The economic value of predicting correlation for asset allocation

He Huang*, Georg Keienburg, and Duane R. Stock

November 29, 2007

Abstract

While there is an extensive literature on modeling the time-varying correlation between risky assets, the economic value of predicting correlation has not yet been explored. This is particularly important for asset allocation, as the correlation between different asset classes is subject to significant time variability. In this paper, we analyze the economic value of predicting correlation in a four asset portfolio including stocks, bonds, commodity futures, and a risk free asset. We employ a simple recursive regression model based on macroeconomic series as predictive variables. Based on an extensive sample from 1974-2007, we find that accounting for predictable correlations in a mean-variance portfolio optimization is economically valuable. Our results are robust to estimation risk, varying degrees of risk aversion, short sale constraints, and different levels of transaction costs.

JEL classification: G11; G14
Keywords: predictability; market timing, performance evaluation, asset allocation, economic value, stock-bond correlation

*Corresponding author

He Huang, Graduate School of Risk Management, University of Cologne, huang@wiso.uni-koeln.de

Georg Keienburg, Graduate School of Risk Management, University of Cologne, keienburg@wiso.uni-koeln.de

Duane Stock, Price College of Business, University of Oklahoma, dstock@ou.edu

The first two authors gratefully acknowledge funding from the German Research Foundation and the latter author also thanks Le Duong for research assistance. We have benefited from helpful comments at the University of Cologne Graduate School of Risk Management seminar. The usual caveats apply.

The economic value of predicting correlation for asset allocation

Abstract

While there is an extensive literature on modeling the time-varying correlation between risky assets, the economic value of predicting correlation has not yet been explored. This is particularly important for asset allocation, as the correlation between different asset classes is subject to significant time variability. In this paper, we analyze the economic value of predicting correlation in a four asset portfolio including stocks, bonds, commodity futures, and a risk free asset. We employ a simple recursive regression model based on macroeconomic series as predictive variables. Based on an extensive sample from 1974-2007, we find that accounting for predictable correlations in a mean-variance portfolio optimization is economically valuable. Our results are robust to estimation risk, varying degrees of risk aversion, short sale constraints, and different levels of transaction costs.

1. Introduction

Astute investors should be constantly searching for ways to improve portfolio performance through return enhancement and risk reduction. While the ability to predict asset returns remains controversial, the ability to measure and predict volatility has increased dramatically in recent years. Of course, portfolio volatility depends on both the variance of individual assets and how returns on these assets are correlated with each other. While there is an extensive literature on modeling the time-varying correlation between risky assets, the economic value of predicting correlation has not yet been explored in the context of asset allocation.

In this paper, our purpose is to answer a central, yet unanswered question: what is the economic value of predicting time-varying correlation for asset allocation? That is, how much would an investor, with a given risk aversion, be willing to pay to switch from a portfolio strategy that applies a naive correlation expectation (e.g. the historical mean) to a portfolio strategy based on more sophisticated, time-varying correlation forecasts. We address this issue by investigating an investor who decides each month on the weights of four asset classes in his portfolio: stocks, bonds, commodity futures, and a risk free asset. We employ a simple recursive regression model based on lagged macroeconomic and financial information to produce forecasts for future correlation and variance. Mean-variance ex-ante optimal portfolios are formed at the end of each month based on these forecasts. This methodology only uses data genuinely available and therefore truly focuses on ex ante predictability. The actual portfolio performance during the following month, with respect to economic value, is then compared to two benchmark strategies. These benchmark strategies are a passive buy and hold strategy maintaining constant portfolio weights, and a naïve correlation timing strategy based on historical mean correlations. By comparing the utility attained when forecasting both correlation and volatility to the utility of a strategy that only forecasts volatility but merely assumes historical correlations, we can isolate the economic value of predicting time-varying correlation.

The importance of our analysis stems from the large interest that dynamic correlation forecasting models receive from practitioners and academics alike. Simple methods such as rolling historical correlations and exponential smoothing are used in practise. This is exemplified by the well-known RiskMetrics™ calculations; see J.P. Morgan (1996) for further details. More complex methods such as varieties of multivariate GARCH have been extensively investigated in the econometric literature. As examples, see Bollerslev, Engle and Wooldridge (1988), Bollerslev (1990), Kroner and Claessens (1991), and surveys by Bollerslev, Chou and Kroner (1992), and Ding and Engle (2001). A different approach is to trace back time-varying correlation to the underlying sensitivity to broad macroeconomic information. Consider for example stock-bond correlation. Since stock and bond returns are differently affected by macroeconomic conditions, such as inflation and stock price to earnings ratio, it follows that the correlation varies with the business cycle. This relation is laid out by studies such as Fama and French (1989), Connolly, Stivers, and Sun (2005), d'Adonna and Kind (2006), Zhou (2006), and Andersen et al. (2007). This business cycle relation can be used to predict future correlation using a regression on lagged macroeconomic information, as done by, e.g., Zhou (2006).

Despite this large body of literature dealing with forecasting and explaining time-varying correlation, the economic value of predicting correlation coefficients for portfolio optimization remains dubious. A sizeable literature deals with predicting correlation for optimizing a stock portfolio, see Elton and Gruber (1973), Elton, Gruber, and Ulrich (1978), and more recently Chan, Karceski and Lakonishok (1999), Engle and Sheppard (2003), Jagannathan and Ma (2003), Ledoit and Wolf (2004), and Elton, Gruber, and Spitzer (2006). Despite all advances, the consensus seems to be that a constant correlation model assuming all pair-wise correlation coefficients equal to the historical mean correlation coefficient performs as well as forecasts produced from more sophisticated models. Similarly, it is common to assume constant correlations, based on historical values, for the purpose of asset allocation. For example, Dopfel (2003) maintains that typically a constant value that is moderately positive (e.g., +0.30 to +0.50) is assumed for the stock-bond correlation. However, given the enormous variability in realized stock-bond correlation,[1] it seems plausible that accounting for time-varying correlations between broad asset classes might be more important than predicting correlation between individual stocks. Hence, predicting correlation might yield an economic value for asset allocation.

While the economic value of predicting correlation has not yet been explored with respect to asset allocation, previous research has investigated the economic value of predicting return and volatility. Marquering and Verbeek (2004) study the economic value of a monthly rebalancing between a risky stock index and a risk free asset by forecasting stock returns and volatility. We go beyond Marquering and Verbeek (2004) by introducing two more risky asset classes: long-term government bonds and commodity futures. Now one needs not only expectations about future return and variance, but also about the correlation between each of the risky assets. Fleming, Kirby, and Ostdiek (2001) investigate the economic value of volatility forecasts for a short-horizon, daily trading strategy in a stock futures, bond futures and commodity futures portfolio. Our paper expands on Fleming, Kirby, and Ostdiek (2001) by isolating the economic value of correlation forecasts from that of variance and covariance forecasts. While Fleming, Kirby, and Ostdiek (2001) demonstrate that forecasting the covariance

[1] An extensive literature has dealt with the recent “decoupling” of stock and bond returns, see Gulko (2002), Dopfel (2003), and Connoly, Stivers, and Sun (2005) among others. While the historical mean stock-bond correlation in the 1980’s and 1990’s has averaged around +0.40, it has been significantly negative throughout the current decade.

matrix between risky assets outperforms a static covariance matrix expectation, it is not clear whether the economic value results from the well-known predictability in conditional individual asset variance or from predictability in their conditional pairwise correlations.

-----------------Please insert Table 1 approximately here -----------------

We address this issue by comparing three separate strategies. These strategies are illustrated by Table 1. Strategy I $(\overline{w})$ is a passive buy and hold strategy that holds constant portfolio weights and is not based on any forecasts. Strategy II $(\hat{\sigma},\overline{\rho})$ is a naïve correlation timing strategy that only forecasts risky asset variances but assumes the correlation to equal historical mean values. Since the historical mean essentially is static,[2] we will refer to this strategy also as static correlation timing. Lastly we consider a full prediction model $(\hat{\sigma},\hat{\rho})$ that we refer to as a dynamic correlation timing strategy. This strategy accounts for time-variability in both, variance and correlation.

Based on a large sample (1974 – 2007) we find that accounting for time-varying correlation enables more optimal portfolios than applying naïve correlation expectations. In our sample, the economic gain of predicting correlation $(\hat{\sigma},\hat{\rho})$ compared to a naïve correlation timing strategy $(\hat{\sigma},\overline{\rho})$ translates into an annual management fee worth 1.26% in our base case. This demonstrates that there is clearly an economic gain from predicting correlation for asset allocation. Most importantly, we find that accounting for both, time-varying conditional asset variances and their time-varying pairwise correlations, yields the highest realized utility. Our

[2] Each month, t, we calculate the historical mean correlation on the sample 1…t. Therefore each month, the sample on which the historical mean correlation is calculated grows by only one month.

results are robust to return estimation risk, varying degrees of risk aversion, short sale constraints, and different levels of transaction costs. Our findings are important, as they justify the increasing efforts academics and practitioners alike put into finding better models for predicting the correlation between asset classes.

The remainder of this paper is structured as follows. In section 2 we develop a trading strategy for a mean-variance investor who rebalances his portfolio each month to maximize expected utility. Section 3 presents the data we use for our empirical analyses and the forecasting models. In section 4 we investigate the impact of estimation errors in correlation on optimal portfolio weights. Section 5 reports the economic value of the dynamic correlation timing strategy and compares it to the benchmark strategies. Section 6 concludes the research.

2. A dynamic mean-variance trading strategy

Consider an investor who holds a portfolio consisting of stocks, bonds, commodity futures, and a risk free asset, which he rebalances each month. At the end of each month, t, the investor maximizes expected utility, $E_t[U_{t+1}]$, for the next month, t+1, by choosing optimal portfolio weights for each asset. Assuming a quadratic utility function, the optimization problem is thus given by

$$\max_{w_{t+1}}\{E_t[U_{t+1}] = r_{t+1}^f + w_{t+1}' \cdot \hat{r}_{t+1}^e - 0.5\gamma \cdot w_{t+1}' \cdot \hat{S}_{t+1}\hat{C}_{t+1}\hat{S}_{t+1} \cdot w_{t+1}\} \qquad (1)$$

where r_{t+1}^f denotes the risk free rate next month, w_{t+1} denotes a 3x1 vector with weights of the three risky assets: stocks, bonds, and commodity futures. These assets will be denoted using the subscripts S, B, C, respectively. For brevity's sake, we will refer to the latter also simply as commodities. It follows, that the risk free asset's weight is $1-w_{t+1}'\cdot\underline{1}$. Here $\hat{r}_{t+1}^{e}$ denotes the vector of expected excess returns for next month, γ is the level of risk aversion, and $\hat{S}_{t+1}$ is a 3x3 matrix containing the expected standard deviations of the three risky assets along the main diagonal and zero otherwise: $\hat{S}_{t+1} = diag\{\hat{\sigma}_{S,t+1}, \hat{\sigma}_{B,t+1}, \hat{\sigma}_{C,t+1}\}$.

