

How to Time the Commodity Market

Devraj Basu, Roel C.A. Oomen and Alexander Stremme *

June 2006

Abstract

Over the past few years, commodity prices have experienced the biggest boom in half a century. In this paper we investigate whether it is possible by active asset management to take advantage of the unique risk-return characteristics of commodities, while avoiding their excessive volatility. We show that observing (and learning from) the actions of different groups of market participants enables an active asset manager to successfully ‘time’ the commodities market. We focus on the information contained in the ‘*Commitment of Traders (COT)*’ report, published by the CFTC. This report summarizes the size and direction of the positions taken by different types of traders in different markets. Our findings indicate that there is indeed significant informational content in this report, which can be exploited by an active portfolio manager. Our dynamically managed strategies exhibit superior out-of-sample performance, achieving Sharpe ratios in excess of 1.0 and annualized alphas relative to the S&P 500 of around 15%.

*Basu is from the Faculty of Finance, Cass Business School, while Oomen and Stremme are from the Department of Finance, Warwick Business School, Coventry CV4 7AL, United Kingdom. Oomen is also a research affiliate of the Department of Quantitative Economics at the University of Amsterdam, The Netherlands. Contact e-mail: roel.oomen@wbs.ac.uk

1 Introduction

Over the past few years, commodity prices have experienced the biggest boom in half a century. Oil prices have almost tripled from around \$25 in 2002 to more than \$70 by the end of 2005. Over the same period, copper prices have more than quadrupled. Unsurprisingly therefore, investment in commodities has been gathering momentum with particularly the fund management industry becoming increasingly attracted to this asset class. However, the sharp decline in prices in May and June 2006 has lead some analysts to forecast the end of the commodities boom. The purpose of this study is to show that, by actively managing a portfolio that includes commodities, it is possible to take advantage of the unique risk-return characteristics of commodities, while avoiding their excessive volatility and hence reducing the risk of substantial losses.

Much of the attraction of commodities appears to be the fact that they seem to produce equity-like returns while having low or even negative correlation with equities, thus potentially providing significant diversification benefits. A recent study by Gorton and Rouwenhorst (2006) supports this view, finding considerable evidence that the inclusion of commodities in a portfolio can improve its risk-return characteristics. The results of this and other similar studies were widely publicized, and received attention well beyond the academic community, which might be partly responsible for the renewed interest of the asset management industry in commodities. At the same time however, there have been calls for caution from academics and practitioners alike. For example, Erb and Harvey (2006) point out that simply adding commodities to an equity portfolio does not necessarily guarantee superior performance, thus emphasizing the need for active management. There is also the issue whether the publicity received by ‘bullish’ studies like that by Gorton and Rouwenhorst (2006) may itself affect investor behavior and thus alter the outlook for commodity investment. These concerns seem to be shared by a number of prominent members of the asset management industry.

For example, Bill Miller of Legg Mason Capital Investment notes in his April 2006 Market Summary¹,

‘The reason to own commodities may be that one believes they provide equity like returns with little correlation with equities. The time to own commodities is (or at least has been) when they are down, when everybody has lost money in them, and when they trade below the cost of production. That time is not now. The data showing the returns of commodities will look very different if you start measuring just after prices have tripled ... I can’t help but be skeptical of the advice to start or increase a position in commodities AFTER the biggest bull move in 50 years.’

The sudden collapse of commodities prices in May and June 2006 seems to validate some of his concerns.

In this paper we investigate whether it is possible by active asset management to take advantage of the diversification benefits of commodities without falling victim of their very high volatility. More specifically, we ask whether observing (and learning from) the actions of different groups of market participants would enable an active asset manager to successfully ‘time’ the commodities market. To this end, we focus on the information contained in the ‘*Commitment of Traders (COT)*’ report, published by the CFTC. This report summarizes the size and direction of the positions taken by different types of traders in different markets (details are given in the next section). Our findings indicate that there is indeed significant informational content in this report, which can be exploited by an active portfolio manager. However, our results also suggest that it is not sufficient to simply ‘follow the lead’ of a specific group of traders, but that it is instead the interaction between the behavior of the different groups that conveys the information necessary for successful market timing.

