

Improved Estimation of the Covariance Matrix of Stock Returns

With an Application to Portfolio Selection

Olivier Ledoit*

January 14, 1999

Abstract

This paper proposes to estimate the covariance matrix of stock returns by an optimally weighted average of two existing estimators: the single-index covariance matrix and the sample covariance matrix. This method is generally known as shrinkage, and it is standard in decision theory and in empirical Bayesian statistics. Our shrinkage estimator can be seen as a way to account for extra-market covariance without having to specify an arbitrary multi-factor structure. For NYSE and AMEX stock returns from 1972 to 1995, it can be used to select portfolios with significantly lower out-of-sample variance than a set of existing estimators, including multi-factor models.

*Assistant Professor of Finance, Anderson Graduate School of Management at UCLA, 110 Westwood Plaza, Los Angeles CA 90095-1481. E-mail: ledoit@ucla.edu; phone: (310) 206 6562; fax: (310) 206 5455. This is based on the first chapter of my Ph.D. thesis at MIT. I wish to acknowledge the help of my classmates and my professors at MIT, including the members of my thesis committee John Heaton, Bin Zhou and especially the chairman Andrew Lo. I got valuable feedback from participants to seminars at MIT, the NBER, UCLA, Washington University in Saint Louis, Yale, Chicago, Wharton and UBC. I have also benefited from comments by Fischer Black, Timothy Crack, Bruce Lehmann, Harry Markowitz, Richard Michaud, Richard Roll, Pedro Santa-Clara and Jay Shanken. Errors are my own.

The objective of this paper is to estimate the covariance matrix of stock returns. This is a fundamental question in empirical Finance with implications for portfolio selection and for tests of asset pricing models.

The traditional estimator – the sample covariance matrix – is seldom used because it imposes too little structure. When the number of stocks N is of the same order of magnitude as the number of historical returns per stock T , the total number of parameters to estimate is of the same order as the total size of the data set, which is clearly problematic. When N is larger than T , the sample covariance matrix is always singular, even if the true covariance matrix is known to be non-singular.¹

These severe problems may come as a surprise, since the sample covariance matrix has appealing properties, such as being maximum likelihood under normality. But this is to forget what maximum likelihood means. It means the most likely parameter values given the data. In other words: let the data speak (and only the data). This is a sound principle, provided that there is enough data to trust the data. Indeed, maximum likelihood is justified asymptotically as the number of observations per variable goes to infinity. It is a general drawback of maximum likelihood that it can perform poorly in small sample. For the covariance matrix, small sample problems occur unless T is at least one order of magnitude larger than N (see Ledoit, 1996).

The cure is to impose some structure on the estimator. The particular form of the structure must be dictated by the problem at hand. In the case of stock returns, a low-dimensional factor structure seems natural. But this leaves two very important questions: *How much structure should we impose? And what factors should we use?*

¹In typical applications, there can be over a thousand stocks to choose from, but rarely more than ten years of monthly data, i.e. $N = 1000$ and $T = 120$.

To address these questions properly, we have to be more specific about how we impose a low-dimensional factor structure. One possible way is to specify a K -factor model with uncorrelated residuals. Then K controls how much structure we impose: the fewer the factors, the stronger the structure. The advantages of this approach are that it is quite familiar to the Finance profession, and that the factors sometimes have economic interpretation. The disadvantages are that there is no consensus on the identity of the factors – except for the first one, which represents a market index –, and there is no consensus on the number of factors K either (Connor and Korajczyk, 1992). In other words, choosing between factor models is very *ad-hoc*. It does not mean that none of them works well, it means that we do not know which one works well *a priori*. For example, if we are interested in selecting portfolios with low out-of-sample variance, in any given data set there may exist a factor model that performs well, but it may be a different one for every data set, and there is no way of telling which one works well without looking out-of-sample, which is cheating. The art of choosing a factor model adapted to a given data set without seeing its out-of-sample fit is just that: an art.

This is why, in this paper, we study another way of imposing factor structure. It is to take a weighted average of the sample covariance matrix with Sharpe's (1963) single-index model estimator. The weight α^* (between zero and one) assigned to the single index model controls how much structure we impose: the heavier the weight, the stronger the structure. This is a well-known technique in Statistics called *shrinkage* and due to Stein (1956): α^* is called the shrinkage intensity, and the single-index model is our choice of shrinkage target. The advantages are that there is strong consensus on the nature of the single factor (a market index), and that there is a way of estimating the optimal shrinkage rate consistently. The estimation of α^* is

the technically challenging part of this paper. It provides a rigorous answer to the question of how much structure we should impose. On any given dataset, there will be a different optimal shrinkage intensity, and our estimation technique will find it without having to look out-of-sample. This takes the ad-hockery out of the task of imposing structure on the covariance matrix of stock returns. It replaces the art of factor selection by a fully automatic procedure.

At this point, it is worth mentioning that the paper is solely concerned with the structure of risk in the stock market, not with the structure of expected returns. Multi-factor models of the covariance matrix can still be very useful if economic arguments tie them up to the cross-section of expected returns, as in the Arbitrage Pricing Theory of Ross (1976). Any discussion of the relationship between risk factors and expected returns is totally outside the scope of the paper. There should be no ambiguity over whether we define “factors” in terms of the mean vector or of the covariance matrix of stock returns: it is always the latter.

