EDHEC RISK AND ASSET
MANAGEMENT RESEARCH CENTRE

393-400 promenade des Anglais
06202 Nice Cedex 3
Tel.: +33 (0)4 93 18 32 53
E-mail: research@edhec-risk.com
Web: www.edhec-risk.com

Exploiting the Informational Content of Hedging Pressure: Timing the Market by "Learning" from Derivatives Traders

June 2007

Devraj Basu
EDHEC Business School

Roel Oomen
Warwick Business School, University of Warwick

Alexander Stremme
Warwick Business School, University of Warwick

EDHEC
BUSINESS SCHOOL
Lille - Nice

Abstract
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 commercial and non-reportable hedging pressure as the predictive variables 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 and gold 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. Our trading strategies can successfully time the market and avoid losses during the burst of the "dot.com" bubble in the second half of 2000, as well as during the bull run that followed. The in-sample results confirm our out-of-sample experiments with p-values of less than 1% in all cases. The predictive variables on their own do not perform nearly as well, indicating that it is linkages between these markets that are important for market timing. The VIX provides a signal to change the weight on the market while hedging pressure indicates the direction. We construct variables that combine both of these features and find that these variables provide the clearest signals for successful market timing.

JEL Classification: C12, C32, G11, G14

Address for correspondence: Alexander Stremme
Warwick Business School - University of Warwick
Coventry, CV4 7AL, United Kingdom
Tel: +44 (0) 2476 522 066
Email: alex.stremme@wbs.ac.uk
phone: +44 (0) 2476 - 522 066
fax: +44 (0) 2476 - 523 779

1. Introduction

The economic value of return predictability has been the focus of much recent research. The in-sample statistical evidence is controversial due to a number of statistical issues while the performance of most predictive variables declines significantly out-of-sample. Several recent studies have instead explored the economic significance of return predictability, by analyzing its effect on the asset allocation decisions of a mean variance investor. Kandel and Stambaugh (1996) and Wachter and Warusawitharana (2006) find that even small levels of predictability can lead to successful market timing, while Avramov and Chordia (2006) show that business cycle variables lead to improvements in out-of-sample portfolio performance even in periods where they have little forecasting power for the equity premium. Campbell and Thompson (2006) show that low out-of-sample R^2 are the norm, but that even such small R^2 are relevant for investors as they can lead to large improvements in portfolio performance. The implication of these studies is that the use of certain forecasting variables could significantly improve the performance of market timing strategies. All of the above studies however, focused on business cycle and macro-economic variables and several studies, both theoretical and empirical, suggest that the actions of market participants, particularly traders, could have a significant influence on the expected return on the market. In particular, since trading on an index such as the S&P 500 is done via the futures and options markets, the actions of traders in the derivatives markets could contain valuable information. We focus on exploiting the informational linkages between spot and derivatives markets using variables that summarize the outlook of traders in these markets. Our strategies are real-time strategies in that they depend on the actions of traders and do not rely on the use of predictive variables that were uncovered ex-post.[1]

Our predictive variables are based on the options and futures markets, where much of the trading on the S&P 500, our risky asset takes place. Our first predictive variable is the level of the VIX, which recent studies (for example Connolly, Stivers, and Sun 2005 and Copeland and Copeland 1999) have found useful for stock-bond diversification as well as market timing. The use of a proxy for volatility is also motivated by Fleming, Kirby, and Ostdiek (2001), and Fleming, Kirby, and Ostdiek (2003), who find that volatility timing has considerable economic value. We exploit linkages between spot and futures markets by extracting information about the net futures position held by different groups of traders from the CFTC"s "Commitment of Traders" report. The CFTC classifies each trader as either "commercial", "non-commercial", or "non-reportable". While the former group mainly consists of large institutional investors who use futures for hedging or risk allocation, the group of "non-reportable" traders includes small investors and hedge funds, using futures mainly for position-taking and speculation. Within each group, we focus on the "hedging pressure" at any given time, defined as the fraction of long futures positions relative to total open interest within this group. This may be regarded as a proxy for non-marketable risks which, in an incomplete market, should affect the expected returns on assets (as implied for example by 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. The effect of hedging pressure on spot markets is less well studied, with only De Roon, Nijman, and Veld (2000) documenting some predictive ability. Our study is the first to explore the informational content of hedging pressure, in particular the differences between different groups of traders, for optimal asset allocation.

Using lagged information from the derivatives markets as predictive instruments, we construct dynamically managed trading strategies in the spot market. These are designed to be *dynamically efficient*, thus using tactical asset allocation to attain a strategic objective (i.e., unconditional efficiency), in contrast to the *conditionally efficient* strategies used in most earlier studies[2] which are myopically optimal (i.e. efficient relative to the conditional moments). As we are interested mainly in market-timing strategies, we focus on the S&P 500 as the risky asset. We also include gold as an investible asset, to allow strategies to move out of the equity markets at times of economic downturn. Each period, the strategies allocate funds between these two assets and a conditionally risk-free asset (proxied by the Federal Funds rate). Predictive information is updated and portfolios are re-balanced at a weekly frequency.

We first investigate whether out-of-sample, the information derived from the derivatives markets can predict extreme market movements and thus help avoid losses in market crashes or periods of high volatility. We find

1 - For example, Cooper, Gutierrez, and Marcum (2006) make the point that rules based on the book-to-market ratio were unlikely to have been used by mutual funds in the 1980s before the Value index was created.

2 - For example Fleming, Kirby, and Ostdiek (2001), and Fleming, Kirby, and Ostdiek (2003), among others.

that pure market-timing strategies based on the VIX and the actions of large commercial hedgers or small speculators would not have incurred any significant losses during the collapse of the "dot.com" bubble in late 2000. The strategies move in and out of the risky asset frequently and for short periods, a characteristic of successful market timing strategies.[3]

In both bull and bear markets, our strategies clearly out-perform the market index (with alphas of between 5 and 12% annually). These strategies seem to allow asset managers to "de-couple" their portfolios from the business cycle during a bear market (with very low betas around 0.15) and thus successfully time the market. Using a quadratic utility specification as in Fleming, Kirby, and Ostdiek (2001), a moderately risk averse investor would be willing to pay annual management fees of between 900 and 1600 basis points to gain access to the optimal strategy over the 2000-2005 period. Adding gold as an additional hedge further improves the performance of timing strategy during the bear market, with the management fee rising to almost 1700 basis points. Both sets of variables are required for successful market timing, strategies based on either the VIX or hedging pressure alone yield considerably inferior performance.

