THE $1/n$ PENSION INVESTMENT PUZZLE

Heath Windcliff* and Phelim P. Boyle†

ABSTRACT

This paper examines the so-called $1/n$ investment puzzle that has been observed in defined contribution plans whereby some participants divide their contributions equally among the available asset classes. It has been argued that this is a very naive strategy since it contradicts the fundamental tenets of modern portfolio theory. We use simple arguments to show that this behavior is perhaps less naive than it at first appears. It is well known that the optimal portfolio weights in a mean-variance setting are extremely sensitive to estimation errors, especially those in the expected returns. We show that when we account for estimation error, the $1/n$ rule has some advantages in terms of robustness; we demonstrate this with numerical experiments. This rule can provide a risk-averse investor with protection against very bad outcomes.

1. INTRODUCTION

There is evidence (see Benartzi and Thaler 2001) that many participants in defined contribution plans use simple heuristic diversification rules in allocating their contributions among the available asset classes. One popular diversification heuristic is often referred to as the $1/n$ rule. Under this rule the investor divides his or her holdings equally among the available assets. We refer to this portfolio as an *equally weighted portfolio.* This strategy has drawn some criticism since it is not an optimal portfolio, in the sense that, in general, it does not lie on the efficient frontier. In fact, some researchers have suggested that pension plans should offer less flexibility to avert some of the poor investment decisions made by uninformed individuals.

In this paper we demonstrate that the $1/n$ portfolio is consistent with the Markowitz efficient portfolios, given a limited set of information. We then argue that, even with all of the available historical information available to investment professionals, in light of the parameter estimation risk the performance of the $1/n$ heuristic is not unreasonable, assuming an appropriate set of investment choices in the plan. As a result, we feel that it is difficult to justify limiting the flexibility of these plans based on the argument that people using the $1/n$ heuristic are not optimally allocating their pension contributions.

There is an extensive literature on the significance of estimation risk in the context of the portfolio optimization problem. Over 20 years ago Jobson and Korkie (1980, 1981) in a series of papers analyzed the problem and proposed some remedies. Since then several contributions to this topic have been made. In the present paper we use methods that have been described in the finance literature. We emphasize that we make no claims that the techniques used in this paper are original, and almost all the ideas presented here can be found in the literature. However, we feel that the paper serves a useful purpose. First, we present a short self-contained treatment that should be of interest to an actuarial audience unfamiliar with the finance literature. Second, we provide an alternative explanation of the $1/n$ puzzle in the context of the asset allocation decision in a defined contribution plan.[1] Third, the $1/n$ rule is sometimes viewed as naive, and we suggest that this judgment may be somewhat premature.

We begin with a brief overview of the classical mean-variance portfolio theory. To start with, we

* Heath Windcliff currently is employed with Exotic Derivatives, TD Securities, e-mail: hawindcl@elora.math.uwaterloo.ca.

† Phelim P. Boyle is Director of the Centre for Advanced Studies in Finance, School of Accountancy, University of Waterloo, 200 University Ave. West, Waterloo, Ontario, Canada NZL 3G1, e-mail: pboyle@uwaterloo.ca.

[1] Benartzi and Thaler (2001) in their influential paper discuss behavioral explanations of this puzzle.

assume that the true expected return vector and the variance-covariance matrix are known. Under these assumptions we provide an example illustrating that the equally weighted portfolio underperforms the set of optimal portfolios generated by mean-variance optimization. If the investor makes particularly simple specifications of the expected returns and covariances, then we show in Section 4 that the equally weighted portfolio is optimal in a mean-variance sense. This leads us to question how naive this simple specification is for the input parameters of the mean-variance analysis. We find that when we take into account the estimation risk when calibrating these parameters to historical data, this simple specification of the mean-variance input parameters can be viewed as being quite reasonable.

Before proceeding to the main results, two additional comments may be helpful. First, there is an obvious risk in using a single set of parameters to build general conclusions on. Hence in Section 6 we use a modified example to explore the sensitivity of the result to parameter specification. Second, while it is convenient to develop the model using the analytical structure of the standard mean-variance framework, the basic intuition can be obtained from the graphs.

2. Mean-Variance Portfolio Theory

The portfolio optimization methods introduced by Markowitz (1952), termed *mean-variance* optimization methods, signified the birth of modern quantitative finance. The key insight provided by this approach is that an investor can use information describing the relationships between assets to construct a portfolio with better risk-return characteristics than if he or she considers the assets individually. As the name implies, the information that mean-variance portfolio theory uses to describe the relationships between the available assets is the expected returns of the assets, their variances, and correlations. We now summarize the basic approach.[2]

Given a market of N available assets, we define a portfolio to be a vector:

$$\mathbf{x} = (x_1, x_2, \ldots, x_N)^T \quad \text{subject to} \quad \mathbf{e}^T\mathbf{x} = 1, \quad (1)$$

where x_i represents the fraction of wealth held in asset i and where $\mathbf{e} = (1, 1, \ldots, 1)^T$ is the $N \times 1$ vector consisting of N 1's. The constraint $\mathbf{e}^T\mathbf{x} = 1$ represents the budget constraint.

Given a risk-tolerance level,[3] $\tau \in [0, \infty)$, the optimal portfolio can be found by solving the parametric quadratic program

$$\max_{\mathbf{x}} (2\tau\boldsymbol{\mu}^T\mathbf{x} - \mathbf{x}^T\boldsymbol{\Sigma}\mathbf{x}) \quad (2)$$

$$\text{subject to} \quad \mathbf{e}^T\mathbf{x} = 1, \quad (3)$$

where

$$\boldsymbol{\mu} = (\mu_1, \mu_2, \ldots, \mu_N)^T, \quad (4)$$

and μ_i is the expected return on asset i, and

$$\boldsymbol{\Sigma} = [\rho_{ij}\sigma_i\sigma_j], \qquad i, j = 1, \ldots, N, \quad (5)$$

where ρ_{ij} is the correlation between assets i and j, and σ_i is the volatility of asset i. We assume that the matrix $\boldsymbol{\Sigma}$ is positive definite.

There are three different ways to formulate this problem, and they all produce the same solutions. First, we can solve for the portfolio that has the smallest variance among all feasible portfolios with the same expected rate of return. Second, we can solve for the portfolio that has the largest rate of return among all feasible portfolios with the same variance. Third, we can formulate the problem as in equation (3), where τ denotes the tradeoff between expect return and variance. We note that when τ equals zero, the optimal solution is the portfolio that has the global minimum variance.

