Exploiting the Informational Content of the Linkages between Spot and Derivatives Markets

Devraj Basu
Roel Oomen
Alexander Stremme

Warwick Business School, University of Warwick

March 2006

address for correspondence:

Roel Oomen
Warwick Business School
University of Warwick
Coventry, CV4 7AL
United Kingdom

e-mail: `roel.oomen@wbs.ac.uk`
phone: +44 (0) 2476 - 528 200
fax: +44 (0) 2476 - 523 779

Exploiting the Informational Content of the Linkages between Spot and Derivatives Markets

March 2006

Abstract

Several recent studies have studied the use of predictive variables in enhancing asset allocation, focusing mainly on various business cycle indicators. While the statistical evidence is mixed, they have found that even small levels of predictability from a statistical viewpoint can have a substantial effect on both the investor's portfolio decision as well as portfolio performance. This study is one of the first to focus on the use of market variables. We use the VIX, open interest and hedging pressure on our base assets, all of which allow us to exploit the linkages between spot and derivatives markets. These variables have been studied separately by both academics and practitioners, but their linkages have not been examined before. Using the S&P 500, gold and Treasury bond index as our base assets, we study the performance of these variables by examining both the out of sample performance of unconditionally efficient portfolios based on our predictive variables as well as their in-sample performance using a statistical test. We find that different sets of variables have predictive power in bear and bull markets. Our trading strategies can successfully time the market and avoid losses during the collapse of the 'dot.com' bubble in the second half of 2000, as well as during the earlier bull run. The in-sample results confirm our out of sample experiments with p-values of less than 1% in all cases. The performance of the strategies deteriorates considerably if we remove any of the predictive variables, indicating that it is the correlations between these variables that drives the strategies.

JEL Classification: C12, C32, G11, G14

1 Introduction

The economic value of return predictability has been the focus of much recent research. While the statistical evidence is mixed, several studies, for example Kandel and Stambaugh (1996), Campbell and Vicera (2001), Avramov and Chordia (2006), or Abhyankar, Basu, and Stremme (2005), have found that even small levels of predictability from a statistical viewpoint can have a substantial effect on both the investor's portfolio decision as well as portfolio performance. Almost all of these studies have focused on predictive variables that are closely linked to the business cycle such as dividend yield and various term structure and macroeconomic variables. This paper focuses on the use of market variables that exploit the linkages between spot, futures and derivatives markets.

Spot and futures market linkages are exploited by using open interest and hedging pressure as the predictive variables. Hedging pressure is defined as the number of long futures contracts divided by the total number of futures contracts and may be regarded as a proxy for non-marketable risks. In an incomplete market the non-marketable risk fraction affects expected returns on assets, as noted by De Roon, Nijman, and Veld (2000). This result is similar to the CAPM with non-tradable assets as discussed in Mayers (1976). The effect of hedging pressure in futures markets has been extensively investigated. For example, Bessembinder (1992) finds that hedging pressure has considerable effect on currency futures risk premia. He suggests that the effects of hedging pressure in this market is larger than that predicted by theoretical models, as firms hedge in the currency markets for a variety of reasons beyond that predicted by theoretical models such as the CAPM with non-traded assets. In recent years hedge funds have become active players in the futures markets and their activities could influence expected returns in futures and spot markets. De Roon, Nijman, and Veld (2000) have documented that hedging pressure has predictive ability for spot markets and our study is the first to explore its effect on asset allocation.

The linkages between the derivatives and spot markets are exploited using the VIX index, a proxy for implied volatility. There is considerable recent evidence[1] that volatility timing has large economic benefits. However all of these studies have used various statistical or econometric proxies for volatility such as realized volatility. Various studies find that implied volatility is a good predictor of future volatility and the VIX index is used by several investment firms in their asset allocation decisions. A recent study by Connolly, Stivers, and Sun (2005) finds that the use of implied volatility as a predictive variable could significantly improve stock-bond diversification.

We study the effect of these variables on market timing, and our risky base assets are the S&P 500, a Treasury bond index, and Gold. In addition we also have a conditionally risk free asset, the Fed Funds rate which may also be regarded as a very short term T-bill. Our predictive variables are the level of the VIX, and for each base asset the corresponding futures market open interest and hedging pressure. While hedging pressure is a directional indicator, open interest measures the size of the outstanding position. Our data is at a weekly frequency.

Our dynamic strategies are unconditionally efficient minimum variance or maximum return strategies. This focus on unconditionally efficient strategies, as opposed to conditionally efficient strategies which have been employed in most other studies, including Fleming, Kirby, and Ostdiek (2001), and Fleming, Kirby, and Ostdiek (2003), is an important aspect of our analysis. Unconditionally efficient strategies are theoretically optimal (Hansen and Richard 1987), utilize the predictive information more effectively (Abhyankar, Basu, and Stremme 2005) and also exhibit a 'conservative response' to extreme values of the predictive variables (Ferson and Siegel 2001), unlike unconditionally efficient strategies that tend to take extreme long or short positions.

We first investigate whether out-of-sample, these linkages predict extreme market movements

[1] For example Fleming, Kirby, and Ostdiek (2001), and Fleming, Kirby, and Ostdiek (2003).

and thus help avoid losses in market crashes or periods of high volatility[2]. We find that pure market timing strategies based on these predictive variables would not have incurred any significant losses during the collapse of the 'dot.com' bubble in late 2000. These strategies seem to allow asset managers to 'de-couple' their portfolios from the business cycle during a bear market (with betas between 0.2 and 0.5) and thus successfully time the market, with a management fee of over 1,100 basis points. The superior performance of these strategies appears to be a bear market phenomenon as the corresponding out of sample strategy during the earlier bull market is a high beta strategy that does not outperform the index. The performance of the strategies deteriorates considerably if we remove any of the predictive variables, indicating that it is the correlations between these variables that drives the strategies. Adding gold as an additional hedge together with its own hedging pressure and open interest further improves the performance of timing strategy during the bear market, with the management fee rising to almost 1300 basis points.

The long run performance of these bear market strategies confirms the out of sample findings. The minimum variance pure market timing strategy using the S&P together with the VIX and its own hedging pressure and open interest as predictive variables underperforms the index over the 1992-2000 period but significantly outperforms it from the end of 2000 to the end of 2005. Investing in the market timing strategy would have yielded a total return of around 400%, while the index had a total return of around 100%. The strategy is low beta (0.34) and high alpha (8.4%). The maximum return strategy matches the index over the first period and achieves a total return of over 1,000% with an alpha of 15.2%. It is however significantly more volatile than the minimum variance strategy and is reflected in the lower management fee of 10.2 percentage points of annual return as management fee for this strategy relative to 18.8 percentage points for the minimum variance strategy. Adding gold and its predictive variables improves the performance of the strategies somewhat, but

[2]We focus on two sample periods for our out sample analysis, the first being the bull run starting at the end of 1992 until April 2000 and then from April 2000 until the end of the sample period (September 2005).

overall the gains are of the same order of magnitude as before.

We next see whether it is possible to benefit from these strategies in a bull run. An out of sample strategy using S&P together with Treasury bond index with the VIX, and hedging pressure and open interest on both assets in the first period achieves a annual return of 60%, an alpha of 42% and a management fee of 16.6 percentage points. This strategy does not perform well over the second period indicating that it is a bull market strategy.