Finally, $\hat{C}_{t+1}$ denotes the predicted (t+1) correlation matrix between the three risky assets:

$$\hat{C}_{t+1} = \begin{pmatrix} 1 & \hat{\rho}_{SB,t+1} & \hat{\rho}_{SC,t+1} \\ \hat{\rho}_{SB,t+1} & 1 & \hat{\rho}_{BC,t+1} \\ \hat{\rho}_{SC,t+1} & \hat{\rho}_{BC,t+1} & 1 \end{pmatrix} \tag{2}$$

Solving the maximization problem yields

$$w_{t+1}^{*} = \gamma^{-1} \cdot \hat{S}_{t+1}^{-1} \cdot \hat{C}_{t+1}^{-1} \cdot \hat{S}_{t+1}^{-1} \cdot \hat{r}_{e,t+1} \tag{3}$$

for risky asset ex-ante optimal weights with any remainder assigned to the risk free asset. As long as the standard deviation forecast matrix $\hat{S}_{t+1}$ and the correlation forecast matrix $\hat{C}_{t+1}$ are positive definite, the vector of optimal weights for the three risky assets w_{t+1}^{*} can be analytically solved by calculating the inverse of $\hat{S}_{t+1}$ and $\hat{C}_{t+1}$.[3]

If we assume that short selling is not feasible for the investor, the optimal weights become $w_{i,t+1}^{ns} = \max\{0, w_{i,t+1}^{*}\}$ for each of the risky assets $i \in \{S,B,C\}$. If we furthermore

[3] Since $\hat{S}_{t+1}$ contains only the (positive) standard deviations along the main diagonal and zeros otherwise, it is positive definite by definition. In order to assure that the expected correlation matrix $\hat{C}_{t+1}$ is positive definite we propose to either not rebalance the portfolio in case of a non-positive definite correlation matrix or, alternatively, to apply techniques such as Higham (2002) to find the nearest positive-definite correlation matrix. However, in our empirical analysis the correlation matrices $\hat{C}_{t+1}$ are positive definite in each month.

disallow borrowing at the risk free rate, and the optimal no-short sale weights exceed one, then we scale them to one by setting $w_{i,t+1}^{ns\&nb} = 1/(w_{t+1}^{ns}{}' \cdot \underline{1}) \cdot w_{t+1}^{ns}$.

All optimal weights hinge on the conditional time t forecasts for t+1's excess returns, asset variances, and cross asset correlations. Since the predictability of returns is questionable (see, among others, Best and Grauer (1991), Bossaerts and Hillion (1999), Chan, Karceski, and Lakonishok (1999), and Ang and Bekaert (2007)) we will assume constant excess return forecasts each month in our empirical analysis, i.e. $\hat{r}_{t+1}^{e} = \overline{r}^{e}$. We choose a value for $\overline{r}^{e}$ based on data that was available before the beginning of our sample, and thus, we focus on ex-ante predictability. We address the issue of estimation risk by also simulating different values of $\overline{r}^{e}$. Further, we will approximate the variance for each of the three risky assets and the three pairwise correlation forecasts with fairly simple functions of observable macroeconomic and financial variables. Let x_t denote a vector of variables that are observed at time t, including a constant. These variables are used to predict the variance of excess returns for each asset i and also the correlation between excess returns $r_{i,t+1}^{e}$ and $r_{j,t+1}^{e}$ of each of the two asset combinations $ij \in \{SB, SC, BC\}$.

Specifically, we assume for each of the three correlations

$$\rho_{ij,\tau+1}^{F} = x_{\tau}{}'\beta_{ij} + \varepsilon_{ij,\tau+1}, \qquad \tau = 1,2,....,t, \tag{4}$$

where $E[x_{\tau}\varepsilon_{ij,\tau+1}] = 0$ and β_{ij} is a vector of unknown coefficients. $\rho_{ij,\tau+1}^{F}$ denotes the Fisher transformation of the respective correlation coefficient $\rho_{ij,\tau+1}$, which transforms the correlation coefficient from the range of (-1, +1) to $(-\infty, +\infty)$.[4] It is a continuous and monotonic function

[4] Fisher transformations of correlation coefficients are known to have standardized normal distribution

and is defined as

$$\rho_{ij,\tau+1}^{f}=\frac{1}{2}\ln\left(\frac{1+\rho_{ij,\tau+1}}{1-\rho_{ij,\tau+1}}\right). \tag{5}$$

In the empirical application, the parameters β_{ij} are estimated recursively by ordinary least squares regression, using information from periods 1 to t. We thus use a recursive regression where estimation is based on a window of expanding size. With the estimate $\hat{\beta}_{ij,t}$ for β_{ij}, we obtain the conditional forecast for the Fisher transformed correlation in period t+1 as

$$\hat{\rho}_{ij,\tau+1}^{f}=x_t\,'\hat{\beta}_{ij,t} \tag{6}$$

which can be reverse transformed into the estimate for the t+1 correlation via

$$\hat{\rho}_{ij,t+1}=\frac{\exp(2\hat{\rho}_{ij,t+1}^{f})-1}{\exp(2\hat{\rho}_{ij,t+1}^{f})+1} \quad . \tag{7}$$

This forecast is updated every period because new information becomes available $(x_{t+1,}x_{t+2},...)$ and because the coefficient estimate $\hat{\beta}_{ij,t}$ is updated as well. Information about future values of $\rho_{ij,t+1}$ or x_t is not used at any point in time, thus a comparison of $\hat{\rho}_{ij,t+1}$ and with the realized correlation $\rho_{ij,t+1}$ provides a genuine measure of ex ante predictability.

In a similar fashion, we apply a linear model for the conditional variance of each risky asset's excess returns, which is explained from a set of variables, y_t, and is potentially different from x_t. Besides macroeconomic and financial variables, y_t can contain lagged dependent variables. Again, the coefficients are estimated recursively using information from observation 1 to t, applying OLS to

asymptotically. It converges to its asymptotic distribution much faster than a lot of other alternative transformations, see Anderson (1984) for details. This approach has recently also been used in d'Addona and Kind (2006).

$$\log \sigma^2_{i,\tau+1} = y_\tau{}'\delta_i + \xi_{i,\tau+1}, \qquad \tau=1,2,...,t, \tag{8}$$

where $\sigma^2_{i,\tau+1}$ denotes asset i's conditional variance in period $\tau+1$ and $E[y_\tau \xi_{i,\tau+1}]=0$. The conditional forecast for period t+1 is thus given by

$$\hat{\sigma}^2_{i,t+1} = \exp(y_t{}'\hat{\delta}_{i,t}) \tag{9}$$

In line with Marquering and Verbeek (2004) we use this approach in order to assure that predicted variance is positive. Again, all forecasts for t+1 are based on information up to period t, and are therefore truly ex-ante forecasts.

The models for conditional variance and correlation are deliberately chosen to be simple, linear models with fixed selections of which variables x_t and y_t are employed for prediction. While recent literature has proposed more complicated non-linear models such as GARCH-type specifications, these techniques were most probably neither available nor computationally feasible for the average investor during most of our sample period. We do not claim that our specification is either "correct" or superior to other specifications, therefore the economic value of our trading strategy might be a conservative estimate. Yet, that does not subtract from our proposition that predicting time-varying correlation is economically valuable compared to a naïve, static correlation forecast.

The main purpose of this paper is to investigate the economic value of dynamic time-varying correlation forecasts. To assess the value of any dynamic strategy, we ask this question: Given a certain degree of risk aversion, γ, how much would an investor be willing to pay to switch over from a passive benchmark strategy? Our benchmark strategy is a constant-weight strategy that holds equal weights in each of the four assets, i.e. the "passive strategy" $(\overline{w})$. We determine the maximum fee, as a percentage of the invested amount that would make an investor indifferent between the passive and the dynamic strategy. To calculate that fee, recall that the ex-post utility of a mean-variance investor in any month t+1 is given by

$$U_{t+1} = r_{t+1}^{f} + w_{t+1}' \cdot r_{t+1}^{e} - 0.5\gamma \cdot w_{t+1}' S_{t+1} C_{t+1} S_{t+1} w_{t+1} - \text{trading costs} \qquad (10)$$

where S_{t+1} and C_{t+1} denote the matrices containing realized standard deviation and the realized correlation matrices, respectively. In line with Andersen et al. (2001) we define each of the realized risky asset variances in any month t+1 as

$$\sigma_{i,t+1}^{2} = \sum_{d=1}^{N_{t+1}} r_{i,d}^{e\ 2} \qquad (11)$$

where N_{t+1} denotes the number of days in month t+1, and $r_{i,d}^{e\ 2}$ denotes the squared excess return of asset i on day d. Similarly the realized monthly correlation between risky asset i and j is defined as

$$\rho_{ij,t+1} = \frac{\sum_{d=1}^{N_{t+1}} r_{i,d}^{e} \cdot r_{j,d}^{e}}{\sqrt{\sum_{d=1}^{N_{t+1}} r_{i,d}^{e\ 2} \cdot \sum_{d=1}^{N_{t+1}} r_{j,d}^{e\ 2}}} \qquad (12)$$

This approach assumes zero as the expected daily excess return, rather than applying the noisy monthly mean value. We make this choice, as daily expected returns are essentially zero. This prevents extreme return realizations to bias our measure. Our approach follows that of Fleming, Kirby and Ostdiek (2001) and Connolly, Stivers and Sun (2005, 2007) among others.

