Predatory or Sunshine Trading? Evidence from Crude Oil ETF Rolls *

Hendrik Bessembinder
University of Utah

Allen Carrion
University of Utah

Laura Tuttle
American University

Kumar Venkataraman
Southern Methodist University

Abstract

We study prices, liquidity, and individual account trading activity around large and predictable ETF "roll" trades in crude oil futures markets to test the implications of the competing theories of predatory and sunshine trading. The results indicate narrower bid-ask spreads, greater limit order book depth, and a larger number of distinct trading accounts providing liquidity on roll versus non-roll dates. We also extend the theory of strategic trading ahead of a known liquidation to show that even a monopolist trader has incentives to effectively provide liquidity rather than follow predatory strategies in a "resilient" market. We estimate that ETFs effectively pay about 30 basis points to complete their roll trades. On balance, the evidence supports the implications of the sunshine trading theories and of our modified theory of strategic trading around a known liquidation, and fail to support the implications of predatory trading models.

This Draft: March 2012

Keywords: Predatory trading; Sunshine Trading; Resiliency; ETFs; Commodity trading.

---

* We thank Jeff Harris, Andrei Kirilenko and Jim Moser for facilitating access to CFTC data, Nela Richardson for comments and assistance in identifying ETF accounts, and Stephen Meek for research assistance. Much of the work herein was conducted while Allen Carrion and Laura Tuttle served as visiting economists and Hendrik Bessembinder and Kumar Venkataraman served as consultants to the Office of Economic Analysis at the Commodity Futures Trading Commission. The views presented here are our own and do not represent the views of the Commodity Futures Trading Commission, its Commissioners or staff.

Predatory or Sunshine Trading? Evidence from Crude Oil ETF Rolls

Abstract

We study prices, liquidity, and individual account trading activity around large and predictable ETF "roll" trades in crude oil futures markets to test the implications of the competing theories of predatory and sunshine trading. The results indicate narrower bid-ask spreads, greater limit order book depth, and a larger number of distinct trading accounts providing liquidity on roll versus non-roll dates. We also extend the theory of strategic trading ahead of a known liquidation to show that even a monopolist trader has incentives to effectively provide liquidity rather than follow predatory strategies in a "resilient" market. We estimate that ETFs effectively pay about 30 basis points to complete their roll trades. On balance, the evidence supports the implications of the sunshine trading theories and of our modified theory of strategic trading around a known liquidation, and fail to support the implications of predatory trading models.

I. Introduction

A trader who gains knowledge that another investor will buy or sell a substantial quantity of a security can potentially profit by trading in the same direction as the investor and reversing the position after the investor transactions are complete. In the case of brokers who are aware of a client's pending orders, the practice is known as "front running." More broadly, the practice has been dubbed "predatory trading" by Brunnermeier and Pedersen (2005). Their model shows that the practice can be disruptive to the markets, in that it causes prices to temporarily overshoot their longer-term equilibrium, and that the predator's profits come at the expense of the other investor.

In this paper, we study individual account trading strategies, overall liquidity levels, and price patterns around the time of large and predictable monthly trades undertaken by eight exchange-traded-funds (ETFs) that are designed to provide returns that track crude oil prices. Rather than holding crude oil inventories, which would entail substantial storage costs, these funds gain exposure to crude oil prices by holding positions in New York Mercantile Exchange (NYMEX) crude oil futures contracts. However, since a NYMEX contract expires each month, the strategy involves "rolling" positions on a monthly basis by selling the expiring contract and purchasing contracts with more distant expiration dates.

To illustrate the relative performance of Oil ETFs versus Crude Oil, we focus on United States Oil Fund (USO), the largest of the ETFs designed to track crude oil prices.[1] Launched in April 2006, USO's assets-under-management grew rapidly. At its peak in early 2009, USO had more than $4.2 billion under management, equating at prevailing prices to over 90,000 NYMEX contracts or 90 million barrels of crude oil. Returns to USO investors have lagged the level of crude oil prices, as displayed on Figure 1. On April 10, 2006, when USO initiated trading, the settlement price of the nearest delivery (May 2006) crude oil futures contract was $68.74, while USO's share price was $68.02. Five and a half years later, on December 12, 2011, the price of the nearest futures contract (for January 2012 delivery) was $97.77, while the USO share price had fallen to $37.88.[2]

[1] Data on ETF's assets-under-management are obtained from Bloomberg.
[2] USO paid no dividends, nor did it split its shares.

Some observers have suggested that predatory trading practices help to explain that the USO share price declined even while crude oil prices rose. For example, the Wall Street Journal reported that "Since the fund (USO) is so big, it is unable to switch in and out of contracts….without moving markets and giving speculators an opportunity to make bets on those moves."[3] The article quotes a trader as stating that "It's like taking candy from a baby" and asserts that the "…. candy comes out of returns of the investors in the fund". Similarly, according to Bloomberg "Professional futures traders exploit the ETFs' monthly rolls to make easy profits at the little guy's expense. Unlike ETF managers, the professionals don't trade at set times. They can buy the next month ahead of the big programmed rolls to drive up the price, or sell before the ETF, pushing down the price investors get paid for expiring futures."[4]

However, predatory trading is not the only potential outcome when investors are aware of large pending orders. Admati and Pfleiderer (1991) present the alternative theory of "sunshine trading," whereby investors who need to make a large transaction and who can credibly signal that their trade is not motivated by private information regarding security value can potentially reduce their transaction costs by preannouncing their intention to trade. The announcement of the pending trade can attract potential natural counterparties as well as additional liquidity suppliers to market, thereby reducing trade execution costs. Consistent with the sunshine trading view, the larger ETFs based on assets-under-management in our sample preannounce the roll dates while the remaining ETFs track commodity indexes that roll on a preannounced schedule.

Studying trading activity and market prices around the time of ETF rolls is of particular interest because the rolls provide opportunities to test the implications of both predatory trading and sunshine trading theories. The Brunnermeier and Pedersen (2005) model focuses on large pending orders that are predictable to potential predators. Commodity ETF roll orders, and in particular those of USO, are large and fully predictable. At the same time, commodity ETFs, being designed to simply mimic commodity

[3] "U.S. Oil Fund Finds Itself at the Mercy of Traders", by Gregory Meyer and Carolyn Cui, The Wall Street Journal, March 6, 2009, page C1.

[4] "ETFs Imperil Investors as Contango, Pre-Roll Conspire", by Peter Robison, Asjylyn Loder and Alan Bjerga. Bloomberg Newsroom, July 22, 2010. The article goes on to quote a trader as saying "I make a living off the dumb money.....These index funds get eaten alive by people like me."

price changes, are unlikely to be perceived as informed traders. Preannounced trades by credibly uninformed traders comprise sunshine trading as envisioned by Admati and Pfleiderer (1991). Our analysis allows us to assess the empirical relevance of the competing predatory trading and sunshine trading theories.

To do so, we employ data on individual orders and trades in crude oil futures made available to us by the Chicago Mercantile Exchange, which owns and operates the NYMEX market. In addition, we use Commodity Futures Trading Commission (CFTC) data that identifies the individual trading accounts associated with each crude oil futures transaction. The former dataset allows us to study posted liquidity in the form of bid and ask quotes, as well as unexecuted displayed depth in the limit order book. The latter dataset allows us to evaluate the strategies used by owners of specific trading accounts around the time of the ETF rolls. Our study of individual orders and trades spans the period March 1, 2008 to February 28, 2009, and therefore includes twelve monthly rolls. We also study daily crude oil settlement price data obtained from the Commodity Research Bureau (CRB) spanning the period from January 1990 through November 2011.

In addition to an empirical study of trading strategies and price movements around the USO rolls, we provide some new analysis of the economics of predatory trading. While the economics issues apply to either buy or sell orders, we will for simplicity focus our discussion on a pending liquidation by the investor. Brunnermeier and Pedersen (2005) analyze predatory trading in a model where trades have permanent effects on price that are proportional to the quantity transacted. In addition, in their model the predator trades simultaneous with, or after, the liquidator. We relax these restrictions to analyze the effects of predatory trading when markets are resilient, in the sense that some or all of the immediate price impact of trades is subsequently reversed, and where the predator can trade ahead of, simultaneous with, or after the liquidator. We focus for simplicity on a monopolist predator.[5]

[5] Brunnermeier and Pedersen (2005) show that competition among predators generally reduces their negative effect.

Our analysis reveals that, when the market is reasonably resilient, the profit maximizing strategy for the predator is to sell before (and, for some parameters after) the period where the liquidator trades, while purchasing during the period that the liquidator trades. In other words, the predator essentially chooses to absorb a portion of the liquidator's order imbalance while it occurs, while offloading the resulting inventory in periods both before and after the liquidation. This analysis shows that a key implication of the Brunnermeier and Pedersen (2005) model, that a profit-maximizing predator will necessarily degrade market quality and impose losses on the investor, need not generalize to settings where the market is resilient in the sense that the price impact of trades is not fully permanent and when the predator can also trade prior to the investor.

Our empirical analysis reveals several findings. First, the oil futures market is indeed resilient. Using CME order book data, we implement a geometric lag regression of price changes on lagged order imbalances to estimate (a) the permanent and temporary component of trading costs and (b) a resiliency parameter that captures the extent to which temporary price impacts persist beyond the period of the order imbalance. The resulting estimates imply that over 99% of the temporary price impact due to an order imbalance is reversed within ten minutes, implying that the limit order book refills rapidly. However, point estimates indicate that the market is marginally less resilient on ETF Roll days than non-Roll days. We also document a reduction in the permanent price impact of order imbalances on ETF roll days, which is consistent with the sunshine trading interpretation. However, our estimates indicate an increase in the temporary price impact of trades on ETF roll days, consistent with the interpretation that liquidity providers must be compensated to absorb the large order imbalances that accompany the rolls.

Second, we find that many conventional measures indicate improved liquidity on roll versus non-roll days. In particular, quoted bid-ask spreads are narrower on roll days, and the quantity of unexecuted orders in the limit order book at prices near the inside is greater on roll days. Further, a larger number of distinct accounts provide liquidity through limit orders on ETF roll days relative to non-roll days. These findings also support Admati and Pfleiderer's (1991) theoretical prediction that preannouncement by a large liquidity trader increases market size by attracting liquidity providers and natural counterparties.

However, market outcomes also reflect the unusually large demand for liquidity, as evidenced by larger intraday buy-sell imbalances and trading volume on Roll days relative to non-roll days. Despite the increased demand, the average effective bid-ask spread on Roll days is marginally smaller, indicating that the imbalance is successfully absorbed by the deeper limit order book.

Third, based on the CFTC data, we find little or no evidence that individual trading accounts use strategies that would reasonably be considered predatory, in the sense of the Brunnermeier and Pedersen (2005) model. In fact, consistent with the simple framework introduced in the paper, we find significant increased usage of a liquidity provision strategy where strategic traders sell the expiring contract the trading day before the roll and offload the resulting position on and after the Roll day, thereby absorbing a portion of the ETF sales during the roll day while shifting a portion of the selling pressure to the preceding day. Our theoretical analysis implies that this strategy mitigates temporary price impacts and improves prices for the rolling ETF.

Fourth, our analysis of daily settlement prices from CRB database indicates that the ETFs do pay to execute their roll trades – about 30 basis points on average per roll, or about 4.4% per year – in the form of settlement prices that diverge on roll days from preceding days. However, this estimate is not excessive relative to the estimated costs of executing large trades in other markets, including large-capitalization equities. Finally, we discuss the relevance of the term structure of crude oil prices and reconcile the lack of evidence in support of predatory strategies with the apparent investment underperformance documented in Figure 1.

The rest of the paper is organized as follows. Section 2 discusses the related literature, the NYMEX market structure, data sources and summary statistics on ETF trading activity on roll days. Section III describes the change in market quality surrounding the Roll days. Section IV presents a model of trading by a monopolist predator and Section V estimates the resiliency of the NYMEX crude oil market. Section VI examines account level activity by strategic traders surrounding the roll while Section VII explains the importance of futures term structure on the performance of USO' stock price.

II. Related Work, Data Sources, and NYMEX market structure

a. Related Literature

A number of studies quantify aspects of risk and return in commodity futures markets.[6] Separately, a large literature has assessed trading strategies and the costs of completing trades, primarily focusing on traditional investment markets such as stocks and bonds.[7] Our study contributes in part by assessing trading costs for larger traders in an important commodity market. However, our most important contributions come in assessing the implications of the theory of predatory trading as well as the theory of sunshine trading, in a setting where both theories potentially apply.

Brunnermeier and Pedersen (2005) model the case where some traders become aware of another trader's need to liquidate a position, and introduce the label "predatory trading" to describe the strategies followed. In their model the predators sell alongside the liquidating trader, before reversing their positions. The predatory trades damage market quality in that they cause the price to temporarily overshoot its long-term equilibrium. Further, the predators' profits come at the expense of lower proceeds to the liquidating trader. Carlin, Lobo, and Vishwanathan (2007) present a multiperiod model to study strategic interactions between large traders. In their model traders typically provide liquidity to each other. However, in situations where the potential profit from following a predatory strategy is sufficiently large to offset lost profits from future liquidity provision, traders can abandon liquidity provision to follow predatory strategies instead. Their model therefore predicts episodic periods of illiquidity attributable to predation.

The model presented by Brunnermeier and Pedersen (2005) assumes that trades have strictly permanent price impacts proportional to the size of the order imbalance. In contrast to this assumption, numerous studies have documented that large financial market trades have both temporary and permanent price impacts. Carlin, Lobo, and Vishwanathan (2007), as well Schoneborn and Schied (2007), model

[6] Recent examples include Hong and Yogo (2012), Gorton, Hayashi, and Rouwenhorst (2008), and Erb and Harvey (2006).

[7] See, for example, Keim and Madhavan (1997) and Jones and Lipson (2001) for evidence from equity markets, and Schultz (2001) and Bessembinder, Maxwell and Venkataraman (2006) for evidence from corporate bonds, and the papers referenced there.

predatory trading when order imbalances have both permanent and temporary price impacts. Schoneborn and Schied show that strategic traders may react to known liquidations by trading in the opposite rather than the same direction as the liquidator. Our model generalizes on that of Schoneborn and Schied in that we also assess the effect of market "resiliency", i.e. the extent to which trades' temporary price impacts spill over to periods subsequent to trade execution. In addition, the models presented in all three papers discussed in this section restrict the strategic traders to trade simultaneous with or after the liquidation, while we allow for trading prior to the liquidation as well.

Admati and Pfleiderer (1991) present an alternative theory relevant to large predictable trades. Their model considers a trader who needs to complete a large transaction and who is not motivated by private information regarding the value of the asset to be traded. They show that a public announcement of the intent to trade, termed "sunshine trading", can attract natural counterparties as well as liquidity suppliers who might not otherwise have been present, and can therefore allow reduce the trade's price impact and allow the trader to achieve a more favorable price.

Our empirical analysis of potential price impacts of roll trades in commodity futures is not entirely unprecedented. Stoll and Whaley (2010) and Mou (2011) study daily prices and rolls undertaken by index investors. In contrast to specialized ETFs that focus on a single commodity, some investors seek to generate returns that match the performance of multi-commodity indices, such as the Standard and Poor's-Goldman Sachs Commodity Index (SP-GSCI). To the extent that these index investors rely on futures positions to track the indices, they also generate periodic roll trades.[8] Stoll and Whaley report that periodic rolls by index investors have little or no effect on a broad cross-section of commodity futures prices, while Mou in contrast reports that significant profits can be earned by investors who trade in advance of the dates that the SP-GSCI index itself begins to track the second- rather than nearest-to-expiration futures contract. Our study is distinguished from these in that we are able to exploit individual trades, with account identifiers, as well as individual limit order book updates. The rich datasets allow us

[8] In addition, a number of authors have assessed whether index investors and other passive long-only investors have affected the level and/or the volatility of commodity prices. See, for example, Brunetti and Buyuksahin (2009), Buyuksahin and Harris (2011), and Kilian (2009).

to identify ETF roll trades with precision, to analyze intraday price and liquidity patterns, and to study the strategies followed by individual accounts around the time of the roll trades, in order to test the implications of predatory vs. sunshine trading theories.

b. The NYMEX Market

The New York Mercantile Exchange is the largest market for crude oil futures in the US. Although the NYMEX continues to operate a physical trading floor, the large majority of crude oil futures contracts are traded on the NYMEX's electronic limit order market, known as Globex. In addition, large traders can negotiate block trades. Though physical delivery is rare (as most contracts are netted against offsetting contracts), each individual NYMEX contract calls for delivery of 1,000 barrels of crude oil at Cushing, Oklahoma, during a designated delivery month. As in other futures markets, transaction prices reflect prices at which oil is to be delivered in the future, not an amount to be paid to enter the contract.

Trading hours for floor trades are 9:00 AM to 2:30 PM New York time. Globex trading occurs around the clock, except for a 45 minute break from 4:15 PM to 5:00 PM New York time. The weighted average price for each contract during the two-minute interval 2:28 to 2:30 PM New York time comprises the contract's "settlement price" for the day, and is used to calculate gains and losses on outstanding futures positions.[9] In particular, long positions receive and short positions pay the change in the settlement price since the prior day (or since the transaction price if entered the same day).

In addition to outright trades that specify a fixed delivery price, the NYMEX allows trading on a "Trade-at-Settlement" (TAS) basis. The price for a TAS trade is the current day settlement price (potentially plus or minus a specified margin), and is generally not known at the time of the trade. NYMEX crude oil settlement prices are widely-used benchmarks for valuing derivative contracts and determining final prices for over-the-counter physical delivery contracts.

