EERI

Economics and Econometrics Research Institute

Liquidity-Induced Dynamics in Futures Markets

Stephen Fagan and Ramazan Gencay

EERI Research Paper Series No 1/2008

EERI
Economics and Econometrics Research Institute
Avenue de Beaulieu
1160 Brussels
Belgium

Tel: +322 299 3523
Fax: +322 299 3523
www.eeri.eu

Liquidity-Induced Dynamics in Futures Markets

Stephen Fagan[*] Ramazan Gençay[†]

October 2007

Abstract

Futures contracts on the New York Mercantile Exchange are the most liquid instruments for trading crude oil, which is the world's most actively traded physical commodity. Under normal market conditions, traders can easily find counterparties for their trades, resulting in an efficient market with virtually no return predictability. Yet even this extremely liquid instrument suffers from liquidity shocks that induce periods of increased volatility and significant return predictability. This paper identifies an important and recurring cause of these shocks: the accumulation of extreme and opposing positions by the two main trader classes in the market, namely hedgers and speculators. As positions become extreme, approaching their historical limits, counterparties for trades become scarce and prices must adjust to induce trade. These liquidity-induced price adjustments are found to be driven by systematic speculative behavior and are determined to be significant.

Key Words: Liquidity, Futures Markets, Return Predictability, Volatility, Trader Positions, Directional Realized Volatility, Hedgers, Speculators, Position Bounds.
JEL No: G0, G1, C1

[*]Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada, V5A 1S6, Email: sfagan@sfu.ca.

[†]Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, British Columbia, Canada, V5A 1S6, Email: gencay@sfu.ca.

1. Introduction

Market liquidity refers to the ability to easily buy or sell a security without causing a significant change in the market price. An essential requirement of liquidity is the ample availability of counterparties who are willing to sell when others want to buy and who are willing to buy when others want to sell. Finance theory often assumes perfect market liquidity where market participants can trade any amount of a security without affecting the price. This implicitly assumes the unlimited presence of counterparties. This assumption, while false, is usually a reasonable simplification in large and active markets. There are, however, occasions where, even in large and active markets, liquidity dries up because of the scarcity of counterparties. Essentially, these are times when, for the current market price, everyone who wants to be long is already long, and everyone who wants to be short is already short. When this occurs, someone wanting to trade will be unable to do so unless the market price adjusts to induce someone to trade.

This paper presents an empirical analysis of these liquidity-induced price adjustments. The market chosen for this study is the New York Mercantile Exchange's light sweet crude oil futures market. This is the most liquid market for trading crude oil which is itself the world's most actively traded physical commodity. Because of its liquidity, this market is taken to be very efficient, with the market prices reflecting the value of crude oil. As such, the futures price is used as an international pricing benchmark and an indicator of world energy prices. Despite this overall liquidity, this paper discovers that there are fairly regular occasions where price dynamics, as investigated through returns and realized volatility, are dominated by the low-liquidity effects resulting from the scarcity of counterparties.

Trade in futures markets is dominated by two classes of traders: large speculators and large hedgers[1]. In the crude oil futures market, these two classes of traders hold, on average, 84% of the outstanding futures contracts. While futures markets are anonymous in the sense that information on who is involved in any given trade is not publicly available, the Commodity Futures Trading Commission (CFTC) provides a weekly snapshot of the holdings of these classes of traders. Under normal market conditions, hedgers and speculators provide each other with ample counterparties since they tend to take opposing positions, as indicated by the sample correlation coefficient of their net long positions of -0.96. However, the number of positions held by these two classes of traders is bounded. Hedgers use futures contracts to reduce the price risk that they face in their business activities, and so the number of contracts they hold is determined by the size of their business interests. This, in

[1] The CFTC requires traders who hold more than 350 futures contracts to report their position and to self-identify themselves as having commercial or non-commercial reasons for trading. This reporting threshold varies across time and across different commodities with the aim of capturing between 70 and 90 percent of the open interest in each market. This paper uses the reporting threshold to identify large traders, and uses the commercial/noncommercial distinction to identify hedgers and speculators. Details of this data is found in Section 2.2.

turn, bounds their positions in the futures market. To use more futures contracts than this bound would be to over-hedge and would increase the price risk they face[2]. Speculators are also constrained by factors such as available capital, trading strategies, as well as leverage and risk limits. As measured by the aggregate net long positions, the class of speculators in the crude oil market tend to hold long rather than short positions[3]. This may be due to the particular distribution of trading strategies in the market, or it may be due to the needs of hedgers that require them to offer favorable prices to induce speculators to take more long positions than they otherwise might.

Under liquid market conditions, hedgers and speculators are not trading near their limits, and so when a trader wishes to trade, it is highly likely that a counterparty will be available to trade with. When traders positions begin to approach their historical limits[4], with hedgers going extremely long and speculators going extremely short, the number of traders willing to go even longer or shorter than they currently are falls dramatically[5]. During these times, larger than normal price movements are seen, with the average realized volatility at its 79th percentile. The notion of directional realized volatility is introduced to capture the difference in contributions to volatility from periods of positive returns versus periods of negative returns. Disaggregating the realized volatility in this manner permits the identification of a significant asymmetry in the cause of high volatility before and after these liquidity shocks.

Returns and position accumulation are shown to have a strong concurrent relationship with, in the crude oil market, large speculators tending to increase their long positions during rising markets while large hedgers tend to increase their short positions. During falling markets, the opposite is observed with large speculators decreasing their long positions as hedgers go long. While this relationship holds concurrently, there is no evidence of predictability of returns from trader positions. However, periods when traders have accumu-

[2]Some hedgers do intentionally take over-hedged positions in futures markets, which are essentially speculative positions. This type of behavior reduces the behaviorial distinction between hedgers and speculators, thereby reducing the effectiveness of the type of analysis performed in this paper. In some markets, like those for equities and currencies, the fuzziness of the distinction between hedgers and speculators is very significant, while in other markets such as those for metals, energy, and agricultural commodities, the distinction is fairly sharp.

[3]Speculators appear to favor being long (agreeing to buy future commodities) and hedgers favor being short (agreeing to sell future commodities) in many futures markets, but there are a few exceptions. For example, in the cotton futures markets, speculators heavily favor being short, and hedgers favor being long.

[4]Since the true bounds of positions taken by large hedgers and speculators would be difficult to determine, historical limits are used as a proxy for this true limit. Since the market has tended to grow over time, true bounds on position sizes is likely to change. To account for this change, a moving window is used to determine historical limits. Details of the constriction of this proxy is found in Section 4.

[5]Note that with speculators strongly tending to be long, it is only when they go extremely short that liquidity-induced dynamics are seen. In fact, when speculators go extremely long, more speculators enter the market which improves liquidity. Details of the differences between these two types of events are found in Section 4.

lated extreme positions are, because of the concurrent relationship, also periods of markets trending in one direction. Since extreme trader positions are near their upper or lower bounds, they inevitably move back towards their average since they cannot become much more extreme than they currently are. This permits prediction of future trader positions, and again because of the concurrent relationship between trader positions and returns, this means that during times of liquidity shocks, returns move in a predictable manner. This temporary return predictability is verified empirically in Section 5.

Within the literature, the study of liquidity has focused mainly on equity (Hasbrouck and Sofianos, 1993; Chordia, Roll, and Subrahmanyam, 2001) and fixed income markets (Huang, Cai, and Song, 2001; Brandt and Kavajecz, 2002). Trader positions have been related to returns (Sanders et. al., 2004) and volatility (Wang, 2002) in futures markets, however, this paper extends the methodology of these studies in three important ways. First, the net position measure is transformed into an index that expresses the current position in relation to recent maximum and minimum values. This index captures the notion that the positions that trader classes can achieve are bounded. Second, rather than examining the overall influence of trader positions, this paper focuses on what happens when the position index approaches its bounds. Finally, the microstructure of these liquidity shocks are examined through the entry and exit of large traders in the market. It is found that the number of hedgers is not significantly changed around times when extreme trader positions are taken. However, speculators who exit the market because of trading loses or to take profits are significant factors in the liquidity-induced dynamics around these times. Interestingly, while these large and sophisticated speculators act in an informed manner during trending markets before and after these liquidity shocks, the majority are on the wrong side of the market at the event date.

The paper is organized as follows. Section 2 discusses the market and CFTC data. Section 3 examines the empirical regularities that differentiate hedgers from speculators. Section 4 analyzes ex post market dynamics around extreme holdings of these two classes of traders. Section 5 analyzes these same dynamics ex ante. A brief summary concludes the paper in Section 6.

2. The Market and CFTC Data

The market under study in this paper is the New York Mercantile Exchange's Light Sweet Crude Oil futures market[6]. This futures contract is the most liquid instrument for trading crude oil, which is the world's most actively traded physical commodity. As a result, the

[6]Detailed information about the crude oil futures market is available from the New York Mercantile Exchange website at www.nymex.com.

futures price of light sweet crude is used as an international pricing benchmark[7] and an indicator of world energy prices. The light sweet variety of crude oil is popular among refineries because of its low sulfur content and its high yield of gasoline, diesel, and other petroleum products. This market is chosen because of its high overall liquidity, as well as the fact that the speculator-hedger distinction is not as fuzzy as in other markets such as the S&P 500 index futures market where contracts can be used as a speculative and hedging instrument by the same fund manager.

Data from two sources are integrated for the analysis: futures market trade data[8] and Commitments of Traders Report (CoT) data from the Commodity Futures Trading Commission (CFTC). The following sections describe these datasets and the construction of the variables of interest.

2.1. Futures Market Data

Futures price data for the NYMEX light sweet crude oil market is collected for the same period for which weekly CoT data is available, namely September 30, 1992 until February 28, 2006 (698 weeks). The data is sampled at the tick level ($5,487,792$ observations) and aggregated to the scales required for variable construction. Since there is trade in several contracts, each with different expiration dates and market prices, a sensible method of constructing a single price series over this period is needed. Since the front contract - the contract with the closest expiration date - usually has the highest trade volume and open interest, the constructed price series is composed almost entirely of prices from the front contract. As the front contract approaches expiration, most traders begin to close out positions in the front contract and enter position in the first-back-month contract. On the first day that the daily number of price ticks, a proxy for the volume of trade, of the first-back-month exceeds that of the front contract, the constructed price series uses the prices from the first-back-month contract from that point onwards.

When the contract providing the 'market price' is replaced by another, a price gap is likely to occur. This price series will give accurate price level information, but price changes will be skewed by the gaps. To get accurate price change data, the series is 'back-adjusted' to remove the gaps. With this back-adjustment, prices are adjusted by determining the difference between prices on the date of rollover and adding or subtracting the difference from all previous prices. This will create a series where price levels are actual for the most recent period but synthetic prior to the date of the last rollover. If this process would create negative prices at any point in the series, a constant is added to all prices in the series to ensure that all prices in the series are positive.

[7]Officially, there are only three crude oil benchmarks: WTI, Brent Blend, and Dubai. WTI, standing for West Texas Intermediate, is used primarily in the U.S. and is the underlying asset for the NYMEX futures contracts.