¹leggmason.com

We construct long-only dynamically managed portfolio strategies and show that these would have allowed a portfolio manager to successfully time the commodity boom, and indeed exit the market before the ‘crash’ of commodity prices in 2006. Our analysis is conducted purely out-of-sample in the sense that the data we use to estimate the relationship between the COT information and asset returns does not overlap the period over which we assess the performance of our strategies. Our findings are similar in spirit to Erb and Harvey (2006), who find that the use of certain market characteristics (such as the term structure of futures prices) in portfolio strategies involving commodities have historically been rewarded with above-average returns.

Our results imply that the CFTC data can be used to construct reliable signals that tell a portfolio manager when to switch from equities into commodities and vice-versa. In particular, even when we estimate our model only until early 2000, the resulting strategy successfully exits the equity market in mid-2000 (just before the collapse of the ‘dot.com’ bubble), and increases its exposure to commodities. Similarly, while our strategies ‘ride’ part of the commodities boom, they remain conservative, and in particular trade out of commodities before the ‘crash’ of May 2006. In summary, the question we address here is not, ‘could you have made money if you had been invested in commodities in the period from 2002 until the end of 2005?’, but rather ‘could you have known when to get into commodities and when to get out?’. With the benefit of hindsight, the answer to the former question is of course ‘yes’, our results provide an affirmative answer even to the second question.

The remainder of this paper is organized as follows. In section 2 we describe the COT report and the predictive variables constructed from it. Section 3 specifies the assets and the asset allocation model. Our results are reported in section 4, and section 5 concludes.

2 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 daily basis in order to ward against 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 are subject to speculative limit rules (e.g. for the S&P 500 futures 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 one. Also, the Commission routinely reviews these classifications. For instance, on June 2, 2006 the CFTC reports²,

‘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 activi-

²Source: www.cftc.gov.

ties, although present, did not constitute a significant part of their overall market position'

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, 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 here is *hedging pressure*, defined as the fraction of long positions within a particular classification. For instance, if commercial hedging pressure is close to 1 then the majority of commercial traders are long futures contracts, which necessarily implies that the non-commercial and non-reportable traders on the other side of the trade are short ³. We thus have three "hedging pressure" variables, i.e. for commercial traders, non-commercial traders, and non-reportable traders.

3 Data and Portfolio Construction

In this paper we consider a long-only portfolio manager who allocates funds between equities, commodities, and the risk-free asset. In particular, we focus on the S&P 500 index, plus

³Note the sum of long or short positions by commercial, non-commercial, and non-reportable firms totals the open interest.

Table 1: COT Summary Statistics

	No. of Traders	Open Interest	Non-Commercial			Commercial		Non-Reportable	
			Long	Short	Spread	Long	Short	Long	Short
<i>Panel A: 2000</i>									
S&P 500	149	403	17	20	2	275	310	108	70
Copper	120	73	16	8	3	40	56	14	7
Oil	165	468	46	30	29	347	364	46	45
<i>Panel B: 2005</i>									
S&P 500	185	677	57	66	7	475	497	137	107
Copper	177	105	36	21	15	40	57	14	12
Oil	252	817	119	109	127	511	507	60	74

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.

copper and oil, and the 1-month CD rate as a conditionally risk-free asset. The predictive variables for the S&P 500 are commercial and non-reportable hedging pressure, which have been shown to work for timing the market (Basu, Oomen, and Stremme 2006), and non-commercial hedging pressure for both copper and oil. Our strategies are executed at a weekly frequency with re-balancing taking place every Monday (following the Friday COT release).

Our managed portfolios are constructed as follows; using historical data, we first estimate a predictive model to capture the relation between our information variables and the returns on the base assets. Based on the estimated model parameters, we then construct dynamic portfolio strategies by specifying a mechanism that maps any given realization of the predictive instruments into a set of portfolio allocation weights. We then compute the returns on the resulting strategies throughout the out-of-sample period, and assess their performance using a variety of industry-standard measures. Specifically, we begin by estimating a

predictive regression of the form,

$$R_t = \mu_0 + BZ_{t-1} + \varepsilon_t, \quad (1)$$

where R_t is the vector of risky asset returns, r_f is the return on the (conditionally) risk-free asset (CD), and Z_{t-1} is the vector of (lagged) predictive instruments. While we assume the residuals ε_t to be serially independent and identically distributed (i.i.d.), we do allow cross-sectional correlation (in other words, the residual variance-covariance matrix Σ of the ε_t is assumed to be constant but not necessarily diagonal). 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) instruments. The return in period t on this strategy is given by, $r_t(\theta) = r_f + (R_t - r_f 1)' \theta_{t-1}$, where $\theta_{t-1} = (\theta_{t-1}^1 \dots \theta_{t-1}^n)'$ is the n -vector of asset allocation weights. One can show⁴ that the weights θ_{t-1} of the efficient dynamic portfolio with given target mean m are given by,

