a sparse estimator of $\mathbf{B}$ depending on $\Omega = [\omega_{j'j}]$,

$$(\hat{\mathbf{B}}, \hat{\Omega}) = \arg \min_{\mathbf{B}, \Omega} \left\{ g(\mathbf{B}, \Omega) + \lambda_1 \sum_{j' \neq j} |\omega_{j'j}| + \lambda_2 \sum_{j=1}^p \sum_{k=1}^q |b_{jk}| \right\}, \quad (2.1)$$

where $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ are tuning parameters.

We selected the lasso penalty on the off-diagonal entries of the inverse error covariance Ω for two reasons. First, it ensures that an optimal solution for Ω has finite objective function value when there are more responses than samples ($q > n$); second, the penalty has the effect of reducing the number of parameters in the inverse error covariance, which is useful when q is large (Rothman et al. 2008). Other penalties such as the ridge penalty could be used when it is unreasonable to assume that the inverse error covariance matrix is sparse. If q is large, estimating a dense Ω means that the MRCE regression method has $O(q^2)$ additional parameters in Ω to estimate compared with doing separate lasso regressions for each response variable. Thus estimating a sparse Ω has considerably lower variability, and so we focus on the lasso penalty on Ω . We show in simulations that when the inverse error covariance matrix is not sparse, the lasso penalty on Ω still considerably outperforms ignoring covariance estimation altogether (i.e., doing a separate lasso regression for each response).

The lasso penalty on $\mathbf{B}$ introduces sparsity in $\hat{\mathbf{B}}$, which reduces the number of parameters in the model and provides interpretation. In classical regression ($q = 1$), the lasso penalty can offer major improvement in prediction performance when there is a relatively small number

of relevant predictors. This penalty also ensures that an optimal solution for $\mathbf{B}$ is a function of Ω . Without a penalty on $\mathbf{B}$ (i.e., $\lambda_2 = 0$), the optimal solution for $\mathbf{B}$ is always $\mathbf{B}^{\text{OLS}}$.

To see the effect of including the error covariance when estimating an L_1 -penalized $\mathbf{B}$, assume that we know Ω and also assume $p < n$. Solving (2.1) for $\mathbf{B}$ with Ω fixed is a convex problem (see Section 2.2) and thus there exists a global minimizer $\mathbf{B}^{\text{opt}}$. This implies that there exists a zero subgradient of the objective function at $\mathbf{B}^{\text{opt}}$ (see theorem 3.4.3, p. 127, in Bazaraa, Sherali, and Shetty 2006). We express this in matrix notation as

$$0 = 2n^{-1} \mathbf{X}^T \mathbf{X} \mathbf{B}^{\text{opt}} \Omega - 2n^{-1} \mathbf{X}^T \mathbf{Y} \Omega + \lambda_2 \Gamma,$$

which gives

$$\mathbf{B}^{\text{opt}} = \hat{\mathbf{B}}^{\text{OLS}} - \lambda_2 (2n^{-1} \mathbf{X}^T \mathbf{X})^{-1} \Gamma \Omega^{-1}, \quad (2.2)$$

where $\Gamma = \Gamma(\mathbf{B}^{\text{opt}})$ is a $p \times q$ matrix with entries $\gamma_{ij} = \text{sign}(b_{ij}^{\text{opt}})$ if $b_{ij}^{\text{opt}} \neq 0$ and otherwise $\gamma_{ij} \in [-1, 1]$ with specific values chosen to solve (2.2). Ignoring the correlation in the error is equivalent to assuming that $\Omega^{-1} = \mathbf{I}$. Thus having highly correlated errors will have greater influence on the amount of shrinkage of each entry of $\mathbf{B}^{\text{opt}}$ than having mildly correlated errors.

2.2 Computational Algorithms

The optimization problem in (2.1) is not convex; however, solving for either $\mathbf{B}$ or Ω with the other fixed is convex. We present an algorithm for solving (2.1) and a fast approximation to it.

Solving (2.1) for Ω with $\mathbf{B}$ fixed at a chosen point $\mathbf{B}_0$ yields the optimization problem

$$\hat{\Omega}(\mathbf{B}_0) = \arg \min_{\Omega} \left\{ \text{tr} \left(\sum_R \Omega \right) - \log |\Omega| + \lambda_1 \sum_{j' \neq j} |\omega_{j'j}| \right\}, \quad (2.3)$$

where $\sum_R = \frac{1}{n} (\mathbf{Y} - \mathbf{X} \mathbf{B}_0)^T (\mathbf{Y} - \mathbf{X} \mathbf{B}_0)$. This is exactly the L_1 -penalized covariance estimation problem considered by (d'Aspremont, Banerjee, and El Ghaoui 2008), (Yuan and Lin 2007), (Rothman et al. 2008), (Friedman, Hastie, and Tibshirani 2008), and Lu (2009, 2010). We use the graphical lasso (glasso) algorithm of (Friedman, Hastie, and Tibshirani 2008) to solve (2.3) since it is fast and the most commonly used algorithm for solving (2.3).

