How Much Foreign Stocks?

Bayesian Approaches to Asset Allocation Can Explain the Home Bias of US Investors

April 2002

Revised January 2003

Ulf Herold*/Raimond Maurer§

* Quantitative Analyst
Metzler Investment GmbH and Johann Wolfgang Goethe University, Frankfurt/Main, Germany
Grosse Gallusstr. 18, 60311 Frankfurt/Main, Germany
Telephone: +49 69 2104 1104
Email: UHerold@Metzler.com

§ Finance Professor
Johann Wolfgang Goethe University, Frankfurt/Main, Germany
Faculty of Business and Economics
Senckenberganlage 31-33 (Uni-PF 58), 60054 Frankfurt/Main, Germany
Telephone: +49 69 798 25227
Email: RMaurer@wiwi.uni-frankfurt.de

How Much Foreign Stocks?

Bayesian Approaches to Asset Allocation Can Explain the Home Bias of US Investors

April 2002

Revised January 2003

ABSTRACT

US investors hold much less foreign stocks than mean/variance analysis applied to historical data predicts. In this article, we investigate whether this home bias can be explained by Bayesian approaches to international asset allocation. In contrast to mean/variance analysis, Bayesian approaches employ different techniques for obtaining the set of expected returns. They shrink the sample means towards a reference point that is inferred from economic theory. We also show that one of the Bayesian approaches leads to the same implications for asset allocation as the mean-variance/tracking error criterion. In both cases, the optimal portfolio is a combination of the market portfolio and the mean/variance efficient portfolio with the highest Sharpe ratio.

Applying the Bayesian approaches to the subject of international diversification, we find that a substantial home bias can be explained when a US investor has a strong belief in the global mean/variance efficiency of the US market portfolio and when he has a high regret aversion of falling behind the US market portfolio. We also find that the current level of home bias can be justified whenever regret aversion is significantly higher than risk aversion.

Finally, we compare the Bayesian approaches to mean/variance analysis in an empirical out-of-sample study. The Bayesian approaches prove to be superior to mean/variance optimized portfolios in terms of higher risk-adjusted performance and lower turnover. However, they do not systematically outperform the US market portfolio or the minimum-variance portfolio.

Keywords: Portfolio selection, asset-pricing models, Bayesian inference, estimation risk, international diversification

I. Introduction

The benefits of international diversification have been subject of a controversial and ongoing debate in the last decades. Classical mean/variance portfolio theory demonstrates that an internationally diversified portfolio has dominant risk/return characteristics compared to a domestic benchmark portfolio. However, the empirical decision behavior of investors is often inconsistent with this normative theory. For example, US investors hold much less foreign stocks than portfolio theory suggests: According to mean/variance analysis, US investors should allocate about 30-40% to foreign equities. The actual allocation of about 8-10% is much below this number (see Lewis [1999] and Sarkar and Li [2002]). This has been termed the "home bias puzzle". To explain the home bias puzzle, some work has been done by using arguments (e.g., based on behavioral finance or non-expected utility theory) which are outside the mean/variance decision criterion. However, surprisingly little effort has been spent on arguments which are consistent with mean/variance analysis, hence establishing a link between normative and descriptive research on international diversification. The key is to focus on the estimation of the input parameters for mean/variance optimization, especially the expected asset returns.

Most published studies demonstrating the benefits from international diversification within the mean/variance approach use sample means derived from historical return data to estimate expected returns. When returns are identically and independently distributed (iid) over time, the sample mean is the best unbiased estimator of the (unknown) mean of the return distribution. While mean/variance analysis is not constrained to historical returns and could be based on subjective forecasts of the future performance of asset classes, the iid assumption is a natural starting point. It is consistent to the efficient market hypothesis. However, when we take estimation errors into account, there are better estimators than the sample mean.[1] In this article, we discuss several Bayesian approaches that employ different estimation techniques to obtain the set of expected returns as an input to the mean/variance decision rule. They have also been developed under the iid setting. Therefore all conclusions we draw are conditional on the iid assumption. If we relax this assumption – and we touch this aspect briefly in one section, there are no general results any more and the conclusions depend on our subjective forecasts of assets' expected returns.[2]

The idea of Bayesian inference is to combine extra-sample, or prior, information with sample returns. The sample means are shrunk towards a reference point that is inferred from economic theory. In this article, we review three Bayesian approaches and investigate whether they can explain the home bias of US investors. The first approach is the Bayes/Stein estimation procedure developed by Jorion (1986). The mean/variance efficient portfolio with the highest Sharpe ratio, the tangency portfolio, is shrunk towards the minimum-variance portfolio. The idea is that there is a tremendous amount of estimation error when estimating expected returns from historical data, therefore expected returns for all asset classes are set to a common value in the prior.

The second approach has been developed only recently in a series of articles by Pastor (2000), Pastor and Stambaugh (1999, 2000) and Wang (2001). The tangency portfolio is now shrunk towards the market portfolio. The prior expected returns are inferred from the CAPM, and the shrinkage effect depends on the degree of sample information in the data and the investor's confidence in the pricing model. Finally, Black and Litterman (1992) also start with the market portfolio. Historical returns do not play any role in their model, they are discarded as worthless for estimating expected returns. Instead, the investor can express his views about future expected returns. These views and the level of conviction that the investor has determine the extent of the deviation from the market portfolio.

Interestingly, the second Bayesian approach of shrinking towards the market portfolio leads to very similar conclusions as the "mean-variance/tracking error" criterion. The idea here is that an investor is not only concerned with expected portfolio return and its variance, but also with underperforming a benchmark portfolio like domestic stocks in our case. Therefore, Chow (1995) extends the classical mean/variance objective function by including an additional term, tracking error variance. We show that this mean-variance/tracking error criterion as well as the second Bayesian approach lead to portfolios which are combinations of the US market portfolio and the global tangency portfolio. Overall, we find that a substantial home bias can be explained when a US investor has a strong belief in the global mean/variance efficiency of the US market portfolio and when he has a high regret aversion of falling behind the US market portfolio. We also find that the current level of home bias can be justified whenever regret aversion is significantly higher than risk aversion.