[9] The exchange's settlement committee retains some discretion to designate settlement prices that differ from the weighted average transaction price in special circumstances, such as unusually light trading.

c. Data Sources

We employ three main datasets. The database from the Commodity Futures Trading Commission (CFTC) includes all completed trades in NYMEX crude oil futures during the period March 1, 2008 to February 28, 2009. The CFTC data includes floor and block trades, as well as trades completed on the Globex electronic market. In addition to information as to trade type, contract, price, and volume, the CFTC data includes an account identification variable for each trade. These account identifiers allow us to assess the number of unique trading accounts active on a given day and to track inventory changes by account during periods of interest. Although buyer and seller accounts are explicitly identified in the CFTC data, the initiator of the trade is not. We use a modified Lee-Ready algorithm to assign trades as buyer- or seller-initiated.[10]

We also obtain for the same time period from the Chicago Mercantile Exchange (owner of the NYMEX) a 5-level deep representation of the limit order book and a record of completed trades for crude oil futures on the electronic GLOBEX market. The CME dataset allows us to construct a continuous record of best bid and ask quotes, as well as the depth of unexecuted orders at and behind the best quotes. A limitation, however, is that the CME data does not include floor trades or negotiated block trades.[11] Third, we obtain from the Commodity Research Bureau (CRB) a daily record of settlement prices, trading volume, and open interest for each NYMEX crude oil contract traded over the longer time interval January 1999 through November 2011. We focus in particular on the daily data after April 2006, when USO was launched.

d. Descriptive Statistics

[10] Specifically, to sign trades, we compare the transaction price with the contemporaneous quote-midpoint (without a 5-second lag) and implement a five-trade look back for the tick-rule.

[11] The CME data has a finer time resolution (to the 100th of a second) as compared to the CFTC data, for which time stamps are truncated to the second. We use the CME transactions to impute centi-second time stamps for CFTC data transactions through an iterative process of matching unique price-quantity pairings. When transactions cannot be perfectly matched between the two datasets, the latest-possible time stamp is imputed within the CFTC dataset to avoid any possible look-ahead bias of matching trades with LOB information.

We rely on the CFTC data to identify the trading activity of eight ETFs for each trading day from March 1, 2008 to February 28, 2009. Aggregate assets-under-management (AUM) for sample ETFs increased significantly from $0.63 billion in March 2008 to $4.66 billion in February 2009. For eight of the twelve months, all ETF trading activity is observed on a single date. For the remaining months, ETF trading activity is observed on multiple days, but a single date in the month accounts for over 90% of ETF trading activity. We designate the date with the largest ETF activity each month as the ETF "roll date" for the month, and we report on trading activity aggregated across the eight ETFs on each roll date.

Table 1 reports on net ETF trading activity as a percentage of market volume for each monthly roll date, for the "front" (nearest to expiration) crude oil contract and for the second nearest to expiration contract. The large majority of roll trades comprise sales of the expiring contract and purchase of the second nearest-to-expiration contract. However, "short" funds trade in the opposite direction, and some funds purchase contracts for more distant maturity dates.[12] This data shows the rapid growth in ETF roll activity during the sample, with net ETF roll-date sales (i.e., selling minus buying activity) of the front contract increasing from 4,455 contracts representing 1% of market volume during the March 2008 roll to 67,882 contracts representing 13% of market volume during the February 2009 roll date. Aggregated across the twelve roll dates, net roll activity comprised 13% of roll-day volume in the front contract. Net ETF roll-date purchases (i.e., buying minus selling activity) accounts for 18% of roll-day volume in the more lightly traded second contract.

Table 1 also reports market volume during the two-minute settlement period. ETF's generally seek to track settlement price changes, and therefore have an interest in completing their trades during the two-minute period used to calculate the settlement price. However, ETF roll volume on average exceeds market volume during the settlement period. This fact suggests that it would be difficult for ETFs to complete their rolls by use of regular trades during the settlement interval.

[12] The Securities and Exchange Commission (SEC) requires the ETFs to provide information on the nature and details of the roll activity. Some ETFs disclose their daily holdings and the exact schedule of roll dates over the next twelve months on the ETFs webpage. Other ETFs more broadly describe their roll schedule and/or the benchmark index that they track in their prospectus. For the majority of the ETFs in our sample, we have obtained the roll dates and specific futures contracts associated with the roll from either the ETFs or the owners of the benchmark index.

Table 2 provides detail regarding the types of trades used by the ETFs and by the overall market. Focusing on market trading activity for front month contract, regular trades are most common, comprising 64% of total activity, followed by spread trades (which combine regular contracts for sale in a given delivery month and purchase in another month), which comprise 21% of activity. TAS trades (regular and negotiated block) comprise 5.5% of overall market activity, while other negotiated block trades comprise 1.3% of total activity. Results are similar for the second month contract, with the exception that spread trades are the most prevalent types in the second month.

In contrast, the ETF trading activity, including trades executed outside the 'roll date' for each month, are primarily TAS trades, comprising 59.5% negotiated block TAS trades and 7.2% TAS trades completed through GLOBEX or the trading floor.[13] Another 25.7% of ETF roll trades are negotiated block transactions at fixed (rather than explicitly settlement-based) prices.[14] However, these trades are virtually all reported after 2:30 PM, when the settlement price is known. Only 0.7% of ETF volume occurs as regular (individual or spread) trades. The heavy use of TAS contracts for ETF rolls reflects the ETF's interest in tracking settlement prices. We have verified that the volume-weighted average execution price for the ETF roll trades match (to the nearest \$.01 per barrel) the corresponding daily settlement prices, for both the front and second contracts, for all twelve rolls, with the lone exception of the front month contract on during the June 2008 roll.

TAS trades can be viewed as allowing the ETF to subcontract the final execution of the roll trades to the TAS counterparty. To the extent that TAS counterparty has a 'natural' offsetting position, the TAS trade enables (consistent with the theory of sunshine trading) both the ETF and counterparty to offset positions at low cost. If the TAS counterparty is a strategic liquidity provider then the counterparty's compensation, if any, is comprised of the difference between roll-day settlement prices (since the ETF's

[13] "Cabinet" trades comprise 6.8% of reported ETF roll volume. Conversations with CFTC data specialists indicate that these are likely miscategorized TAS trades.

[14] Virtually all ETF trades were negotiated blocks at fixed prices for the first nine rolls in the sample. In contrast, virtually all ETF trades were TAS trades during the final three rolls in the sample.

and hence their counterparties trade at prices matching the settlement prices), and the average price of the counterparties' offsetting trades.

III. Measuring Liquidity on Roll and Non-Roll Days

Carlin, Lobo, and Viswanathan (2007) present a model of large traders who most often find it in their interest to cooperate by providing liquidity to each other. However, they show that in situations where the short term profits are large enough, these traders may abandon liquidity provision to instead engage in predatory trading strategies. In Brunnermeier and Pedersen (2005) predation occurs ahead of any known liquidation. The sunshine trading theory of Admati and Pfleiderer (1991), in contrast, implies that the preannouncement of trading intentions, particularly by a credibly non-informed trader, can attract additional liquidity suppliers to market.

We provide some preliminary evidence on the relevance of predatory versus sunshine trading theories by simple comparisons, reported on Table 3, of average trading activity and market outcomes across roll and non-roll days in our primary (March 2008 to February 2009) sample. The basic unit of observation for results reported on Table 3 is one minute. Each measure is computed for each minute of the trading day between 9:00 AM and 3:00 PM EST, and results are then averaged across minutes. To account for intraday patterns in liquidity, we compare Roll and non-Roll market quality each minute and report the Wilcoxon signed rank test of the hypothesis that the median is equal across Roll and non-Roll days.

The results verify that trading activity (measured based on the CFTC data) is substantially greater on roll days, averaging 855 contracts per minute in the front contract and 402 contracts per minute in the second contract, compared to 648 contracts in the front month and 274 contracts in the second contract on non-roll days. Figure 2 displays average trading volume by minute for roll and non-roll days. Most notable is the spike in trading activity at the time of the daily settlement, particularly on roll days.

Next, we examine various measures of price pressure, bid-ask spreads, and book depth on Roll and non-Roll days using the NYMEX order data. To assess the magnitude of shocks to liquidity demand we define a "trade imbalance" measure for each minute based on the buyer-initiated (a buy market order executes against a standing sell limit order) less seller-initiated (a sell market order executes against a standing buy limit order) volume, and standardize the measure by subtracting the mean and dividing by the standard deviation of imbalance observed during the same minute on roll and non-roll days. In the front contract the net trade imbalance is on average negative on Roll days but positive on non-roll days, and the difference in medians is statistically significant (t-statistic = -1.82). The positive trade imbalance is consistent with the arrival of large traders (ETFs, or the counterparties to their TAS contracts) who demand liquidity with market sell orders in the front contract on roll days. Interestingly, we do not detect a net order imbalance in the second-nearest-to-delivery contract on roll days. This result reflects either that the ETF counterparties do not offset their positions on roll days, or that they do not primarily use marketable orders to do so.

The evidence also indicates that liquidity provision is more competitive on roll days. Quoted bid-ask spreads (the difference between the lowest limit price for unexecuted sell orders and the highest limit price for unexecuted buy orders) on roll days decline from an average of 1.17 basis points to 1.13 basis points in the front contract, and from 1.45 basis points to 1.42 basis points in the second contract. Figure 3 displays average quoted spreads by minute for roll and non-roll days. The patterns indicate smaller intraday quoted spreads for the majority of minutes on roll days. These declines, while modest, are statistically significant, and must be evaluated in light of the substantial increase in liquidity demand, which might have been anticipated to widen spreads.

We also assess liquidity supply by computing the "depth" of unexecuted orders in the limit order book. In particular, we determine the total volume of unexecuted sell (ask depth) and buy (bid depth) orders at prices within four ticks of the most competitive prices. Bid depth for the front contract, which is relevant for those seeking to sell, increases from an average of 46.5 contracts on non-roll days to 51.9 contracts on roll days (t-statistic = 13.17), while ask depth (relevant for those seeking to buy) increases from an average of 44 contracts to 49.2 contracts (t-statistic = 14.60). Figures 4 display average bid and ask depths by minute for roll and non-roll days and support increased liquidity provision in front contracts throughout the day. However, both the bid and ask depth for the second contract decline modestly, by about 1 contract, on roll days.

Next, we assess average effective bid-ask spreads by minute, where the effective spread for a given trade is twice the excess of the trade price over the bid-ask midpoint for those trades initiated by buy market orders and twice the excess of the midpoint over the trade price for those trades initiated by sell market orders. Effective spreads differ from quoted spread for trades that are large enough that they do not fully execute at the quoted price. Effective spreads also reflect the effect of fully hidden or iceberg orders in the limit order book. The average effective spread for each minute is computed as the volume-weighted mean of effective spreads for individual trades. The results indicate modest reductions in effective spreads as well, from an average of 2.06 basis points to 1.96 basis points (t-statistic = -5.64) for the front contract, and from 2.39 basis points to 2.29 basis points for the second contract (t-statistic = -4.60). Figure 5 displays average effective spreads by minute for roll and non-roll days.

Finally, we use the CFTC data to assess the number of distinct accounts that supply liquidity on roll and non-roll days. An account is deemed to supply liquidity if at least one market sell order is executed against a buy limit order posted by the account for the front contract and at least one market buy order is executed against a sell limit order posted by the account for the second contract. We find that an average of 874 distinct accounts provide liquidity in the front contract on roll days, compared to 681 accounts on non-roll days (t-statistic = 2.23). For the second contract the number of liquidity-supplying increases modestly, from 153 on non-roll days to 167 on roll days.

We also use the Ansari-Bradley test to evaluate the hypothesis that the one-minute observations on each variable are drawn from the same statistical distribution on Roll and non-Roll days. The results indicate strong rejection (p-value < .0001) of the null hypothesis in favor of the alternative that the measures are more volatile on roll days for trading activity per minute, standardized trade imbalances, bid and ask-side depth, and effective bid-ask spreads. These results are consistent with the view that more liquidity is both demanded and supplied on roll days. Interestingly, the data do not indicate more volatility in quoted spreads on roll days.

These simple comparisons of liquidity measures across roll and non-roll days provide results consistent with the sunshine trading theory of Admati and Pfleiderer (1991), who predict that the announcement of an upcoming trade by a credibly non-informed trader will attract additional liquidity suppliers to market. In particular, despite a large increase in trading activity and trade imbalances on roll days, quoted spreads and effective spreads (which measure costs of trading) decline, while quantities of unexecuted orders on the relevant side of the limit order book and the number of distinct accounts providing liquidity both increase. The latter findings provide direct support for Admati and Pfliederer's (1991) prediction that preannouncement increases the size of the market by attracting liquidity providers.

IV. A Model of Trading by a Monopolist Predator

In this section, we present a simple model of trading around the time of a known liquidation. Our main finding is that the Brunnermeier and Pederson (JF, 2005) implication that predatory trading is disruptive to the market in that it causes prices to overshoot their long run equilibrium, and is costly to the party who must liquidate in that it reduces their sales proceeds, do not generalize to a resilient market. Two important features of the Brunnermeier and Pederson model are that the price impacts of trades are entirely permanent, and that the predator does not trade prior to the liquidation. We show here that relaxation of these assumptions can lead to the conclusion that the presence of predator (who we refer to by the more benign label "strategic trader") can improve market quality and liquidator proceeds. To make this point in the simplest manner possible we focus on a monopolist predator in the absence of uncertainty.

III.A. The Setting

The investor must liquidate a known quantity, Q_L. A strategic trader is aware of the liquidation, and can trade before, simultaneous with, or after the liquidation, and chooses transaction quantities to maximize profits.[15] Strategic trades sum to zero across periods; that is the strategic trader ultimately does not absorb the liquidators position. The liquidator's position, as well as order imbalances generated by the strategic trader, are absorbed by the limit order book, the dynamics of which are described in the next section.

III.B Trade Prices when the market is resilient.

The model of trade prices presented here is an extension of that presented in Chapter 15 of Hasbrouck (2007). The private information conveyed by trades is measured by a permanent price impact parameter, λ, and the security value evolves according to $V_t = V_0 + \lambda Q_t$, where q_i is signed order flow in period i and $Q_t = \sum_{i=1}^{t} q_i$ is the accumulated order flow since base period 0.

Trades also have temporary price impacts. The immediate temporary price impact, γq_i, is proportional to the signed order flow, reflecting that small orders execute at quotes that differ from security value and that larger orders walk up the limit order book. The temporary price impact of the time i trade potentially persists beyond time i, if the limit order book is not refilled instantaneously. Specifically, the trade price at time t depends on the time t order flow as well as prior order flow according to:

$$P_t = V_t + \gamma A_t = V_0 + \lambda Q_t + \gamma A_t \tag{1}$$

where $A_t = \sum_{j=0}^{t-1} \theta^j q_{t-j}$.

Defining $M_t = V_0 + \lambda Q_{t-1} + \gamma\theta A_{t-1}$ as the midpoint, we can also express the trade price as

$$P_t = M_t + (\lambda+\gamma)q_t.$$

Here, A_t is a weighted sum of orders from time 0 to t, and the parameter θ measures the resiliency of the market. If θ = 0 the market is completely resilient in the sense that the temporary price impact of the

[15] As noted Brunnermeier and Pederson (JF, 2005) focus on predator trades that occur simultaneous with or after the liquidation. Carlin, Lobo, and Viswanathan (2007) study a predator constrained to trade during the same time interval as the liquidator.

order at time t affects only the time t price. A completely resilient market requires that the limit order book refill instantaneously after an order is executed. If so, the midpoint equals the pre-trade value, and $P_t = V_{t-1} + (\lambda+\gamma)q_t$. If $0 < \theta < 1$ the limit order book takes time to refill, and the temporary effect of the time t order flow extends beyond time t. The temporary price impact is dampened as the limit order book is refilled, and more rapidly for smaller θ. When $\theta > 0$ the midpoint differs from underlying value as a function of recent order imbalances. If $\theta = 1$ the temporary impact is never reversed, and thus is indistinguishable from permanent impact. In this case $A_t = Q_t$ and $P_t = V_0 + (\lambda+\gamma)Q_t$.

III.C Market Prices and Outcomes

Trading occurs during three intervals: before, during, and after the investor's liquidation. We will interpret each interval as comprising a trading day. Let Q_p, Q_d, and Q_a denote net signed order flow (sum of liquidator and strategic trading) during the "pre", "during", and "after" days, respectively. For simplicity we consider the case where trading during each day is spread evenly across N periods.

With per-period signed order flow q constant across periods we can express $A_t = qY_t$, where $Y_t = \sum_{j=0}^{t-1} \theta^j$ is a parameter that captures the cumulative impact of current and prior trades in a less than perfectly resilient market. Define also $\overline{Y}_N = \sum_{t=1}^{N} Y_t / N$, which is the average of Y_t across periods 1 to N. Let V_1 denote the value of the security at the beginning of the "pre" day. Given these simplifications, the average trade price across the N periods of the "pre" day is:

$$\overline{P}_p = V_1 + I_0 Q_p, \text{ where } I_0 = \left[\lambda \frac{(N+1)}{2N} + \frac{\overline{Y}_N \gamma}{N} \right] \tag{2}$$

The average trade price for the "during" day is:

$$\overline{P}_d = V_1 + I_1 Q_p + I_0 Q_d, \text{ where } I_1 = \left[\lambda + \frac{Y_N^2 \gamma \theta}{N^2} \right] \tag{3}$$

And the average trade price for the "after" day is:

$$\overline{P}_a = V_1 + I_2 Q_p + I_1 Q_d, + I_0 Q_a, \text{ where } I_2 = \left[\lambda + \frac{Y_N^2 \gamma \theta^{(N+1)}}{N^2} \right] \tag{4}$$

It may be useful to consider how the I_0, I_1 and I_2 parameters, which measure the effects on average trade prices within each day of current day, prior day and second prior day order imbalances, are affected by the resiliency parameter, θ. When θ = 1 we have $I_1 = I_2 = (\lambda+\gamma)$, and $I_0 = \left[[\gamma + \lambda] \frac{(N+1)}{2N} \right]$. Note that in this case the temporary and permanent price impacts are indistinguishable, and the full price impact (λ+γ) of the current day order imbalance carries over to subsequent days. The effect of the current day order imbalance on the same day mean price is smaller by the factor (N+1)/2N, which approaches ½ as the number of periods in the day becomes large. The smaller impact on the average current day price reflects that the full impact is not manifest until trading is completed at the end of the day.

In contrast, when θ < 0, so that the market is resilient, the effects of trades on average prices within each day is reduced, and more so the smaller is θ, i.e. the more resilient is the market. In particular, when θ = 0 we have $I_1 = I_2 = \lambda$, and $I_0 = \left[\lambda \frac{(N+1)}{2N} \right] + \frac{\gamma}{N}$. In the fully resilient market price impacts are smaller, with the reduction attributable to a reduced effect on average prices of the temporary price impact parameter, γ. In the fully resilient market the temporary price impact has no effect beyond the current day. Within the same day, the effect of the order imbalance of the average price is reduced from (N+1)/2N in the non-resilient market to 1/N in the fully resilient market. With 0 < θ < 1, i.e., when the market is less-than-perfectly resilient, the effect of trades on average daily prices are larger than the case when θ = 0, but smaller than the case when θ =1.