Muirhead (1987) reviews the large literature on shrinkage estimators of the covariance matrix in finite-sample statistical decision theory. All these estimators suffer from at least two severe drawbacks, either of which is enough to make them ill-suited to stock returns: (i) they break down when $N > T$; (ii) they do not exploit the *a priori* knowledge that stock returns tend to be positively correlated to one another. Ledoit (1996) shows that the solution to the first problem is to define the optimal shrinkage intensity by minimizing a loss function that does not involve the inverse of the covariance matrix, but he does not address the second problem. Frost and Savarino (1986) show that the solution to the second problem is to use a shrinkage target that incorporates a market factor, but they ignore without justification the correlation between estimation error on the shrinkage target and on the sample covariance matrix, and they are still exposed to the

first problem. The contribution of our paper to the literature is to combine the advantages of Ledoit (1996) and Frost and Savarino (1986) while eliminating their respective drawbacks. This extension is achieved by substantially modifying the asymptotic framework of Ledoit (1996). Our technique is so general that it is applicable to any shrinkage target.

A noteworthy innovation is that the optimal shrinkage intensity depends on the correlation between estimation error on the sample covariance matrix and on the shrinkage target. Intuitively, if the two of them are positively (negatively) correlated, then the benefit of combining the information that they contain is smaller (larger). The introduction of this correlation term resolves a deep logical inconsistency in earlier empirical Bayesian literature, where the prior is estimated from sample data, yet at the same time is assumed to be independent from sample data.

We test the performance of our shrinkage estimator on stock returns data for portfolio selection. Using NYSE and AMEX stocks from 1972 to 1995, we find that our estimator yields portfolios with significantly lower out-of-sample variance than a set of well-established competitors, including multi-factor models.

A potential application of this shrinkage estimator will be to run predictive cross-sectional Generalized Least Squares regressions of average stock returns on betas and various stock attributes. GLS regressions have more power than the OLS regressions currently being run, and their economic interpretation is clearer because they are directly related to portfolio selection (see Kandel and Stambaugh, 1995).

The first section presents our shrinkage estimator of the covariance matrix. The following section presents empirical evidence on its out-of-sample performance for portfolio selection. The last section concludes.

1 Shrinkage Estimator of the Covariance Matrix

This section presents the covariance matrix estimator that we recommend for stock returns.

1.1 Statistical Model

Let \mathbf{X} denote an $N \times T$ matrix of T observations on a system of N random variables representing T historical returns on a universe of N stocks.

Assumption 1 *Stock returns are independent and identically distributed (iid) through time.*

Even though actual stock returns do not verify Assumption 1, it is an acceptable first-cut approximation. This means that we abstract from lead-lag effects (Lo and MacKinlay, 1990), nonsynchronous trading (Shanken, 1987), and autoregressive conditional heteroskedasticity (Bollerslev, Engle and Wooldridge, 1988). There is room for improving the performance of our estimator by relaxing Assumption 1. For example, Ledoit and Santa-Clara (1997) have recently developed a technique for estimating the full-blown Multivariate GARCH(1,1) model even for large number of assets.

Assumption 2 *The number of stocks N is fixed and finite, while the number of observations T goes to infinity.*

This is the classic asymptotic framework. It stands in sharp contrast to the non-standard asymptotic framework of Ledoit (1996) where N went to infinity with T . Being able to transpose the basic intuition of Ledoit (1996) to the classic asymptotic framework is the key to extending his results to any shrinkage target, including the single-index model. This fundamental change not only makes the method much more general, but it also makes the proofs much more transparent.

Assumption 3 *Stock returns have finite fourth moment:*

$$\forall i, j, k, l = 1, \dots, n \quad \forall t = 1, \dots, T \quad E[|x_{it}x_{jt}x_{kt}x_{lt}|] < \infty \quad (1)$$

This is so that we can apply the Central Limit Theorem to sample variances and covariances.

Note that stock returns are not assumed to be normally distributed in this paper.

1.2 Sample Covariance Matrix

The sample mean vector \mathbf{m} and the sample covariance matrix \mathbf{S} are defined by:

$$\mathbf{m} = \frac{1}{T} \mathbf{X} \mathbf{1} \quad (2)$$

$$\mathbf{S} = \frac{1}{T} \mathbf{X} \left(\mathbf{I} - \frac{1}{T} \mathbf{1} \mathbf{1}' \right) \mathbf{X}' \quad (3)$$

where $\mathbf{1}$ denotes a conformable vector of ones and \mathbf{I} a conformable identity matrix.² Equation (3) shows why the sample covariance matrix is not invertible when $N \geq T$: the rank of \mathbf{S} is at most equal to the rank of the matrix $\mathbf{I} - \mathbf{1} \mathbf{1}' / T$, which is $T - 1$. Therefore when the dimension N exceeds $T - 1$ the sample covariance matrix is rank-deficient. Intuitively, the data do not contain enough information to estimate the unrestricted covariance matrix.

²Vectors (matrices) are denoted in lower (upper) case boldface.

1.3 Single-Index Covariance Matrix Estimator

Sharpe's (1963) single-index model assumes that stock returns are generated by:

$$x_{it} = \alpha_i + \beta_i x_{0t} + \varepsilon_{it} \quad (4)$$

where residuals ε_{it} are uncorrelated to market returns x_{0t} and to one another and have variance δ_{ii} . The covariance matrix implied by this model is:

$$\Phi = \sigma_{00} \beta \beta' + \Delta \quad (5)$$

where σ_{00} is the variance of market returns, β is the vector of slopes, and Δ is the diagonal matrix containing residual variances. Call ϕ_{ij} the (i, j) entry of Φ .