Solving (2.1) for $\mathbf{B}$ with Ω fixed at a chosen point Ω_0 yields the optimization problem

$$\hat{\mathbf{B}}(\Omega_0) = \arg \min_{\mathbf{B}} \left\{ \text{tr} \left[\frac{1}{n} (\mathbf{Y} - \mathbf{XB})^T (\mathbf{Y} - \mathbf{XB}) \Omega_0 \right] + \lambda_2 \sum_{j=1}^p \sum_{k=1}^q |b_{jk}| \right\}, \quad (2.4)$$

which is convex if Ω_0 is nonnegative definite. This follows because the trace term in the objective function has the Hessian $2n^{-1}\Omega_0 \otimes \mathbf{X}^T\mathbf{X}$, which is nonnegative definite because the Kronecker product of two symmetric nonnegative definite matrices is also nonnegative definite. A solution can be efficiently computed using cyclical-coordinate descent analogous to that used for solving the single output lasso problem (Friedman, Hastie, and Tibshirani 2007). We summarize the optimization procedure in Algorithm 1. We use the ridge penalized least squares estimate $\hat{\mathbf{B}}^{\text{RIDGE}} = (\mathbf{X}^T\mathbf{X} + \lambda_2\mathbf{I})^{-1}\mathbf{X}^T\mathbf{Y}$ to scale our test of parameter convergence since it is always well defined (including when $p > n$).

Algorithm 1: Given Ω and an initial value $\mathbf{B}^{(0)}$, let $\mathbf{S} = \mathbf{X}^T\mathbf{X}$ and $\mathbf{H} = \mathbf{X}^T\mathbf{Y}\Omega$.

Step 1: Set $\hat{\mathbf{B}}^{(m)} \leftarrow \hat{\mathbf{B}}^{(m-1)}$. Visit all entries of $\hat{\mathbf{B}}^{(m)}$ in some sequence and for entry (r, c) update $\hat{b}_{rc}^{(m)}$ with the minimizer of the objective function along its coordinate direction given by

$$\hat{b}_{rc}^{(m)} \leftarrow \text{sign} \left(\hat{b}_{rc}^{(m)} + \frac{h_{rc} - u_{rc}}{s_{rr}\omega_{cc}} \right) \left(\left| \hat{b}_{rc}^{(m)} + \frac{h_{rc} - u_{rc}}{s_{rr}\omega_{cc}} \right| - \frac{n\lambda_2}{s_{rr}\omega_{cc}} \right)_+, \text{ where } u_{rc} = \sum_{j=1}^p \sum_{k=1}^q \hat{b}_{jk}^{(m)} s_{rj}\omega_{kc}.$$

Step 2: If $\sum_{j,k} |\hat{b}_{jk}^{(m)} - \hat{b}_{jk}^{(m-1)}| < \varepsilon \sum_{j,k} |\hat{b}_{jk}^{\text{RIDGE}}|$ then stop, otherwise go to Step 1.

A full derivation of Algorithm 1 is found in the Appendix. Algorithm 1 is guaranteed to converge to the global minimizer if the given Ω is nonnegative definite. This follows from the fact that the trace term in the objective function is convex and differentiable and the penalty term decomposes into a sum of convex functions of individual parameters (Tseng 1988; Friedman, Hastie, and Tibshirani 2007). We set the convergence tolerance parameter $\varepsilon = 10^{-4}$.

In terms of computational cost, we need to cycle through pq parameters, and for each compute u_{rc} , which costs at most $O(pq)$ flops, and if the least sparse iterate has v non-zeros, then computing u_{rc} costs $O(v)$. The worst case cost for the entire algorithm is $O(p^2q^2)$.

Using (2.3) and (2.4) we can solve (2.1) using blockwise coordinate descent, that is, we iterate minimizing with respect to $\mathbf{B}$ and minimizing with respect to Ω .

Algorithm 2 (MRCE): For fixed values of λ_1 and λ_2 , initialize $\mathbf{B}^{(0)} = \mathbf{0}$ and $\Omega^{(0)} = \Omega(\hat{\mathbf{B}}^{(0)})$.

Step 1: Compute $\hat{\mathbf{B}}^{(m+1)} = \hat{\mathbf{B}}(\hat{\Omega}^{(m)})$ by solving (2.4) using Algorithm 1.

Step 2: Compute $\Omega^{(m+1)} = \Omega(\hat{\mathbf{B}}^{(m+1)})$ by solving (2.3) using the glasso algorithm.

Step 3: If $\sum_{j,k} |\hat{b}_{jk}^{(m+1)} - \hat{b}_{jk}^{(m)}| < \varepsilon \sum_{j,k} |\hat{b}_{jk}^{\text{RIDGE}}|$ then stop, otherwise go to Step 1.

Algorithm 2 uses blockwise coordinate descent to compute a local solution for (2.1). Steps 1 and 2 both ensure a decrease in the objective function value. In practice we found that for certain values of the penalty tuning parameters (λ_1, λ_2) , the algorithm may take many iterations to converge for high-dimensional data. For such cases, we propose a faster approximate solution to (2.1).

Algorithm 3 (Approximate MRCE): For fixed values of λ_1 and λ_2 ,

Step 1: Perform q separate lasso regressions each with the same optimal tuning parameter $\hat{\lambda}_0$ selected with a cross-validation procedure. Let $\hat{\mathbf{B}}_{\hat{\lambda}_0}^{\text{lasso}}$ denote the solution.

Step 2: Compute $\hat{\mathbf{\Omega}} = \hat{\mathbf{\Omega}}(\hat{\mathbf{B}}_{\hat{\lambda}_0}^{\text{lasso}})$ by solving (2.3) using the glasso algorithm.

