In Defense of Portfolio Optimisation What If We Can Forecast?

David Allen[1], Colin Lizieri[2] and Stephen Satchell[3]

Abstract:
We challenge academic consensus that estimation error makes mean-variance portfolio strategies inferior to passive equal-weighted approaches. We demonstrate analytically, via simulation and empirically that investors endowed with modest forecasting ability benefit substantially from an MV approach. An investor with some forecasting ability improves expected utility by increasing the number of assets considered. We frame our study realistically using budget constraints, transaction costs and out-of-sample testing for a wide range of investments. We derive practical decision rules to choose between passive and mean variance optimisation results and generate results consistent with much financial market practice and the original Markowitz formulation.

[1] Pembroke College, University of Cambridge
[2] Department of Land Economy, University of Cambridge: author for correspondence, cml49@cam.ac.uk
[3] University of Sydney Business School and Trinity College, University of Cambridge.

In Defense of Portfolio Optimisation What If We Can Forecast?

1. Introduction

Markowitz's influential *Portfolio Selection* (1952) underpins many advances in financial economics and remains one of the most widely used quantitative approaches to portfolio construction in the industry[i]. Nonetheless, the approach of allocating capital to risky assets by maximising expected return for a given level of risk has been subject to numerous criticisms[ii]. One, that the mean-variance optimisation model requires either normally distributed returns or investor quadratic utility, sits in contradiction both to Markowitz's original formulation and his subsequent research[iii]. Normal distributions are *not* required; it suffices that the distribution be characterized by location and scale[iv]. A second, often attributed to Michaud (1989), is the phenomenon of error maximisation, where return and covariance forecast errors are magnified in the estimated portfolio weights which may lead to poor out-of-sample performance. Given doubts (both normative and positive) about forecast ability and the presence of noise in return series, some authors have concluded that mean-variance is inferior to more passive strategies such as the equal-weighted 1/N approach. This paper examines this choice, exploring the relationship between forecasting ability and investor welfare. We demonstrate, analytically, by way of simulation, and empirically, that mean variance outperforms 1/N with very modest forecasting ability, contrary to views expressed in much of the literature, even accounting for the presence of error. We do not claim that we can eliminate estimation error in any practical way. Nonetheless, our results present a much more compelling case for the use of portfolio optimisation techniques. Our findings have an intuitive logic, to the extent that the implications may seem obvious: if you have good forecasting skill, make use of it, if not, then equal-weighting or risk-weighting a portfolio may be preferable particularly in the absence of good estimation windows for mean and covariance. We provide a formal demonstration of this intuition.

Our analytic approach allows us to contrast mean-variance with 1/N allocation and explicitly identify the drivers of expected utility. We apply the research to a range of realistic investment problems and re-examine the influential model of DeMiguel, Garlappi and Uppal (2009) [DGU hereinafter] which has supported advocacy of 1/N. We extend our analytic results in a high dimension simulation framework (in the spirit of Kane, Kim and White, 2010) which provides strong confirmatory results. We then step away from the theoretical iid normal framework to conduct out-of-sample empirical tests using rolling portfolio rebalancing with short estimation windows and realistically large sets of potential assets. Mean variance performs well in this context.

First, we provide a utility maximising framework that unifies estimation error and forecasting ability effects, building on Grinold's (1989) work relating forecasting ability to the information ratio. Our results show that the Kan and Zhou (2007) result that increasing the number of assets leads to performance deterioration does not necessarily hold when one allows for forecasting, even in the presence of estimation error. Second, we provide a closed form solution for the amount of skill required for mean variance to outperform 1/N *ex ante*. This provides a decision rule allowing investors to choose between approaches. Third, our empirical results show that mean variance can, indeed, outperform 1/N strategies in realistic high dimension, short estimation period windows: exactly the conditions where estimation error effects should be at their most severe. Before proceeding with our analytic, simulation and empirical analysis, we briefly review the existing empirical and analytic literature on mean variance portfolio approaches.

2. *Mean Variance, Errors in Estimation and Forecasting Ability*

The mean variance approach seeks to over-weight assets with low correlations, high expected returns and low relative variance. We note Michaud's (1989) argument that in empirical settings, such assets may be subject to substantial estimation errors, coining the term "error maximisation" for the resultant portfolios, where over-weighting leads to most assets being driven from the portfolio resulting in "corner solutions". As Scherer (2002) notes, with two highly correlated assets, the optimisation algorithm will tend to take long positions in the asset with the higher expected returns and short positions in the other, even though return difference may be within error margins. Numerous studies, have suggested that mean variance does not outperform passive benchmarks, including 1/N approaches, when tested in an out-of-sample context[v]. One important contribution comes from DGU, who compare fourteen models for portfolio allocation (many of which were designed to reduce the impact of estimation error) over seven data sets and conclude that "*none is consistently better than the 1/N rule*", a result that holds using Sharpe ratio, certainty equivalent and turnover measures. They suggest their results cast significant doubt on the utility of such active portfolio investment strategies.

The majority of those studies critical of the mean variance approach take, as their starting point, *ex post* sample estimates as forecasts of expected returns and covariance, often over short time series, despite substantial evidence that sample moments have low predictive power. In Markowitz's original formulation, the proposed E-V rule starts from the *expected* return:

> '*To use the E-V rule in the selection of securities, we must have procedures for finding reasonable m and s. These procedures, I believe should combine statistical techniques and the judgement of practical men* [sic].' (Markowitz, 1952, p91].

Some sixty years after the original formulation, he reiterated this view:

> '*Judgment plays an essential role in the proper application of risk-return analysis for individual and institutional portfolios. For example, the estimates of mean, variance, and covariance of a mean-variance analysis should be forward-looking rather than purely historical.*' (Markowitz, 2010, p7).

A branch of the literature uses longer time series to generate sample estimates in place of short windows. For example, Kritzman, Page and Turkington (2010), using long-term historical averages and covariances as inputs suggest that mean variance generates higher out of sample Sharpe ratios than the 1/N approach[vi]. Furthermore, Kritzman *et al.* demonstrate that virtually any reasonable return forecasts applied to mean-variance optimization would outperform 1/N, and that DGU, by using rolling five year means as forecasts, were, in effect, assuming a negative forecasting ability.

More commonly, a range of models and approaches have been used to reduce the impact of estimation error, including Bayesian shrinkage techniques, resampling and imposing weight constraints[vii]. While some of these techniques have been shown to reduce the impact of estimation error and accord with industry practice, we will not employ them, in order to provide a pure evaluation of the performance of the mean variance approach in the face of errors and noise.

A smaller body of research has examined the performance of mean variance using one-period ahead return forecasts – which is closer in spirit to Markowitz's original proposal. Most of the research combining linear forecasting with mean variance does show outperformance relative to passive benchmarks[viii]. However, typically, such research uses small asset sets (where the effect of estimation error is less extreme). In practical applications, higher dimensions are needed – for example for equity funds benchmarked on the Russell 3000 or MSCI index - and it is here that estimation error will become more problematic. We therefore focus our simulation and empirical work on high dimensional problems where the impact of estimation-error is most severe.

The expected utility of using mean variance with forecasting will be dependent on the quality of those forecasts. Grinold (1989) and Grinold and Kahn (1999) develop expressions for expected utility. The best known is Grinold and Kahn's "fundamental law of active management" that relates the information ratio (IR) - the active return divided by the active risk - to the information coefficient (IC) – the correlation between forecast and realised returns – and the number of stock positions:

$$IR = IC\sqrt{N}$$

They conclude '*it takes only a modest amount of skill to win so long as that skill is deployed frequently and across a large number of stocks*' (Grinold and Kahn, 1999, p162).

At issue, however, is how large the sample period and number of assets must be for active management to outperform passive strategies. Zhou (2008) argues that over 10,000 months of data are required for mean variance to achieve 90% of true maximum utility. DGU argue that for a portfolio containing 25 assets, 3,000 months of estimation data are required for mean variance to outperform 1/N; this rises to 12,000 months for a portfolio of 100 assets. They conclude that this infeasible requirement means that I/N should be preferred to mean variance.

The diametrically opposed conclusions of Grinold and Kahn and DGU result from the differing underlying assumptions. DGU assume estimation error but no forecasting ability. Thus increasing the number of assets increases estimation error effects, damaging expected utility. By contrast, Grinold and Kahn assume forecasting ability but not estimation error. Hence, increasing breadth of investment opportunities increases utility. We seek to reconcile these differences by considering both forecasting ability *and* estimation error in the same modelling approach.

From the 1980s, academics increasingly reported evidence of partial predictability of returns using publicly available information and of anomalies that cast doubt on a pure market efficiency model. The burgeoning literature will be familiar and we will not discuss it in depth here. Nonetheless, documented relationships between dividend yield and equity returns, inflation, interest rates and credit spreads and autocorrelation[ix] led Fama and French (1988) to conclude '*there is much evidence that stock returns are predictable*', their commitment to the efficient market paradigm notwithstanding. It thus seems reasonable to assume *some* forecasting ability. However, those forecasts will be noisy and, hence, estimate error must be considered.

We compare the performance of mean variance with 1/N for a number of reasons. First, since the 1/N rule does not require any estimations, it carries no estimation risk. Second, it is easy

to apply in the market and hence represents a practical strategy for investors. Moreover, others have used 1/N as a basis to benchmark evaluations of mean variance, so we follow in that tradition. Finally, there is behavioural evidence that investors tend to equally weight investment choices in the absence of specific skills: for example, Benartzi and Thaler (2001) document this for US investors choosing defined contribution investment plans. They are not alone. Markowitz, questioned about his own allocations for his retirement investments in his TIAA-CREF account confessed:

> '*I should have computed the historic covariances of the asset classes and drawn an efficient frontier. Instead...I split my contributions fifty-fifty between bonds and equities*'. (p 115, Zweig, 1998)

It is the case that 1/N will be optimal in certain contexts; if the distribution of future returns is independent of current information, and each asset has the same mean, variance and pair-wise correlation, then the mean-variance portfolio will coincide with the 1/N portfolio; in practice 1/N will be very close to optimal for cases not too different from these assumptions. We turn now to consider mean variance utility with both forecasting ability and error estimation.