This model can be estimated by running a time-series regression of stock i 's returns on the market. Call b_i the slope estimate and d_i the residual variance estimate. Then the single-index model yields the following estimator for the covariance matrix of stock returns:

$$\mathbf{F} = s_{00} \mathbf{b} \mathbf{b}' + \mathbf{D} \quad (6)$$

where s_{00} is the sample variance of market returns, \mathbf{b} is the vector of slope estimates, and \mathbf{D} is the diagonal matrix containing residual variance estimates. Call f_{ij} the (i, j) entry of \mathbf{F} . We need to make two technical assumptions.

Assumption 4 $\Phi \neq \Sigma$

Assumption 5 *The market portfolio has strictly positive variance.*

The exact composition of the market portfolio is not as critical here as it is for the CAPM (Roll, 1977). All we need is for x_{0t} to explain a lot of the variance of most stocks, and any broad-based market index would do. As a matter of fact, equal-weighted indices are better at explaining stock market variance than value-weighted indices, yet another departure from CAPM intuition. The assumption that residuals are uncorrelated to one another should theoretically preclude the market from being a portfolio of the N stocks in the sample. But since the market model is misspecified anyway (Assumption 4), it does not matter. For simplicity, we take the market index as the equal-weighted portfolio of the N stocks in the sample:

$$x_{0t} = \frac{1}{N} \sum_{i=1}^N x_{it} \quad (7)$$

1.4 General Form of the Shrinkage Estimator

At one extreme the market index covariance matrix comes from a one-factor model, while at the other extreme the sample covariance matrix can be interpreted as an N -factor model (each stock being a factor, there are no residuals). The intuition of the profession has always been that the best model lies somewhere between these two extremes. For example, Rosenberg (1974) stresses the importance of extra-market covariance, while Jobson and Korkie (1980) document the poor performance of the sample covariance matrix. Until now, this intuition has been expressed mostly through K -factor models with $1 < K < N$.

If we abstract ourselves from the Finance context and look at the broader picture from a statistician's point of view, we see another way of capturing the same intuition. The key is to recognize that the market model covariance matrix has a lot of bias coming from a stringent

and misspecified structural assumption, but little in the way of estimation error, and that the opposite is true of the sample covariance matrix: it is unbiased (asymptotically) but has a lot of estimation error. A fundamental principle of statistical decision theory is that there exists an interior optimum in the trade-off between bias and estimation error. Since Stein's (1956) seminal work, we know that one way of attaining this optimal trade-off is simply to take a properly weighted average of the biased and unbiased estimators. This is called *shrinking* the unbiased estimator full of estimation error towards a fixed target represented by the biased estimator. For example, Stein (1956) showed that shrinking sample means towards a constant can, under certain circumstances, improve accuracy. Efron and Morris (1977) provide a general introduction to shrinkage, and Jorion (1986) shows its importance in the context of portfolio selection.

Here this well-established statistical method suggests taking a weighted average of the market model covariance matrix and the sample covariance matrix. This is an alternative way of capturing the intuition that led to the development of multi-factor models. The main advantage of this alternative is that it does not require knowledge of the number and nature of factors (beyond the obvious market factor).

1.5 Formula for the Optimal Shrinkage Intensity

The only technical difficulty is to find the shrinkage intensity. This section defines the optimal shrinkage intensity and gives an explicit formula for it as a function of the true (and unobservable) covariance matrix. We will then produce a *bona-fide* estimator in the next section.

Since N is fixed and T goes to infinity, \mathbf{S} is consistent but \mathbf{F} is not, therefore the optimal shrinkage intensity vanishes asymptotically. Standard arguments imply that it is of the order

constant/ T (see Ledoit, 1996). Therefore the only interesting question is to find the constant, which we call α . The shrinkage covariance matrix can be written as: $\frac{\alpha}{T}\mathbf{F} + \left(1 - \frac{\alpha}{T}\right)\mathbf{S}$. In this framework, shrinkage can be thought of as an asymptotic first-order correction.³

We have to choose the objective according to which the shrinkage intensity α/T is asymptotically “optimal.” All existing shrinkage estimators from finite-sample statistical decision theory and also Frost and Savarino’s (1986) break down when $N \geq T$ because their objectives involve the inverse of the covariance matrix. One of the contributions of Ledoit (1996) was to recognize this fact, and then propose an objective that did not depend on the inverse of the covariance matrix. This objective is extremely intuitive: it is a quadratic measure of distance between true and estimated covariance matrices based on the Frobenius norm.

Definition 1 *The Frobenius norm of the $N \times N$ symmetric matrix \mathbf{Z} with entries $(z_{ij})_{i,j=1,\dots,N}$ and eigenvalues $(\lambda_i)_{i=1,\dots,N}$ is defined by:*

$$\|\mathbf{Z}\|^2 = \text{Trace}(\mathbf{Z}^2) = \sum_{i=1}^N \sum_{j=1}^N z_{ij}^2 = \sum_{i=1}^N \lambda_i^2 \quad (8)$$

Our objective is based on the quadratic loss function $\left\| \frac{\alpha}{T}\mathbf{F} + \left(1 - \frac{\alpha}{T}\right)\mathbf{S} - \mathbf{\Sigma} \right\|^2$. This loss vanishes at the rate constant/ T , therefore the right quantity to consider is: $\left\| \frac{\alpha}{T}\mathbf{F} + \left(1 - \frac{\alpha}{T}\right)\mathbf{S} - \mathbf{\Sigma} \right\|^2 T$. One possibility would be to minimize the limit of its expectation. But without further assumptions little is known about the limiting behavior (or even the existence!) of this expectation. Therefore we must focus instead on the expectation of the limiting distribution. This is analogous to the familiar notion of asymptotic variance. We denote the asymptotic expectation by AsyE. Finally,