Step 3: Compute $\hat{\mathbf{B}} = \mathbf{B}(\hat{\mathbf{\Omega}})$ by solving (2.4) using Algorithm 1.

The approximation summarized in Algorithm 3 is only iterative inside its steps. The algorithm begins by finding the optimally tuned lasso solution $\hat{\mathbf{B}}_{\hat{\lambda}_0}^{\text{lasso}}$ (using cross-validation to select the tuning parameter $\hat{\lambda}_0$), then computes an estimate for $\mathbf{\Omega}$ using the glasso algorithm with $\hat{\mathbf{B}}_{\hat{\lambda}_0}^{\text{lasso}}$ plugged in, and then solves (2.4) using this inverse covariance estimate. Note that one still must select two tuning parameters (λ_1, λ_2) . The performance of the approximation is studied in Section 3.

2.3 Tuning Parameter Selection

For the MRCE methods, the tuning parameters λ_1 and λ_2 could be selected using K -fold cross-validation, where validation prediction error is accumulated over all q responses for each fold. Specifically, select the optimal tuning parameters $\hat{\lambda}_1$ and $\hat{\lambda}_2$ using

$$(\hat{\lambda}_1, \hat{\lambda}_2) = \arg \min_{\lambda_1, \lambda_2} \sum_{k=1}^K \|\mathbf{Y}^{(k)} - \mathbf{X}^{(k)} \mathbf{B}_{\lambda_1, \lambda_2}^{(-k)}\|_F^2,$$

where $\mathbf{Y}^{(k)}$ is the matrix of responses with observations in the k th fold, $\mathbf{X}^{(k)}$ is the matrix of predictors of observations in the k th fold, and $\mathbf{B}_{\lambda_1, \lambda_2}^{(-k)}$ is the estimated regression coefficient matrix computed with observations outside the k th fold, with tuning parameters λ_1 and λ_2 . We have found in simulations that λ_2 , which controls the penalization on the regression coefficient matrix, has greater influence on prediction performance than λ_1 , which controls the penalization of the inverse error covariance matrix.

3. Simulation Study

3.1 Estimators

We compare the performance of the MRCE method, computed with the exact and the approximate algorithms, to other multivariate regression estimators that produce sparse estimates of $\mathbf{B}$. We report results for the following methods:

- *Lasso*: Perform q separate lasso regressions, each with the same tuning parameter λ .
- *Separate lasso*: Perform q separate lasso regressions, each with its own tuning parameter.
- *MRCE*: The solution to (2.1) (Algorithm 2).
- *Approx. MRCE*: An approximate solution to (2.1) (Algorithm 3).

The ordinary least squares estimator $\mathbf{B}^{\hat{\text{OLS}}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$ and the Curds and Whey method of (Breiman and Friedman 1997) are computed as a benchmark for low-dimensional models (they are not directly applicable when $p > n$).

We select the tuning parameters by minimizing the squared prediction error, accumulated over all q responses, of independently generated validation data of the same sample size ($n = 50$). This is similar to performing the cross-validation approach described in Section 2.3, and is used to save computing time for the simulations. For the MRCE methods, the two tuning parameters are selected simultaneously.

3.2 Models

In each replication for each model, we generate an $n \times p$ predictor matrix $\mathbf{X}$ with rows drawn independently from $N_p(0, \Sigma_X)$ where $\Sigma_X = [\sigma_{Xij}]$ is given by $\sigma_{Xij} = 0.7^{|i-j|}$. This model for the predictors was also used by (Yuan et al. 2007) and (Peng et al. 2010). Note that all of the predictors are generated with the same unit marginal variance. The error matrix $\mathbf{E}$ is generated independently with rows drawn independently from $N_q(0, \Sigma_E)$. We consider two models for the error covariance:

- AR(1) error covariance: $\sigma_{Eij} = \rho_E^{|i-j|}$, with values of ρ_E ranging from 0 to 0.9.
- Fractional Gaussian Noise (FGN) error covariance:

$$\sigma_{Eij} = 0.5 \left((|i-j|+1)^{2H} - 2|i-j|^{2H} + (|i-j|-1)^{2H} \right),$$

with values of the Hurst parameter $H = 0.9, 0.95$.

The inverse error covariance for the AR(1) model is a tri-diagonal sparse matrix while its covariance matrix is dense, and thus this error covariance model completely satisfies the regularizing assumptions for the MRCE method, which exploits the correlated error and the sparse inverse error covariance. The FGN model is a standard example of long-range dependence and both the error covariance and its inverse are dense matrices. Varying H gives different degree of dependence, from $H = 0.5$ corresponding to an iid sequence to $H = 1$ corresponding to a perfectly correlated one. Thus the introduction of sparsity in the inverse error covariance by the MRCE method should not help; however, since the errors are highly correlated the MRCE method may still perform better than the lasso penalized regressions for each response, which ignore correlation among the errors. The sample size is fixed at $n = 50$ for all models.

We generate sparse coefficient matrices $\mathbf{B}$ in each replication using the matrix element-wise product,

$$\mathbf{B} = \mathbf{W} * \mathbf{K} * \mathbf{Q},$$

where $\mathbf{W}$ is generated with independent draws for each entry from $N(0, 1)$, $\mathbf{K}$ has entries with independent Bernoulli draws with success probability s_1 , and $\mathbf{Q}$ has rows that are either all one or all zero, where p independent Bernoulli draws with success probability s_2 are made to determine whether each row is the ones vector or the zeros vector. Generating $\mathbf{B}$ in this manner, we expect $(1 - s_2)p$ predictors to be irrelevant for all q responses, and we expect each relevant predictor to be relevant for $s_1 q$ of the response variables.