In this paper we assume that all of the assets are risky. First, we consider the case where there are no restrictions on the portfolio weights. In other words, short selling is permitted. In this case we can solve equations (2)–(3) analytically (see Panjer et al. 1998) to find the optimal portfolio

$$\mathbf{x}^*(\tau) = \mathbf{x}_{min} + \tau\Delta\mathbf{x}_{risk}, \quad (6)$$

where

$$\mathbf{x}_{min} = \frac{1}{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}}\boldsymbol{\Sigma}^{-1}\mathbf{e}, \quad (7)$$

$$\Delta\mathbf{x}_{risk} = \boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} - \frac{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}}\boldsymbol{\Sigma}^{-1}\mathbf{e}. \quad (8)$$

[2] See Panjer et al. (1998) for more details.

[3] See Panjer et al. (1998) for a definition and interpretation of risk tolerance.

The portfolio $\mathbf{x}_{min}$ represents the minimum-variance portfolio in this market, while $\Delta\mathbf{x}_{risk}$ represents a self-financing portfolio adjustment[4] that optimally trades off risk versus reward in this market. Notice that $\mathbf{x}_{min}$ does not depend on the expected returns, while $\Delta\mathbf{x}_{risk}$ depends on both the expected returns and the covariances of the assets.

The expected return on the optimal portfolio is given by

$$\mu[\tau] = \boldsymbol{\mu}^T\mathbf{x}^*(\tau), \tag{9}$$

and the variance of the return on the optimal portfolio is given by

$$(\sigma[\tau])^2 = (\mathbf{x}^*(\tau))^T\boldsymbol{\Sigma}\mathbf{x}^*(\tau), \tag{10}$$

where the optimal weights are obtained from equation (6). As the risk tolerance parameter τ varies, the points $(\sigma[\tau]^2, \mu[\tau])$ trace out the top half of a parabola in mean-variance space. This provides a nice geometric interpretation of the optimal portfolios, and the top half of the parabola is known as the *efficient frontier.* An investor who cares only about expected return and variance will want to hold a portfolio that is on this efficient frontier. The precise location will depend on how he or she trades off expected return and variance, and this will be determined by the investor's risk tolerance, τ.

We can use these explicit solutions to give an intuitive geometric interpretation of the parameter τ. The solution with $\tau = 0$ corresponds to the portfolio with the global minimum variance σ^2_{min} given by

$$\sigma^2_{min} = \frac{1}{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}}.$$

The expected return on this portfolio is

$$\mu_{min} = \frac{\boldsymbol{\mu}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}}{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}}.$$

With a little algebra we can derive the following expression for the expected return on the efficient portfolio:

$$\mu(\tau) = \mu_{min} + \tau\lambda,$$

where the constant λ is given by

$$\lambda = \boldsymbol{\mu}^T\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} - \frac{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}{\mathbf{e}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}}\,\boldsymbol{\mu}^T\boldsymbol{\Sigma}^{-1}\mathbf{e}.$$

As λ and μ_{min} are constants, higher values of τ correspond to higher values of the expected return. We can obtain a similar relation among the variances:

$$(\sigma[\tau])^2 = (\sigma_{min})^2 + \tau^2\lambda.$$

From the last two equations we obtain

$$\frac{\mu(\tau) - \mu_{min}}{(\sigma[\tau])^2 - (\sigma_{min})^2} = \frac{1}{\tau}.$$

This means that in expected return-variance space the slope of the straight line that joins the minimum variance portfolio to the portfolio on the efficient frontier corresponding to τ is equal to 1 over τ. This corresponds to the investor's risk aversion, so its reciprocal, τ, corresponds to the investor's risk tolerance.

Since we will be analyzing the investment decisions of employees in defined contribution pension plans where short sales are not permitted, we impose the condition

$$\mathbf{x} \geq \mathbf{0}. \tag{11}$$

In this case we do not obtain an analytical solution for the optimal portfolio weights. However, we can solve equations (2)–(3) using numerical quadratic programming methods. Details are given by Best and Grauer (1990). When the asset weights are restricted to be non-negative, we rule out some of the solutions that were feasible for the unrestricted case. Thus the efficient portfolio in the case when there is no short selling lies inside (to the southeast of) the efficient portfolio for the case when there is short selling.[5]

In practice the decision maker will not know the true values of the expected return vector $\boldsymbol{\mu}$ and the variance-covariance matrix $\boldsymbol{\Sigma}$. Markowitz (1952) suggested that $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ should be estimated using forward-looking projections, but usually these parameters are estimated from histori-

[4] Self-financing means that the weights sum to zero. It is readily checked that $\mathbf{e}^T\Delta\mathbf{x}_{risk} = 0$.

[5] The efficient portfolio with positive weights may sometimes coincide with the efficient portfolio when short selling is permitted over a certain segment. Best and Grauer (1992) analyze this.

cal data. The historical estimates for $\boldsymbol{\Sigma}$ tend to be somewhat robust, but the estimates for $\boldsymbol{\mu}$ are very noisy, even in a stationary market, as illustrated by Broadie (1993). This estimation risk has important consequences for the optimal portfolio selection problem. Black and Litterman (1992) suggested using reverse optimization and implied expected returns in lieu of the historically estimated returns.

3. MEAN-VARIANCE EFFICIENT PORTFOLIOS FOR A FIVE-ASSET UNIVERSE

We use a five-asset universe to illustrate some of these ideas. We assume that an investor wishes to select an efficient portfolio based on five different asset classes. This mirrors that of a hypothetical pension plan participant where the plan sponsor allows participants to allocate their holdings among five asset classes. To be specific, we assume that the returns on these asset classes correspond to portfolios that track the following indices:

- S&P 500 Large Cap (Large Cap)
- S&P Mid Cap 400 (Mid Cap)
- Russell 2000 Small Cap (Small Cap)
- Morgan Stanley World Equity excluding the United States (World Index)
- Lehman Brothers long-term government bond index (Long Bond Index).

Holden and Van Derhei's (2003) descriptive study of pension plan holdings indicates that these asset choices are broadly representative of those offered to many pension plan participants within their 401k plan.

We assume that the annualized expected returns, volatilities, and correlations of the five assets in the pension plan are given in Table 1. These descriptive statistics were estimated using 15 years of historical data from February 1981 through September 1997. We then adjusted the expected returns of each asset downward by 5% per annum to account for current market conditions. Since this study is meant only to be representative of the conditions that pension plan participants face, we assume that the market parameters given in these tables describe the true distribution of future returns. In other words, the expected returns, variances, and covariances reflected in Table 1 are assumed to *correspond to the true population parameters,* and thus there is no estimation risk. In Section 5 we will investigate the impact of parameter estimation risk on the mean-variance optimal portfolios.