Let Q_L denote the liquidator sales on the "during" day. We can describe the strategic trader's order flows by a pair of proportionality parameters ρ_d and ρ_p, defined so that positive values indicate trading in the same direction as the liquidator. Including the requirement that the strategic order trader order flow sums to zero across the three days, the imbalance absorbed by the limit order book (the sum of liquidator and strategic order flow) each day is:

$$Q_p = -\rho_p Q_L \;, \tag{5}$$

$$Q_d = -(1+\rho_d)Q_L \,, \tag{6}$$

and $Q_a = Q_L(\rho_d + \rho_p)$. (7)

Given these assumptions, market outcomes depend on average daily prices. In particular, the liquidator's proceeds depend only on the average price during the liquidation day:

$$LP = Q_L \overline{P}_d \,. \tag{8}$$

The quantity liquidated is ultimately absorbed by the limit order book. The total acquisition cost paid by limit order traders is:

$$AC = Q_L \rho_p \overline{P}_p + Q_L(1-\rho_d)\overline{P}_d - Q_L(\rho_p + \rho_d)\overline{P}_a \,. \tag{9}$$

The strategic trader's profits as a function of average daily prices can be stated as:

$$SP = Q_L\left[\rho_p(\overline{P}_p - \overline{P}_a) + \rho_d(\overline{P}_d - \overline{P}_a)\right] , \tag{10}$$

which reveals that the strategic profits are driven by differences in average trade prices across days. Using expressions (5) to (7) in expressions (2) to (4) for prices and substituting into (10), strategic profits can also be stated in terms of trading and price impact parameters as:

$$SP = Q_L^2\left[\rho_p^2(I_2 - 2I_0) + \rho_d^2(I_1 - 2I_0) + \rho_p\rho_d(I_2 - 2I_0) + I_1\rho_p + (I_1 - I_0)\rho_d\right] \tag{11}$$

Straightforward calculus reveals that strategic trader profits are maximized when:

$$\rho_d^* = \frac{2I_0 - I_1}{4I_1 - 6I_0 - I_2} \tag{12}$$

and

$$\rho_p^* = \frac{-I_1}{2(I_2 - 2I_0)} + \frac{I_1 - 2I_0}{8I_1 - 12I_0 - 2I_2} \tag{13}$$

For comparison to the results of Brunnermeier and Pedersen, we also consider the case where the strategic trader does not transact on the "pre" day. Then (11) is restated as:

$$SP = Q_L^2\left[\rho_d^2(I_1 - 2I_0) + (I_1 - I_0)\rho_d\right] \tag{11A}$$

and the restricted first order condition for constrained profit maximization is:

$$\rho_d^* = \frac{I_0 - I_1}{2(I_1 - 2I_0)} \tag{12A}$$

III.D Illustration of the Outcomes

Table 4 illustrates the outcomes of this analysis when the strategic trader chooses quantities to maximize profit. The illustration includes an initial price (V_0) equal to \$100, N = 32 fifteen minute periods within each eight-hour trading day, Q_L = 20 units liquidated, temporary price impact, $\gamma = 0.5$, permanent price impact, $\lambda = 0.015$, for values of the resiliency parameter, θ, ranging from zero to one. These parameters imply a permanent price impact of 20*.015/100 = 30 basis points from the liquidation, which is in line with multi-day estimates of EFT roll price impacts reported in Section VI below. The temporary price impact if the full 20 units were brought to market in a single period would be 20*0.5 = 10%. Of course the liquidation would actually spread over at least the 32 periods of the "during" day.

The column of Table 4 labeled "base" reports outcomes, including both the liquidator's proceeds (LP, from expression 8) and limit order traders' acquisition cost (AC, from expression 9) for the 20 units, when the strategic trader is absent. In the absence of *any* price impacts the liquidator proceeds would be 20*100 = 2000. Given the positive price impact parameters, liquidator proceeds in the absence of strategic trading declines with θ, i.e. as the market is less resilient, from \$1990.7 when $\theta = 0$ to \$1893.8 when $\theta = 1$.

The next set of columns report outcomes when the strategic trader selects quantities to maximize profits, according to expressions (12) and (13), followed by columns that report outcomes when the strategic trader maximizes profits subject to a no-trading-prior constraint (expression 12A). The columns labeled "Pre", "During", and "After" report the profit maximizing strategic trader order flow as a proportion of the liquidator sale, with positive values indicating trading in the same direction as the liquidator. The column "SP" reports the strategic traders profit from expression (10).

Despite our use of simplifying assumptions (in particular, that trading is level across periods within each day), our results confirm those of Brunnermeier and Pedersen when we assume, as they did, that trades' price impacts are fully permanent ($\theta = 1$). When constrained to not trade prior to the liquidation, the strategic trader sells simultaneous with the liquidator, and reverses by buying the following day. The profit maximizing quantity transacted by the strategic trader is, given these parameters, 7.75 times the size of the liquidation itself. Consistent with Brunnermeier and Pedersen, the liquidator's proceeds are reduced dramatically from $1893 in the base case to $1071 when the strategic trader is present.

The consistency of results with those of Brunnermeier and Pedersen hinge crucially on the permanence of the price impacts. Reducing the θ parameter, i.e. introducing resiliency to the market, rapidly reduces the extent to which the profit-maximizing strategic trader sells alongside the liquidator. When $\theta = 0.96$, as well as for all smaller values of θ, the profit-maximizing strategic trader *buys* during the period when the liquidator sells, thereby absorbing a portion of the liquidator's order imbalance. Further, for the same set of parameter values, the liquidator's proceeds are *larger* when the strategic trader is present than when the strategic trader is absent.

We next assess outcomes when the strategic trader can also transact before the liquidator. Here, we see that the strategic trader acts as a liquidity supplier, in the sense that she buys during the day when the liquidator sells, for all values of θ. Depending on market resiliency, the strategic trader absorbs between 33% and 46% of the liquidator's sales in the "during" interval. In addition, the strategic trader follows a strategy that casual observers might refer to as frontrunning, in the sense that she sells during the period preceding the liquidator's sale for all for all values of θ. However, the liquidator receives higher proceeds in the presence of the strategic trader than in the base case with no strategic trading for all values of θ less than 0.94. The strategic trader's actions reduce liquidator proceeds only when θ approaches one, i.e. as price impacts become permanent.

Figures 6A and 6B display period-by-period trade prices around the 20 unit liquidation, for market resiliency parameters, θ, of 0.0 and 0.98, respectively. Outcomes for intermediate resiliency parameters are generally similar – only when θ approaches 1 are the effects altered. The first third of the individual observations pertain to the pre-liquidation day, the second third pertain to the during-liquidation day, and the final third pertain to the post-liquidation day.

Figure 6A applies when the market is fully resilient ($\theta = 0$), so that the period t price deviates from value only due to contemporaneous order flow. The value of the liquidated asset decreases from \$100 to \$99.70, due to the permanent price impact of the liquidation. Here, in the absence of strategic trading, the price remains at 100 until the liquidation begins. The price drops immediately to \$99.68 due to temporary price impact as the liquidation begins, and continues to decline to \$99.39 due to accumulated permanent price impact as the liquidation is completed. After the liquidation is completed the price returns to its new equilibrium of \$99.70.

Notably, strategic trading *reduces* price impacts relative to the no-strategic trading benchmark. This occurs because the strategic trader absorbs a portion of the sales during the liquidation, and shifts the selling pressure to the preceding day. As a consequence overall price impact is reduced. The minimum price with strategic trading is \$99.47 (as the liquidation ends) compared to a minimum price of \$99.39 without strategic trading.

Figure 6B applies when $\theta = 0.98$. Here, 98% of the temporary price impact in a given period persists to the next period. Overall price impacts are greater in the less-resilient market. In the absence of strategic trading the price is pushed as low as \$92.27, before recovering. In contrast, in the presence of the strategic trader the minimum price is \$92.75. In general, as the market becomes less resilient (as θ increases) from over the range $\theta = 0.00$ to $\theta = 0.98$ the price impact of the liquidation increases, but is always smaller with strategic trading than without. Further, over the range $\theta = 0.00$ to $\theta = 0.96$ liquidator proceeds are larger with strategic trading than without. Only when θ approaches 1, i.e. as the market loses all resiliency, does strategic trading lead to larger price impacts or reduced liquidator proceeds.

The results here therefore confirm those of Brunnermeier and Pedersen, but show that their central results are specific to their assumed structure. In particular, the strategic trader causes the price to overshoot its long term equilibrium and reduces the liquidators proceeds *only* when the market is non-resilient, i.e. when trades price impacts are close to permanent. In contrast, in a resilient market the profit maximizing strategic trader will choose to absorb a portion of the liquidator's order imbalance during the liquidation period, and thereby reduces the price impact of the liquidator's order imbalance.

These insights are relevant to ETF crude oil rolls. There is little reason to think that such rolls convey private information that would lead to permanent changes in crude oil prices. Further, the crude oil futures markets are very active: on roll days an average of 850 front-month contracts trade each minute. To the extent that the limit order book refills quickly after order imbalances, i.e. the market is resilient, the liquidator benefits from the presence of a strategic trader who anticipates the liquidation. We next report direct estimates of the resiliency of the crude oil market.

V. Estimating the Resiliency of the NYMEX Crude Oil Market

The analysis is the preceding section underscores that the incentives of a predatory or strategic trader who becomes aware of the trading intentions of another large investor depend crucially on whether the price impact of trades is permanent or temporary, and in the case of temporary price impacts, on the resiliency of the market. To assess the incentives and likely market impacts of strategic trading around ETF rolls therefore requires estimates of trades' permanent and temporary price impacts as well as market resiliency in the NYMEX crude oil markets.

Based on expression (1), which quantifies trade prices as a function of current and prior order imbalances, we estimate these parameters with geometric lag regressions of the form:

$$P_t - M_1 = \alpha + \gamma \sum_{j=t-k}^{t} \theta^{t-j} q_j + \lambda \sum_{j=t-k}^{t} q_j^* + \varepsilon_t \qquad (13)$$

where P_t is the time t trade price, M_1 is the quote midpoint at the beginning of trading day (9:00 A.M. New York time), q_j is the signed order imbalance at time j, and q^*_j is the residual from a fifth-order auto-regression of q_j. The use of q^*_j instead of q_j to estimate the permanent price impact of order imbalances follows Madhavan, Richardson, and Roomans (1997) and Huang and Stoll (1997), and accommodates the presence of positive autocorrelation in order imbalances. The specific-choice of the fifth-order auto-regression specification follows Sadka (2006). We implement (13) using time periods defined as either one or five seconds. For time periods with multiple trades P_t is measured as the last trade price during the time period, and Q_j is measured as the net order imbalance (excess of buyer-initiated over seller-initiated trades, measured in contracts) during the interval. The geometric lag expression (equation 13) is implemented using NYMEX order data by Generalized Method of Moments (GMM), using SAS Proc Model with a Bartlett Kernel set equal to the lag length plus one.

Table 4, Panel A presents the regression parameters estimated using five second time intervals, and using sixty lags of signed order imbalance. We report results for the front and second-to-maturity contracts, and for the full sample as well as separately for roll and non-roll days. While we are primarily interested in assessing the resiliency of the crude oil futures markets, it is also of interest to compare parameter estimates across roll and non-roll days. Consistent with results reported on Table 1 and Table 3, the coefficient estimates based on full sample of roll and non-roll days indicate that the front contract is more liquid than second contract. In particular, the front contract has smaller permanent price impact (5.6 cents versus 8.4 cents), smaller temporary price impact (2.1 cents versus 5.7 cents) and is more resilient (0.959 versus 0.976).

The parameter estimates for roll and non-roll days provide some evidence of market stress on roll days due to heightened liquidity demand. The estimated temporary price impact is larger on roll days (0.038 vs. 0.020 for the front contract), as is the estimated resiliency parameter (0.986 vs. 0.959 for the front contract). However, the standard errors of the estimates are sufficiently large that differences in estimates across roll and non-roll days are not statistically significant.[16]

The estimated permanent price impact for trades in the front contract is lower on roll days than on non-roll days (5.4 cents versus 5.8 cents), and the difference is statistically significant (p-value = 0.03).[17] A decrease in permanent price impact is consistent with the reasoning that ETF roll trades are viewed as

[16] We also estimate expression (13) over one-second rather than five-second time intervals, with results that are generally similar to those discussed (Table 5, Panel B). Consistent with results of the 5-second specification, the estimated temporary price impact for the front contract increases (p-value of difference = 0.12) and market resiliency declines (the θ estimate increases) on roll days (p-value of difference = 0.07).

[17] Note that the price impact is estimated based on the surprise order imbalance rather than the level of the order imbalance. The ETF rolls studied here involve an average of over 20,000 contracts, but the large size of this imbalance is not a surprise to market participants. For this reason the estimated permanent price impact does not map directly into the Section IV theory, which did not allow for predictable order imbalances.

uninformed. It may be surprising that the estimated permanent price impact on roll days is positive, if ETF rolls comprise non-informed trading. However, while ETF trades are large, they still comprise a minority of trading activity on roll days, and informed traders may be present. The model of Admati and Pfleiderer (1991) implies that informed traders may in fact prefer to trade at times when liquidity traders are also more active, in order to camouflage their trades and take advantage of reduced price impacts. Consistent with the reasoning, the estimated permanent price impact of trades in the second-nearest-to-delivery contract is actually larger (0.149 vs. 0.082) on roll than non-roll days.

Recall that the resiliency parameter, θ, measures the proportion of the temporary price impact attributable to current interval order imbalance that persists to the next period. The analysis in Section IV shows that the resiliency parameter is crucial in assessing the trading patterns that will maximize strategic trader profits. The estimates of θ reported on Table V uniformly exceed 0.95. At first glance these estimates may appear to imply a non-resilient market. However, the estimates must be interpreted in light of the time interval that defines a period. We estimate expression (13) over short five second intervals to increase the effective sample size. Focusing on the front-month resiliency estimate of 0.959, for example, the estimated proportion of the temporary impact that persists after one minute is 0.959^{12} = .605, the proportion that persists after five minutes is 0.959^{60} = .081, the proportion that persists after ten minutes is 0.959^{120} = .007, and the proportion that persists after 15 minutes is 0.959^{180} = .0005. Stated alternately, these estimates imply that over 90% of the temporary price impact caused by an order imbalance is reversed within five minutes, and 99.95% is reversed within fifteen minutes, indicating that the crude oil futures markets are indeed quite resilient.

Among the θ estimates reported on Table 5, the one that implies the *least* resilient market is the 0.993 estimate obtained for the second-nearest-to-expiration contract on roll days. A resiliency estimate of 0.993 for five second intervals implies that less than 30% (0.993^{180} = 0.283) of the temporary price impact persists for fifteen minutes. The numerical illustrations of the theoretical analysis presented in Section IV are calibrated to a fifteen minute period. The empirical estimates obtained here therefore indicate that resiliency parameters of 0.3 or less at a fifteen minute interval are relevant in evaluating strategic trader incentives.

VI. Account-level analysis

The theoretical model presented in Section IV shows that the profit-maximizing strategy for a trader who is aware of a large upcoming transaction depends on the resiliency of the underlying market. The empirical results reported in Section V indicate that the oil futures market is indeed resilient, and that the analysis should focus on resiliency parameters θ of 0.3 or less when evaluating trading strategies over fifteen minute periods.

We next report results obtained when we rely on the CFTC transaction data, which includes unique trading accounts identifiers, to assess strategies followed by individual trading accounts in the period surrounding ETF rolls. The CFTC data identifies the trading accounts associated with both the buy and sell side of each transaction. For this analysis we define the “During” interval as 9 A.M. to market closure at 4:15 P.M. New York time on the ETF roll day, the “Before” interval as midnight on Day -3 (three trading days prior) to 9 A.M. on roll day, and the “After” interval as the reopening of trading at 5 P.M. on the roll day through Midnight on Day +3 (three trading days after) the roll day.

The theoretical analysis focuses on a strategic trader who does not ultimately absorb the liquidator’s position. Since strategic traders might also transact for reasons unrelated to the roll we relax this definition for the empirical implementation. Specifically, we identify an account as potentially being a strategic trader if the absolute value of the net change in the account’s inventory as a fraction of the account’s total activity in the days surrounding the roll is less than 0.25. The ETF's natural counterparties (i.e., accounts that hold or can be induced by price concessions to hold opposite positions as the ETF) are unlikely to be classified as strategic traders, since their inventory change to total trading ratio is likely to exceed the threshold of 0.25.[18]

We categorize each account identified as a possible strategic trader as following one of twelve possible trading strategies, as described in Panel A of Table 6. Those strategic traders whose signed position change on the roll day itself (Day 0) is of the opposite sign as the ETF's order flow are deemed to be effectively following liquidity provision strategies, labeled ST1 to ST5. Those strategic traders whose signed position change on the roll day is in the same direction as the ETF order flow are, consistent with the analysis of Brunnermeier and Pedersen (2005) deemed to be following potentially predatory strategies, labeled ST8 to ST12.[19] The categories ST6 and ST7 represent accounts without any trading activity on the roll day.

We further categorized strategic traders into one of five sub-strategies within liquidity provision (ST1-ST5) and predatory trading (ST8-ST12) categories on the basis of the account's change in net positions in the before and after periods. For example, a liquidity providing account that trades in the opposite (same) direction as the ETF in the before interval and who also trades in the same (opposite) direction as the ETF in the after interval is placed in categories ST1 (ST3). Our objective in identifying the sub-categories is to assess the relative importance of strategies implied by various theories. For example, the Brunnermeier and Pederson (2005) predatory trading strategy involves no trading before the ETF, trading the same direction as the ETF during, and trading opposite the ETF in the after period, which we label ST11. Our analysis in Section IV above implies that the strategic trader will trade in the same direction as the ETF before the roll begins and in the opposite direction as the ETF during the roll (ST3, ST4, and ST5). Further, for resiliency parameters, θ, less than about 0.3 the strategic trader will trade the opposite direction as the ETF during the after period as well (ST3).