³In the introduction of the paper, the notation was to call α^* the weight placed on the single-index model, thus $\alpha^* = \alpha/T$.

we are ready to state the objective as:

$$\min_{\alpha} \text{AsyE} \left[\left\| \frac{\alpha}{T} \mathbf{F} + \left(1 - \frac{\alpha}{T}\right) \mathbf{S} - \boldsymbol{\Sigma} \right\|^2 T \right] \quad (9)$$

This optimization program has a unique solution. It is available in closed form, as shown by the following theorem.

Theorem 1 *Let π denote the sum of asymptotic variances of errors on entries of the sample covariance matrix: $\pi = \sum_{i=1}^N \sum_{j=1}^N \text{AsyVar} [\sqrt{T} s_{ij}]$. Similarly, let ρ denote the sum of asymptotic covariances of errors on entries of the single-factor covariance matrix with errors on entries of the sample covariance matrix: $\rho = \sum_{i=1}^N \sum_{j=1}^N \text{AsyCov} [\sqrt{T} f_{ij}, \sqrt{T} s_{ij}]$. Finally, let γ measure the misspecification of the single-factor model: $\gamma = \sum_{i=1}^N \sum_{j=1}^N (\phi_{ij} - \sigma_{ij})^2$. Then the optimal shrinkage rate that solves Equation (9) is:*

$$\alpha = \frac{\pi - \rho}{\gamma} \quad (10)$$

Proof of Theorem 1 The squared Frobenius norm is a quadratic form on the space of N -dimensional symmetric matrices whose associated inner product is: $\mathbf{Z}_1 \circ \mathbf{Z}_2 = \text{Trace}(\mathbf{Z}_1 \mathbf{Z}_2)$.

Let Θ denote the objective to minimize in Equation (9). The optimal shrinkage rate α verifies $d\Theta/d\alpha = 0$.

$$\begin{aligned} \Theta &= \text{AsyE} \left[\left\| \frac{\alpha}{T} \mathbf{F} + \left(1 - \frac{\alpha}{T}\right) \mathbf{S} - \boldsymbol{\Sigma} \right\|^2 T \right] \\ &= \text{AsyE} \left[\left\{ \frac{\alpha^2}{T^2} \|\mathbf{F}\|^2 + \left(1 - \frac{\alpha}{T}\right)^2 \|\mathbf{S}\|^2 + \|\boldsymbol{\Sigma}\|^2 \right. \right. \end{aligned}$$

$$\begin{aligned}
& +2\frac{\alpha}{T} \left(1 - \frac{\alpha}{T}\right) \mathbf{F} \circ \mathbf{S} - 2\frac{\alpha}{T} \mathbf{F} \circ \boldsymbol{\Sigma} - 2 \left(1 - \frac{\alpha}{T}\right) \mathbf{S} \circ \boldsymbol{\Sigma} \Big\} T \Big] \\
\frac{d\Theta}{d\alpha} &= 2\text{AsyE} \left[\frac{\alpha}{T} \|\mathbf{F}\|^2 - \left(1 - \frac{\alpha}{T}\right) \|\mathbf{S}\|^2 + \left(1 - 2\frac{\alpha}{T}\right) \mathbf{F} \circ \mathbf{S} - \mathbf{F} \circ \boldsymbol{\Sigma} + \mathbf{S} \circ \boldsymbol{\Sigma} \right] \\
&= 2\text{AsyE} \left[\frac{\alpha}{T} \{ \|\mathbf{F}\|^2 + \|\mathbf{S}\|^2 - 2\mathbf{F} \circ \mathbf{S} \} - \{ \|\mathbf{S}\|^2 - \mathbf{F} \circ \mathbf{S} + \mathbf{F} \circ \boldsymbol{\Sigma} - \mathbf{S} \circ \boldsymbol{\Sigma} \} \right] \\
\alpha &= \frac{\text{AsyE} \left[\{ \|\mathbf{S}\|^2 - \mathbf{F} \circ \mathbf{S} + \mathbf{F} \circ \boldsymbol{\Sigma} - \mathbf{S} \circ \boldsymbol{\Sigma} \} T \right]}{\text{AsyE} \left[\|\mathbf{F}\|^2 + \|\mathbf{S}\|^2 - 2\mathbf{F} \circ \mathbf{S} \right]} \\
&= \frac{\text{AsyE}[(\mathbf{S} - \mathbf{F}) \circ (\mathbf{S} - \boldsymbol{\Sigma})T]}{\text{AsyE}[\|\mathbf{F} - \mathbf{S}\|^2]} \\
&= \frac{\text{AsyE}[\|\mathbf{S} - \boldsymbol{\Sigma}\|^2 T] - \text{AsyE}[(\mathbf{F} - \boldsymbol{\Phi}) \circ (\mathbf{S} - \boldsymbol{\Sigma})T]}{\|\boldsymbol{\Phi} - \boldsymbol{\Sigma}\|^2} \\
&= \frac{\pi - \rho}{\gamma}
\end{aligned}$$

where we have used $\text{AsyE}[\mathbf{F}] = \boldsymbol{\Phi}$, $\text{AsyE}[\mathbf{S}] = \boldsymbol{\Sigma}$ and $\text{AsyE}[\|\mathbf{F} - \boldsymbol{\Phi}\|^2] = \text{AsyE}[\|\mathbf{S} - \boldsymbol{\Sigma}\|^2] = 0$.