Since each active trading strategy is associated with transaction costs, we also need to account for the transaction costs from rebalancing the risky assets' weights. For the sake of simplicity, we assume trading costs to be symmetrical for buy and sell transactions and we assume that transaction costs are a fixed, linear proportion of the amount traded. Let tc denote a

3x1 vector containing the proportional transaction costs for each of the risky assets. We can measure the ex-post realized average utility of a trading strategy k throughout our sample as

$$\bar{U}_k = \frac{1}{T}\sum_{t=0}^{T-1}\left(r^f_{t+1} + w_{t+1}'\cdot r^e_{t+1} - 0.5\gamma\cdot w_{t+1}'S_{t+1}C_{t+1}S_{t+1}w_{t+1} - \left|w_{t+1}-w_t\right|'\cdot tc\right).^{5} \qquad (13)$$

We express a dynamic strategy's economic value as the maximum monthly fee Δ_k that would make a mean-variance investor indifferent between the dynamic strategy k and a passive equal-weights strategy $(\bar{w})$. To determine Δ_k we need to solve the equation:

$$\frac{1}{T}\sum_{t=0}^{T-1}\left(r^f_{t+1} + w_{t+1}'\cdot r^e_{t+1} - 0.5\gamma\cdot w_{t+1}'S_{t+1}C_{t+1}S_{t+1}w_{t+1} - \left|w_{t+1}-w_t\right|'\cdot tc - \Delta_k\right) = \bar{U}_{\bar{w}} \qquad (14)$$

It follows, that Δ_k can be easily expressed as the difference between two average utility levels, $\bar{U}_k$ and $\bar{U}_{\bar{w}}$. Also, we can express the economic value of a dynamic strategy k over an alternative dynamic strategy l as the difference between Δ_k and Δ_l, which corresponds to the difference between the two average utility levels, $\bar{U}_k$ and $\bar{U}_l$.

3. Data and forecasting methodology

The primary data used is in this study is daily excess returns for stocks, bonds and commodity futures. The sample covers the time period January 1974 until March 2007. More specifically, we include the dividend adjusted return on the S&P 500 index (denoted by $r_{S,t}$), the

[5] In line with Marquering, Verbeek (2004) we neglect weight changes during month t due to asset price movements. That is, we calculate the traded amount as $\left|w_{t+1}-w_t\right|$, rather than subtracting the risky assets' weights at the *end of* period t from next month's weights. This simplification only marginally changes numerical results, not qualitatively. But it allows us to consider an equal-weight strategy that does not incur any trading costs.

return on a long term (10year) US Treasury bond (denoted by $r_{B,t}$) and the return on the Commodity Research Bureau futures index (CRB) (denoted by $r_{C,t}$)[6]. We will refer to these returns as the stock, bond and commodity returns. All returns were converted to excess returns (denoted by $r^e_{S,t}$, $r^e_{B,t}$, $r^e_{C,t}$) using the monthly three month Treasury bill rate (denoted by r_t^f).[7] Also, the three month Treasury bill is included in our analysis as the risk free asset, that we will refer to as cash. We adjust S&P 500 index returns for dividend yields and adjust for weekends, holidays and trading days when at least one of the markets was closed. S&P 500 index and dividend yield data was provided by Thomson Financial Datastream, while Treasury bond yields and 3 month Treasury bill returns were obtained from the Federal Reserve Bank of St. Louis. The CRB index was provided by the Commodity Research Bureau and obtained through Datastream.

Table 2 presents some summary statistics for the monthly stock, bond and commodity returns and excess returns. For the full sample period, February 1974 to March 2007, the dividend adjusted excess return on the S&P 500 exceeds the risk free rate by 0.55% per month on average. The average excess return on the ten-year Treasury bond is approximately 0.22% per month, and the CRB index excess returns is slightly lower at 0.21% per month. Returns for the three month T-bill is on average 0.48% per month.

-----------------Please insert Table 2 approximately here -----------------

[6] The CRB is the oldest commodity futures index. It is an equal weighted index that assumes a long-only position in a variety of commodity futures contracts that are regularly rolled over.

[7] Since a futures contract does not require an initial investment, we assume that the investor of the CRB holds a collateral position in three month T-Bills in the amount of the futures contracts' par value. Therefore we consider the CRB index return to be excess returns.

As expected, the standard deviation for monthly stock excess returns is highest at 4.40%. The average monthly volatility for bonds is 2.45% and 3.22% for commodities. Interestingly, the standard deviations of all three assets are lower for the second sub-period despite higher returns.

Monthly asset correlation as presented in Panel D is calculated using monthly excess returns. For stocks and bonds, we find an average correlation of ρ_{SB} = 0.18 over the full sample, while ρ_{SC} = 0.03 and ρ_{BC} = -0.13. Noticeable is the shift in correlation over the two sub samples. While correlation between stocks and bonds is highly positive in the years 1974-1990, it becomes slightly negative in the second half of our sample. This shift has received some attention lately, see Gulko (2002) and Connolly, Stivers, and Sun (2005, 2007) among others. It is commonly explained with the flight-to-quality hypothesis, where rising stock market uncertainty tends to decrease the co-movement between stock and bonds. The sign of the correlation between stocks and commodities also varies. Note that ρ_{BC} is always negative throughout both subsamples. Figure 1 shows the 12-month moving average of all three pairwise asset correlations.

-----------------Please insert Figure 1 approximately here -----------------

As discussed in the previous section, we assume an investor who forms expectations each month about next month's risky asset returns, volatility, and pairwise correlations in order to choose optimal portfolio weights. In our empirical analysis we allow for a calibration period of 71 months for the investor to form his initial expectations. Therefore the forecast period is January 1980 – March 2007. All performance evaluations will be applied to this forecast sample.

As stated in section 2, evidence for return predictability is questionable, and it is connected with substantial estimation risk. Therefore we do not apply any dynamic forecasts for returns. Instead we apply constant values for expected returns, which an investor could have plausibly chosen at the beginning of our sample. At the beginning of our forecast sample in 1980, a potential choice could have been to use the long term return forecast for stocks and long-term government bonds[8] presented by Ibbotson and Sinquefield (1976). Ibbotson and Sinquefield (1976) develop a simulation model to forecast probability distributions of returns for stocks and bonds for the period 1976-2000. This simulation takes into account not only historical data, but also inflation expectations and the yield curve. Using their long term return predictions, an investor in 1980 would have expected a constant monthly excess return for stocks of $\overline{r}_S^e$=0.503%, and $\overline{r}_B^e$=0.099% for long term government bonds.[9] While Ibbotson and Sinquefield (1976) provide a natural choice for stock and bond return expectations, we are unaware of any 1970's study which predicts long term commodity returns. Therefore we have to rely on past returns to form expectations for commodity returns. However the 1970's produced enormous economic changes, such as the shift in Federal Reserve interest rate policy, the elimination of the gold standard, and most importantly for commodity returns, the oil crisis in 1973. Therefore we exclude all observations past the oil crisis, and apply the average monthly excess return $\overline{r}_C^e$=0.092% for the CRB index over the time period September 1956 until December 1972 as our constant expected return.[10] We address any potential ex-post selection bias, arising from

[8] Ibbotson and Sinquefield (1976) consider long term government bonds of 20 years constant maturity while we use bonds of ten year maturity for our analysis. However, given the very high co-movement between 10- and 20-year treasuries, this seems like a reasonable choice.
[9] Ibbotson and Sinquefield (1976) report expected annual returns of 13.0% for stocks, 8% for bonds and 6.8% for T-bills.
[10] The CRB futures index is not available before September 1956.

choosing the return forecasts, by also conducting our empirical analysis with various simulated return expectations.

While we rely on constant return expectations, we employ a simple recursive regression model to forecast asset correlation and variance as outlined in section 2. Pesaran and Timmerman (1995) propose a simple model to predict stock returns by relying on publicly available macroeconomic and financial data. This is appealing, since the business cycle has been identified as a major determinant for stock price movement as early as in Angas (1936). For bond returns, Fama and French (1989) find that variation in bond returns can be explained by the same macroeconomic variables that possess predictive power for stock returns. Only recently, Gorton and Rouwenhorst (2006) analyze properties of commodity futures as an asset class. An important finding is that the business cycle is a major determinant for commodity returns. Given the strong evidence for the business cycle's effect on each individual asset return, it is reasonable to assume a relation between macroeconomic variables and pairwise correlations between the aforementioned asset classes as well as volatility in their returns.

We assume that a simple model including macroeconomic and financial variables has some ability to predict asset correlation. The explanatory variables we include in the prediction model are based on those proposed by Pesaran and Timmerman (1995): the last two lags for the three month Treasury bill yield, the price-earnings ratio on the S&P 500, dividend yield on the S&P 500, inflation, industrial production, the last two lags for the 12-month Treasury bill yield, and monetary growth. Additionally, since Fama and French (1998) stress the importance of default spreads as a business cycle indicator, we add the default spread, defined as Moody's Baa corporate bond yield minus Aaa yields. While these variables were originally chosen to account only for time-varying stock and bond returns, it seems like a good choice for our purpose, given that all these variables were chosen to represent the macroeconomic cycle. We stress that it is not the goal of this study to choose the best prediction model, particularly not to make that choice based on ex-post knowledge of a model's performance. Rather, we strive to show that any reasonable forecast for asset correlations increases portfolio utility compared to a naïve, static forecast. Since macroeconomic data is frequently revised, see for example Hautsch and Hess (2007) for a discussion of the revision process for employment data, we include macroeconomic variables only with two months lags. Since it is a safe assumption that after two months the data is not further revised, as maintained by Marquering and Verbeek (2004), this approach ensures that we use the same macroeconomic data that was available to an investor during our sample.

-----------------Please insert Table 3 approximately here -----------------

Table 3 represents the in-sample regression for the Fisher transformed asset correlation coefficients on the lagged macroeconomic and financial variables. We analyze whether there is in-sample explanatory power in macroeconomic information in order to justify the use of such a forecasting model. The ultimate test for the economic value of a strategy that applies this model, however, has to be an out-of-sample test. Regression results for the stock-bond correlation indicate a high in sample explanatory power for the macroeconomic series which we employ. The explanatory power for stock-commodity and bond-commodity correlation is somewhat weaker. This is indicated by a high adjusted R^2 of about 0.36 for the stock-bond correlation and about 0.13 for both stock-commodities and bond-commodities correlation. Although not all explanatory variables are individually significant, all three F-statistics in Table 3 strongly reject the hypothesis that the set of employed macroeconomic variables contains no explanatory power.

Because there is no a priori reason to exclude any variables that may forecast co-movement of returns from predicting second moments, we include all aforementioned variables in the variance regressions. In order to account for the heteroscedasticity known from financial time series data of return volatility, one lag of the dependent variable is additionally included. Table 4 reports in-sample regression results for the S&P 500, 10-year T-Bond and the CRB Futures Index return variance. Again, the results indicate high explanatory power for most of the employed independent variables. Due to the importance of the lagged dependent variable, R^2 values are even higher than for the correlation regressions.

-----------------Please insert Table 4 approximately here -----------------

While the in-sample regression results indicate a high degree of explanatory power for the macroeconomic series we employ, it is not clear what the out-of-sample predictability is. Also, we again stress that ultimately we test the model's performance by investigating its economic value in a utility maximizing framework. Nevertheless, it is interesting to look at the model's out-of sample statistical performance, as statistical significance would suggest practical application.