[8]The data provider for the crude oil futures market tick data is TickData, www.tickdata.com.

The notion of the return on a futures contract is not clear since the initial value of a contract is zero. In practice, holders of futures contracts must commit some of their wealth to a margin account as a 'performance bond' to ensure their ability to pay any amounts owing to their counterparty on a daily basis. At a minimum, the trader must commit the exchange-mandated margin amounts[9] to this account. If any amount less than the futures price is held in this account, then the futures exchange could issue a 'margin call,' requiring the addition of more funds to the margin account, in the event that the value of the margin account falls below a specified value. In order to avoid these complications, this analysis assumes that an amount of cash equal to the futures price is kept in a non-interest paying margin account[10]. To get an accurate measure of return, the use of both the unadjusted price series, P_t, and the back-adjusted series, P_t^b, are needed. The unadjusted price series gives accurate price level information while the adjusted series gives accurate price change information. The return on a futures contract, r_t, between times $t-1$ and t is defined by

$$r_t = \frac{P_t^b - P_{t-1}^b}{P_{t-1}}. \tag{1}$$

The returns analyzed in this paper are weekly returns, calculated from prices on the close of each Tuesday - the 'as of' date of the CoT report which is released three days later on Friday.

In addition to returns, realized volatility (Dacorogna et. al., 2001; Andersen and Bollerslev, 1998) is analyzed as a measure of the true volatility of this return series. Realized volatility is sampled on a weekly basis and constructed from 5-minute intraday data. Let there be h five-minute intervals during which the market is open in one week (currently $h = 270$)[11]. Then realized volatility is calculated as

$$RV_t = \sum_{i=1}^{h} \left(r^{(h)}_{t-1+i(\frac{1}{h})} \right)^2 \tag{2}$$

where $r_t^{(h)}$ is the five-minute return at time t. Figure 1 presents plots of the unadjusted price series, the weekly returns, and the realized volatility.

[9]As of October 24, the initial margin for a contract of 1000 barrels of light sweet crude oil was 5400 USD. With a futures price of $87.65 per barrel, this minimum margin requirement was only 6% of the value of the underlying oil at the futures price.

[10]It is usually possible to keep US Treasury securities instead of cash in a futures margin account, thereby allowing the margin to earn a risk-free interest rate. The assumption of non-interest paying margin accounts is made for simplicity and is not significant since we would be interested in returns in excess of a risk-free rate in either case.

[11]The current day trading hours for NYMEX crude oil futures is from 10:00 until 14:30 EST. Thus, one trading day has 54 five-minute intervals, and one trading week has 270 five-minute intervals.

2.2. CFTC Trader Data

The Commodity Futures Trading Commission (CFTC) has been releasing weekly Commitments of Traders Reports[12] since September 30, 1992, and at various lower frequencies since 1924[13]. These reports provide a breakdown of the open interest in the American futures markets as of each Tuesday. Open interest refers to the total number of futures contracts that have been entered into and not yet exited through a transaction or delivery. Total open interest is broken down according to the type of trader and by the type of positions (long or short) that each type of trader holds. There are three types of traders identified by the CFTC: reporting commercial, reporting non-commercial, and non-reporting. A trader is classified as a reporting trader if they hold positions in excess of a threshold[14] established by the CFTC on a market-by-market basis. A commercial trader is one who self-identifies as being "engaged in business activities hedged by use of the futures and options markets."[15] The non-reportable positions are all positions that are held by traders whose total position size is below the CFTC's threshold. In what follows, reporting commercial traders are considered large hedgers, reporting non-commercial traders are considered large speculators, and non-reporting traders are considered small traders who may be hedging or speculating.

Thus, the disaggregation of open interest performed by the CFTC occurs along three dimensions: reporting versus non-reporting, hedging versus speculating, and long versus short (and versus spreads[16] for non-commercials). The seven components of the total open

[12]The Commitments of Traders Reports are published every Friday on the CFTC website at http://www.cftc.gov/cftc/cftccotreports.htm.

[13]Early predecessors of the CoT reports were prepared by the U.S. Department of Agriculture's Grain Futures Administration and were released yearly. As of June 30, 1962, CoT data began to be released every month.

[14]The current CFTC threshold for the crude oil futures market is 350 contracts. The threshold will vary across time and markets in an effort to capture between 70 and 90 percent of the open interest in the market. One of the main reasons for collecting this data is to detect and deter attempts at market manipulation by large traders. It is illegal to manipulate the price of any commodity in interstate commerce (eg. cornering a market), and the CoT reporting system allows the CFTC staff to identify large positions that could pose a threat to orderly trading.

[15]From CFTC Form 40, Statement of Reporting Trader, page 7. The reliability of these classifications is not perfect and varies from market to market. There are times when hedgers take speculative positions, and when speculators hedge their positions. In some markets, such as equity index futures markets, the reliability is quite questionable, and other, such as the crude oil futures market, where the classification is expected to be more reliable.

[16]The spread number in the COT report measures the extent to which a speculator holds equal long and short positions. For example, if a speculator was long 100 contracts and short 70 contracts, then the spread number will be 70 and the long number will be 30.

interest (TOI) that are detailed in the CoT reports are related in the following manner:

$$\overbrace{\overbrace{[S_{Long} + S_{Short} + 2(S_{Spread})]}^{speculators} + \overbrace{[H_{Long} + H_{Short}]}^{hedgers}}^{reporting} + \overbrace{[NR_{Long} + NR_{Short}]}^{nonreporting} = 2(TOI) \quad (3)$$

where S, H and NR refer to positions held by large speculators, large hedgers and non-reporting traders, respectively, with subscripts (long, short, and spread) indicating the type of positions. On the left side of equation (3), each contract is counted twice since both long and short positions are counted even though a long and a short position constitute only one contract. Consequently, the total open interest is doubled on the right side of the equation.

As can be seen from the sample CoT report[17] in Figure 2, the report also lists the number of traders in each category. Note that the sum of the numbers of traders from each category exceeds the total number of traders due to the fact that a single trader, such as a spreading speculator, can be counted in more than one category[18]. Also of note is that the CoT is available in two versions: a futures-only version, and a futures and options version that converts options positions into equivalent number of futures positions. Since the current study is interested in the relationship between trader positions and futures market dynamics, the futures-only version is used.

Various measures of trader positions are constructed to illuminate the relationships between trader classes and market dynamics. These measures are defined as they are introduced in the next three sections.

3. Hedgers and Speculators

The large majority of open interest is held by large hedgers and speculators. In the sample under study, an average of 84% of open interest is held by large traders that must report their holdings to the CFTC. Of the reportable holdings, large hedgers dominate the large speculators by holding, on average, 68% of the total open interest compared with 16% for the speculators. Equation (2) demonstrates the calculation of the percent of open interest (POI) held by large speculators. The POI for hedgers[19] is calculated similarly, but without the spread component.

$$POI_t^S = \frac{S_{Long,t} + S_{Short,t} + 2(S_{Spread,t})}{2(TOI_t)} \quad (4)$$

[17] See CFTC (2007) for a history of the CoT reports and a full description of CoT variables.

[18] A trader is counted in each category in which the trader holds a position. For example, a hedger who is long December 2008 crude oil while also being short June 2009 crude oil will be counted only once for the number of total traders, but will be counted in both the 'commercial long' and the 'commercial short' categories.

[19] $POI_t^H = \frac{H_{Long,t} + H_{Short,t}}{2(TOI_t)}$

where S refers to positions held by large speculators with the subscripts (long, short, and spread) indicating the type of positions. From the time series of POI presented in Figure 3, the larger size of the positions of hedgers relative to speculators is clear - hedgers hold between 60% and 80% of the outstanding positions, speculators hold between 5% and 35% of the outstanding positions, and small traders hold the remaining positions.

Another feature of POI evident in Figure 3 is the inverse relationship between movements in positions of hedgers and speculators. Within the sample, the percent of open interest held by large speculators and hedgers, POI_t^S and POI_t^H, have a correlation coefficient of -0.40, indicating that when one trader is entering a position, there is a tendency for a trader of another class to be closing a position. This observation is consistent with the proposition that different trader types provide each other with liquidity.

The most popular variable in the analysis of trader positions is the net long measure, denoted NL_t^C for trader class C. The net long measures are calculated by

$$NL_t^S = S_{Long,t} - S_{Short,t} \quad \text{and} \quad NL_t^H = H_{Long,t} - H_{Short,t}. \tag{5}$$

Summary sample statistics for the net long positions by trader types are presented in Table 1. The positive mean, median, and skewness of the net long positions of speculators indicates a tendency of speculators to go long, while hedgers tend to go short. This observation reinforces the proposition that multiple trader types increase market liquidity since if speculators were not in the market, hedgers would have a harder time finding counterparties for their short positions. A second observation can be made from the sample kurtosis, both of which indicate thinner than Gaussian tails. If trader positions are bounded due to some budgetary or risk constraints, the distribution of their positions would be expected to be finite or, at most, thin.

As can be seen in Figure 4, the net long position of large hedgers and speculators tend to move in opposing directions. The strong sample correlation of -0.96 is not unexpected since a net change for one trader type requires another trader type to take an opposing position. What is interesting, however, is that the relationship between small traders and large traders is much weaker, with $\rho(NL^S, NL^{small}) = 0.543$ and $\rho(NL^H, NL^{small}) = -0.755$. Since the small trader class is a mix of small speculators and hedgers, the behaviorial regularities associated with trader classes are weaker, though the signs of the correlations indicate that a position taken by small traders is more likely a speculation than a hedge.

3.1. Net Long Positions and Market Returns

Before studying at the dynamics induced by traders taking extreme positions, it is useful to first examine the relationships between the market variables and the entire distribution of positions. There is a positive concurrent correlation of 0.23 between market returns and speculative net long positions. Given the strong negative correlation of -0.96 between the net long positions of hedgers and those of large speculators, it is not surprising to find a negative correlation of -0.24 between returns and hedger net long positions. Thus, large speculators are net buyers in rising markets, while large hedgers are net sellers.

The next step is to consider the lead-lag relationship between positions and returns. The relevant questions here are 'Do returns lead trader positions?' and 'Do trader positions lead returns?' A useful way to begin this analysis is to fit a vector autoregressive model to the data and performing Wald tests of Granger causality. For the estimation, two two-dimensional time series

$$X_t = (x_{1t}, x_{2t}) = (r_t, NL_t^S) \quad \text{and} \quad Z_t = (z_{1t}, z_{2t}) = (r_t, NL_t^H) \tag{6}$$

will be fitted to two separate l-lag vector autoregressive (VAR(l)) models. This is done, rather than using a single three-dimensional time series, in order to reduce the high levels of multicollinearity between NL_t^S and NL_t^H, and because of their highly negative correlation of -0.97 model misspecification should not be very significant.