3.3 Performance Evaluation

We measure performance using model error, following the approach of (Yuan et al. 2007), which is defined as

$$\text{ME}(\hat{\mathbf{B}}, \mathbf{B}) = \text{tr}[(\hat{\mathbf{B}} - \mathbf{B})^T \sum_{\mathbf{x}} (\hat{\mathbf{B}} - \mathbf{B})].$$

We also measure the sparsity recognition performance using true positive rate (TPR) and true negative rate (TNR),

$$\text{TPR}(\hat{\mathbf{B}}, \mathbf{B}) = \frac{\#\{(i, j): \hat{b}_{ij} \neq 0 \text{ and } b_{ij} \neq 0\}}{\#\{(i, j): b_{ij} \neq 0\}}, \quad (3.1)$$

$$\text{TNR}(\hat{\mathbf{B}}, \mathbf{B}) = \frac{\#\{(i, j): \hat{b}_{ij} = 0 \text{ and } b_{ij} = 0\}}{\#\{(i, j): b_{ij} = 0\}}, \quad (3.2)$$

Both the true positive rate and true negative rates must be considered simultaneously since OLS always has perfect TPR and $\hat{\mathbf{B}} = \mathbf{0}$ always has perfect TNR.

3.4 Results

The model error performance for AR(1) error covariance model is displayed in Figure 1 for low-dimensional models, and in Figure 2 and Table 1 for high-dimensional models. Standard errors are omitted in the figures because of visibility issues, and we note that they are less than 4% of the corresponding average model error. We see that the margin by which MRCE and its approximation outperform the lasso and separate lasso in terms of model error increases as the error correlation ρ_E increases. This trend is consistent with the analysis of the subgradient equation given in (2.2), since the manner by which MRCE performs lasso shrinkage exploits highly correlated errors. Additionally, the MRCE method and its approximation outperform the lasso and separate lasso more for sparser coefficient matrices. We omitted the exact MRCE method for $p = 60$, $q = 20$, and $p = q = 100$ because these cases

were computationally intractable. For a single realization of the model with $p = 20$, $q = 60$, and $\rho_E = 0.9$, using the tuning parameters selected with cross-validation, MRCE took 4.1 seconds, approximate MRCE took 1.7 seconds, lasso took 0.5 seconds, and separate lasso took 0.4 seconds to compute on a workstation with a 2 GHz processor with 4 GB of RAM. All of the sparse estimators outperform the ordinary least squares method by a considerable margin. The Curds and Whey method, although designed to exploit correlation in the responses, is outperformed here because it does not introduce sparsity in $\mathbf{B}$.

The model error performance for FGN error covariance model is reported in Table 2 for low-dimensional models and in Table 3 for high-dimensional models. Although there is no sparsity in the inverse error covariance for the MRCE method and its approximation to exploit, we see that both methods are still able to provide considerable improvement over the lasso and separate lasso methods by exploiting the highly correlated error. As seen with the AR(1) error covariance model, as the amount of correlation increases (i.e., larger values of H), the margin by which the MRCE method and its approximation outperform competitors increases.

We report the true positive rate and true negative rates in Table 4 for the AR(1) error covariance models and in Table 5 for the FGN error covariance models. We see that as the error correlation increases (larger values of ρ_E and H), the true positive rate for the MRCE method and its approximation increases, while the true negative rate tends to decrease. While all methods perform comparably on these sparsity measures, the substantially lower prediction errors obtained by the MRCE methods give them a clear advantage over other methods.

4. Example: Predicting Asset Returns

We consider a dataset of weekly log-returns of nine stocks from 2004, analyzed by (Yuan et al. 2007). We selected this dataset because it is the most recent dataset analyzed in the multivariate regression literature. The data are modeled with a first-order vector autoregressive model,

$$\mathbf{Y} = \tilde{\mathbf{Y}}\mathbf{B} + \mathbf{E},$$

where the response $\mathbf{Y} \in \mathbb{R}^{T-1 \times q}$ has rows $\mathbf{y}_2, \dots, \mathbf{y}_T$ and the predictor $\tilde{\mathbf{Y}} \in \mathbb{R}^{T-1 \times q}$ has rows $\mathbf{y}_1, \dots, \mathbf{y}_{T-1}$. Here $\mathbf{y}_t$ corresponds to the vector of log-returns for the nine companies at week t . Let $\mathbf{B} \in \mathbb{R}^{q \times q}$ denote the transition matrix. Following the approach of (Yuan et al. 2007), we use log-returns from the first 26 weeks of the year ($T = 26$) as the training set, and the log-returns from the remaining 26 weeks of the year as the test set. Prediction performance is measured by the average mean squared prediction error over the test set for each stock, with the model fitted using the training set. Tuning parameters were selected with 10-fold CV.

Average test squared error over the 26 test points is reported in Table 6, where we see that the MRCE method and its approximation have somewhat better performance than the lasso

and separate lasso methods. The lasso estimate of the transition matrix $\mathbf{B}$ was all zeros, yielding the null model. Nonetheless, this results in prediction performance comparable (i.e., within a standard error) to the FES method of (Yuan et al. 2007) (copied directly from table 3 on page 341), which was shown to be the best of several competitors for these data. This comparable performance of the null model suggests that the signal is very weak in this dataset. Separate lasso, MRCE, and its approximation estimated 3/81, 4/81, and 12/81 coefficients as nonzero, respectively.