3. *Mean Variance Utility: An Analytic Approach*

In this section, we set out two models which describe the factors that drive investors' expected utility. Our focus is the interplay between estimation error, forecasting ability and the budget constraint. Our models are based on one-period ahead mean-variance optimisation, which is consistent with commercially available portfolio optimisation software and follows a long line of academic research[x].

For simplicity and consistency with prior research, we assume a constant *IC* across all assets. Previous literature has suggested values of IC in the range zero to 0.1 (which equates to R^2 values of 0% to 1%). This implies very little forecasting ability, which is intuitively problematic, since the mean-variance optimiser treats expected returns as certain. However, it has been pointed out to us by a referee that any individual is likely to have varying *IC*'s depending on their expertise in different markets and sectors. This point favours 1/N investing where only N need vary with the investment set considered.

The starting point of our model is a set of return forecasts, the covariances between those returns and a covariance vector relating forecast to actual returns. In setting up the model, we make several assumptions that are common in the portfolio strategy literature, including independence and normality of forecast errors and that the covariance matrix is known (we relax this assumption in sections 4 and 5). We directly model the budget constraint faced by investors[xi] and use the constrained weights to derive optimum allocations. We confine the mathematical derivations to the appendix (with further details available from the authors). The analysis allows us to present two important models which explore the benefits of a mean-variance strategy for an investor with some forecasting ability but faced with estimation error and the existence of a budget constraint.

Model 1
The unconditional expected utility of the mean-variance investor under the assumptions of multiple forecasts, one for each asset, with a constant forecasting ability level, IC, estimation error, and in the presence of a budget constraint, using the unconditional covariance matrix, is given by

$$E[U] = \frac{\alpha + (N - 1)\left(IC^2 - \frac{1}{T}\right) - \frac{(\beta - \lambda)^2}{\gamma}}{2\lambda} \tag{1}$$

where $\alpha = \mu'\Sigma^{-1}\mu$ refers to the squared Sharpe ratio of the unconstrained mean-variance portfolio, $\beta = \mu'\Sigma^{-1}i$, $\gamma = i'\Sigma^{-1}i$, λ is the level of risk aversion, μ is a vector of returns, Σ is the covariance matrix, i is a vector of ones, and T refers to the length of the estimation window.

This model brings together two strands of the portfolio literature – the impact of estimation error and the role of forecasting while, by incorporating the budget constraint, it is possible to provide a fair comparison with the 1/N allocation rule. Equation (1) shows that expected utility is positively related to the Sharpe ratio of the unconstrained Markowitz portfolio and to forecasting level, IC. Expected utility is also positively related to the length of the estimation window, T, since expected utility increases as estimation error decreases. The relationship between the number of assets, N, and expected utility is, however, more complex. Holding other inputs constant, the impact of increasing the number of assets on utility depends on the level of forecasting ability relative to the number of estimation periods. For example, if forecasting ability is sufficiently large relative to T, utility will increase with an increase in available assets. This helps reconcile the contradictory results of Grinold (1989), where the information ratio increases with number of assets, and DGU whose results suggest expected utility declines with additional assets. [xii]

How much forecast ability is needed for mean-variance to generate superior returns to 1/N? We define the critical level of forecasting ability, IC^*, as the level of forecasting ability where the expected utility of mean-variance is equal to the expected utility of 1/N. If $IC < IC^*$, then the investor would be better off employing 1/N. [xiii]

Setting equation (1) above to the expected utility of the 1/N investor, $V_{1/N}$, and solving for IC, we arrive at the forecasting level required for mean-variance to outperform 1/N. This provides a useful decision rule for practitioners. If their forecasting ability exceeds IC^*, then they should employ mean-variance, otherwise they are better off employing 1/N.

$$IC^* = \sqrt{\frac{2\lambda V_{1/N} - \alpha + \frac{N-1}{T} + \frac{(\beta - \lambda)^2}{\gamma}}{N - 1}} \tag{2}$$

It is worth noting that as the length of our data, T, increases, so IC^* falls and, hence, we need less forecasting skill to beat 1/N. However, IC^* is increasing in risk aversion[xiv], which may push more conservative investors towards 1/N. From Model 1, we are also able to derive an expression for expected utility for a simplified case where volatility and pairwise correlation are constant across assets. This allows us to shed further light on what drives expected utility. Our Model 2 states[xv]:

Model 2
The unconditional expected utility of the mean-variance investor under the assumptions of multiple forecasting variables, a, with a constant forecasting ability level, IC, a constant pairwise correlation, ρ, a constant volatility, σ, a cross-sectional dispersion of mean returns,

σ_u^2, *estimation error, and in the presence of a budget constraint, using the unconditional covariance matrix, is given by*

$$E[U] = \frac{(N-1)\left(\frac{\sigma_u^2}{\sigma^2}(1-\rho) - \frac{1}{T} + IC^2\right) + O(1)}{2\lambda} \tag{3}$$

This result shows that expected utility falls as individual asset volatility and the correlation between assets increases. However, importantly, as the cross-sectional dispersion of expected returns increases, so too does utility. This is consistent with the empirical results of Petajisto (2013). Investors using a strategy based on a portfolio optimiser benefit from heterogeneity in returns as measured by cross-sectional mean dispersion (σ_u^2). This, then, is a key variable for active managers as it provides a measure of investment opportunity.

3.1 Applying the Models

We now use our Model 1 to investigate the expected utility of an investor using a mean variance approach, with some defined forecasting ability, and in the presence of estimation error and a budget constraint.

We compare our results directly with those of DGU's 2009 study, by using the same data sets with the sole exception that we did not have access to their ten S&P sectors. We anticipate that the results would be very similar for the ten Fama-French industries that we do include. When we refer to DGU in this section, we are referring to their analytic model of utility and not to their simulation and rolling portfolio rebalancing work which we discuss in sections 4 and 5. We show that for most plausible assumptions, mean variance allocation outperforms the passive 1/N rule.

The *International* data set includes eight developed market MSCI indices and the MSCI World index. The *Industries* data set includes ten US value-weighted industry portfolios. The *MKT/SMB/HML* data set includes the market, size, and book-to-market long/short factor portfolios. The '*FF-1*' data set includes the 20 portfolios formed by the intersection of the Fama-French size portfolios with the book-to-market portfolios, and the market portfolio, MKT. The *FF-3* data set augments the *FF-1* data set with the size, SMB, and book-to-market, HML, factor portfolios. The *FF-4* data set augments *FF-3* with the momentum factor, *UMD*. We use the same inception points as in DGU

Table 1 shows that mean-variance generates higher utility than 1/N even if forecasting ability is zero for every single data set. These results follow DGU in using truncated data sets for the majority of the tests (omitting almost 40 years of data prior to 1963). Adding these data expands the estimation period and makes the conclusion still more decisive. Our results pertain to a risk aversion of 5. However, our conclusions are robust to different choices.

<Table 1 about here>

We have attempted to reconcile these results with those of DGU. Their analytical approach requires estimates of the squared Sharpe ratios of the 1/N and mean variance portfolios. These estimates are taken from their existing datasets which have between three and 24 assets. However, as they allow the portfolios to grow to up to 100, they leave these estimates constant; the available set of assets has greatly expanded while the distribution of returns

remains unchanged. In practice, were we to add more countries or use more granular definitions of sectors, the Sharpe ratios of the mean-variance and 1/N portfolios would surely change.

A more natural approach would be to use the actual number of assets, the available estimation window, and the observed Sharpe ratios for different investment problems. As Table 2 shows, with such a revised approach, use of the DGU model now shows mean-variance outperforming 1/N whatever the assumption about investor knowledge. The DGU result, then, seems to relate to the specific way they have applied their model.

<Table 2 about here>

Table 3 sets out means, correlations, cross-sectional dispersion, and unconstrained Sharpe ratios for value weighted industry portfolios of increasing granularity, drawn from Ken French data (with observations running from 1926 to 2013). As the number of industries included increases, from five to forty-eight, the mean pairwise correlation falls by 28%, from 0.80 to 0.58 and the cross-sectional dispersion doubles. This result is intuitive (as granularity increases, idiosyncratic differences become more important[xvi]) and casts serious doubt on the validity of assuming a constant multivariate distribution of returns as asset numbers increase. These lower correlations and higher cross-sectional dispersions drive an increase in the mean-variance expected Sharpe ratio from 0.16 to 0.37 and, hence, increase the likelihood of mean variance outperforming 1/N. As shown in Table 3, the DGU model indicates that mean-variance will outperform 1/N in all cases, even though it does not allow for any forecasting ability.

<Table 3 about here>

We acknowledge a limitation of our analytical approach in that, unlike DGU, we ignore the effect of estimation error in the covariance matrix. In theory, however, we can increase the sampling frequency of covariance estimator to reduce estimation error to any arbitrary level. Cochran's (1934) theorem suggests that, under normality, the sample variance based on T observations follows a scaled chi-squared distribution and the standard error of the sample error tends asymptotically to zero. Empirical support for this has been found in equity markets (Andersen, Bollerslev, Diebold and Ebens, 2001).

The principle that high-frequency data can be used to eliminate estimation error in the covariance matrix is limited in practice due to microstructure issues, thin trading, and departures from normality, which all place a limit on precision (Hansen and Lunde, 2006). As a result, lower frequency data in conjunction with factor models tends to be used to estimate the covariance matrix for investment problems involving large numbers of assets.

4. Mean Variance Performance: A Simulation Approach

In this section, we compare the expected performance of mean-variance and 1/N in the presence of forecasting ability using simulation. Simulation allows us to consider cases where closed form solutions for expected utility are not readily available. We relax the assumption that estimation error in the covariance matrix can be eliminated and employ the single-index model of Sharpe (1964) consistent with Treynor and Black (1973), and Kane *et al.* (2010). We also examine the optimal weight relation used in the empirical work of DGU to better

understand the poor performance of the mean-variance approach documented by these authors, together with the Treynor-Black and minimum-variance portfolios.