Our findings suggest that the actions of small speculators, in conjunction with the VIX, have the most predictive power during the post 2000 bear market, consistent with the findings in Barber, Odean, and Zhu (2006) and Hvidkjaer (2006). Indeed, a perusal of the position sizes shows that non-reportable hedging pressure increased in the second half of 2000 and thus seemed to provide a good contrary indicator for market returns in this period. The actions of large commercial hedgers as a predictive variable also help improve portfolio performance, although not as much. However, it is the interaction between these two groups of traders that contains the most predictive information.

Our strategies also seem to perform well during a bull market. A market-timing strategy over the 2002-2005 period when the market recovered after the "dot.com" crash, using non-reportable hedging pressure, would have generated an alpha of just over 5%, while the strategy with gold added and both hedging pressures had an alpha of nearly 6%.

In summary, we find that the information contained in derivatives trades does possess considerable predictive power for the spot market. However, our findings indicate that neither the VIX nor hedging pressure alone are sufficient for successful market timing. It appears that the VIX, unsurprisingly, predicts periods of increased market movement, while hedging pressure helps predict direction. Strategies using the VIX only tend to move into the market prior to both upswings as well as downturns, while hedging pressure "tells" managers when to seek shelter in the risk-free asset or in gold. To test this interpretation, we construct a composite variable that combines the VIX with hedging pressure and find that the resulting market-timing strategies perform almost as well as strategies based on all variables. We also find that the correlation of the portfolio weights with the combined variable is much higher than that with each of the variables on their own. Finally, to assess the statistical significance of our results, we perform an in-sample analysis. The increase in portfolio performance is significant at the 1% level in all cases considered.

The remainder of this paper is organized as follows. Section 2 describes our methodology, while the predictive variables are described in Section 3. The results of our empirical analysis are presented in Section 4, while Section 5 describes and discusses the performance of our directional variables. Section 6 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

3 - In contrast, strategies based on changes in interest rates tend to "get out of" the market too late and then stay out for too long, performing worse than the index itself.

of the former are determined 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.[4] This is an important consideration in particular with respect to transaction costs.

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.

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, we use a variety of standard *ex-post* portfolio performance measures. These include Sharpe ratios, Jensen"s alpha, and information ratios.

<u>Measures of Statistical Significance</u>

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 to measure the extent to which the optimal use of predictive information expands the efficient frontier, and hence the opportunity set available to the investor. We show that the difference in (squared) slopes of the frontiers with and without the optimal use of predictability has a known (χ^2) distribution, and 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.

<u>Measures of Economic Value</u>

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 his 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. Data Description

As base assets, we use weekly returns on the S&P 500 and gold. The data are obtained from the CBOE.[5] 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. The second is hedging pressure for different categories of traders, which is defined as total long positions divided by the sum of long and short positions. These data are obtained from the CFTC"s "Commitment of Trader" report (see following section).

3.1. The CFTC"s "Commitment of Traders" Report

The *Commodity and Futures Trading Commission* (CFTC) is a regulatory body that is entrusted with preserving the US futures markets" key economic role, namely that of price discovery and risk sharing or hedging. The CFTC runs a comprehensive market surveillance program that monitors trading activity in US futures markets on a

4 - See also Ferson and Siegel (2001) or Abhyankar, Basu, and Stremme (2005).
5 - www.cboe.com

daily basis in order to ward off price manipulation, market squeezes, and other abusive practices. In doing so, it compiles a large-trader report (LTR) showing the futures and options positions that a "reporting firm" (i.e., clearing members, futures commission merchants, foreign brokers, and traders) holds over and above specific reporting levels set by the commission. For instance, for S&P 500 futures, a reporting firm needs to maintain a position in excess of 1,000 contracts for it to be included in the LTR. The aggregate position of all reporting firms typically makes up for about 70-90 percent of the open interest and thus gives a fairly comprehensive overview of trading activity in any given futures market monitored by the CFTC. Futures positions that are held for other purposes than pure hedging, such as speculation or arbitrage, are subject to speculative limit rules (for the S&P 500 futures, for example the speculative limit is 20,000 contracts). Because of this, the CFTC classifies reporting firms as "commercial" when a trader uses a particular futures contract for hedging as defined in the Commission"s regulations and as "non-commercial" otherwise. Of course, reporting firms that trade in multiple markets simultaneously may be classified as commercial in one and as non-commercial in another. Also, the Commission routinely reviews these classifications. For instance, on June 2, 2006 the CFTC reports, *"As part of this ongoing review process, Commission staff recently interviewed several traders who were classified as Commercial and had significant open contracts in at least one of the foreign currency futures markets. Based on information obtained in these interviews, several traders have been reclassified from Commercial to Noncommercial because their hedging or risk management activities, although present, did not constitute a significant part of their overall market position".*[6]

The daily LTR are not publicly available. Instead, the CFTC publishes a commitment of traders (COT) report every Friday afternoon at 3:30 p.m. Eastern time which details the preceding Tuesday"s aggregate number of long and short positions of commercial and non-commercial traders, together with the residual long and short positions of the "non-reportable" firms,[7] for markets in which 20 or more reportable firms are active. For non-commercial firms, the COT report also includes the number of long contracts that are offset by short contracts, known as "spreading" positions. Table 1 contains some summary statistics of the COT data for the assets we focus on below. As pointed out by Haigh, Hranaiova, and Overdahl (2005), among others, this data is highly aggregated. Still, in this paper we show that it contains valuable information for the purpose of portfolio allocation decisions in general and the timing of commodity markets in particular. The variable we focus on is *hedging pressure*, defined as the fraction of long positions within a particular classification. We thus have three hedging pressure variables, i.e., for commercial traders, non-commercial traders, and non-reportable traders.