Finally, we compare the risk-adjusted performance of the Bayesian approaches to mean/variance analysis in a realistic out-of-sample study. The Bayesian approaches prove to be superior to mean/variance optimization. They produce portfolios that exhibit a higher risk-adjusted return (measured by the Sharpe ratio), while reducing turnover. However, the Bayesian approaches are not systematically superior to heuristic strategies like the market portfolio or the minimum-variance portfolio. The empirical results hence confirm the well-known fact that it is hard to estimate expected returns from historical returns alone.

II. Mean/variance analysis and the home bias of US investors

Markowitz (1959) mean/variance efficiency is the classic paradigm of portfolio theory for allocating capital among risky assets. Markowitz shows how to construct efficient portfolios.

The minimum variance frontier comprises all portfolios that have minimum variance for a given level of expected return. The mean/variance efficient frontier is the upward sloping portion of the minimum variance frontier. Inputs are expected returns for each asset, volatility of returns around expected returns, and the correlations between all asset returns. The optimization algorithm takes these inputs as parameters of *known* probability distributions. However, in reality, they are *estimates* of parameters of unknown probability distributions. This issue of estimation risk was the primary reason for the development of the Bayesian approaches discussed in section III.

Most studies about international diversification come to the conclusion that foreign stocks account for 30-40% in the mean/variance efficient portfolio with the highest Sharpe ratio (tangency portfolio); see, e.g., Britton-Jones (1999) and Lewis (1999). These studies apply mean/variance analysis to historical data. The optimal weight depends on the sample period and the assets under consideration. In our empirical study, we use the MSCI USA index as representing US stocks and either an aggregate index or individual country indices for foreign stocks. In the "aggregate index case," we use the MSCI EAFE index, which comprises non-U.S. developed equity markets on a market value weighted basis. In the "individual country indices case," foreign stocks are represented by the rest of the G7 countries (Canada, France, Germany, Italy, Japan, UK) and eight emerging markets (Argentina, Brazil, Chile, Mexico, Hongkong, Korea, Singapore, Thailand). The G7 country returns are proxied by MSCI performance indices, as well as Hongkong and Singapore. Indices for Argentina, Brazil, Chile, Mexico, Korea, and Thailand are taken from the International Finance Corporation (IFC). The sample period is 1/1976 to 3/2002. We use unhedged returns with the USD as numeraire currency, and for calculating the Sharpe ratio, we use the 3month Treasury bill as a proxy for the risk-free rate. We impose short-selling restrictions because most investors face these constraints.

In the aggregate index case, foreign stocks account for 39% in the MVP and 24% in the tangency portfolio. The fact that the weight of foreign stocks is higher in the MVP than in the TP is somewhat counter-intuitive, but it can be explained with Figure 1. In our sample period, US stocks have dominated EAFE stocks in mean/variance terms.[3] One reason for this is that part of the EAFE volatility is due to the foreign currency exposure. Another reason is that the EAFE index is probably highly mean/variance inefficient, as Wilcox (1994) argues. We therefore include individual country indices in a second step. Despite the fact, that US stocks have been a dominant asset class, foreign stocks are attractive because they provide diversification benefits. The tangency portfolio (comprised of US and EAFE stocks) lies above the line connecting US equities and the risk-free rate; this line represents portfolios invested in US equities and cash. When splitting foreign equities into individual developed countries and adding emerging markets, the allocation to foreign equities sharply increases to about 45% in the MVP and more than 70% in the TP, as can be seen in Figure 3 below.[4] Overall, mean/variance analysis based on historical data cannot explain the home bias of US investors observed in practice.

>> Figure 1 ABOUT HERE <<

III. Bayesian approaches to international asset allocation

1. Overview of Bayesian approaches

The most crucial input for asset allocation is the set of expected returns. Expected returns can be estimated from historical returns, derived from a forecasting model, or inferred from an asset-pricing model. In each case, there is a substantial amount of uncertainty attached to these estimates. The problem with mean/variance analysis is that it utilizes only a single set of estimates: asset allocation is based only on sample means in the iid setting (or otherwise only on personal judgments about the future performance of asset classes). Furthermore, the estimates of expected returns are treated as if they were the true values. A better approach would be to assess the information content of the different information sources and then combine them into a single estimate. This exactly is the basic idea of Bayesian statistics.

Bayesian techniques shrink the estimators for the means from their sample means to some prior values, depending on the degree of estimation error in the sample. They thereby produce a new, combined set of estimates for the expected returns. There are several possibilities how to specify the prior means; see Figure 2. Jorion (1986) shrinks the tangency portfolio to the minimum-variance portfolio, while Pastor (2000) shrinks it towards the market portfolio. Black and Litterman (1992) also shrink towards the market portfolio, but instead of using sample means, their model is designed to incorporate investor's views about future expected returns. We will now investigate the implications of these approaches on the subject of international diversification. We discuss Pastor's approach in most detail, as it has been developed only recently and has not yet received as widespread attention as the two other ones.

2. Shrinking towards the minimum-variance portfolio

As estimates of expected returns are prone to estimation errors, Jorion (1986, 1991) shrinks mean/variance efficient portfolios towards the MVP. The MVP is less vulnerable to estimation risk as it does not make use of any information about expected returns. The rationale is that we can estimate the covariance matrix from the sample returns quite precisely, while estimation errors in the means are tremendous.[5] Jorion sets the prior means to a common value across all N assets. He specifies the prior as

[1] $\mu \sim N(1\mu_0, \Sigma/\phi)$,

where μ is the N×1 vector of expected returns, 1 is a vector of ones, μ_0 is the expected return of the MVP, Σ denotes the N×N covariance matrix, and ϕ determines the prior precision. Jorion further assumes that asset returns are multivariate normally distributed, $R \sim N(\mu, \Sigma)$. Combining the prior distribution with the sample likelihood function yields the posterior vector of expected returns

[2] $$\mu_T = \frac{\phi}{\phi + T} 1\mu_0 + \frac{T}{\phi + T} \overline{R},$$

where $\overline{R}$ denotes the N×1 vector of sample means and T is the number of observations. Jorion further demonstrates that ϕ can be estimated from the data:

[3] $$\hat{\phi} = \frac{N + 2}{(\overline{R} - \mu_0 1)' \Sigma^{-1} (\overline{R} - \mu_0 1)},$$

where the denominator measures the observed dispersion of the sample means around the common mean.[6]

As equation [2] shows, Jorion shrinks the sample means towards the MVP mean return. The longer the sample history, T, the weaker is the shrinkage. In the extreme, $T \to \infty$, the investor will use the sample means, $\mu_T = \overline{R}$, i.e. the Bayes/Stein estimator includes the sample mean as a special case. At the other extreme, with no uncertainty in the prior, $\phi \to \infty$, the Bayes/Stein approach results in the MVP. In the (more interesting) cases in between these extremes, the Bayes/Stein approach shrinks the TP towards the MVP. In fact, it can be shown that the optimal portfolio under Bayes/Stein is a combination of MVP and TP.[7]

Applying the Bayes/Stein estimation procedure to our sample leads to mixed results (see Figure 3). In the aggregate index case, the Bayes/Stein portfolio is unexpectedly close to the MVP. With individual country indices, it lies in the middle between TP and MVP. The allocation to foreign stocks is significantly reduced from more than 70% to roughly 55%. However, the home bias effect cannot be explained.