To assess the statistical significance of our results, we use a test developed in Abhyankar, Basu, and Stremme (2005). This test is based on the difference in the slope of the efficient frontiers with and without the optimal use of predictive information. As this test has a known statistical distribution (both in finite sample as well as asymptotically), we can assess whether the expansion of the frontier due to predictability is statistically significant or just a product of sampling error. The in-sample results confirm our out of sample experiments. In our first period, the VIX together with hedging pressure and open interest on the S&P and gold does not significantly expand the investor's opportunity set with p-values between 8 and 11%. In contrast in the second period the Sharpe ratios of the dynamic strategy rise sharply with the difference in Sharpe ratios being significant at the 1% level both for the S&P by itself and the S&P together with gold. Conversely in the first period the VIX together with hedging pressure and open interest on the S&P and the bond index leads to a Sharpe ratio of 2.43, the difference having a p-value of 0.

The remainder of this paper is organized as follows. Section 2 describes our methodology. The results of our empirical analysis are presented in Section 3. Section 4 concludes.

2 Measuring and Exploiting Predictability

In this section, we describe our methodology for measuring the economic gains and statistical significance of return predictability. Details are provided in Appendix A.

2.1 Dynamically Efficient Trading Strategies

Most of the existing literature on predictability and market-timing focuses on *'myopically optimal'* (conditionally efficient) strategies. In contrast, we focus here on *'dynamically optimal'*, i.e. unconditionally efficient strategies, as studied in Ferson and Siegel (2001), and Abhyankar, Basu, and Stremme (2005). While the portfolio weights of the former are determined *ex-post* on the basis of the conditional return moments, the weights of the latter are determined *ex-ante* as functions of the predictive instruments. In this sense, dynamically optimal strategies are truly actively managed, while myopically optimal strategies can be though of as sequences of one-step-ahead efficient *static* portfolios. Because dynamically optimal strategies are designed to be efficient with respect to their long-run unconditional moments, they display a more 'conservative' response to changes in the predictive instruments[3]. This is an important consideration in particular with respect to transaction costs.

Studies that have examined market-timing using predictive variables such as the short rate (Breen, Glosten, and Jagannathan 1989), or time-variation in the conditional Sharpe ratio Whitelaw (2005), have employed naive portfolio strategies related to conditional efficiency. In contrast, unconditionally efficient strategies optimally utilize *both* the predictive information *and* the time-variation in the conditional Sharpe ratio (Cochrane 1999), and can significantly outperform naive market-timing strategies.

We provide precise specifications of the weights of dynamically efficient strategies in Appendix A.1. In our empirical applications, we consider both efficient minimum-variance strategies (designed to track a given target average return), as well as efficient maximum-return strategies (designed to track a given target volatility). The former are particularly useful in risk management as they provide portfolio insurance against crashes and periods of excess volatility. The latter can be though of as 'active alpha' strategies, designed to achieve maximum performance at a tolerable level of risk.

[3]See also Ferson and Siegel (2001) or Abhyankar, Basu, and Stremme (2005).

2.2 Measuring the Value of Return Predictability

We employ a variety of performance criteria to measure the gains due to the optimal use of return predictability. To assess the long-run performance of dynamically managed strategies based on business cycle and sentiment indicators, we use a variety of standard *ex-post* portfolio performance measures. These include Sharpe ratios, Jensen's alpha, and information ratios.

Measures of Statistical Significance

To capture the incremental gains due to the optimal use of predictive information, we compare the performance of optimally managed portfolios with that of traditional *'fixed-weight'* strategies, i.e. those for which the weights do not depend on the instruments). We wish measure the extent to which the optimal use of predictive information expands the efficient frontier, and hence the opportunity set available to the investor. Because the location of the global minimum-variance (GMV) portfolio is virtually unaffected by the introduction of predictive variables (see also Figure 1), we use the asymptotic slope of the frontier (i.e. the Sharpe ratio relative to the zero-beta rate associated with the mean of the GMV) as such a measure. Because we can show that the difference in (squared) slopes of the frontiers with and without the optimal use of predictability has a known (χ^2) distribution, we are able to assess the statistical significance of any gains due to predictability. A precise definition of our test statistic is given in Appendix A.2.

Measures of Economic Value

In addition to our statistical tests, we also employ a utility-based framework to assess the economic value of return predictability. Following Fleming, Kirby, and Ostdiek (2001), we consider a risk averse investor whose preferences over future wealth are given by a quadratic von Neumann-Morgenstern utility function. Consider now an investor who faces the decision whether or not to acquire the skill and/or information necessary to implement the active portfolio strategy that optimally exploits predictability. The question is, how much of their expected return would the investor be willing to give up (e.g. pay as a management fee) in return for having access to the superior strategy? Put differently, by how much does the return on the inferior strategy have to be increased to make the investor indifferent between the optimal and the inferior strategy. A precise definition of this premium is given in Appendix A.3.

3 Empirical Analysis

3.1 Data Description

As base assets, we use weekly returns on the S&P 500, a Treasury bond index, and gold. The data are obtained from the CBOE (`www.cboe.com`). Returns are calculated based on closing prices and we record them on the date they are made public. For the return on the risk-free asset we use the Federal Funds rate.

We use two types of predictive instruments. The first one is the VIX, a derivative market indicator which aggregates implied volatilities of S&P 500 (SPX) index options. Second, for each base asset, we include two futures market indicators, namely open interest and 'hedging pressure'. Open interest is defined as the total of all futures contracts entered into and not yet closed out by an offsetting transaction. The aggregate of all long open interest is equal to the aggregate of all short open interest. Hedging pressure is defined as total commercial[4] long positions divided by the sum of long and short commercial positions. All data are

[4]When an individual reportable trader is identified to the Commission, the trader is classified either as 'commercial' or 'non-commercial'. All of a trader's reported futures positions in a commodity are classified as commercial if the trader uses futures contracts in that particular commodity for hedging as defined in the Commission's regulations (1.3(z)). A trading entity generally gets classified as a "commercial" by filing a statement with the Commission (on CFTC Form 40) that it is commercially "...engaged in business activities hedged by the use of the futures or option markets".

obtained from the CFTC web site (`www.cftc.gov`).

3.2 Empirical Analysis

We estimate a linear predictive model of the form,

$$(R_t - r_f 1) = (\bar{\mu} - r_f 1) + B \cdot Z_{t-1} + \varepsilon_t, \tag{1}$$

where Z_{t-1} is the vector of (lagged) predictive instruments, and R_t is the vector of risky asset returns. We assume that the residuals ε_t are serially independent and independent of Z_{t-1}. This implies that the conditional variance-covariance matrix Σ does not depend on Z_{t-1}. However, because we estimate (1) jointly across all assets, we do not assume the ε_t to be cross-sectionally uncorrelated, i.e. we do not assume Σ to be diagonal.

In-Sample Results

We focus on two sample periods, the first covering the period from October 1992 until April 2000, and the second covering the period from April 2000 until the end of our sample period (September 2005). The former captures the bull run leading up to and including the 'dot.com' bubble, while the latter covers the subsequent collapse of the bubble. The results of our in-sample estimation are reported in Tables 1 and 2, respectively. We first investigate pure market timing, using only the S&P 500 and the risk-free asset (Panel B1 in Tables 1 and 2). Next, we add either the bond index or gold[5] (Panel B2).