In order to test the out-of-sample predictive power for the correlation regression models, we run recursive regressions with a growing window size as described in section 2. For each month, t, during the forecast sample (Jan 1980 – Mar 2007) we predict next month's, t+1, pairwise correlations. We then compare the root mean squared error ($RMSE = \frac{1}{N_t}\sum\sqrt{\left(\rho_{t+1} - \hat{\rho}_{t+1}\right)^2}$) of the dynamic macroeconomic model ($\hat{\rho}$) to a naïve correlation expectation $(\bar{\rho})$; the mean historical correlation which is based on ex-ante periods 1 through t. We report results for this out-of-sample test in Table 5.

-----------------Please insert Table 5 approximately here -----------------

The dynamic model based on macroeconomic variables produces much lower statistical errors in terms of RMSE, whereas the performance of the naïve forecast is comparably worse for all three correlations. The difference is more pronounced for the stock-bond correlation, as the variability for the stock-bond correlation is much larger than for the other two correlations, as can be seen in Figure 1. These results indicate that dynamic correlation forecasts based on a regression on macroeconomic variables yields superior predictions for future correlation than a naïve, static forecast. Whether these better predictions will result in an economic gain will be addressed in section 5.

4. Marginal effects of predicting correlation on optimal portfolio weights

The goal of this paper is to investigate whether there is economic value in predicting the correlation between different asset classes. One step towards this goal is to analyze how ex-ante optimal portfolio weights in a mean-variance framework depend on correlation forecasts. We therefore investigate how optimal weights change when there is a marginal adjustment in predicted asset correlation. This analysis shows the effect of prediction errors in estimating correlation. As outlined in section 2, optimal portfolio weights depend on the expected characteristics (return, volatility, correlation) of all risky assets in a portfolio. The connection between predicted values and optimal portfolio weights is more straightforward with respect to an asset's expected return and volatility. If an asset's expected return is adjusted upwards, its weight in the ex-ante optimal portfolio will increase. Similarly, if an asset's expected volatility increases, its weight will decrease. On the other hand, it is not evident how the expected correlation between two risky assets affects their portfolio weights, and by how much.

Recall that equation (4) represents the vector of optimal asset weights, w_i^*, as a function of all expected characteristics of all risky assets. For the sake of brevity, we restrict our analysis to the stock weight, w_S^*, in the ex-ante optimal portfolio. We can rewrite equation (4) to show w_S^* as a function of the other optimization inputs:

$$w_S^* = \frac{1}{\gamma \cdot \det(\hat{\Sigma})} \cdot \begin{bmatrix} (\hat{\sigma}_B^2 \cdot \hat{\sigma}_C^2 - \hat{\rho}_{BC}{}^2 \cdot \hat{\sigma}_B^2 \cdot \hat{\sigma}_C^2) \cdot \hat{\mathrm{r}}_{\mathrm{S}}^{\mathrm{e}} \\ +(\hat{\rho}_{SC} \cdot \hat{\sigma}_S \cdot \hat{\sigma}_C^2 \cdot \hat{\rho}_{BC} \cdot \hat{\sigma}_B - \hat{\rho}_{SB} \cdot \hat{\sigma}_S \cdot \hat{\sigma}_B \cdot \hat{\sigma}_C^2) \cdot \hat{\mathrm{r}}_{\mathrm{B}}^{\mathrm{e}} \\ +(\hat{\rho}_{SB} \cdot \hat{\sigma}_S \cdot \hat{\sigma}_B^2 \cdot \hat{\rho}_{BC} \cdot \hat{\sigma}_C - \hat{\rho}_{SC} \cdot \hat{\sigma}_S \cdot \hat{\sigma}_C \cdot \hat{\sigma}_B^2) \cdot \hat{\mathrm{r}}_{\mathrm{C}}^{\mathrm{e}} \end{bmatrix} \qquad (15)$$

where $\hat{\Sigma} = \hat{S}' \cdot \hat{C} \cdot \hat{S}$ denotes the expected covariance matrix. Note that $\hat{\Sigma}$ itself is a function of all predicted asset characteristics. It follows that the function w_S^* in equation (16) is neither a linear function nor a monotonic function, particularly not in the level of any expected correlation. If w_S^* is derived with respect to any of the predicted asset characteristics,[11] say $\hat{\rho}_{SC}$, $\frac{\partial w_S^*}{\partial \hat{\rho}_{SC}}$ will yield the marginal impact of adjusting $\hat{\rho}_{SC}$ on the optimal stock weight w_S^* with all other forecasts held equal. Likewise one can derive the marginal impact for expected returns, e.g. $\frac{\partial w_S^*}{\partial \hat{r}_S}$, and volatility, e.g. $\frac{\partial w_S^*}{\partial \hat{\sigma}_S}$. We will refer to these variables as the marginal impact of prediction errors

[11] Since the covariance matrix Σ is positive definite, this function is differentiable.

(on optimal weights) and report these values in Table 6.

-----------------Please insert Table 6 approximately here -----------------

Table 6 reports the marginal impact on the optimal stock weight w_S^* for prediction errors in expected return $\hat{r}_s^e$, expected standard deviation $\hat{\sigma}_s$, and expected stock-bond ($\hat{\rho}_{SB}$) and stock-commodity ($\hat{\rho}_{SC}$) correlation. In column 1, all expected values are set to equal our sample mean. It shows that for the given setup a marginal $1/100$ standard deviation increase of the expected stock correlation with bonds or with commodity futures results in approximately a 0.001 drop of the stock weight in the ex-ante optimal portfolio. All other optimization inputs are held equal. This is comparably lower than the impact of an adjustment in expected volatility or returns. In our setting, a marginal increase of the expected stock volatility will lead to about a five times higher adjustment of optimal weights, while an adjustment of expected return of the same magnitude would lead to a fifty times larger adjustment. However, the marginal impact of correlation forecasts varies with the other input values, particularly with the level of expected correlation. In column 2, we decrease the level of expected stock-bond correlation by 0.30. Now the marginal impact of prediction errors in expected stock-bond correlation is more than twice as high (-0.0022), while the other marginal effects are about the same. Similarly the marginal impact of predicted stock-commodity futures correlation doubles as the level of predicted stock-commodity futures correlation is lowered by 0.30, see column 3.

This example demonstrates that it is not possible to generalize how errors in correlation forecasts will impact the weights assigned to each asset. The ex-ante optimal weight in a mean-variance framework is neither linear nor monotonic in the expected correlations, but it is a complex function of all assets' expected characteristics. It follows that statistical forecasting errors cannot be directly translated into economic value, as marginal adjustments of expected correlations lead to non-linear weight adjustments of varying magnitude. However, the impact on portfolio weights for prediction errors in correlations is typically smaller than the impact of return and volatility estimations. But this marginal impact varies with, e.g., the level of correlation forecasts. Given the complex link between predicted correlation and optimal ex-ante weights, it is imperative to look beyond statistical measures for correlation forecasting models. Therefore the remainder of this paper shall investigate the economic value of predicting correlation in more detail.

5. Economic value of dynamic correlation timing

Good statistical out-of-sample prediction power does not necessarily lead to higher economic value in a utility optimization framework. Therefore this section shall investigate the economic value of predicting correlation. As stated in section 2, economic value of a dynamic strategy could be best described as the amount of money an investor would be willing to pay to switch over from a passive, constant-weight strategy. This section compares how a dynamic correlation timing strategy $(\hat{\sigma},\hat{\rho})$ that accounts for time-varying correlation and variances compares to two benchmark strategies. One is a passive, constant-weight strategy $(\overline{w})$ that maintains equal weights in each asset and the second is a naïve correlation timing strategy $(\hat{\sigma},\overline{\rho})$, that only predicts individual asset volatility and assumes future correlations to equal the historical mean. By comparing the dynamic correlation strategy with the naïve correlation strategy, we can point to the economic value of predicting correlation, as both strategies only differ in their expected correlation. We will begin this analysis by presenting a base case, which is based on rather typical assumptions for transaction costs, the investor's risk aversion, and short sale constraints (section 5.1). Then, we vary these assumptions in order to test the robustness of our results (section 5.2). Also, we address the issue of return estimation risk by employing bootstrapped return estimations (section 5.3). Finally we compare the economic value of correlation predictions to that of predicting individual asset volatility (section 5.4).

5.1 The base case

In the base case, we assume a moderately risk averse investor with a risk aversion coefficient, γ, of 5. Each month he rebalances his portfolio to maximize his ex-ante expected utility based on expected asset returns, variances, and correlations, as described in section 2. Furthermore, we assume a medium level of transaction costs of 0.5% for stocks and 0.1% for both bonds and commodity futures.[12] Those transaction costs are not actively taken into account for portfolio optimization, but rather reduce the ex post utility according to equation (14). The investor faces neither short sale nor borrowing constraints. Monthly expected stock and bond excess returns are assumed to be constant according to the expectations derived in Ibbotson and Sinquefield (1976). Expected commodity excess returns equals historical long term average return for the time period September 1956 to December 1972. Forecasts for variances and correlations are based on recursive regressions on macroeconomic and financial variables, as outlined in section 3.

---

[12] We adopt the trading costs for stocks and bonds from Pesaran and Timmermann (1995). Since futures trading incurs very low transaction cost 0.1% is a safe assumption for commodities.

-----------------Please insert Table 7 approximately here -----------------

Table 7 reports the results in the base case setting for the three benchmark strategies during the full sample January 1980 - March 2007 (Panel A), during the first half (Panel B) and the second half of the sample (Panel C). For the full sample period we additionally report results for four buy and hold strategies that invest in only one single asset. That is, the investor holds only stocks, only bonds, only commodities and only cash (T-bills). Since all forecasts are based on ex-ante values, the comparison is strictly out-of-sample. Column 1 reports the average monthly realized utility according to equation (14). For example, an investor following a buy and hold strategy in stocks would have attained an average monthly utility of 0.0056 over the full sample period. This utility value can be interpreted as equivalent to a risk free investment with a total return of 0.0056 per month. The utility for a “stocks only” strategy is slightly higher than for a 100% cash investment. The “commodities only” strategy performs worst among the passive, single asset strategies while a “bond-only” investment performs best. As expected, a passive strategy, i.e. a buy and hold strategy that maintains 25% in each of the assets, outperforms all single asset investments. Therefore, since single-asset strategies are obviously not a reasonable choice for a mean-variance investor, they will be excluded from any future analysis. The passive, equal-weights strategy shall serve as our benchmark for a passive buy and hold strategy and we refer to it simply as the passive strategy $(\bar{w})$ hereafter. It will also be the benchmark strategy for calculating the economic value of the dynamic strategies.