We replicate Kane *et al.'s* research design, simulating the activity of a hypothetical investment manager. Under this framework, the manager employs N security analysts, each assigned to a single stock. Stock selection problems typically involve higher dimensions, shorter data histories, and lower cross-sectional dispersion than asset allocation problems due to the dominant equity market factor. For these reasons, the stock selection problem should provide a more rigorous test of the benefits of the mean-variance approach.

The simulation is set up as follows. We generate 180 monthly market returns with a monthly excess return of 0.71% and a standard deviation of 4.33%. We randomly generate betas, β_i, for the N analysed stocks from a normal distribution with a mean of one and a standard deviation of 0.3. We generate 180 abnormal returns, $z_{i,t}$, for each stock with a mean of zero and a cross-sectional variance of $\sigma_{z,i}^2$, where $\sigma_{z,i}$ is sampled from a lognormal cross-sectional distribution with μ=1.98, σ=0.36.

Forecasts, z_i^f, of the abnormal returns are generated using:

$$z_i^f = r_{i,0} + r_{i,1}z_i + v_i$$

where v_i is independent of z_i and is normally distributed with a mean of zero and a variance $\sigma_{v,i}^2$, The variance, $\sigma_{v,i}^2$, is a function of the precision of the ith forecaster. It is assumed that $r_{i,0}$=0 and $r_{i,0}$=1; however these parameters are unknown to the investment manager and must be estimated by regression. The level of error variance is set to give the desired level of forecasting ability.

$$\sigma_{v,i} = \sqrt{\left(\frac{(\sigma_{z,i}^2(1 - R_i^2))}{R_i^2} \right)} \tag{4}$$

We consider two skill levels: almost zero skill (IC=0.001), and modest skill (IC =0.07). We consider investment universe sizes of 10, 100 and 500[xvii].

At the end of the 60th month, our investment manager uses the previous 60 months of data to calculate the market mean return and volatility, stock betas, realised abnormal returns, and expected abnormal returns. With these estimates the manager computes the portfolio weights for the mean-variance investor (where we assume a budget-constrained investor with asset weights constrained to sum to one); Treynor-Black weights as in Kane *et al.* which are a constrained variant of the original Treynor-Black model[xviii]; DGU's approach to determining optimal weighting for a mean-variance investor; minimum-variance; and 1/N. We then shift the estimation window forward one month at a time, repeating the exercise until we have 120 months of out-of-sample performance for each model. The process is then repeated 1000 times for each level of forecasting ability and for each universe size, N, enabling us to evaluate the statistical significance of the differences in relative performance

We compute the average mean return, standard deviation, Sharpe ratio and M^2 across simulations. M^2 converts the Sharpe ratio to a measure of risk-adjusted return premium to the market, as proposed by Modigliani and Modigliani (1997)[xix].

4.1 Simulation Results

The results presented in this section make a strong case that mean variance performs better than 1/N for the stock selection problems where there is modest forecast ability. For consistency with prior research, we primarily focus on realised Sharpe ratios as a measure of risk-adjusted return. Table 4, panel A presents the results for an investor with almost zero forecasting ability. 1/N delivers a higher Sharpe ratio than mean-variance whatever the size of the investment universe – a difference statistically significant at the 99% level in all cases. The 1/N rule also outperforms the other allocation models. Consistent, then, with the conclusions of DGU, the findings imply that, in the absence of any forecasting ability, investors are best off adopting the 1/N investment strategy.

The DGU mean-variance portfolios tend to perform poorly. This appears to relate to extreme weightings as noted above and is likely to result from the process by which the authors force the asset weights to sum to unity[xx]. Indeed, in DGU, the authors note that their mean-variance weight in the *International* data set ranges from -148195% to +116828%. This may suggest that the empirical results of DGU could be driven by their portfolio construction method, as also suggested by Kirby and Ostdiek (2012).

<Table 4 about here>

We also examined the relationship between Sharpe ratio and the size of the asset universe for the different strategies. For 1/N, the Sharpe ratio increases with N but plateaus quickly as diversification benefits are realised. While Treynor-Black and minimum variance portfolios also show improving Sharpe ratios, 1/N retains superior risk-adjusted performance. For mean-variance, Sharpe ratio declines with N. Consistent with the analytic model, then, for low levels of forecasting ability, gains from forecasting are offset by increases in estimation error and, in the absence of skill, irrespective of universe size, investors are better off employing the passive investment strategy.

We now consider an investor with modest forecasting abilities, equating to an information coefficient of 0.07 (Table 4 Panel B). Our mean-variance investor now performs at least as well as 1/N for every universe size and level of risk aversion. The difference in Sharpe ratio is statistically significant in all save one case. Now, consistent with our analytical model, the Sharpe ratios of mean-variance increase with the size of our investment universe. Even these very low levels of forecasting ability are sufficient to overcome the additional estimation error induced by increasing the size of the universe. With 500 assets, the mean-variance Sharpe ratios are more than five times larger than those of 1/N. The M^2 of mean-variance portfolios are very high, in excess of 2.5% per month. In context, Malkiel (2013) reports annual mutual fund fees averaging 0.9%.

Our results cast doubt on DGU's conclusion that mean-variance is unworkable in higher dimensions. Grinold and Kahn (1999) develop a simple binary model that relates the information coefficient to the number of forecasts that are directionally correct (the "hit rate"). Our IC of 0.07 equates to a hit rate of just 53.5% a month, or an R^2 of 0.5%. It is remarkable that such modest levels of forecasting ability can generate meaningful gains in utility. The uplift in utility is, tangible and it is a benefit that increases with the size of the asset universe.

To conclude this section, we revisit the impact of estimation error in the covariance matrix. We model two investors: one knows the true population covariance matrix, the other must estimate it using a single index method. Table 5 compares results where the investor estimating the matrix has medium risk aversion and modest forecasting ability (IC=0.07). As can readily be seen, the Sharpe ratios, M^2 and certainty equivalents are very similar irrespective of whether the true covariance matrix is known. By implication the adverse effects of estimation error are largely offset by use of a reasonable factor model. This provides further support for our analytic results where the covariance matrix was assumed to be known.

<Table 5 about here>

The analytic and summary results presented above suggest that substantial gains are available to an investor by using a mean-variance approach provided that the investor has some modest forecasting ability. However, these finding rest on some simplifying assumptions that may be violated in practice. We have assumed that returns are independently normally distributed, that forecasts are independent across stocks and that there are zero trading costs. Moreover, in the simulation results our forecasts are generated using a single index model which is aligned to our data generation process. While these are valuable simplifications for modelling purposes, they may introduce a bias that favours mean-variance. In the next section, we relax those assumptions by moving to an empirical setting where none of those properties can be assumed to hold. If mean-variance still generates superior performance, this will provide still stronger benefits of the approach where investors do possess some forecasting ability.

5. Out of Sample Empirical Evaluation

In this section, we evaluate the performance of mean-variance in an out-of-sample portfolio rebalancing framework that is intended to replicate the problems faced by institutional fund managers. The most comprehensive work in this area is by DGU (2009) who evaluate twelve extensions to mean-variance designed to reduce estimation error for seven data sets. In five of those sets, they find that 1/N generates higher Sharpe ratios and certainty equivalents than the mean variance approach. Given the data span employed, these results cannot be dismissed lightly: DGU argue that some 6,000 months of history would be required for a portfolio of 50 stocks to outperform 1/N. However, given our findings in the previous section, it may be that this result comes from the particular portfolio allocation procedure used. In an international asset allocation context, it has already been shown that mean-variance outperforms passive benchmarks if we condition our forecasts on fundamental variables (Solnik, 1993). We look to extend this work by considering large universes with some 1500 to 3000 stocks to evaluate per month, with covariance matrix estimation windows of as little as 60 months, where estimation error issues are most acute, to reproduce the type of practical problems facing, for example, a small cap manager with a Russell 2000 benchmark and securities with short trading histories. We use 25 years of data across three regions: Asia ex-Japan, Europe and the US.

In our set up, a hypothetical manager makes return forecasts at month end, estimates the covariance matrix using a single index model and derives optimal portfolio weights, rebalancing her portfolio. We acknowledge that superior estimates of risk can be attained using factor models, such as the Fama and French (1993) 3-Factor model (rarely used by practitioners), macroeconomic risk models, or statistical factor models. Fundamental factor models such as those developed by BARRA and Axioma are routinely used by practitioners

to reduce estimation error. The results that we present for mean-variance based on one factor should therefore be conservative. We note also that our models involve calculating a matrix inverse and that the work of Fan et al. (2008) implies that, relative to a k factor model, we incur large losses in accuracy in this context. Rather than choosing a k factor model amongst many possible, we chose the simplest and most parsimonious model to demonstrate the impact of forecast ability. However, a referee has pointed out to us that our single factor specification may be helping our improved performance by reducing estimation error in the covariance matrix; we acknowledge this possibility.

While in previous sections we endowed the investor with a given level of forecasting ability, here our investor develops forecasts using a range of fundamental and price-based variables that had been documented in literature prior to the out-of-sample tests. Given our prior results and its prominence in the finance industry, as an active strategy we focus on mean-variance, comparing it to 1/N and a minimum variance strategy.

We estimate univariate models by regressing the local currency excess returns against a range of variables that have been suggested in the literature. In this, we follow Solnik (1993) who suggests this approach is equivalent to the currency-hedged risk premium with interest-rate parity. We estimate separate models for the three regions, using an expanding window of data with an initial size of 120 months (1990-1999), making out-of-sample prediction more difficult. At the end of each month, we add a further month of data and re-estimate the regression coefficients. As in Moskowitz *et al.* (2012), data are stacked for all stocks and dates and a pooled panel regression is run in each region for each forecasting variable. The forecasting variables are Winsorised at the 5% and 95% level to mitigate the effect of erroneous data:

$$r_{i,t,local} = a_{T+1} + b_{T+1}X_{i,t} + e_{i,t} \tag{8}$$

where $r_{i,t,local}$ is the excess stock return in local currency, and $X_{i,t}$ is the forecasting variable. We then use the coefficient from (8) to generate a one month ahead forecast, $r^f_{i,T+1}$, for each stock in the investment universe. We follow Welch and Goyal (2008) in providing in- and out-of-sample correlations of our forecasting models. We estimate stock betas and idiosyncratic variances using the trailing five years of daily local excess returns and use the FTSE World Europe, FTSE World Asia-Pacific ex Japan and FTSE USA as our benchmark proxies.