As predictive instruments we include the level of the VIX, and for each set of assets the corresponding open interest and hedging pressure variables. For the S&P alone over the first sample period (Table 1) these variables do not possess enough predictive power (with an R^2 of only 1.7%). This is reflected by our statistical test where, although the Sharpe ratio rises from 0.86 to 1.27, the p-value is only 8%. When we add the bond index, we find that the

[5]The return on gold is derived from the quoted price per troy ounce of Gold Bullion.

Sharpe ratio increases to 2.43, the difference having a p-value of 0. The maximal R^2 in this case is also considerably higher at 5.3%.

The in-sample results are quite different in the second sample period (Table 2). For the S&P alone, the R^2 rises sharply to 4.6%, and the Sharpe ratios increases from 0.25 to 1.62 (with a p-value of less than 1%). The maximal R^2 rises to 7.23% when gold is added to the asset universe, with Sharpe ratios more than tripling from 0.56 to 2.09 (the change being highly significant with a p-value of almost zero).

To illustrate the overall effect of our predictive variables, we also estimated the model on the full sample (from 1992 to 2005), using all three assets. Figure 1 shows the efficient frontiers with (solid line) and without (dashed line) the optimal use of predictive information.

3.3 How to Beat a Bear Market

Given the in-sample results it is natural to ask whether the increased predictive power of the instruments during the bear market could have been exploited by portfolio managers. We thus conduct an out-of-sample experiment, estimating the predictive model using data until April 2000 (pre-dating the collapse of the 'dot.com' bubble by several months), then constructing dynamically efficient portfolio strategies on the basis of the model estimates, and studying their performance over the remainder of the sample period. We focus here on strategies that are designed to minimize volatility at a given level of mean return, as we are interested in avoiding losses during this period. The results are reported in Panel B of Table 3.

We first consider a market timing strategy with the S&P 500, using the VIX, and hedging pressure and open interest on SPX futures as predictive variables. The fixed-weight strategy mirrors the decline of the S&P over this period, with a mean return of -0.75% per annum and a Sharpe ratio of -0.29. Our dynamically optimal strategy on the other hand has a much higher mean of 9.69% with a Sharpe ratio of 0.60. It achieves an (annualized) alpha of 8%

relative to the S&P by successfully 'de-coupling' from the market index (with a beta of only 0.21). An investor with a risk aversion coefficient of 5 would be willing to pay a management fee of over 1,100 basis points per annum to switch to this strategy.

The cumulative returns on both strategies, and the portfolio weights of the dynamic strategy, are shown in Panel (A) of Figure 2. Even using only market-timing, the dynamic strategy avoids sharp losses during the collapse of the 'dot.com' bubble and, and although designed to minimize variance, has a buy-and-hold return of over 50% over the 2000-2005 period, while the S&P sustains losses of over 20% during this period. The strategy, as illustrated by the portfolio weight shown in the bottom part of the figure, is very much a 'long-short' strategy and avoids the crash by shorting the S&P.

When we add gold as an additional asset, and include its hedging pressure and open interest as predictors, the performance of the strategies improves further, with the mean of increasing to 10.2% and the Sharpe ratio rising to 0.69. Although designed to minimize variance, the strategy is again a low-beta, high-alpha strategy (with a beta of 0.28 and an alpha of 8.9%). The management fee now increases to almost 1,300 basis points. The corresponding cumulative returns and portfolio weights are shown in Panel (C) of Figure 2.

Performance in Bull Markets

To see that this performance is a bear market phenomenon we conduct the same experiment during the bull market 1992-2000 period. We estimate the model during the first half of the period and run the strategy over the second half. In the case of using the S&P only, although the dynamic minimum-variance strategy has a higher mean and Sharpe ratio than the index, the volatility is almost twice as high at 34.2%. Even the minimum-variance strategy has now become a high-beta strategy (with a beta of 1.7), and as such is not fundamentally different from an index-tracking strategy. Given the high risk of the strategy, a moderately risk averse investor would not be willing to pay a positive management fee. The performance deteriorates even further when gold is added to the asset universe, achieving a much lower Sharpe ratio than the index itself.

3.4 How to Benefit From a Bull Run

During the 1992-2000 bull market, we consider dynamic strategies designed to maximize return rather than minimize variance. A pure market-timing strategy, while achieving an overall gain above the index, however experiences very sharp losses of more than 50% (see Panel A of Figure 3). In fact, the dynamic maximum-return strategy has a volatility of 67.2%, making it unacceptably risky. The picture changes dramatically when the bond index is included (Panel B). The dynamic strategy now outperforms the S&P. The out-of-sample maximum-return strategy has a mean return of 60.5% (and an alpha of 42%), but its volatility is only marginally higher than that of the index. The strategy has a low beta of 0.44 and maintains an almost constant moderate long position in the index, while actively managing a large position in the Treasury term spread (with a long position in T-bills and a short position of almost equal size in the bond index).

Unlike in a bear market, adding Gold does not provide any benefit at all in a bull market. Gold together with the S& only, while achieving a cumulative return similar to the index (Panel C of Figure 3), has a volatility of almost twice that of the index. In contrast, adding gold to the S&P and the bond index improves the strategy's performance only very marginally. In fact, in this case the position in gold is virtually constant and slightly negative. In other words, while gold is a valuable vehicle for hedging the risks in falling markets, it is not of any value in a bull market.

3.5 Transaction Costs

As our strategies are dynamically managed, the issue of transaction costs is of natural importance. Because the strategies are implementable using futures contracts, we assume transaction costs of 5 basis points (i.e. 5 cents per $1 transaction). We find that for all of our minimum-variance strategies, transaction costs are below 50 basis points per year, while for the maximum-return strategies are at most 160 basis points. While the latter may seem high, it is insignificant in comparison with cumulative (annualized) returns of 64.2%.

An important aspect in this is our focus on unconditionally efficient (as opposed to conditionally efficient) strategies. Unconditionally efficient strategies are theoretically optimal (Hansen and Richard 1987), utilize the predictive information more effectively (Abhyankar, Basu, and Stremme 2005) and also exhibit a conservative response to extreme values of the predictive variables (Ferson and Siegel 2001). In contrast, the portfolio weights of conditionally efficient strategies are much more volatile, often requiring extreme long or short positions. Moreover, the conditionally efficient strategies significantly underperform their unconditionally efficient counterpart out-of-sample, confirming the above findings. For example, the conditionally efficient minimum-variance strategy achieves a Sharpe ratio of only 0.23, compared with 0.60 for the unconditionally efficient strategy. On the other hand, the conditionally efficient strategy incurs transaction costs of almost 980 basis points per annum (compared with only 37 basis points for the corresponding unconditionally efficient strategy). Conversely, while conditionally efficient maximum-return strategies tend to be slightly cheaper than their unconditionally efficient counterparts, they also achieve a considerably lower return. We thus see that investors, and indeed empirical studies, utilizing the conditionally efficient strategy would have failed to observe any real economic gains from return predictability while those implementing the unconditionally efficient strategy would have.