Table 1: COT Summary Statistics

	No. of Traders	Open Interest	Non-Commercial			Commercial		Non-Reportable	
			Long	Short	Spread	Long	Short	Long	Short
Panel A: 2000 S&P 500	149	403	17	20	2	275	310	108	70
Panel B: 2005 S&P 500	185	677	57	66	7	475	497	137	107

This table reports the average number of traders, plus the open interest and long, short and spread positions (in 000"s of contracts) within each classification

4. Empirical Analysis

In this section we first describe the out-of-sample results over the period May 2000-end 2005, followed by the out-of-sample performance over the 2002-2005 period. We confirm these findings with an in-sample analysis over the entire sample period (1992-2005).

4.1 How to Beat a Bear Market

In order to ascertain whether a portfolio manager could benefit from using dynamically managed strategies with these predictive variables, we 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

6 - Source: www.cftc.gov.

7 - Note that the sum of long or short positions by commercial, non-commercial, and non-reportable firms totals the open interest.

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 2.

We first consider a pure market-timing strategy using the S&P 500 as the risky asset. 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. We try various combinations of the predictive variables, first using the VIX and each of the different hedging pressure variables separately. The results are reported in Table 2. We find that commercial hedging pressure, together with the VIX, leads to a considerable improvement, with the mean return increasing to 8.4% with a Sharpe ratio of 0.44 and information ratio of 0.97 while non-commercial hedging pressure leads to a marginal improvement, with a mean return of 4.5% and a Sharpe ratio of 0.12. The most dramatic improvement results from using non-reportable hedging pressure, with the mean return rising sharply to 13.67%, the Sharpe ratio jumping to 1.12, and the information ratio to 1.30. The cumulative returns and the portfolio weights of this strategy are shown in Figure 2. Even using only market-timing, the dynamic strategy avoids sharp losses during the collapse of the "dot.com" bubble and, although designed to minimize variance, has a buy-and-hold return of over 80% 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. The strategy has a low beta (0.13) and high annualized alpha (11.37%). An investor with a risk aversion coefficient of 5 would be willing to pay a management fee of over 1,600 basis points per annum to switch to this strategy. The strategies move in and out of the risky asset for short periods and frequently, similar to the recommendations of professional market timers who are investment advisors. This category of market timers seems to have had the greatest success, found by as Chance and Helmer (2001). A market-timing strategy based on the short rate, by contrast, moves into the risk-free asset in the first half of 2001 after the biggest falls in the market and remains in the risk-free asset for the life of the strategy and the performance of this strategy is worse than that of the index.

These findings suggest strongly that the actions of small speculators have the most predictive power during the post-2000 bear market. Indeed a perusal of the position sizes shows that non-reportable hedging pressure increased in the second half of 2000 and thus seemed to provide a good contrary indicator, reflected in the fact that the weight on the risky asset is negatively correlated (-0.39) with it. The actions of large commercial hedgers also help improve portfolio performance although there is almost no correlation between the risky asset weight and non-commercial hedging pressure. Non-commercial hedging pressure, which reflects the activities of commercial speculators, mainly hedge funds presumably, has the power to improve portfolio performance. A possible explanation is that the activities of commercial speculators are so diverse that they provide no clear signal.

We next add a safe asset, gold, to the asset set. The mean return on gold over this period was 9.49% compared to -2.54% for the S&P. The performance of the fixed-weight strategy in fact declines when we add gold with the mean return now -2.07%. However for the dynamically managed strategies, adding gold improves the performance using non-reportable hedging pressure, while the performance of the strategy utilizing commercial hedging pressure declines relative to using the index alone. It is clear from this that our strategies are not just about moving into gold, as one strategy has a considerably higher mean return (12.93%) while the other has a considerably lower (4.87%) mean return than gold. Combining commercial and non-reportable hedging pressure leads to the best performance, as can be seen from Table 2. Figure 2 shows the cumulative return and portfolio weights of the strategy. The unconstrained strategy has a buy and hold return of over 90% over the 2000-2005 period with a Sharpe ratio of 1.28 and an information ratio of 1.38, while the constrained or long-only strategy has a cumulative return of 50% with lower volatility with a Sharpe ratio of 1.03 and an information ratio of 1.30. The main effect of adding gold seems to stabilize the portfolio weights in that extreme long short positions are eliminated. This strategy is again low beta (0.04 for unconstrained, 0.16 for constrained) and high alpha (11.08% for unconstrained, 8.39% for constrained) with the management fee for an investor with a risk aversion coefficient of 5 being 1,779 basis points for the unconstrained and 1,118 basis points for the constrained strategy.

4.2. Performance over a Bull Run

Our analysis thus far has shown that our conditional asset allocation strategies estimated during a bull run (1992-2000) perform well during a bear market. It is interesting to see how well the strategies perform during an up market when estimated over a bear market. To this end we focus on the bull run towards the end of our sample period, estimate our predictive regression over April 2000-April 2002, and study the performance of the strategies considered above from May 2002 to the end of the sample (September 2005). The results are reported at the bottom of Table 2. The mean return on the index over this period was 5.74% and the Sharpe ratio was 0.21. The minimum variance market timing strategy using commercial hedging pressure and the VIX has a Sharpe ratio of 0.53 and an information ratio of 0.57, and has a relatively high beta (0.79) with a modest alpha (2.06%). When we replace commercial hedging pressure with non-reportable hedging pressure the results are very similar, except that the strategy is now low beta (0.08). Adding gold and combining commercial and non-reportable hedging pressure leads to the best performing strategy again with the minimum variance strategy having a mean return of 6.58% and a low volatility of 5.14% leading to a Sharpe ratio of 0.92 and an information ratio of 0.98. In contrast, although gold has a considerably higher mean return of 12.08% it has much higher volatility of 14.39%, leading to a lower Sharpe ratio of 0.68. The cumulative return and portfolio weights of this strategy are shown in Figure 3, from which we can see that the portfolio weights are rarely less than zero or greater than one, i.e it is very close to being a long only strategy. It is again a low beta (0.15) and moderate alpha (4.21%) strategy. The cumulative return of the long only maximum return strategy is also shown in Figure . The cumulative buy and hold return of this strategy is almost 50% with a mean return of 13.44% higher than that for gold and a volatility of 12.35% more than 2% lower than that of gold. These strategies thus perform well in both bull and bear markets.