[18] As an illustration, suppose an account sells 1,000 contracts before, sells 1,000 contracts during, and buys 1,500 contracts after the roll. The absolute value of net change in account's inventory is 500 contracts, while total trading activity is 4,500 contracts. Since the ratio of (500/4500) < 0.11, we classify the account as a strategic trader.

[19] Specifically, for the front contract, a strategic trader whose inventories increase (i.e., net buyer) on the roll day is categorized as a liquidity provider, while a strategic trader whose inventories decrease (i.e., net sell) on roll day is categorized as a predatory trader.

Having assigned each strategic account to one of the 12 strategies at each roll period, we aggregate the *strategic volume* across all trading accounts associated with a strategy, where the strategic volume is simply the round trip volume associated with an account surrounding the roll.[20] Note that for each identified strategy there is a complementary strategy involving opposite trading patterns. For example, ST1 and ST12 are complimentary, in that ST1 involves trading against, against, and with the ETF during the three intervals, while ST12 involves the opposite pattern of trading with, with, and against the ETF during the three intervals.

Since some strategies might be more common than others for reasons entirely unrelated to ETF rolls, we focus on *normalized strategic volume*, which is the strategic volume in a category less the strategic volume in the complementary strategy. Also, to assess whether particular strategies are used more frequently around ETF rolls as compared to other periods, we calculate normalized strategic volume for each of the strategies on all usable non-roll days. A non-roll day is usable for this purpose if the interval three days prior to three days after does not overlap with the three days prior or after an actual ETF roll day.

In Panel B of Table 6 we report regression coefficients obtained when normalized strategic volume for all days (ETF roll and non-roll days) is regressed on an indicator variable that equals one for ETF roll days, as well as separate indicator variable for February 6, 2009 roll. This date is of particular interest as it was the subject of a CFTC enforcement action related to the execution and reporting of a pair of large (over 29,000 contracts in each of the front month and second month crude oil contracts) block TAS trades.[21] It is therefore of interest to understand whether certain strategies were particularly prevalent on this date. The FEB6 coefficient measures abnormal trading activity on Feb 6, 2009 relative to the other ETF roll days. Columns (1) to (6) report results when the dependent variable is normalized strategic volume in categories ST1 to ST6.[22]

[20] Building on the illustration above, the strategic volume for an account that sells 1,000 contracts before, sells 1,000 contracts during, and buys 1,500 contracts after the roll is 1,500 contracts.

[21] http://www.cftc.gov/PressRoom/PressReleases/pr5816-10.

[22] Note that results for strategies ST7 to ST12 would simply be the opposite of results for ST1 to ST6, since they are the complements of the first six strategies.

To interpret the results of this analysis, recall that categories ST1 to ST6 involve liquidity provision, while their complimentary strategies ST7 to ST12 potentially involve predation. Therefore, increases in normalized strategic volume for categories ST1 to ST6 at the roll would indicate increase liquidity provision, while decreases in normalized strategic volume for categories ST1 to ST6 at the roll would indicate increased use of predatory strategies.

Several results reported on Panel B of Table 6 are noteworthy. First, the statistically significant intercepts in columns (3) and (4) for the front contract imply that strategic trading in category ST4 is more prevalent and in ST3 less prevalent relative to complementary strategies on non-roll days. Second, we estimate statistically significant positive coefficients for ST3 on the roll day indicator for both the front and second contracts, and significant (for the second contract) or nearly significant (for the front contract) negative coefficients for ST1 on the roll day. Since ST1 and ST3 are both liquidity providing strategies the results may appear to conflict. Note, though, that ST3 involves trading in the same direction as the upcoming ETF trades prior to the roll, while ST1 involves trading the opposite direction prior to the ETF roll. The analysis presented in Section IV and illustrated in Table 4 indicates that the profit maximizing strategy always involves (for any resiliency parameter) trading the opposite direction as the ETF in the period prior to the roll, as in ST3, ST4, and ST5. Further, Table 4 indicates that the most profitable strategy involves trading opposite the ETF in the period after the ETF roll if the market is resilient, with θ less than about 0.3. That is, for resiliency parameters in line with the empirical estimates reported in the preceding section, ST3 is indeed the more profitable strategy. We therefore interpret these results as indicating a substitution by strategic traders toward the most profitable liquidity provision strategies around the time of the ETF roll.

We estimate a negative and significant roll day coefficient for ST6, indicating increased usage of the complimentary strategy ST7, which involves trading the same direction as the ETF prior to the roll and reversing the position after the roll. This strategy might also be interpreted as predatory. However, the result is not observed for the second contract, which would imply that the predator was taking on more risk than necessary (since an outright rather than a spread position would be held for at least a trading day). Further, the alternative strategy of absorbing some of the ETF imbalance would have been more profitable in general, and the February 6 indicator variable provides an opposite coefficient estimate.

With regard to other estimated coefficients for February 6, only the coefficient for ST2, which is a liquidity provision strategy, is significant, and only in the second contract. This analysis therefore does not provide evidence of disruptive predatory strategies around the controversial February 6, 2009 roll.

The results presented here indicate that the NYMEX oil futures market is resilient, with temporary price impacts of trades being fully reversed within minutes of an order imbalance shock. Evidence based on limit order book depth, quoted spreads, and effective spreads, are also consistent with increased competition from liquidity providers on roll date. Further, the results of this analysis of trading by individual accounts also points towards increased liquidity provision around the roll.

Estimates reported in Table V indicate that order imbalances do have a permanent price impact, even around the roll. A strategic trader who provides liquidity by absorbing order imbalances during the roll must offset the inventory (by trading in the same direction as the ETF roll) either before or after the roll. Given long-lived price impacts it is more profitable to conduct offsetting trades ahead of (as in ST3 or ST4) than after (as in ST1 or ST2) the roll. Our analysis implies that ST3 will be the most profitable strategy, and the evidence supports that strategic traders indeed provide liquidity at the roll while using this strategy.[23] Importantly, our analysis also indicates that the temporary price impacts associated with the large trades are reduced and the liquidator (ETF) proceeds are enhanced by strategic trading of this type. This stands in sharp contrast to the implications of the Brunnermeier and Pedersen (2005) model, which would have predicted the use of strategies ST8 to ST12, larger price impacts, and decreased liquidator (ETF) proceeds.

VII. Estimating the Cost of Executing the Roll

[23] To expand on this reasoning, note that the ETFs (or equivalently, the TAS counterparties of ETF) actively sell the front contract and actively buy the second month on the roll day. The positive permanent price effect implies that ETF trading would cause front contract prices to decrease and second contract prices to increase on the roll day. For this reason, strategic liquidity providers in front month contracts are better off building a short position (at higher prices) before the roll day. Along similar lines, strategic liquidity providers in the second month contracts are better off building a long position (at lower prices) before the roll day.

As noted, the ETFs completed (with one minor exception) all of their roll transactions in the March 2008 to February 2009 period at weighted average prices equal to the NYMEX settlement prices for the same contracts on the same day. However, this result does not imply that the ETFs incurred no cost to complete the rolls, as the settlement prices themselves may have been affected by the offsetting trades of the ETF counterparties or by strategic traders.

The roll transaction most often consists of the sale of the nearest to expiration contract and the purchase of the 2^{nd} nearest to expiration contract. It is common in the literature on trading costs to estimate the cost of executing a trade by comparing the transaction price to a pre-trade benchmark price. Let F_{1T} denote the transaction price at which the 1^{st} nearby contract is sold and F_{1B} denote the benchmark price for the nearby contract. Similarly, let F_{2T} and F_{2B} denote the price at which the second nearby contract is purchased and the benchmark price for the second nearby contract, respectively. The proportional execution cost for the sale of the nearby contract can be measured as $\ln(F_{1B}/F_{1T})$, while the proportional execution cost for the purchase of the second nearby contract can be measured as $\ln(F_{2T}/F_{2B})$. Summing the two proportional cost measures, the combined cost can be measured as:

Proportional roll cost = $S_T - S_B$,

where $S_T = \ln(F_{2T}/F_{1T})$ and $S_B = \ln(F_{2B}/F_{1B})$. The notation S reflects that the log ratio of the second nearest futures price to the nearest futures price can be viewed as the slope of the futures term structure. The implication is that the proportional cost of the roll can be measured by the futures term slope implied by the transaction prices in excess of the future term slope implied by the benchmark prices. Since this analysis requires only daily settlement prices, we extend the time interval studied to include the interval April 2006 to October 2011, where the former date coincides with the launching of USO, the largest ETF designed to track crude oil prices.

We identify from the USO web site the date of each USO roll during this period, and report on Table 7 estimated roll costs. Estimates are based on comparisons of the term slope implied by settlement prices on the USO roll date to the term spread implied by settlement prices ranging from one to ten days prior to the USO roll.[24] The daily record of settlement prices is obtained from CRB database. Estimated rolls costs vary from 10 basis points (prior day benchmark) to approximately 30 basis points (benchmark 7 to 9 days prior to the roll). While the sample size is relatively small and the estimates are noisy, cost estimates based on benchmarks from five to seven days prior to the roll differ significantly from zero (t-statistics range from 2.15 to 2.58).

These estimates of ETF roll trading costs are not particularly large as compared to estimated costs for institutional trades in U.S. equities in recent years. For example, Amber, Irvine, Puckett and Venkataraman (2011) report that *one-way* institutional trading costs in large-capitalization equities average around 15 to 20 basis points in 2008-09. The average institutional order size in their sample is around 1.5% of average daily volume (ADV) in the stock, while the ETF roll activity as percentage of ADV averages (Table 1) 13% for the front month and 18% for the second month. In addition, the estimated ETF roll cost reported here reflects a round-trip trade, as the ETF sells the nearby contract and buys the second nearby contract. These estimates indicate that, while the ETFs do effectively pay non-trivial execution costs to have their large trades executed, the costs are not large relative to institutional trading costs in large-capitalization equities. The modest trade costs for the large ETF rolls reflect the resiliency of the crude oil futures market, and the effectiveness of the "sunshine trading" strategy where preannouncement attracts liquidity suppliers, including strategic traders, as well as natural counterparties.

VIII. Explaining Crude Oil ETF Stock Price Performance

In general the analysis reported here provides little evidence that strategic traders engage in predatory trading of the type described by Brunnermeier and Pedersen (2005) around crude oil ETF rolls. This result is consistent with the theoretical analysis presented here implying that strategic traders have incentives to trade in a more benign manner in a resilient market, relative to a market with purely

[24] Prior to March 2009, the USO roll was accomplished on a single date two weeks prior to the expiration of the nearest-to-delivery contract. From March 2009 the USO roll has been spread over four days, beginning two weeks prior to the expiration of the nearby contract (source: USO webpage). For this analysis we compute S_T for the single roll date prior to March 2009, and we average across terms slopes observed for the four roll dates thereafter.

permanent price impacts as considered by Brunnermeier and Pedersen. It is also consistent with the reasoning that the pre-announcement of the roll attracts additional liquidity providers, as in the sunshine trading theory of Admati and Pfleiderer (1991). What then explains the poor performance of USO stock (the largest of the crude oil ETFs) compared to the level of crude oil prices, as documented in Figure 1?

The following insights apply. First, USO and other crude oil ETFs were not designed to deliver returns that track changes in the level of crude oil prices, either spot or future, but changes in prices of individual futures contracts. A recent USO "Fact Sheet" indicates that USO's investment objective is to deliver returns that reflect "changes in percentage terms of ….the futures contract on light, sweet crude oil traded on the New York Mercantile Exchange that is the near month contract to expire, less USO's expenses." Second, investors who hold spot crude oil do not earn returns that match the change in spot prices, as they incur costs of storing and insuring the crude oil. Indeed, crude oil ETFs take positions in futures contracts rather than holding spot crude oil to avoid incurring such storage costs. In contrast, ETFs designed to deliver returns that depend on the prices of commodities with lower storage costs (e.g. precious metals) often hold physical inventories.[25]

We formalize these arguments and quantify their relative importance as follows. Let P_t denote the date t spot price and $F_t(m)$ denote the date t futures price for delivery at date t+m. Futures prices are linked to spot prices by the no-arbitrage "cost-of-carry" relation:

$$F_t(m) = P_t e^{S_t m}, \qquad (14)$$

where $S_t = r_t + c_t$ is the continuously compounded per-period cost of carrying inventory, including forgone interest, r_t, and other storage costs, c_t. Non-interest storage costs include costs of renting storage tanks, insurance, etc., and may at times be offset in part or full by "convenience yields" that reflect the option value of holding inventory. Applying (14) to futures for delivery at dates t+m and t+n, the per-period cost of carrying inventory can be inferred as:

[25] For example, the Wall Street Journal recently reported that the SPDR Gold Shares ETF holds over 41 million troy ounces of physical gold inventory. See http://blogs.wsj.com/marketbeat/2012/03/02/etfs-hold-more-gold-than-italy-france/.

$$S_t = \frac{\ln(F_t(m)/F_t(n))}{(m-n)} \tag{15}$$

Expression (15) implies that the cost of carrying inventory is revealed by the slope of the futures term structure. Using expression (15) with the daily CRB data, we compute the cost of storage implied by the settlement prices of the first and second nearest-to-expiration crude oil contracts for each trading day from January 1, 1999 to October 20, 2011. We report results on Table 8, for the full sample and for subsamples. We focus in particular on the subsample beginning April 10, 2006 (when USO was launched). For the full sample the mean implied storage cost (multiplied by 250 to convert to an annual equivalent) is 0.33%, which an associated t-statistic of 0.85. In contrast, during the post-USO subsample the mean implied storage cost was 16.20%, with an associated t-statistic of 21.73.

A positive term slope, whereby futures prices exceed spot prices, and more so for more distant delivery dates, characterize what practitioners often refer to as a "contango" market. The cost-of-carry relation implies that contango will be observed only when net storage costs for the marginal holder of inventory are positive.[26] To see that the positive term slope represents marginal storage costs, recognize that S_t also represents the pre-storage-cost daily return to a strategy of purchasing crude oil for delivery at date t+n at price $F_t(n)$ and simultaneously selling the same oil for delivery at date t+m at price $F_t(m)$. Positive arbitrage profits are available if oil can be stored from date t+n to date t+m for a per-period cost less than S_t.[27]

As of April 10, 2006, the nearest-to-expiration futures price (for May delivery) was $68.74, and the associated spot price implied by expression (14) was $68.31. By October 20, 2011, the nearest to expiration futures price (for November delivery) had risen to $85.30. Despite the increase in oil prices, USO's share price decreased from $68.02 on April 10, 2006 to $33.33 on October 20, 2011. To assess the reasons for this decrease, first define:

---

[26] Pirrong (2011) documents that the collapse in crude oil prices during the recent financial crisis was accompanied by large increases in physical crude oil inventories and in the marginal cost of carrying inventory.

[27] The cost-of-carry relation is a no-arbitrage condition for the *marginal* holder of inventory. Those who can store a commodity for lower cost can earn profits. Anecdotal accounts (e.g. http://blogs.reuters.com/great-debate/2010/07/22/contango-and-the-real-cost-of-carry/) indicate entry by non-traditional firms (e.g. hedge funds) into the oil storage business in recent periods.

$$U_{t+1} = \ln\left[\frac{P_{t+1}}{P_t e^{S_t}}\right] . \tag{16}$$

The denominator of expression (16) is the time t expectation of the time t+1 spot price, assuming the absence of any return premium, i.e. if the spot price is expected to appreciate enough to offset the marginal cost of storage. U_{t+1} can therefore be interpreted as the ex post return premium in the day t+1 spot price, comprised of the ex ante premium (if any) and the ex post price shock. (If the only cost of storage was foregone interest then U_{t+1} would simply be the return in excess of the interest rate). We construct a daily time series of spot prices implied by expression (14), relying on the nearest to expiration futures price and the previously computed daily storage cost estimates. From this series of implied spot prices we compute the time series of realizations of U_t, the daily ex post return premium in spot oil prices.

Table 7 reports the resulting time series mean of U_t (annualized by multiplying by 250). For the full sample, the ex post spot premium is 5.62% per year. During the 1990s, the ex post spot premium was only 4.69% per year. In the period before USO was launched, January 2000 through April 9, 2006, the ex post spot premium surged to 23.02% per year. In contrast, after April 10, 2006, when USO was active, the mean ex post spot return premium was -12.22% per year. While none of these means are statistically significant (reflecting the high variability of spot price changes) the accumulated effect is nevertheless economically important. The negative ex post premium for the period when USO was active implies that the appreciation of oil spot oil prices during the sample period was *less* than sufficient to compensate for the cost of carrying inventory.

To assess the performance of futures returns in comparison, apply expression (14) to futures at dates t and t+1 to obtain the following expression for the continuously compounded daily return on a given futures contract:

$$\ln\left[\frac{F_{t+1}(m-1)}{F_t(m)}\right] = \ln\left[\frac{P_{t+1} e^{S_{t+1}(m-1)}}{P_t e^{S(m)}}\right] \tag{17}$$

Using (16), and defining $\Delta S = S_{t+1} - S_t$, the futures return can be expressed as:

$$\ln\left[\frac{F_{t+1}(m-1)}{F_t(m)}\right] = U_t + (m-1)\Delta S \quad . \tag{18}$$

From (16), the continuously compounded growth in the spot price from date t to t+1 can be written as:

$$\ln\left[\frac{P_{t+1}}{P_t}\right] = U_t + S_t \quad . \tag{19}$$

Comparing expressions (18) and (19) yields several insights. First, and most important, for a given cost of carry ($\Delta S = 0$), the rate of appreciation in the spot price exceeds that of the futures price by S_t, the cost of carrying inventory. Stated alternately, spot price appreciation will exceed changes in prices of individual futures contracts in Contango markets, and vice versa in "backwardated" markets (where the implied cost of carry is negative, presumably due to large "convenience yields"). As noted, the marginal cost of carry was large and significant (16.2% per year) during the USO sample period. Large underperformance of futures price changes relative to spot price changes is therefore implied by the no-arbitrage cost-of-carry relation.

Second, the cost of carry has no direct implication for futures returns, as S_t does not appear in expression (18).[28] In particular, the fact that the rolling a futures position in a contango market involves buying the second contract at a price higher than the selling price for the expiring contract has no direct implication for returns to roll strategies.[29] Third, the futures return does depend on ΔS. For a given a spot price, increases in the cost of carry improve futures returns. Finally, both futures and spot returns are equally affected by U_t, the ex post premium in the spot price. The data indicate that the time series average of U has been negative in the period since USO was launched.