Our dataset begins on 1/1990 and ends 12/2014. Our tests use the constituents of the S&P Broad Market indices from the three regions and include all publically listed equities with float-adjusted market values in excess of $100million, and are free from survivorship bias. Table 6 shows summary statistics from the out-of-sample test period from 1/2000 for all stocks with at least 120 months of history. Returns are in local currency and include dividends; the values are consistent with prior literature.

<Table 6 about here>

We use seven individual forecasting variables, drawn from prior literature and conventionally defined[xxi]. Price Momentum (PM) is defined as the twelve month price changes ending one period before estimation (see Moskowitz et al., 2012); Earnings Momentum is the (EM) is defined as the change in IBES consensus forecast earnings per share over the last three months, divided by the current price; Price Reversal (PR) is captured measuring the price

change one month previous; dividend yield (DY), earnings yield (EY) and book to market ratio (BM) are defined conventionally; finally we define return on equity (ROE) as the last reported earnings per share over the book value per share. All Factset Fundamentals data are lagged by three months to eliminate any potential look ahead bias.

For portfolio selection, the objective function seeks to maximise the expected return less risk aversion, multiplied by expected risk and trading costs with weights constrained:

$$\max_{w_t} U_p(w_t) \quad w_t'r_{T+1}^f - \frac{\lambda}{2}w_t'\hat{\Sigma}_{T+1}w_t - |w_t - w_{t,pre}|'TC \tag{9}$$

$$\text{subject to } w_t'i = 1; \ w_{min} \leq w_t \ \leq w_{max}$$

where w_t is the vector of optimal weights, r_{T+1}^f is a vector of return forecasts, λ is the coefficient of risk aversion, $\hat{\Sigma}_{T+1}$ is the estimated single-index covariance matrix, $w_{t,pre}$ are the pre-rebalanced weights, TC is a vector of trading costs, and w_{min} and w_{max} are the minimum and maximum weight vectors. In fuller results, available from the authors, we use three levels of risk aversion to give a range of *ex post* portfolio variances. Here, we report the mid-range risk aversion. We use a 0.5% transaction cost assumption across all stocks both as an optimisation input and to compute net returns – a choice consistent with DGU, Balduzzi and Lynch (1999) and Kirby and Ostdiek (2012). Since we are using local currency excess returns there are no hedging costs to consider. To ensure optimal portfolios are feasible in practice we constrain minimum and maximum weights to:

$$w_{min,i,t} = \frac{-ADV_{i,t}}{Fund\ Size}; \qquad w_{max,i,t} = \frac{ADV_{i,t}}{Fund\ Size}$$

where $ADV_{i,t}$ is the average daily volume of stock i over the previous thirty trading days at month t. We use a fund size of $500 million at each point in time. This ensures that the positions are feasible for a realistic institutional portfolio. Almost all of the existing literature does not account for market liquidity. We feel that this constraint is important to ensure that our results are achievable in practice.

We use a number of metrics to evaluate portfolio performance. We calculate net of trading costs mean excess return; the Sharpe ratio, M^2 and certainty equivalent. To test whether differences between the mean variance and 1/N Sharpe ratios are statistically significant, we use the Jobson-Korkie (1981) t-statistic with Memmel's (2003) correction and, similarly, use Greene (2002) to test for significant difference in certainty equivalent[xxii].

Table 7 provides coefficient estimates, t-statistics and in- and out-of-sample correlations for the seven forecasting models for the three regions. With the exception of the book-to-market ratio in Europe and the reversal variable in Asia, the signs of the coefficients are consistent with literature and significant. The out-of-sample correlations between predicted and observed values are positive for 16 of the 21 cases. However, it is striking how small the correlations are: if mean variance outperforms 1/N it will indicate that only a very modest level of forecasting skill is required. The out-of-sample R statistics are typically higher in Europe and Asia than in the US, suggesting that the latter is more informationally efficient.

<Table 7 about here>

We show the results for a medium level of risk aversion for the three regions in Table 8. Panel A shows results for the Asia ex Japan region. The Sharpe ratios are statistically larger for mean-variance for all bar the book-to-market ratio model, as are the M^2 and certainty equivalent results. These results hold for other levels of risk aversion. The average Sharpe ratio is substantially higher than that of 1/N and the standard deviation is lower, despite the higher gross exposure. We note also that the minimum-variance portfolio (with no forecast involved) outperforms 1/N.

<Table 8 about here>

These findings are echoed in Panel B, showing results for the European region, with significantly larger Sharpe ratios for all models (and all levels of risk aversion). Typically, the Sharpe ratios are three to four times larger than that for 1/N and the other metrics confirm the outperformance of the forecast-based MV strategies. As before, the 1/N strategy has a higher variance than any of the model-based mean variance approaches.

Panel C shows the results for the US models. While the results are similar in form, with the MV models exhibiting higher Sharpe ratios (on average double that of 1/N), there are few statistically significant differences. Closer inspection reveals that this is driven in part by the low correlation between the MV and 1/N strategies, which leads to a higher standard error for the difference in Sharpe ratios. Given that the 1/N strategy maps to the market portfolio, the low correlation provides diversification benefits for the mean variance strategies. As noted above, the out-of-sample correlations of the US model forecasts with observed returns are substantially lower than for the European and Asian models and, consistent with our analytic and simulation results, lower levels of forecasting ability translate into lower levels of utility. Nevertheless, across all seven models and for three different levels of risk aversion, mean variance delivers higher Sharpe ratios in 18 of 21 cases (with only the price momentum model delivering a lower figure). These results, confirming the analytic and simulation findings, stand in marked contrast to the DGU contention of the superiority of 1/N strategies and point to substantial benefits from a forecast-driven MV approach.

6. Conclusions

Much recent finance literature has contended that equal weighted investment strategies are preferable to mean-variance optimisation. The basis for this conclusion seems to rest on two strong assumptions: first, that the investor has no forecasting ability; and, second, that estimation error in the covariance matrix is irreducible. Allowing for forecasting ability is consistent with Markowitz's (1952, 1959) original formulation of portfolio theory and the extensive literature on capital market anomalies.

Our analytic results present a compelling case for the mean-variance portfolio approach, and stand in contrast to De Miguel *et al.* (2009) who argue that vast amounts of data are required for mean-variance to outperform 1/N. We show that only a modicum of forecasting ability is required for mean-variance to outperform 1/N: in many of the investment problems we consider, mean-variance is favoured even in the absence of forecasting ability. We also show that applying the DGU model in the most basic way produces results that favour mean-variance.

In our simulation study we show for the stock selection problem, in the absence of forecasting ability, investor welfare decreases as the size of the investment universe

increases, and 1/N should be preferred to mean-variance. With only a modest level of forecasting ability, this conclusion is reversed, investor welfare increases as the size of the investment universe grows and mean-variance tends to outperform 1/N.

Our empirical findings support the analytic and simulation based results. Even with short covariance matrix estimation windows of 60 months and using simple univariate forecasting models, mean-variance outperforms 1/N across almost all models and regions with in many instances the superior performance being statistically significant.

Overall, our results present a much more compelling case for mean-variance optimised portfolios. DeMiguel *et al.* (2009) conclude that there are 'many miles to go' before the promised benefits of optimal portfolio choice can be realised out of sample. Our findings suggest that we may have already arrived.
.

References

.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens, 2001, The distribution of realised stock return volatility, *Journal of Financial Economics* 61, 43-76.

Balduzzi, P., and A.W. Lynch, 1999, Transaction costs and predictability: Some utility cost calculations, *Journal of Financial Economics* 52, 47-78.

Benartzi, S., and R. H. Thaler, 2001, Naive diversification strategies in defined contribution saving plans, *American Economic Review* 91, 79-98.

Campbell, J. Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373-399.

Campbell, J. Y., and R. J. Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195-228.

.

Chow, G., Jacquier, E., Kritzman, M. and Lowry, K., 1999, Optimal portfolios in good times and bad, *Financial Analysts Journal,* 55:3, 65-73.

Cochran, W.G., 1934, The distribution of quadratic forms in a normal system, with applications to the analysis of covariance, *Mathematical Proceedings of the Cambridge Philosophical Society,* 30, 178-184.

DeMiguel, V., L. Garlappi, and R. Uppal, 2009, Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?, *Review of Financial Studies* 22, 1915-1953.

Fama, E. F., and K. R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3-25.

Fama, E. F., and K. R. French, 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23-49.

Fama, E. F., and K. R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics,* 33, 3-56.

Fisher, K. and Statman, M. (1997) The mean-optimization puzzle: security portfolios and food portfolios, *Financial Analysts Journal,* 53:4, 41-50.

Fleming, J., C. Kirby, and B. Ostdiek, 2001, The economic value of volatility timing, *Journal of Finance* 56, 329-352.

Fleming, J., C. Kirby, and B. Ostdiek, 2003, The economic value of volatility timing using 'Realised' Volatility, *Journal of Financial Economics* 67, 473-509.

Frost, P. A., and J. E. Savarino, 1988, For better performance - constrain portfolio weights, *Journal of Portfolio Management* 15, 29-34.

Grinold, R. C., 1989, The fundamental law of active management, *Journal of Portfolio Management* 15, 30-37.

Grinold, R.C. and R. Kahn, 1999, *Active Portfolio Management* (McGraw-Hill).

Hansen, P. R., and A. Lunde, 2006, Realised variance and market microstructure noise, *Journal of Business & Economic Statistics* 24, 127-161.

Herold, U., and R. Maurer, 2006, Portfolio choice and estimation risk. A comparison of Bayesian to heuristic approaches, *Astin Bulletin* 36, 135-160.