We report the estimate of the unit lag coefficient matrix $\mathbf{B}$ for the approximate MRCE method in Table 7, which is the least sparse estimate, identifying 12 nonzero entries. The estimated unit lag coefficient matrix for separate lasso, MRCE, and approximate MRCE all identified the log-return for Walmart at week $t - 1$ as a relevant predictor for the log-return of GE at week t , and the log-return for Ford at week $t - 1$ as a relevant predictor for the log-return of Walmart at week t . The FES does not provide any interpretation.

We also report the estimate for the inverse error covariance matrix for the MRCE method in Table 8. A nonzero entry (i, j) means that we estimate that ε_i is correlated with ε_j given the other errors (or ε_i is partially correlated with ε_j). We see that AIG (an insurance company) is estimated to be partially correlated with most of the other companies, and companies with similar products are partially correlated, such as Ford and GM (automotive), GE and IBM (technology), as well as Conoco Phillips and Exxon (oil). These results make sense in the context of financial data.

5. Summary and Discussion

We proposed the MRCE method to produce a sparse estimate of the multivariate regression coefficient matrix $\mathbf{B}$. Our method explicitly accounts for the correlation of the response variables. We also developed a fast approximate algorithm for computing MRCE which has roughly the same performance in terms of model error. These methods were shown to outperform q separate lasso penalized regressions (which ignore the correlation in the responses) in simulations when the responses are highly correlated, even when the inverse error covariance is dense.

Although we considered simultaneous L_1 -penalization of $\mathbf{B}$ and $\mathbf{\Omega}$, one could use other penalties that introduce less bias instead, such as SCAD (Fan and Li 2001; Lam and Fan 2009). In addition, this work could be extended to the situation when the response vector samples have serial correlation, in which case the model would involve both the error covariance and the correlation among the samples.

Acknowledgments

We thank Ming Yuan for providing the weekly log-returns dataset. We also thank the associate editor and two referees for their helpful suggestions. This research has been supported in part by the Yahoo Ph.D. student fellowship (A. J. Rothman) and National Science Foundation grants DMS-0805798 (E. Levina), DMS-0705532 and DMS-0748389 (J. Zhu).

Appendix: Derivation of Algorithm 1

The objective function for Ω fixed at Ω_0 is now

$$f(\mathbf{B}) = g(\mathbf{B}, \Omega_0) + \lambda_2 \sum_{j=1}^p \sum_{k=1}^q |b_{jk}|.$$

We can solve for $\mathbf{B}$ with cyclical coordinate descent. Express the directional derivatives as

$$\begin{aligned} \frac{\partial f^+}{\partial \mathbf{B}} &= \frac{2}{n} \mathbf{X}^T \mathbf{X} \mathbf{B} \Omega - \frac{2}{n} \mathbf{X}^T \mathbf{Y} \Omega + \lambda_2 \mathbf{1}(b_{ij} \geq 0) - \lambda_2 \mathbf{1}(b_{ij} < 0), \\ \frac{\partial f^-}{\partial \mathbf{B}} &= -\frac{2}{n} \mathbf{X}^T \mathbf{X} \mathbf{B} \Omega + \frac{2}{n} \mathbf{X}^T \mathbf{Y} \Omega - \lambda_2 \mathbf{1}(b_{ij} > 0) + \lambda_2 \mathbf{1}(b_{ij} \leq 0), \end{aligned}$$

where the indicator $\mathbf{1}(\cdot)$ is understood to be a matrix. Let $\mathbf{S} = \mathbf{X}^T \mathbf{X}$ and $\mathbf{H} = \mathbf{X}^T \mathbf{Y} \Omega$ and

$u_{rc} = \sum_{j=1}^p \sum_{k=1}^q b_{jk} s_{rj} \omega_{kc}$. To update a single parameter b_{rc} we have the directional derivatives,

$$\begin{aligned} \frac{\partial f^+}{\partial b_{rc}} &= u_{rc} - h_{rc} + n\lambda_2 \mathbf{1}(b_{ij} \geq 0) - n\lambda_2 \mathbf{1}(b_{ij} < 0), \\ \frac{\partial f^-}{\partial b_{rc}} &= -u_{rc} + h_{rc} - n\lambda_2 \mathbf{1}(b_{ij} > 0) + n\lambda_2 \mathbf{1}(b_{ij} \leq 0). \end{aligned}$$

Let b_{rc}^0 be our current iterate. The unpenalized univariate minimizer $\hat{b}_{rc}^*$ solves

$$\hat{b}_{rc}^* s_{rr} \omega_{cc} - b_{rc}^0 s_{rr} \omega_{cc} + u_{rc} - h_{rc} = 0,$$

implying $\hat{b}_{rc}^* = b_{rc}^0 + \frac{h_{rc} - u_{rc}}{s_{rr} \omega_{cc}}$. If $\hat{b}_{rc}^* > 0$, then we look leftward and by convexity the penalized

minimizer is $\hat{b}_{rc} = \max(0, \hat{b}_{rc}^* - \frac{n\lambda_2}{s_{rr} \omega_{cc}})$. Similarly if $\hat{b}_{rc}^* < 0$, then we look to the right and by

convexity the penalized univariate minimizer is $\hat{b}_{rc} = \min(0, \hat{b}_{rc}^* + \frac{n\lambda_2}{s_{rr} \omega_{cc}})$, thus

