ARTICLE IN PRESS

J. Int. Financ. Markets Inst. Money xxx (2017) xxx–xxx

Contents lists available at ScienceDirect

Journal of International Financial Markets, Institutions & Money

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

Mean-variance versus naïve diversification: The role of mispricing ☆

Cheng Yan [a,*], Huazhu Zhang [b]

[a] Durham University Business School, Millhill Lane, Durham DH1 3LB, UK
[b] Nomura International Plc, 1 Angel Lane, London EC4R 3AB, UK

ARTICLE INFO

Article history:
Received 29 October 2016
Accepted 21 December 2016
Available online xxxx

JEL classification:
C44
D81
G11
G12

Keywords:
Finance
Portfolio choice
Mean-variance
1/*N* naïve diversification
Mispricing

ABSTRACT

We compare the equal-weight naïve 1/*N* portfolio with mean-variance strategies from the perspective of mispricing (alpha) and provide three new findings. First, we analytically show that the 1/*N* rule approaches the *ex ante* mean-variance efficient portfolio in the absence of mispricing. With mispricings, mean-variance strategies can overcome the difficulty brought by the imprecise parameter estimates and outperform 1/*N* by exploiting the mispricing. Second, with mispricings the 1/*N* rule is unlikely to outperform mean-variance strategies even when *N* is large, since mean-variance strategies have more opportunities to exploit mispricings. Third, minimum-variance strategies do not exploit mispricings and underperform the 1/*N* rule.

1. Introduction

Mean-variance analysis is the cornerstone of modern finance. Markowitz (1952) provides a rigorous framework to consider the risk-return tradeoff, and a methodology to construct optimal portfolios. Although the mean-variance analysis is used pervasively in the academia, the main difficulty in its practical implementations stems from the estimation error or parameter uncertainty problem (Brandt, 2009). Good estimates of the first and second moments are necessary for mean-variance optimization to provide reasonable portfolio weights. An alternative to the mean-variance framework is the naïve equal-weight portfolio investing 1/*N* of total wealth in each of the *N* assets, which can be found in the ancient Babylonian Talmud 1500 years ago and has been observed for individual investors in modern times (Benartzi and Thaler, 2001; Huberman and Jiang, 2006; Brown et al., 2007). The 1/*N* rule does not require parameter estimation and it has been shown that the mean-variance strategies cannot beat the 1/*N* rule in a strand of literature including, among others, DeMiguel et al. (2009), casting doubts upon the practical usefulness of the Markowitz framework.

☆ We are also grateful to Hee-Joon Ahn, Simon Hayley, Stewart Hodges, Xiao Qiao, Yeqin Zeng, and participants in the 2015 Paris Financial Management Conference (PFMC-2015) for helpful comments and discussions. We also thank Victor DeMiguel, Ken French, Lorenzo Garlappi and Raman Uppal for the data. The views expressed herein are those of the authors and do not necessarily reflect or represent those of Nomura.

* Corresponding author.
E-mail addresses: cheng.yan@durham.ac.uk (C. Yan), huazhu.zhang@hotmail.com (H. Zhang).

http://dx.doi.org/10.1016/j.intfin.2016.12.005
1042-4431/

We evaluate the performance of the 1/N rule relative to a broad set of mean-variance strategies and provide three new findings. We present an analytical expression to understand the performance of the 1/N rule. If the Capital Asset Pricing Model (CAPM) holds, for instance, the market portfolio coincides with the *ex ante* tangency portfolio, which has the highest possible Sharpe ratio. With low idiosyncratic volatility relative to market volatility and a large number of assets, the Sharpe ratio of the 1/N rule approaches that of the market portfolio. In this case, the 1/N rule is likely to outperform sample-based mean-variance strategies, which are plagued by estimation errors. Our analytical expression provides an explanation for the excellent performance of the 1/N rule in DeMiguel et al. (2009) without resorting to simulations.

We show that the mean-variance strategies can beat the simple 1/N rule when the CAPM does not hold, even with a large N. Deviations relative to the CAPM (mispricings or alphas) imply the market portfolio is no longer mean-variance optimal. Whereas the Sharpe ratio of the 1/N rule still approaches that of the market portfolio, the mean-variance strategies can exploit the mispricing to form portfolios with higher Sharpe ratios. Holding N constant, for sufficiently large mispricings, mean-variance strategies will outperform the 1/N rule. As the number of assets N increases, there is a tradeoff between precisely estimating the covariance matrix and exploiting mispriced assets. Our simulations show that, given sufficiently large deviations from the cross-sectional asset-pricing model, an increase in the number of securities will cause mean-variance strategies to outperform the 1/N rule. This result overturns the findings in DeMiguel et al. (2009)[1] but is consistent with Huberman and Jiang (2006).[2] Although we use the CAPM as a benchmark model in our analysis, our results hold under more general models including the Fama and French (1992, 1993) and Carhart (1997) models.