The VAR(l) model has the following specification

$$X_t = \varphi_0 + \Phi_1 X_{t-1} + \cdots + \Phi_l X_{t-l} + \epsilon_t \tag{7}$$

where ϵ_t is a zero mean white noise bivariate process with a time invariant covariance matrix, which is equivalent to the following two equation system

$$r_t = \varphi_{1,0} + \Sigma_{i=1}^{l} \varphi_{1,i} r_{t-i} + \Sigma_{i=1}^{l} \varphi_{1,j} NL_{t-j}^S + \epsilon_{1,t} \tag{8}$$

$$NL_t^S = \varphi_{2,0} + \Sigma_{i=1}^{l} \varphi_{2,i} r_{t-i} + \Sigma_{i=1}^{l} \varphi_{2,j} NL_{t-j}^S + \epsilon_{2,t}. \tag{9}$$

The number of lags is chosen by trying several lag lengths and choosing the model that minimizes the Schwarz-Bayesian information criterion (BIC). For the case of both $X_t = (r_t, NL_t^S)$ and $Z_t = (r_t, NL_t^H)$ a lag length of 1 is selected. Table 2 contains the results of the least squares estimation of the VAR(1) model for $X_t = (r_t, NL_t^S)$. The fit for the return model is quite poor indicating lack of predictability, whereas the fit for the NL^S model is significantly better with highly significant coefficients for both lagged returns and positions sizes. These observations can be confirmed with formal tests for Granger non-causality[20]. The p-value of the test of the null hypothesis that NL^S fails to Granger cause r is 0.1024, and so the null is not rejected. On the other hand, the p-value of the test of the null hypothesis that r fails to Granger cause NL^S is 0.0016, and so the null is strongly rejected.

[20] Granger causality does not imply true causality, but rather implies forecasting ability. A variable y_1 is said to Granger cause y_2 provided that y_1 is found to be useful in predicting y_2; otherwise y_1 fails to Granger cause y_2. In the current one lag model, y_1 fails to Granger cause y_2 provided that the mean square error of a forecast of $y_{2,t}$ based on $y_{2,t-1}$ is the same as the mean square error of a forecast of $y_{2,t}$ based on $y_{2,t-1}$ and $y_{1,t-1}$. Granger non-causality hypotheses are easily tested using the Wald statistic.

Given the strong correlation between, NL^S and NL^H, the least squares estimation of the VAR(1) model for $Z_t = (r_t, NL_t^H)$ is expected to be similar to those for $X_t = (r_t, NL_t^S)$ given above. Table 3 reports these estimation results. As expected, the null hypothesis that NL^H fails to Granger cause r is not rejected, with a p-value of 0.2424. The p-value of the test of the null hypothesis that r fails to Granger cause NL^H is 0.0175, and so the null is rejected.

As economic theory suggests, return predictibility should not be easy to find. This has been confirmed in this section, with a simple VAR(1) model demonstrating no predictive ability from trader positions. Evidence is found, however, that differentiates the class of large hedgers from that of large speculators. Speculators are net buyers in rising markets, and hedgers are net sellers. As will be shown in Section 4.1, this concurrent relationship between the accumulation of trader positions and market returns will induce return predictability when positions become extreme and must revert back towards more moderate levels.

4. Extreme Positions and Market Dynamics - An Ex Post Analysis

When speculators and hedgers acquire positions that are in extreme opposition to each other, that is when one group is extremely long while the other is extremely short, the capacity for these trader groups to continue adding to these positions is limited by the availability of counterparties. Since speculators tend to take long positions and hedgers tend to take short positions, there is a particular shortage of counterparties when the positions of speculators become extremely short and those of hedgers become extremely long. At such times, speculators who are willing to take long positions have either already taken them or have left the market, thereby causing the availability of counterparties to fall, and with it, liquidity dries up. These liquidity shocks are studied in the remainder of the paper.

To study the market dynamics as trader positions approach their bounds, an event study methodology is used. The first step in this methodology is to define the extreme event that is being studied. To do this, the notion of trader positions, as expressed by the net long measure described in the previous section, must be extended to capture the notion that the positions that trader classes can achieve are bounded. Since the market under study has been growing over the sample period, an absolute value of the position bounds for all times is inappropriate. Instead, a local measure, $P_t^C(\tau)$ is employed which uses the maximum and minimum NL_t^C values achieved by trader class C within a moving window of τ periods. In this section, a centered window is used, thereby using information from the recent past and future. Consequently, in this ex post study, we can study return regularities, but no claim of predictability can be made from the results. In the next section, an ex ante analysis is presented in which only past information is used in the construction of variables, allowing a discussion of predictability. The advantage of an ex post analysis is that a better proxy for position bounds is permitted. That is, extreme positions are easier to identify when future information is used, whereas in the ex ante study, a large position may get labeled as extreme and then be dwarfed by an even larger position a few periods later. This trader position measure is expressed as

$$P_t^C(\tau) = \frac{NL_t^C - \min\{NL_{t'}^C\}}{\max\{NL_{t'}^C\} - \min\{NL_{t'}^C\}} \tag{10}$$

where $t' \in \{t-(\tau-1)/2, \cdots, t-1, t, t+1, \cdots t+(\tau-1)/2\}$ for an odd-valued τ. In an ex ante form that is presented in Section 5, using only past information, this variable is similar to the market sentiment indicator of Briese (1990). For the current study, τ is taken to be around nine months (39 weeks). Clearly, values of $P_t^C(\tau)$ take values between 0 and 1, with values achieving 1 when net long positions exceed all values within the nine month centered moving window. Figure 5 plots this position index for large speculators and hedgers.

The events of interest occur when one trader class is extremely long and the other class is extremely short. To identify the times of these events, two indicator variables are constructed. $SLHS_t$ identifies times when speculators are extremely long and hedgers are extremely short, and $HLSS_t$ identifies times when speculators are extremely short and hedgers are extremely long. These variables are defined to be

$$SLHS_t = \begin{cases} 1 & \text{if} \quad P_t^S(\tau) - P_t^H(\tau) = \max[P_{t'}^S(\tau) - P_{t'}^H(\tau) : t' \in [t - \frac{(\tau-1)}{2}, t + \frac{(\tau-1)}{2}]] \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

$$HLSS_t = \begin{cases} 1 & \text{if} \quad P_t^S(\tau) - P_t^H(\tau) = \min[P_{t'}^S(\tau) - P_{t'}^H(\tau) : t' \in [t - \frac{(\tau-1)}{2}, t + \frac{(\tau-1)}{2}]] \\ 0 & \text{otherwise} \end{cases} \tag{12}$$

An event of type $SLHS$ occurs when the difference between the position indices of speculators and hedgers is larger than that difference at any other time within a centered window of length τ. For example, if $P_t^S(\tau) = 1$ and $P_t^H(\tau) = 0$, then speculators are extremely long and hedgers are extremely short, and so $SLHS_t = 1$ captures the extreme opposition of the positions of these two classes of traders. Within the sample period there are 20 events of each type.

4.1. Extreme Positions and Returns

It was shown in Section 3 that a strong concurrent relationship exists between trader positions and market returns. With rising markets, speculators tend to be net buyers and hedgers tend to be net sellers. With falling markets, speculators tend to be net sellers and hedgers tend to be net buyers. This relationship implies that extreme events are often preceded by periods of trending prices. Once trader positions are near their upper or lower bounds, they inevitably move back towards their average since they cannot become much more extreme than they currently are. When trader positions revert away from their extremes, the price trend often ends or even reverses.

In the case of $SLHS$-type events, where speculators are extremely long and hedgers extremely short, there is a significant increase in prices that begins around seven weeks before the extreme event resulting, on average, in a 8.8% cumulative abnormal return over that period, where the abnormal return is defined to be the return minus the sample mean return of 0.36%. After the event, the upward trend is reversed and prices tend to fall for the next eight weeks. Thus, $SLHS$-type events tend to coincide with local price maximums. In the case of $HLSS$-type events, where speculators are extremely short and hedgers extremely long, there is a significant decrease in prices that begins around six weeks before the extreme event resulting, on average, in a -9.3% cumulative abnormal return over that period. The price trend reverses after the event date and tends to rise for the next seven weeks. Thus, $HLSS$-type events tend to coincide with local price minimums. Both of these cases are seen in Figure 6.

To evaluate the effect of trader positions on returns for the week before and after those positions become extreme, the following model is estimated

$$r_t = \alpha + \beta_1 SLHS_t + \beta_2 SLHS_{t-1} + \gamma_1 HLSS_t + \gamma_2 HLSS_{t-1} + \varepsilon_t \tag{13}$$

where ε_t is assumed to be a white noise process, $SLHS_t$ is an indicator variable for times when large speculators are extremely long and large hedgers are extremely short, and $HLSS_2$ is an indicator variable for times when large speculators are extremely short and large hedgers are extremely long. The result of this estimation is found in Table 4. All of the estimated variable coefficients except for γ_1 are significant. Since the return, r_t, measures the price change between times t and $t-1$, β_1 is the change in price leading up to an $SLHS$-type event occurring at time t. Similarly, β_2 is the change in price that occurs when the $SLHS$-event occurred at time $t-1$, that is, the change in price after such an event. Estimates of β_1 and β_2 are positive and negative, respectively, indicating that prices tend to rise during the week before an $SLHS$-type event and fall the following week. Conversely, estimates of γ_1 and γ_2 are negative and positive, respectively, indicating that prices tend to fall during the week before an $HLSS$-type event and rise the following week. These results are consistent with the price trends and reversals around extreme events seen in Figure 6.

4.2. Extreme Positions and Realized Volatility

In addition to return regularities, liquidity-induced dynamics should also have an impact on the size of the returns. To examine this conjecture, this section examines the behavior of the volatility of market returns, as measured by the realized volatility, around times of extreme and opposing trader positions. Realized volatilities are sampled weekly and constructed with five-minute returns, as described in Section 2.1. Figure 7 plots the average realized volatility around the event dates. There is a clear difference between the two classes of extreme events. The average realized volatility before and immediately after type-$SLHS$ extreme positions, where speculators are long and hedgers are short, is significantly lower than the overall sample average[21]. In contrast, it is significantly higher before and just after type-$HLSS$ extreme positions, where speculators are short and hedgers are long. In this case, the average realized volatility on the event date is at the 79th percentile of the entire sample and is even higher for several weeks before the event date.

To further evaluate the effect of these extreme events on return volatility, the following model is estimated:

$$\ln(RV_t) = \alpha + \beta \ln(RV_{t-1}) + \gamma HLSS_t^* + \varepsilon_t \tag{14}$$

where RV_t is the realized volatility described in Section 2.1, $HLSS_t^*$ is an indicator variable for whether a $HLSS$-type extreme event occurs within the next five weeks, and ε_t is assumed to be a white noise process. This new indicator is defined more formally as

$$HLSS_t^* = \begin{cases} 1 & \text{if} \quad HLSS_t \vee HLSS_{t+1} \vee HLSS_{t+2} \vee HLSS_{t+3} \vee HLSS_{t+4} \vee HLSS_{t+5} \\ 0 & \text{otherwise.} \end{cases} \tag{15}$$

The results of this estimation are found in Table 5. The estimated coefficient of $HLSS_t^*$ indicates that a 6.6% increase in the realized volatility is expected in each of the five weeks before the event date, indicating that $HLSS$-type extreme events can be economically important causes of high-volatility periods.