4.3. In-Sample Analysis

We now examine the in-sample performance of the predictive variables and strategies over the entire sample period (October 1992-September 2005). This is to assess the statistical significance of our results, to ensure that our out-of-sample results are not merely an artifact of the time period chosen. The regression results are in Table 1.

Pure Market-Timing Strategies

Using commercial hedging pressure and the VIX does enlarge the investor"s opportunity set with the optimal Sharpe ratio rising to 1.20 from 0.36with the difference having a p-value of .0002, with an R^2 of 2.48%. The minimum variance strategy with a target mean of 15% achieves a mean of 14.95% and a Sharpe ratio of 0.98 reasonably close to the *ex-ante* value. It is again a low beta (.35) and high alpha (8.31%) strategy.

When we replace commercial hedging pressure with non-reportable hedging pressure the optimal Sharpe ratio increases slightly to 1.26 with a p-value of .0001 and a higher R^2 of 2.74%. The minimum variance strategy almost achieves its target mean at 14.99%, and the strategy"s Sharpe ratio is 1.07. It is again low beta (0.30) and high alpha 8.61%.

Adding Gold as an Investible Asset

If we now add gold to the asset set and consider both commercial and non-reportable hedging pressure the optimal Sharpe ratio rises to 1.46 from 0.36 with the difference having a p-value of .0003, and a maximum R^2 of 3.68%. The minimum variance strategy has a Sharpe ratio of 1.33 and is again a low beta and high alpha strategy.

4.4. 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. We find that for all of our minimum-variance strategies, transaction costs are below 60 basis points per year, while for the maximum-return strategies they are at most 140 basis points. While the latter may seem high, it is insignificant in comparison with cumulative (annualized) returns of 32.2%. 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 eficient strategies significantly under-perform their unconditionally efficient counterpart out-of-sample, confirming the above findings. For example, the conditionally efficient minimum-variance strategy using all assets and variables achieves a Sharpe ratio of 0.42, compared with 1.12 for the unconditionally efficient strategy. On the other hand, the conditionally efficient strategy incurs transaction costs of almost 195 basis points per annum (compared with only 55 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, unlike those implementing the unconditionally efficient strategy, investors and indeed empirical studies utilizing the conditionally efficient strategy would have failed to observe any real economic gains from return predictability.

5. How do the Strategies work?

We first analyze the bear market strategies. To understand how the strategies work we calculate the correlations between the portfolio weights and the predictive variables. We find that for the pure market timing strategies the weights on the risky asset are positively correlated with the VIX (0.84 for commercial hedging pressure and 0.62 for non-reportable hedging pressure) implying that the high levels of the VIX are a signal to invest in the market, i.e., higher risk implies higher market returns, which is consistent with our estimation period being a bull market. This by itself is not enough in a bear market and market-timing strategies using the VIX alone do not perform well. The success of these strategies depends crucially on the additional signals provided by hedging pressure. The risky asset weight for the strategy utilizing commercial hedging pressure is positively correlated with commercial hedging pressure, but the correlation is quite low (0.10), while that with non reportable hedging pressure has a higher negative correlation with it (-0.32). The strategy utilizing non-commercial hedging pressure performs better showing that it is a reliable contrarian indicator, but neither hedging pressure variable does well without the VIX. Our findings thus suggest that both sets of variables are required for successful market timing, with the VIX providing a signal to change the weight on the market, while hedging pressure provides the signal to increase or decrease.

5.1 Construction of VHP and VHPN

Motivated by this we construct a directional variable based on both these indicators. We first consider deviations from the 1st percentile for commercial hedging pressure and 99th percentile for non reportable hedging pressure, in order to eliminate the influence of outliers for each variable. Commercial hedging pressure is positively correlated with index returns and the difference between its minimum and the 1st percentile is almost 7% of its range while non-reportable hedging pressure is negatively correlated with index returns and the difference between its maximum and the 99th percentile is over 11% of its range, indicating that the effect of outliers could be quite considerable. We then multiply each of these with the level of the VIX and label them VHP and VHPN respectively. The variable, VHP, constructed from commercial hedging pressure has relatively low correlation with the VIX (0.22) and higher correlation with commercial hedging pressure (0.72) while VHPN constructed from non-reportable hedging pressure also has relatively low negative correlation with the VIX (-0.33) and relatively correlation with non-reportable hedging pressure (0.65). This pattern of correlations suggest that these variables are directional in nature, providing signals to move in or out of the market.

We first examine the pattern of correlations between these variables and the strategies. For the strategy utilizing commercial hedging pressure the correlation between the risky asset weight and VHP is 0.62, the same order of magnitude as that of the VIX while the correlation between it and VHPN for the strategies with non-commercial hedging pressure is -0.79, higher than that with it and the VIX. This indicates quite strongly that these signed variables are what appear to drive the pure market timing strategies. In the case of gold and both hedging pressures, the SP weight is positively correlated with VHP (0.66) which is higher than that for the VIX (0.63) and commercial hedging pressure (0.26). It is negatively correlated with VHPN (-0.78) higher in magnitude than the VIX or non-reportable hedging pressure (-0.37). The weight on gold has a strong positive correlation with both VHPN and non-reportable hedging pressure (0.66 and 0.73 respectively), low negative correlation with VHP and commercial hedging pressure (-0.19 and -0.25 respectively) and low correlation with the VIX (0.08). Thus VHPN and non-reportable hedging pressure provides the strongest signal to invest in the "safe" asset while the VIX provides virtually no signal for such an investment, suggesting that its role in optimal asset allocation is quite different from the "fear gauge" perception of market investors.