Ingersoll, J.E., 1987. *Theory of financial decision making* (Rowman & Littlefield Publishers).

Jagannathan, R., and T. S. Ma, 2003, Risk reduction in large portfolios: Why imposing the wrong constraints helps, *Journal of Finance* 58, 1651-1683.

Jobson, J. D., and B. M. Korkie, 1981, Performance hypothesis-testing with the Sharpe and Treynor measures, *Journal of Finance* 36, 889-908.

Kan, R., and G. Zhou, 2007, Optimal portfolio choice with parameter uncertainty, *Journal of Financial and Quantitative Analysis* 42, 621-656.

Kane, A., T. Kim, and H. White, 2010, Forecast precision and portfolio performance, *Journal*

of Financial Econometrics 8, 265-304.

Kirby, C., and B. Ostdiek, 2012, It's all in the timing: Simple active portfolio strategies that outperform naive diversification, *Journal of Financial and Quantitative Analysis* 47, 437-467.

Kritzman, M., 2006, Are optimizers error maximizers? Hype versus reality?, *Journal of Portfolio Management* 32, 66-69.

Kritzman, M. and Markowitz, H. (2017) An interview with Nobel laureate Harry M. Markowitz, *Financial Analysts Journal,* 73:4, 16-21.

Kritzman, M., S. Page, and D. Turkington, 2010, In defense of optimization: The fallacy of 1/N, *Financial Analysts Journal* 66, 31-39.

Levy, H., and H. M. Markowitz, 1979, Approximating expected utility by a function of mean and variance, *American Economic Review* 69, 308-317.

Lintner, J., 1975, Inflation and security returns, *Journal of Finance* 30, 259-280.

Lo, A. W., and A. C. MacKinlay, 1988, Stock market prices do not follow random walks: Evidence from a simple specification test, *Review of Financial Studies* 1, 41-66.

Malkiel, B. G., 2013, Asset management fees and the growth of finance, *Journal of Economic Perspectives* 27, 97-108.

Markowitz, H. M., 1952, Portfolio selection, *Journal of Finance* 7, 77-91.

Markowitz, H.M., 1959. *Portfolio selection: Efficient diversification of investments* (John Wiley and Sons, Inc.).

Memmel, C., 2003, Performance hypothesis testing with the Sharpe ratio, *Finance Letters* 1, 21-23.

Michaud, R., 1989, The Markowitz Optimization Enigma: Is 'Optimized' Optimal?, *Financial Analysts Journal* 45, 31-42.

Michaud, R., 1998. *Efficient Asset Management* (Harvard Business School Press, New York).

Modigliani, F., and L. Modigliani, 1997, Risk adjusted performance, *Journal of Portfolio Management* 23, Winter, 45-54.

Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen, 2012, Time series momentum, *Journal of Financial Economics* 104, 228-250.

Petajisto, A., 2013, Active share and mutual fund performance, *Financial Analysts Journal* 69, 73-93.

Poterba, J. M., and L. H. Summers, 1988, Mean reversion in stock-prices: evidence and implications, *Journal of Financial Economics* 22, 27-59.

Rozeff, M. S., 1984, Dividend yields are equity risk premiums, *Journal of Portfolio Management* 11, 68-75.

Scherer, B., 2002, Portfolio resampling: review and critique, *Financial Analysts Journal* 58, 98-109.

Sharpe, W. F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425-442.

Shiller, R. J., 1984, Stock prices and social dynamics, *Brookings Papers on Economic Activity* 457-510.

Solnik, B., 1993, The performance of international asset allocation strategies using conditioning information, *Journal of Empirical Finance* 1, 33-55.

Treynor, J. L., and F. Black, 1973, How to use security analysts to improve portfolio selection, *Journal of Business* 46, 66-86.

Zhou, G., 2008, On the fundamental law of active portfolio management: What happens if our estimates are wrong? *Journal of Portfolio Management* 34, 26-33.

Zweig, J., 1998, Five investing lessons from America's top pension fund, *Money Magazine,* January 1998.

Table 1 - Utility of mean-variance versus 1/N using Model 1

Table 1 indicates whether the expected utility of the mean-variance investor exceeds the expected utility of the 1/N investor for each data set using Model 1. The estimation window, T, refers to the number of months used to estimate the expected return vector. A tick mark signifies that the expected utility of the mean-variance approach exceeds the expected utility of 1/N using Model 1, a cross mark signifies the reverse.

	Internat.	Industry	MKT/SMB/HML	FF-1-Factor	FF-4-Factor
Assets, N	9	11	3	21	24
Estimation window, T	379	497	497	497	497
Information coefficient, IC					
$IC = 0$	✔	✔	✔	✔	✔
$IC = 0.025$	✔	✔	✔	✔	✔
$IC = 0.05$	✔	✔	✔	✔	✔
$IC = 0.075$	✔	✔	✔	✔	✔
$IC = 0.10$	✔	✔	✔	✔	✔

Table 2 - Utility of mean-variance and 1/N using Proposition 1 of DeMiguel et al., 2009

Table 2 indicates whether the expected utility of the mean-variance investor exceeds the expected utility of the 1/N investor for each data set using Model 1 (DeMiguel et al., 2009). $SR_{1/N}$ and SR_{mv} refer to the in-sample monthly Sharpe ratios of the 1/N and mean-variance portfolios. The estimation window, T, refers to the number of months used to estimate the expected return vector. A tick mark signifies that the expected utility of the mean-variance approach exceeds the expected utility of 1/N using the three conditions of DeMiguel et al. (2009).

	S&P Sectors	Internat	Industry	MKT/SMB/HML	FF-1-factor	FF-4-factor
Assets, N	11	9	11	3	21	24
Estimation Window, T	276	379	497	497	497	497
$SR_{1/N}$	0.19	0.13	0.14	0.22	0.16	0.18
SR_{mv}	0.39	0.21	0.21	0.29	0.51	0.54
Conditions						
μ unknown, Σ is known	✔	✔	✔	✔	✔	✔
μ known, Σ is unknown	✔	✔	✔	✔	✔	✔
μ unknown, Σ is unknown	✔	✔	✔	✔	✔	✔

Table 3 - Expected utility of mean-variance versus 1/N using Proposition 1 of DeMiguel et al., 2009

Table 3 indicates whether the expected utility of the mean-variance investor exceeds the expected utility of the 1/N investor for each data set using Model 1 (DeMiguel *et al.*, 2009). $\sigma_{\bar{x}}$ refers to the cross-sectional dispersion in mean returns. $\bar{\rho}_{i \neq j}$ refers to the average pairwise correlation. $S_{1/N}$ and S_{mv} refer to the in-sample monthly Sharpe ratios of the 1/N and mean-variance portfolios. The estimation window, T, refers to the number of months used to estimate the expected return vector and/or covariance matrix. A tick mark signifies that the expected utility of the mean-variance approach exceeds the expected utility of 1/N using the three conditions of DeMiguel *et al.* (2009).

	5 Industries	10 Industries	30 Industries	48 Industries
Assets, N	5	10	30	48
Estimation Window, T	1050	1050	1050	1050
$\sigma_{\bar{x}}$	0.07	0.09	0.12	0.14
$\bar{\rho}_{i \neq j}$	0.80	0.72	0.64	0.58
$S_{1/N}$	0.13	0.14	0.13	0.15
S_{mv}	0.16	0.19	0.24	0.37
Conditions				
μ unknown, Σ is known	✓	✓	✓	✓
μ known, Σ is unknown	✓	✓	✓	✓
μ unknown, Σ is unknown	✓	✓	✓	✓

Table 4: Simulation Summary

Table 4 provides the average mean, standard deviation, Sharpe ratio and M^2 derived through simulation using the algorithm of Kane et al. (2010), over 1000 simulations of 120 months each. Mean-variance refers to the mean-variance weights with a budget constraint of 100%, Treynor-Black† refers to the Treynor-Black model with a margin requirement as in Kane et al. (2010), DGU (2009) refers to the optimal weight relation of DeMiguel et al. (2009), and 1/N refers to the equally weighted portfolio. * (**) indicates the metric is statistically significantly different at the 95% (99%) level.

Panel A: Almost Zero Forecasting Ability
Results for an investor with almost zero forecasting ability (R^2=0.000001, IC=0.001)

	Model	Mean	Standard-deviation	Sharpe ratio	M^2
N=10	Mean-variance: λ=20	0.56**	4.91**	0.11**	-0.25**
	Mean-variance: λ=100	0.54**	4.47**	0.12**	-0.22**
	Treynor-Black†	0.59**	10.64**	0.06**	-0.46**
	DGU (2009)	8.64	4.06E2**	0.01**	-0.67**
	Minimum-variance	0.53**	4.45**	0.12**	-0.22**
	1/N	0.70	5.05	0.14	-0.14
N=100	Mean-variance: λ=20	0.47**	7.64**	0.06**	-0.48**
	Mean-variance: λ=100	0.32**	2.9**	0.11**	-0.27**
	Treynor-Black†	0.67*	6.79**	0.10**	-0.31**
	DGU (2009)	-41.47*	9.69E2**	0.01**	-0.68**
	Minimum-variance	0.28**	2.43**	0.11**	-0.24**
	1/N	0.71	4.4	0.16	-0.04
N=500	Mean-variance: λ=20	0.88**	16.46**	0.05**	-0.51**
	Mean-variance: λ=100	0.35**	3.83**	0.09**	-0.35**
	Treynor-Black†	0.68**	4.93**	0.14**	-0.14**
	DGU (2009)	-14.08*	4.94E2**	0.02**	-0.65**
	Minimum-variance	0.22**	1.61**	0.14**	-0.14**
	1/N	0.71	4.34	0.16	-0.03

Panel B: Modest Forecasting Ability
Results for an investor with medium forecasting ability (R^2=0.005, IC=0.071)