The efficient frontier for portfolios comprised of these assets, assuming that short selling is not permitted, is plotted in Figure 1. This figure shows the optimal tradeoff between reward, measured in units of expected portfolio return, and risk, measured in units of standard deviations. Note that the lowest point on extreme southwest corner of the efficient frontier corresponds to the case when the portfolio is entirely invested in the Long Bond so that the expected return is 5.47%, and the standard deviation is 5.58%. On the other hand, the highest point on the extreme northeast corner of the efficient frontier corresponds to the case when the portfolio is entirely invested in the Mid Cap fund so that the expected return is 12.88%, and the standard deviation is 15.49%. In this figure we also show the risk-reward tradeoff for the equally weighted portfolio. Since the risk-reward tradeoff for the equally weighted portfolio lies below the efficient frontier, it is not optimal in the sense that preferred portfolios exist that simultaneously have greater expected return and less risk, for example, the frontier portfolio with $\tau = 0.2$ that is labeled in Figure 1.

Table 1
Market Data (Assumed True)

Asset Name	μ	σ	ρ				
Large Cap	.1165	.1449	1.0000	0.9180	0.8341	0.4822	0.3873
Mid Cap	.1288	.1549	0.9180	1.0000	0.9386	0.4562	0.3319
Small Cap	.0968	.1792	0.8341	0.9386	1.0000	0.4206	0.2002
World Index	.0921	.1715	0.4822	0.4562	0.4206	1.0000	0.2278
Long Bond Index	.0547	.0558	0.3873	0.3319	0.2002	0.2278	1.0000

Note: Expected returns, volatilities, and correlations for distribution of returns of assets available in the pension plan described in this section.

Figure 1
Risk-Reward Tradeoff

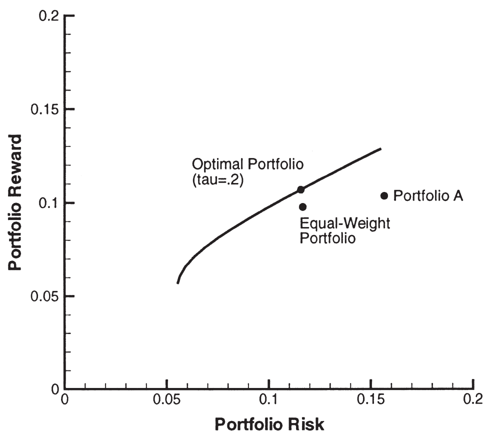

Note: Measured in units of standard deviations of portfolio returns and expected returns, respectively, between portfolios on the efficient frontier and the equally weighted portfolio. Descriptive statistics of the assets available in the pension plan are given in Table 1. Point labeled as Portfolio A will be described in a later section.

4. When Is the Equally Weighted Portfolio Optimal?

Since there is empirical evidence that some investors use the equally weighted portfolio, it is of interest to find a set of simple assumptions that justify the equally weighted portfolio. The motivation here is find the assumptions that would lead to investors selecting the equally weighted portfolio. We show that the equally weighted portfolio is optimal in a very simple market where the assets are indistinguishable and uncorrelated.[6] Note that this implies that all the assets have the same expected return and the same volatility, and that the pairwise correlations are all zero. In this case we can see from symmetry that the $1/n$ portfolio is optimal.

Consider the simple market described above where the asset returns are indistinguishable and uncorrelated. Using a zero subscript to denote this special market, we can write the expected return vector and variance-covariance matrix as

$$\boldsymbol{\mu}_0 = \bar{\mu}\mathbf{e}, \tag{12}$$

$$\boldsymbol{\Sigma}_0 = \bar{\sigma}^2\mathbf{I}, \tag{13}$$

where $\bar{\mu}$ and $\bar{\sigma}$ are the representative levels of the expected returns and asset volatilities, respectively. For any choice of $\bar{\mu}$, and for any nonzero choice of $\bar{\sigma}$,[7] the optimal portfolios, obtained using equations (7)–(8), are given by

$$\mathbf{x}_{min,0} = \frac{1}{N}\mathbf{e}, \tag{14}$$

$$\Delta\mathbf{x}_{risk,0} = \mathbf{0}. \tag{15}$$

In other words, investors who share this particularly simple view of the market should select the equally weighted portfolio for any risk-tolerance level, τ. It is also easy to see that this portfolio is also optimal when short selling is not allowed. The portfolio in this case will consist of a single point in mean standard deviation space.

It is well known that the this simple description of the market is not accurate. Chan, Karceski, and Lakonishok (1999) find that for a sample of 500 U.S. stocks the average correlation is 28%. The lowest correlation was minus 37%, and the highest 92%. Chan et al. also demonstrate that correlation between two stocks is higher when they belong to the same industry than when they belong to different industries. Hence the simple distributional assumptions that lead to the equally weighted portfolio being optimal are not supported by the empirical evidence.

We also saw in Section 3 that the equally weighted portfolio does not lie on the efficient frontier and hence is suboptimal. However, this analysis assumed that the true population parameters are known. In the following section we examine the impact of parameter estimation on the mean-variance optimal portfolios.

5. An Introduction to Parameter Estimation Risk

The efficient frontier computed in Section 3 assumed that we knew the both the expected return vector, $\boldsymbol{\mu}$, and the variance-covariance matrix, $\boldsymbol{\Sigma}$. Although Markowitz (1952) in his original paper suggested using a discount dividend model to specify the expected returns, many practitioners

[6] These are not the most general conditions under which we obtain this result. If the assets all have the same correlation coefficient as well as identical means and variances, we would still find that the equally weighted portfolio is optimal.

[7] This is consistent with our assumption that all of the assets available in the pension plan are risky.

estimate the moments of the future asset return distributions using historical data. A number of papers (Michaud 1998; Broadie 1993; Best and Grauer 1991; Britten-Jones 1999; Bengtsson and Holst 2002) have explored the impact of parameter estimation risk on mean-variance optimal portfolios.

We denote the mean-variance optimal portfolio based on the true population parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ by $\mathbf{x}_{\mathbf{true}}(\tau) = \mathbf{x}_{\mathbf{true}}$. We drop the τ just to simplify the notation, with the understanding that $\mathbf{x}_{\mathbf{true}}$ is still a function of τ. Suppose that the estimates of the expected return and variance-covariance matrix based on a sample of M historical observations are denoted by $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$. Denote the mean-variance optimal portfolio based on these estimated parameters by $\mathbf{x}_{\mathbf{est}}(\tau) = \mathbf{x}_{\mathbf{est}}$. Following Jobson and Korkie (1981) and Broadie (1993), we can distinguish three frontiers:

- *The true frontier.* In this case the expected return and variance are computed as

$$\mathbf{x}_{\mathbf{true}}^{\mathbf{T}}\boldsymbol{\mu}, \qquad \mathbf{x}_{\mathbf{true}}^{\mathbf{T}}\boldsymbol{\Sigma}\mathbf{x}_{\mathbf{true}}. \tag{16}$$