□

The weight placed on the single-factor model increases in the error on the sample covariance matrix (through π) and decreases in the misspecification of the market model (through γ). The most palatable interpretation of this solution is a geometric one: $\frac{\alpha}{T}\mathbf{F} + \left(1 - \frac{\alpha}{T}\right)\mathbf{S}$ is (asymptotically) the orthogonal projection of the true covariance matrix $\boldsymbol{\Sigma}$ onto the line joining market model and sample covariance matrices (see Figure 1).

Note also the appearance of the term ρ that measures covariance between estimation errors of \mathbf{S} and \mathbf{F} . It was neither present in the work of Ledoit (1996) nor Frost and Savarino (1986). It is easier to understand within the empirical Bayesian interpretation of our shrinkage estimator: we can say that we have a prior based on the market model which we combine with sample information. Estimating the prior \mathbf{F} from the same dataset as \mathbf{S} violates the pure Bayesian principle that prior and sample information should be independent. This is typical of the empirical

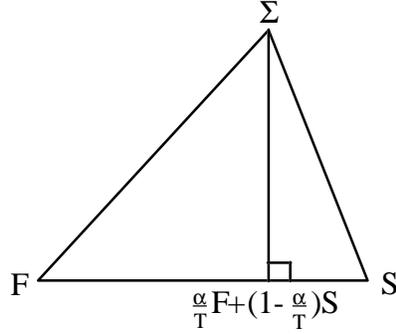


Figure 1: Geometric Interpretation of Theorem 1. The notion of orthogonality among N -dimensional symmetric matrices is defined by the inner product associated with the Frobenius norm.

Bayesian approach. Yet this violation is often ignored, and there is an “art” to choosing the prior so that the violation is not too damaging. We get rid of the need for such artistry by explicitly taking into account the correlation between prior and sample information through ρ .

Finally, the reader can verify that Theorem 1 is totally general: nowhere did we use the fact that \mathbf{F} is a single-index model estimator. Equation (10) stays the same as long as \mathbf{F} is an asymptotically biased estimator of the covariance matrix.

1.6 A Consistent Estimator for the Optimal Shrinkage Intensity

Note that $\frac{\alpha}{T}\mathbf{F} + \left(1 - \frac{\alpha}{T}\right)\mathbf{S}$ is not a *bona fide* estimator because α depends on unobservables. Therefore we need to find a consistent estimator for $\alpha = (\pi - \rho)/\gamma$. This task is considerably facilitated by our choice of a quadratic objective in Equation (9), since in general estimating quadratic functionals is an easy task.

We can decompose π into $\pi = \sum_{i=1}^N \sum_{j=1}^N \pi_{ij}$ where $\pi_{ij} = \text{AsyVar} \left[\sqrt{T} s_{ij} \right]$, ρ into $\rho = \sum_{i=1}^N \sum_{j=1}^N \rho_{ij}$ where $\rho_{ij} = \text{AsyCov} \left[\sqrt{T} f_{ij}, \sqrt{T} s_{ij} \right]$, and γ into $\gamma = \sum_{i=1}^N \sum_{j=1}^N \gamma_{ij}$ where $\rho_{ij} = (\phi_{ij} - \sigma_{ij})^2$. The usual asymptotic theory provides consistent estimators for π_{ij} , ρ_{ij} and γ_{ij} .

Lemma 1 A consistent estimator for π_{ij} is given by:

$$p_{ij} = \frac{1}{T} \sum_{t=1}^T \{(x_{it} - m_i)(x_{jt} - m_j) - s_{ij}\}^2 \quad (11)$$

Proof of Lemma 1 p_{ij} is the usual estimator for the asymptotic variance of s_{ij} . It converges to $\text{Var}[(x_{i1} - \mu_i)(x_{j1} - \mu_j)]$, which is equal to π_{ij} . \square

Let m_0 denote the sample mean of market returns and s_{i0} the sample covariance of stock i 's returns with the market.

Lemma 2 On the diagonal a consistent estimator of ρ_{ii} is given by $r_{ii} = p_{ii}$, and for $i \neq j$ a consistent estimator of ρ_{ij} is given by $r_{ij} = \frac{1}{T} \sum_{t=1}^T r_{ijt}$ where:

$$r_{ijt} = \frac{s_{j0}s_{00}(x_{it} - m_i) + s_{i0}s_{00}(x_{jt} - m_j) - s_{i0}s_{j0}(x_{0t} - m_0)}{s_{00}^2} (x_{0t} - m_0)(x_{it} - m_i)(x_{jt} - m_j) - f_{ij}s_{ij} \quad (12)$$