[28] However, expression (18) does not rule out negative covariation between the cost of carry and futures returns, which has in fact been documented in a number of commodity markets. See, for example, Liu and Tang (2010), Szymanowska, De Roon, Nijman, and Van den Goorbergh (2011), and the papers referenced therein. The data reported on Table 8 are consistent with such negative covariation, in that futures returns are positive and the cost of storage negative during the January 2000 to April 2006 period, while futures returns are negative and the cost of storage positive during the April 2006 to October 2011 period.

[29] In contrast, the financial press and some academic authors have incorrectly asserted that a roll trade in a contango market generates an immediate loss. For example, Mou (2011, page 13) claims that the excess of the nearest-delivery futures price over the more distant-delivery futures price "is the amount of gain (or loss) per unit of the commodity when rolling futures forward. "

That USO returns trailed the growth of spot prices is to be anticipated in light of expressions (18) and (19) above, which imply that futures price appreciation will be less than spot price appreciation in contango markets. But did excessive roll costs also contribute? To shed some additional light on this issue, we compute daily time series of futures returns, $\ln(F_{t+1}(m-1)/ F_t(m))$. The preceding expression applies to futures prices for a contract with a fixed delivery date, t+m. To compute a continuous time series of futures returns also requires "rolling" to more distant contracts as contracts mature.

We compute three time series of daily futures returns. The first, denoted "Return 1", is based on the current and prior day prices of the nearest-to-expiration contract, for all days including the last day of trading for the expiring contract. The second, denoted "Return 2", is based on the current and prior day prices of the second-nearest-to-expiration contract, for all days including the last day of trading for the expiring contract. The third, denoted "Return Benchmark" is based on the current and prior day prices of the nearest-to-expiration contract when the nearest contract is two or more weeks from expiration, and on the current and prior day prices of the second-nearest-to-expiration contract from then until the nearest contract expires. These benchmark returns reflect the timing of USO's rolls. If the USO roll adversely affects market prices at the time of the roll, e.g. due to predatory trading or due to imperfectly elastic liquidity supply, then the benchmark returns will be lower than the other two return series**.**

Table 8 reports mean (x250) daily returns on the three futures return series, for the full sample as well as for the USO period. For the latter interval we also report the mean continuously compounded return to the USO ETF. Each mean return is negative during the USO period.[30] However, the returns are highly variable, and mean returns do not differ significantly from zero (t-statistics range from -0.51 to -0.85). It is nevertheless instructive to compare the return series. The mean benchmark futures return is -14.11% per year. By comparison, the mean Return 1 is -8.94% per year and the mean Return 2 is

[30] Note though that the annualized mean of each of the three futures return series exceeded 20% during the preceding interval January 1, 2000 to April 9, 2006.

-9.71% per year. The benchmark return is 5.17% per year less than Return 1 and 4.40% less than Return 2, and the latter differential is statistically significant (t-statistic = -2.65).[31]

The prices that comprise the benchmark return are potentially affected by the fact that USO rolls its positions two weeks prior to expiration. Return 2 and the benchmark return are identical (since each reflects a position in the 2nd nearest to delivery contract) during the two weeks prior to expiration. The statistically significant differential in Return 2 relative to benchmark returns therefore reflects better performance of the 2nd nearest-to-expiration contract relative to the nearest-to-expiration contract, i.e. widening of the spread, in the half month preceding the USO roll. The differential of 4.40% per year equates to 37 basis points per roll, and comprises an estimate of the cost of completing the roll.[32] Analogously, the excess of Return 1 over the benchmark return reflects better performance of the nearest-to-expiration contract relative to the second-to-nearest expiration contract, i.e. narrowing of the spread, in the two weeks after the USO roll. This differential of 5.17% per year equates to 43 basis points per roll, and comprises an additional (though noisier) estimate of the roll cost.

The mean annualized return (daily mean x250) to the USO ETF during the sample was -12.79%. In contrast to the perception created by inspection of Figure 1, the USO ETF actually outperformed the -14.11% mean annual return on a crude oil futures benchmark return constructed to reflect changes in the prices of the futures positions held by USO. This outperformance reflects that USO was able to earn interest on margin positions.

We summarize this analysis as follows. The actual annualized appreciation in implied spot prices during the USO sample period April 10, 2006 to October 20, 2011 was 3.98%.[33] From expression (19), the mean spot appreciation can be decomposed into the sum of the mean cost-of-carry (S = 16.20%) and the mean ex post spot market return premium (U = -12.22%). From (18), the mean futures return equals

[31] That the former is not significant reflects the greater volatility of the Return 1 series, which in turn reflects in part the high volatility in the settlement price of the expiring contract on the final trading day.

[32] The same comparison during the 1/1/00 to 4/9/06 period gives an estimated roll cost of 4.57%. USO was not active during this period, but index funds tracking the Goldman Sachs Commodity Index were (see Mou (2011)). In contrast, the Return2 and Benchmark return means are almost identical during the 1990s, before index trading became popular.

[33] The appreciation was from $68.51 to $85.30 across 1393 trading days, implying daily mean continuously compounded growth of .016% or growth per 250 days of 3.98%.

the mean ex-post spot premium if the futures term slope is constant. The futures Return 2 series mean is -9.71% per year, which slightly exceeds the mean ex-post spot return premium due to variation in the cost of carry (ΔS) during the sample. The benchmark futures return of -14.11% is lower, which plausibly reflects temporary distortions of futures prices, i.e. trade execution costs, associated with the roll. Finally, USO's actual share price return of -12.79% per year slightly exceeds the benchmark futures return. Interestingly, USO's actual return, which incorporates the effect of management fees, interest, roll costs, and changes in the cost-of-carry, is very similar to the ex post spot crude oil return premium of -12.22% per year.

IX. Conclusions

In this paper, we study trading strategies, liquidity, and price patterns around the time of large and predictable monthly trades undertaken by eight exchange-traded-funds (ETFs) that are designed to provide returns that track crude oil prices. Aggregated across the twelve roll dates in our sample period, net roll activity comprised 13% of roll-day volume in the front contract and 18% of roll-day volume in the second contract. We assess the implications of the theory of predatory trading as well as the theory of sunshine trading, in a setting where both theories potentially apply.

Our study is distinguished from related work in that we are able to exploit individual trades with account identifiers associated with each crude oil futures transactions using data from the Commodity Futures Trading Commission (CFTC), as well as individual limit order book updates provided by the Chicago Mercantile Exchange. Our analysis reveals that the oil futures market is indeed resilient. We estimate that virtually all of the temporary price impact caused by an order imbalance is reversed within 15 minutes. However there is some evidence that the market is marginally less resilient and temporary price impacts are marginally larger on ETF roll days than non-Roll days.

We develop a simple model to demonstrate the economics of strategic trading in a resilient market. For the range of resiliency parameters that we estimate for the oil futures market, the model predicts that the strategic trader will choose to act as a liquidity provider by absorbing a portion of the liquidator's order imbalance on Roll day, while offloading the resulting inventory in periods before the Roll day. Further, for the same set of parameter values, the liquidator's proceeds are larger when the strategic trader is present than when the strategic trader is absent.

Evidence based on limit order book depth, quoted spreads, and effective spreads, are consistent with increased competition from liquidity providers on roll date. Furthermore, using unique trading account identifiers in the CFTC dataset, we categorize strategic traders into sub-strategies that are consistent with liquidity provision and predatory trading. Consistent with our model, we document a statistically significant increase in a liquidity provision strategy wherein traders provides liquidity to the ETF during the roll day and shifts the selling pressure to the preceding day.

Nonetheless, we estimate that ETFs effectively pay about 30 basis points to complete their roll trades. While the ETF roll associated execution costs are non-trivial, we note that the costs are not large relative to institutional trading costs in large-capitalization equities. The modest trade costs for the large ETF rolls reflect the resiliency of the crude oil futures market, and the effectiveness of the "sunshine trading" strategy where preannouncement attracts liquidity suppliers, including strategic traders, as well as natural counterparties.

Overall, we find little evidence that strategic traders engage in predatory trading of the type described by Brunnermeier and Pedersen (2005) around crude oil ETF rolls. Our results are consistent with the theoretical analysis presented here implying that strategic traders have incentives to trade in a more benign manner in a resilient market. It is also consistent with the reasoning that the pre-announcement of the roll attracts additional liquidity providers, as in the sunshine trading theory of Admati and Pfleiderer (1991).

References

Admati, A.R., and P. Pfleiderer, 1991, Sunshine trading and financial market equilibrium, Review of Financial Studies, Vol 4 (3), 443-481.

Anand, A., Irvine, P., Puckett. A., and K. Venkataramanan, 2011, Market Crashes and Institutional Trading, working paper, Syracuse University.

Attari, M., Mello, A.S. and M. Ruckes, 2005, Arbitraging arbitragers, Journal of Finance, Vol. 60 (5), 2471-2511.

Bessembinder, H., Maxwell, W.H., and K. Venkataraman, Market transparency, liquidity externalities and institutional trading costs in corporate bonds, Journal of Financial Economics, 2006, Vol 82, 251-288.

Brunetti, C., and B. Buyuksahin, 2009, Is speculation destabilizing?, working paper, Johns Hopkins University.

Brunnermeier, M., and L.H. Pedersen, 2005, Predatory trading, Journal of Finance, Vol 60 (4), 1825-1863.

Buyuksahin, B., and J. Harris, 2011, Do speculators drive crude oil prices?, The Energy Journal, 32, 167-202.

Cai, F., 2009, Trader exploitation of order flow information during the LTCM crisis, The Journal of Financial Research, Vol. 32 (3), 261-284.

Carlin, B., Lobo, M.S. and S. Viswanathan, 2007, Episodic liquidity crises: Cooperative and predatory trading, Journal of Finance, 62 (5), 2235-2270.

Erb, C. and C. Harvey, 2006, The strategic and tactical value of commodity futures, Financial Analysts Journal, 62, 69-97.

Fama, E., and K. French, 1987, Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage, Journal of Business, Vol. 60 (1), 55-73.

Fama, E., and K. French, 1988, Business cycles and the behavior of metals prices, Journal of Finance, Vol 43 (5), 1075-1093.

Madhavan, A., Richardson, M., and M. Roomans, 1997, Why do security prices change: a transaction level analysis of NYSE-listed stocks, Review of Financial Studies, 10, 1035-1064.

Gorton, G., F. Hayashi, and G. Rouwenhorst, 2007, The fundamentals of commodity futures returns, Yale University working paper.

Hong, H. and M. Yogo, 2012, What does futures market open interest tell us about the macroeconomy and asset prices?, Journal of Financial Economics, forthcoming.

Hasbrouck, J., 2007, Empirical Market Microstructure: The institutions, economics and econometrics of securities trading. Oxford University Press.

Huang, R., and H. Stoll, 1997, The components of the bid-ask spread: a general approach, Review of Financial Studies, 10, 995-1034.

Jones, C., and M. Lipson. 2001. Sixteenths: Direct evidence on institutional execution costs. *Journal of Financial Economics* 59: 253-278.

Keim, D., and A. Madhavan. 1995. Anatomy of the trading process: Empirical evidence on the behavior of institutional traders. *Journal of Financial Economics* 37: 371-398.

Kilian, L. 2009, Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, American Economic Review 99, 1053-1069.

Labuszewski, J.W., Nyhoff, J., Co, R. and P.E. Peterson, 2010, The CME Group Risk Management Handbook, John Wiley and Sons.

Mou, Y., 2011, Limits to Arbitrage and Commodity Index Investment: Front-running the Goldman Roll, working paper, Columbia University.

Pirrong, C., 2011, An evaluation of the performance of oil price benchmarks during the financial crisis, working paper, University of Houston.

Sadka, R., 2006, Momentum and post-earnings announcement drift anomalies: The role of liquidity risk, Journal of Financial Economics, 80, 309-349.

Schoneborn, T., and A. Schied, 2007, Liquidity in the face of adversity: Stealth vs. sunshine trading, predatory trading vs. liquidity provision, working paper, Technical University Berlin.

Schultz, P., 2001. Corporate bond trading costs: a peek behind the curtain. Journal of Finance 56, 677–698.

Stoll, H., and R.E. Whaley, 2010, Commodity Index Investing and Commodity Future Prices, Journal of Applied Finance, Issue 1, 1-40.

Szymanowska, M., F. De Roon, T. Nijman, and R. Van den Goorbergh, 2011, An Anatomy of Commodity Futures Risk Premia, working paper, Erasmus University.

Table 1. ETF Trading activity on Roll Days

Reported are daily trading volume (in contracts) in NYMEX Oil Futures market and during the two-minute settlement period on ETF Roll days. Also reported are cumulative trading volume attributable to eight sample ETFs on the ETF roll day. ETF roll day is identified as the day with the highest ETF sample trading volume aggregated across all ETFs in each month. We rely on Commodity Futures Trading Commission (CFTC) dataset to identify the trading activity of eight ETFs during the period March 1, 2008 to February 28, 2009. The CFTC dataset includes all completed trades in NYMEX crude oil futures, including floor and block trades, as well as trades completed on the GLOBEX electronic market. ETF Roll dates are identified based on the CFTC dataset and independently verified by the ETF or the owner of the benchmark index that the ETF tracks. Reported are net ETF selling activity (sells minus buys) for expiring (front) contract and net ETF buying activity (buys minus sells) for next-to-expiring (second) contract on ETF Roll days.

	Front Contract on Roll Date				Second Contract on Roll Date			
Roll date	ETF Selling Activity (contracts)	Market Trading Volume (contracts)	ETF %	Market Trading Volume During Settlement	ETF Buying Activity (contracts)	Market Trading Volume (contracts)	ETF %	Market Trading Volume During Settlement
3/5/08	4,455	414,308	1%	16,756	5,362	205,827	3%	10,449
4/8/08	5,632	307,800	2%	16,338	5,694	165,544	3%	15,775
5/6/08	5,122	331,913	2%	11,933	5,139	129,110	4%	6,632
6/6/08	8,779	508,749	2%	18,139	6,228	231,984	3%	11,112
7/8/08	7,208	382,404	2%	15,378	8,055	154,453	5%	13,299
8/6/08	6,289	307,994	2%	16,189	6,293	140,471	4%	13,489
9/8/08	11,961	317,923	4%	18,581	11,439	142,644	8%	14,791
10/7/08	9,119	342,917	3%	21,235	9,097	193,234	5%	15,414
11/6/08	13,031	292,018	4%	6,756	14,665	87,869	17%	3,578
12/5/08	23,725	327,140	7%	27,508	23,751	157,572	15%	22,765
1/6/09	49,852	331,307	15%	9,145	42,409	183,802	23%	7,659
2/6/09	67,882	518,382	13%	32,674	58,764	318,960	18%	29,187
Sum	213,055	4,382,855	13%	210,632	196,896	2,111,470	18%	164,150

Table 2. Trade Types of ETF Trades and the Overall Market

Reported are trade types for ETF trades and overall (non-ETF) trading activity in NYMEX Oil Futures market. We rely on Commodity Futures Trading Commission (CFTC) dataset to identify the trading activity of eight ETFs during the period March 1, 2008 to February 28, 2009. The CFTC dataset includes all completed trades in NYMEX crude oil futures, including floor and block trades, as well as trades completed on the GLOBEX electronic market for expiring (front) contract and next-to-expiring (second) contract. Spread trades combine regular contracts for sale in a given delivery month and purchase in another month. Trade-at-settlement (TAS) trades are executed during the trading day at a to-be-determined settlement price.

	Front Month				Second Month			
Trade Type	ETF Trading Activity	% of ETF	Market Trading Activity	% of Market	ETF Trading Activity	% of ETF	Market Trading Activity	% of Market
Option on Future			112,690	0.2%			129,293	0.4%
Option Spread Ratio			225,767	0.3%			205,843	0.7%
Option Spread Conversion			152,450	0.2%			36,417	0.1%
Exchange For Physical			171,631	0.2%			8,644	0.0%
Crack Spread			2,045,089	2.9%			822,321	2.7%
Crack Cross			108,081	0.2%			49,502	0.2%
Trade-at-settlement	15,870	7.2%	3,485,249	4.9%	14,937	6.9%	1,186,320	3.9%
Cabinet	14,966	6.8%	352,729	0.5%	16,564	7.6%	234,238	0.8%
Block Trade	56,670	25.7%	906,990	1.3%	59,887	27.5%	567,582	1.9%
Block TAS Trades	130,951	59.5%	401,856	0.6%	122,126	56.1%	289,605	1.0%
Regular Outright	219	0.1%	45,032,772	63.7%	2,862	1.3%	8,625,010	28.7%
Intra-Commodity Spread	1,239	0.6%	14,987,761	21.2%	991	0.5%	16,690,443	55.5%
Regular Outright Cross			1,821,314	2.6%	7	0.0%	356,066	1.2%
Intra-Commodity Spread Cross			643,962	0.9%	7	0.0%	711,425	2.4%
Other	314	0.1%	255,111	0.4%	414	0.2%	144,121	0.5%
Total	220,229	100%	70,703,452	100%	217,795	100%	30,056,830	100%

Table 3. Average Market Quality Measures on ETF Roll and non-Roll days

Reported are market quality measures on ETF Roll days and non-roll days in the NYMEX Oil Futures market. ETF roll day is identified as the day with the highest ETF sample trading volume in each month. We rely on Commodity Futures Trading Commission (CFTC) datafor trading volume and number of liquidity providing accounts. The spread and imbalance measures are based on the Chicago Mercantile Exchange's Datamine database for GLOBEX electronic market. Roll dates are based on trading activity of eight ETFs in the CFTC database during the period March 1, 2009 to February 28, 2009. Market quality is calculated each minute of the day and then averaged across ETF-roll and non-roll days. Reported are the Wilcoxon signed rank t-statistic and p-value with the null hypothesis of zero difference in the mean. Trading volume from CFTC database includes all completed trades in NYMEX crude oil futures, including floor and block trades, as well as trades completed on the GLOBEX. Trade imbalance is the signed difference between buyer and seller initiated volume standardized by subtracting the mean and dividing by the standard deviation of imbalance during the same minute (across roll and non-roll days). Trades are signed as buyer- or seller-initiated based on a modified Lee and Ready (1991) algorithm. Quoted bid-ask spread (in basis points) is the difference between the lowest limit price for unexecuted sell orders and the highest limit price for unexecuted buy orders. Depth is the total volume of unexecuted sell (ask depth) and buy (bid depth) orders at prices within four ticks of the most competitive prices. Effective spread (in basis points) for a given trade is twice the excess of the trade price over the bid-ask midpoint for those trades initiated by buy market orders and twice the excess of the midpoint over the trade price for those trades initiated by sell market orders.

