Journal of Banking & Finance 37 (2013) 3018–3034

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Size matters: Optimal calibration of shrinkage estimators for portfolio selection

Victor DeMiguel [a], Alberto Martin-Utrera [b,*], Francisco J. Nogales [b]

[a] Department of Management Science and Operations, London Business School, United Kingdom
[b] Department of Statistics, Universidad Carlos III de Madrid, Spain

A R T I C L E I N F O

Article history:
Received 29 May 2012
Accepted 22 April 2013
Available online 7 May 2013

JEL classification:
C14
G11

Keywords:
Portfolio choice
Estimation error
Shrinkage intensity
Out-of-sample evaluation
Bootstrap

A B S T R A C T

We carry out a comprehensive investigation of shrinkage estimators for asset allocation, and we find that *size matters*—the shrinkage intensity plays a significant role in the performance of the resulting estimated optimal portfolios. We study both portfolios computed from shrinkage estimators of the moments of asset returns (*shrinkage moments*), as well as *shrinkage portfolios* obtained by shrinking the portfolio weights directly. We make several contributions in this field. First, we propose two novel calibration criteria for the vector of means and the inverse covariance matrix. Second, for the covariance matrix we propose a novel calibration criterion that takes the condition number optimally into account. Third, for shrinkage portfolios we study two novel calibration criteria. Fourth, we propose a simple multivariate smoothed bootstrap approach to construct the optimal shrinkage intensity. Finally, we carry out an extensive out-of-sample analysis with simulated and empirical datasets, and we characterize the performance of the different shrinkage estimators for portfolio selection.

1. Introduction

The classical mean–variance framework for portfolio selection proposed by Markowitz (1952) formalizes the concept of investment diversification, and it is widely used nowadays in the investment industry. To compute mean–variance portfolios, one needs to estimate the mean and covariance matrix of asset returns. One possibility is to replace these quantities with their sample estimators, but these are obtained from historical return data and contain substantial estimation error. As a result, mean–variance portfolios computed from sample estimators perform poorly out of sample; see, for instance, Jobson and Korkie (1981), Best and Grauer (1991), Broadie (1993), Britten-Jones (1999) and DeMiguel et al. (2009).

One of the most popular approaches to combat the impact of estimation error in portfolio selection is to use shrinkage estimators, which are obtained by "shrinking" the sample estimator towards a target estimator.[1] The advantage is that while the shrinkage target is usually biased, it also contains less variance than the sample estimator. Thus it is possible to show under general conditions that there exists a shrinkage *intensity* for which the resulting *shrinkage* estimator contains less estimation error than the original sample estimator; see James and Stein (1961). The key then is to characterize the optimal trade-off between the sample estimator (low bias), and the target (low variance). In other words, shrinkage estimators can help reduce estimation error, but the shrinkage intensity (*size*) matters.

In this paper, we make an extensive investigation of shrinkage estimators for portfolio selection. We study both portfolios computed from shrinkage estimators of the moments of asset returns (*shrinkage moments*), as well as *shrinkage portfolios* obtained by shrinking directly the portfolio weights computed from the original (un-shrunk) sample moments.

* Corresponding author. Tel.: +34 91 624 9859; fax: +34 91 624 9848.
E-mail address: amutrera@est-econ.uc3m.es (A. Martin-Utrera).

[1] Other approaches proposed to combat estimation error in portfolio selection include: Bayesian methods (Barry, 1974; Bawa et al., 1979), Bayesian methods with priors obtained from asset pricing models (MacKinlay and Pastor, 2000; Pastor, 2000; Pastor and Stambaugh, 2000), robust optimization methods (Cornuejols and Tutuncu, 2007; Goldfarb and Iyengar, 2003; Garlappi et al., 2007; Rustem et al., 2000; Tutuncu and Koenig, 2004), Bayesian robust optimization (Wang, 2005), robust estimation methods (DeMiguel and Nogales, 2009), and imposing constraints (Best and Grauer, 1992; Jagannathan and Ma, 2003; DeMiguel et al., 2009).

0378-4266/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2013.04.033

Constructing shrinkage estimators is a three-step procedure. First, define the shrinkage target. Second, choose the calibration criterion that determines the shrinkage intensity. Third, use the available data to estimate the shrinkage intensity that optimizes the calibration criteria. Our work contributes mainly to the last two steps by proposing new calibration criteria, and providing parametric and nonparametric approaches to compute the shrinkage intensity. The shrinkage targets we consider are in general similar to those considered in the existent literature.