  The true frontier is unattainable since in practice we do not know the true parameters and hence cannot realistically compute $\mathbf{x}_{\mathbf{true}}$.
- *The estimated frontier.* In this case the expected return and variance are computed as

$$\mathbf{x}_{\mathbf{est}}^{\mathbf{T}}\hat{\boldsymbol{\mu}}, \qquad \mathbf{x}_{\mathbf{est}}^{\mathbf{T}}\hat{\boldsymbol{\Sigma}}\mathbf{x}_{\mathbf{est}}. \tag{17}$$

  The estimated frontier is based on the portfolio $\mathbf{x}_{\mathbf{est}}$, which is derived using the estimated weights and hence can be calculated. However, it does not provide a sensible comparison since in reality the future returns will be drawn from the true distribution and not the distribution given by the sample estimates, $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$.
- *The actual frontier.* In this case the expected return and variance are computed as

$$\mathbf{x}_{\mathbf{est}}^{\mathbf{T}}\boldsymbol{\mu}, \qquad \mathbf{x}_{\mathbf{est}}^{\mathbf{T}}\boldsymbol{\Sigma}\mathbf{x}_{\mathbf{est}}. \tag{18}$$

  The actual frontier depicts the performance of the portfolios constructed using the sample estimates, $\mathbf{x}_{\mathbf{est}}$, under the true distribution of returns, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

A representative profile of the relationships among the true, estimated, and actual frontiers is illustrated in Figure 2. In this case and in the example considered by Broadie (1993), the estimated frontier lies above the true frontier,[8] but this need not always be the case. The true frontier must by construction dominate the actual frontier. Although we have depicted a situation where the estimated frontier dominates the true frontier, this need not always be the case. We note that the true frontier is the unattainable ideal, the estimated frontier is illusory, and the actual frontier is the most realistic one for many comparison purposes.

Figure 2
Representative Profile of Relationships among True, Estimated, and Actual Frontiers

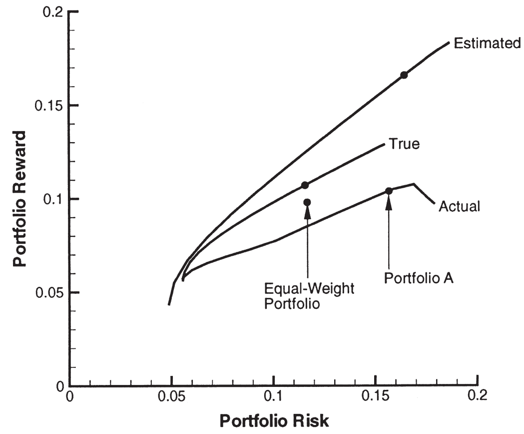

Note: Points on the frontiers correspond to risk tolerance level of $\tau = 0.2$. Estimates of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are given in Table 2.

The simplest way to illustrate this is by using a numerical example. Suppose the asset return distributions given in Table 1 follow a joint multivariate normal distribution. Also suppose that the parameters values given in Table 1 are the true population parameters. Hence the means, volatilities, and correlations given in Table 1 are the population values. Of course, an investor who is attempting to construct the mean-variance optimal portfolio is not able to observe directly the true descriptive statistics of the return distributions. Instead, an investor can observe a sample of returns generated by the assets. For example, suppose that the investor takes a sample of five years of monthly return data (60 observations)

[8] Broadie (1993) found this to be a typical profile based on his analysis.

Table 2
Market Data (Estimated)

	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\rho}$				
Large Cap	.0991	.1468	1.0000	0.9208	0.8266	0.4017	0.5018
Mid Cap	.1602	.1620	0.9208	1.0000	0.9413	0.4188	0.4641
Small Cap	.1830	.1867	0.8266	0.9413	1.0000	0.4386	0.2925
World Index	.0042	.1747	0.4017	0.4188	0.4386	1.0000	0.1863
Long Bond Index	.0440	.0489	0.5018	0.4641	0.2925	0.1863	1.0000

Note: Estimates of expected returns, volatilities, and correlations for asset return distributions obtained from one scenario using a sample of 60 monthly observations.

and uses the sample estimates as proxies to describe the future asset return distributions. Representative sample estimates for a typical numerical scenario are provided in Table 2.

Comparing the results provided in Table 2 with the true parameters from Table 1, we see that although the volatilities and correlations can be estimated with reasonable accuracy, the estimates of the expected returns are very poor. Specifically, we note that the sample returns seriously overestimate the mean for the Small Cap Index and underestimate the mean for the World Index. To see the effect that this will have on the portfolios generated using mean-variance optimization, in Table 3 we compare the true optimal portfolio and the reportedly optimal portfolio constructed using the estimated parameters, both using a risk-tolerance level of $\tau = 0.2$. We find that the portfolio constructed using the estimated parameters has overweighted the Small Cap Index and underweighted the World Index. This demonstrates the error-maximizing property (discussed by Broadie 1993) of the mean-variance optimization procedure. Notice that assets whose sample returns have been optimistically biased are overweighted in the recommended portfolio.

We say that a portfolio dominates another if it has both a higher return and a lower standard deviation. In Table 3 we see that if the investor knows the true market distributions, then the optimal portfolio for $\tau = 0.2$ dominates the equally weighted portfolio. However, we cannot always decide whether the portfolio obtained from mean-variance optimization with $\tau = 0.2$ using the estimated parameters is better than, or worse than, the equally weighted portfolio. This portfolio is labeled as Portfolio A in Figures 1 and 2. We see that it offers a somewhat higher expected return, but it has a higher standard deviation (15.9%) than the equally weighted portfolio (11.7%).

6. THE NUMERICAL EXAMPLE REVISITED

The data for the numerical example were based on actual returns. However, it is potentially dangerous to reach general conclusions from the results of a single example. A referee noted that a special feature of our data set could be viewed as unrealistic and that this feature may drive some of the results.[9] In this section we discuss this point and make some additional calculations to assess the robustness of the performance of the equally weighted portfolio to the input parameters.

The input population parameters for the numerical example are given in Table 1. We generated a sample of 60 monthly returns for asset class from the true population distribution and used the sample to obtain estimates for the expected return vector $\hat{\mu}$ and the variance covariance vector $\hat{\Sigma}$. Then we computed the optimal portfolio weights, $\mathbf{x}_{\mathbf{est}}$, using $\hat{\mu}$ and $\hat{\Sigma}$. These estimated weights were given in Table 3 for $\tau = 0.2$. Note that the amount allocated to the Small Cap asset class in this optimization is 60.81%. However, the true expected return on this asset class is 9.68%, which means that it is out of line with the return profiles of the other major equity indices, especially the Large Cap Index and the Mid Cap Index, both of which have lower standard deviations and higher expected returns. In addition, all three indices are quite strongly correlated, so the Small Cap Index affords only limited

[9] We thank the referee for insightful comments on this point.