$$\theta_{t-1} = \frac{w(m) - r_f}{1 + H_{t-1}^2} \Sigma^{-1}(\mu_{t-1} - r_f 1), \quad (2)$$

where $\mu_{t-1} = \mu_0 + BZ_{t-1}$ is the vector of conditional expected asset returns, H_{t-1} is the maximum conditional Sharpe ratio⁵, and $w(m)$ is a constant that depends on the target mean m . In this paper, we use the constrained version of (2), for which the weights θ_{t-1} are required to be non-negative (i.e. no short selling), and their sum not to exceed 1 (i.e. no borrowing). From (2), it is clear that the conditional Sharpe ratio H_{t-1} plays a key role in the behavior of the optimal strategy. Moreover, it can be shown that the maximum (squared) *unconditional* Sharpe ratio is given by $\lambda_*^2 = E(H_{t-1}^2)$, which implies that it is variability in the conditional Sharpe ratio which leads to increases in the unconditional Sharpe ratio. The goal of dynamic asset allocation in this setting is to find predictive instruments that lead to high levels of inter-temporal variation in the conditional Sharpe ratio.

⁴See Ferson and Siegel (2001).

⁵It is well known that $H_{t-1}^2 = (\mu_{t-1} - r_f 1)' \Sigma^{-1} (\mu_{t-1} - r_f 1)$.

4 Results

We perform two out-of-sample experiments. The first estimates the predictive regression until April 2000 and then uses the estimates as inputs to construct the dynamically managed strategy. The performance of this strategy is then assessed in the out-of-sample period from May 2000 until May 2006. Over this period, the S&P 500 had a negative annualized return (-0.8%) while copper and oil had annualized returns of 34% and 24% respectively, with Sharpe ratios of 1.31 and 0.52, indicating that oil was much more volatile than copper.

A long-only manager who intended to form a static (fixed-weight) portfolio of these assets, using historical estimates of means, volatilities and correlations of these assets but none of the predictive variables, would have only generated an average annual return of 1.4% with a volatility of 10.1%, leading to a negative Sharpe ratio of -0.14 . In contrast, the optimal use of the predictive information contained in hedging pressure variables in a long-only strategy would have increased performance dramatically, with a return of 16.8% and a volatility of 10.3% (as compared to 20% for copper and 36% for oil), leading to a Sharpe ratio of 1.23, close to that of copper. The strategy has a high annualized alpha (relative to the S&P 500) of 13.9% and a low CAPM beta of about 0.1. The cumulative returns and weights of this strategy are shown in Figure 1. It is clear from the figure that the strategy is primarily conservative with a significant fraction invested in the risk-free asset much of the time. It weathers the collapse of the ‘dot.com’ bubble by remaining fully invested in the risk free asset. It then trades off the S&P against copper, keeping a relatively stable position in oil. Interestingly, around the beginning of 2005, the weight on copper drops to zero and the strategy invests larger and larger fractions in the S&P. Finally, in May 2006 the strategy entirely abandons the equity and commodities markets and becomes fully invested in the safe asset, thus successfully avoiding the losses incurred by the risky assets. It seems that the strategy follows the advice of Bill Mason (without ‘being aware’ of it)!

To check the robustness of our results, we conduct the same experiment over the period from May 2002 until May 2006. All the risky assets have returns considerably greater than

the risk-free CD, with Sharpe ratios of 0.27 for the S&P, and 1.99 and 0.85 for copper and for oil, respectively. The performance of the long-only static strategy now improves, with a mean of 7.3% and a Sharpe ratio of 0.45. The dynamically managed strategy's performance improves more dramatically, with the mean return increasing to 18.5% and the Sharpe ratio rising to 1.64. During this period, the strategy is almost entirely de-coupled from the equity index (with a beta of virtually zero), allowing it to achieve an active annual alpha of 15.8%. The cumulative return and weights are shown in Figure 2, and are very similar to those in Figure 1. Our strategies thus seem to work well in both bull and bear markets.

We thus conclude that dynamic asset allocation using hedging pressure would have helped a long only portfolio manager time the commodity boom and would also have told him to exit before the market experienced a significant correction.