The correlation patterns are very similar over the bull run. The weight on the risky asset for the pure market timing strategy using commercial hedging pressure has a higher correlation with VHP (0.89) than either the VIX (0.73) or commercial hedging pressure (0.28), while that for non-reportable hedging pressure has a correlation of -0.67 with VHPN comparable in magnitude to that for the VIX (0.76) and higher than that for non-reportable hedging pressure (-0.13).

5.2. Performance of Strategies using VHP and VHPN

To further confirm our findings we run our market-timing strategies using VHP and VHPN. The results using VHP are considerably better than those with commercial hedging pressure, with the Sharpe ratio rising to 1.30 from 0.44, with both the average return increasing from 8.39 to 11.78 and the volatility reducing to 7.14% from 12.08%. Using VHPN we get results very similar to those with the VIX and non-reportable hedging pressure with a Sharpe ratio of 1.12 compared to 1.20. Adding gold and using both VHP and VHPN leads to a Sharpe ratio of 1.24, very close to that obtained from using the VIX and both hedging pressures (1.28).

A comparison of the risky asset weights involving commercial hedging pressure and VHP (Fig 4) shows that strategy using VHP goes short the market during the period of sharpest losses (mid 2000-mid 2001), while that using commercial hedging pressure is occasionally long in the market over this period, just before short lived rallies. The signed variable VHP thus seems to be a more reliable indicator of longer term market movements, as is clear from comparing the top and middle graphs in Panel (B) of Figure 4.

All of these strategies are highly correlated with their predictive variables. The market-timing strategy using VHP has a correlation of 0.93 with the risky asset weight, in contrast to 0.62 for the less successful strategy involving the VIX and commercial hedging pressure.

The strategy involving VHPN has a correlation of -0.88 with its risky asset weight similar to that for the strategy involving non reportable hedging pressure, and both strategies achieve a similar level of performance.

The strategy with gold and both variables has high correlations on the index weight with VHP and VHPN (0.77 and -0.88) while the weight on gold has reasonable positive correlation with VHPN (0.61), but low correlation with VHP, consistent with our earlier findings.

It is significant to note that the correlation of the index weights on these strategies with the VIX is much lower than the ones involving hedging pressure. The correlations are around 0.25 for all the three strategies, again illustrating the fact that successful market-timing strategies require a directional indicator such as those constructed here.

Our strategies perform similarly over the bull market. Replacing the VIX and commercial and non-reportable hedging pressure with VHP and VHPN respectively produces very similar results with the Sharpe ratios using VHP and VHPN being slightly higher in all cases.

The strategies with gold and both hedging pressures over the bull run behave very similarly to the bear market ones with the correlations with the SP weight and VHP and VHPN similar in magnitude to that with the VIX (0.89, -0.64 and 0.71) respectively and higher than that with either hedging pressure, while the weight on gold is positively correlated with both VHPN and non-reportable hedging pressure (0.47 and 0.59 respectively) and has low correlations with all of the others (less than 0.10 in magnitude).

Thus, in summary, the strategy using VHP outperforms that with commercial hedging pressure, while that with VHPN performs as well as that with non reportable hedging pressure. VHP and VHPN are highly correlated with all the strategies while the VIX is highly correlated only with the strategies involving itself.

Finally we focus on the in-sample performance of strategies using VHP and VHPN. The pure market timing strategies achieve Sharpe ratios very similar to those with the VIX and hedging pressure, with p-values much less than 1% in all cases. Adding gold and using both the variables leads to the highest Sharpe ratio of 1.26, with a p-value of 0.63%.

6. 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 on portfolio performance. This paper focuses on the use of market variables that exploit the links between spot, futures, and derivatives markets, as opposed to the business cycle indicators employed in most earlier studies.

Spot and futures market links are exploited by using commercial and non-reportable hedging pressure as the predictive variables, representing the actions of hedgers and small speculators, while the links between the derivatives and spot markets are exploited using the VIX index, a proxy for implied volatility. Using the S&P 500 and gold 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.

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 bull run that followed. 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 drive the strategies. The VIX provides a signal to change the weight on the market, while hedging pressure indicates the direction. We construct variables that combine both of these features and find that these variables provide the clearest signals for successful market timing, while strategies based on these variables perform at least as well and sometimes out-perform strategies based on the VIX and hedging pressure.

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*.

- Barber, B., T. Odean, and N. Zhu (2006): "Do Noise Traders Move Markets?," working paper, UC Davis and UC Berkeley.

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

- Campbell, J., and S. Thompson (2006): "Predicting Excess Stock Returns Out-of-Sample: Can Anything Beat the Historical Average?," forthcoming, *Review of Financial Studies*.

- Chance, D., and M. Helmer (2001): "The Performance of Professional Market Timers: Daily Evidence from Executed Strategies," *Journal of Financial Economics*, 62, 377-411.

- 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.

- Cooper, M., R. Gutíerrez, and B. Marcum (2005): "On the Predictability of Stock Returns in Real Time," *Journal of Business* 78(2), 469-500.

- Copeland, M., and T. Copeland (1999): "Market Timing: Style and Size Rotation Using the VIX," *Financial Analysts Journal*, 55(2), 73-82.

- 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.

- Haigh, M. S., J. Hranaiova, and J. Overdahl (2005): "Price Volatility, Liquidity Provision and the Role of Managed Money Traders in Energy Futures Markets," working paper, US Commodity Futures Trading Commission.

- 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.

- Hvidkjaer, S. (2006): "Small Trades and the Cross-Section of Stock Returns," working paper, University of Maryland.

- 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.

- Wachter, J., and M. Warusawitharana (2006): "Predictable Returns and Asset Al-location: Should a Skeptical Investor Time the Market?," working paper, University of Pennsylvania.