	Model	Mean	Standard-deviation	Sharpe ratio	M^2
N=10	Mean-variance: λ=20	0.85**	5.25**	0.16**	-0.04**
	Mean-variance: λ=100	0.6**	4.52**	0.14	-0.16
	Treynor-Black†	1.25**	10.81**	0.12**	-0.21**
	DGU (2009)	32.27**	6.06E2**	0.05**	-0.50**
	Minimum-variance	0.54**	4.49**	0.12**	-0.21**
	1/N	0.70	5.05	0.14	-0.14
N=100	Mean-variance: λ=20	3.83**	9.69**	0.39**	0.95**
	Mean-variance: λ=100	1.00**	3.17**	0.32**	0.63**
	Treynor-Black†	1.56**	6.35**	0.25**	0.33**
	DGU (2009)	63.9**	8.55E2**	0.14**	-0.11**
	Minimum-variance	0.30**	2.45**	0.12**	-0.21**
	1/N	0.71	4.4	0.16	-0.04
N=500	Mean-variance: λ=20	17.98**	21.53**	0.84**	2.84**
	Mean-variance: λ=100	3.78**	4.78**	0.79**	2.65**
	Treynor-Black†	1.86**	4.92**	0.38**	0.9**
	DGU (2009)	1.25E2**	1.14E3**	0.28**	0.45**
	Minimum-variance	0.23**	1.62**	0.14**	-0.13**
	1/N	0.71	4.34	0.16	-0.03

Table 5 - Simulation summary, covariance matrix is unknown or known

Table 5 provides the average mean, standard deviation, Sharpe ratio, M^2, certainty equivalent, and maximum gross exposure, derived through simulation using the algorithm of *Kane et al.* (2010), for an investor with medium forecasting ability ($R^2=0.005$, IC=0.071), over 1000 simulations of 120 months each, using a λ of 60. . Rows marked 'Unknown' pertain to cases where the covariance is estimated using the single-index model with an estimation window of 60 months. Rows marked 'Known' pertain to cases where the investor knows the true population covariance matrix.

	Covariance matrix	Mean	Standard deviation	Sharpe ratio	M^2	Certainty equivalent	Maximum gross exposure
N=10	Unknown	0.64	4.58	0.14	-0.13	0.11	1.48
	Known	0.63	4.38	0.14	-0.12	0.14	1.39
N=50	Unknown	0.92	3.77	0.24	0.31	0.56	3.04
	Known	0.79	3.36	0.23	0.27	0.50	3.03
N=100	Unknown	1.48	4.06	0.36	0.82	1.06	4.61
	Known	1.29	3.56	0.36	0.82	0.97	4.61
N=500	Unknown	6.15	7.51	0.82	2.77	4.73	16.81
	Known	5.61	6.58	0.86	2.92	4.52	16.68

Table 6 - Summary statistics: Stock returns

Table 6 provides the mean, median, 5[th], and 95[th] percentile for the mean, standard deviation, skewness, kurtosis, minimum, and maximum of monthly stock returns in local currency for the period 1/2000 to 12/2014.

Asia	Mean	Median	5th Percentile	95th Percentile
Mean	1.0	1.0	-0.1	2.4
Standard deviation	10.7	10.1	5.6	17.1
Skewness	0.4	0.2	-0.6	1.9
Kurtosis	7.1	5.1	3.0	15.5
Minimum	-33.4	-31.9	-55.3	-16.7
Maximum	45.4	37.3	17.5	92.8
Europe				
Mean	0.7	0.8	-0.7	1.8
Standard deviation	10.2	9.7	5.8	16.4
Skewness	0.2	0.1	-0.7	1.6
Kurtosis	6.6	5.1	3.3	14.1
Minimum	-33.7	-31.8	-59.2	-16.3
Maximum	42.5	35.8	17.4	90.1
USA				
Mean	1.0	1.0	-0.2	2.2
Standard deviation	12.4	11.3	6.1	21.9
Skewness	0.5	0.2	-0.6	2.3
Kurtosis	7.6	5.2	3.3	18.5
Minimum	-38.7	-36.7	-66.0	-18.2
Maximum	56.2	43.9	19.7	136.8

Table 7 - Predictive regression summaries
Table 7 provides the coefficient, b_{t+1}, t-statistic, and in-sample and out-of-sample R for each variable. The in-sample R is given by the final regression for the period 1/1990 to 12/2014. The out-of-sample R refers to the out-of-sample correlation of the forecasts with excess returns for the period 1/2000 to 12/2014.

Asia ex-Japan	DY	EY	BM	PM	EM	Rev	ROE
Coefficient	0.43	0.37	0.00	0.45	0.55	0.01	0.51
t-statistic	18.44	15.37	0.15	19.68	22.55	0.38	21.23
% positive	100%	100%	27%	100%	100%	63%	100%
R in-sample	0.020	0.023	0.019	0.021	0.031	0.013	0.004
R out-of-sample	0.021	0.024	-0.022	0.024	0.030	-0.002	0.022
Europe	**DY**	**EY**	**BM**	**PM**	**EM**	**Rev**	**ROE**
Coefficient	0.13	0.35	-0.04	0.47	0.47	-0.17	0.36
t-statistic	10.35	26.07	-3.21	37.19	34.57	-13.88	28.42
% positive	98%	100%	66%	100%	100%	0%	100%
R in-sample	0.016	0.025	-0.008	0.033	0.028	-0.002	0.005
R out-of-sample	0.017	0.028	-0.008	0.030	0.030	-0.005	0.023
USA	**DY**	**EY**	**BM**	**PM**	**EM**	**Rev**	**ROE**
Coefficient	0.20	0.59	0.03	0.35	0.44	-0.23	0.53
t-statistic	16.69	48.46	2.53	29.81	34.93	-19.72	44.03
% positive	100%	100%	93%	100%	100%	0%	100%
R in-sample	0.004	0.013	0.004	0.009	0.008	-0.012	0.004
R out-of-sample	0.006	0.016	0.012	-0.005	0.004	0.010	0.010

Table 8: Mean Variance versus 1/N: Summary Performance Statistics

Table 8 provides the out-of-sample monthly summary statistics for the mean-variance portfolios for the period 1/2000 to 12/2014. * (**) denotes statistical significance at the 95% (99%) level. Further results are available in Appendices. For the forecast models, EY is earnings yield, DY is dividend yield, BM is book to market ratio, PM is price momentum, Rev is price reversal, EM is earnings momentum and ROE is return on equity. Results presented are for medium risk aversion ($\lambda = 60$) with further results available in appendices.

Panel A: Asia ex-Japan Region

Model	Gross Return	Net Return	Standard Deviation	Sharpe Ratio	M^2	Certainty equivalent
EY	1.12	1.00	3.88	0.29*	0.50*	0.97*
DY	1.16	1.03	4.04	0.29*	0.50*	1.00*
BM	1.01	0.89	3.91	0.26	0.37	0.86*
PM	1.16	1.02	4.09	0.28*	0.48*	0.99*
Rev	1.06	0.94	3.91	0.27*	0.42*	0.91*
EM	1.30	1.13	3.78	0.34**	0.74**	1.15**
ROE	1.12	0.99	4.02	0.28*	0.46*	0.96*
Min-var.	1.06	0.92	3.91	0.27*	0.43*	0.91*
1/N	0.55	0.49	5.48	0.09	-0.30	0.25

Panel B: European Region

Model	Gross Return	Net Return	Standard Deviation	Sharpe Ratio	M^2	Certainty Equivalent
EY	1.16	1.06	3.60	0.32**	0.64**	1.03**
DY	0.91	0.84	3.31	0.27*	0.44*	0.8*
BM	0.79	0.72	3.35	0.23	0.27*	0.67*
PM	1.11	0.97	4.96	0.22	0.23	0.86
Rev	0.93	0.86	3.44	0.27*	0.42*	0.81
EM	1.51	1.33	4.18	0.36*	0.81*	1.34**
ROE	1.00	0.92	3.39	0.30*	0.53*	0.89*
Min-var.	0.9**	0.82**	3.33	0.27*	0.42*	0.79*
1/N	0.41	0.36	5.32	0.07	-0.40	0.13

Panel C: US Region

Model	Gross Return	Net Return	Standard Deviation	Sharpe Ratio	M^2	Certainty Equivalent
EY	2.13	1.98	7.05	0.30	0.56	1.64
DY	1.10	1.05	4.70	0.23	0.26	0.87
BM	1.07	1.04	3.53	0.30	0.57	0.95*
PM	0.43	0.24	8.88	0.05	-0.52	-0.36
Rev	1.33	1.22	4.70	0.28	0.48	1.11
EM	1.55	1.21	8.07	0.19	0.09	0.89
ROE	1.62	1.49	7.11	0.23	0.24	1.11
Min-var.	1.13	1.08	3.62	0.31	0.6	0.99*
1/N	0.73	0.67	6.22	0.11	-0.23	0.34

What If We Can Forecast?
Portfolio Optimisation, Estimation Error and Forecast Ability.

Appendices:

Appendix A: Analytic Proof of Model 1

Appendix B: Model Parameters and Estimation Inputs

Appendix A: Analytic Proof of Model 1

Assume we have forecasts a (N x 1), of returns r_{t+1} (N x 1), a mean vector, μ (N x 1), a covariance matrix, Σ (N x N), and a covariance vector, C (N x1), between forecast and realised returns. In common with other work in this area (Grinold, 1989), we treat forecasts a as iid normally distributed variables with a mean of zero and a standard deviation of one. For mathematical convenience, we assume that the true covariance matrix is known. To determine the expression for expected utility, we substitute the conditional moments into this optimal weight relation, and then calculate the expected return and risk. Note that in the main text, we used IC to denote the information coefficient, while hereafter we use d to avoid confusion with the multiplication of the identity matrix, I, and the covariance vector, C.

Assuming constant forecasting ability, and given $\bar{x}$, the historical sample mean of returns, the multivariate pdf is given by:

$$
\begin{pmatrix} r_{t+1} \\ a \\ \bar{x} \end{pmatrix} \sim N \left[\begin{pmatrix} \mu \\ 0 \\ \mu \end{pmatrix}, \begin{bmatrix} \Sigma & C & 0 \\ C & I_N & 0 \\ 0 & 0 & \Sigma/T \end{bmatrix} \right] \tag{A1}
$$

The conditional moments of r_{t+1} given a and $\bar{x}$ are given by

$$
\mu^* = \bar{x} + \Sigma^{\frac{1}{2}} da \tag{A2}
$$

$$
\Sigma^* = \Sigma \tag{A3}
$$

where we have replaced μ by $\bar{x}$.