A drawback of the Bayes/Stein approach is that it imposes the prior assumption that all assets have the same expected return, irrespective of their risk profile. It would be more economically sound to link the prior expected return to the systematic risk of an asset class. This is done by the Bayesian approach that we consider next.

3. Shrinking towards the market portfolio: Incorporating an asset-pricing model

In a series of articles, Pastor (2000), Pastor and Stambaugh (1999, 2000) and Wang (2001) have suggested to incorporate an asset-pricing model for asset allocation. They combine the results of mean/variance optimization and the implications of an asset-pricing model, again using Bayesian inference. Their motivation is that mean/variance analysis, on the one side, only utilizes the data and bases portfolio selection on the first and second sample moments. It, however, completely ignores the potential usefulness of an asset-pricing model. Basing asset allocation only on an asset-pricing model, on the other side, makes no use at all of the time series of returns. In general, a model will neither be a perfect description of reality (by construction) nor will it be completely useless for decision making. To quote Pastor (2000, p. 179):

> "By definition, every model is a simplification of reality. Hence, even if the data fail to reject the model, the decision maker may not necessarily want to use the model as a dogma. At the same time, the notion that models implied by finance theory could be entirely worthless seems rather extreme. Hence, even if the data reject the model, the decision maker may want to use the model at least to some degree."

In practice, investors employ both approaches, but for different purposes. E.g., when they decide *within* an asset class like US equities to employ a passive instead of an active manager, they (implicitly) follow the CAPM because this model implies to invest in an index fund. For asset allocation decisions *across* asset classes, e.g. the split between stocks and bonds, they often perform mean/variance optimization.

The approach developed by Pastor et al. combines the sample information and the implications of an asset-pricing model. It is flexible to accommodate single- and multi-factor models, but we consider only the CAPM. In this case, the prior expected excess returns are set equal to the expected excess returns implied by the CAPM, as given by

[4] $\mu_{CAPM} = \beta\mu_M$,

where β is the N×1 beta vector and μ_M denotes the risk premium of the market portfolio. As in Jorion's approach, the posterior expected returns are a weighted average of the prior expected returns and the sample means:

[5] $\mu_T = \omega\mu_{CAPM} + (1-\omega)\bar{r}$,

where $\bar{r}$ denotes the sample means of the excess returns (over the risk-free rate) and ω denotes the shrinkage factor. The sample means are shrunk towards the implied CAPM excess returns, hence the tangency portfolio is shrunk towards the market portfolio. The shrinkage factor is a measure for the weight which is assigned to the CAPM and is given by

[6] $$\omega = \frac{s^2/\sigma^2}{s^2/\sigma^2 + T/(1+SR^2)},$$

where SR is the Sharpe ratio of the market portfolio and s^2 is the average variance of the residual terms in the multivariate regression of the assets' excess returns on the excess returns of the market portfolio. σ measures the dispersion of the assets' alphas, i.e. the deviation of the assets' expected returns from the values implied by the CAPM. It is a measure of the investor's prior uncertainty in the CAPM. Thus, the degree of shrinkage depends on how much confidence the investor has into the validity of the CAPM and on the strength of the violations of the CAPM in the historical data. The optimal portfolio is approximately a linear average of the market portfolio and the tangency portfolio with a fraction of ω invested into the market portfolio.[8]

Figure 4 displays the asset allocations for varying degrees of confidence in the domestic CAPM. The shrinkage factor, ω, is shown in the first row, and the prior uncertainty, σ, that the investor has in the pricing model, is included in the second row. If the investor has an extremely strong belief in the global mean/variance efficiency of the US market portfolio, he sets σ to zero and ends up with the US market portfolio. If he has no confidence at all in the domestic CAPM, he will choose the allocation on the right side and allocate about 25% to foreign stocks in the aggregate index case (or 70% in the individual indices case). If he has some confidence, he will choose a portfolio in the middle.

>> Figure 4 ABOUT HERE <<

The information displayed in Figure 4 can be used in different ways. Of course, the investor could specify his prior uncertainty in the domestic CAPM, and then choose the according portfolio. However, it might be difficult for him to quantitatively determine σ. One useful approach is therefore to start with the current allocation and check whether the implied σ is consistent to one's belief. E.g., the average US investor has allocated about 8-10% to foreign stocks. This allocation implies a moderately high confidence into the domestic CAPM, because the investor assigns a weight of about 65% to the model and only 35% to the data in the aggregate index case. Hence, the investor might want to question whether he is really that confident in the model (or whether he is even more confident). Another possibility is to estimate σ from the data. This is called an "empirical Bayes approach" because σ is the prior parameter and is usually specified exogenously in Bayesian statistics. Estimating σ from the data will lead to a shrinkage weight of 0.5, as Wang (2001) has shown, implying that the model and the data have equal impact on asset allocation. Loosely speaking, the investor will invest half of his wealth into the US market portfolio and half into the tangency portfolio. A US investor should then hold about 13% in foreign equities in the aggregate index case (or 40% in the in the individual indices case).

To summarize, the Bayesian approach in this section can explain the home bias effect. It shows that a US investor must have a strong belief in the global mean/variance efficiency of the US market portfolio in order to justify the current level of home bias. As a final remark, note that in line with Pastor (2000), the domestic CAPM serves as an anchor (or, shrinkage target) for our analysis. Instead, we could have started with the international CAPM. The global tangency portfolio would then have been shrunk towards the world market portfolio. The reason to employ the domestic CAPM was to explore the home bias effect, in particular to investigate the strength of the belief into the global mean/variance efficiency of the US market portfolio that an investor must have to justify the current level of home bias.