Appendix

A.1. Dynamically Efficient Strategies

To specify a dynamically managed trading strategy, we denote by $\theta_{t-1}^k = \theta^k(Z_{t-1})$ 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,

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

where r_t^k 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_t^k 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[8] that the weights of any unconditionally efficient managed strategy can be written as,

$$\theta_{t-1}^* = \frac{w - r_{t-1}^f}{1 + H_{t-1}^2} \cdot \Sigma_{t-1}^{-1} \left(\mu_{t-1} - r_f 1 \right). \tag{2}$$

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 $w \in \mathrm{IR}$ is a constant. By choosing $w \in \mathrm{IR}$ in (2) 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 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[9] that up to a first-order approximation, the (squared) maximum slope of the dynamically managed frontier is given by,

$$\lambda_*^2 \approx E(H_{t-1}^2), \quad \textit{where} \quad H_{t-1}^2 = (\mu_{t-1} - r_f 1)' \Sigma_{t-1}^{-1} (\mu_{t-1} - r_f 1). \tag{3}$$

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 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 (2), 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.[10] 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 hypothesis is that predictability does not matter, i.e., $\Omega = 0$. As the set of fixed-weight strategies is contained in the set of

8 - See Ferson and Siegel (2001), or Abhyankar, Basu, and Stremme (2005).
9 - See Abhyankar, Basu, and Stremme (2005).
10 - In the case of a single risky asset, this was shown by Jagannathan (1996).

dynamically managed strategies, we always have $\Omega \geq 0$. In the linear predictive setting used in our empirical analysis, one can show[11] that under the null, the test statistic

$$T \cdot \Omega \quad \text{is distributed as} \quad \chi^2_{N \times K} \quad \text{asymptotically}$$

where N is the number of assets, 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 this 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

$$E(r_t^* - \delta) - \frac{\gamma}{2(1+\gamma)}E\big((r_t^* - \delta)^2\big) = E(r_t) - \frac{\gamma}{2(1+\gamma)}E(r_t^2), \tag{4}$$

where r_t^* is the optimal strategy and r_t is a fixed-weight strategy that does not take pre-dictability 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.

11 - See for example Abhyankar, Basu, and Stremme (2005).

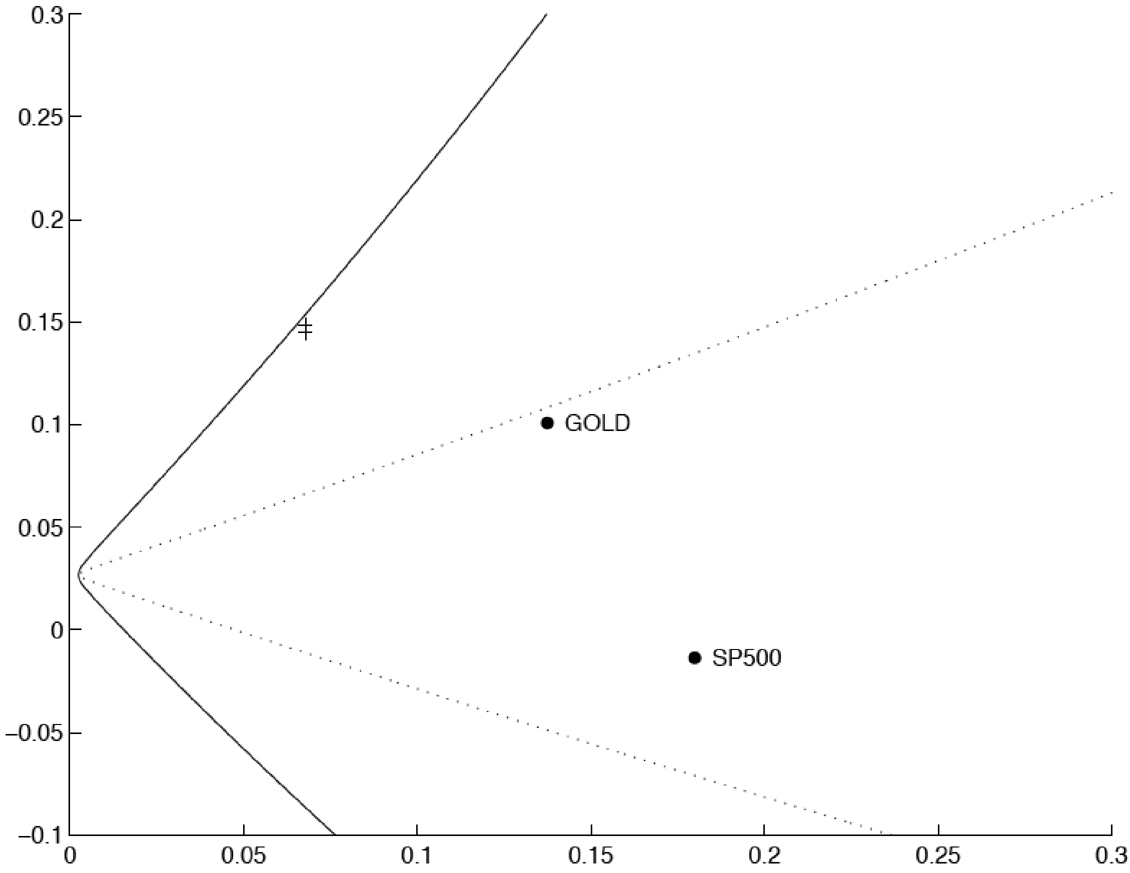

Figure 1: Efficient Frontier)

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 and gold), and all predictive instruments (the VIX, commercial and non-reportable hedging pressure on the the S&P 500). 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	GOLD	
Panel (A) Summary Statistics				
Average Return	3.99%	10.33%	2.75%	
Volatility	0.24%	16.63%	13.21%	
Sharpe Ratio		0.357	−0.091	
Instrument	Regression Coefficients			
Panel (B1) S&P 500 only				
VIX		0.0485		
CHP		0.0309		
NRHP		−0.0378		
			R^2	2.78%
			Fixed-Weight Sharpe Ratio	0.358
			Optimal Sharpe Ratio	1.272
			(p-Value)	(0.02%)
Panel (B2) S&P 500 +Gold				
VIX		0.0485	−0.0163	
CHP		0.0309	0.0660	
NRHP		−0.0378	0.0476	
			Maximum R^2	3.68%
			Fixed-Weight Sharpe Ratio	0.363
			Optimal Sharpe Ratio	1.456
			(p-Value)	(0.03%)

Table 1: In-Sample Results (October 1992 - September 2005)

This table reports the summary of the full sample estimates for the entire sample period (October 1992 to September 2005). 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 commercial and non-reportable hedging pressure. The p-values are obtained from the asymptotic χ^2-distribution of the test statistic Ω (see Appendix A.2).