It is typical in the literature to employ an approximation that is akin to the unconditional covariance and we follow this tradition here.

The optimal mean-variance weights in the presence of a budget constraint with known parameters is given by

$$
\omega = \frac{1}{\lambda} \Sigma^{-1} \mu - \frac{(\beta - \lambda)}{\lambda \gamma} \Sigma^{-1} \imath \tag{A4}
$$

where $\beta = \mu' \Sigma^{-1} i, \gamma = i' \Sigma^{-1} i$, and λ is the level of risk aversion.

Proof: see Jorion (1985), p 267 and Ingersoll (1987), Chapter 4, Appendix A, p 68.

To determine the joint impact of estimation error and forecasting ability on expected utility in the presence of a budget constraint, we substitute the conditional moments in (A2) and (A3) into the optimal weight relation as follows.

$$
w = \frac{1}{\lambda} \left(\Sigma^{-1} \left(\bar{x} + \Sigma^{\frac{1}{2}} da \right) - \frac{\left(\bar{x}' \Sigma^{-1} \imath + da' \Sigma^{-\frac{1}{2}} i - \lambda \right)}{\gamma} \Sigma^{-1} \imath \right)
$$

The expected return of the mean-variance portfolio conditional on a and $\bar{x}$ is then

$$
E[r_p | a, \bar{x}] = \frac{1}{\lambda} \left(\Sigma^{-1} \left(\bar{x} + \Sigma^{\frac{1}{2}} da \right) - \frac{\left(\bar{x}' \Sigma^{-1} \imath + da' \Sigma^{-\frac{1}{2}} i - \lambda \right)}{\gamma} \Sigma^{-1} \imath \right)' \left(\mu + \Sigma^{\frac{1}{2}} da \right)
$$

In the above we use the historic mean if we are considering past values or the population mean if we are considering future values. Taking expectations over both conditioning variables, we have

$$E[r_p] = \frac{1}{\lambda}\left((\alpha + Nd^2) - \frac{\beta(\beta - \lambda) + d^2 tr\left(\Sigma^{-\frac{1}{2}}ii'\Sigma^{-\frac{1}{2}}\right)}{\gamma}\right)$$

where $\alpha = \mu'\Sigma^{-1}\mu$ refers to the squared Sharpe ratio of the unconstrained mean-variance portfolio

Now $tr\left(\Sigma^{-\frac{1}{2}}ii'\Sigma^{-\frac{1}{2}}\right) = tr(i'\Sigma^{-1}i)$
$$= tr(\gamma) = \gamma \text{ , so}$$
$$E[r_p] = \frac{1}{\lambda}\left((\alpha + (N-1)d^2) - \frac{\beta(\beta - \lambda)}{\gamma}\right)$$

The expected risk of the portfolio conditional on a is:

$$E[\sigma_p^2|a,\bar{x}] = \frac{1}{\lambda^2}\left(\left(\bar{x} + \Sigma^{\frac{1}{2}}da\right) - \frac{\left(\bar{x}'\Sigma^{-1}i + da'\Sigma^{-\frac{1}{2}}i - \lambda\right)}{\gamma}i\right)'\Sigma^{-1}\left(\left(\bar{x} + \Sigma^{\frac{1}{2}}da\right)\right.$$
$$\left. - \frac{(\bar{x}'\Sigma^{-1}i + da'\Sigma^{-\frac{1}{2}}i - \lambda)}{\gamma}i\right)$$

Taking expectations over both conditioning variables, we have

$$= \frac{1}{\lambda^2}\left(\alpha + \frac{N}{T} + Nd^2 - 2\frac{\left(\beta^2 + \frac{tr(i'\Sigma^{-1}i)}{T} - \beta\lambda + d^2 tr\left(\Sigma^{-\frac{1}{2}}ii'\Sigma^{-\frac{1}{2}}\right)\right)}{\gamma}\right.$$
$$\left. + \frac{(\beta - \lambda)^2 + \frac{tr(i'\Sigma^{-1}i)}{T} - d^2 tr\left(\Sigma^{-\frac{1}{2}}ii'\Sigma^{-\frac{1}{2}}\right)}{\gamma}\right)$$

where T refers to the length of the estimation window.

Now $tr(i'\Sigma^{-1}i)$ and $tr\left(\Sigma^{-\frac{1}{2}}ii'\Sigma^{-\frac{1}{2}}\right)$ both equal γ, thus
$$E[\sigma_p^2] = \frac{1}{\lambda^2}\left(\alpha + (N-1)\left(d^2 - \frac{1}{T}\right) + \frac{\lambda^2 - \beta^2}{\gamma}\right)$$

Hence we have:

Model 1
$$E[U] = \frac{\alpha + (N-1)\left(d^2 - \frac{1}{T}\right) - \frac{(\beta - \lambda)^2}{\gamma}}{2\lambda}$$

For the 1/N investor, the weights are given as follows

$$\omega = \frac{1}{N}\iota$$

where ι is a vector of ones.

Expected utility for this case is then given by

$$V_{1/N} = E\left[U_{w=\frac{1}{N}\iota}\right] = \frac{\iota'\mu}{N} - \frac{\lambda}{2N^2}\iota'\Sigma\iota \tag{A5}$$

which is non-stochastic.

We can now state the following decision rule. The mean-variance approach should be preferred to 1/N, if

$$\frac{\alpha + (N-1)\left(IC^2 - \frac{1}{T}\right) - \frac{(\beta - \lambda)^2}{\gamma}}{2\lambda} - V_{1/N} > 0 \tag{A6}$$

We can now solve for the critical level of forecasting ability required to outperform 1/N.

Corollary 1
The critical forecasting ability, IC^*, required for the mean-variance and 1/N investor to have the same expected utility under the assumptions of multiple forecasting variables, a (M=N), with a constant forecasting ability level, estimation error, and in the presence of a budget constraint, using the unconditional covariance matrix, assuming the term in the square root is positive, is given by

$$IC^* = \sqrt{\frac{2\lambda V_{\frac{1}{N}} - \alpha + \frac{N-1}{T} + \frac{(\beta - \lambda)^2}{\gamma}}{N-1}}$$

Appendix B: Model Parameters and Estimation Inputs

This appendix sets out the formation of model parameters and estimation of inputs for the simulation tests.

First, we generate monthly market returns with a monthly excess return of 0.71% and a standard deviation of 4.33%. We randomly generate our stock betas, β_i, for the N analysed stocks from a normal distribution with a mean of 1.0 and a standard deviation of 0.3. We generate abnormal returns, $z_{i,t}$, with a mean of 0 and a cross-sectional variance of $\sigma_{z,i}^2$, where $\sigma_{z,i}$ is sampled from a lognormal cross-sectional distribution with $\mu=1.98$ and $\sigma=0.36$. We employ a rolling estimation window of 60 months which is common in the literature.

Forecasts, z_i^f, of the abnormal returns are simulated using:

$$z_i^f = r_{i,0} + r_{i,1}z_i + v_i \qquad (B1)$$

where v_i is independent of z_i and is normally distributed with a mean of zero and a variance $\sigma_{v,i}^2$, The variance, $\sigma_{v,i}^2$, is a function of the precision of the i^{th} forecaster. It is assumed that $r_{i,0}=0$ and $r_{i,1}=1$, however these parameters are unknown to the investment manager and must be estimated.

It can easily be seen that $z_i^f \sim N(0, \sigma_{z,i}^2 + \sigma_{v,i}^2)$ and $IC^2 = R_i^2 = [\sigma_{z,i}^2 + \sigma_{v,i}^2]^{-1}\sigma_{z,i}^2$. The level of error variance can then be set to give the desired level of forecasting ability.

$$\sigma_{v,i} = \sqrt{\left(\frac{\left(\sigma_{z,i}^2(1 - R_i^2)\right)}{R_i^2}\right)} \qquad (B2)$$

We consider four skill levels: almost zero skill ($R^2=0.000001$, IC=0.001), low skill ($R^2=0.001$, IC=0.032), medium skill ($R^2=0.005$, IC=0.071), and high skill ($R^2=0.01$, IC =0.1). We consider investment universe sizes, N, of 10, 50, 100, 300, and 500.

B1 Estimation of model inputs

At the beginning of each period, the investment manager uses the preceding T months to estimate the following:

    I.   **Market parameters:** estimated as sample values:

$$\hat{E}(r_{m,T+1}) = \hat{\mu}_{m,T+1} = \frac{1}{T}\sum_{t=1}^{T} r_{m,t}$$

$$\widehat{Var}(r_{m,T+1}) = \hat{\sigma}_{m,T+1}^2 = \frac{1}{T}\sum_{t=1}^{T} (r_{m,t} - \hat{\mu}_{m,T+1})^2$$

    where $r_{m,t}$ is the excess return of the market in period t.

    II.   **Stock betas:** estimated by OLS regression of excess stock returns on excess market returns:

$$r_i = \hat{a}_{i,T+1} + \hat{\beta}_{i,T+1}r_m + \hat{\varepsilon}_i \qquad i = 1, \dots, N$$

    where r_i is the vector of stock returns for stock i (T x 1).