3.6 Long Run Portfolio Performance

Next we investigate the long-run performance of the different strategies over the entire sample period.

Pure Market Timing

Focusing first on pure market-timing strategies, we find that overall the predictive variables do enlarge the investor's opportunity set with the Sharpe ratios increasing to 1.22 from 0.36, the difference being significant at the 1% level . The dynamic minimum-variance portfolio strategy tracks the target variance very closely, although the ex-post Sharpe ratio is not quite as high as the ex-ante value at 0.98. It is again a low-beta (0.34) and high-alpha (8.4%) strategy. A risk-averse investor (we used a moderate coefficient of risk aversion of 5 would be willing to give up 18.8 percentage points of annual return as management fee for the optimal strategy.

The dynamic maximum-return strategy in contrast slightly overshoots its volatility target, but achieves a much higher mean relative to the fixed-weight strategy (24.1% compared to 9.7%). This strategy performs as well as the index during the bull run, even though there is quite a bit of levering in and out of the index during this period and spectacularly outperforms the index after 2000. The strategy achieves an incredible (annualized) alpha of 15.2%, while its volatility is only about 2% higher than that of the index itself. An investor with a risk aversion coefficient of 5 would be willing to give up 10.2 percentage points of annual return as management fee for this strategy.

Adding Other Assets

In line with our out-of-sample results, adding gold or the bond index to the asset set (and their respective open interest and hedging pressure to the predictive variables) further improves performance. For example, adding gold increases the Sharpe ratio with predictability from 1.22 to 1.61. Relative to the fixed-weight case, the increase in Sharpe ratio is more than four-fold, with a p-value indistinguishable from zero. The results obtained by adding the bond index are very similar (with the Sharpe ratio increasing to 1.58). Including all assets further increases the Sharpe ratio to 2.02.

The performance of the corresponding dynamically efficient strategies match the theoretical predictions closely. The *ex-post* risk and return of both the minimum-variance and the maximum-return strategies are plotted in Figure 1, relative to the theoretical efficient frontier. Figure 4 shows the cumulative return and portfolio weights of the minimum-variance strategy that uses all assets (and their associated predictive instruments). The strategy achieves an (annualized) cumulative return of 16.4% at a volatility of only just over 5%.

The corresponding maximum-return strategy yields a return of 38.4% per annum at a moderate volatility of 15.2%.

3.7 How do the Strategies Work?

First, we note (Table 3) that our dynamic strategies tend to have much lower betas (between 0.2 and 0.5) than the fixed weight ones. Thus our variables seem to allow asset managers to 'de-couple' their portfolios from the business cycle and thus successfully time the market.

An inspection of the coefficients on the predictive variables (Tables 1 and 2) shows the coefficients on hedging pressure are very different in the two sample periods. In the bull market the beta is quite low (0.06) while it is 0.32 over the second period. The coefficients on the VIX are lower in the second period than in the first period but the difference is not particularly great (0.0006 versus 0.0004).

To gain insights into how the strategies work over the two periods, we examine the correlations between the weights on the risky asset and the various predictive variables for the two out of sample strategies. The correlation between the weights on the S&P and its hedging pressure for the strategy in the first period is 0.21 while it is -0.13 for the bear market strategy. This clearly shows how the nature of the strategy changes in the two periods and it appears to be the main driver of the success of the dynamic strategy in the second period. It seems that in the second period hedging pressure is driven by institutional investors and is high, for example when investors are covering short options positions and could be a leading indicator of market declines. Our strategy seems to be able to exploit this fact. In the bear market when hedging pressure has much greater predictive power, our strategy decreases the weight on the risky asset when the number of long positions increases and this significantly improves its performance. The correlation with open interest also changes from positive to negative in the two periods (0.06 to -0.38). The correlations with the VIX are positive and considerably higher over the second period (0.90 versus 0.62) but the correlation with changes in the VIX is negative in the second period (-0.13) while it is positive in the first period (0.15). Thus in the bear market a rising VIX causes the strategy to decrease holdings in the risky asset while the opposite is the case in the bull market.

4 Conclusions

The economic value of return predictability has been the focus of much recent research. While the statistical evidence is mixed, several studies have found that even small levels of predictability from a statistical viewpoint can have a substantial effect on both the investor's portfolio decision as well as portfolio performance. This paper focuses on the use of market variables that exploit the linkages between spot, futures and derivatives markets, as opposed to the business cycle indicators employed in most of the earlier studies.

Spot and futures market linkages are exploited by using hedging pressure as the predictive variable while the linkages between the derivatives and spot markets are exploited using the VIX index, a proxy for implied volatility. Using the S&P 500, gold and a 5 year Treasury Bond as our base assets, we study the performance of these variables by examining both the out of sample performance of unconditionally efficient portfolios based on our predictive variables as well as their in-sample performance using a statistical test. We find that different sets of variables have predictive power in bear and bull markets. Our trading strategies can successfully time the market and avoid losses during the bursting of the 'dot.com' bubble in the second half of 2000, as well as during the earlier bull run. The in-sample results confirm our out of sample experiments with p-values of less than 1% in all cases. The performance of the strategies deteriorates considerably if we remove any of the predictive variables, indicating that it is the correlations between these variables that drives the strategies.

References

Abhyankar, A., D. Basu, and A. Stremme (2005): "The Optimal Use of Asset Return Predictability: An Empirical Analysis," working paper, Warwick Business School.

Avramov, D., and T. Chordia (2006): "Predicting Stock Returns," forthcoming, Journal of Financial Economics.

Bessembinder, H. (1992): "Systematic Risk, Hedging Pressure and Risk Premiums in Futures Markets," *Review of Financial Studies*, 5, 637–667.

Breen, W., L. Glosten, and R. Jagannathan (1989): "Economic Significance of Predictable Variation in Stock Index Returns," *Journal of Finance*, 44, 1177–1189.

Campbell, J., and L. Vicera (2001): "Who Should Buy Long-Term Bonds?," *American Economic Review*, 91, 99–127.

Cochrane, J. (1999): "Portfolio Advice for a Multi-Factor World," *Economic Perspectives, Federal Reserve Bank of Chicago*, 23(3), 59–78.

Connolly, R., C. Stivers, and L. Sun (2005): "Stock Market Uncertainty and the Stock-Bond Return Relation," *Journal of Financial and Quantitative Analysis*, 40(1), 161–194.

De Roon, F., T. Nijman, and C. Veld (2000): "Hedging Pressure Effects in Futures Markets," *Journal of Finance*, 55(3), 1437–1456.

Ferson, W., and A. Siegel (2001): "The Efficient Use of Conditioning Information in Portfolios," *Journal of Finance*, 56(3), 967–982.

Fleming, J., C. Kirby, and B. Ostdiek (2001): "The Economic Value of Volatility Timing," *Journal of Finance*, 56(1), 329–352.

Fleming, J., C. Kirby, and B. Ostdiek (2003): "The Economic Value of Volatility Timing using 'Realized' Volatility," *Journal of Financial Economics*, 67, 473–509.

Hansen, L., and S. Richard (1987): "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55(3), 587–613.

Jagannathan, R. (1996): "Relation between the Slopes of the Conditional and Unconditional Mean-Standard Deviation Frontier of Asset Returns," in *Modern Portfolio Theory and its Applications: Inquiries into Asset Valuation Problems*, ed. by S. Saito et al. Center for Academic Societies, Osaka, Japan.

