

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection Lee Kong Chian School Of
Business

Lee Kong Chian School of Business

2-2010

An Analysis of Extreme Price Shocks and Illiquidity among Systematic Trend Followers

Wing Bernard LEE

Singapore Management University, bernardlee@smu.edu.sg

Shih-Fen CHENG

Singapore Management University, sfcheng@smu.edu.sg

Annie KOH

Singapore Management University, anniekoh@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/lkcsb_research



Part of the [Finance and Financial Management Commons](#), and the [Portfolio and Security Analysis Commons](#)

Citation

LEE, Wing Bernard; CHENG, Shih-Fen; and KOH, Annie. An Analysis of Extreme Price Shocks and Illiquidity among Systematic Trend Followers. (2010). 1-32.

Available at: https://ink.library.smu.edu.sg/lkcsb_research/1870

This Working Paper is brought to you for free and open access by the Lee Kong Chian School of Business at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection Lee Kong Chian School Of Business by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

An Analysis of Extreme Price Shocks and Illiquidity Among Systematic Trend Followers

Bernard Lee
Sim Kee Boon Institute for Financial Economics
Singapore Management University
bernardlee@smu.edu.sg

Shih-Fen Cheng
School of Information Systems
Singapore Management University
sfcheng@smu.edu.sg

Annie Koh
Lee Kong Chian School of Business
Singapore Management University
anniekoh@smu.edu.sg

June 15, 2010

Abstract

We construct an agent-based model to study the interplay between extreme price shocks and illiquidity in the presence of systematic traders known as trend followers. The agent-based approach is particularly attractive in modeling commodity markets because the approach allows for the explicit modeling of production, capacities, and storage constraints. Our study begins by using the price stream from a market simulation involving human participants and studies the behavior of various trend-following strategies, assuming initially that their participation will not impact the market. We notice an incremental deterioration in strategy performance as and when strategies deviate further and further from the theoretical strategy of lookback straddles (Fung and Hsieh [3]), due to the negative impacts of transaction costs and imperfect execution. Next, the trend followers are allowed to participate in the market, trading against zero intelligence computer traders making randomized bids and offers. We notice that market prices begin to break down as the percentage of trend followers in the market reaches 80%. In addition, in a market dominated by smart traders, it becomes increasingly difficult for any of them to generate profits using what is supposed to be a “long gamma” strategy. After all, trading is a zero-sum game: It is not feasible for any “long gamma” trader to generate a consistent profit unless someone else is willing to be on the other side of his/her trades. In any such market dominated by smart traders with low liquidity and extreme price instability, one proposed solution (as proposed earlier by the U.S. Commodity Futures Trading Commission) is to control position size limits, by either decreasing them (in the original proposal) or increasing them (for completeness in our analysis). Based on our simulation results, we have found no evidence supporting that such a solution will be effective; in fact, doing so will only lead to erratic price behavior as well as a variety of practical issues when imposing such changes to position size limits. An alternative proposal is to intervene in the market direct/indirectly, such as by using a market maker to inject/reduce liquidity. Our simulation results show evidence that injecting and reducing liquidity by the market maker can both be effective. However, a market maker can accumulate a large negative P&L by buying in a one-sided, falling market in which it is the only bidder, or vice versa. Therefore, in practice, no market maker may volunteer to participate in any such market rescue efforts unless governments are willing to underwrite some of its large potential losses. In short, direct/indirect intervention by controlling liquidity is not a panacea, and there are practical limits to its effectiveness.

1 Background and Motivation

Commodity trading has been one of the backbones of economic activities since the dawn of the human civilization. Over the centuries, the trading of commodities has become more and more complex, evolving from primitive barter exchange (direct exchange of goods or services without monetary instruments) to more sophisticated forward contracts between producers and consumers (agreements to buy or sell at a fixed price at a future time), and then, in recent years, to formal futures exchanges with clearing houses guaranteeing transactions.

For most commodities, excessive volatility will hurt either producers or consumers (usually both) and this might have dire consequences: producers and consumers would react to volatile market shocks, potentially making the market even more volatile and straining supplies and/or demands. Eventually, the market can break down completely and the whole supply chain that depends on the commodity will be adversely impacted. Therefore, no matter how a particular commodity is traded, reining in excessive volatility of its price has always been an important policy objective for the authority overseeing its trading. Roughly speaking, factors contributing to the volatility of the commodity prices can be classified into three major categories: 1) natural/seasonal fluctuations in supply and demand, 2) disruptions along the commodity supply chain due to unexpected factors such as acts of nature, and 3) speculation.

For the first two sources of volatilities, while their occurrences are mostly inevitable (but predictable to some degree), there are wide varieties of financial instruments (e.g., options and futures) and physical facilities (e.g., the construction of additional storage for strategic commodities like crude oil and grains) that can be utilized to reduce the likely impact when struck by such unexpected events. However, as commodities become more and more popular as an asset class, the influence of speculation on commodity prices has become increasingly more significant. The influx of hedge funds, commodity trading advisors, and institutional investors into commodity markets has changed the dynamics of the market drastically. They are often seen by the public as convenient scapegoats behind the spectacular commodity booms and busts in the past two years.