From a mean variance investor’s point of view, active timing of portfolio volatility clearly outperforms passive strategies in terms of utility, even if the expected correlation is based on a naïve prediction. For example, the naive correlation timing strategy $(\hat{\sigma},\overline{\rho})$ yields a utility gain over the passive strategy of 0.00084 (0.00767 - 0.00683) per month. We refer to this value as the economic value of the dynamic strategy and report the values in column 2. Using the dynamic correlation timing strategy $(\hat{\sigma},\hat{\rho})$, which accounts for both time-varying correlation and volatility, results in even higher economic value of 0.0019 (0.0087 - 0.0068). To put this value into perspective, an investor in the given setting would be willing to pay an annual management fee of 2.3% to switch from a passive, equal-weights investment to this dynamic strategy. This clearly outperforms the naive correlation timing strategy, which would be only worth a 1.0% annual fee to the same investor. Since the difference between the dynamic strategy and the naïve strategy lies only in the dynamic correlation prediction, the additional economic value over the naïve correlation strategy of 1.3% per year could be interpreted as the economic value of predicting correlation.

Column 3 reports the monthly portfolio excess returns while column 4 reports portfolio return standard deviation. Among the buy and hold strategies, the stock-only portfolio yields the highest return (0.65%) but also highest standard deviation (4.27%). For the volatility timing strategies, portfolio returns are highest for the dynamic correlation timing strategy at 0.67% per month. This is even higher than a stock-only investment, but still comes at considerably lower volatility (3.30%). Column 5 reports the Sharpe ratio, defined as the ratio of the mean excess return on the portfolio divided by the standard deviation. The highest Sharpe ratio (0.2021) for the dynamic correlation strategy suggests that investors following this strategy are better rewarded for the risk that they take. Contrary to the utility ranking in which the naïve correlation strategy dominates the passive strategy, the passive strategy yields a higher Sharpe ratio (0.1814) than the naïve correlation strategy (0.1718). This contradiction might be due to the following reasons. First, the optimization function maximizes mean-variance utility, as this paper is concerned with utility, and does not include the Sharpe ratio as a target value for optimization. Second, it is known that the Sharpe ratio is an inappropriate measure for evaluating dynamic strategies as the risk of a dynamic strategy is typically overestimated by the unconditional sample standard deviation, see Marquering and Verbeek (2004). Still, the dynamic correlation strategy yields the highest Sharpe ratio among the benchmark strategies.

Panel B reports results for the first subsample spanning January 1980 until August 1993, while Panel C reports results for the second half, September 1993 until March 2007. Remember that the summary statistics reveal significant changes between the two subsamples. The second subsample is marked by much higher excess returns and lower volatility for the risky assets and a much lower risk free rate. This change is well exemplified by the passive, equal-weights strategy's performance. The passive strategy yields an average monthly utility of 0.0080 during the first subsample, and marginally outperforms the naïve correlation timing strategy. But in the second subsample it performs far worse than both active volatility timing strategies and only yields an average utility of 0.0057. Despite the radical changes between the two subsamples, the dynamic correlation timing strategy consistently outperforms both the passive and the naïve correlation strategy in terms of economic value and Sharpe ratio.

In order to visualize the time-varying weights in the ex-ante optimal portfolio we plot optimal portfolio weights for each asset according to the dynamic correlation strategy during the sample period.

-----------------Please insert Figure 2 approximately here -----------------

Figure 2 reveals that there is considerable time-variation in portfolio weights. This particularly applies for the stock weight, as stocks have the highest volatility on average but also the highest variability in volatility. Since all risky assets' risk, return, and correlation expectations are favorable, cash is shorted throughout most of the sample, with average cash weight being -35.4%. The risky assets' weights in the ex-ante optimal portfolio are 76.6%, 23.0%, and 35.9% on average for stocks, bonds, and commodities, respectively.

5.2 Sensitivity analysis

In the previous subsection we have considered a base case, which assumes realistic and moderate assumptions for transaction costs, the investor's risk aversion, and trading constraints. Next, we test if the main result – that there is economic value in predicting asset correlations – is robust to variations of the base case. We report results for various configurations in Table 8.

-----------------Please insert Table 8 approximately here -----------------

When considering a dynamic trading strategy, transaction costs are an important consideration. Compared to the benchmark of a passive, constant-weights strategy $(\overline{w})$, an investment strategy based on time-varying forecasts is likely to incur considerably higher transaction costs and may not be as profitable as the constant-weights strategy when transaction costs are appropriately taken into account. This consideration is especially important when dealing with more than one asset. For example the dynamic correlation strategy $(\hat{\sigma},\hat{\rho})$ incurs total portfolio reshifting of 46.9% on average each month, while the naïve correlation strategy $(\hat{\sigma},\overline{\rho})$ only reshifts 36.0%. The passive strategy assumes constant weights and therefore does not incur any reshifting. In the basic setting, transaction costs are assumed to be at a medium level of 0.5% for stocks and 0.1% for bonds and commodity futures. We additionally report the impact of varying transaction costs on the realized utility of the presented trading strategies in Panel A of Table 8. While the passive strategy is unaffected by transaction costs, both volatility timing strategies' performances change. Assuming zero transaction costs or low costs of 0.1% for all assets, the volatility timing strategies clearly outperform the passive strategy. The dynamic correlation strategy has the highest realized utility of 0.0100 for zero and 0.0095 for low transaction costs, which is considerably more than for the naïve correlation strategy. As transaction costs increase, realized utility for the active strategies decreases. Assuming high transaction costs of 1% for stocks, 0.5% for bonds and 0.1% for commodity futures, the passive strategy (0.0068) outperforms the naïve correlation strategy (0.0063). Still, the dynamic correlation strategy (0.0070) yields the highest utility among the benchmark strategies.

We additionally test if our results are applicable for different representative investors. While all risk-averse investors seek to avoid risk, different investors have different levels of risk aversion. In the base case, we have considered a moderately risk averse investor represented by a risk aversion coefficient of $\gamma=5$. Panel B of Table 8 additionally reports the realized utility for varying levels of risk aversion. The average realized utility of the dynamic correlation strategy clearly trumps the performance of naive correlation timing. For an aggressive investors ($\gamma=1$) the average monthly utility difference is as high as 0.0047 (0.0242-0.0195). This translates into an additional annual management fee of 6.6% (22.3%-15.7%) that the aggressive investor would pay for time-varying correlation predictions. A conservative investor with $\gamma=10$, on the other hand, values both active volatility timing strategies less. Such an investor would be willing to pay an annual management fee of 0.6% for the dynamic correlation timing strategy, while he would be indifferent between a passive investment and the naïve correlation strategy. Additionally, the better performance for a dynamic correlation strategy compared to both benchmark strategies even holds for extremely risk averse investors, with γ of as high as 20.

Finally, as institutional investors often face trading constraints, Panel C reports results when short sales of risky assets are not allowed (middle column) and when additionally borrowing cash is not permitted (right column). While the passive strategy is not affected by either constraint, both volatility timing strategies are affected. The short sale constraint for risky assets – stocks, bonds, and commodities – seems not to be a big concern. While the dynamic correlation timing strategy's performance remains virtually unchanged, the naïve correlation strategy's performance even marginally increases after imposing short sale constraints. Both active strategies, however, seem to rely heavily on borrowing cash. After imposing short sale and borrowing constraints both active strategies' performances drop considerably. Still, the dynamic correlation strategy yields the highest average utility (0.0075), followed by the naïve strategy with an average utility of 0.0073, compared to the passive strategy's average utility of 0.0068.

5.3 Impact of return estimation risk

As explained in section 3 both active volatility timing strategies predict volatility whereas the dynamic correlation strategy additionally forecasts correlation. The expected return on the other hand, has been set to a constant value for the previous analyses. The stock and bond return expectations are taken from Ibbotson and Sinquefield (1976). Expected commodity returns were assumed to equal long term historical values for the CRB index. Even though we assume that our returns would have been a reasonable ex-ante assumption, the choice of the expected return values might be affected by an ex-post selection bias. Since optimal portfolio weights are strongly sensitive to the expected return estimates as shown in section 4, it is imperative to assure that our assertions hold for different return expectations. As a robustness check, we therefore additionally consider a range of expected return estimates generated by a simple bootstrap approach. Following Fleming, Kirby, and Ostdiek (2001), we begin by randomly drawing 1,000 bootstrapped samples. First, we draw randomly with replacement from the sample of actual excess returns. That way, we generate 1,000 artificial return series of 398 monthly stock, bond, and commodity returns. Then, we compute the mean excess return $\overline{r}_i^e$ over each return series and use them, along with our dynamic conditional correlation and volatility predictions, to compute optimal portfolio weights. Finally, we apply these weights to the actual returns and conduct our utility evaluation. This approach allows us to approximate the estimation risk an investor would face deciding for a reasonable return forecast.

-----------------Please insert Table 9 approximately here -----------------

Table 9 reports the average utility for our base case investor. Column 1 reports average monthly utility, averaged over the 1,000 bootstrapped return estimates. Again, the dynamic correlation timing strategy $(\hat{\sigma},\hat{\rho})$ outperforms naïve correlation forecasts $(\hat{\sigma},\overline{\rho})$, which in turn outperforms the passive strategy $(\overline{w})$. Column 2 reports the relative hit ratio for each of the three benchmark strategies. A hit is accounted for when the respective strategy yields the highest utility compared to both other strategies. Out of the 1,000 cases we observe, the dynamic correlation strategy attains the highest utility for 731 of the 1,000 return estimations, while the passive investment yields the highest utility for 217 return estimations. Interestingly, the naïve correlation strategy, which differs from the dynamic correlation strategy only in its expected correlation but not in variance predictions, yields even less hits than the passive strategy, 52. This demonstrates the importance of accounting for time-variation in both correlations and variances.

5.4 Comparison of predicting correlation versus predicting variance

In the previous subsections, we find it economically profitable to use a portfolio strategy that predicts both individual asset variances and pairwise asset correlations, when compared to a naïve correlation timing strategy that only predicts volatility. Similarly, it is well-known from Fleming, Kirby, and Ostdiek (2001) and Marquering and Verbeek (2004) that forecasting individual asset volatility also contains economic value. Therefore it is reasonable to assume, that a naïve estimate of both characteristics will yield the worst results, while predicting only volatility or only correlation will yield higher utility, and finally, predicting both characteristics will lead to most optimal portfolios. In this subsection we investigate these assumptions. This analysis also reveals how much predicting time-varying correlation contributes to finding optimal portfolio weights, when compared to the contribution of predicting individual asset volatility. To separate these effects, we compare the utility for five trading strategies that differ in which characteristics are predicted and which characteristics are based on naïve expectations. We report results in Table 10.