$\hat{b}_{rc} = \text{sign}(\hat{b}_{rc}^*) (|\hat{b}_{rc}^*| - \frac{n\lambda_2}{s_{rr} \omega_{cc}})_+$. Also if $\hat{b}_{rc}^* = 0$, which has probability zero, then both the loss and penalty part of the objective function are minimized and the parameter stays at 0. We can write this solution as

$$\hat{b}_{rc} = \text{sign} \left(b_{rc}^0 + \frac{h_{rc} - u_{rc}}{s_{rr} \omega_{cc}} \right) \left(\left| b_{rc}^0 + \frac{h_{rc} - u_{rc}}{s_{rr} \omega_{cc}} \right| - \frac{n\lambda_2}{s_{rr} \omega_{cc}} \right)_+.$$

References

- Anderson T. Estimating Linear Restrictions on Regression Coefficients for Multivariate Normal Distributions. *The Annals of Mathematical Statistics*. 1951; 22:327–351. 948.
- Bair E, Hastie T, Paul D, Tibshirani R. Prediction by Supervised Principal Components. *Journal of the American Statistical Association*. 2006; 101(473):119–137. 949.
- Bazaraa, MS.; Sherali, HD.; Shetty, CM. *Nonlinear Programming: Theory and Algorithms*. 3rd. NJ: Wiley; 2006. 950
- Bedrick E, Tsai C. Model Selection for Multivariate Regression in Small Samples. *Biometrics*. 1994; 50:226–231. 948.
- Breiman L, Friedman JH. Predicting Multivariate Responses in Multiple Linear Regression. *Journal of the Royal Statistical Society, Ser B*. 1997; 59:3–54. with discussion. 949,953.
- Brown P, Vannucci M, Fearn T. Bayes Model Averaging With Selection of Regressors. *Journal of the Royal Statistical Society, Ser B*. 2002; 64:519–536. 948.
- d'Aspremont A, Banerjee O, El Ghaoui L. First-Order Methods for Sparse Covariance Selection. *SIAM Journal on Matrix Analysis and Its Applications*. 2008; 30(1):56–66. 949,951.
- Fan J, Li R. Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties. *Journal of the American Statistical Association*. 2001; 96(456):1348–360. 960.
- Friedman J, Hastie T, Tibshirani R. Pathwise Coordinate Optimization. *The Annals of Applied Statistics*. 2007; 1(2):302–332. 951,952.
- Friedman J, Hastie T, Tibshirani R. Sparse Inverse Covariance Estimation With the Graphical Lasso. *Biostatistics*. 2008; 9(3):432–441. 949,951. [PubMed: 18079126]
- Fujikoshi Y, Satoh K. Modified AIC and C_p in Multivariate Linear Regression. *Biometrika*. 1997; 84:707–716. 948.
- Izenman AJ. Reduced-Rank Regression for the Multivariate Linear Model. *Journal of Multivariate Analysis*. 1975; 5(2):248–264. 948.
- Lam C, Fan J. Sparsistency and Rates of Convergence in Large Covariance Matrices Estimation. *The Annals of Statistics*. 2009; 37:4254–4278. 960.
- Lu Z. Smooth Optimization Approach for Sparse Covariance Selection. *SIAM Journal on Optimization*. 2009; 19(4):1807–1827. 951.
- Lu Z. Adaptive First-Order Methods for General Sparse Inverse Covariance Selection. *SIAM Journal on Matrix Analysis and Applications*. 2010; 31:2000–2016. 951.
- Obozinski, G.; Wainwright, MJ.; Jordan, MI. Technical Report. Vol. 761. UC Berkeley, Dept. of Statistics; 2008. Union Support Recovery in High-Dimensional Multivariate Regression. 949
- Peng J, Zhu J, Bergamaschi A, Han W, Noh DY, Pollack JR, Wang P. Regularized Multivariate Regression for Identifying Master Predictors With Application to Integrative Genomics Study of Breast Cancer. *The Annals of Applied Statistics*. 2010; 4:53–77. 949,953. [PubMed: 24489618]
- Reinsel, G. *Elements of Multivariate Time Series Analysis*. 2nd. New York: Springer; 1997. 947
- Reinsel, G.; Velu, R. *Multivariate Reduced-Rank Regression: Theory and Applications*. New York: Springer; 1998. 948
- Rothman AJ, Bickel PJ, Levina E, Zhu J. Sparse Permutation Invariant Covariance Estimation. *Electronic Journal of Statistics*. 2008; 2:494–515. 949-951.
- Tibshirani R. Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society Ser B*. 1996; 58:267–288. 949.
- Tseng, P. Technical Report LIDS-P 1840. Massachusetts Institute of Technology, Laboratory for Information and Decision Systems; 1988. Coordinate Ascent for Maximizing Nondifferentiable Concave Functions. 952
- Turlach BA, Venables WN, Wright SJ. Simultaneous Variable Selection. *Technometrics*. 2005; 47(3): 349–363. 949.
- Witten DM, Tibshirani R. Covariance-Regularized Regression and Classification for High-Dimensional Problems. *Journal of the Royal Statistical Society Ser B*. 2009; 71(3):615–636. 949.
- Yuan M, Lin Y. Model Selection and Estimation in Regression With Grouped Variables. *Journal of the Royal Statistical Society Ser B*. 2006; 68(1):49–67. 949.