Realized volatility is a measure of the size of price movements, but it says nothing of the direction of those movements. Yet, the ups and downs of prices might not equally contribute to the return volatility in any given period. During a swiftly rising market, one might expect that the numerous buyers would make selling easy, but buying may be difficult in the sense that it would induce significant slippage. Similarly, during a market crash, selling would be difficult, but buying would be easy. To investigate the notion that liquidity shocks may be asymmetric in this sense, the concept of directional realized volatility is introduced. Directional realized volatility disaggregates the realized volatility into a component that measures the contribution of positive returns to the realized volatility, denoted DRV_t^P, and another component that measures the contribution of negative returns, denoted DRV_t^N. The two components, which are additive variance-preserving transformations of realized volatility, are defined as follows:

$$DRV_t^P = \sum_{i=1}^{h} \left[I\left(r_{t-1+i(1/h)}^{(h)} \geq 0 \right) \cdot r_{t-1+i(1/h)}^{(h)} \right]^2 \tag{16}$$

[21] For Figure 7 and the following estimations, a single outlier realized volatility value was removed from the series. The outlier was on 24 March 1998 and was over 19 standard deviations away from the sample mean. There was no large price adjustment on that week, so the value was replaced with an average of the preceding and following values.

$$DRV_t^N = \sum_{i=1}^{h} \left[I \left(r_{t-1+i(1/h)}^{(h)} < 0 \right) \cdot r_{t-1+i(1/h)}^{(h)} \right]^2 \tag{17}$$

where $r_{t-1+i(1/h)}^{(h)}$ is the five-minute return as defined in Section 2.1 and $I(\cdot)$ is an indicator function taking the value 1 if the condition in the brackets is satisfied, and 0 otherwise. The positive (negative) directional realized volatility sums the squared five-minute returns that are positive (negative).

Figure 8 plots the average positive and negative directional realized volatilities for the weeks around $HLSS$-type events, when hedgers are extremely long and speculators extremely short. It is clear that the realized volatilities before and after the event date are composed of two nonsymmetric components. Before the event, the majority of the volatility comes from negative returns as is seen from the increase in DRV_t^N before the event date in Figure 8(a). After the event date, DRV_t^N falls sharply, but DRV_t^P, in Figure 8(b), peaks thereby keeping the overall realized volatility higher than average. This asymmetric composition of return volatility indicates that potential ill-effects of liquidity shocks do not affect all market participants equally. High negative directional volatility would be bad for sellers, while high positive directional volatility would be bad for buyers. Further, a significant difference between the negative and positive volatilities indicates a trending prices.

4.3. Extreme Positions and the Number of Traders

So far, little has been said of the composition of the groups of large hedgers and speculators. The Commitments of Traders Reports list the number of large traders in each of these groups, thereby allowing an analysis of their composition around the times that they acquire extreme positions. For the class of $SLHS$-type events, where speculators are extremely short and hedgers extremely long, the left panel of Figure 9 shows that, on average, traders are entering the market before the event date and then leave the market afterwards. The right panel shows that traders are leaving the market before and slightly after an $HLSS$-type event, and then quickly begin entering again, raising the number of traders back to its pre-event level. This observation partially explains why low volatility is associated with $SLHS$-type events, and high volatility is associated with $HLSS$-type events. The number of traders in a market has long been viewed as a proxy for liquidity since fewer traders reduces the availability of counterparties. In other words, when there are many traders in a market, the wide distribution opinions, hedging requirements, available capital, and other factors means that finding someone who wants to take the opposite position that you want is highly likely.

Two questions arise from the observation that traders leave or enter the market before and after extreme positions are taken: Who is it that is leaving the market, and why are they leaving? A partial answer to the first question is seen in Figure 10 which shows that the strong majority of the traders that systematically leave and enter the the market are speculators, and not hedgers. The fact that hedgers do not enter or exit the market in a regular fashion around these extreme position events is due to the fact that their positions are not dependent on the profitability of their trades, but are rather determined by their business interests.

The number of speculators can be further disaggregated into the number of speculators who are long and the number that are short. Figure 11 shows these numbers, thereby giving a fuller answer to the question of who is leaving the market and why they are leaving. The right panel of the figure indicates that virtually all of the traders that are leaving the market before the $HLSS$-type events are speculators that were long. Recalling the relationship between positions and returns seen in Figure 6, these traders were long in a falling market, so it is likely that they were leaving the market to avoid further loses. At the same time leading up to the event date, the number of speculators that are short is actually increasing. These traders who are selling short in a falling market are realizing significant profits, and immediately after the event date they begin to take their profits by closing their short positions. This is seen in Figure 11 where the number of speculators who are short after the $HLSS$-type event dates steadily falls. Also after the event date, the traders who were long and had left the market begin to re-enter the market establishing new long positions. Thus, there seems to be regularities in the entry and exit of traders around the event dates that are consistent with speculative profit-taking and with capital constraints forcing traders to exit losing trades. A similar argument can be made to explain the regularities around $SLHS$-type events in the left panel of Figure 11.

Figure 11 can also explain the difference between the volatility levels around $SLHS$ and $HLSS$-type events. In both cases, while the changes in the number of speculators who are short moves in opposition to the changes in the number of speculators who are long, the magnitude of the later changes are more than twice as large. Thus, the movements in and out of the market by the speculators who are long dominate the overall change in the number of traders in the market. Around $HLSS$-type events, the number of long speculators is low thereby causing the overall number of traders to be low. This, in turn, reduces liquidity and causes an increase in volatility. In contrast, around $SLHS$-type events, the number of long speculators increases thereby increasing the overall number of traders. Here liquidity increases and volatility falls.

To check whether these regularities in the entry and exit of speculators plays a role in the microstructure of returns and volatilities, the following model is estimated,

$$r_t = \alpha + \beta_1 \triangle T^{spec}_{long,t} + \beta_2 \triangle T^{spec}_{short,t} + \varepsilon_t \tag{18}$$

where $\triangle T^{spec}_{long,t}$ is the change since the previous period in the number of speculators who are long, $\triangle T^{spec}_{short,t}$ is the change in the number of speculators who are short, and ε_t is a white noise process. The estimation results in Table 6 indicate that an increase in the number of long speculators is associated with a significantly higher return.

The effect of the change in the number of speculators on realized volatility is analyzed via the following model,

$$RV_t = \alpha + \beta RV_{t-1} + \delta_1 \bigtriangleup T^{spec}_{long,t} + \delta_2 \bigtriangleup T^{spec}_{short,t} + \varepsilon_t \tag{19}$$

The estimation results in Table 7 indicate that increasing the number of long speculators decreases the expected realized volatility. In the estimation of models (18) and (19), it is the change in the number of speculators that are long which are significant. The estimated coefficient of $\bigtriangleup T^{spec}_{short,t}$ has the expected sign, with more short speculators associated with lower returns and lower volatility, but they are not significant.

4.4. Informed and Uninformed Traders

If a speculative trader were aware of the date of a significant reversal of price movements, then they would take positions that would profit from this price action. An informed trader[22] would therefore want to be long on a $HLSS$-event date and short on a $SLHS$-event date. However, these events are essentially defined as times when large speculators are on the wrong side of the market. On average, at $SLHS$-event dates, only 36% of large speculators are short, and only 38% are long on $HLSS$-event dates. Thus, an essential feature of these liquidity shocks is that a significant number of traders be uninformed of their timing.

Figure 11 does show that many speculators act informed with respect to the price trends both before and after the event dates. The longer the trend has been in effect, the more speculators trade in that direction. However, some traders enter the trend too late as is indicated by the spike at time zero. Since only aggregate trader numbers are reported in the CoT reports, it is impossible to tell when a specific trader both enters and exits the market. Consequently, the distribution of profits within a class of traders is unknown.

With the majority of large speculators being on the wrong side of the market at $SLHS$ and $HLSS$-event dates, some conclusions about uninformed traders can be drawn. First, it would be inappropriate to model these uninformed traders as being equally likely to take long or short positions, since that would imply that on average the majority of speculators would be on the right side of the market. Second, uninformed trade appears to be dependent on past trades since they appear to want to follow a trend that is about to end. Finally, there appears to be degrees to which a trader is informed with traders entering a trend throughout its duration.

[22] Since a hedger's trading behavior is determined by business interests and risk management strategies, their actions would be independent of knowledge of these event dates. In this sense, hedgers act as if they are uninformed traders.

5. Extreme Positions and Returns - An Ex Ante Analysis

In the previous section, future information was employed in determining what was to be called an extreme event. As a result, while return regularities were identified, nothing could be said of return predictability. In this section, the dynamics of returns are examined around extreme events, similar to the analysis on Section 4.2, however only past information will be used to identify extreme trader positions.

Backward looking windows of six months (26 weeks) were used to construct position indices for both speculators and hedgers. As in Section 4, if the difference between the speculator index and the hedger index is at a six month maximum, then an $SLHS$-type event is said to have occurred. In this case, speculators are extremely long and hedgers extremely short. To identify times when speculators are extremely short and hedgers are extremely long, an $HLSS$-type event is said to have occurred if the difference in position indices is a six month minimum. Without the use of future information, it is impossible to know that a current maximum will not be exceeded in the next period. Despite this, Table 8 indicates that even without future information, both $SLHS$-type and $HLSS$-type events are significant as both explanatory and predictive variables.

As with the ex post case, this predictability is cause by the mean reversion of trader positions away from their extremes. In the ex post analysis, mean reversion was a certainty after an extreme event. In the ex ante case, mean reversion is not certain, but the probability that the net positions will revert towards their means exceeds the probability that they will become even more extreme. For speculators, their net positions revert towards their mean within four weeks following an extreme event 83% of the time[23]. For hedgers, mean reversion occurs within four weeks 77% of the time. With this expected mean reversion and the established concurrent relationship between trader positions and returns, return predictability is possible.

Figure 12 plots the averaged abnormal return for this ex ante case. Both $SLHS$ and $HLSS$-type events do not significantly influence the average cumulative mean returns until five weeks after the event dates.

6. Conclusions

This paper has shown that extreme trader positions do indeed cause liquidity shocks, even in otherwise liquid markets. The fact that high volatility is associated with increased return

[23]This percentage is calculated from the formula

$$\frac{1}{40}\left[\sum_t I\left[NL_t^S - NL_{t+4}^S > 0\right] \cdot I\left[SLHS_t\right] + \sum_t I\left[NL_t^S - NL_{t+4}^S < 0\right] \cdot I\left[HLSS_t\right]\right]$$

since there are 40 extreme events in the entire sample.

predictability in equity markets has been demonstrated by LeBaron (1992). This paper has found the same regularity in futures markets, and has further identified the underlying mechanism that induces some of this regularity. As speculators and hedgers accumulate extreme and opposing positions, the availability of counterparties falls inducing a fall in liquidity. However, differences in the behavior of speculators in rising versus falling markets leads to increased volatility only when speculators are extremely net short. Since returns and position accumulation are shown to have a strong concurrent relationship, when extreme trader positions move back towards their average as they inevitably do, returns move in a predictable manner.