4. Shrinking towards the market portfolio: Incorporating investor's views

Black and Litterman (1992) also use the CAPM as their prior assumption. With this respect, their approach is similar to that of Pastor et al. However, while the approach of Pastor et al. accommodates any asset-pricing model, Black and Litterman base their analysis only on the CAPM. Furthermore, in an international context, they employ the international CAPM version developed by Black (1990). According to Black's universal hedging theory, each investor (no matter what his numeraire currency) should hold the world market portfolio, partly hedged against currency risk. The hedge ratio is identical across all investors and is determined empirically; it takes on a value of about 80%.

The important difference of the Black/Litterman model to the approaches discussed so far is that historical returns play no role in their approach to estimate expected returns. Instead, the investor can express his views about future expected returns. The views are articulated in terms of probability statements like, e.g., “the expected return of foreign stocks is 2%-4% lower than that of domestic stocks over the next 12 months, with a probability of 90%”. These views are combined with the prior expected values (equilibrium excess returns implied by the CAPM) in a Bayesian way,[9] and they determine the extent of the deviation from the market portfolio. If an investor is very bearish on foreign stocks (i.e., he beliefs that their expected return is much below the expected return implied by the CAPM), the model will result in a significantly lower allocation to foreign stocks compared to their weight in the world market portfolio. The output of the Black/Litterman model therefore strongly depends on the investor’s views. It does not lead to a “unique” solution, in contrast to the approaches of Jorion and Pastor. As we do not make use of any forecasts of expected returns in this paper, we do not apply the Black/Litterman approach in the empirical out-of-sample study below.

IV. The mean-variance/tracking error (MVTE) criterion

Under certain assumptions, the mean/variance criterion is consistent to expected utility maximization.[10] The fact that investors usually deviate from mean/variance efficient portfolios is partly due to estimation risk which has been discussed above. Another reason is that they might have different objectives than expected utility maximization. Clarke et al. (1994) argue that an investor is not only averse to loosing money. He is also averse to losses relative to a specific benchmark portfolio like an all-cash strategy or the market portfolio. Similarly, portfolio managers are not only concerned with the prospect of losing money but also about falling behind their peer groups.

In this line, Chow (1995) argues that neither the mean-variance nor the mean-tracking error criterion yields satisfactory results. Investors usually show aversion to both absolute risk (the risk of loosing money) and relative risk (the risk of not meeting a benchmark or their peer group). Hence, Chow extends the classical mean/variance objective function

[7] $\Phi = \mu_P - \lambda_A \sigma_P^2$,

where μ_P denotes expected portfolio (excess) return, λ_A denotes risk aversion, and σ_P denotes portfolio volatility in the following way:

[8] $\Phi' = \mu_P - \lambda_A \sigma_P^2 - \lambda_R \psi_P^2$,

where λ_R denotes regret aversion and ψ_P denotes tracking error, i.e. the standard deviation of portfolio returns in excess of the benchmark. Chow refers to this formula as "mean-variance/tracking error" (MVTE) criterion.[11] Hence, an investor is looking for a portfolio with high expected return but low absolute risk and low relative risk. The ratio of risk aversion to regret aversion determines whether the portfolio will be tilted more towards the mean/variance efficient or towards the benchmark portfolio. In fact, we prove in Appendix 2 that the optimal portfolio is a linear weighted average of the tangency portfolio and the benchmark portfolio, where the weight on the benchmark is given by $\lambda_R/(\lambda_A+\lambda_R)$.

To make the results comparable to the Bayesian approach above, we set the parameters in the following way. First, we set risk aversion, λ_A, equal to the implied risk aversion by holding the tangency portfolio.[12] Then, with no regret aversion (λ_R=0), the investor will choose the tangency portfolio. On the other hand, with infinite regret aversion ($\lambda_R \to \infty$), he will invest in the benchmark portfolio. Second, we set the benchmark portfolio equal to US equities as we have done in the previous section. For the different purpose of defining regret aversion relative to the peer group of an average US investor, one could use the average US investor's actual holdings of foreign equities to specify the benchmark.

Figure 5 displays the allocations to US and foreign equities for varying levels of regret aversion. The first row shows the regret aversion parameter, λ_R, the second row shows tracking error, ψ_P (in %), and the third row contains the proportion invested in the benchmark portfolio, ϕ. As mentioned above, ϕ is equal to $\lambda_R/(\lambda_A+\lambda_R)$. It corresponds to the shrinkage factor, ω, in the Bayesian approach of Pastor et al. For a regret aversion of λ_R=1.69 or a tracking error of 59bp, the portfolio is identical to the portfolio under the empirical Bayes approach in the aggregate index case. Equivalently, there also exists a portfolio in the individual indices case (λ_R=1.87) which is the counterpart to the empirical Bayes portfolio. Appendix 1 summarizes some further issues with respect to this duality.

>> Figure 5 ABOUT HERE <<

To summarize, both the Bayesian approach discussed in section III.3. and the MVTE criterion lead to similar implications for asset allocation, although they are based on different objective functions: Both result in portfolios which are linear weighted combinations of the global tangency portfolio and the US market portfolio. In both approaches, there is (at least) one parameter that the investor must specify: either the prior degree of confidence in an asset-pricing model or the level of risk relative to a pre-specified benchmark portfolio he is willing to take. Practical experience will show which approach will be preferred by investors. Of course, both approaches might be combined in practice because they provide different points of view about the benefits of international diversification and might help the investor to better understand the implications of his investment decision. E.g., the investor could first employ the Bayesian approach and derive a preliminary asset allocation. He could then compute his implied regret aversion using the MVTE criterion and then decide whether he is willing to incur the resulting tracking error.

Finally, we note that the MVTE criterion can explain the home bias puzzle. It shows that a US investor ends up with an allocation of 12.2% into foreign stocks if his regret and risk aversion are of equal magnitude. Empirical studies, however, show that in general, regret aversion is considerably higher than risk aversion. E.g., when regret aversion equals three times risk aversion,[13] the investor will set $\phi=0.75$ and hence allocate only about 9% abroad. This number is consistent with the home bias observed in practice.