Kandel, S., and R. Stambaugh (1996): "On the Predictability of Stock Returns: An Asset Allocation Perspective," *Journal of Finance*, 51(2), 385–424.

Mayers, D. (1976): "Nonmarketable Assets, Market Segmentation and the Level of Asset Prices," *Journal of Financial and Quantitative Analysis*, 11(1), 1–12.

Whitelaw, R. (2005): "Time Varying Sharpe Ratios and Market Timing," working paper, New York University.

A Appendix

A.1 Dynamically Efficient Strategies

To specify a dynamically managed trading strategy, we denote by $\theta^k_{t-1} = \theta^k(Z_{t-1})$ the fraction of portfolio wealth invested in the k-th risky asset at time $t-1$, given as a function of the vector Z_{t-1} of (lagged) predictive instruments. The return on this strategy is given by,

$$r_t(\theta) = r_f + \sum_{k=1}^{n} (\, r^k_t - r_f \,)\theta^k_{t-1}, \tag{2}$$

where r^k_t is the return on the k-th risky asset, and r_f denotes the return on the risk-free Treasury bill. The difference in time indexing indicates that, while the return r_f on the risk-free asset is known at the beginning of the period, the returns r^k_t on the risky assets are uncertain *ex-ante* and only realized at the end of the period. Note however that we do *not* assume r_f to be unconditionally constant. It can be shown[6] that the weights of any unconditionally efficient managed strategy can be written as,

$$\theta^*_{t-1} = \frac{w - r^f_{t-1}}{1 + H^2_{t-1}} \cdot \Sigma^{-1}_{t-1} \, (\, \mu_{t-1} - r_f 1 \,). \tag{3}$$

Here, μ_{t-1} and Σ_{t-1} are the conditional (on Z_{t-1}) mean vector and variance-covariance matrix of the base asset returns, and $w \in \mathbb{R}$ is a constant. By choosing $w \in \mathbb{R}$ in (3) appropriately, we can construct efficient strategies that track a given target expected return or target volatility.

A.2 Measures of Statistical Significance

To measure the economic gain due to predictability, we measure the extent to which the optimal use of predictive information expands the unconditionally efficient frontier, i.e. the

[6] See for example Ferson and Siegel (2001), or Abhyankar, Basu, and Stremme (2005).

opportunity set available to the investor. In the absence of an unconditionally risk-free asset, the efficient frontier is described by three parameters, the location (mean and variance) of the GMV, and the asymptotic slope of the frontier (i.e. the maximum Sharpe ratio relative to the zero-beta rate corresponding to the mean of the GMV). Note however that because of the low volatility of T-bill returns, the location of the GMV will be virtually unaffected by the introduction of predictive instruments (see also Figure 3). Therefore, we focus here on the change in asymptotic slope of the frontier as a measure of predictability. Denote by λ_* the slope of the frontier *with* optimal use of predictability, and by λ_0 the slope in the fixed-weight case (without making use of predictive information). In a slight abuse of terminology, we often refer to λ_* and λ_0 simply as Sharpe ratios.

One can now show[7] that up to a first-order approximation, the (squared) maximum slope of the dynamically managed frontier is given by,

$$\lambda_*^2 \approx E(\, H_{t-1}^2 \,), \quad \text{where} \quad H_{t-1}^2 = (\, \mu_{t-1} - r_f 1 \,)' \, \Sigma_{t-1}^{-1} \, (\, \mu_{t-1} - r_f 1 \,). \tag{4}$$

Here, μ_{t-1} and Σ_{t-1} are the conditional mean vector and variance-covariance matrix of the base asset returns. The error in the above approximation is of the order $\text{var}(\, H_{t-1}^2 \,)$. To obtain the corresponding expression for λ_0, we simply replace μ_{t-1} and Σ_{t-1} by their unconditional counterparts.

Note that H_{t-1} is the *conditional* Sharpe ratio, once the realization of the conditioning instruments is known. From (3), it is clear that H_{t-1} plays a key role in the behavior of the optimal strategy. Moreover, the above result shows that the maximum unconditional Sharpe ratio is given by the unconditional second moment of the conditional Sharpe ratio[8]. Consequently, time-variation in the conditional Sharpe ratio improves the ex-post risk-return trade-off for the mean-variance investor, a point also noted by Cochrane (1999).

To measure the effect of predictability, we define the test statistic $\Omega = \lambda_*^2 - \lambda_0^2$. Our null

[7]See, for example, Abhyankar, Basu, and Stremme (2005).

[8]In the case of a single risky asset, this was shown by Jagannathan (1996).

hypothesis is that predictability does not matter, i.e. $\Omega = 0$. As the set of fixed-weight strategies is contained in the set of dynamically managed strategies, we always have $\Omega \geq 0$. In the linear predictive setting (1) used in our empirical analysis, one can show[9] that under the null, the test statistic

$$\frac{T-K-1}{K} \cdot \Omega \quad \text{is distributed as} \quad F_{K,T-K-1} \quad \text{in finite samples,}$$

and $T \cdot \Omega$ is distributed as χ^2_K asymptotically. Here, K is the number of instruments in Z_{t-1}, and T is the number of time-series observations. This result allows us to assess the statistical significance of the economic gains due to predictability.

A.3 Measures of Economic Value

In addition to our statistical tests, we also employ a utility-based framework to assess the economic value of return predictability. Following Fleming, Kirby, and Ostdiek (2001), we consider a risk averse investor whose preferences over future wealth are given by a quadratic von Neumann-Morgenstern utility function. They show that, if relative risk aversion γ is assumed to remain constant, the investor's expected utility can be written as,

$$\bar{U} = W_0\big(E(\,r_t\,) - \frac{\gamma}{2(1+\gamma)} E(\,r_t^2\,)\big),$$

where W_0 is the investor's initial wealth and r_t is the return on the portfolio they hold. Consider now an investor who faces the decision whether or not to acquire the skill and/or information necessary to implement the active portfolio strategy that optimally exploits predictability. The question is, how much of their expected return would the investor be willing to give up (e.g. pay as a management fee) in return for having access to the superior strategy? To solve this problem, we need to find the solution δ to the equation

$$E(\,r_t^* - \delta\,) - \frac{\gamma}{2(1+\gamma)} E\big(\,(r_t^* - \delta)^2\,\big) = E(\,r_t\,) - \frac{\gamma}{2(1+\gamma)} E(\,r_t^2\,), \tag{5}$$

[9]See for example Abhyankar, Basu, and Stremme (2005).

where r_t^* is the optimal strategy and r_t is a fixed-weight strategy that does not take predictability into account. The solution δ represents the management fee (as a fraction of portfolio returns) that the investor would be willing to pay in order to gain access to the superior strategy. Graphically, the premium can be found in the mean-variance diagram by plotting a vertical line downwards, starting from the point that represents the optimal strategy r_t^*, and locating the point where this line intersects the indifference curve through the point that represents the inferior strategy r_t.