References

[1] Bradley, J. (1968). *Distribution-Free Statistical Tests*. Prentice Hall, New Jersey.

[2] Brandt, B. and K. Kavajecz (2002). Price Discovery in the U.S. Treasury Market: The Impact of Orderflow and Liquidity on the Yield Curve. *University of Pennsylvania Working Paper*, Philadelphia.

[3] Briese, S.E. (1990). Commitments of Traders as a Sentiment Indicator. *Technical Analysis of Stocks and Commodities*, Vol. 8, No. 5, May, 199-204.

[4] Campbell, J., A. Lo, and A. C. MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton University Press, Princeton.

[5] Chordia, T., R. Roll and A. Subrahmanyam (2001). Market Liquidity and Trading Activity. *Journal of Finance*, 56, 501-530.

[6] Chordia, T., A. Sarkar, and A. Subrahmanyam (2003). An Empirical Analysis of Stock and Bond Market Liquidity. *Federal Reserve Bank of New York Staff Reports*, No. 164.

[7] Commodity Futures Trading Commission (2007). *Backgrounder*, Washington.

[8] Dacorogna, M. M., R. Gençay, U. A. Müller, R. B. Olsen and O. V. Pictet (2001). *An Introduction to High Frequency Finance*. Academic Press, San Diego.

[9] Gorton, G., F. Hayashi, and K.G. Rouwenhorst (2007). The Fundamentals of Commodity Futures Returns. *Yale International Center for Finance Working Paper*, No. 07-08.

[10] Haigh, M., J. Hranaiova, and J. Overdahl (2005). Price Dynamics, Price Discovery and Large Futures Trader Interactions in the Energy Complex. *CFTC Working Paper*, Washington.

[11] Hamilton, J.D. (1994). *Time Series Analysis*. Princeton University Press, Princeton.

[12] Hasbrouck, J. and G. Sofianos (1993). The Trades of Market Makers: An Empirical Analysis of NYSE Specialists. *Journal of Finance*, Vol. 48, No. 5, 1565-1593.

[13] Huang, R., J. Cai, and X. Wang (2001). Inventory Risk-Sharing and Public Information-Based Trading in the Treasury Note Interdealer Broker Market. *University of Notre Dame Working Paper*, Notre Dame.

[14] Miffre, J. (2002). The Predictability of Futures Returns: Rational Variation in Required Returns or Market Inefficiency? *Applied Financial Economics*, 12, 715-724.

[15] LeBaron, B. (1992). Some Relationships Between Volatility and Serial Correlation in Stock Market Returns. *Journal of Business*, 65, 199-219.

[16] Sanders, D., K. Boris, and M. Manfredo (2004). Hedgers, Funds, and Small Speculators in the Energy Futures Markets: An Analysis of the CFTC's Commitments of Traders Reports. *Energy Economics*, 26, 425-445.

[17] Sheskin, D.J. (2000). *Handbook of Parametric and Nonparametric Statistical Procedures*. Chapman & Hall, New York.

[18] Wang, C. (2003). Investor Sentiment, Market Timing, and Futures Returns. *Applied Financial Economics*, 13, 891-898.

[19] Zivot, E. and J. Wang (2006). *Modeling Financial Time Series With S-Plus - Second Edition*. Springer, New York.

Trader Class	Min	Max	Median	Mean	Std Dev	Skewness	Kurtosis
Speculators	$-71,928$	$88,712$	$8,361$	$8,471$	$30,544$	0.06207	2.777
Hedgers	$-103,854$	$94,868$	$-8,103$	$-7,804$	$39,144$	-0.05284	2.617

Table 1: Summary Sample Statistics for the Net Long Measure of Position Holdings for Large Hedgers and Speculators. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006, as reported in the CFTC Commitments of Traders Reports.

$$r_t = \varphi_0 + \varphi_1 r_{t-1} + \varphi_2 NL^S_{t-1} + \epsilon_{1,t}$$

	Intercept	r_{t-1}	NL^S_{t-1}
Coefficients	341.77	-0.05	0.06
Std Error	18.47	0.04	0.04
t-stat	18.49	-1.32	1.63
p-value	0.0000	0.1848	0.1029

$\overline{R}^2 = 0.0024$

$$NL^S_t = \beta_0 + \beta_1 r_{t-1} + \beta_2 NL^S_{t-1} + \epsilon_{2,t}$$

	Intercept	r_{t-1}	NL^S_{t-1}
Coefficients	97.71	0.09	0.63
Std Error	14.20	0.03	0.03
t-stat	6.88	3.15	21.46
p-value	0.0000	0.0017	0.0000

$\overline{R}^2 = 0.4262$

Table 2: Results of Fitted VAR(1) Model for Returns and Positions of Speculators. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

$$r_t = \varphi_0 + \varphi_1 r_{t-1} + \varphi_2 NL^H_{t-1} + \epsilon_{1,t}$$

	Intercept	r_{t-1}	NL^H_{t-1}
Coefficients	377.32	−0.05	−0.04
Std Error	21.70	0.04	0.04
t-stat	17.38	−1.20	−1.17
p-value	0.0000	0.2290	0.2428
$\overline{R}^2 = 0.0006$			

$$NL^H_t = \beta_0 + \beta_1 r_{t-1} + \beta_2 NL^H_{t-1} + \epsilon_{1,t}$$

	Intercept	r_{t-1}	NL^H_{t-1}
Coefficients	157.82	−0.07	0.62
Std Error	17.01	0.03	0.03
t-stat	9.28	−2.38	20.78
p-value	0.0000	0.0177	0.0000
$\overline{R}^2 = 0.4012$			

Table 3: Results of Fitted VAR(1) Model for Returns and Positions of Hedgers. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

$$r_t = \alpha + \beta_1 SLHS_t + \beta_2 SLHS_{t-1} + \gamma_1 HLSS_t + \gamma_2 HLSS_{t-1} + \varepsilon_t$$

	Intercept	$SLHS_t$	$SLHS_{t-1}$	$HLSS_t$	$HLSS_{t-1}$
Coefficient	0.000	0.019	−0.031	−0.018	0.024
Std Error	0.002	0.007	0.007	0.012	0.011
t-stat	0.095	2.530	−4.331	−1.465	2.130
p-value	0.924	0.012	0.000	0.143	0.034
R^2	0.025				

Table 4: Regression Model of the Effect of Current and Past Extreme Trader Positions on Weekly Returns. The variable $SLHS_t$ is an indicator variable for times when large speculators are extremely long and large hedgers are extremely short, and similarly $HLSS_t$ is an indicator variable for times when large speculators are extremely short and large hedgers are extremely long. Heteroskedasticity robust standard errors and p-values are reported. The error term is assumed to be from a white noise process. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

$$\ln(RV_t) = \alpha + \beta \ln(RV_{t-1}) + \gamma HLSS_t^* + \varepsilon_t$$

	Intercept	$\ln(RV_{t-1})$	$HLSS_t^*$
Coefficient	-2.254	0.629	0.066
Std Error	0.712	0.116	0.027
t-stat	-3.167	5.434	2.464
p-value	0.002	0.000	0.014
R^2	0.465		

Table 5: Regression Model of the Effect of Extreme Trader Positions on Weekly Realized Volatility. RV_t is the realized volatility described in Section 2.1, $HLSS_t^*$ is an indicator variable for whether a $HLSS$-type event occurs within the next five weeks, and ε_t is assumed to be a white noise process. $HLSS_t$-type events occur when large speculators are extremely short and large hedgers are extremely long. Heteroskedasticity robust standard errors and p-values are reported. The error term is assumed to be from a white noise process. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

$$r_t = \alpha + \beta_1 \triangle T_{long,t}^{spec} + \beta_2 \triangle T_{short,t}^{spec} + \varepsilon_t$$

	Intercept	$\triangle T_{long,t}^{spec}$	$\triangle T_{short,t}^{spec}$
Coefficient	0.000	0.003	-0.001
Std Error	0.002	0.000	0.000
t-stat	-0.080	10.771	-1.524
p-value	0.937	0.000	0.128
R^2	0.173		

Table 6: Regression Model of the Effect of Changes in the Number of Speculators on Weekly Returns. The variable $\triangle T_{short,t}^{spec}$ is the change in the number of speculators who are short, and $\triangle T_{long,t}^{spec}$ is the change in the number of speculators who are long. Heteroskedasticity robust standard errors and p-values are reported. The error term is assumed to be from a white noise process. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

$$100 \cdot RV_t = \alpha + \beta \cdot 100 \cdot RV_{t-1} + \delta_1 \triangle T^{spec}_{long,t} + \delta_2 \triangle T^{spec}_{short,t} + \varepsilon_t$$

	Intercept	$100 \cdot RV_{t-1}$	$\triangle T^{spec}_{long}$	$\triangle T^{spec}_{short}$
Coefficient	0.147	0.430	-0.002	-0.003
Std Error	0.030	0.124	0.001	0.002
t-stat	4.910	3.467	-2.107	-1.450
p-value	0.000	0.001	0.036	0.147
R^2	0.195			

Table 7: Regression Model of the Effect of Changes in the Number of Speculators on Weekly Realized Volatility. The variable $\triangle T^{spec}_{short,t}$ is the change in the number of speculators who are short, and $\triangle T^{spec}_{long,t}$ is the change in the number of speculators who are long. Heteroskedasticity robust standard errors and p-values are reported. The error term is assumed to be from a white noise process. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

$$r_t = \alpha + \beta_1 SLHS_t + \beta_2 SLHS_{t-1} + \gamma_1 HLSS_t + \gamma_2 HLSS_{t-1} + \varepsilon_t$$

	Intercept	$SLHS_t$	$SLHS_{t-1}$	$HLSS_t$	$HLSS_{t-1}$
Coefficient	0.001	0.039	-0.022	-0.039	0.018
Std Error	0.002	0.004	0.005	0.006	0.006
t-stat	0.248	8.899	-4.388	-6.318	2.926
p-value	0.804	0.000	0.000	0.000	0.004
R^2	0.107				

Table 8: Regression Model of the Effect of Current and Past Extreme Trader Positions on Weekly Returns Using Only Past Information. The variable $SLHS_t$ is an indicator variable for times when large speculators are extremely long and large hedgers are extremely short, and similarly $HLSS_t$ is an indicator variable for times when large speculators are extremely short and large hedgers are extremely long. Heteroskedasticity robust standard errors and p-values are reported. The error term is assumed to be from a white noise process. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