Table 3

Portfolios and Portfolio Performances Under True Market Data

Portfolio	True Market μ	True Market σ	Asset Holdings: Large Cap	Mid Cap	Small Cap	World Index	Bond Index
x_{true}	.1070	.1156	.0000	.6796	.0000	.0529	.2675
x_{est}	.1048	.1590	.0000	.3311	.6081	.0000	.0608
Equal weight	.0978	.1167	.2000	.2000	.2000	.2000	.2000

Note: Expected returns, volatilities, and holding of the true optimal portfolio for risk tolerance of $\tau = 0.2$, the actual portfolio obtained by mean-variance optimization with $\tau = 0.2$ using estimated parameters reported in Table 2, and the equally weighted portfolio. Expected returns and volatilities given in columns 2 and 3 are computed using the true population parameters of asset return distributions specified in Table 1.

diversification benefit. Because of this, the Small Cap Index is not held in the true efficient frontier expect for a very small percentage holding when τ is zero or very close to zero.

However, the sampled data produce a relatively high expected return for the Small Cap Index, 18.30%. This high expected return means that the allocation to the Small Cap Index is very high, since $\mathbf{x}_{est}(3)$ is over 60%. When we realize that the true efficient portfolio hardly ever includes this asset class, we see the source of the deviation from optimality. Our input assumptions concerning the expected return on the Small Cap Index drive this conclusion. It is therefore desirable to modify the example to see how the equally weighted portfolio performs under other, more realistic, assumptions.

Hence we redo the example using revised input assumptions by modifying those in Table 1. Specifically we increase the expected return on the Small Cap Index from 9.68% by 4% to 13.68%. All the other input assumptions in Table 1 are unchanged. Note that under this revised assumption the Small Cap Index now has the highest expected return as well as the highest volatility. It is convenient to use graphs to summarize the performance of the equally weighted portfolio under these revised assumptions. In Figure 3 the true efficient portfolio is shown as the topmost curve. The equally weighted portfolio is marked with a plus sign on the graph. We also generated 100 samples of simulated data, where each sample consisted of 60 months' returns. For each sample we computed the sample expected return vector $\hat{\mu}$ and the sample variance covariance vector $\hat{\Sigma}$. Using these inputs we computed the optimal portfolio weights to generate the actual efficient frontier. Armed with these weights we use the true population expected returns and covariance matrix to obtain the actual frontier. In this way we obtain 100 actual frontiers. The lower curve in Figure 3 represents the average of these 100 actual frontiers. We note that the equally weighted portfolio lies very close to this average frontier. This result would seem to suggest that the equally weighted portfolio does not offer much of an advantage in this particular example.

However, there is another feature of the equally weighted portfolio that may be appealing to some investors. A risk-averse investor may be concerned about avoiding the really bad outcomes. In Figure 4 we depict the four worst outcomes of the 100 frontiers. We see that the equally weighted portfolio now performs well relative to these bad outcomes. In this example the equally weighted portfolio does not dominate the average outcome, but it does perform better than the worst case outcomes.

7. PARAMETER ESTIMATION ERROR

In the preceding section we saw that portfolios generated using mean-variance optimization can lie quite far away from the efficient frontier when there is parameter estimation error. In this section we provide a statistical comparison of the performances of mean-variance optimal portfolios with parameter estimation error and the equally weighted portfolio. Assuming that the true market parameters are as given in Table 1,[10] we generate scenarios of market data from which we obtain estimates of the market parameters. Using these estimated parameters we generate

[10] Note that we used the original market parameters here rather than the revised ones in Section 6. This is reasonable since our main focus here is on the relationship between the length of the historical window and the precision of the estimates.

Figure 3
Performance of Equal Weight Portfolio Versus Average at Optimized Portfolios Using Estimated Market Data

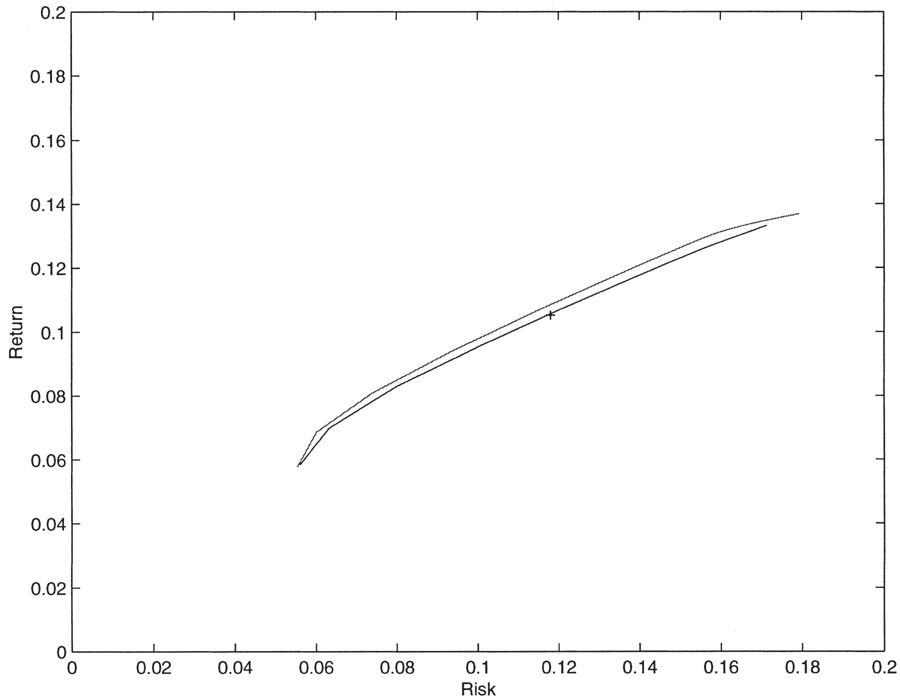

Note: Top curve denotes true frontier based on true population parameters. Lower curve represents average of 100 actual frontiers based on 100 simulated samples. Equally weighted portfolio is denoted by + sign. Values of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the same as those given in Table 1 except that expected return on the Small Cap Index is 4% higher.

mean-variance optimal portfolios and then study their performance under the true market parameters.

In Figure 5 we investigate the mean-variance optimal portfolios computed using estimated parameters for risk tolerances of $\tau = 0$ and $\tau = 0.2$. In the first panel when $\tau = 0$, we see that risk-reward tradeoffs of the estimated portfolios, $\mathbf{x}_{est}$, are all close to the true global minimum-variance portfolio. This illustrates that the global minimum-variance portfolio can be approximated quite well with the estimated parameters using five years of data (60 monthly samples).