From our results, it seems that copper on its own would have out-performed our strategy in both periods. The question hence arises whether it would not have been optimal to simply invest all funds in copper. While it is true that with the benefit of hindsight, a copper-only strategy would have achieved similar performance, our dynamic strategy is forward-looking and does not rely on hindsight. Moreover, the performance of copper is specific to the chosen sample periods. Over other periods, we find that our strategy often out-performs individual commodities, sometimes by a wide margin, but virtually never under-performs.

4.1 How do the Strategies Work?

In order to further analyze how these strategies work we examine the correlations between the weights on the various base assets and the different hedging pressure variables, for the out-of-sample period from May 2000 until May 2006. We first consider the weight on the S&P 500 and its associated measures of commercial and non-reportable hedging pressure. The correlation with commercial hedging pressure is positive (0.39), indicating that increases in commercial hedging pressure lead the strategy to increase investment in the S&P. However, the reverse is true for non-reportable hedging pressure, where the correlation is negative and

larger (-0.66) which indicates that this variable is a reliable contrarian indicator, i.e. it tells us when to get out of the market.

Non-commercial hedging pressure on both copper and oil are positively correlated with the weight on the respective asset (0.44 and 0.57). As non-commercial hedging pressure represents the activities of large speculators such as hedge funds, increases in their long positions means is interpreted as a bullish signal for the respective commodity. Finally, the analyze correlations between commercial and non-reportable hedging pressure on the S&P and the sum of weights on copper and oil. We find that non-reportable hedging pressure has a considerable positive correlation (0.55) with this sum, indicating that not only does it provide a signal to exit the equity market, it also tells the strategy to move into commodities (rather than the risk-free asset). The correlation with commercial hedging pressure is low (-0.12).

It is also interesting to note that the strategy exits copper around the second half of 2005, which coincides with a drop in non-commercial hedging pressure for copper. In fact, hedging pressure for copper was between 0.7 and 0.8 until May 2005, after which it dropped sharply and remained mostly in the range between 0.5 and 0.6 . The strategy responds to this shift by moving completely into the safe asset in May 2006, most likely due to a rise in S&P non-commercial hedging pressure.