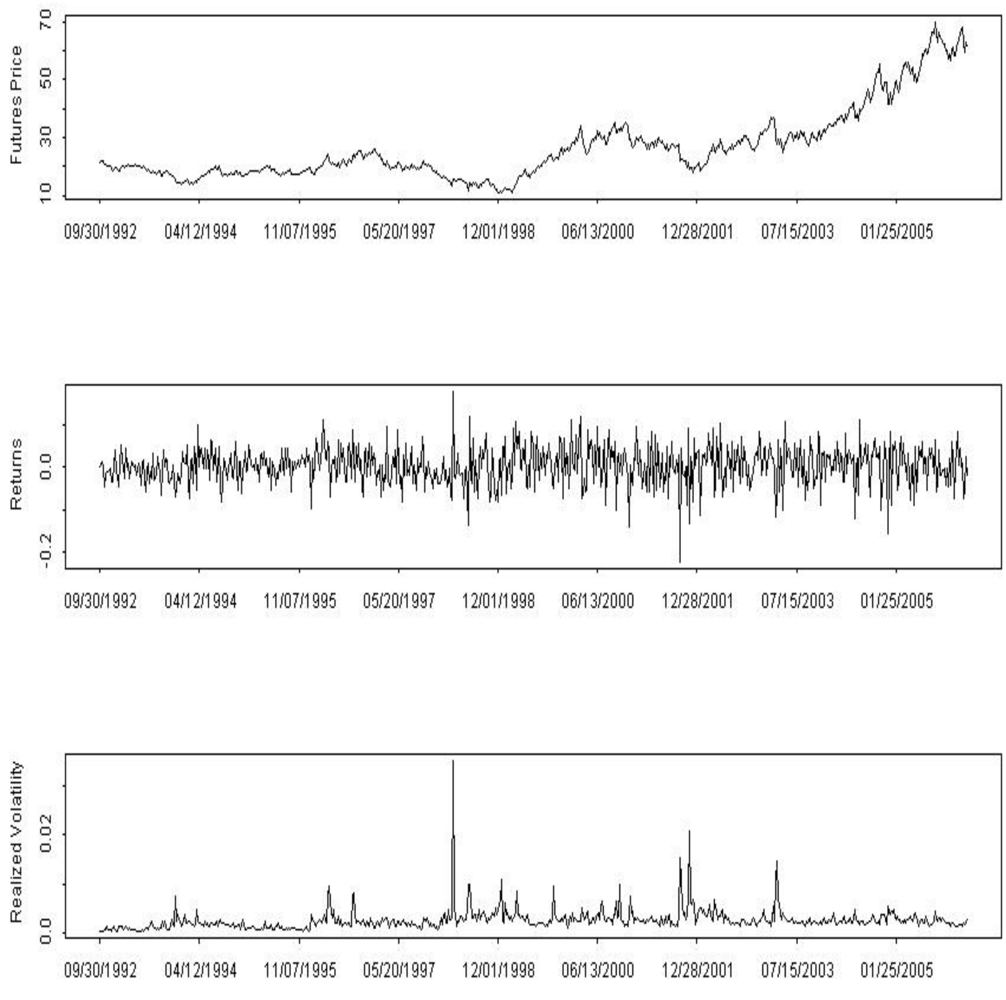

Figure 1: Weekly Prices (top frame), Returns (middle frame), and Realized Volatility (bottom frame) for NYMEX Light Sweet Crude Oil Futures. Prices are the closing price on each Tuesday. Returns are based on an unleveraged investment of the value of the underlying commodity at the futures price of the previous Tuesday. Realized volatility is also sampled each Tuesday, and is constructed by summing the squared five-minute returns over the previous week. The sample period is from September 30, 1992 to February 28, 2006 with 698 observations. Source: TickData, www.tickdata.com.

```
CRUDE OIL, LIGHT SWEET - NEW YORK MERCANTILE EXCHANGE                Code-067651
FUTURES ONLY POSITIONS AS OF 07/17/07                              |
-------------------------------------------------------------------| NONREPORTABLE
      NON-COMMERCIAL        |    COMMERCIAL    |      TOTAL        |   POSITIONS
----------------------------|------------------|-------------------|-----------------
  LONG  | SHORT  |SPREADS |  LONG  | SHORT   |  LONG  | SHORT    |  LONG  | SHORT
-------------------------------------------------------------------------------------
(CONTRACTS OF 1,000 BARRELS)                         OPEN INTEREST:    1,549,425
COMMITMENTS
 246,844  137,421  310,801  911,229  1015492  1468874  1463714   80,551   85,711

CHANGES FROM 07/10/07 (CHANGE IN OPEN INTEREST:      3,053)
  -2,201      663   12,566  -10,804  -18,584     -439   -5,355    3,492    8,408

PERCENT OF OPEN INTEREST FOR EACH CATEGORY OF TRADERS
    15.9      8.9     20.1     58.8     65.5     94.8     94.5      5.2      5.5

NUMBER OF TRADERS IN EACH CATEGORY (TOTAL TRADERS:      324)
     107       94      127       91      102      268      269
```

Figure 2: Example of a CFTC Commitments of Traders (CoT) Report. Released every Friday, the CoT report summarizes the positions of various classes of traders as of the previous Tuesday. The commercial class refers to large hedgers, and the non-commercial class refers to large speculators. Non-reportable positions are those held by small traders. The Commitments of Traders Reports are published every Friday on the CFTC website at http://www.cftc.gov/cftc/cftccotreports.htm.

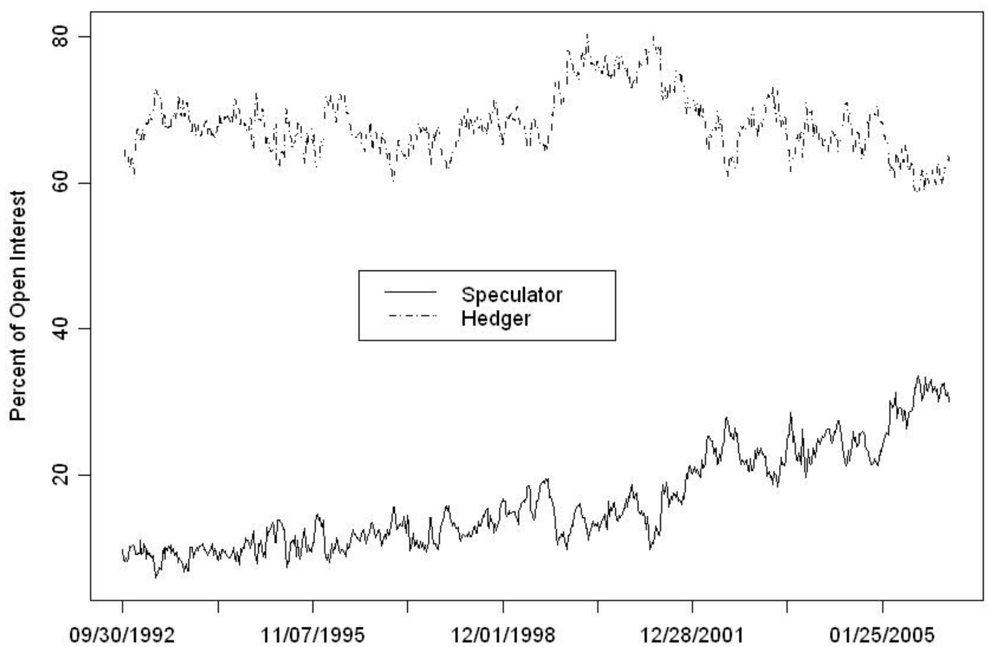

Figure 3: Percent of Open Interest for Large Hedgers and Speculators. Information is from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports.

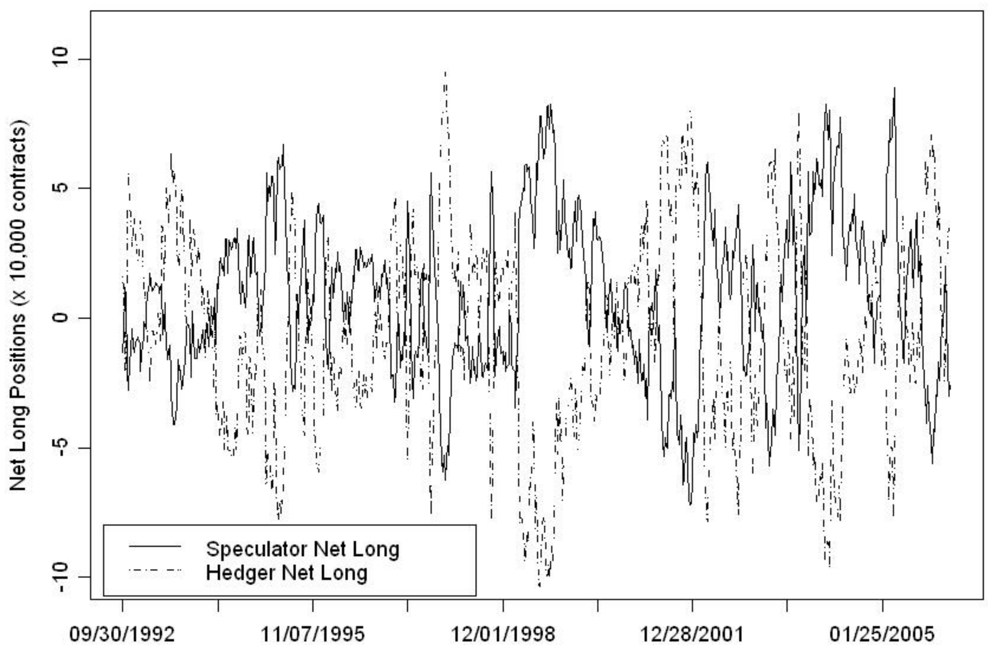

Figure 4: Net Long Positions of Large Hedgers and Speculators. The net long positions are the number of long positions minus the number of short positions within the trader class. Values are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006, as reported in the CFTC Commitments of Traders Reports.

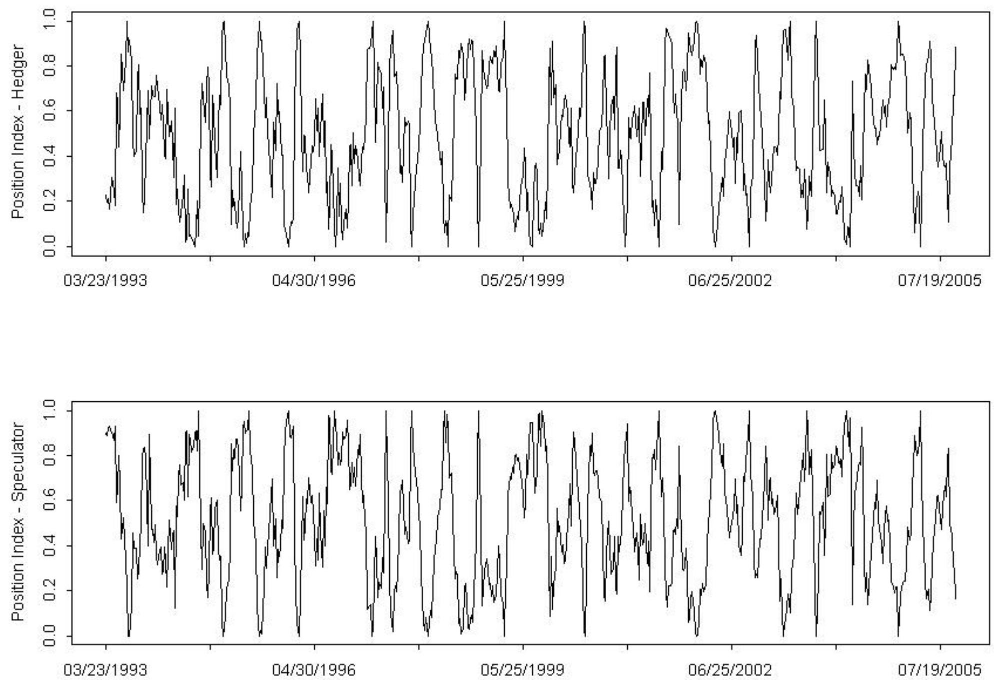

Figure 5: Trader Position Index, $P_t^C(\tau)$, for Hedgers and Speculators with $\tau = 39$ Weeks. The position index is a measure of how long a trader class relative to the maximum and minimum net long positions of that class over the centered window of length τ weeks. A value of 1 indicates that a new τ-period net long maximum has been reached. A value of 0 indicates that a new τ-period net long minimum (a net short maximum) has been reached. The upper panel is the position index for large hedgers, and the lower panel is the position index for large speculators. Values are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006, as reported in the CFTC Commitments of Traders Reports.