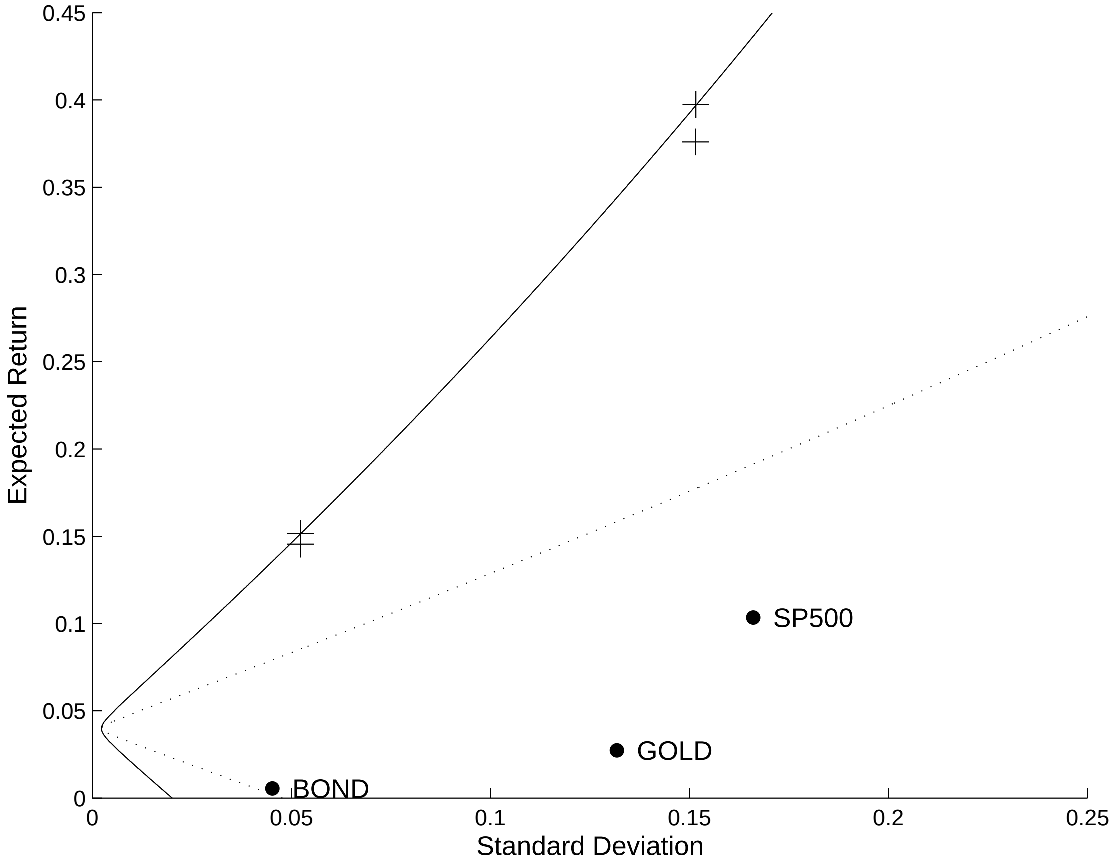

Figure 1: Efficient Frontier)

This graph show the unconditionally efficient frontiers with (solid line) and without (dashed line) the optimal use of predictability. We use all available assets (the S&P 500, the bond index, and gold), and all predictive instruments (the VIX, and open interest and hedging pressure for each asset). Also shown are the ex-post mean and variance ('+') of the maximum-return and minimum-variance strategies. The circles indicate the performance of these strategies net of transaction costs.

	Asset				
	RF	S&P500	BOND	GOLD	
Panel (A) Summary Statistics					
Average Return	4.9%	19.5%	−0.4%	−2.0%	
Volatility	0.1%	15.5%	4.4%	12.8%	
Sharpe Ratio		0.84	−1.19	−0.53	
Instrument	Regression Coefficients				
Panel (B1) S&P 500 only					
VIX		0.0006			
Open Interest		0.0000			
Hedging Pressure		0.0624			
				Maximum R^2	1.7%
				Fixed-Weight Sharpe Ratio	0.86
				Optimal Sharpe Ratio	1.28
				(p-Value)	(8.14%)
Panel (B2) S&P 500 + Bond					
VIX		0.0894	0.0006		
Open Interest		0.0242	0.0000		
Hedging Pressure		0.0586	−0.0073		
				Maximum R^2	5.3%
				Fixed-Weight Sharpe Ratio	1.68
				Optimal Sharpe Ratio	2.43
				(p-Value)	(0.03%)

Table 1: In-Sample Results (Bull Market, 1992 – April 2000)

This table reports the summary of the sub-sample estimates for the period leading up to and including the 'dot.com' bubble. The table reports the mean-variance performance (Panel A) of the assets themselves, the coefficients of the predictive regression and the theoretically optimal Sharpe ratios (Panel B). Panel (B1) focuses on pure market-timing, while in Panel (B2) we add the bond index to the asset set. The predictive instruments in all cases are the VIX, and open interest and hedging pressure for each of the assets. The p-values are obtained from the asymptotic χ^2-distribution of the test statistic Ω (see Appendix A.2).

	Asset				
	RF	S&P500	BOND	GOLD	
Panel (A) Summary Statistics					
Average Return	2.7%	−2.0%	2.0%	10.2%	
Volatility	0.3%	18.0%	4.6%	13.8%	
Sharpe Ratio		−0.26	−0.14	0.51	
Instrument	Regression Coefficients				
Panel (B1) S&P 500 only					
VIX		0.0004			
Open Interest		0.0000			
Hedging Pressure		0.3154			
				Maximum R^2	4.7%
				Fixed-Weight Sharpe Ratio	0.25
				Optimal Sharpe Ratio	1.62
				(p-Value)	(0.35%)
Panel (B2) S&P 500 + Gold					
VIX		0.0007		−0.0004	
Open Interest		0.0000		0.0000	
Hedging Pressure		0.2750		0.0065	
				Maximum R^2	7.2%
				Fixed-Weight Sharpe Ratio	0.56
				Optimal Sharpe Ratio	2.09
				(p-Value)	(0.06%)

Table 2: In-Sample Results (Bear Marker, April 2000 – 2005)

This table reports the summary of the sub-sample estimates for the period following the collapse of the 'dot.com' bubble. The table reports the mean-variance performance (Panel A) of the assets themselves, the coefficients of the predictive regression and the theoretically optimal Sharpe ratios (Panel B). Panel (B1) focuses on pure market-timing, while in Panel (B2) we add gold to the asset set. The predictive instruments in all cases are the VIX, and open interest and hedging pressure for each of the assets. The p-values are obtained from the asymptotic χ^2-distribution of the test statistic Ω (see Appendix A.2).