-----------------Please insert Table 10 approximately here -----------------

Strategies I through III remain unchanged from the previous analyses. Again, we use a passive, equal-weight strategy ($\overline{w}$) as the benchmark. This strategy yields realized utility of

0.0068 and per definition has an economic value of zero. Also the naïve correlation timing strategy $(\hat{\sigma},\bar{\rho})$ that only predicts volatility and the dynamic correlation strategy $(\hat{\sigma},\hat{\rho})$ that predicts both volatility and correlation are included in this analysis. Additionally, we introduce a "naïve volatility" strategy $(\bar{\sigma},\hat{\rho})$ that expects each asset's variance to equal previous month's value, i.e. $\hat{\sigma}_{t+1} = \sigma_t$. This seems to be a rather natural choice for naïve volatility timing, due to the well-known heteroscedasticity in asset returns.[13] Lastly, we introduce a truly naïve strategy $(\bar{\sigma},\bar{\rho})$ that is completely based on naïve predictions for both correlation and volatility.

This completely naive strategy $(\bar{\sigma},\bar{\rho})$ yields realized utility of 0.0064 which translates into a negative economic value of -0.0004. This shows that naïve predictions are not appropriate for a mean-variance portfolio optimization. As a consequence, the resulting portfolio might perform worse than a passive diversification strategy assigning constant, equal weights to all assets. Naïve correlation timing $(\hat{\sigma},\bar{\rho})$ yields a monthly economic value of 0.0009 on average. Comparing the difference between this strategy's utility with the completely naive $(\bar{\sigma},\bar{\rho})$ strategy's utility yields a measure for the economic value added by time-varying volatility predictions. Similarly we can consider the difference in realized utility between the naïve volatility strategy $(\bar{\sigma},\hat{\rho})$ and the naïve $(\bar{\sigma},\bar{\rho})$ strategy as the economic value of predicting correlation. It seems that predicting individual asset variances is more important (utility for strategy $(\hat{\sigma},\bar{\rho})$: 0.0077) for finding optimal portfolios than predicting the correlation between these assets (utility for strategy $(\bar{\sigma},\hat{\rho})$: 0.0072). However, our analysis is only based on one specific correlation forecasting methodology and one specific forecast for asset variances. Thus, it is not clear how other prediction models would perform. Most importantly, our analyses show

[13] Using historical means for the naive volatility estimate yields much worse results.

that predicting both volatility and correlations yields most optimal results, as the utility for the fully dynamic strategy $(\hat{\sigma}, \hat{\rho})$ clearly outperforms all other strategies. The importance of predicting both characteristics becomes even more evident, when we consider the Sharpe ratio. Out of all four active volatility timing strategies, only the fully dynamic strategy $(\hat{\sigma}, \hat{\rho})$ attains a higher Sharpe ratio than the passive, equal weights strategy. In conclusion, it seems apparent that it is indispensable to apply sophisticated predictions for both variance and correlation, rather than only predicting just one of the characteristics.

6. Conclusion

Despite decades of research in modeling time-varying correlation between risky assets, it is still common practice to assume static values for correlation in the context of portfolio optimization. This is particularly problematic in the context of asset allocation, as the correlation between different asset classes is subject to significant time variability, with macroeconomic conditions being a major determinant of cross-asset class correlation. The question arises whether predicting correlation in the context of asset allocation yields an economic gain over naïve estimates such as the historical mean correlation. We address this question by investigating the economic value of different trading strategies in a mean-variance optimization framework for a portfolio of stocks, bonds, commodity futures, and a risk free asset. We define economic value as the indifference fee as a percentage of the invested amount that would make a given investor indifferent between two trading strategies. By comparing two trading strategies that rely on the same expectations for returns and individual asset volatilities, but only differ in whether correlation is predicted or not, we can isolate the economic value of time-varying correlation predictions.

Based on a large sample (1974 – 2007) we find that accounting for time-varying correlation enables more optimal portfolios than applying naïve correlation expectations. In our sample, the economic gain of predicting correlation over a naïve strategy relying on historical mean correlations translates into an annual management fee worth 1.26%. This demonstrates that there is clearly an economic gain from predicting correlation for asset allocation. Most importantly, we find that accounting for both, time-varying conditional asset variances and their time-varying pairwise correlations, yields the highest realized utility. Our results are robust to return estimation risk, varying degrees of risk aversion, short sale constraints, and different levels of transaction costs. Our findings are important, as they justify the increasing efforts academics and practitioners alike put into finding better models for predicting the correlation between asset classes. Also, our results suggest that investors allocating portfolio weights to different asset classes should employ sophisticated predictions for the cross-asset class correlation rather than rely on naïve, historical values.

References

Anderson, T. W., 1984, An Introduction to Multivariate Statistical Analysis, 2nd edition, John Wiley and Sons, Inc.

Andersen, T.G., Bollerslev, T., Diebold, F.X. and Ebens, H., 2001, The Distribution of Realized Stock Return Volatility. *Journal of Financial Economics* 61, 43-76.

Andersen, T.G, Bollerslev, T., Diebold, F.X. and Vega, C., 2007. Real-Time Price Discovery in Stock, Bond and Foreign Exchange Markets. *Journal of International Economics*, forthcoming.

Ang, A. and Bekaert, G., 2007. Stock Return Predictability: Is it There? *Review of Financial Studies* 20, 651-707.

Angas, L.L.B., 1936. Investment for Appreciation. Forecasting Movements in Security Prices. Technique of Trading in Shares for Profit. London, England: McMillan.

Best, M. J. and Grauer, R., 1991. On the sensitivity of mean-variance efficient portfolios to changes in asset means: Some analytical and computational results. *Review of Financial Studies* 4, 315-342.

Bollerslev, T.P., 1990. Modelling the Coherence in Short Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model. *Review of Economics and Statistics* 72, 498-505.

Bollerslev, T.P., Chou, R. and Kroner, K., 1992. ARCH modeling in finance. *Journal of Econometrics* 52, 5-59.

Bollerslev, T.P., Engle, R.F. and Wooldridge, J.M., 1988. A capital asset pricing model with time-varying covariances. *Journal of Political Economy* 96, 116-131.

Bossaerts, P. and Hillion, P., 1999. Implementing statistical criteria to select return forecasting models: What do we learn? *Review of Financial Studies* 12, 405–428.

Chan, L. K. C., Karceski, J. and Lakonishok, J., 1999. On portfolio optimization: forecasting covariances and choosing the risk model. *Review of Financial Studies* 12, 263–278.

Connolly, R., Stivers, C. and Sun, L., 2005. Stock Market Uncertainty and the Stock-Bond Return Relation. *Journal of Financial and Quantitative Analysis* 40, 161-194.

Connolly, R., Stivers, C. and Sun, L., 2007. Commonality in the Time- Variation of Stock-Bond and Stock-Stock Return Comovements, *Journal of Financial Markets* 10, 192-218.

d'Addona, S. and Kind, A.H., 2006. International Stock–Bond Correlations in a Simple Affine

Asset Pricing Model. *Journal of Banking and Finance* 30, 2747-2765.

Ding, Z. and Engle, R.F., 2001. Large Scale Conditional Covariance Matrix Modelling, Estimation and Testing. *Academia Economic Papers* 29, 157-184.

Dopfel, F.E., 2003. Asset allocation in a lower stock-bond correlation environment. *Journal of Portfolio Management* 30, 25-38.

Elton, E. J. and Gruber, M. J., 1973. Estimating the dependence structure of share prices – implications for portfolio selection. *Journal of Finance* 28, 1203–1232.

Elton, E. J. and Gruber, M. J. and Spitzer, J., 2006. Improved Estimates of Correlation Coefficients And Their Impact on the Optimum Portfolios. *European Financial Management* 12, 303-318.

Elton, E. J., Gruber, M. J. and Ulrich, T., 1978. Are betas best?. *Journal of Finance* 33, 1375–1384.

Engle, R. and Sheppard, K., 2003. Theoretical and empirical properties of dynamic correlation multivariate garch. Working Paper, New York University 2003.

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

Fleming, J., Kirby, C. and Ostdiek, B., 2001. The Economic Value of Volatility Timing. *Journal of Finance* 56, 329-352.

Gorton, G. and Rouwenhorst, K.G., 2006. Facts and Fantasies about Commodity Futures. *Financial Analysts Journal*, 62, 47-68.

Gulko, L., 2002. Decoupling: if the US treasury repays its debt, what then?. *Journal of Portfolio Management* 28, 59–66.

Hautsch, N., Hess, D., 2007. Bayesian Learning in Financial Markets: Testing for the Relevance of Information Precision in Price Discovery. *Journal of Financial and Quantitative Analysis* 42, 189-208.

Higham, N. J., 2002. Computing the Nearest Correlation Matrix - A Problem from Finance. *IMA Journal of Numerical Analysis* 22, 329-343.

Ibbotson, R.C. and Sinquefield, R.A. 1976. Stocks, Bonds, Bills and Inflation: Simulations of the Future (1976-2000). *Journal of Business* 49, 313-338.

Jagannathan, R. and Ma, T., 2003. Risk reduction in large portfolios: why imposing wrong constraints helps. *Journal of Finance* 58, 1651–1684.

Kroner, K.F. and Claessens, S., 1991. Optimal Dynamic Hedging Portfolios and the Currency Composition of External Debt. *Journal of International Money and Finance* 10, 131-148.

Ledoit, O. and Wolf, M., 2004. Honey, I shrunk the sample covariance matrix'. *Journal of Portfolio Management* 30, 110-119.

Marquering, W. and Verbeek, M., 2004. The Economic Value of Predicting Stock Index Returns and Volatility. *Journal of Financial and Quantitative Analysis* 39, 407-429.

Pesaran, M. H. and Timmermann, A., 1995. Predictability of stock returns: Robustness and economic significance. *Journal of Finance* 50, 1201–1228.

J.P.Morgan, 1996, RiskMetrics™ Technical Document, 4-th Edition, December 1996.

Zhou, Y., 2006. Business Cycle Asymmetry in the Stock-Bond Correlation: A Study of the US and the UK from 1855 to 2001. *Working paper*, Cornell University.