Market Quality	Roll		Non-Roll			Wilcoxon signed rank		Ansari-Bradley Z Statistic	
	Mean	Median	Mean	Median	Difference:	T-Stat	P-value	Z-Stat	P-value
Trading Volume per Minute (contracts)									
- Front contract	855.4	716.4	648.0	580.7	207.4	12.4	(0.000)	5.8	(0.000)
- Second contract	402.3	283.9	273.7	239.8	128.7	8.4	(0.000)	8.8	(0.000)
Standardized Trade Imbalance									
- Front contract	-0.02	-0.02	0.01	0.01	-0.03	-1.82	(0.034)	19.7	(0.000)
- Second contract	0.00	0.01	0.00	0.00	0.01	0.31	(0.378)	17.9	(0.000)
Quoted Spreads (in basis points)									
- Front contract	1.13	1.12	1.17	1.16	-0.04	-5.6	(0.000)	-10.1	(0.000)
- Second contract	1.42	1.39	1.45	1.42	-0.03	-3.1	(0.001)	12.3	(0.000)
Near-inside Bid Depth (contracts)									
- Front contract	51.9	51.0	46.4	46.0	5.5	13.17	(0.000)	10.4	(0.000)
- Second contract	23.9	23.3	24.8	24.6	-1.0	-6.37	(0.000)	9.6	(0.000)
Near-inside Ask Depth (contracts)									
- Front contract	49.2	48.2	44.0	43.9	5.2	14.60	(0.000)	5.7	(0.000)
- Second contract	20.5	19.9	21.8	21.6	-1.3	-7.96	(0.000)	9.0	(0.000)
Liquidity Supplying Accounts (Number)									
- Front contract	874		681		193	2.23	(0.001)		
- Second contract	167		153		14	0.39	(0.500)		
Effective Spread (in basis points)									
- Front contract	1.96	1.83	2.06	2.01	-0.10	-5.64	(0.000)	9.4	(0.000)
- Second contract	2.29	2.06	2.39	2.32	-0.10	-4.60	(0.000)	8.7	(0.000)

Table 4. Strategic Trading around a known liquidation – Numerical Outcomes

Reported are the outcomes based on the strategic trader's optimal choice of quantities to maximize profits (expression (13)). The initial price (V0) equal to \$100, N=32 periods within each interval, and the liquidation amount is 20 units. The parameter λ=0.015 implies that the cumulative permanent price impact is of the liquidation is 30 basis points. The column 'LP' refers to the liquidator's proceeds for the 20 units. The columns 'Pre', 'During' and 'After' report the profit maximizing strategic trader's order flow as a proportion of the liquidator sale, with positive values indicating trading in the same direction as the liquidator. The column 'SP' reports the strategic traders profit from expression (9). The column 'AC' reports the limit order traders' (or natural counterparties) acquisition cost for the 20 units.

Table 4: Numerical Outcomes from Closed Form Solutions -- Strategic Trading Around a Known 20 Unit Liquidation, Lambda = .015, Gamma = 0.5

	Base	Optimal Unconstrained Strategic Trading						Optimal Strategic Trading, Constrained No Prior					
Theta	LP = AC	Pre	During	After	SP	LP	AC	Pre	During	After	SP	LP	AC
0.00	1990.7	0.40	-0.33	-0.07	1.8	1991.4	1993.1	0.00	-0.13	0.13	0.2	1991.9	1992.1
0.02	1990.5	0.40	-0.33	-0.07	1.8	1991.3	1993.1	0.00	-0.13	0.13	0.2	1991.8	1992.0
0.04	1990.4	0.39	-0.33	-0.06	1.8	1991.2	1993.0	0.00	-0.14	0.14	0.2	1991.7	1992.0
0.06	1990.3	0.39	-0.33	-0.06	1.8	1991.2	1993.0	0.00	-0.14	0.14	0.3	1991.6	1991.9
0.08	1990.1	0.39	-0.33	-0.05	1.8	1991.1	1992.9	0.00	-0.14	0.14	0.3	1991.5	1991.8
0.10	1990.0	0.38	-0.33	-0.05	1.8	1991.0	1992.8	0.00	-0.14	0.14	0.3	1991.4	1991.7
0.12	1989.8	0.38	-0.33	-0.04	1.8	1991.0	1992.8	0.00	-0.15	0.15	0.3	1991.3	1991.6
0.14	1989.7	0.37	-0.33	-0.04	1.8	1990.9	1992.7	0.00	-0.15	0.15	0.3	1991.2	1991.5
0.16	1989.5	0.37	-0.33	-0.04	1.9	1990.8	1992.6	0.00	-0.15	0.15	0.3	1991.1	1991.4
0.18	1989.3	0.36	-0.33	-0.03	1.9	1990.7	1992.6	0.00	-0.15	0.15	0.3	1990.9	1991.3
0.20	1989.2	0.36	-0.33	-0.03	1.9	1990.6	1992.5	0.00	-0.15	0.15	0.4	1990.8	1991.2
0.22	1989.0	0.36	-0.33	-0.02	1.9	1990.5	1992.4	0.00	-0.16	0.16	0.4	1990.7	1991.1
0.24	1988.8	0.35	-0.33	-0.02	1.9	1990.4	1992.3	0.00	-0.16	0.16	0.4	1990.5	1990.9
0.26	1988.6	0.35	-0.33	-0.01	2.0	1990.3	1992.2	0.00	-0.16	0.16	0.4	1990.4	1990.8
0.28	1988.3	0.34	-0.33	-0.01	2.0	1990.1	1992.1	0.00	-0.16	0.16	0.4	1990.2	1990.7
0.30	1988.1	0.34	-0.33	-0.01	2.0	1990.0	1992.0	0.00	-0.16	0.16	0.5	1990.0	1990.5
0.32	1987.9	0.34	-0.33	0.00	2.0	1989.9	1991.9	0.00	-0.17	0.17	0.5	1989.9	1990.4
0.34	1987.6	0.33	-0.33	0.00	2.1	1989.7	1991.8	0.00	-0.17	0.17	0.5	1989.7	1990.2
0.36	1987.3	0.33	-0.33	0.01	2.1	1989.5	1991.6	0.00	-0.17	0.17	0.6	1989.5	1990.0
0.38	1987.0	0.32	-0.33	0.01	2.1	1989.4	1991.5	0.00	-0.17	0.17	0.6	1989.2	1989.8
0.40	1986.7	0.32	-0.34	0.02	2.2	1989.2	1991.3	0.00	-0.17	0.17	0.6	1989.0	1989.6
0.42	1986.4	0.31	-0.34	0.02	2.2	1989.0	1991.2	0.00	-0.18	0.18	0.6	1988.8	1989.4
0.44	1986.0	0.31	-0.34	0.03	2.3	1988.8	1991.0	0.00	-0.18	0.18	0.7	1988.5	1989.2
0.46	1985.6	0.31	-0.34	0.03	2.3	1988.5	1990.8	0.00	-0.18	0.18	0.7	1988.2	1988.9
0.48	1985.2	0.30	-0.34	0.03	2.4	1988.3	1990.6	0.00	-0.18	0.18	0.8	1987.9	1988.7
0.50	1984.8	0.30	-0.34	0.04	2.4	1988.0	1990.4	0.00	-0.18	0.18	0.8	1987.6	1988.4
0.52	1984.3	0.30	-0.34	0.04	2.5	1987.7	1990.2	0.00	-0.19	0.19	0.9	1987.2	1988.1
0.54	1983.8	0.29	-0.34	0.05	2.6	1987.4	1989.9	0.00	-0.19	0.19	0.9	1986.8	1987.8
0.56	1983.3	0.29	-0.34	0.05	2.7	1987.0	1989.6	0.00	-0.19	0.19	1.0	1986.4	1987.4
0.58	1982.7	0.28	-0.34	0.05	2.7	1986.6	1989.3	0.00	-0.19	0.19	1.0	1986.0	1987.0
0.60	1982.0	0.28	-0.34	0.06	2.8	1986.2	1989.0	0.00	-0.19	0.19	1.1	1985.5	1986.6
0.62	1981.3	0.28	-0.34	0.06	2.9	1985.7	1988.6	0.00	-0.19	0.19	1.2	1984.9	1986.1
0.64	1980.5	0.27	-0.34	0.06	3.1	1985.2	1988.2	0.00	-0.20	0.20	1.2	1984.3	1985.5
0.66	1979.6	0.27	-0.34	0.07	3.2	1984.6	1987.8	0.00	-0.20	0.20	1.3	1983.7	1985.0
0.68	1978.7	0.27	-0.34	0.07	3.3	1983.9	1987.3	0.00	-0.20	0.20	1.4	1982.9	1984.3
0.70	1977.6	0.27	-0.34	0.07	3.5	1983.2	1986.7	0.00	-0.20	0.20	1.5	1982.1	1983.6
0.72	1976.4	0.26	-0.34	0.08	3.7	1982.3	1986.0	0.00	-0.20	0.20	1.6	1981.1	1982.7
0.74	1975.0	0.26	-0.34	0.08	3.9	1981.3	1985.3	0.00	-0.20	0.20	1.7	1980.0	1981.7
0.76	1973.4	0.26	-0.34	0.08	4.2	1980.2	1984.4	0.00	-0.20	0.20	1.8	1978.8	1980.6
0.78	1971.6	0.26	-0.34	0.08	4.5	1978.9	1983.4	0.00	-0.20	0.20	1.9	1977.4	1979.3
0.80	1969.6	0.26	-0.34	0.08	4.8	1977.4	1982.2	0.00	-0.20	0.20	2.1	1975.7	1977.8
0.82	1967.1	0.26	-0.34	0.08	5.2	1975.5	1980.8	0.00	-0.20	0.20	2.2	1973.7	1975.9
0.84	1964.2	0.27	-0.35	0.08	5.7	1973.3	1979.0	0.00	-0.20	0.20	2.3	1971.3	1973.6
0.86	1960.8	0.27	-0.35	0.08	6.3	1970.5	1976.8	0.00	-0.19	0.19	2.4	1968.4	1970.8
0.88	1956.6	0.29	-0.35	0.07	7.1	1966.9	1973.9	0.00	-0.19	0.19	2.4	1964.7	1967.1
0.90	1951.4	0.30	-0.36	0.06	8.1	1962.0	1970.1	0.00	-0.18	0.18	2.3	1959.9	1962.2
0.92	1944.9	0.33	-0.37	0.04	9.6	1955.1	1964.7	0.00	-0.16	0.16	1.9	1953.5	1955.4
0.94	1936.7	0.38	-0.38	0.00	12.1	1944.2	1956.3	0.00	-0.12	0.12	1.1	1944.1	1945.3
0.96	1926.1	0.48	-0.41	-0.07	17.5	1923.8	1941.3	0.00	-0.04	0.04	0.1	1928.7	1928.8
0.98	1912.2	0.74	-0.46	-0.28	36.2	1868.3	1904.4	0.00	0.22	-0.22	2.9	1893.0	1896.0
1.00	1893.8	16.17	-0.33	-15.83	1648.5	-1401.1	247.4	0.00	7.75	-7.75	386.7	1070.6	1457.2

Table 5. Regression estimates of permanent and temporary price impact and the resiliency of the market

Reported are estimates of the permanent price impact (λ), the temporary price impact (γ) and the resiliency of the market (θ) in the NYMEX crude oil markets for the full sample and separately on ETF roll and non-roll days. The analysis relies on trades and limit order book data from Chicago Mercantile Exchange's Datamine database on GLOBEX electronic market. Roll dates are based on trading activity of eight ETFs in the CFTC database during March 1, 2009 to February 28, 2009. We estimate these parameters with geometric lag regressions of the following form:

$$P_t - M_1 = \alpha + \gamma \sum_{j=t-k}^{t} \theta^{t-j} q_j + \lambda \sum_{j=t-k}^{t} q_j^* + \varepsilon_t,$$

where P_t is the time t trade price, M_1 is the quote midpoint at the beginning of trading day (i.e., 9:00 A.M. EST), Q_j is the signed order imbalance at time j, and Q^*_j is the residual from a fifth-order auto-regression of Q_j, following the specification in Sadka (2006). Expression (13) is estimated for an one-second model with 75 lags of order imbalance (Panel B) and a five-second model with sixty lags of order imbalance (Panel A). For time periods with multiple trades P_t is measured as the last trade price and Q_j is measured as net trade imbalance during the period. The geometric lag expression (13) is estimated by Generalized Method of Moments (GMM), using SAS Proc Model with a Bartlett Kernel set equal to the lag length plus one. Reported are the difference between regressions coefficients estimated on ETF-roll and non-roll days and test of statistical significance of the difference.

	Number of observations	alpha (α)	Lambda (λ)	Gamma (γ)	Theta (θ)	R^2
Panel A: Time interval = 5 seconds; Lags = 60						
Front Contract: Full sample	1,086,363	25.643	0.056	0.021	0.959	53.16%
Non-Roll Days	836,963	18.942	0.058	0.020	0.959	52.84%
Roll	48,412	50.860	0.054	0.038	0.986	63.11%
Difference		31.918 ***	-0.004 **	0.018	0.027	
p-value		(0.00)	(0.03)	(0.21)	(0.11)	
Second Contract: Full sample	1,084,194	-8.196	0.084	0.057	0.982	15.47%
Non-Roll Days	835,194	-9.860	0.082	0.048	0.977	16.61%
Roll	48,378	23.505	0.149	0.151	0.993	9.49%
Difference		33.365 ***	0.067 ***	0.104	0.016	
p-value		(0.00)	(0.00)	(0.20)	(0.42)	
Panel B: Time interval = 1 second; Lags = 75						
Front Contract: Full sample	5,261,609	25.84	0.051	0.038	0.976	53.64%
Non-Roll Days	4,047,759	19.180	0.052	0.036	0.975	53.28%
Roll	237,349	53.110	0.050	0.063	0.990	64.35%
Difference		33.930 ***	-0.002 ***	0.027	0.015 *	
p-value		(0.00)	(0.00)	(0.12)	(0.07)	
Second Contract: Full sample	5,184,068	-7.410	0.076	0.070	0.994	16.46%
Non-Roll Days	3,987,888	-8.792	0.075	0.060	0.993	17.44%
Roll	213,335	36.440	0.143	0.182	0.996	9.48%
Difference		45.232 ***	0.069 ***	0.122	0.004	
		(0.00)	(0.00)	(0.18)	(0.77)	

Table 6: Strategic Trading surrounding the ETF Roll

Reported in Panel A are patterns associated with twelve strategic trading strategies associated with the ETF-roll. To be identified as a strategic trader, the absolute value of net change in the (non-ETF) account's inventory to the account's total activity surrounding the roll must be less than 0.25. The "During" period is defined as between 9 A.M. and 5 P.M. EST on the ETF roll day, the "Before" period is defined from Midnight on Day -3 (three trading days prior) to 9 A.M. on roll day, and the "After" period is defined from 5 P.M. on roll day to Midnight on Day +3 (three trading days after) relative to the roll day. A strategic trader whose signed position change on roll day moves *against* ETF's inventory change is deemed a liquidity provider (Strategies ST1 to ST5) while a strategic trader whose signed position change on roll day moves *with* the ETF's inventory change is deemed a predatory trader (Strategies ST8 to ST12). Categories ST6 and ST7 correspond to trading patterns with no trading activity on the roll day. Strategic traders are further classified into one of five sub-strategies within liquidity provision (ST1-ST5) and predatory trading (ST8-ST12) based on the account's change in net positions in the Before and After period. Also identified below is the complementary strategy where strategic traders pursue an opposite trading pattern surrounding the ETF Roll. Panel A reports the direction of ETF activity and those for each strategy for expiring (front) contract and next-to-expiring (second) contract on ETF Roll days.

Panel A: Direction of ETF and Strategic Trading surrounding the ETF Roll

Strategy	Trading Pattern (relative to ETF)			Front month			Second month			Complement strategy
	Before	During	After	Before	During	After	Before	During	After	
ETF*				none	sell	none	none	buy	none	
ST 1	against	against	with	buy	buy	sell	sell	sell	buy	ST 12
ST 2	none	against	with	none	buy	sell	none	sell	buy	ST 11
ST 3	with	against	against	sell	buy	buy	buy	sell	sell	ST 10
ST 4	with	against	none	sell	buy	none	buy	sell	none	ST 9
ST 5	with	against	with	sell	buy	sell	buy	sell	buy	ST 8
ST 6	against	none	with	buy	none	sell	sell	none	buy	ST 7
ST 7	with	none	against	sell	none	buy	buy	none	sell	ST 6
ST 8	against	with	against	buy	sell	buy	sell	buy	sell	ST 5
ST 9	against	with	none	buy	sell	none	sell	buy	none	ST 4
ST 10	against	with	with	buy	sell	sell	sell	buy	buy	ST 3
ST 11	none	with	against	none	sell	buy	none	buy	sell	ST 2
ST 12	with	with	against	sell	sell	buy	buy	buy	sell	ST 1

Table 6, Panel B presents regressions coefficients of normalized strategic volume for all days (ETF roll and Control days) on an ETF roll day indicator variable and an indicator variable for February 6, 2009 roll. The analysis relies on the CFTC dataset. The *strategic volume* is aggregated across all (non-ETF) trading accounts associated with a specific strategy, identified in Panel A, where strategic volume is simply the round trip volume associated with an account surrounding the roll. We define *Normalized strategic volume* as the strategic volume in a strategy less strategic volume in complementary strategy. Such normalization accounts for abnormal trading volume and liquidity associated with a roll day and allows comparison of trading activity across ETF roll and non-roll days. We also calculate Normalized strategic volume for each strategy on every usable non-roll days. A usable non-roll day (Control Day 0) is defined as a day when Days [-3, +3] relative to Control Day 0 does not have any overlap with [-3,+3] relative to ETF roll day. The Roll-day indicator variable takes the value of one for an ETF-roll day, and equals zero otherwise. The Feb 6 indicator variable takes a value of one for the Feb 6 2009 ETF-roll (subject to a CFTC enforcement action) and equals zero otherwise.