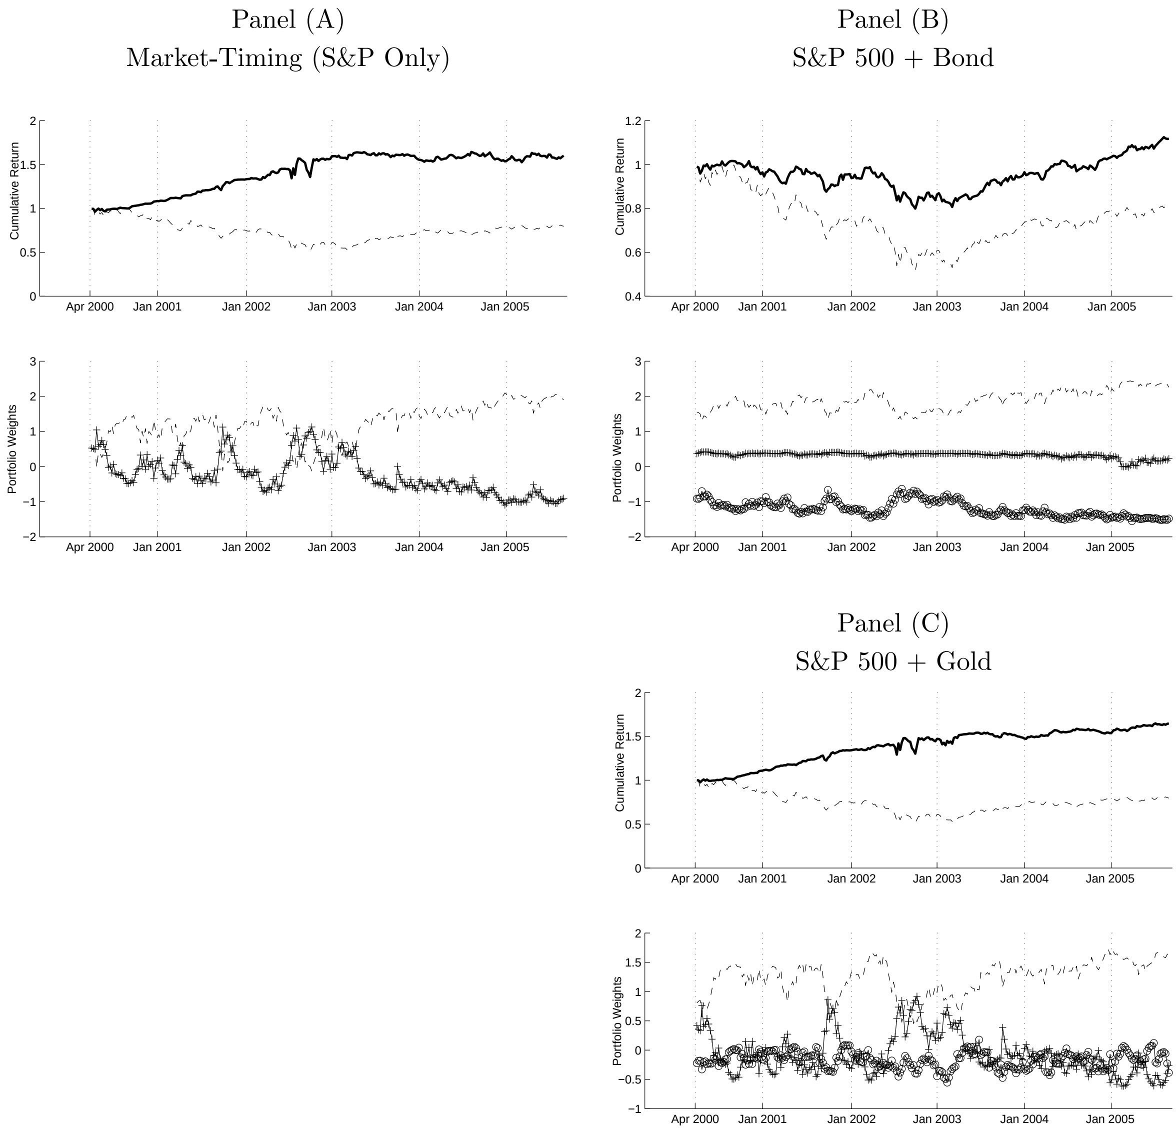

Figure 2: Out-of-Sample Performance (Bear Market, April 2000 – 2005)

These graphs show the performance and portfolio weights of three different strategies during the period following the collapse of the ‘dot-com’ bubble (from 2000 until 2005). In each panel, the top graph shows the cumulative return of the minimum-variance strategy (solid line) and the market index (dashed line), normalized to have unit value in April 2000. The bottom graph shows the portfolio weights on the risk-free asset (dashed line), the market index (‘+’), and the bond index or gold (‘∘’), respectively. Panels (A) focuses on market-timing (i.e. allocating between the risk-free asset and the index), while Panels (B) and (C) add the bond index or gold, respectively. The predictive instruments in all cases are the VIX, and open interest and hedging pressure for each of the assets.