5 Conclusion

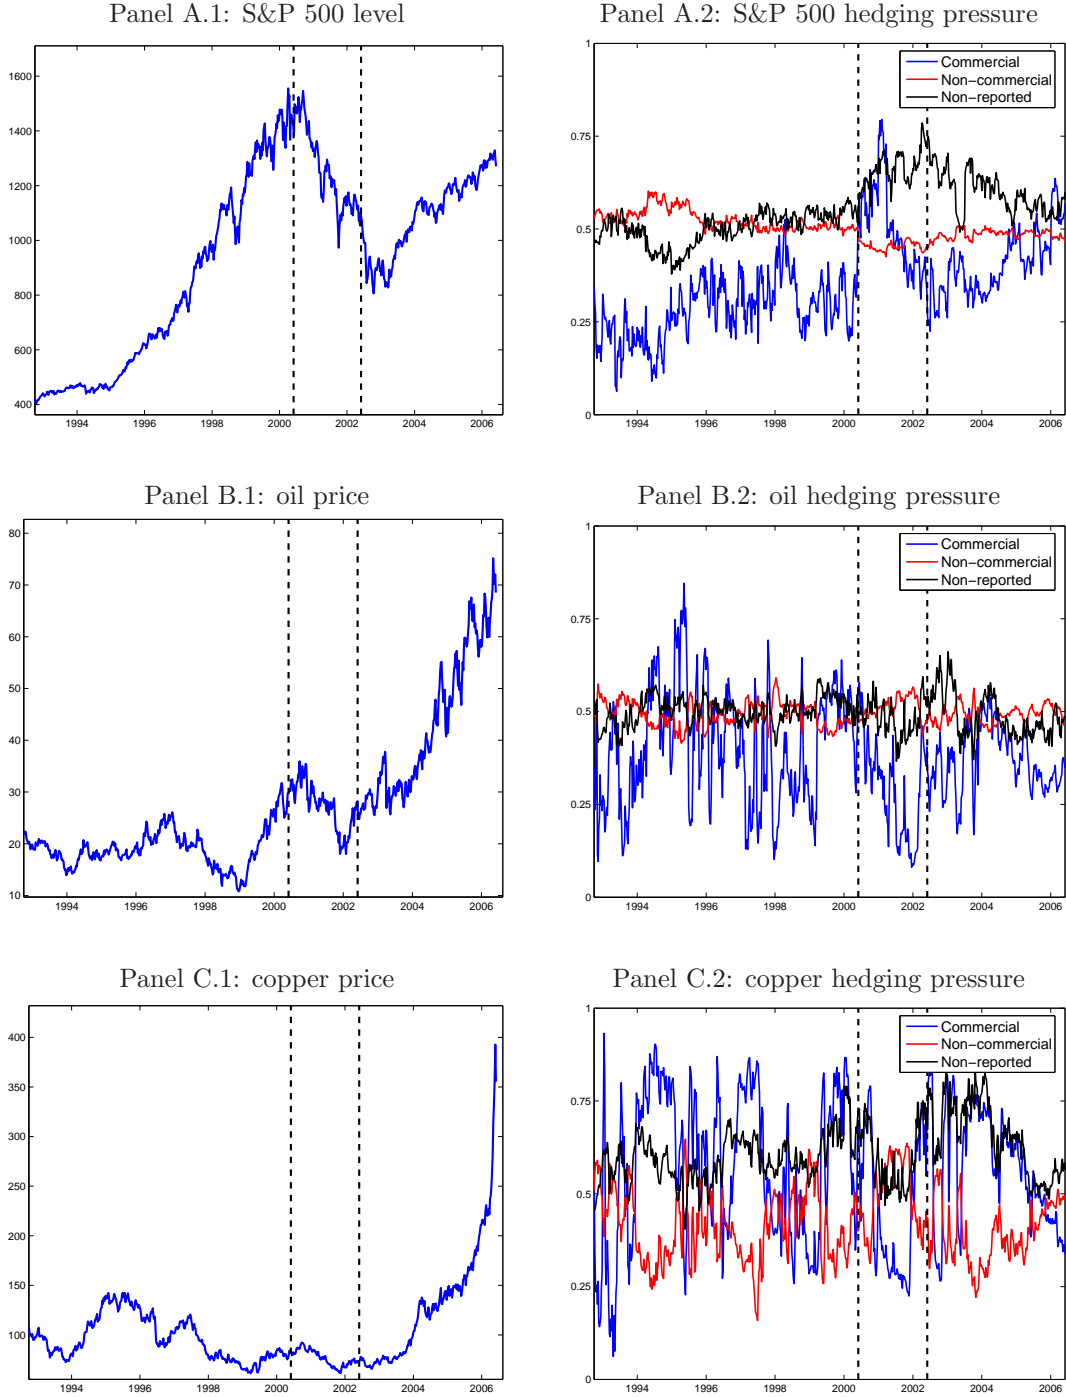
We show that the information contained in the CFTC's 'Commitment of Traders' report could have helped a long-only portfolio manager to successfully time the commodity boom. The key predictive variables are various kinds of hedging pressure which summarize the activities of various classes of market participants. Our dynamic strategy exited copper around in mid-2005 and switched completely into the risk-free asset in May 2006 thus avoiding the latest market downturn. It will be interesting to see how this strategy continues to evolve

in the coming months.