We consider three shrinkage estimators of the moments of asset returns. First, we consider a shrinkage estimator of the vector of means similar to those considered before by Jorion (1986) or Frost and Savarino (1986). Unlike these authors, however, we define our estimator a priori as a convex combination of the sample mean and a target element, and we calibrate the shrinkage intensity to minimize the expected quadratic loss—a criterion that distinguishes our work from that of the aforementioned papers. We provide a closed-form expression for the optimal shrinkage intensity under the assumption that returns are independent and identically distributed (iid), but without imposing any further assumptions on the return distribution. Second, we consider the shrinkage covariance matrix proposed by Ledoit and Wolf (2004b), and we implement the same calibration criterion, the expected quadratic loss. Unlike Ledoit and Wolf (2004b), however, we provide a closed-form expression of the optimal shrinkage intensity *for finite samples* by assuming that returns are iid *normal*. Third, we consider a shrinkage estimator of the inverse covariance matrix that is a convex combination of the inverse of the sample covariance matrix and the identity matrix. This estimator is similar to those considered by Frahm and Memmel (2010) and Kourtis et al. (2012), but our contribution is to consider a different calibration criterion for the shrinkage intensity: the expected quadratic loss. Moreover, under iid normal returns, we provide a closed-form expression of the true optimal shrinkage intensity that minimizes the expected quadratic loss. Finally, we propose a new calibration criterion for the shrinkage covariance matrix that takes into account not only the expected quadratic loss but also its condition number. The condition number gives a bound for the sensitivity of the computed portfolio weights to estimation errors in the mean and covariance matrix of asset returns, and thus calibrating the shrinkage covariance matrix so that its condition number is relatively small helps to reduce the impact of estimation error in portfolio selection. Indeed, our experiments with simulated and empirical data demonstrate the advantages of using this criterion for the construction of minimum-variance portfolios.

We investigate three different shrinkage portfolios. The first is obtained by shrinking the sample mean–variance portfolio towards the sample minimum-variance portfolio and it is closely related to the three-fund portfolio of Kan and Zhou (2007); the second is obtained by shrinking the sample mean–variance portfolio towards the equally-weighted portfolio as in Tu and Zhou (2011); and the third is obtained by shrinking the sample minimum-variance portfolio towards the equally-weighted portfolio, similar to DeMiguel et al. (2009). We contribute to the literature by considering, in addition to the utility and variance criteria, two novel calibration criteria: the expected quadratic loss minimization criterion, and the Sharpe ratio maximization criterion. We study the expected quadratic loss criterion because of its good performance in the context of shrinkage covariance matrices (see Ledoit and Wolf, 2004a); and we consider the Sharpe ratio criterion because it is a particular case of the expected utility criterion and it is a relevant performance measure for investors.

For both types of shrinkage estimators, moments and portfolio weights, we propose a multivariate nonparametric smoothed bootstrap approach to estimate the optimal shrinkage intensity. This approach does not impose any assumption on the distribution of asset returns. To the best of our knowledge, this is the first work to consider such a nonparametric approach for shrinkage estimators within the context of portfolio optimization.

Finally, we evaluate the out-of-sample performance of the portfolios obtained from shrinkage moments, as well as that of the shrinkage portfolios on the six empirical datasets listed in Table B.1. For portfolios computed from shrinkage moments, we identify two main findings. First, the shrinkage estimator of the vector of means calibrated with our proposed criterion improves the out-of-sample performance of the resulting mean–variance portfolios. Second, taking the condition number of the estimated covariance matrix into account improves the quality of its shrinkage estimators. For shrinkage portfolios we identify two main findings. First, we find that for those shrinkage portfolios that make use of the sample mean, the best calibration criterion is the portfolio variance minimization criterion. Second, for shrinkage portfolios that ignore the sample mean, the best calibration criterion is to minimize the expected quadratic loss. Finally, for both shrinkage moments and shrinkage portfolios, we find that the nonparametric bootstrap approach to estimate the optimal shrinkage intensity tends to work better than the parametric approach based on normality.

Summarizing, we contribute to the literature of shrinkage estimators for portfolio selection in the following aspects: first, we propose new calibration criteria for shrinkage estimators of moments of asset returns. Second, we consider new calibration criteria for shrinkage portfolios. Concretely, we consider a expected quadratic loss minimization criterion, as well as a Sharpe ratio maximization criterion. Third, we study a multivariate nonparametric approach to compute the optimal shrinkage intensity when returns are iid. Finally, we carry out a comprehensive empirical investigation of shrinkage estimators for portfolio selection on six empirical datasets.

The paper is organized as follows. Section 2 introduces all the considered shrinkage estimators for portfolio selection. Section 3 characterizes the optimal shrinkage intensities when asset returns are iid normal. Section 4 proposes a smoothed bootstrap approach

Table B.1
List of datasets. This table list the various datasets analyzed, the abbreviation used to identify each dataset, the number of assets N contained in each dataset, the time period spanned by the dataset, and the source of the data. The dataset of CRSP returns (SP100) is constructed in a way similar to Jagannathan and Ma (2003), with monthly rebalancing: in January of each year we randomly select 100 assets as our asset universe for the next 12 months.

#	Dataset	Abbreviation	N	Time period	Source
1	5 Industry Portfolios representing the US stock market	5Ind	5	01/1972–06/2009	K. French[a]
2	10 Industry Portfolios representing the US stock market	10Ind	10	01/1972–06/2009	K. French
3	38 Industry Portfolios representing the US stock market	38IndP	38	01/1972–06/2009	K. French
4	48 Industry Portfolios representing the US stock market	48Ind	48	01/1972–06/2009	K. French
5	100 Fama and French Portfolios of firms sorted by size and book to market	100FF	100	01/1972–06/2009	K. French
6	100 randomized stocks from S&P 500	SP100	100	01/1988–12/2008	CRSP[b]

[a] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.
[b] CRSP, The Center for Research in Security Prices.