Proof of Lemma 2 On the diagonal $f_{ii} = s_{ii}$, therefore $\rho_{ii} = \pi_{ii}$ can be consistently estimated by $r_{ii} = p_{ii}$. When $i \neq j$ we have: $f_{ij} = b_i b_j s_{00} = s_{i0}s_{j0}/s_{00}$, therefore the ‘‘delta’’ method yields:

$$\rho_{ij} = \text{AsyCov} \left[\sqrt{T} \frac{s_{i0}s_{j0}}{s_{00}}, \sqrt{T} s_{ij} \right] \quad (13)$$

$$\begin{aligned} &= \frac{s_{j0}}{s_{00}} \text{AsyCov} \left[\sqrt{T} s_{i0}, \sqrt{T} s_{ij} \right] + \frac{s_{i0}}{s_{00}} \text{AsyCov} \left[\sqrt{T} s_{j0}, \sqrt{T} s_{ij} \right] \\ &\quad - \frac{s_{i0}s_{j0}}{s_{00}^2} \text{AsyCov} \left[\sqrt{T} s_{00}, \sqrt{T} s_{ij} \right] \end{aligned} \quad (14)$$

The usual estimator for $\text{AsyCov} [\sqrt{T}s_{k0}, \sqrt{T}s_{ij}]$ where $k = 0, k = i$ or $k = j$ is:

$$\frac{1}{T} \sum_{t=1}^T \left\{ (x_{kt} - m_k)(x_{0t} - m_0) - s_{k0} \right\} \left\{ (x_{it} - m_i)(x_{jt} - m_j) - s_{ij} \right\} \quad (15)$$

Plugging this estimator into Equation (14) and rearranging yields Equation (12). \square

Lemma 3 *A consistent estimator for $\gamma_{ij} = (\phi_{ij} - \sigma_{ij})^2$ is its sample counterpart $c_{ij} = (f_{ij} - s_{ij})^2$.*

Proof of Lemma 3 This is because f_{ij} and s_{ij} are consistent estimators for ϕ_{ij} and σ_{ij} respectively. \square

Now it is easy to construct an estimator for the optimal shrinkage rate.

Theorem 2 *$a = (p - r)/c$ is a consistent estimator for the optimal shrinkage rate $\alpha = (\pi - \rho)/\gamma$.*

Proof of Theorem 2 Under Assumption 4 ($\gamma > 0$), combining the results of Lemmas 1-3 proves the theorem. \square

Using this notation, the shrinkage estimator that we recommend for the covariance matrix is:

$$\boxed{\hat{\mathbf{S}} = \frac{a}{T} \mathbf{F} + \left(1 - \frac{a}{T}\right) \mathbf{S}} \quad (16)$$

Different shrinkage targets would lend themselves just as well to the estimation of π , ρ and γ .

The formula for r would need to be readjusted on a case-by-case basis, but the formulas for p and c would not change.

2 Empirical Results

We present empirical evidence on the performance of the shrinkage estimator defined in the last section. We compare it to existing estimators in terms of its ability to select portfolios of stocks with low out-of-sample variance.

2.1 Portfolio Selection

Consider a universe of N stocks whose returns are distributed with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Markowitz (1952) defines the problem of portfolio selection as:

$$\begin{array}{l} \min_{\mathbf{w}} \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \\ \text{subject to} \end{array} \left\{ \begin{array}{l} \mathbf{w}'\mathbf{1} = 1 \\ \mathbf{w}'\boldsymbol{\mu} = q \end{array} \right. \quad (17)$$

where $\mathbf{1}$ denotes a conformable vector of ones and q is the expected rate of return that is required on the portfolio. The well-known solution is:

$$\mathbf{w} = \frac{C - qB}{AC - B^2}\boldsymbol{\Sigma}^{-1}\mathbf{1} + \frac{qA - B}{AC - B^2}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} \quad (18)$$

$$\text{where } A = \mathbf{1}'\boldsymbol{\Sigma}^{-1}\mathbf{1}, \quad B = \mathbf{1}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} \quad \text{and} \quad C = \boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$$

Equation (18) shows that optimal portfolio weights depend on the inverse of the covariance matrix. This sometimes causes difficulty if the covariance matrix estimator is not invertible or if it is numerically ill-conditioned, which means that inverting it amplifies estimation error tremendously (Michaud, 1989). The shrinkage estimator $\hat{\mathbf{S}}$ is the weighted average of two positive semi-definite

matrices, one of which (F) is invertible, therefore it is invertible. It inherits the good-conditioning of the single-index model estimator, *not* the ill-conditioning of the sample covariance matrix.

In practice, the covariance matrix is estimated from historical data available up to a given date, optimal portfolio weights are computed from this estimate, then the portfolio is formed on that date and held until the next rebalancing occurs. The performance of a covariance matrix estimator is measured by the variance of this optimal portfolio *after* it is formed. It is a measure of out-of-sample performance, or of predictive ability. An estimator that overfits in-sample data can turn out to work very poorly for portfolio selection, which is why imposing some structure is beneficial.

The other input into portfolio selection is the vector of expected returns. It is sometimes argued that estimating the covariance matrix well is less important than estimating the expected returns well. We believe that this view is profoundly misguided. First, the essence of mean-variance analysis is that there is a trade-off between risk and return, therefore any reduction in risk translates into an increase in expected returns. Second, having a good estimator of the covariance matrix helps us estimate more precisely the excess return associated with, for example, beta, size or book-to-market, by constructing portfolios that load on these characteristics and have low variance (this is related to the efficiency gain in running GLS cross-sectional regressions of stock returns on these characteristics). Third, it is not the role of just the statistician to determine expected returns, it is also the role of the economist, and of the portfolio manager who is paid to generate valuable private information about future stock price movements; whereas only statistics can generate information about the covariance matrix. This justifies why it is perfectly legitimate to concentrate on the covariance matrix alone without worrying about expected returns, as we do

here.