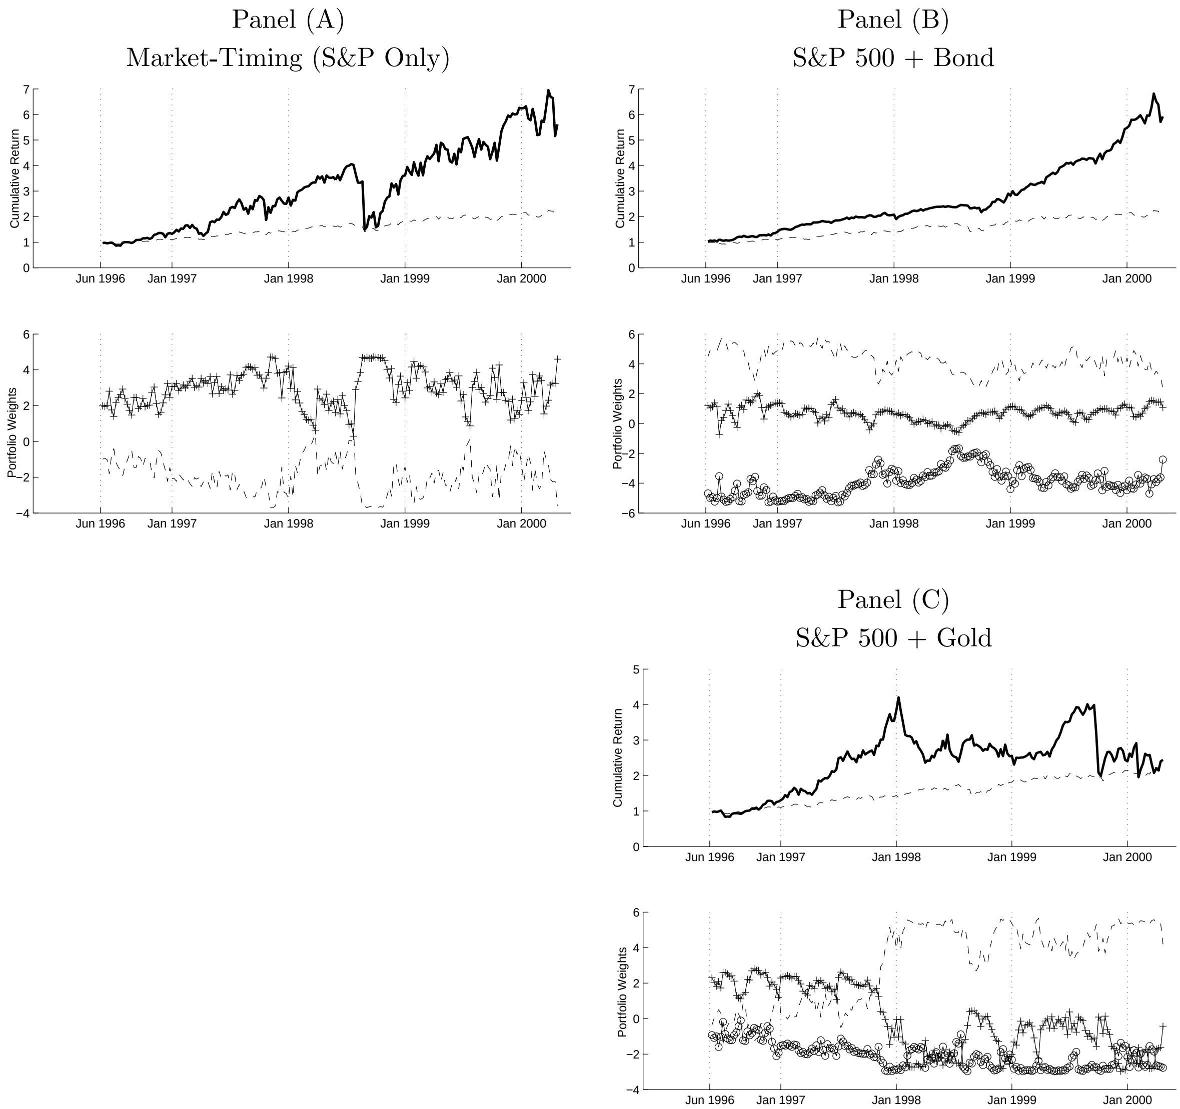

Figure 3: Out-of-Sample Performance (Bull Market, June 1996 – 2000)

These graphs show the performance and portfolio weights of three different strategies during the period leading up to and including the 'dot-com' bubble (from 1996 until 2000). In each panel, the top graph shows the cumulative return of the maximum-return strategy (solid line) and the market index (dashed line), normalized to have unit value in June 1996. The bottom graph shows the portfolio weights on the risk-free asset (dashed line), the market index ('+'), and the bond index or gold ('∘'), respectively. Panels (A) focuses on market-timing (i.e. allocating between the risk-free asset and the index), while Panels (B) and (C) add the bond index or gold, respectively. The predictive instruments in all cases are the VIX, and open interest and hedging pressure for each of the assets.

Panel (A) Bull Market (Maximum-Return Strategy)

	S&P Only		S&P + Bond	
	Fixed-Weight	Optimal	Fixed-Weight	Optimal
Expected Return	35.4%	98.9%	54.0%	60.5%
Volatility	31.3%	67.2%	29.0%	20.1%
Sharpe Ratio	0.80	0.95	1.32	2.10
CAPM Beta	1.60	3.29	1.33	0.44
Jensen's Alpha	0.0%	13.1%	18.9%	42.4%
Management Premium		−43.2%		16.63%

Panel (B) Bear Market (Minimum-Variance Strategy)

	S&P Only		S&P + Gold	
	Fixed-Weight	Optimal	Fixed-Weight	Optimal
Expected Return	−0.7%	9.7%	−2.1%	10.2%
Volatility	11.9%	11.0%	10.8%	10.1%
Sharpe Ratio	−0.29	0.60	−0.44	0.69
CAPM Beta	0.66	0.21	0.54	0.27
Jensen's Alpha	0.0%	8.0%	−2.0	8.8%
Management Premium		11.1%		12.9%

Table 3: Out-of-Sample Portfolio Performance

This table reports the out-of-sample performance of the maximum-return strategies (Panel A) during the bull run, and the minimum-variance strategy (Panel B) during the bear market, respectively, both with ('optimally managed') and without ('fixed-weight') the optimal use of predictability. The base assets are the risk-free asset, the S&P 500 index and, in the right-hand side column, the bond index (in Panel A) or gold (in Panel B), respectively. All figures are annualized. .

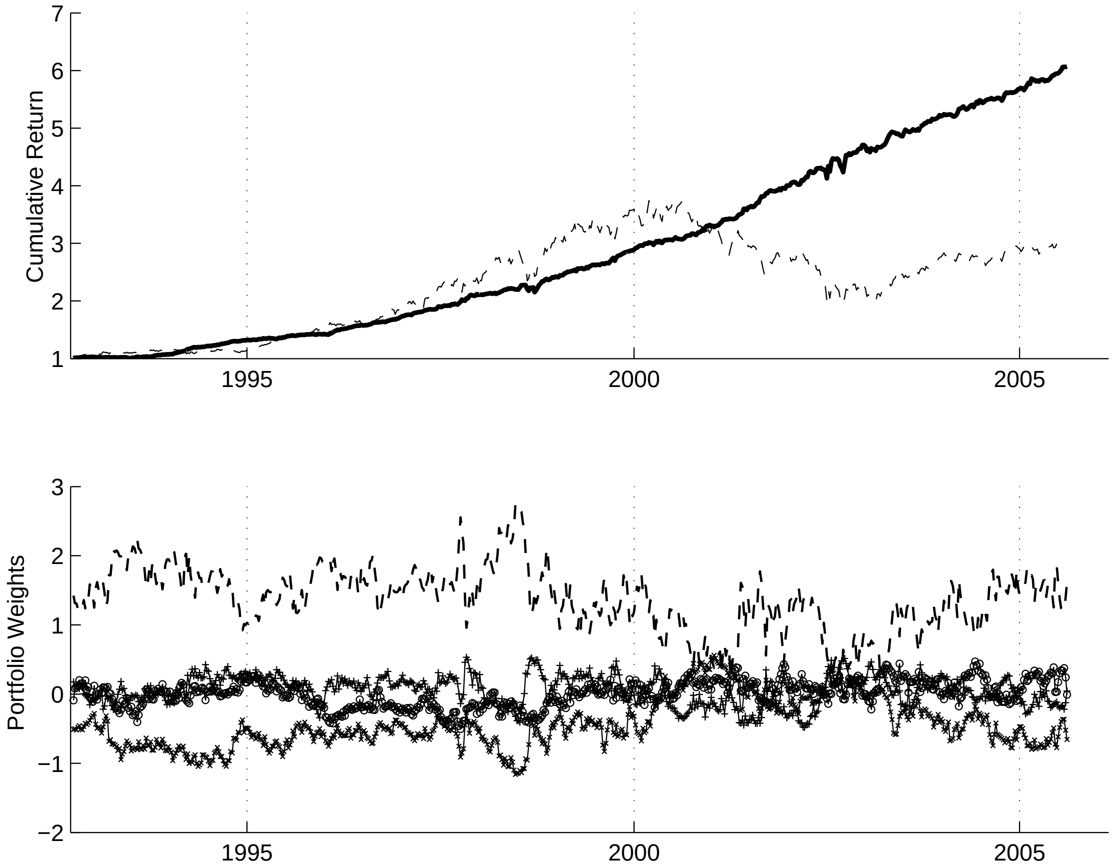

Figure 4: In-Sample Portfolio Performance)

This graph show the performance of the market index (dashed line) and the minimum-variance strategy (solid line) over the entire sample period (from 1992 to 2005). The top graph shows the cumulative returns, normalized to have unit value in October 1992. The bottom graph shows the portfolio weights on the risk-free asset (dashed line), the market index ('+'), and the bond index (×') and gold ('○'), respectively. The predictive instruments are the VIX, and open interest and hedging pressure for each of the assets.