V. Out-of-sample study

In this section, we perform an out-of-sample study and compare the risk-adjusted performance of the Bayesian approaches to mean/variance analysis. We also measure the monthly turnover that the strategies generate. Mean/variance optimization leads to unstable and extreme portfolio weights over time due to the estimation risk problem. Mean/variance optimized portfolios also lack of diversification. The Bayesian approaches incorporate estimation risk into portfolio selection and, thus, should reduce turnover. When shrinking towards the market portfolio, the optimal portfolios will also be more diversified, as the market portfolio shows a high degree of diversification by construction.

Using rolling windows, we construct portfolios based on past returns and hold them fixed for one future period. By comparing the out-of-sample Sharpe ratios of the two Bayesian approaches, we can evaluate whether employing asset-pricing models for portfolio selection leads to superior results. The out-of-sample study also gives an indication whether Bayesian approaches which utilize only past return data are useful for tactical asset allocation.

We consider the same asset classes as before. We take the viewpoint of a US investor and calculate monthly excess returns by subtracting the 3month T-Bill rate from the total returns of the EAFE and the individual country indices. The total sample period is from 1/1976 to 3/2002. A rolling window of 120 months is used to estimate the optimization input parameters. In the first run, portfolio weights are based on the estimation period from 1/76 to 12/85. Using the returns of 1/86, the first out-of-sample portfolio returns can be calculated. Then the estimation period is rolled one month forward, and the next portfolio composition is based on 2/76 to 1/86. This procedure results in a total of 195 out-of-sample returns. At each point of time, a total of four portfolios is optimized: the Bayes/Stein portfolio (labeled "BayesMVP"), the portfolio under the Bayesian approach by Pastor et al., where the shrinkage factor is determined by the empirical Bayes approach ("BayesBM"), the mean/variance tangency portfolio ("TP"), and the MVP ("MVP"). We compare these to the US market portfolio ("BM") and the equally-weighted portfolio ("EWP"). Short-selling is not permitted.

Figure 6 presents the annualized mean returns, standard deviations and Sharpe ratios of the out-of-sample strategies. The average monthly one-way turnover is displayed in the last row. When using the aggregate indices (Panel A), the BayesBM strategy shows a Sharpe ratio of 0.599, which is slightly higher than the Sharpe ratio of US stocks (0.578). The Sharpe ratio of the BayesMVP strategy is approximately equal to the benchmark portfolio. Both Bayesian approaches are clearly better than the EWP, which invests half of the wealth in US stocks and half in the EAFE index and which is rebalanced monthly. The Bayesian approaches are also better than the tangency portfolio. The Sharpe ratios are a little higher and monthly turnover is reduced. When extending the universe (Panel B), the two Bayesian approaches show a considerably higher performance than the tangency portfolio with less turnover. However, they cannot improve on the market portfolio. The best strategy is the MVP.

>> Figure 6 ABOUT HERE <<

The main findings of the empirical study can be summarized in three points.[14] First, incorporating estimation risk into portfolio selection clearly leads to superior performance compared to the classical tangency portfolio. Second, at least for our sample period and assets studied, employing asset-pricing models for portfolio selection cannot enhance performance. Third, neither of the two Bayesian approaches can (significantly) outperform the US market portfolio. Therefore, the empirical results confirm the well-known fact that it is hard to estimate expected returns (as well as expected alphas) from historical returns alone. As several previous studies have shown, stock returns are partly predictable (see, e.g., Ferson/Harvey [1993]). For tactical asset allocation strategies, it is promising to relax the iid assumption and use instrumental variables to predict stock returns instead of relying on historical returns alone. To construct portfolios from these forecasted mean returns, the Black/Litterman model can be employed.

VI. Conclusion

Mean/variance analysis applied to historical data cannot explain the home bias of US investors. It utilizes only one information set – sample means – and hence is vulnerable to estimation errors. In contrast, Bayesian approaches combine different information sources into a single set of expected returns. While Jorion (1986) shrinks the tangency portfolio towards the minimum-variance portfolio, Pastor (2000), Pastor/Stambaugh (1999, 2000) and Wang (2001) shrink it towards the market portfolio. Interestingly, their Bayesian approach leads to the same implications for asset allocation as the mean-variance/tracking error (MVTE) criterion suggested by Chow (1995). In both cases, the optimal portfolio is a combination of the US market portfolio and the global tangency portfolio. Therefore, the MVTE criterion, which has been a mere description of how investors behave so far, can be given a theoretical foundation. There are arguments to use it in a normative way.

Applying these approaches to the subject of international diversification, we find that the Bayes/Stein procedure developed by Jorion results in portfolios which are still heavily exposed to foreign stocks and hence cannot account for the home bias effect. The Bayesian approach of shrinking towards the market portfolio as well as the MVTE criterion, on the other side, can explain the low allocation to foreign stocks observed in practice. They show that the current home bias is justified when a US investor has a strong belief in the global mean/variance efficiency of the US market portfolio and when his regret aversion of falling behind the US market portfolio exceeds his risk aversion.

Finally, we compare mean/variance analysis and the Bayesian approaches to asset allocation in an empirical out-of-sample study. We find that both Bayesian approaches lead to a clearly superior performance and to lower turnover compared to the tangency portfolio. However, they cannot beat systematically the minimum-variance or the market portfolio. The empirical study hence confirms that it is hard to estimate expected returns from historical returns alone. For tactical asset allocation, some exogenous information needs to be included.