Panel B: Normalized Strategic Volume Regressions	(1)	(2)	(3)	(4)	(5)	(6)
Front Month Contract						
Intercept	-306	-52	-851	368	3	-244
t(Intercept)	-0.86	-0.76	-2.14	4.07	0.01	-1.48
Roll_day	-2122	166	2805	254	-222	-1154
t(PAT_day)	-1.91	0.78	2.26	0.90	-0.26	-2.25
FEB6	-777	402	4479	-418	2766	4165
t(FEB6)	-0.21	0.57	1.10	-0.45	0.98	2.47
Second Month Contract						
Intercept	-89	28	-396	-79	102	-83
t(Intercept)	-0.42	0.44	-1.52	-0.87	0.59	-0.59
PAT_day	-1560	-43	2111	-49	-78	50
t(PAT_day)	-2.37	-0.22	2.59	-0.17	-0.14	0.11
FEB6	1277	2678	-1005	1402	957	76
t(FEB6)	0.59	4.11	-0.38	1.51	0.54	0.05

Table 7: Estimated Cost of Completing the USO Roll.

The Table reports the mean estimated proportional roll cost across 66 monthly rolls occurring from April 2006 to October 2011. The roll cost each month is estimated as the term slope on the roll day (two weeks before expiration) less the term slope the indicated number of days prior. From March 2009 onward the roll occurs over four days, and the roll cost is estimated as the term slope averaged across the four roll days less the term slope the indicated number of days before the roll begins. The term slope on day T is defined as $100*\ln(F_{2T}/F_{1T})$, where F_{1T} is the day T settlement price for the crude oil futures contract nearest to expiration and F_{2T} is settlement price of the second nearest-to-expiration crude oil futures contract. Daily settlement prices are obtained from Commodity Research Bureau (CRB).

Benchmark is	Mean Cost	Std. Error	t-stat	P-value
1 Day Prior	0.0980	0.0696	1.41	0.1639
2 Days Prior	0.1559	0.0857	1.82	0.0736
3 Days Prior	0.1754	0.1150	1.53	0.1320
4 Days Prior	0.1602	0.1046	1.53	0.1306
5 Days Prior	0.2107	0.0981	2.15	0.0355
6 Days Prior	0.2340	0.1012	2.31	0.0239
7 Days Prior	0.2861	0.1107	2.58	0.0120
8 Days Prior	0.2743	0.1383	1.98	0.0515
9 Days Prior	0.3190	0.1651	1.93	0.0578
10 Days Prior	0.2075	0.2724	0.76	0.4490

Table 8: Understanding the Crude Oil ETF Stock Price Performance.

The Table reports on the performance of various futures benchmarks between January 1990 and October 2011. All data except USO ETF share prices are obtained from the Commodity Research Bureau (CRB). Daily futures "Return 1" series is based on the current and prior day prices of the nearest-to-expiration contract, for all days including the last day of trading for the expiring contract. Daily futures "Return 2" series is based on the current and prior day prices of the second-nearest-to-expiration contract, for all days including the last day of trading for the expiring contract. Daily futures "Return Benchmark" series is based on the current and prior day prices of the nearest-to-expiration contract when the nearest contract is two or more weeks from expiration, and on the current and prior day prices of the second-nearest-to-expiration contract from then until the nearest contract expires. Also reported are annualized USO ETF returns since USO's inception date. Cost of storage is the futures term slope implied by the settlement prices of the first and second month crude oil futures contracts for each trading day. Ex post Spot Premium is based on a daily time series of spot prices implied by expression (14), relying on the nearest to expiration futures price and the daily cost of. From this series of implied spot prices we compute the time series of realizations of U_t, the daily ex post return premium in spot oil prices. Appreciation in the implied spot price is the summations of the ex post spot premium and the cost of storage. Multiplying by 250 annualizes the daily returns.

	4/10/06 to 10/20/11		**1/1/00 to 4/9/06**		**1/1/90 to 12/31/99**		**1/1/90 to 10/20/11**	
	Days	**1393**	**Days**	**1564**	**Days**	**2510**	**Days**	**5467**
Variable	Mean (x250)	T-stat	Mean (x250)	T-stat	Mean (x250)	T-stat	Mean (x250)	T-stat
Appreciation in Implied Spot Price (S(t)+U(t))	3.98%		15.45%		1.14%		5.95%	
Cost of Storage (term slope S(t))	16.20%	21.73	-7.57%	-11.97	-3.55%	-6.27	0.33%	0.85
Expost Spot Premium (U(t))	-12.22%	-0.68	23.02%	1.49	4.69%	0.37	5.62%	0.65
Futures Return 1	-8.94%	-0.51	26.02%	1.79	2.41%	0.21	6.27%	0.77
Futures Return 2 U(t)+((M-1)*ΔS)	-9.71%	-0.60	25.37%	1.86	5.24%	0.52	7.19%	0.98
Futures Benchmark Return	-14.11%	-0.85	20.80%	1.46	4.84%	0.45	4.58%	0.60
Benchmark less Return 1	-5.17%	-1.19	-5.22%	-2.17	2.43%	0.90	-1.69%	-0.94
Benchmark less Return 2	-4.40%	-2.65	-4.57%	-2.82	-0.40%	-0.15	-2.61%	-1.92
USO ETF Return	-12.79%	-0.80						

Figure 1: Daily United States Oil Fund (USO) share price and Front Month NYMEX Crude Oil Price

The figure represents the daily USO share price and the front month NYMEX crude oil price over five and a half year period since the inception of USO (April 12, 2006).

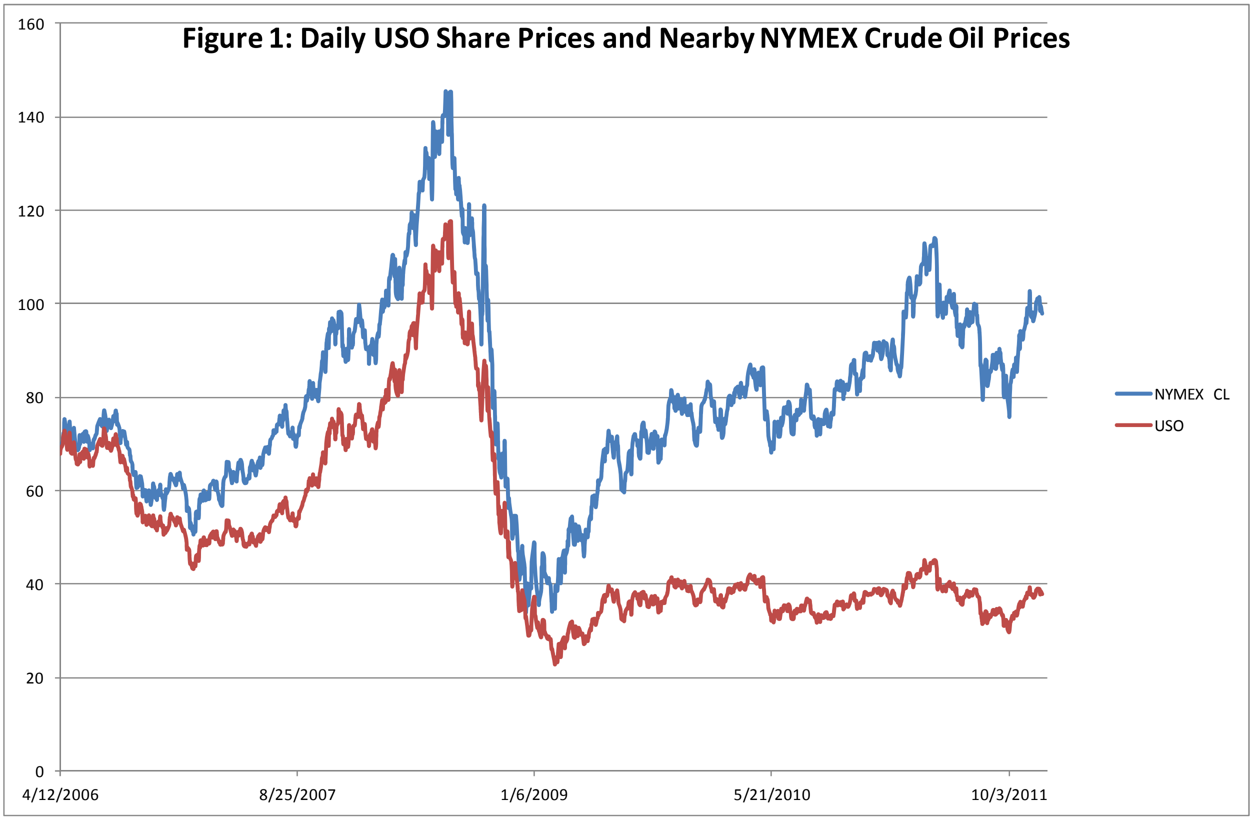

Figure 2: Trading volume and trade imbalance on ETF roll and non-roll days.

Reported are intra-day patterns in trading volume and trade imbalance on ETF Roll days and non-roll days in the NYMEX Oil Futures market. We rely on Commodity Futures Trading Commission (CFTC) data for trading volume and Chicago Mercantile Exchange's Datamine database for imbalance measures. Roll dates are based on trading activity of eight ETFs in the CFTC database during the period March 1, 2008 to February 28, 2009. Market quality is calculated each minute of the day and then averaged across ETF-roll and non-roll days. Trading volume from CFTC database includes all completed trades in NYMEX crude oil futures, including floor and block trades, as well as trades completed on the GLOBEX. Trade imbalance is the signed difference between buyer and seller initiated volume standardized by subtracting the mean and dividing by the standard deviation of imbalance during the same minute (across roll and non-roll days). Trades are signed as buyer- or seller-initiated based on a modified Lee and Ready (1991) algorithm.

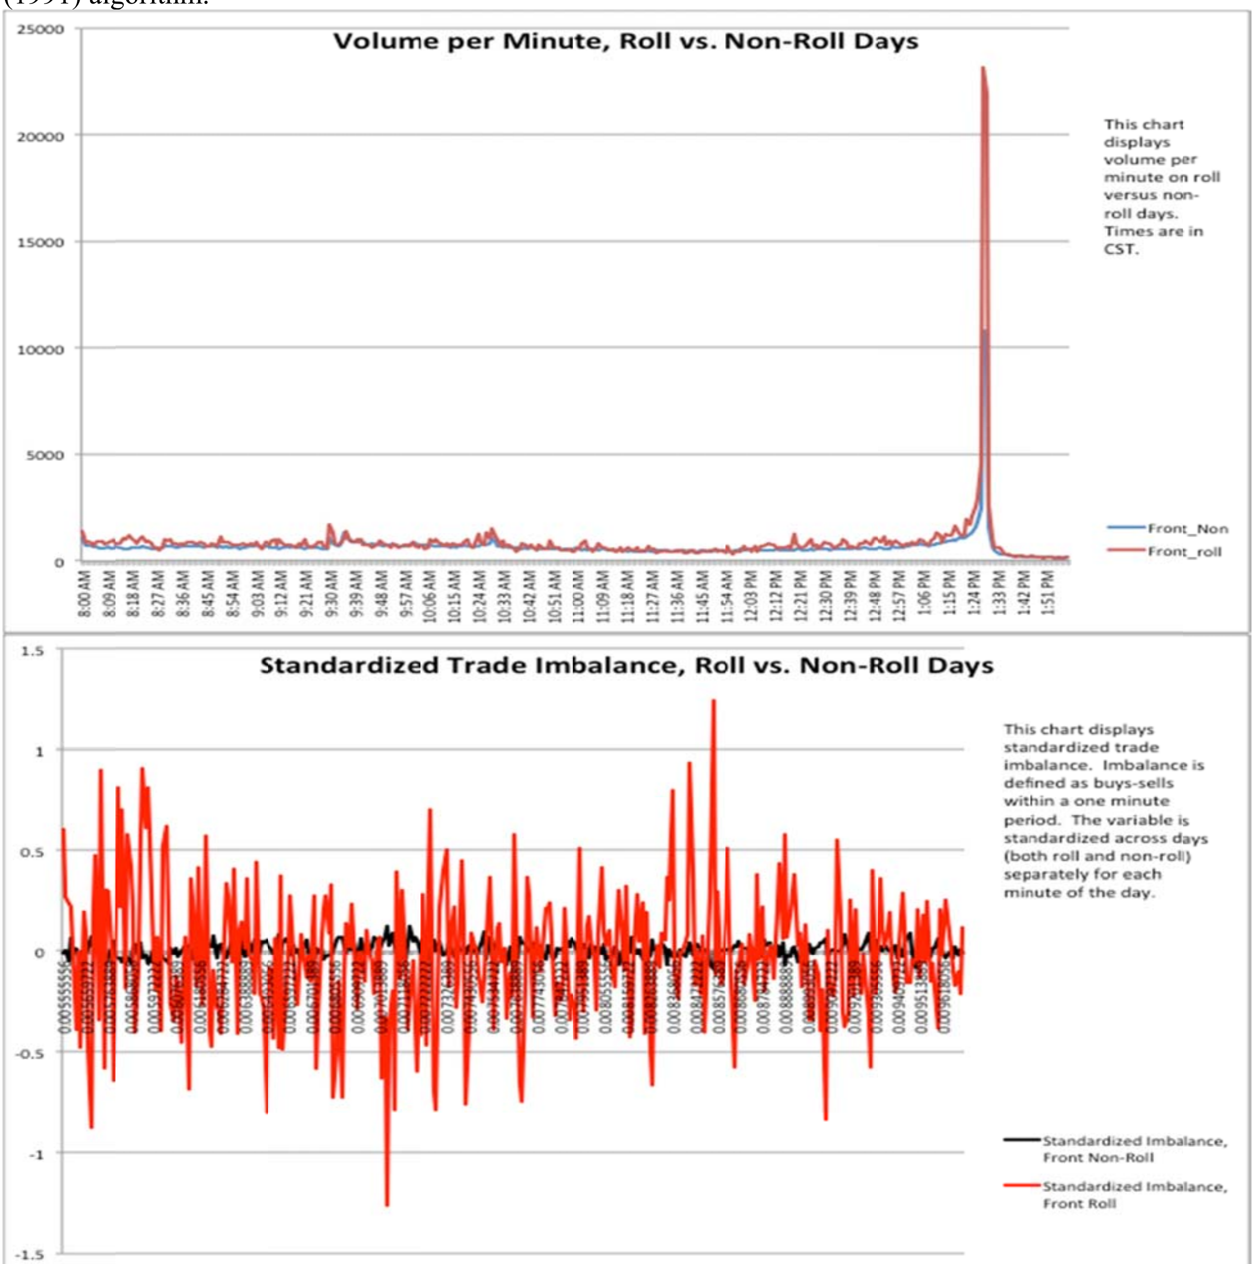

Figure 3: Intraday Quoted Spreads on ETF roll and non-roll days.

Reported are quoted spreads on ETF Roll days and non-roll days in the NYMEX Oil Futures market. We rely Chicago Mercantile Exchange's Datamine database on GLOBEX electronic market for quoted spread measure. Roll dates are based on trading activity of eight ETFs in the CFTC database during the period March 1, 2008 to February 28, 2009. Market quality is calculated each minute of the day and then averaged across ETF-roll and non-roll days. Quoted bid-ask spread (in basis points) is the difference between the lowest limit price for unexecuted sell orders and the highest limit price for unexecuted buy orders.

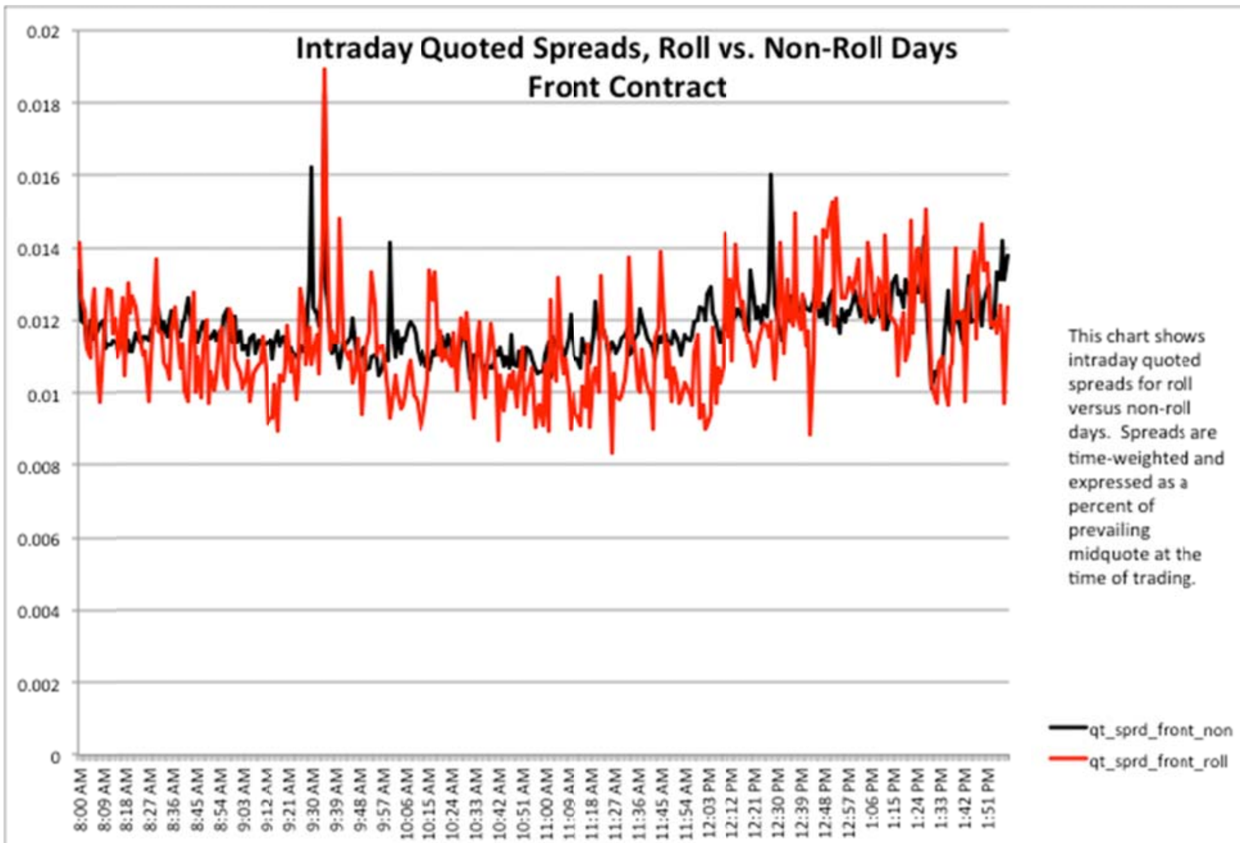

Figure 4: Intraday Limit Order Book depth on ETF roll and non-roll days.