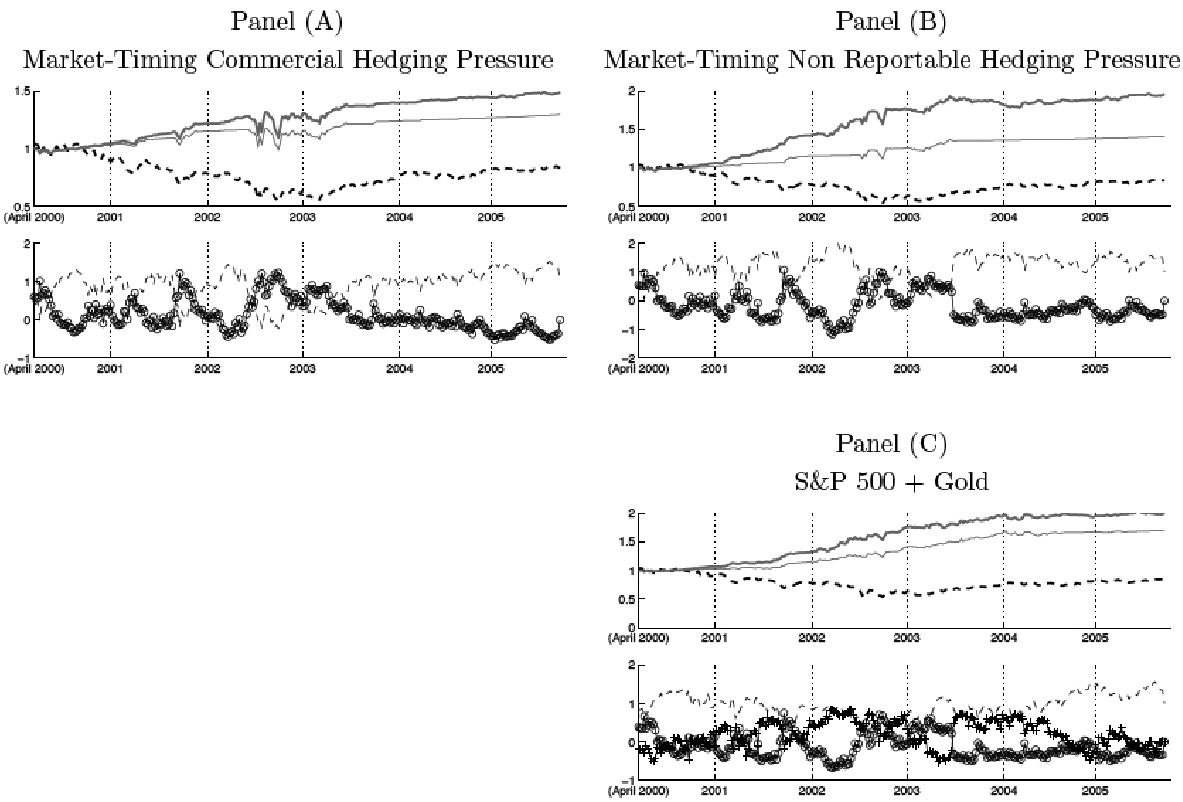

Figure 2: Out-of-Sample Performance (Bear Market, April 2000 - 2005)

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 (o), respectively. Panels (A) and (B) focus on market-timing (i.e., allocating between the risk-free asset and the index) using the VIX and commercial and non reportable hedging pressure respectively, while in Panels (C) we add the bond index, with the predictive variables being the VIX and both hedging pressures.

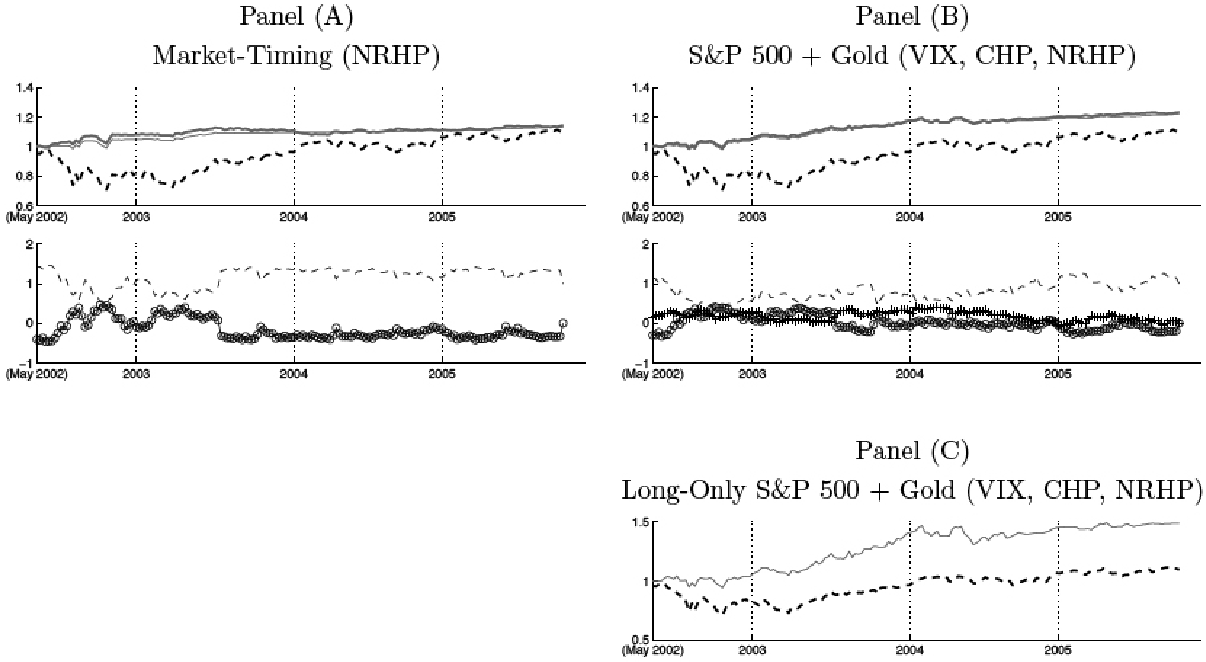

Figure 3: Out-of-Sample Performance (Bull Market, May 2002 - 2005)