VII. Appendix

1. Portfolio compositions under the MVTE criterion and the Bayesian approach of incorporating an asset-pricing model

Here, we summarize several technical notes about the relation between portfolios under the MVTE criterion and the Bayesian approach of shrinking the tangency portfolio towards the market portfolio. First, the optimal portfolio weights under the Bayesian approach are not linear in the shrinkage factor. E.g. in Panel A of Figure 4, the portfolio with ω=0.5 has a portion of 12.8% allocated into foreign stocks, which is slightly more than 12.2% (which is half the weight of foreign stocks in the portfolio with ω=0). Second, the composition of the tangency portfolios under both the Bayesian approach and the MVTE criterion slightly differ. This can be seen by comparing the portfolio on the right hand-side in Figure 4 and Figure 5. The reason is that for portfolio optimization under the MVTE criterion, the (unadjusted) sample covariance matrix is used, while in the Bayesian approach, the predictive covariance matrix is employed. In the latter one, some degrees-of-freedom adjustments are made. Third, the shrinking factor, ω, indicates the weight which is assigned to the asset-pricing model and does not equal the fraction of wealth which is invested in the benchmark portfolio, as explained above. Hence, for a given value of ω, the portfolio under the Bayesian approach will differ from the portfolio under the MVTE criterion with the same value of ϕ (except for ω=ϕ=0). However, there does exist a portfolio under the MVTE criterion with a different value of ϕ which is identical to the portfolio under the Bayesian approach. E.g., the portfolio with ϕ=0.43 in Panel B of Figure 5 is identical to the portfolio with ω=0.5 in Panel B of Figure 4. Therefore, also the Bayesian approach to asset allocation leads to linear weighted combinations of the tangency portfolio and the benchmark portfolio. This result holds irrespective whether short-selling restrictions are included or not.

2. Decomposition of the optimal portfolio under the MVTE criterion

Here, we prove that the optimal portfolio under the MVTE criterion is a linear weighted average of the tangency portfolio and the benchmark portfolio. We show that the weight on the benchmark is given by $\lambda_R/(\lambda_A+\lambda_R)$ with λ_A being absolute risk aversion and λ_R being regret aversion.

The aim is to maximize the extended objective function

[A-1] $$\max_{h_P}\left(\mu_P - \lambda_A \sigma_P^2 - \lambda_R \psi_P^2\right) = \max_{h_P}\left(h_P'\mu - \lambda_A h_P' V h_P - \lambda_R (h_P - h_B)' V (h_P - h_B)\right),$$

where h_P and h_B denote the N×1 vector of portfolio weights and benchmark weights, respectively, μ is the N×1 vector of expected excess returns and V is the N×N covariance matrix of returns.

Differentiating [A-1] with respect to h_P and setting the expression to zero yields

$$\mu - 2\lambda_A V h_P - 2\lambda_R V (h_P - h_B) = 0$$

Hence,

$$h_P = \frac{1}{2(\lambda_A + \lambda_R)} V^{-1}\left(\mu + 2\lambda_R V h_B\right) = \frac{1}{2(\lambda_A + \lambda_R)} V^{-1}\mu + \frac{\lambda_R}{\lambda_A + \lambda_R} h_B$$

Recognizing that $2\lambda_A = 1'V^{-1}\mu$ (see Footnote 12) and that the weights of the tangency portfolio are given by $h_{TP} = \frac{V^{-1}\mu}{1'V^{-1}\mu}$, it follows

[A-2] $$h_P = \frac{\lambda_A}{\lambda_A + \lambda_R} h_{TP} + \frac{\lambda_R}{\lambda_A + \lambda_R} h_B .$$

As h_B and h_{TP} are fully invested portfolios, h_P is also fully invested.

1 Technically, „better“ means that the sample mean is not admissible, i.e. it does not have a uniformly lower risk function than all other estimators. The risk function is a core concept of statistical decision theory; see, e.g., Mood et al. (1974).

2 Replacing sample means with personal judgments about the future performance of assets is appropriate only when the iid assumption is violated. We stay within the classical framework of iid returns, except in Section III.4., where we will deal with the issue of how incorporating subjective views about expected returns.

3 Of course, a sample period can be found where the EAFE index showed a stronger performance than the MSCI index for the U.S. (e.g., from 1976 to 1990). We base our analysis on the longest available sample period because of our assumption of iid returns. Thus, the longer the sample period, the lower estimation errors.

4 Panel B of Figure 3 also confirms the well-known fact that mean/variance optimized portfolios generally exhibit a very low degree of diversification. Mean/variance efficient portfolios therefore are not implemented in practice, before making adjustments. We do not focus on this diversification issue here; instead we are more interested in the overall weight of foreign stocks in the portfolio.

5 Estimation errors in the means also have a much severe impact on optimal portfolios compared to errors in the variances and correlations. See, e.g., Chopra and Ziemba (1993).

6 With the assumption that Σ is known, the predictive covariance matrix is given by $\Sigma+\Sigma/(\phi+T)$ and is in general not different from Σ for practical purposes. Therefore we use Σ for portfolio optimization.

7 This holds when portfolio weights are not constrained, and the reason is that Jorion specifies the prior covariance matrix as proportional to the asset covariance matrix, as can be seen in [1].

8 This holds only approximately because only the predictive mean is linear in ω but the covariance matrix is not. We refer readers to the original articles of Pastor (2000) and Pastor and Stambaugh (2001) for the equation of the predictive covariance matrix.

9 Although in their original paper, Black and Litterman (1992) refer to the “mixed estimation strategy” of Theil (1971), their model can also be developed within a Bayesian modeling framework.

10 This holds either if the investor’s utility function is quadratic, or the distribution of asset returns is (multivariate) normal.

11 The MVTE criterion is a special case of a multiple-benchmark optimization, with one risky benchmark and the risk-free rate as an additional benchmark. See Wang (1999) for a more general framework. See also Wilcox (1994), who sticks to the classical objective function [7], but computes covariances with respect to a hybrid benchmark consisting of cash and the risky benchmark. This is equivalent to the MVTE criterion.

12 The implied risk aversion of the tangency portfolio is given by $(1/2)\mu' V^{-1} 1$, where μ is the vector of expected excess returns, V is the covariance matrix of returns and 1 is a vector of ones. See Lee (2000).

13 This is a conservative assumption. E.g., Waring et al. (2000) find that regret aversion is about six times greater than risk aversion. This is the case for an investor who is exposed to 12% volatility (portfolio consisting of stocks and bonds) and 2% tracking error. With regret aversion equal to risk aversion, the investor would have to increase the tracking error to more than 10%, which is way beyond what is observed in practice.

14 These results also hold for a rolling period length of 90 months. When shortening the length to 60 months or less, all optimized portfolios show a poorer performance. The same result is true when extending the rolling period length beyond 150 months. This indicates that 5 years of monthly data are not sufficient to reduce estimation errors and more than 12 years do not capture the time-variation of returns. The results are similar for quarterly data.