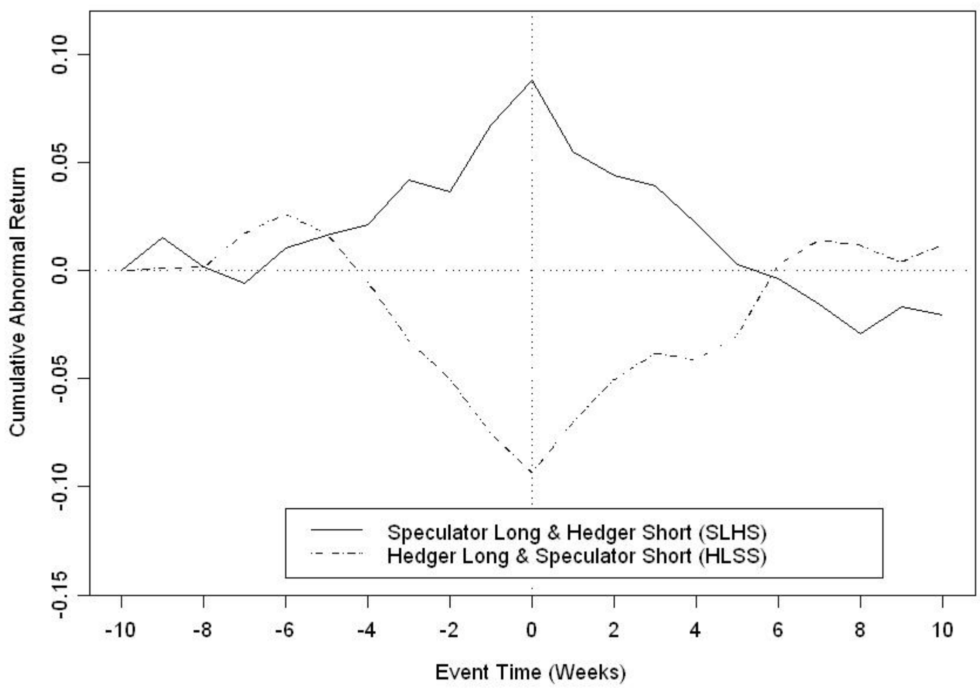

Figure 6: Averaged Cumulative Abnormal Returns. Returns are assumed to have a constant mean. Time relative to the event date, measured in weeks, are on the horizontal axis. Events at time zero (indicated by a vertical dotted line) are extreme opposing positions by large hedgers and speculators. When speculators are long and hedgers are short (solid line), the event date coincides with a local price top. When speculators are short and hedgers long, the event date coincides with a local price bottom. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Within the sample period, there are 20 extreme events of each type for a total of 40 extreme events. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

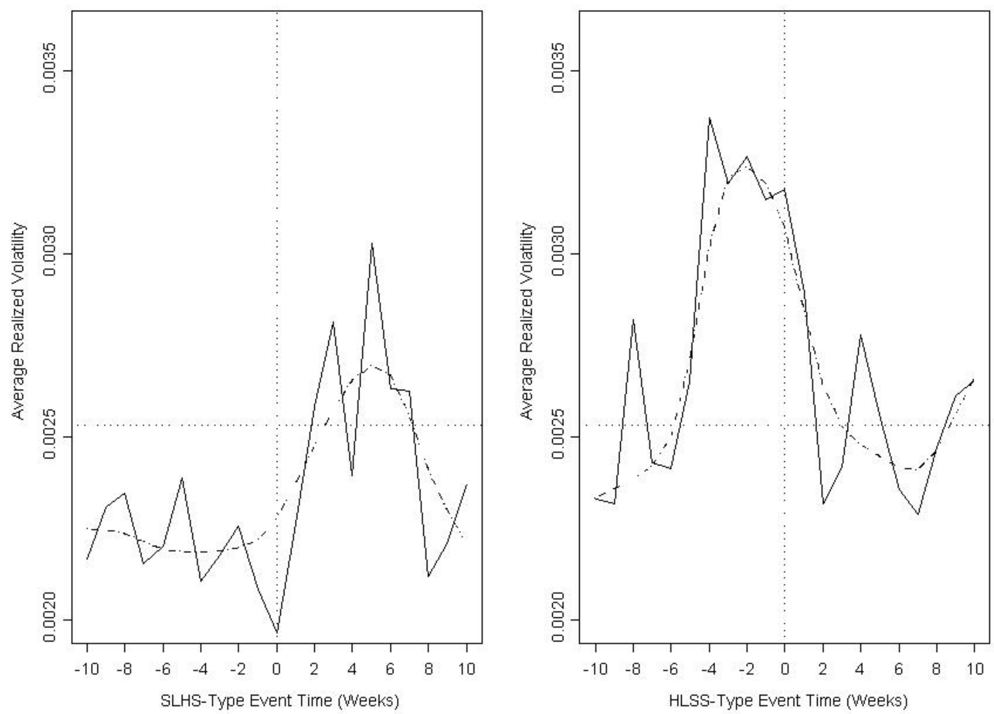

Figure 7: Average Realized Volatility Around Type $SLHS$-Type Event Dates and Type $HLSS$-Type Event Dates. $SLHS$-type events are when speculators are extremely long and hedgers extremely short. $HLSS$-type events are when speculators are extremely short and hedgers extremely long. The dashed lines are a smoothed versions of the average realized volatility computed by using a method of running medians known as 4(3RSR)2H with twicing. The dotted horizontal line is the average realized volatility over the entire sample. Times relative to the event date, measured in weeks, are on the horizontal axis from 10 weeks before traders take an extreme position until 10 weeks after. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

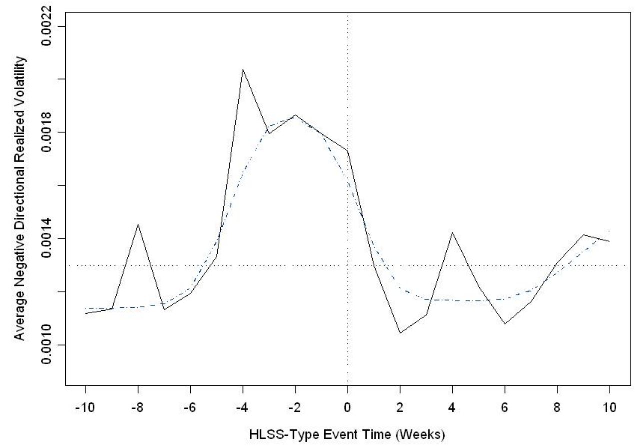

(a)

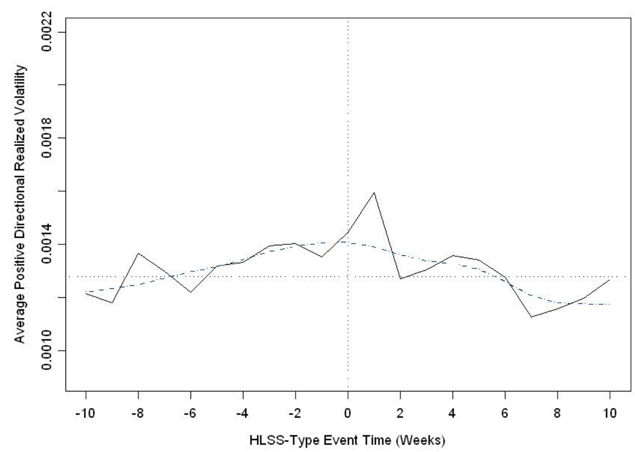

(b)

Figure 8: Average Directional Realized Volatility Around Type $HLSS$-Type Event Dates. $HLSS$-type events are when speculators are extremely short and hedgers extremely long. Panel (a) plots the average negative directional realized volatility, and panel (b) plots the positive directional realized volatility. The dashed lines are a smoothed versions of the average realized volatility computed by using a method of running medians known as 4(3RSR)2H with twicing. The dotted horizontal lines in the respective panels are the average negative and positive directional realized volatility over the entire sample. Times relative to the event date, measured in weeks, are on the horizontal axis from 10 weeks before traders take an extreme position until 10 weeks after. Note that the spike in panel (a) at time -1 week is spurious and is caused by a single large value. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.

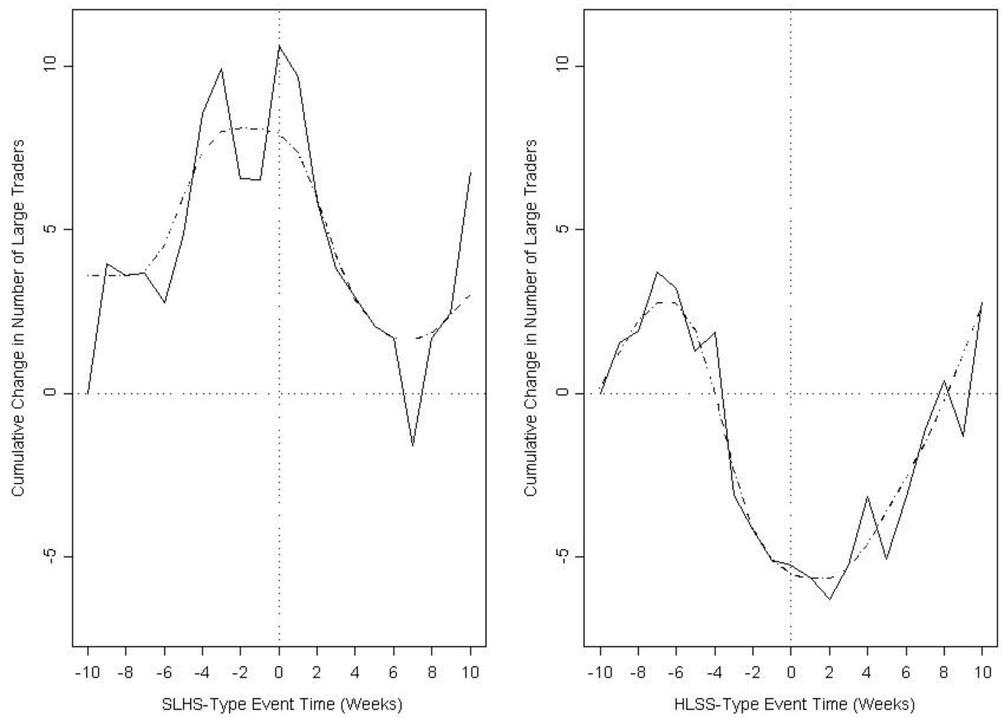

Figure 9: Averaged Cumulative Change in the Number of Large Traders Around Dates when Speculators and Hedgers Have Extreme and Opposing Positions. $SLHS$-type events are when speculators are extremely long and hedgers extremely short. $HLSS$-type events are when speculators are extremely short and hedgers extremely long. The dashed line is a smoothed version of the series, computed by using a method of running medians known as 4(3RSR)2H with twicing. Times relative to the event date, measured in weeks, are on the horizontal axis from 10 weeks before traders take an extreme position until 10 weeks after. The left panel shows that traders are entering the market before an $SLHS$-type event, and then leave the market afterwards. The right panel shows that traders are leaving the market before and slightly after an $HLSS$-type event, and then begin entering again. The statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. The number of traders in the market is reported in the CFTC Commitments of Traders Reports.