- Yuan M, Lin Y. Model Selection and Estimation in the Gaussian Graphical Model. *Biometrika*. 2007; 94(1):19–35. 949,951.
- Yuan M, Ekici A, Lu Z, Monteiro R. Dimension Reduction and Coefficient Estimation in Multivariate Linear Regression. *Journal of the Royal Statistical Society Ser B*. 2007; 69(3):329–346. 948,953,954,958,959.

Figure 1.

Average model error versus AR(1) correlation ρ_E , based on 50 replications with $n = 50$, $p = q = 20$, and $s_2 = 1$.

Figure 2.

Average model error versus AR(1) correlation ρ_E , based on 50 replications with $n = 50$ and $s_1 = 0.1, s_2 = 1$.

Table 1

Model error for the AR(1) error covariance models of high dimension, with $p = q = 100$, $s_1 = 0.5$, and $s_2 = 0.1$. Averages and standard errors in parentheses are based on 50 replications with $n = 50$.

ρE	lasso	sep.lasso	ap.MRCE
0.9	58.79 (2.29)	59.32 (2.35)	34.87 (1.54)
0.7	59.09 (2.22)	59.60 (2.30)	60.12 (2.02)

Table 2

Model error for the FGN error covariance models of low dimension. Averages and standard errors in parentheses are based on 50 replications with $n = 50$. Tuning parameters were selected using a 10^3 resolution.

p	q	H	s_1, s_2	OLS	lasso	sep.lasso	MRCE	ap.MRCE	C&W
20	20	0.95	0.1, 1	14.51 (0.69)	2.72 (0.10)	2.71 (0.11)	1.03 (0.02)	1.01 (0.03)	9.86 (0.46)
20	20	0.90	0.1, 1	14.49 (0.53)	2.76 (0.09)	2.77 (0.09)	1.78 (0.05)	1.71 (0.05)	10.29 (0.36)
20	20	0.95	0.5, 1	14.51 (0.69)	9.89 (0.26)	8.94 (0.21)	3.63 (0.09)	4.42 (0.16)	11.72 (0.45)
20	20	0.90	0.5, 1	14.49 (0.53)	10.01 (0.21)	9.03 (0.18)	6.11 (0.14)	6.34 (0.13)	12.29 (0.34)

Table 3

Model error for the FGN error covariance models of high dimension. Averages and standard errors in parentheses are based on 50 replications with $n = 50$. Tuning parameters were selected using a 10^x resolution.

p	q	H	s_1, s_2	OLS	lasso	sep.lasso	MRCE	ap.MRCE
20	60	0.95	0.1, 1	46.23 (2.04)	8.56 (0.36)	8.63 (0.37)	3.31 (0.19)	3.20 (0.18)
20	60	0.90	0.1, 1	45.41 (1.42)	8.60 (0.24)	8.69 (0.25)	5.31 (0.15)	5.03 (0.14)
60	20	0.95	0.1, 1	NA	11.15 (0.35)	11.23 (0.36)	–	4.84 (0.12)
60	20	0.90	0.1, 1	NA	11.14 (0.30)	11.21 (0.30)	–	7.44 (0.16)
100	100	0.95	0.5, 0.1	NA	58.28 (2.36)	58.86 (2.44)	–	31.85 (1.26)
100	100	0.95	0.5, 0.1	NA	58.10 (2.27)	58.63 (2.36)	–	47.37 (1.68)

Table 4

True positive rate/true negative rate for the AR(1) error covariance models, averaged over 50 replications; $n = 50$. Standard errors are omitted (the largest standard error is 0.04 and most are less than 0.01). Tuning parameters were selected using a 10^x resolution.

p	q	ρ_E	s_1, s_2	lasso	sep.lasso	MRCE	ap.MRCE
20	20	0.9	0.1, 1	0.83/0.72	0.82/0.74	0.95/0.59	0.94/0.62
20	20	0.7	0.1, 1	0.83/0.71	0.82/0.73	0.89/0.60	0.89/0.63
20	20	0.5	0.1, 1	0.83/0.70	0.81/0.73	0.86/0.62	0.87/0.63
20	20	0	0.1, 1	0.84/0.70	0.82/0.72	0.85/0.63	0.85/0.64
20	20	0.9	0.5, 1	0.86/0.44	0.87/0.44	0.93/0.42	0.91/0.45
20	20	0.7	0.5, 1	0.85/0.47	0.87/0.42	0.86/0.51	0.86/0.52
20	20	0.5	0.5, 1	0.83/0.52	0.87/0.44	0.83/0.54	0.85/0.48
20	20	0	0.5, 1	0.84/0.50	0.87/0.43	0.84/0.51	0.82/0.56
20	60	0.9	0.1, 1	0.83/0.70	0.80/0.74	0.94/0.58	0.93/0.61
20	60	0.7	0.1, 1	0.84/0.71	0.81/0.73	0.89/0.61	0.89/0.62
20	60	0.5	0.1, 1	0.84/0.70	0.82/0.73	0.86/0.64	0.86/0.64
20	60	0	0.1, 1	0.83/0.71	0.81/0.74	0.85/0.63	0.85/0.65
60	20	0.9	0.1, 1	0.79/0.76	0.79/0.76	–	0.89/0.66
60	20	0.7	0.1, 1	0.79/0.76	0.78/0.76	–	0.85/0.65
60	20	0.5	0.1, 1	0.79/0.76	0.79/0.76	–	0.83/0.66
60	20	0	0.1, 1	0.79/0.76	0.79/0.76	–	0.81/0.66
100	100	0.9	0.5, 0.1	0.77/0.81	0.76/0.82	–	0.87/0.72
100	100	0.7	0.5, 0.1	0.78/0.81	0.76/0.82	–	0.82/0.72