References

Black/F. (1990): Equilibrium exchange rate hedging, in: Journal of Finance 45, 899-908

Black, F./Litterman, R. (1992): Global Portfolio Optimization, in: Financial Analysts Journal, September/October, 8-43

Chopra, V./Ziemba, W.T. (1993): The effect of errors in means, variances, and covariances on optimal portfolio choice, in: Journal of Portfolio Management 19, No. 3, 6-11

Chow, G. (1995): Portfolio selection based on return, risk, and relative performance, in: Financial Analysts Journal, March/April, 54-60

Clarke, R./Krase, S./Stratman, M. (1994): Tracking errors, regret, and tactical asset allocation, in: Journal of Portfolio Management, Summer, 16-24

Ferson, W.E./Harvey, C.R. (1993): The risk and predictability of international equity returns, in: Review of Financial Studies 6, No. 3, Fall, 527-566

Jorion, P. (1986): Bayes-Stein estimation for portfolio analysis, in: Journal of Financial and Quantitative Analysis 21, No. 3, 279-292

Jorion, P. (1991): Bayesian and CAPM estimators of the mean: Implications for portfolio selection, in: Journal of Banking and Finance 15, No. 3, 717-727

Lee, W. (2000): Theory and Methodology of Tactical Asset Allocation, New York: Wiley

Lewis, K. (1999): Trying to explain home bias in equity and consumption, in: Journal of Economic Literature 37, 571-608

Markowitz, H. (1959): Portfolio selection: Efficient diversification of investments, New York

Mood, A.M./Graybill, F.A./Boes, D.C. (1974): Introduction to the theory of statistics, 3rd ed., New York: McGraw-Hill

Pastor, L. (2000): Portfolio selection and asset pricing models, in: Journal of Finance 55, No. 1, 179-223

Pastor, L./Stambaugh, R.F. (1999): Costs of equity capital and model mispricing, in: Journal of Finance 54, No. 1, 67-121

Pastor, L./Stambaugh, R.F. (2000): Comparing asset pricing models: An investment perspective, in: Journal of Financial Economics 56, No. 3, 335-381

Sarkar, A./Li, K. (2002): Should US investors hold foreign stocks?, in: Federal Reserve Bank of New York Current Issues in Economics and Finance 8, No. 3, March, pp. 1-6

Theil, H. (1971): Principles of econometrics, New York: Wiley&Sons

Wang, M. (1999): Multiple-benchmark and multiple-portfolio optimization, in: Financial Analysts Journal January/February, 63-72

Wang, Z. (2001): A shrinkage approach to asset allocation using asset-pricing models, Working Paper Columbia University

Waring, B./Whitney, D./Pirone, J./Castille, C. (2000): Optimizing manager structure and budgeting manager risk, in: Journal of Portfolio Management, Spring, 90-104

Wilcox, J. (1994): EAFE is for wimps, in: Journal of Portfolio Management, Spring, 68-75

Figure 1: Mean/variance efficient frontier of US and EAFE stocks

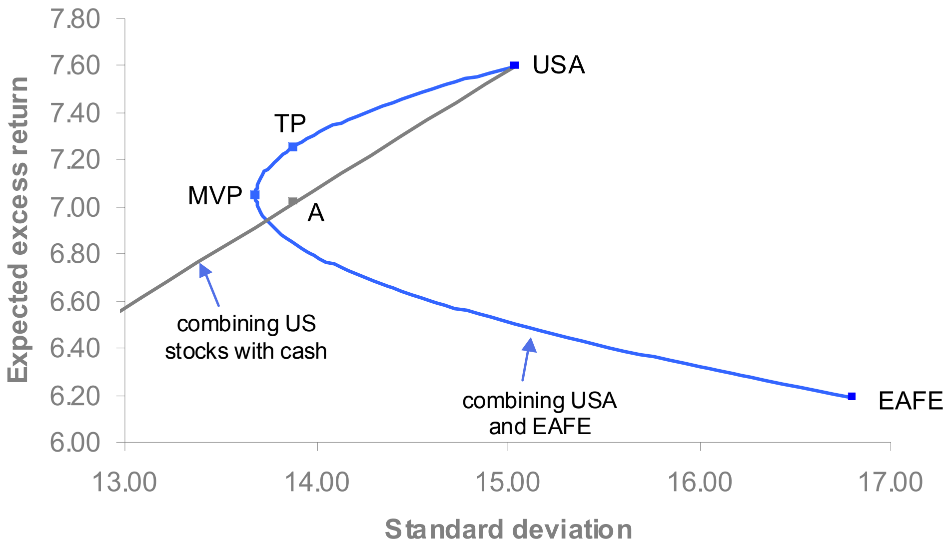

Figure 2: Bayesian approaches employ different information sets and shrinkage targets