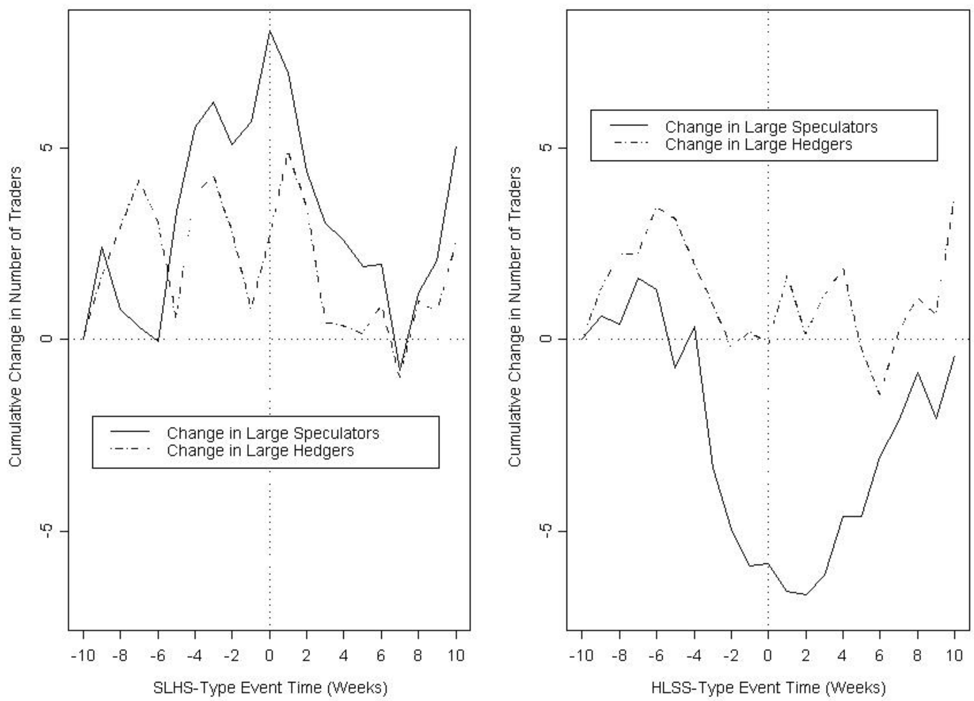

Figure 10: Averaged Cumulative Change in the Number of Speculators and Hedgers Around Dates when Speculators and Hedgers have Extreme and Opposing Positions. $SLHS$-type events are when speculators are extremely long and hedgers extremely short. $HLSS$-type events are when speculators are extremely short and hedgers extremely long. The solid and dashed lines represent the cumulative change in the number of speculators and hedgers, respectively. Times relative to the event date, measured in weeks, are on the horizontal axis from 10 weeks before traders take an extreme position until 10 weeks after. The left panel shows that traders are entering the market before an $SLHS$-type event, and then leave the market afterwards. The right panel shows that traders are leaving the market before and slightly after an $HLSS$-type event, and then begin entering again. The statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. The number of traders in the market is reported in the CFTC Commitments of Traders Reports.

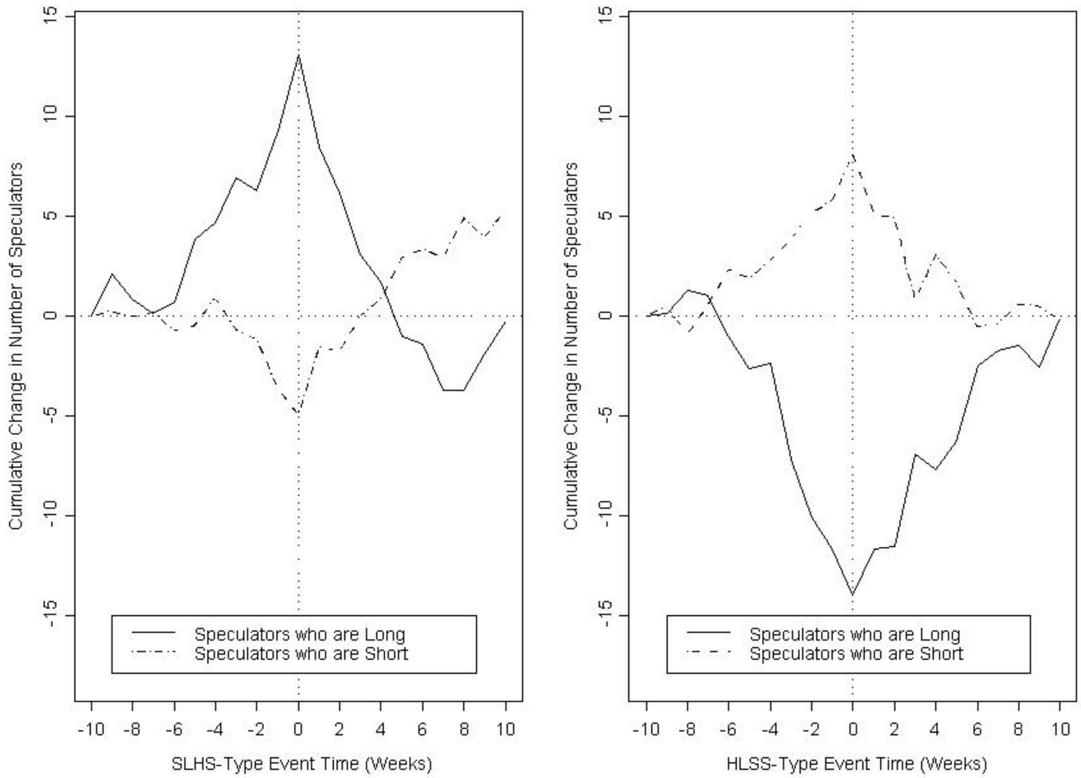

Figure 11: Averaged Cumulative Change in the Number of Speculators Around Dates Where Hedgers are Extremely Long and Speculators Extremely Short, disaggregated into to speculators who are long and speculators who are short. $SLHS$-type events are when speculators are extremely long and hedgers extremely short. $HLSS$-type events are when speculators are extremely short and hedgers extremely long. The solid and dashed lines represent the cumulative change in the number of speculators that are long and that are short, respectively. Times relative to the event date, measured in weeks, are on the horizontal axis from 10 weeks before traders take an extreme position until 10 weeks after. The dotted vertical line indicates the event time. The averaged cumulative change in the total number of large hedgers. The statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. The number of traders in the market is reported in the CFTC Commitments of Traders Reports.

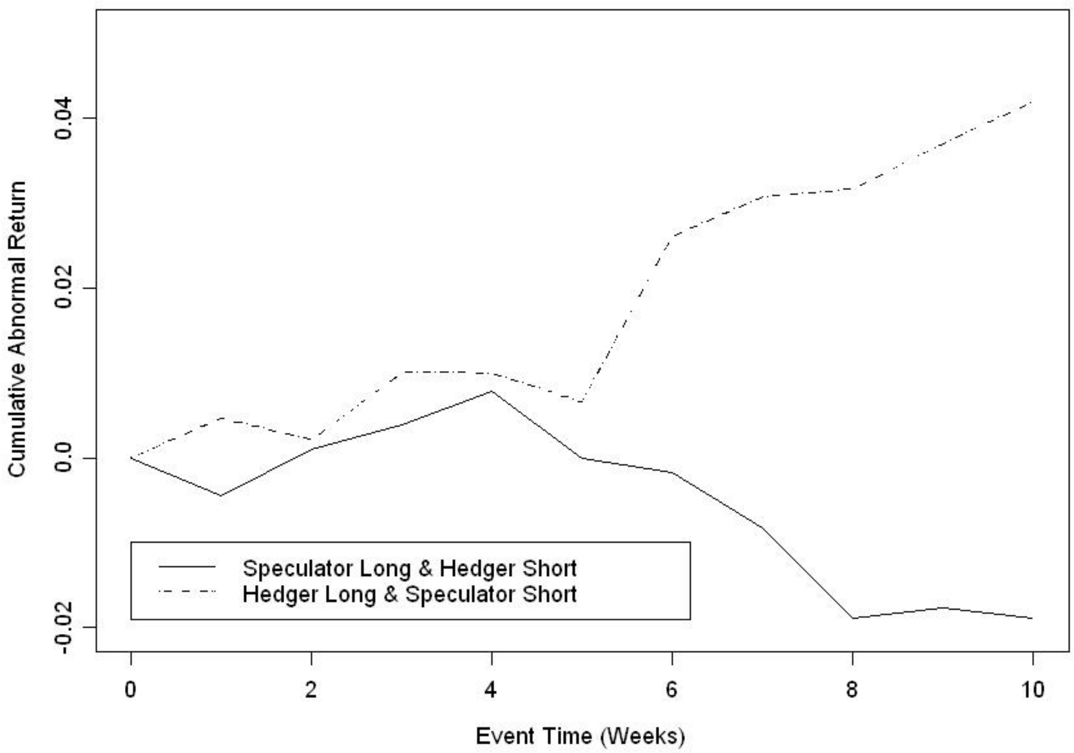

Figure 12: Averaged Cumulative Abnormal Return Using Only Future Information. Returns are assumed to have a constant mean. Time relative to the event date, measured in weeks, are on the horizontal axis. Events at time zero (indicated by a vertical dotted line) are extreme opposing positions by large hedgers and speculators. When speculators are long and hedgers are short (solid line), the event date coincides with a local price top. When speculators are short and hedgers long, the event date coincides with a local price bottom. Statistics are calculated from 698 weekly observations from September 30, 1992 until February 28, 2006. Trader positions are reported in the CFTC Commitments of Traders Reports and the weekly returns are calculated from tick-level price data from TickData, www.tickdata.com.