References

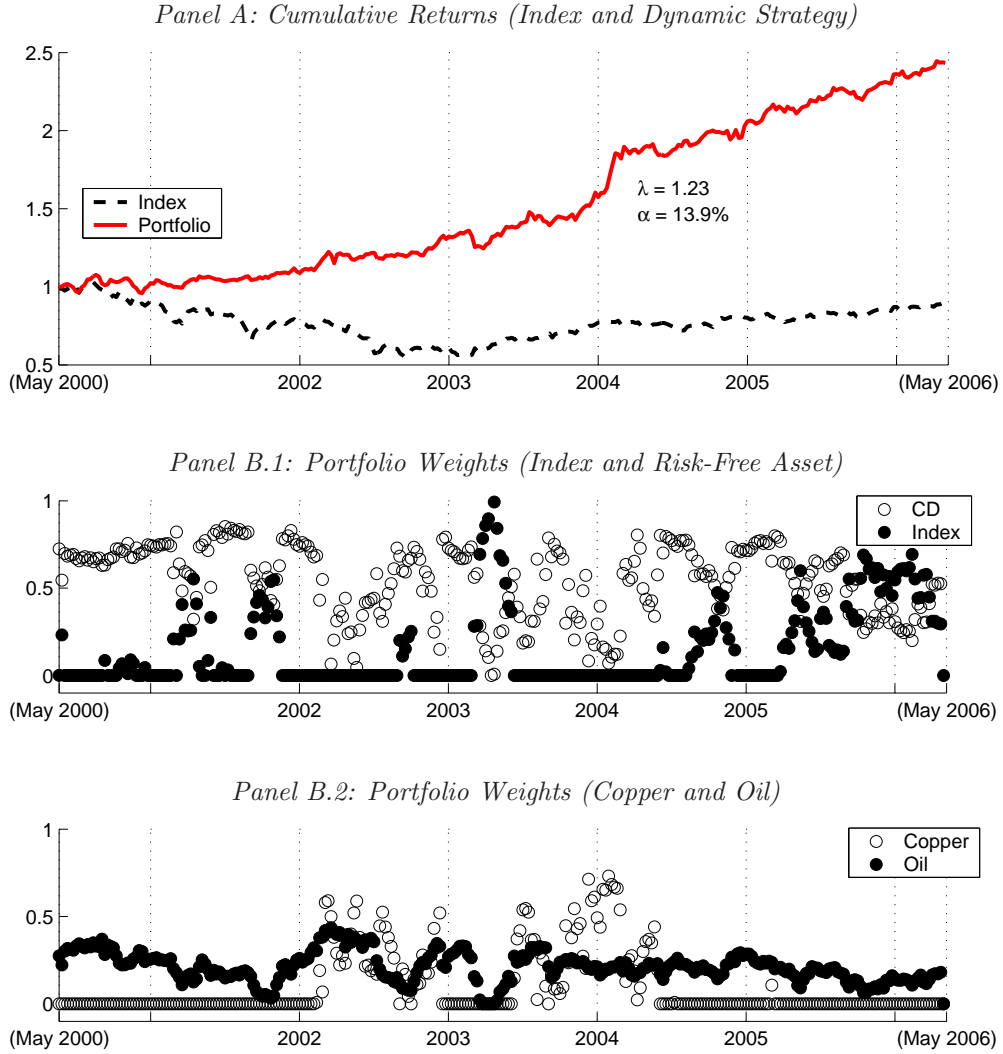
- BASU, D., R. OOMEN, and A. STREMME (2006): “Exploiting the Linkages between Spot and Derivatives Markets,” working paper, Warwick Business School.
- ERB, C., and C. HARVEY (2006): “The Strategic and Tactical Value of Commodity Futures,” *Financial Analysts Journal*, 62(2), 69–97.
- FERSON, W., and A. SIEGEL (2001): “The Efficient Use of Conditioning Information in Portfolios,” *Journal of Finance*, 56(3), 967–982.
- GORTON, G., and K. ROUWENHORST (2006): “Facts and Fantasies about Commodity Futures,” *Financial Analysts Journal*, 62(2), 47–68.
- 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.

Figure 1: Time Series of Asset Prices and Hedging Pressure




These graphs show the time series of prices and the three hedging pressure variables (non-commercial, commercial and non-reportable), for each of the three assets (S&P 500, copper and oil). The dashed lines indicate the start dates of our two out-of-sample periods (May 2000 and May 2002).

Figure 2: Out-of-Sample Performance (May 2000 – May 2006)



These graphs show the performance and portfolio weights of our dynamically managed long-only strategy during the period following the collapse of the ‘dot-com’ bubble (from May 2000 until May 2006). Panel A shows the cumulative returns of the long-only maximum-return strategy (solid line) and the market index (dashed line), normalized to have unit value in May 2000. Also reported are the annualized Sharpe ratio (λ) and CAPM alpha (α) of the strategy. Panels B.1 and B.2 show the portfolio weights of the strategy on the risk-free asset and the S&P (B.1), and on copper and oil (B.2). The predictive variables are commercial and non-reportable hedging pressure for the S&P, and non-commercial hedging pressure for copper and oil.

Figure 3: Out-of-Sample Performance (May 2002 – May 2006)



These graphs show the performance and portfolio weights of our dynamically managed long only strategy during the period during the bull run after the collapse of the ‘dot.com’ bubble (from May 2002 until May 2006). Panel A shows the cumulative returns of the long-only maximum-return strategy (solid line) and the market index (dashed line), normalized to have unit value in May 2002. Also reported are the annualized Sharpe ratio (λ) and CAPM alpha (α) of the strategy. Panels B.1 and B.2 show the portfolio weights of the strategy on the risk-free asset and the S&P (B.1), and on copper and oil (B.2). The predictive variables are commercial and non-reportable hedging pressure for the S&P, and non-commercial hedging pressure for copper and oil.