2.2 Data

Stock returns are extracted from the Center for Research in Security Prices (CRSP) monthly database. The same procedure is repeated for every year from $t = 1972$ to $t = 1994$. We use data from August of year $t - 10$ to July of year t to estimate the covariance matrix of stock returns.⁴ Then on the first trading day in August of year t we build a portfolio with minimum variance (according to this covariance matrix estimate) under certain constraints. We hold this portfolio until the last trading day in July of year $t + 1$, at which time we liquidate it and start the process all over again. Thus, the *in-sample* period goes from August of year $t - 10$ to July of year t , and the *out-of-sample* period goes from August of year t to July of year $t + 1$. The main quantity of interest is the out-of-sample standard deviation of this investment strategy over the 23-year period from August 1972 to July 1995. This is a predictive test, in the sense that our investment strategy does not require any hindsight.

In August of year t , we consider the universe of common stocks traded on the New York Stock Exchange (NYSE) and the American Stock Exchange (AMEX) with valid CRSP returns for the last 120 months and valid Standard Industrial Classification (SIC) codes. The resulting number of stocks varies across years between $N = 909$ and $N = 1314$.

We consider two minimum variance portfolios: the global minimum variance portfolio, and the portfolio with minimum variance under the constraint of having 20% expected return. In both cases short sales are allowed, and no additional restriction is imposed (except that weights

⁴We rebalance every August because the earliest AMEX stock returns available from CRSP are in August 1962.

sum up to one). It is common practice in the investment community to impose client-specific constraints on portfolio weights or to minimize active risk with respect to an exogenously specified benchmark, but we abstract from that in order to have a “maximum stress” test of the performance of the covariance matrix estimator. For expected returns, we just take the average realized return over the last 10 years. This may or may not be a good predictor of future expected returns, but our goal is not to predict expected returns: it is only to show what kind of reduction in out-of-sample variance our method yields under a fairly reasonable linear constraint.

2.3 Competing Estimators

We run a horse race between various covariance matrix estimators proposed in the literature.

Identity The simplest model is to assume that the covariance matrix is a scalar multiple of the identity matrix. This is the assumption implicit in running an Ordinary Least Squares (OLS) cross-sectional regression of stock returns on stock characteristics, as Fama and MacBeth (1973) and their successors do. Interestingly, it yields the same weights for the minimum variance portfolios as a two-parameter model where all variances are equal to one another and all covariances are equal to one another. This two-parameter model is discussed by Jobson and Korkie (1980), and by Frost and Savarino (1986).

Constant Correlation Elton and Gruber (1973) recommend a model where every pair of stocks has the same correlation coefficient. Thus, there are $N + 1$ parameters to estimate: the N individual variances, and the constant correlation coefficient.

Pseudo-Inverse It is impossible to use the sample covariance matrix directly for portfolio selection when the number of stocks N exceeds the number of historical returns T , which is the case here. The problem is that we need the inverse of the sample covariance matrix, and it does not exist. One possible trick to get around this problem is to use the pseudo-inverse, also called generalized inverse or Moore-Penrose inverse. Replacing the inverse of the sample covariance matrix by the pseudo-inverse into Equation (18) yields well-defined portfolio weights.

Market Model This is the single-factor covariance matrix of Sharpe (1963), which is defined in Section 1.3.

Industry Factors This refinement of the single-index model assumes that market residuals are generated by industry factors:

$$x_{it} = \alpha_i + \beta_i x_{0t} + \sum_{k=1}^K c_{ik} z_{kt} + \varepsilon_{it} \quad (19)$$

where K is the number of industry factors, c_{ik} is a dummy variable equal to one if stock i belongs to industry category k , z_{kt} is the return to the k^{th} industry factor in period t , and ε_{kt} denotes residuals that are uncorrelated to the market, to industry factors, and to each other. Every stock is assigned to one of the 48 industries defined by Fama and French (1997). This high number of factors is similar to the one used by the company BARRA to produce commercial multi-factor estimates of the covariance matrix (Kahn, 1993).⁵ Industry factor returns are defined as the return to an equally-weighted portfolio of the stocks from this industry in our sample.

⁵In addition, BARRA use proprietary methods, including other factors that are not industry-based, therefore this is not a test of their performance.

Principal Components An alternative approach to multi-factor models is to extract the factors from the sample covariance matrix itself using a statistical method such as principal components. Some investment consultants such as Advanced Portfolio Technologies successfully use a refined version of this approach (Blin, 1997). Since principal components are chosen solely for their ability to explain risk, fewer factors are necessary, but they do not have any direct economic interpretation.⁶ A sophisticated test by Connor and Korajczyk (1993) finds between four and seven factors for the NYSE and AMEX over 1967-1991, which is in the same range as the original test by Roll and Ross (1980). The number of factors that we use here is five.

Shrinkage Towards Identity The original shrinkage estimator of Ledoit (1996) used a scalar multiple of the identity matrix as shrinkage target. This may not be very appropriate for stock returns, since stocks have different variances and mainly positive covariances. Nevertheless, we include this estimator.

Shrinkage Towards Market Model This is the estimator \hat{S} developed in Section 1.6.

2.4 Out-of-Sample Standard Deviations

For every one of the eight estimators described in the previous section, we compute the out-of-sample (annualized) standard deviation of the minimum variance portfolios as per Section 2.2.

The results are in Table 1.

We can see that naive diversification (one dollar in every stock) performs the worst, while the shrinkage estimator developed above performs the best. The t -statistics of whether \hat{S} yields

⁶Except for the first factor, which is highly correlated with the market index.