Tables

Table 1: Benchmark strategies

Strategies	Weights depend on
Passive	
I: 25% equal weights	$\overline{w}$
Volatility timing strategies	
II: Naive correlation timing	$w(\hat{\sigma}, \overline{\rho})$
III: Dynamic correlation timing	$w(\hat{\sigma}, \hat{\rho})$

This table presents the three benchmark portfolio strategies that we compare throughout this paper. Strategy I is a buy and hold strategy allocating constant, equal weights to all four asset classes. Strategy II performs volatility forecasts but assumes correlation to equal historical means. Strategy III uses dynamic volatility and dynamic correlation forecasts.

Table 2: Summary statistics

	(1) Feb 1974 - Mar 2007 n=398	(2) Feb 1974 - Aug 1990 n=199	(3) Sep 1990 - Mar 2007 n=199
Panel A: Monthly S&P 500			
Mean return	1.04%	1.08%	0.99%
Standard deviation	4.39%	4.83%	3.91%
Mean excess return	0.55%	0.43%	0.67%
Excess return standard deviation	4.40%	4.86%	3.90%
Panel B: Monthly 10-year T-bonds			
Mean return	0.70%	0.76%	0.64%
Standard deviation	2.45%	2.80%	2.05%
Mean excess return	0.22%	0.12%	0.32%
Excess return standard deviation	2.45%	2.81%	2.04%
Panel C: Monthly CRB futures index			
Mean return	0.70%	0.75%	0.64%
Standard deviation	3.20%	3.74%	2.57%
Mean excess return	0.21%	0.11%	0.31%
Excess return standard deviation	3.22%	3.76%	2.58%
Panel D: Monthly asset correlations			
Mean correlation stocks and bonds	0.1813	0.2991	-0.0230
Mean correlation stocks and commodities	0.0343	-0.0104	0.1132
Mean correlation bonds and commodities	-0.1282	-0.1751	-0.0386

This table reports summary statistics for the full sample (column 1), the first (column 2), and the second half of the sample (column 3). Panel A, B, C report mean monthly (excess) returns and standard deviation of monthly (excess) returns for stocks, bonds, and commodities. Panel D reports mean correlation of monthly returns for each two-asset combination. The stock index return is measured by the return on the S&P 500 index adjusted for dividends. Bond returns are the monthly return of the 10 year constant maturity US Treasury note. Commodity returns are returns of the CRB commodity futures index, assuming a fully collateralized position in three month US Treasury bills. Excess returns are calculated by subtracting the three month T-bill rate.

Table 3: In-sample correlation regressions

Explanatory Variables	(1) ρ^F_{SB}		(2) ρ^F_{SC}		(3) ρ^F_{BC}	
	Coeff.	p-value	Coeff.	p-value	Coeff.	p-value
Intercept	0.909	0.004	-0.226	0.135	-0.040	0.794
Price Earnings Ratio(-1)	-0.024	0.007	0.000	0.916	-0.008	0.054
Dividend Yield(-1)	0.009	0.839	-0.035	0.199	-0.088	0.001
Inflation(-2)	-2.329	0.000	-0.329	0.444	0.681	0.154
Industrial Production(-2)	-0.268	0.649	-0.167	0.772	-0.319	0.560
3-month treasury yield(-1)	-3.146	0.001	-0.031	0.975	0.487	0.614
3-month treasury yield (-2)	1.316	0.159	-0.633	0.510	-0.710	0.436
12-month treasury yield(-1)	4.334	0.000	0.396	0.665	-0.455	0.642
12-month treasury yield (-2)	-2.141	0.044	0.504	0.576	0.886	0.337
Money stock(-2)	-3.940	0.000	0.614	0.325	1.143	0.048
Default spread(-1)	-2.133	0.129	2.502	0.002	2.316	0.013
Sample size	398		398		398	
R^2	0.372001		0.156776		0.146637	
Adjusted R^2	0.355689		0.134874		0.124472	
Prob(F-statistic)	0.0000		0.0000		0.0000	

This table reports in-sample regression results for the Fisher transformation of pairwise asset correlation coefficients, $\rho^F_{ij,t}$, on a set of lagged macroeconomic and financial variables, x_t, assuming $\rho^F_{ij,t+1} = x_t'\beta_{ij} + \varepsilon_{ij,t+1}$. Column 1, 2, and 3 present results for the stock-bond, stock-commodities, and bond-commodity correlations, respectively. We employ macroeconomic variables with a lag of two month to avoid any look-ahead bias, and financial variables with a lag of one month. The independent variables include: price earnings ratio on the S&P 500, dividend yield on the S&P 500, change in industrial production, inflation rate, 3-month and 12-month treasury yields (both in percent), growth in money stock, and the default spread (yield for BAA corporate bonds minus AAA rates). We report Newey-West corrected p-values. Adjusted R^2 denotes the R^2 adjusted for the degrees of freedom.

Table 4: In-sample variance regressions

Explanatory Variables	(1) Log(σ_S^2)		(2) Log(σ_B^2)		(3) Log(σ_C^2)	
	Coeff.	p-value	Coeff.	p-value	Coeff.	p-value
Intercept	-3.985	0.000	-4.691	0.00	-4.208	0.000
Log(Var(-1))	0.551	0.000	0.529	0.02	0.425	0.000
Price Earnings Ratio(-1)	0.031	0.000	0.021	0.55	-0.003	0.655
Dividend Yield(-1)	-0.042	0.455	-0.036	0.52	-0.118	0.012
Inflation(-2)	1.393	0.081	-0.543	0.49	4.232	0.000
Industrial Production(-2)	-1.443	0.087	-0.726	0.55	-0.673	0.339
3-month treasury yield(-1)	-1.860	0.192	-1.169	0.11	-4.799	0.001
3-month treasury yield(-2)	1.219	0.385	2.720	0.08	3.350	0.013
12-month treasury yield (-1)	3.065	0.071	3.405	0.02	5.375	0.002
12-month treasury yield (-2)	-1.682	0.317	-3.993	0.00	-4.067	0.013
Money stock(-2)	2.080	0.065	-3.766	0.00	3.542	0.001
Default spread(-1)	0.565	0.761	6.267	0.00	1.089	0.472
Sample size	398		398		398	
R^2	0.493723		0.541831		0.439547	
Adjusted R^2	0.47922		0.528706		0.423492	
Prob(F-statistic)	0.0000		0.0000		0.0000	

This table reports in-sample regression results for the logarithm of asset variances, $\sigma_{i,t}^2$, on a set of lagged macroeconomic and financial variables, y_t, assuming $\log \sigma_{i,t+1}^2 = y_t'\delta_i + \xi_{i,t+1}$. Column 1, 2, and 3 present results for the stock, bond, and commodity variance, respectively. We employ macroeconomic variables with a lag of two month to avoid any look-ahead bias, and financial variables with a lag of one month. The independent variables include: price earnings ratio on the S&P 500, dividend yield on the S&P 500, change in industrial production, inflation rate, 3-month and 12-month treasury yields (both in percent), growth in money stock, and the default spread (yield for BAA corporate bonds minus AAA rates). We report Newey-West corrected p-values. Adjusted R^2 denotes the R^2 adjusted for the degrees of freedom.

Table 5: Out-of-sample test for correlation predictions

Root mean squared error	(1) $\hat{\rho}_{SB}$	(2) $\hat{\rho}_{SC}$	(3) $\hat{\rho}_{BC}$
Naïve correlation model $(\overline{\rho})$	0.2947	0.2137	0.2214
Dynamic correlation model $(\hat{\rho})$	0.2289	0.1947	0.2021

This table presents out of sample tests for correlation predictions. Each month we forecast the pairwise correlation between stocks, bonds and commodities for next month. We compare a dynamic correlation model against forecasts produced by a naïve correlation model, which assumes correlations equal to historical mean values. Column 1 reports root mean squared errors for the stock-bond correlation. Column 2 and 3 report root mean squared errors for the stock-commodity and bond-commodity correlation.

Table 6: Marginal impact of prediction errors on optimal stock weights

Marginal weight impact of expected stock characteristics	(1) $\hat{\rho}_{SB}$ = 0.181 $\hat{\rho}_{SC}$ = 0.034	(2) $\hat{\rho}_{SB}$ = -0.119 $\hat{\rho}_{SC}$ = 0.034	(3) $\hat{\rho}_{SB}$ = 0.181 $\hat{\rho}_{SC}$ = -0.266
Returns ($\frac{\partial w_S^*}{\partial \hat{r}_S} \cdot \sigma_{r_S} / 100$)	0.0472	0.0461	0.0501
Volatility ($\frac{\partial w_S^*}{\partial \hat{\sigma}_S} \cdot \sigma_{\sigma_S} / 100$)	-0.0050	-0.0056	-0.0059
Correlation with bonds ($\frac{\partial w_S^*}{\partial \hat{\rho}_{SB}} \cdot \sigma_{\rho_{SB}} / 100$)	-0.0010	-0.0022	-0.0010
Correlation with commodities ($\frac{\partial w_S^*}{\partial \hat{\rho}_{SC}} \cdot \sigma_{\rho_{SC}} / 100$)	-0.0008	-0.0009	-0.0019

This table reports the marginal impact of an adjustment in predicted stock characteristics on the stock weight in an ex-ante optimal portfolio. We differentiate the optimal stock-weight function with respect to expected stock return, standard deviation, as well as its correlation with bonds and commodity futures. All optimization inputs are set to the sample mean in column 1. In column 2 we lower ρ_{SB} by 0.3, and in column 3 we lower ρ_{SC} by 0.3, with all other values held equal. The risk aversion coefficient γ is set to 5. The marginal impact is then multiplied with 1/100 of a standard deviation in the respective stock characteristics.