For investors who are taking on risk, as in the case when $\tau = 0.2$ shown in panel b, the mean-variance optimal portfolios computed using estimated parameters can lie far away from the efficient frontier when only five years of data are used. In fact, in panel b we find that there are many cases where the equal-weight portfolio dominates the reportedly optimal portfolios generated using estimated parameters. This is because the equally weighted behavioral portfolio does not require the specification of the market parameters using historical data and hence does not suffer from parameter sampling risk.

If the market parameters are stationary, we can reduce the parameter estimation risk by sampling a longer history of asset returns. In panel c we use 20 years (240 monthly samples) of observed data to estimate the expected returns and covariances and see that portfolios now lie much closer to the efficient frontier. It is very unlikely that the market parameters would remain stationary over a 20-year window, making this brute force approach inappropriate.

It has been noted by Best and Grauer (1991) and Broadie (1993) that, comparatively speaking, it is more difficult to obtain reliable estimates of the expected returns than it is to obtain reliable estimates of the variance-covariance matrix using historical data. We illustrate that it is the mis-

Figure 4
Performance of Equal Weight Portfolio Versus Worst of Optimized Portfolios Using Estimated Market Data

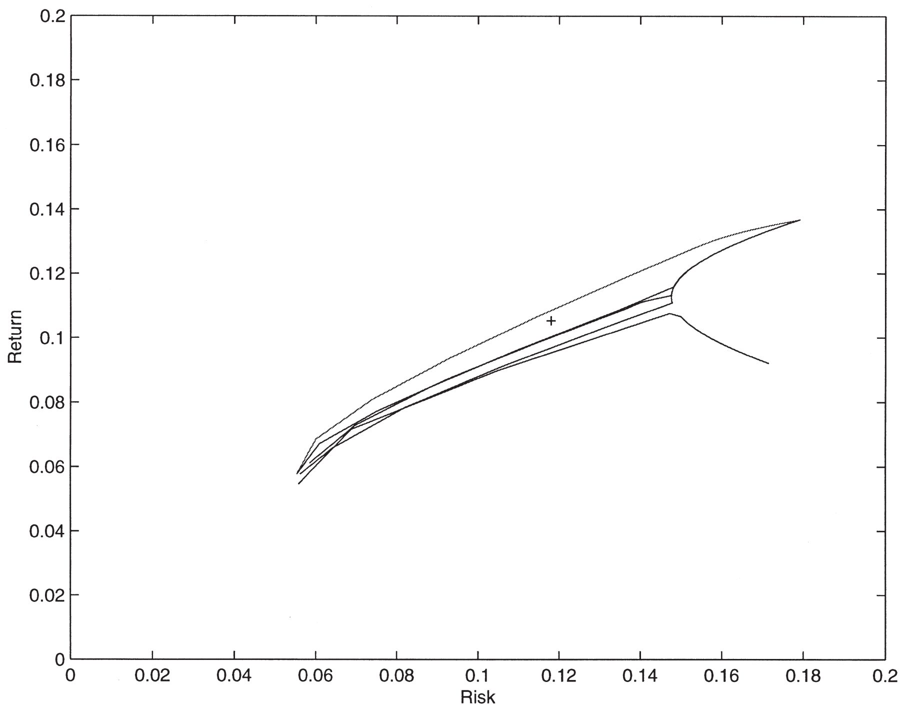

Note: True frontier is topmost curve, and the other four curves represent four worst actual frontiers based on 100 simulated samples. Equally weighted portfolio is denoted by + sign. Estimates of μ and Σ are the same as those given in Table 1 except that expected return on the Small Cap Index is 4% higher.

specification of the means that causes the actual performances of the estimated portfolios to lie off of the efficient frontier in panel d. In this panel we show that the computed portfolios lie very close to the efficient frontier if the vector of expected returns is assumed to be known, while estimating the variance-covariance matrix using only five years of data. This implies that investors should take extra care to specify accurately their views of the expected returns on the assets. We will discuss a procedure for improving the estimates of expected returns in Section 8.

In summary, we have observed that, at least for the true market parameters studied in this paper, using a reasonable amount of historical data (five years), the performance of the equally weighted portfolio is comparable with the performance of the mean-variance optimal portfolios once we account for the impact of parameter estimation risk. Simply using a longer window of historical samples is not a realistic solution to this problem because of the nonstationarity of the market parameters. We also demonstrated that the misspecification of the expected returns has the most impact on the performance of the mean-variance optimized portfolios.

8. Using Shrinkage Estimators to Improve Estimates of Returns

We have seen that the misspecification of the expected returns sometimes has a dramatic negative impact on the performance of the mean-variance optimized portfolios. In this section we use a Bayesian shrinkage model suggested by Jorion (1986) to quantify how much useful information is contained in a sample of historical returns. This, in turn, will allow us to quantify the amount of information that is disregarded by investors who utilize the simple, equally weighted diversification heuristic.