Not all mean-variance strategies are able to beat the 1/N rule. Estimation errors in the sample means have a greater influence on the performance of mean-variance strategies than the ones in the sample covariance matrix. As a result, the literature has shifted attention from mean-variance strategies to minimum-variance strategies (Green and Hollifield, 1992; Jagannathan and Ma, 2003; Ledoit and Wolf, 2003; DeMiguel et al., 2009).[3] However, Wang et al. (2015) suggest that it is difficult to find a strategy under the minimum-variance framework that reliably outperforms the naïve 1/N strategy. In our simulations, the 1/N rule consistently outperforms several variations of the minimum-variance portfolio, including the true minimum-variance portfolio based on population moments. This is not surprising, as the minimum-variance portfolios are designed to have the lowest feasible variance, but not necessarily the highest Sharpe ratio.

We confirm our simulation results through an empirical investigation using the size and book-to-market portfolios, the Fama-French factors, and the industry portfolios. Although DeMiguel et al. (2009) find that the mean-variance strategies can hardly beat the 1/N rule, their data is from July 1963 to November 2004, and excludes the Global Financial Crisis (GFC) in the late 2000s during which mispricings may have been the largest. Using an extended sample from July 1963 through December 2015, we find that a number of mean-variance strategies are able to outperform the 1/N rule.

The central intuition for our findings is based on the tradeoff between the exploitation of mispricing and sampling variation in estimated parameters when comparing mean-variance strategies against the 1/N rule. In the absence of mispricing, estimation errors cause the mean-variance strategies to under-perform the 1/N rule. Mispricings provide mean-variance strategies an advantage over the 1/N rule in that mean-variance strategies can benefit from mispricing through intelligently changing the portfolio weights to increase expected returns. This advantage and the disadvantage from estimation errors both increase with the number of investable assets, and the former dominates given sufficiently large mispricings. Such a tradeoff does not apply to minimum-variance strategies, which do not exploit mispricing to increase expected returns. By construction, the minimum-variance portfolios are only concerned about risk and ignore the information from the expected returns.

Our paper most closely relates to DeMiguel et al. (2009), Tu and Zhou (2011) and Wang et al. (2015). DeMiguel et al. (2009) compare the 1/N rule against mean-variance strategies and find that the mean-variance strategies can hardly beat the 1/N rule. We uncover the important role of the zero mispricing in their study with a closed-form expression, and overturn their result that the mean-variance strategies cannot beat the 1/N rule when N is large by introducing deviations from the cross-sectional model. Whereas Tu and Zhou (2011) advocate the better performance of their newly proposed combination rules under non-zero mispricing, we ask if other mean-variance strategies also outperform and investigate the size of mispricing required for outperformance relative to the 1/N rule. Wang et al. (2015) suggest that the minimum-variance strategy cannot outperform the naïve 1/N strategy in a two-asset case when hedging the underlying returns with futures. We extend their asset allocation exercise to more assets and confirm their findings in a more general case.

Our first result from simulations—that, mean-variance strategies can beat the simple 1/N rule in the presence of mispricing—was originally suggested by Tu and Zhou (2011). We include this result here for two reasons. First, it provides a very useful springboard for our two other contributions, namely the analysis of the impact of the number of investable assets (N) and the performance of the minimum-variance strategies. Second, we are able to offer a theoretical reasoning framework

[1] DeMiguel et al. (2009) note "*What is N? That is, for what number and kind of assets does the 1/N strategy outperform other optimizing portfolio models? The results show that the naive 1/N strategy is more likely to outperform the strategies from the optimizing models when: (i) N is large, because this improves the potential for diversification, even if it is naive, while at the same time increasing the number of parameters to be estimated by an optimizing model; (ii) the assets do not have a sufficiently long data history to allow for a precise estimation of the moments.*"

[2] Huberman and Jiang (2006) note in their abstract that "*Records of over half a million participants in more than 600 401(k) plans indicate that participants tend to allocate their contributions evenly across the funds they use, with the tendency weakening with the number of funds used*".

[3] Minimum-variance strategies can be seen as a special case of the mean-variance strategies. For instance, DeMiguel et al. (2009) note "*Also, although this strategy does not fall into the general structure of mean-variance expected utility, its weights can be thought of as a limiting case of Eq. (3), if a mean-variance investor either ignores expected returns or, equivalently, restricts expected returns so that they are identical across all assets*".