These graphs show the performance and the portfolio weights of three different strategies, during the bull market following the collapse of the "dot.com" bubble (from 2002 until 2005). In each panel, the top graph shows the cumulative return of the strategy (solid line) and the market index (dashed line), normalized to have unit value in May 2002. The bottom graph shows the portfolio weights on the risk-free asset (dashed line), the market index (+), and gold (o), respectively. Panels (A) and (B) focus on minimum-variance strategies, while Panel (C) shows a long-only maximum-return strategy. Panel (A) focuses on market-timing only, while in Panels (B) and (C) gold is added to the asset set. In Panel (A), the predictive instruments are the VIX and commercial hedging pressure (CHP), while in Panels (B) and (C) also non-reportable hedging pressure (NRHP) is added.

Panel (A) Bear Market (May 2000 – Sept 2005)

	S&P Only			S&P + Gold	
	Fixed Weight	Optimal		Fixed Weight	Optimal
		CHP	NRHP		ALL
Expected Return	-0.74%	8.39%	13.66%	-2.07%	13.06%
Volatility	11.93%	12.08%	9.04%	10.80%	8.04%
Sharpe Ratio	-0.288	0.443	1.200	-0.444	1.282
CAPM Beta	0.657	0.477	0.125	0.535	0.040
Jensen's Alpha	0.0%	8.19%	11.37%	-0.02%	11.08%
Information Ratio	0.0	0.973	1.300	−0.416	1.382
Management Premium		9.09%	16.24%		17.79%

Panel (B) Bull Run (May 2002 – Sept 2005)

	S&P Only			S&P + Gold	
	Fixed Weight	Optimal		Fixed Weight	Optimal
		CHP	NRHP		ALL
Expected Return	−0.37%	4.40%	4.24%	0.72%	6.58%
Volatility	9.51%	5.52%	5.15%	9.73%	5.14%
Sharpe Ratio	−0.212	0.528	0.486	−0.095	0.920
CAPM Beta	−0.510	0.786	0.083	−0.520	0.152
Jensen's Alpha	0.0%	2.06%	2.20%	0.01	4.21%
Information Ratio	0.0	0.565	0.448	0.752	0.981
Management Premium		11.89%	6.29%		7.62%

Table 2: Out-of-Sample Portfolio Performance

This table reports the out-of-sample performance of the efficient strategies during the collapse of the "dot. com" bubble (Panel A, estimation period 1992 - April 2000, evaluation period May 2000 - end 2005), and the subsequent bull market (Panel B, estimation period 2000 - April 2002, evaluation period May 2002 - end 2005). Reported are the out-of-sample performance of the portfolios with ("optimally managed") and without ("fixed-weight") the use of predictive information. The base assets are the risk-free asset, the S&P 500 index and, in the right-hand side column, gold. The predictive instruments are the VIX, and either commercial (CHP) and non-reportable (NRHP) hedging pressure in the left-hand column, and all three instruments in the right-hand column. All figures are annualized.

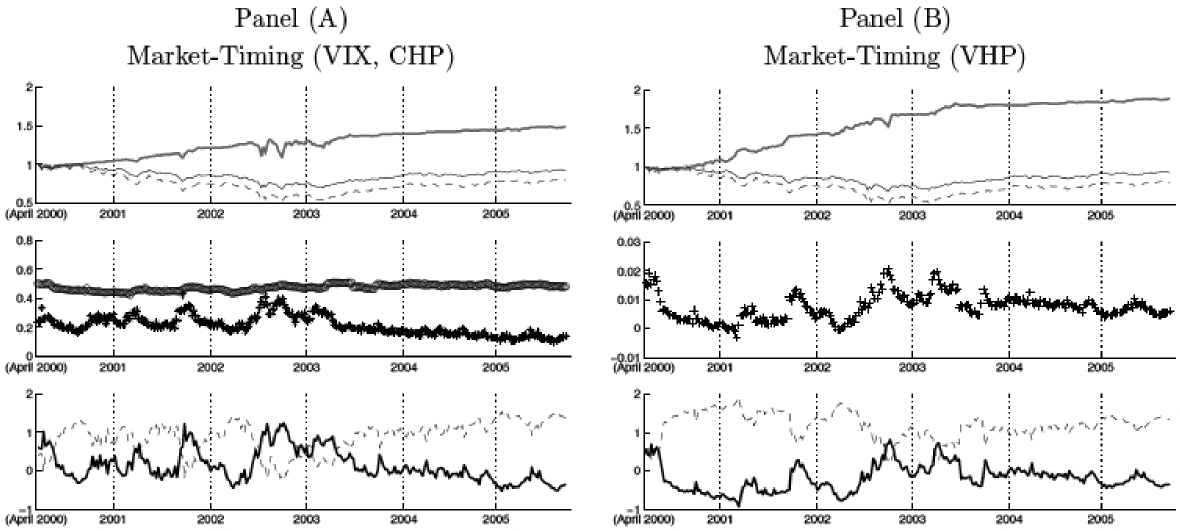

Figure 4: Comparison with Directional Indicator

These graphs show the performance, the values of the predictive instrument(s), and the portfolio weights of two market-timing 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 strategy (solid line) and the market index (dashed line), normalized to have unit value in April 2000. The middle graph shows the time series of the predictive variable(s), and the bottom graph shows the portfolio weights on the risk-free asset (dashed line), and the market index (+). In Panel (A) the instruments are the VIX and commercial hedging pressure (CHP), while Panel (B) uses the compound directional variable VHP (see Section 5.1).