Reported are limit order book depth on ETF Roll days and non-roll days in the NYMEX Oil Futures market. We rely on Chicago Mercantile Exchange's Datamine database on GLOBEX electronic market for the depth measure. Roll dates are based on trading activity of eight ETFs in the CFTC database during the period March 1, 2008 to February 28, 2009. Market quality is calculated each minute of the day and then averaged across ETF-roll and non-roll days. Depth is the total volume of unexecuted sell (ask depth) and buy (bid depth) orders at prices within four ticks of the most competitive prices.

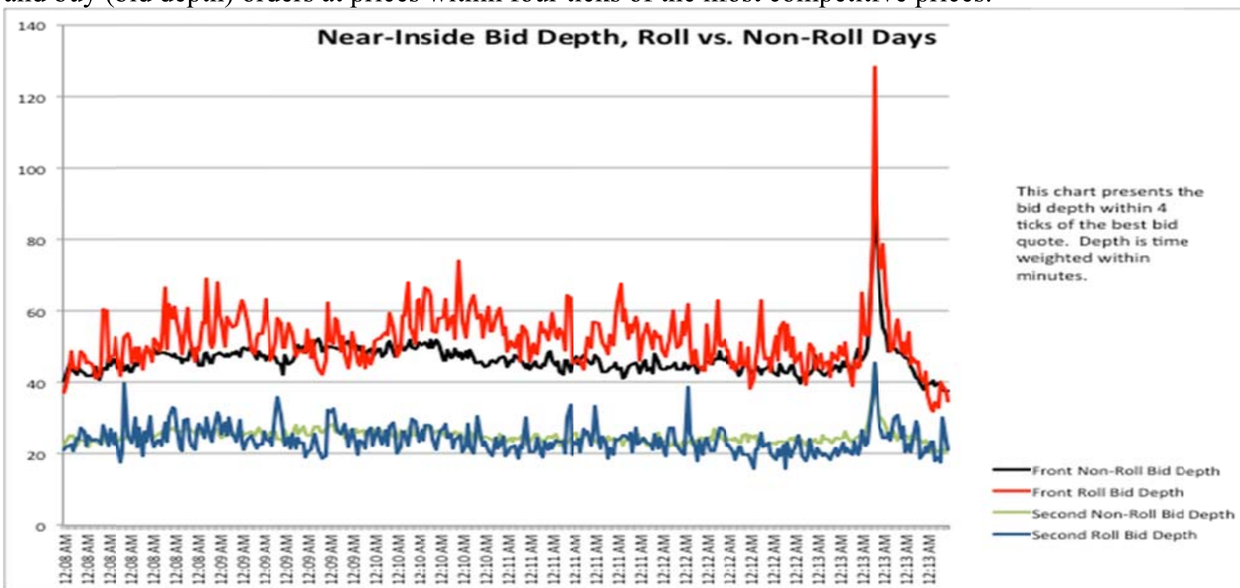

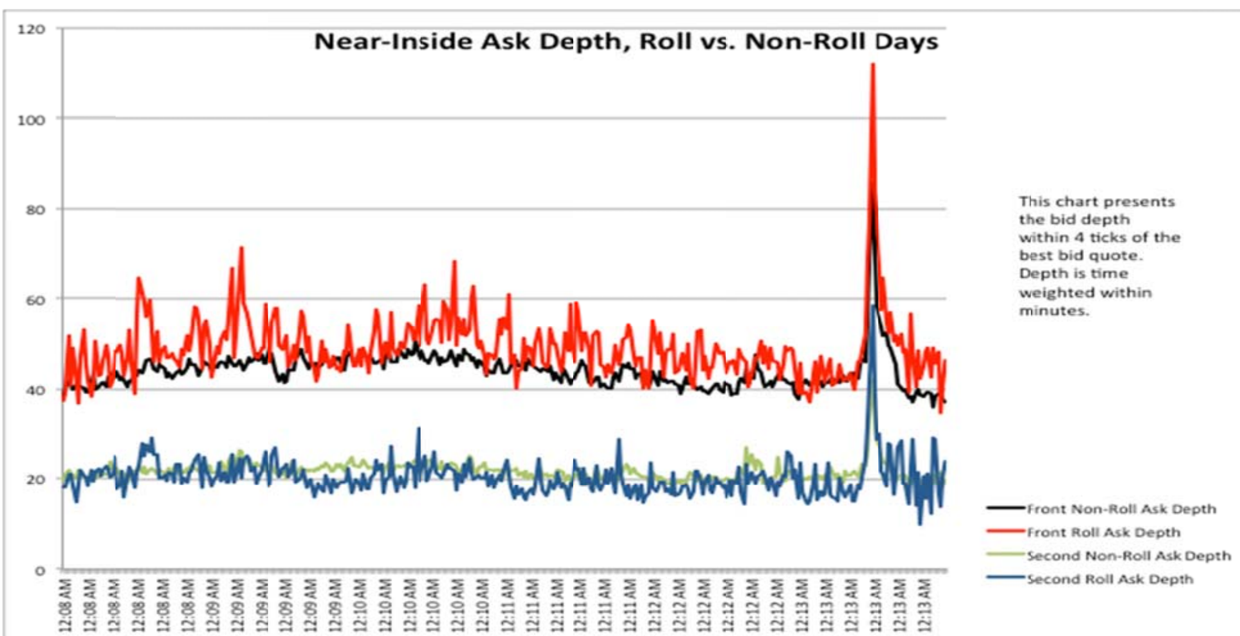

Figure 5: Intraday Effective Spreads on ETF roll and non-roll days.

Reported are effective spreads on ETF Roll days and non-roll days in the NYMEX Oil Futures market. We rely on Chicago Mercantile Exchange's Datamine database on GLOBEX electronic market for the spread measure. Roll dates are based on trading activity of eight ETFs in the CFTC database during the period March 1, 2008 to February 28, 2009. Market quality is calculated each minute of the day and then averaged across ETF-roll and non-roll days. Effective spread (in basis points) for a given trade is twice the excess of the trade price over the bid-ask midpoint for those trades initiated by buy market orders and twice the excess of the midpoint over the trade price for those trades initiated by sell market orders.

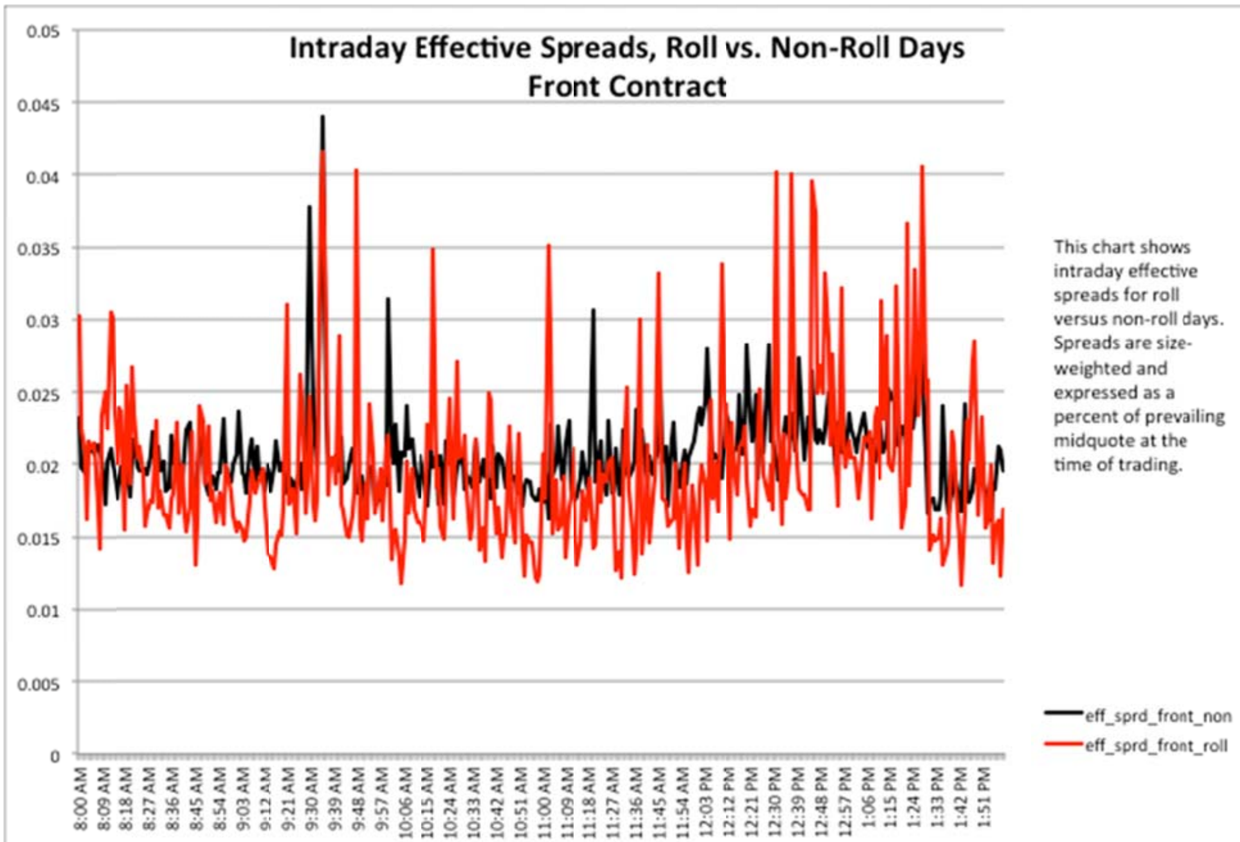

Figure 6: Strategic trading, resiliency and the evolution of price.

Figures 5A and 5B display period-by-period trade prices around a large, pre-announced liquidation. The illustration includes an initial price (V_0) equal to \$100, N = 32 fifteen minute periods within each eight-hour trading day, Q_L = 20 units liquidated, temporary price impact, $\gamma = 0.5$, permanent price impact, $\lambda = 0.015$,, for market resiliency parameters, θ, of 0.0 (Figure 5A) and 0.98 (Figure 5B), respectively. Outcomes for intermediate resiliency parameters are generally similar – only when θ approaches 1 are the effects altered. Table 4 illustrates the outcomes of a broader analysis when the strategic trader chooses quantities to maximize profit. The first third of the individual observations pertain to the pre-liquidation day, the second third pertain to the during-liquidation day, and the final third pertain to the post-liquidation day. The small-diamond line illustrates the evolution of price when the strategic trader is absent. The large-square line illustrates the evolution of price when the strategic trader selects quantities to maximize profits, according to expressions (12) and (13).

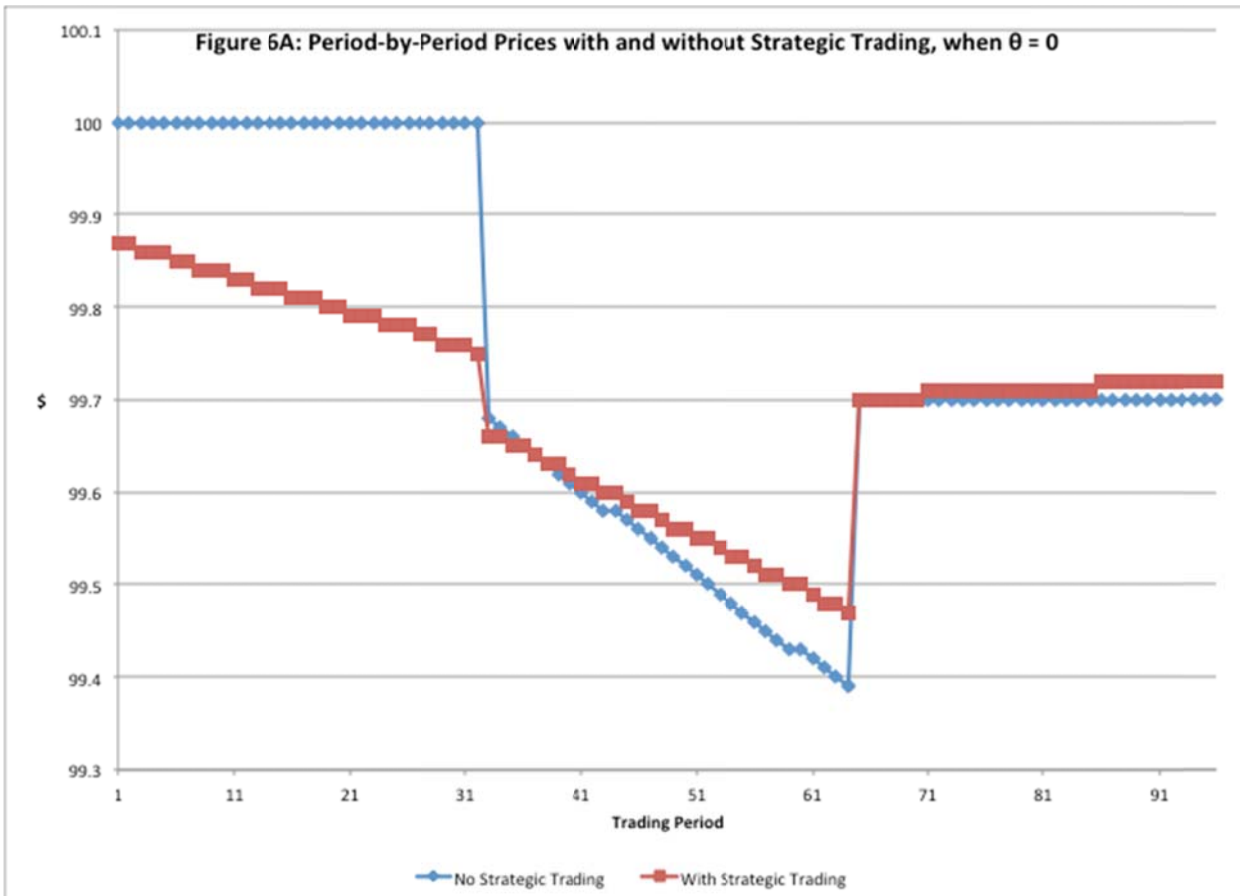

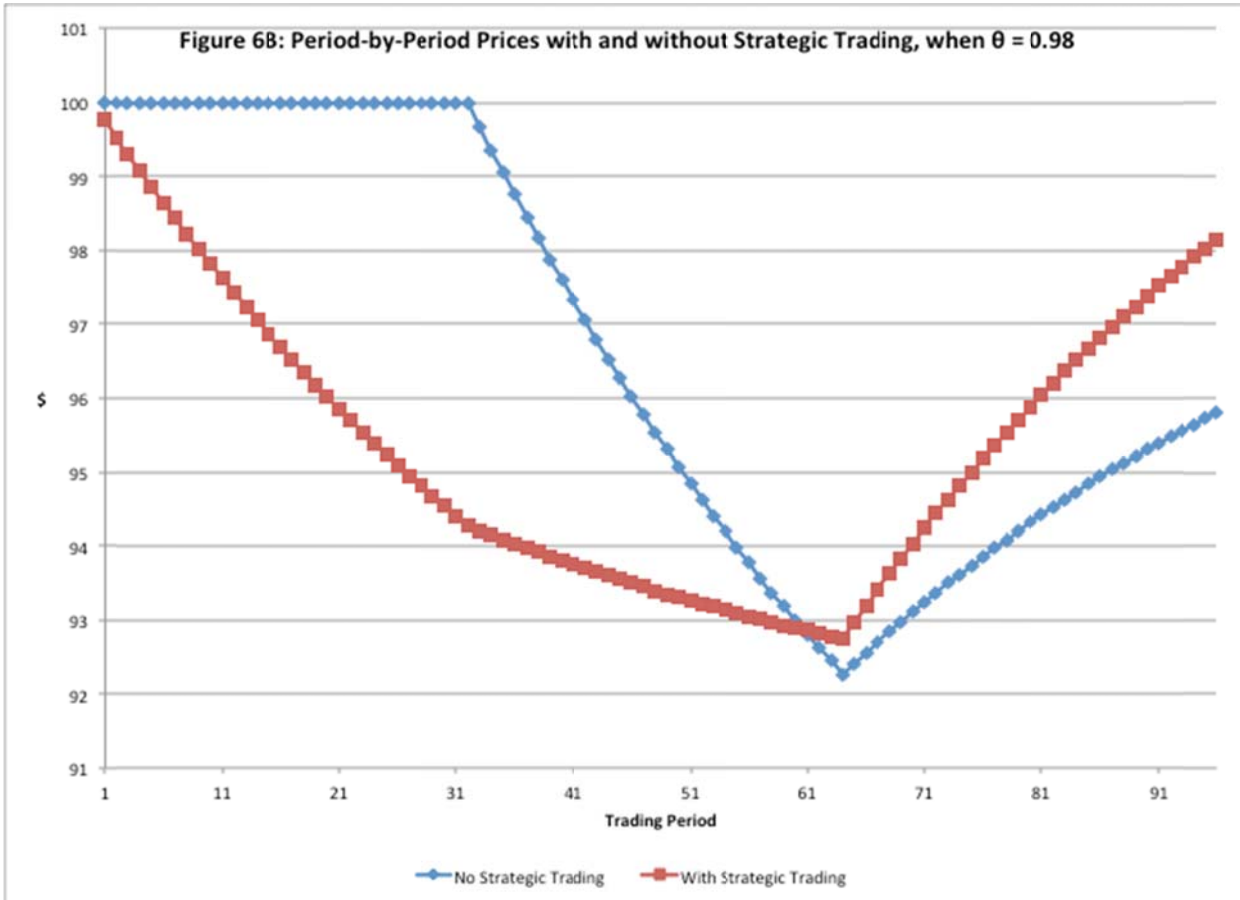
Figure 6B: Period-by-Period Prices with and without Strategic Trading, when θ = 0.98
101
100
99
98
97
96
95
94
93
92
91
$
1
11
21
31
41
51
61
71
81
91
Trading Period
No Strategic Trading
With Strategic Trading