Table 5

True positive rate/true negative rate for the FGN error covariance models averaged over 50 replications; $n = 50$. Standard errors are omitted (the largest standard error is 0.04 and most are less than 0.01). Tuning parameters were selected using a 10^x resolution.

p	q	H	s_1, s_2	lasso	seplasso	MRCE	ap.MRCE
20	20	0.95	0.1, 1	0.83/0.72	0.81/0.75	0.94/0.55	0.93/0.59
20	20	0.90	0.1, 1	0.84/0.71	0.83/0.73	0.90/0.59	0.89/0.61
20	20	0.95	0.5, 1	0.87/0.40	0.87/0.45	0.93/0.39	0.92/0.39
20	20	0.90	0.5, 1	0.86/0.43	0.87/0.45	0.88/0.51	0.90/0.43
20	60	0.95	0.1, 1	0.83/0.70	0.81/0.73	0.93/0.55	0.93/0.58
20	60	0.90	0.1, 1	0.83/0.70	0.81/0.73	0.90/0.58	0.90/0.60
60	20	0.95	0.1, 1	0.79/0.76	0.79/0.76	–	0.89/0.66
60	20	0.90	0.1, 1	0.79/0.76	0.78/0.76	–	0.87/0.65
100	100	0.95	0.5, 0.1	0.77/0.81	0.75/0.82	–	0.87/0.72
100	100	0.90	0.5, 0.1	0.77/0.81	0.75/0.82	–	0.83/0.71

Table 6

Average testing squared error for each output (company) $\times$ 1000, based on 26 testing points. Standard errors are reported in parentheses. The results for the FES method were copied from table 3 of (Yuan et al. 2007).

	OLS	sep.lasso	lasso	MRCE	ap.MRCE	FES
Walmart	0.98 (0.27)	0.44 (0.10)	0.42 (0.12)	0.41 (0.11)	0.41 (0.11)	0.40
Exxon	0.39 (0.08)	0.31 (0.07)	0.31 (0.07)	0.31 (0.07)	0.31 (0.07)	0.29
GM	1.68 (0.42)	0.71 (0.17)	0.71 (0.17)	0.71 (0.17)	0.69 (0.17)	0.62
Ford	2.15 (0.61)	0.77 (0.25)	0.77 (0.25)	0.77 (0.25)	0.77 (0.25)	0.69
GE	0.58 (0.15)	0.45 (0.09)	0.45 (0.09)	0.45 (0.09)	0.45 (0.09)	0.41
ConocoPhillips	0.98 (0.24)	0.79 (0.22)	0.79 (0.22)	0.79 (0.22)	0.78 (0.22)	0.79
Citigroup	0.65 (0.17)	0.61 (0.13)	0.66 (0.14)	0.62 (0.13)	0.62 (0.13)	0.59
IBM	0.62 (0.14)	0.49 (0.10)	0.49 (0.10)	0.49 (0.10)	0.47 (0.09)	0.51
AIG	1.93 (0.93)	1.88 (1.02)	1.88 (1.02)	1.88 (1.02)	1.88 (1.02)	1.74
AVE	1.11 (0.14)	0.72 (0.12)	0.72 (0.12)	0.71 (0.12)	0.71 (0.12)	0.67

Table 7

Estimated coefficient matrix **B** for approximate MRCE.

	Wal	Exx	GM	Ford	GE	CPhil	Citi	IBM	AIG
Walmart	0	0	0	0	0	0	0.123	0.078	0
Exxon	0	0	0	0	0	0	0	0	0
GM	0	0	0	0	0	0	0	0	0
Ford	-0.093	0.035	0.012	0	0	0	0	-0.040	-0.010
GE	0	0	0	0	0	0.044	0	0	0
ConocoPhillips	0	0.007	0	0	0	0	0	-0.005	0
Citigroup	0	0	0.025	0	0	0	0	0	0
IBM	0	0	0	0	0	0	0	0	0
AIG	0	0	0.031	0	0	0	0	0	0

Table 8

Inverse error covariance estimate for MRCE.

	Wal	Exx	GM	Ford	GE	CPhil	Citi	IBM	AIG
Walmart	1810.0	0	-378.0	0	0	0	0	0	-10.8
Exxon	0	4409.2	0	0	0	-1424.1	0	0	-8.4
GM	-378.0	0	2741.3	-1459.2	-203.5	0	-363.7	-56.0	-104.9
Ford	0	0	-1459.2	1247.4	0	0	0	0	0
GE	0	0	-203.4	0	2599.1	0	-183.7	-1358.1	-128.5
CPhillips	0	-1424.1	0	0	0	2908.2	0	0	-264.3
Citigroup	0	0	-363.7	0	-183.7	0	4181.7	0	-718.1
IBM	0	0	-56.1	0	-1358.1	0	0	3353.5	-3.6
AIG	-10.8	-8.4	-104.9	0	-128.5	-264.3	-718.1	-3.6	1714.2