Figure 5

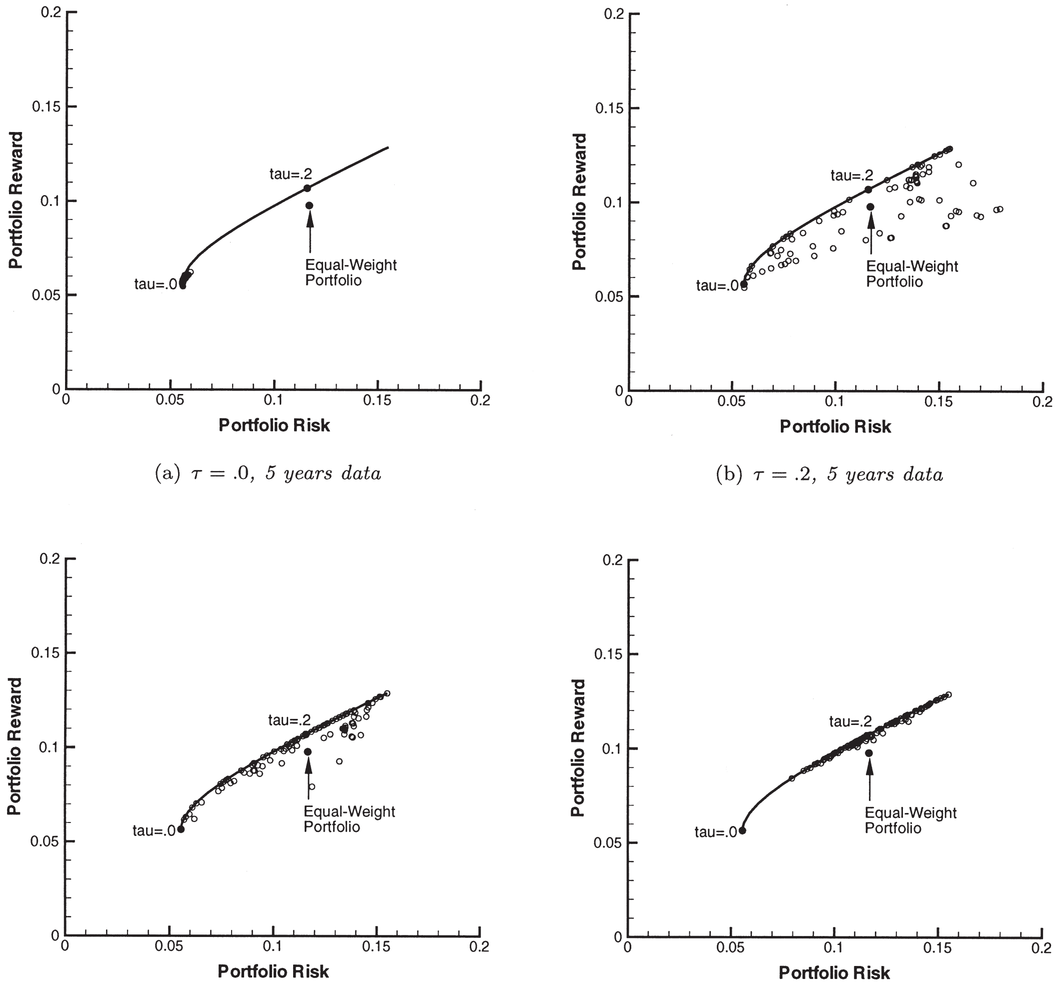

(a) $\tau = .0$, *5 years data*

(b) $\tau = .2$, *5 years data*

(c) $\tau = .2$, *20 years data*

(d) $\tau = .2$, *5 years data, Expected returns known*

Note: Actual risk-reward tradeoff under true market parameters, measured in units of standard deviations and expected returns, respectively, of a sample of 100 mean-variance optimal portfolios computed using parameters estimated from a history of observed returns.

In Section 4 we defined a market

$$\boldsymbol{\mu}_0 = \bar{\mu}\mathbf{e}, \tag{19}$$

$$\boldsymbol{\Sigma}_0 = \bar{\sigma}^2\mathbf{I}, \tag{20}$$

in which the assets are statistically indistinguishable. In this market the optimal risk-reward tradeoff is obtained with the equally weighted portfolio, a result that is independent of the specified average return $\bar{\mu}$ and average volatility $\bar{\sigma}$. At the other extreme we have an investor who uses a sample of historical returns to estimate the relevant market parameters, $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$. We have seen that our estimate of the expected returns, $\hat{\boldsymbol{\mu}}$, can be particularly noisy and, due to the error-maximizing property of the mean-variance optimization, can result in poor portfolio selection.

For simplicity we assume that the sample variance-covariance matrix, $\hat{\boldsymbol{\Sigma}}$, provides an accurate estimate of the true covariances. We focus on quantifying the benefits of including information contained in a sample of historical data when specifying the expected return vector. Consider the Bayes-Stein weighted learning model described by Jorion (1986):

$$\hat{\boldsymbol{\mu}}_{BS} = \boldsymbol{\mu}_0 + \alpha(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_0), \tag{21}$$

Table 4
Mean and Standard Deviation of Weighting Parameter α

Months of Data Used	α Mean	α Std Dev.
60	0.4995	0.1611
240	0.7028	0.1048

Note: Defined as equation (22) using 5 years and 20 years of historical data. Statistics were estimated from a sample of 1,000 scenarios.

where

$$\alpha = 1 - \frac{N+2}{N+2+M(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_0)^T \boldsymbol{\Sigma}^{-1}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_0)}, \quad (22)$$

and where N is the number of assets and M is the number of historical data points in the sample. The parameter, $\alpha \in [0, 1]$, specifies the relative importance of the prior information, $\boldsymbol{\mu}_0$, and the information contained in the historical sample, $\hat{\boldsymbol{\mu}}$. Note that when $M \to \infty$, $\alpha \to 1$ and all the weight is placed on the sample estimate $\hat{\mu}$ from the historical data. The estimator in equation (21) is often called a shrinkage estimator because the estimated expected returns tend to $\boldsymbol{\mu}_0$ as $\alpha \to 0$. Although it is possible to consider other priors, in this paper the prior, $\bar{\mu}$, is chosen to be the mean of the returns on the minimum-variance portfolio:

$$\bar{\mu} = \frac{\hat{\boldsymbol{\mu}}^T \boldsymbol{\Sigma}^{-1} \mathbf{e}}{\mathbf{e}^T \boldsymbol{\Sigma}^{-1} \mathbf{e}}. \quad (23)$$

One reason for selecting this prior is that the weights on the minimum variance portfolio depend only on the variance covariance matrix $\hat{\boldsymbol{\Sigma}}$ and do not depend on the expected returns. In practice $\hat{\boldsymbol{\Sigma}}$ is much more stable.

As an aside we note the resemblance between the Bayes-Stein estimator in equation (21) and the linear credibility approach for estimating insurance premiums.

In our case we are interested in determining the confidence level, α, in a set of historical data. If α is consistently low, then there is little information contained in the sample mean, and investors who establish equally weighted portfolios are, in fact, doing something that is somewhat reasonable.[11] In Table 4 we show the mean and standard deviation of the weighting parameter, α, for 60 and 240 months of historical data. Using five years of data, the Bayes-Stein estimator places approximately equal weight on the prior and the sample means, indicating that there is surprisingly little useful information about the expected returns contained in the sample means of the historical data, even if the parameters are stationary. As more historical data are used, the quality of the sample mean improves and α increases, but, of course, the benefits will be spurious if the market parameters are not stationary.

In Figure 6 we plot the risk-reward tradeoff for a sample of portfolios constructed using the sample means and those constructed using the Bayes-Stein estimators. We find that more of the portfolios lie along the efficient frontier when the shrinkage estimators are used.

9. CONCLUSIONS

In this paper we have explored some of the rationale behind the equally weighted portfolio that is popular with some defined contribution pension plan participants. If the parameters of the return distribution are assumed to be known, this simple heuristic leads to portfolios that lie below the efficient frontier and consequently are suboptimal. We show that the equally weighted portfolio is optimal in a market where the assets are indistinguishable and uncorrelated. This leads us to question how much investors lose by implicitly specifying the market parameters in this simple manner.

In practice we do not know the true future asset return distributions, and often these are estimated using samples of historical data while assuming that the parameters are stationary. If we take into account the impact of parameter estimation risk, then the computed portfolios using mean-variance optimization no longer dominate the equally weighted portfolio. Using a Bayes-Stein learning model we show that, using a sufficiently small window of historical data over which we might hope that the parameters are stationary, little weighting is placed on the sample

[11] However, as pointed out by Brennan and Torous (1999), they are still missing out on some of the diversification benefits that can be obtained by using the correlations between assets, which can be estimated with adequate precision from a reasonable number of historical samples.