    III.   **Realised abnormal returns:** given by the residuals of the regressions in (II):

$$\hat{z}_{i,t} = r_{i,t} - \hat{\beta}_{i,T+1}r_{m,t}; \qquad i = 1, \dots, N$$

    IV.   **Expected abnormal return and realised skill:** given by regressing realised abnormal returns calculated in (III) on forecast returns:

$$\hat{z}_{i,t} = \hat{c}_{i,0} + \hat{c}_{i,1}z_{i,t}^f + \hat{e}_{i,t}; \qquad i = 1, \dots, N$$

V. **Forecast alphas:** given by adjusting the analysts' forecasts by the level of realised skill estimated in (IV):

$$\hat{\vartheta}_{i,T+1} = \hat{c}_{i,0} + \hat{c}_{i,1} z^f_{i,T+1}; \qquad i = 1, \ldots, N$$

VI. **Covariance matrix:** we estimate the covariance matrix of excess returns using the single-index model of Sharpe (1964):

$$\hat{\Sigma}_{T+1} = \hat{\beta}'_{T+1} \hat{\sigma}^2_{m,T+1} \hat{\beta}_{T+1} + \hat{\Delta}_{T+1}$$

where $\hat{\sigma}^2_{m,T+1}$ and $\hat{\beta}'_{T+1}$ are estimated in (I) and (II) above respectively, and $\hat{\Delta}_{T+1}$ is the diagonal matrix of abnormal variances.

To evaluate the loss in utility from the use of the single-index model relative to when the covariance matrix is known, we also employ the population covariance matrix calculated using the population values for the beta vector, market risk, and the idiosyncratic variance matrix:

$$\Sigma = \beta \sigma^2_m \beta' + \Delta$$

B2 Estimation of model weights

In what follows, when we refer to the mean-variance investor, we are referring to the budget constrained investor where the asset weights sum to unity. This sets a level playing field with the Treynor-Black models, the DeMiguel *et al.* (2009) expression, minimum-variance, and 1/N.

B2.1 Mean-variance weights

The budget constrained mean-variance weights are estimated using as:

$$\omega_{mv} = \frac{1}{\lambda} \Sigma^{*-1} \mu^* - \frac{(\beta^* - \lambda)}{\lambda \gamma^*} \Sigma^{*-1} i$$

where $\Sigma^* = \hat{\Sigma}_{T+1}$, and $\mu^* = \hat{\beta}'_{T+1} \hat{\mu}_{m,T+1} + \hat{\vartheta}_{T+1}$, $\beta^* = \mu^{*\prime} \Sigma^{*-1} i$, $\gamma^* = i' \Sigma^{*-1} i$
We consider risk aversions of 20, 60, and 100 in order to provide a broad range of realised risk levels.

B2.2 DeMiguel *et al.* (2009) mean-variance weights

The optimal mean-variance expression of DeMiguel *et al.* (2009) is given by

$$w = [|i' \Sigma^{*-1} \mu^*|]^{-1} \Sigma^{*-1} \mu^* \tag{B3}$$

where Σ^* and μ^* are as defined above.

B2.3 Treynor-Black weights

The Treynor-Black model, as discussed in Kane *et al.* (2010), consists of two stages. In the first stage, the optimal active portfolio, A, with weights h, is selected to maximise the expected information ratio:

$$\max_h IR_A(h) \frac{h' \hat{\vartheta}_{T+1}}{\sqrt{h' \hat{\Delta}^{-1}_{T+1} h}} \text{ subject to } h'i = 1; \ h^* = [\hat{\vartheta}_{T+1}' \hat{\Delta}^{-1}_{T+1} i]^{-1} \hat{\Delta}^{-1}_{T+1} \hat{\vartheta}_{T+1}; r_A = h^{*\prime} \mu^* \tag{B4}$$

where h^* is the vector of optimal weights, $\hat{\vartheta}_{T+1}$ is the N x 1 vector of expected abnormal returns, and $\hat{\Sigma}_{T+1}$ and μ^* are defined as above.

In the second stage, the overall portfolio, P, is determined by mixing the active portfolio, A, with the market portfolio, M, giving, $r_p = w r_a + (1-w) r_m$, to maximise the expected Sharpe ratio of the overall portfolio:

$$\max_h SR_P(h) = \frac{E(r_p(w))}{\sqrt{var(r_p(w))}}; \quad w^* = \frac{\vartheta_A \sigma^2_{m,T+1}}{(1 - \beta_A) \vartheta_A \sigma^2_{m,T+1} + \hat{\mu}_{m,T+1} h^{*\prime} \hat{\Delta}^{-1}_{t,T+1} h^*}$$

where $\vartheta_A = h^{*\prime}\hat{\vartheta}_{T+1}$, $\beta_A = h^{*\prime}\hat{\beta}'_{T+1}$, and $\hat{\Sigma}_{T+1}$, $\hat{\mu}_{m,T+1}$ and $\hat{\sigma}^2_{m,T+1}$ are as defined above.

The overall portfolio is then given by the vector

$$w_{TB,T+1} = (h^*_{1,T+1}, \dots, h^*_{N,T+1}, 1 - w_{T+1}) \tag{B5}$$

Kane *et al.* (2010), however, do not employ this pure Treynor-Black model, and instead rescale the weights whenever the total short exposure exceeds 200%. Specifically, when $w_{TB,S,T+1} > 2$, the rescaled long and short asset weights are given by:

$$w^*_{TB,L,i,T+1} = \frac{3}{2}|w_{TB,L,T+1}|^{-1}w_{TB,L,i,T+1}; \qquad i = 1, \dots, N_L$$

$$w^*_{TB,S,j,T+1} = \frac{1}{2}|w_{TB,S,T+1}|^{-1}w_{TB,S,j,T+1}; \qquad j = 1, \dots, N_S$$

where $w_{TB,L,T+1}$ and $w_{TB,S,T+1}$ equal the sum of the long and short positions respectively, and $w_{TB,L,i,T+1}$ and $w_{TB,S,i,T+1}$ are elements of the vectors of long and short weights respectively.

This adjustment rescales the total long exposure to 300%, and the total short exposure to 200%, and is designed to mimic a margin constraint. We also examine the performance of the pure Treynor-Black weights in order to determine the effect of the margin constraint.

End Notes

[i] See for example, Fabozzi, Focardi & Jonas, 2007; EDHEC European Investment Practices Survey, 2008.

[ii] As an example, see Fisher and Statman (1997).

[iii] Markowitz (1959) and Levy and Markowitz (1979) show that quadratic approximation provides a reasonable and robust working assumption for a broad range of utility functions and return distributions, but in neither the 1952 nor the 1959 research is it asserted that normality or quadratic utility hold, nor are they requirements for the model. See also Kritzman and Markowitz (2017). An alternative approach broadly within this framework would be to have regime dependent risk aversion, see Chow et al., 1999.

[iv] We note also there are accepted Bayesian approaches to deal with the parameter uncertainty issue.

[v] For example, Jobson and Korkie, 1981; Jagannathan and Ma, 2003.

[vi] Kritzman (2006) also demonstrates that even with substantial misallocation induced by estimation error, there are only relatively small changes in portfolio *ex ante* distributions and, hence, on a limited reduction in expected utility.

[vii] See, *inter alia,* Black and Litterman, 1991; Michaud, 1998; Frost and Savarino, 1998.

[viii] Significant contributions include: Solnik, 1993; Herold and Maurer, 2006; Campbell and Shiller, 1988; Fama and French, 1988; Lintner, 1975; Fama and Schwert, 1977; Campbell, 1987; Poterba and Summers, 1988; Lo and Mackinlay, 1988.

[ix] Ball, 1978; Rozeff, 1984; Shiller, 1984.

[x] Among authors using a myopic agent approach are Fleming, Kirby, and Ostdiek, 2001; 2003; Jagannathan and Ma, 2003; DeMiguel et al. 2009; and Kirby and Ostdiek, 2012.

[xi] DGU divide the unconstrained mean-variance weights by their sum to ensure that the weights satisfy the budget constraint. Our results show that this seemingly innocuous approach has some unusual effects. We instead incorporate the budget constraint directly into the optimisation problem using the result given by Ingersoll (1987).

[xii] We should sound a note of caution: these results hold when alpha, beta and gamma are held constant. However, adding assets to the investment universe will change the return and covariance matrices and it may be that marginal assets may be more prone to forecast errors.

[xiii] Formally:

$$IC^* \equiv inf\{IC: L_{MV}(w^*, \hat{w}) = L_{1/N}(w^*, w^{1/N})\}$$

where $L_y(w^*, \hat{w})$ is the expected utility loss of using weights, $\hat{w}$, instead of the optimal weights w^* based on strategy *y*. Here, we equate the expected utility loss of using weights, w^, instead of the optimal weights w^* based on the different portfolios. We could equally well define it in terms of expected utility.

[xiv] Provided $V_{1/N}$ is positive.

[xv] See Appendix for more details. Full proof available from the authors.

[xvi] As an illustration, the correlation between the "soda" and "gold" industries in the 48 industry classification is just 0.06; their parent sectors in the five industry classification have a correlation of 0.88.

[xvii] Fuller results are available from the authors, which also include further portfolio weighting strategies, asset numbers and risk aversion ratios, confirming our overall conclusions.

[xviii] As shown in the expanded results, the very poor performance of Treynor-Black appears to reflect high volatility and extreme returns driven by extreme weights. This holds whatever the forecasting ability. It was presumably for this reason that Kane et al. imposed constraints that limit long positions to 300% and short positions to 200% (in effect a margin constraint). The Kane *et al.* adjusted model performs better than the pure Treynor Black model but it is problematic that an essentially arbitrary constraint is required to ensure reasonable performance.

[xix] Mathematically, $M^2 = SR_p \sigma_m - \mu_m$ where SR_p is the portfolio Sharpe ratio, μ_m is the mean market return, and σ_m is the market standard deviation. We set μ_m=0.73%, and σ_m=4.27% as in Kane *et al.* (2010).

[xx] DGU force the asset weights to sum to unity by dividing the unconstrained mean-variance weights by the absolute value of the sum of the weights as follows:

$$w = \frac{\Sigma^{*-1}\mu^*}{[|i'\Sigma^{*-1}\mu^*|]}$$

The DGU relation coincides with the tangency portfolio and is independent of the level of risk aversion and it can readily be seen that it can lead to extreme results. It is applied to all their models. The high resultant maximum exposures are stark – results available from authors.

[xxi] Full data sources and prior literature supporting their inclusion are available in the appendices.

[xxii] We acknowledge that improved tests on Sharpe ratios could be accomplished using bootstrapping procedures as in Ledoit and Wolf (2008) although these involve some implementation costs.