Table 7: Economic value of benchmark trading strategies

		(1)	(2)	(3)	(4)	(5)
Panel A: Jan 1980 - Mar 2007		***Utility***	***EV***	***Excess Return***	**STD**	***Sharpe Ratio***
Passive Strategies						
	100% Stocks	0.0056	-0.0012	0.65%	4.27%	0.1516
	100% Bonds	0.0063	-0.0005	0.30%	2.57%	0.1179
	100% Commodities	0.0043	-0.0026	0.15%	2.82%	0.0544
	100% Risk free	0.0047	-0.0021	0.0%	0.25%	0.0000
I:	25% Equal weights $(\overline{w})$	0.0068	-	0.28%	1.52%	0.1814
Volatility timing strategies						
II:	Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0077	0.0008	0.51%	2.99%	0.1718
III:	Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0087	0.0019	0.67%	3.30%	0.2021
Panel B: Jan 1980 - Aug 1993		**Utility**	**EV**	***Excess Return***	**STD**	**Sharpe Ratio**
Passive Strategies						
I:	25% equal weights $(\overline{w})$	0.0080	-	0.24%	1.67%	0.1458
Volatility timing strategies						
II:	Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0075	-0.0005	0.36%	3.29%	0.1094
III:	Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0086	0.0007	0.51%	3.49%	0.1462
Panel C: Sep 1993 - Mar 2007		**Utility**	**EV**	***Excess Return***	**STD**	**Sharpe Ratio**
Passive Strategies						
I:	25% equal weights $(\overline{w})$	0.0057	-	0.31%	1.36%	0.2278
Volatility timing strategies						
II	Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0079	0.0021	0.67%	2.67%	0.2496
III:	Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0088	0.0030	0.82%	3.11%	0.2645

This table presents the performance of the three benchmark trading strategies (I, II, III). Strategy I is a buy and hold strategy allocating constant, equal weights to all four asset classes. Strategy II performs volatility forecasts but assumes correlation to equal historical means. Strategy III uses dynamic volatility and dynamic correlation forecasts. Panel A reports results for the full sample, and additionally presents results for 4 passive strategies which fully invest in one single asset class. Panel B and C report results for the first and second half of the sample. Column 1 reports average monthly realized utility. Column 2 shows the economic value of each strategy, measured as the utility gain over the passive strategy I. Column 3 reports average monthly excess returns, while column 4 reports the standard deviation. Column 5 presents the Sharpe ratio, defined as excess returns divided by the sample standard deviation.

Table 8: Sensitivity analysis

Panel A: Transaction costs	Utility			
	Zero	Low	Medium	High
Passive Strategies				
I: 25% Equal weights $(\overline{w})$	0.0068	0.0068	0.0068	0.0068
Volatility timing strategies				
II: Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0088	0.0084	0.0077	0.0063
III: Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0100	0.0095	0.0087	0.0070

Panel B: Risk aversion	Utility			
	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$	$\gamma = 20$
Passive Strategies				
I: 25% Equal weights $(\overline{w})$	0.0073	0.0068	0.0062	0.0049
Volatility timing strategies				
II: Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0195	0.0077	0.0062	0.0055
III: Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0242	0.0087	0.0067	0.0057

Panel C: Trading constraints	Utility		
	Unrestricted	No short sales	No short sales & no borrowing
Passive Strategies			
I: 25% Equal weights $(\overline{w})$	0.0068	0.0068	0.0068
Volatility timing strategies			
II: Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0077	0.0079	0.0073
III: Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0087	0.0086	0.0075

This table presents average realized utility for the three benchmark strategies under varying alternative assumptions. Results are based on the full sample January 1980 – March 2007. Strategy I is a buy and hold strategy allocating constant, equal weights to all four asset classes. Strategy II performs volatility forecasts but assumes correlation to equal historical means. Strategy III uses dynamic volatility and dynamic correlation forecasts. Panel A varies transaction costs, assuming transaction cost levels: zero (0% for stocks / 0% for bonds / 0% for commodities), low (0.1% / 0.1% / 0.1%), medium (0.5% / 0.1% / 0.1%), and high (1.0% / 0.5% / 0.1%). Panel B varies the investor's risk aversion coefficient, γ, assuming levels of: 1, 5, 10, and 20. Panel C compares the trading strategies' realized utility under different trading constraints, when either short sales or short sales and borrowing is not allowed.

Table 9: Return estimation risk

		Utility	Hit ratio
Passive Strategies			
I:	25% Equal weights $(\overline{w})$	0.0068	21.7%
Volatility timing strategies			
II:	Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0072	5.2%
III:	Dynamic correlation $(\hat{\sigma}, \hat{\rho})$	0.0076	73.1%

This table presents average realized utility for the three benchmark strategies under varying return expectations. We use a bootstrap procedure to simulate different ex ante information sets, by drawing randomly with replacement from the actual set of excess returns. This way, we generate 1,000 artificial return series of 398 monthly stock, bond, and commodity returns. Then, we compute the mean excess returns over each return series and use these values as the constant expected returns. This approach allows us to calculate realized utilities for each benchmark strategy for 1,000 different constant return expectations. Column 1 reports the mean utility averaged over 1,000 return expectations for each strategy. Strategy I is a buy and hold strategy allocating constant, equal weights to all four asset classes. Strategy II performs volatility forecasts but assumes correlation to equal historical means. Strategy III uses dynamic volatility and dynamic correlation forecasts. Column 2 reports the hit ratio for each strategy. Hit ratio is defined as the percentage of return expectations for which the respective strategy yield the highest realized utility.

Table 10: Economic value of predicting correlation vs. predicting volatility

		(1) **Utility**	(2) **EV**	(3) **Excess Return**	(4) **STD**	(5) **Sharpe Ratio**
Passive Strategies						
I:	25% Equal weights $(\overline{w})$	0.0068	-	0.28%	1.52%	0.1814
Volatility timing strategies						
	Naïve strategy $(\overline{\sigma}, \overline{\rho})$	0.0064	-0.0004	0.52%	3.72%	0.1403
II:	Naive correlation $(\hat{\sigma}, \overline{\rho})$	0.0077	0.0009	0.51%	2.99%	0.1718
	Naive volatility $(\overline{\sigma}, \hat{\rho})$	0.0072	0.0004	0.68%	4.12%	0.1655
III:	Dynamic strategy $(\hat{\sigma}, \hat{\rho})$	0.0087	0.0019	0.67%	3.30%	0.2021

This table presents the performance of five trading strategies, which differ in what asset characteristics are predicted and what characteristics are set to naïve values. Strategy I $(\overline{w})$ is a buy and hold strategy allocating constant, equal weights to all four asset classes. Strategy II $(\hat{\sigma}, \overline{\rho})$ performs volatility forecasts but assumes correlation to equal historical means. Strategy III $(\hat{\sigma}, \hat{\rho})$ uses dynamic volatility and dynamic correlation forecasts. The naïve volatility strategy $(\overline{\sigma}, \hat{\rho})$ predicts correlation, but assumes asset volatility to equal previous month's value. The naïve strategy $(\overline{\sigma}, \overline{\rho})$ assumes naïve expectations for both correlation and volatility. Column 1 reports average monthly realized utility. Column 2 shows the economic value of each strategy, measured as the utility gain over the passive strategy I. Column 3 reports the average excess return per month while column 4 reports the standard deviation. Column 5 presents the Sharpe ratio, defined as excess returns divided by the sample standard deviation.

Figures

Figure 1: 12-month moving average asset correlations

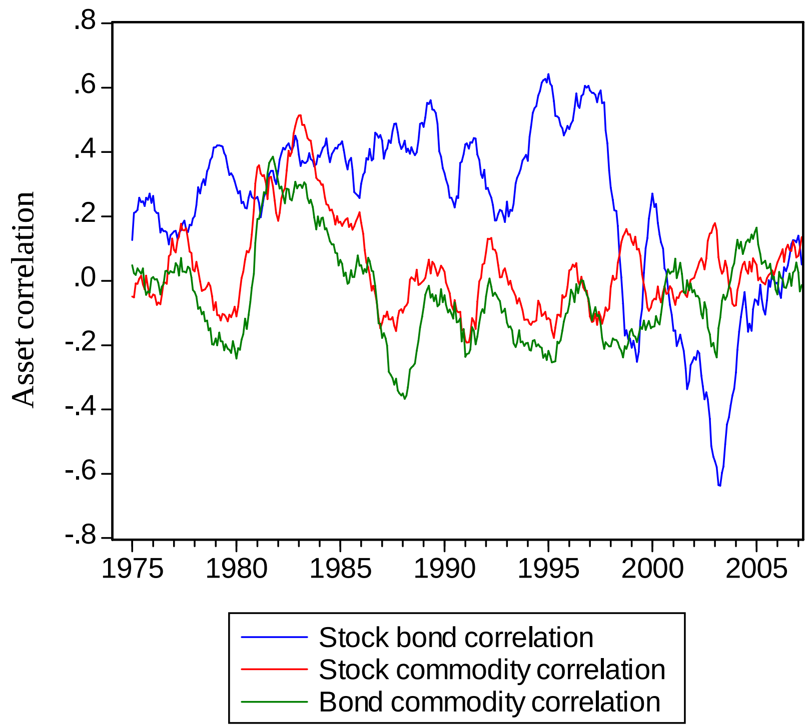

This figure reports the 12-month moving average of asset correlations throughout the sample period January 1980 - March 2007 for the stock-bond, stock-commodity, and bond-commodity correlation.

Figure 2: Portfolio weights for the dynamic correlation strategy

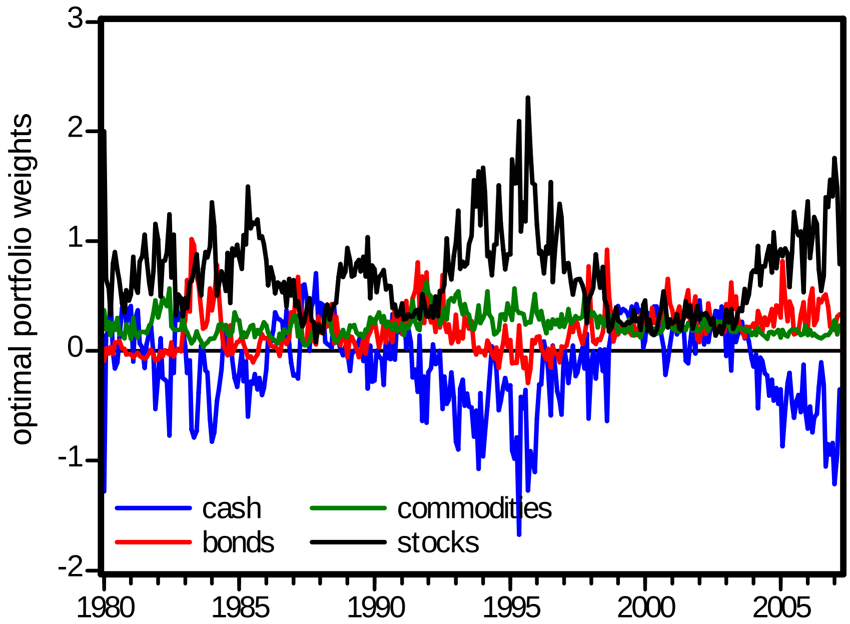

This figure reports monthly optimal portfolio weights for each asset class, using the dynamic correlation timing strategy $(\hat{\sigma}, \hat{\rho})$, throughout the sample period January 1980 until March 2007.