Figure 6

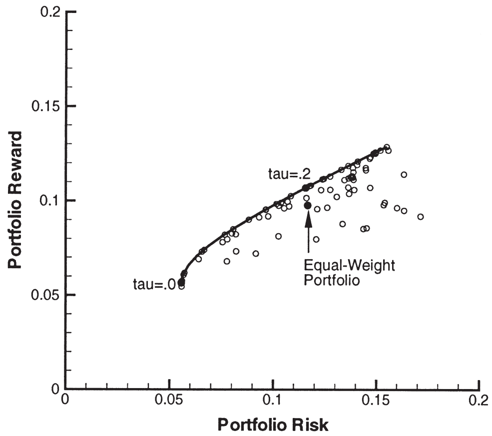

(a) *Sample Estimate*

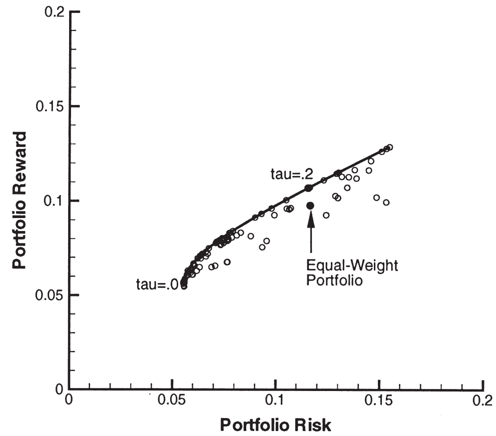

(b) *Bayes-Stein Estimate*

Note: Actual risk-reward tradeoff under true market parameters, measured in units of standard deviations and expected returns, respectively, of a sample of 100 mean-variance optimal portfolios computed using parameters estimated from a history of observed returns. Risk tolerance level was $\tau = 0.2$, and 60 months of historical data were available.

means. The prior for the expected returns used in the Bayes-Stein model is identical in form to the naive specification implicitly used by investors who utilize the $1/n$ heuristic. This indicates that it is surprisingly difficult to improve on this simple diversification rule.

The numerical results given in this paper, and the conclusions based on them, were derived from a set of market parameters that we feel are representative of the choices offered to many investors within their pension plan. The advantages of the equally weighted portfolio when there is estimation risk have been noted by other authors using different data sets and different time periods. For example, Jobson and Korkie (1980) state that "naive formation rules such as the equal weight rule can outperform the Markowitz rule." Michaud (1998) also notes because of that estimation risk, "an equally weighted portfolio may often be substantially closer to the true MV optimality than an optimized portfolio."

Defined contribution plans are becoming the dominant vehicle for providing pension income. In this connection the portfolio strategies of participants are a critical factor since the asset allocation decision determines the ultimate benefits available under these plans. Here we examined the $1/n$ rule and provided a justification. Of course, the performance of the equally weighted heuristic will depend on the asset choices available to the plan members. Given that some participants use the equally weighted heuristic to select their portfolios, perhaps pension plan sponsors should take this into account when selecting the available asset classes.

We emphasize that we have not carried out a comprehensive analysis of the performance of the $1/n$ rule. There are many possible topics of future research. These include

- An investigation of the impact of using different asset classes on the equally weighted portfolio
- Generating expected returns from a market equilibrium model such as the Capital Asset Pricing Model
- Extensions to the multiperiod case.

These questions are left for future research.

ACKNOWLEDGMENTS

The authors acknowledge support from the Natural Sciences and Engineering Research Council of Canada. They also appreciate the comments of the referees on an earlier version.

REFERENCES

BENARTZI, SHLOMO, AND RICHARD THALER. 2001. "Naive Diversification Strategies in Defined Contribution Retirement Saving Plans," *American Economic Review* 91(1): 79–98.

Bengtsson, Christoffer, and Jan Holst. 2002. "On Portfolio Selection: Improved Covariance Matrix Estimation for Swedish Asset Returns." Working paper, Lund University.

Best, Michael, and Robert Grauer. 1990. "The Efficient Set Mathematics When the Mean Variance Problem Is Subject to General Linear Constraints," *Journal of Economics and Business* 42: 105–20.

———. 1991. "On the Sensitivity of Mean-Variance Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results." *Review of Financial Studies* 4(2): 315–42.

———. 1992. "Positively Weighted Minimum-Variance Portfolios and the Structure of Asset Expected Returns," *Journal of Financial and Quantitative Analysis* 27(4): 513–37.

Black, Fischer, and Robert Litterman. 1992. "Global Portfolio Optimization," *Financial Analysts Journal* 48(5): 28–43.

Brennan, Michael, and Walter Torous. 1999. "Individual Decision Making and Investor Welfare." *Economic Notes by Banca Monte dei Paschi di Siena SpA* 28(2): 119–43.

Britten-Jones, Mark. 1999. "The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights," *Journal of Finance* 54(2): 655–71.

Broadie, Mark. 1993. "Computing Efficient Frontiers Using Estimated Parameters," *Annals of Operations Research: Financial Engineering* 45: 21–58.

Chan, Louis K. C., Jason Karceski, and Josef Lakonishok. 1999. "On Portfolio Optimization: Forecasting Covariances and Choosing the Risk Model." *Review of Financial Studies* 12(5): 937–74.

Holden, Sarah, and Jack VanDerhei. 2003. "401(k) Plan Asset Allocation, Account Balances, and Loan Activity in 2001," *ICI Perspective* 9(2).

Jobson, J. Dave, and Bob M. Korkie. 1980. "Estimation for Markowitz Efficient Portfolios," *Journal of the American Statistical Association* 75(371): 544–54.

———. 1981. "Putting Markowitz Theory to Work," *Journal of Portfolio Management* 7: 70–74.

Jorion, Phillippe. 1986. "Bayes-Stein Estimation for Portfolio Analysis," *Journal of Financial and Quantitative Analysis* 21(3): 279–92.

Markowitz, Harry M. 1952. "Portfolio Selection," *Journal of Finance* 7: 77–91.

Michaud, Richard. 1998. *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation.* Cambridge, MA: Harvard Business School Press.

Panjer, Harry H., et al. 1998. *Financial Economics: With Applications to Investments, Insurance, and Pensions.* Schaumburg, Ill.: Actuarial Foundation.

Discussions on this paper can be submitted until January 1, 2005. The authors reserve the right to reply to any discussion. Please see the Submission Guidelines for Authors on the inside back cover for instructions on the submission of discussions.