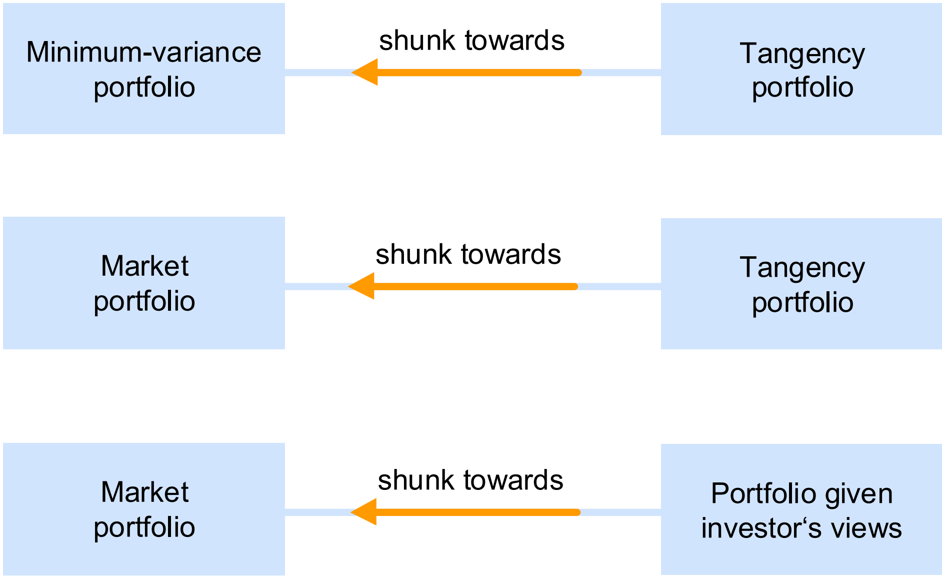

Figure 3: MVP, TP, and Bayes/Stein portfolio

a) Aggregate index case

	MVP	B/S	TP
USA	61.03%	61.74%	75.68%
EAFE	38.97%	38.26%	24.32%

b) Individual country indices case

	MVP	B/S	TP
USA	54.65%	45.82%	27.98%
CAN	0.00%	0.00%	0.00%
JAP	11.75%	4.65%	0.00%
FRA	0.00%	0.00%	3.39%
GER	11.19%	6.41%	0.00%
ITA	5.92%	3.17%	0.00%
UK	6.28%	14.49%	18.93%
ARG	0.76%	4.01%	7.91%
BRA	0.63%	2.77%	5.16%
CHI	7.80%	15.03%	23.25%
MEX	0.00%	0.13%	3.29%
HK	0.00%	0.00%	4.07%
SIN	0.00%	0.00%	0.00%
KOR	0.64%	3.52%	6.03%
THA	0.37%	0.00%	0.00%

Figure 4: Shrinking towards the US market portfolio

a) Aggregate index case

Implication of domestic CAPM

Estimating σ from the data („empirical Bayes approach") implies ω=0.5

Mean/variance optimization

ω	1.00	0.90	0.75	0.65	0.50	0.25	0.10	0.00
σ	0.00	0.28	0.48	0.61	0.83	1.43	2.48	∞
USA	100.00%	97.31%	93.40%	90.87%	87.18%	81.30%	77.93%	75.75%
EAFE	0.00%	2.69%	6.60%	9.13%	12.82%	18.70%	22.07%	24.25%

σ: prior uncertainty in the CAPM
ω: shrinkage factor (weight assigned to asset-pricing model)

b) Individual country indices case

ω	1.00	0.95	0.90	0.75	0.50	0.25	0.10	0.00
σ	0.00	0.49	0.71	1.23	2.13	3.70	6.40	∞
USA	100.00%	95.16%	90.50%	77.51%	58.80%	43.17%	35.05%	30.09%
CAN	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
JAP	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
FRA	0.00%	0.23%	0.44%	1.05%	1.93%	2.66%	3.04%	3.27%
GER	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
ITA	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
UK	0.00%	1.27%	2.49%	5.90%	10.82%	14.92%	17.06%	18.36%
ARG	0.00%	0.53%	1.04%	2.47%	4.53%	6.25%	7.14%	7.69%
BRA	0.00%	0.35%	0.68%	1.61%	2.95%	4.07%	4.65%	5.01%
CHI	0.00%	1.56%	3.07%	7.26%	13.30%	18.35%	20.97%	22.57%
MEX	0.00%	0.22%	0.44%	1.03%	1.89%	2.60%	2.97%	3.20%
HK	0.00%	0.27%	0.54%	1.27%	2.33%	3.22%	3.68%	3.96%
SIN	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
KOR	0.00%	0.40%	0.79%	1.88%	3.45%	4.75%	5.43%	5.85%
THA	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

σ: prior uncertainty in the CAPM
ω: shrinkage factor (weight assigned to asset-pricing model)

Figure 5: Allocation to foreign stocks for varying levels of regret aversion

a) Aggregate index case

	Benchmark portfolio							Tangency portfolio
λ_R	∞	35.88	17.00	5.67	1.89	1.69	0.63	0.00
ψ_P	0.00	0.06	0.11	0.28	0.56	0.59	0.84	1.12
ϕ	1.00	0.95	0.90	0.75	0.50	0.47	0.25	0.00
USA	100.00%	98.78%	97.57%	93.92%	87.84%	87.18%	81.76%	75.68%
EAFE	0.00%	1.22%	2.43%	6.08%	12.16%	12.82%	18.24%	24.32%

λ_R: regret aversion
ϕ: fraction invested in the benchmark portfolio
ψ_P: tracking error

b) Individual country indices case

λ_R	∞	22.54	7.51	2.50	1.87	0.83	0.28	0.00
ψ_P	0.00	0.42	1.05	2.09	2.40	3.14	3.77	4.19
ϕ	1.00	0.90	0.75	0.50	0.43	0.25	0.10	0.00
USA	100.00%	92.80%	82.00%	63.99%	58.80%	45.99%	35.18%	27.98%
CAN	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
JAP	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
FRA	0.00%	0.34%	0.85%	1.69%	1.94%	2.54%	3.05%	3.39%
GER	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
ITA	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
UK	0.00%	1.89%	4.73%	9.47%	10.83%	14.20%	17.04%	18.93%
ARG	0.00%	0.79%	1.98%	3.96%	4.53%	5.93%	7.12%	7.91%
BRA	0.00%	0.52%	1.29%	2.58%	2.95%	3.87%	4.64%	5.16%
CHI	0.00%	2.32%	5.81%	11.62%	13.30%	17.44%	20.92%	23.25%
MEX	0.00%	0.33%	0.82%	1.64%	1.88%	2.47%	2.96%	3.29%
HK	0.00%	0.41%	1.02%	2.03%	2.33%	3.05%	3.66%	4.07%
SIN	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
KOR	0.00%	0.60%	1.51%	3.01%	3.45%	4.52%	5.42%	6.03%
THA	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

λ_R: regret aversion
ϕ: fraction invested in the benchmark portfolio
ψ_P: tracking error

Figure 6: Out-of-sample results

a) Aggregate indices (USA, EAFE)

	BM	TP	BayesBM	BayesMVP	MVP	EWP
Mean return	9.02	9.08	9.07	8.64	8.05	7.12
Volatility	15.60	16.23	15.14	15.03	14.67	14.75
Sharpe ratio	0.578	0.559	0.599	0.575	0.549	0.483
Turnover	0.00	2.22	1.55	1.20	1.03	0.89

b) Individual country indices (7 developed and 8 emerging markets)

	BM	TP	BayesBM	BayesMVP	MVP	EWP
Mean return	9.02	6.96	8.12	8.97	10.13	12.16
Volatility	15.60	18.28	16.48	16.80	15.14	18.60
Sharpe ratio	0.578	0.381	0.493	0.534	0.669	0.654
Turnover	0.00	6.44	4.83	4.91	3.66	2.86