	Std. Deviation Unconstrained	Std. Deviation Constrained
Identity	17.75 (0.44)	17.94 (0.42)
Constant Correlation	14.27 (0.19)	16.30 (0.29)
Pseudo-Inverse	12.37 (0.23)	13.73 (0.32)
Market Model	12.00 (0.16)	13.77 (0.27)
Industry Factors	10.84 (0.17)	12.32 (0.23)
Principal Components	10.31 (0.16)	11.30 (0.22)
Shrinkage To Identity	10.21 (0.17)	11.11 (0.21)
Shrinkage To Market	9.55 (0.15)	10.43 (0.20)

Table 1: Risk of Minimum Variance Portfolios. “Unconstrained” refers to the global minimum variance portfolio, while “constrained” refers to the minimum variance portfolio with 20% expected return. Standard deviation is measured out-of-sample at the monthly frequency, annualized through multiplication by $\sqrt{12}$, and expressed in percents. Standard errors on these standard deviation estimates are reported in parenthesis.

portfolios with lower variance than its seven competitors range from 2.73 (against the other shrinkage estimator) to 7.39 (against the constant correlation model). They are all significant at the one percent level.

To assess economic significance, the rule of thumb is that a decrease of two basis points in standard deviation corresponds to an increase of one basis points in expected returns, using standard numbers for the risk-return tradeoff.⁷ For example, gains over the two multi-factor models are 43 and 94 basis points respectively in terms of average returns for the constrained portfolio. By this metric, improvement over the two multi-factor models and Ledoit’s (1996) earlier shrinkage method is reasonable. Improvement over the other estimators, including the sample covariance matrix and the single-index model (the ones which we are combining together), is large.

⁷Specifically: 8.5% market risk premium and 17% market standard deviation.

2.5 Weight Distribution

We also report descriptors of the weights of the global minimum variance portfolio: turnover, short interest, lowest and highest weight. When a stock has missing out-of-sample observations we assume that it earns the riskfree rate. Thus, it is important to report the cash position, defined as the total amount invested in stocks with missing observations. These weight descriptors are in Table 2.

	Turnover	Short Interest	Lowest Weight	Highest Weight	Cash Position
Identity	6	0	0.09	0.09	2.50
Constant Correlation	24	68	-0.17	2.86	-1.40
Pseudo-Inverse	96	99	-1.16	1.12	1.46
Market Model	23	51	-0.41	2.10	0.13
Industry Factors	43	89	-1.08	2.81	0.54
Principal Components	49	80	-0.84	2.95	0.78
Shrinkage To Identity	71	113	-1.23	1.19	0.99
Shrinkage To Market	61	98	-1.01	3.81	0.78

Table 2: Weight Descriptors. These are expressed in percents and averaged over the 23 years in our sample. A short interest of 68% means that for every dollar invested in the portfolio we short 68 cents worth of stocks, while buying \$1.68 worth of other stocks. Annual turnover above 100% is possible because of short sales.

If we wanted to do portfolio selection in practice, all these minimum variance portfolios would have too much of a short interest to be really attractive. But absolutely no effort was made to control for this characteristic, therefore there is considerable room for improving it without substantially degrading the performance reported in Table 1. In our view, Table 2 indicates that the weight distribution generated by the shrinkage estimator is acceptable overall.

The cash position due to missing out-of-sample observations is always very small, therefore the standard deviations in Table 1 do indeed correspond to portfolios almost fully invested in

equities. Otherwise, it would be easy to get low risk simply by holding cash.

2.6 Shrinkage Intensity

Figure 2 shows how the estimate of the optimal shrinkage intensity evolves through the 23 years in our sample.

Figure 2: Optimal Shrinkage Intensity Estimate. This is the weight a/T placed on the single-factor covariance matrix, as defined by Theorem 2.

It is always between zero and one, which is what we would expect. It is remarkably stable through time. In particular, this implies that there is very little estimation error, as predicted by Theorem 2. It is fairly high: around 80%. This means that there is four times as much estimation error in the sample covariance matrix as there is bias in the single-factor model. In theory a/T is an asymptotically negligible correction, but we see that in practice it makes a big difference.

In spite of the large number of stocks, the computations are remarkably fast. Using a desktop personal computer, it took us about five minutes to compute the optimal shrinkage intensity,

the covariance matrix estimator $\hat{\mathbf{S}}$, and the minimum variance portfolio weights for $N = 1,000$ stocks.

The term r in the formula for the optimal shrinkage intensity, which is one of the major innovations of this paper, turns out to be critical in practice. If we had omitted it, all the shrinkage intensities would have been well above one and thus meaningless.

3 Conclusion

We have developed a flexible method for imposing some structure into a large-dimensional estimation problem. Through a key change in the asymptotic framework, we have been able to derive much more general results than the earlier paper by Ledoit (1996), and at the same time simplify the proofs a lot. The big picture is that every procedure for estimating a large number of parameters should have a dose of shrinkage, and that the method developed here is an appealing way to do it.

This paper also reduces our dependence on multi-factor models, which are surrounded by unresolved questions about the number of factors and their identity. There has been much debate over whether factors should have an economic interpretation or should explain a lot of the variation in stock returns. Ideally, they should do both. By this (admittedly stringent) criterion, there is one obvious factor: the market. We are not saying that extra-market covariance is negligible, but that it lacks strong factor structure. This is precisely why we have developed a way to account for extra-market covariance without fitting it into an arbitrary factor structure.