Another danger that can be caused by speculative trades is potentially massive defaults due to misplaced bets. This has occurred numerous times when rogue traders placed huge bets on commodities, or when out-sized bets went wrong without sufficient cash available to cover margin requirements. The high-profile failure of the hedge fund Amaranth in 2006 can be attributed to one single trader's misplaced gamble on the natural gas market. Amaranth, with over \$9 billion in assets, borrowed \$8 for every \$1 of its own fund in order to place its bets. The size of such bets could easily distort prices and destabilize markets both before and after the gambles went bad.

Not surprisingly, speculative activities at such order of magnitude have provoked public debates on how to best manage the market impact of speculators. For example, the U.S. Commodity Futures Trading Commission (CFTC) (<http://www.cftc.gov/>) has recently engaged in the drafting of regulations to curb speculation by imposing position size limits. The resulting policy debates focus on whether speculations might cause market instability; if so, whether imposing position limits would be an effective policy solution.

There are two primary challenges in conducting a rigorous analytical study on the above two issues: 1) to come up with proper models for speculative trading, and 2) to construct an environment that allows the analysis of how regulatory policies may impact different trading strategies. In this paper, we adopt an agent-based framework to address these two challenges. Our ultimate goal in this paper is to discuss the effectiveness of controlling position limits when the market is pushed to extreme conditions in which participants are dominated by speculators. The advantages of our agent-based approach are:

- The agent-based framework allows us to build the market from the bottom up. We can therefore take advantage of past research and implement well-known strategies for different roles (e.g., traders trading for liquidity, speculators behaving like hedge funds, and market makers).
- Many different aspects of the trading simulation are easily configurable. We can easily test different market compositions by changing the mix of agents. Also, different market scenarios can be tested easily.
- An agent-based simulation approach allows us to observe the “complete” market evolution, not just a final analytical solution. This allows us to perform more in-depth time-series analyses like the stability and volatility of price history. We can also obtain the profit-and-loss sequences for all agent types.

The rest of paper is structured as follows: Section 2 will go over a brief background on the agent-based approach that we adopt. Section 3 will provide the descriptions for all implemented agent strategies. Section 4 will highlight the theoretical upper bounds of our agent strategies. Section 5 will discuss how we set up our experiment and present our findings. Sections 6 and 7 will examine the impacts of changing position limits and liquidity levels on market stabilities. Finally, Section 8 will discuss the potential implications of our results and conclude the paper.

2 The Agent-based Computational Approach

During the past decade, with advances in both hardware and software, significant progress was made in enabling large-scale study on computational economics. In particular, “Agent-based Computational Economics” (ACE) (see [6], [9], and [10] for comprehensive reviews) has become an alternative approach in studying the behaviors of complex economical systems. The merit of ACE is probably best explained in Tesfatsion’s own words [10]:

The defining characteristic of ACE models is their constructive grounding in the interactions of agents... Starting from an initially specified system state, the motion of the state through time is determined by endogenously generated agent interactions.

These features from ACE are exactly what we are looking for in our simulation platform.

2.1 Simulation Platform

The agent-based simulation platform that we have adopted for our study is derived from a generic market trading game server, AB3D [8]. The generic design of AB3D allows us to easily construct market mechanism, create agent strategies, and compose market scenarios.

To simplify the simulation design and focus on analyzing the primary issue that we care most about, we assume that all trades are handled by a standard Continuous Double Auction (CDA), and all participating agents can buy *long* or sell *short* at any given moment. Although we will be modeling a commodity futures exchange, we assume for now that we ignore the daily settlement of futures contracts. This implies that *margin calls* will not be modeled and the cash flows will be computed only when a transaction is made (to establish or close out a certain position). However, we do require that all agents should not exceed their pre-defined position limits at any time.

2.2 Agent Strategies

As stated earlier, one major goal of this paper is to assess the impact of the level of speculation on market stability. To achieve such a goal, we must clearly define what trading strategies each agent may use. As highlighted in Section 1, coming up with a precise definition of the behavior of a representative speculator is a key challenge. This is so because the term “speculation” is quite general, and it can potentially refer to a wide spectrum of trading strategies.

To avoid unnecessary complications, we follow the argument of Fung and Hsieh [3], and adopt a popular trading strategy: the trend-following strategy commonly used by “Commodity Trading Advisors” (CTA) that replicates lookback options straddles as the speculative strategy. An options straddle is constructed by acquiring both call (rights to buy) and put (rights to sell) options on the same asset at the same time. The “lookback” option is a special type of option that allows its holder to exercise it at its most favorable price from the point of purchase to the point of exercise. Therefore, a lookback call (put) option would allow its owner to buy (sell) the underlying asset at the lowest (highest) price seen after the option is bought. By holding both lookback call and put options, the trader can pocket the difference between the highest and the lowest asset prices seen during the lifespan of the options (after deducting the costs of the call and put options).

To construct controlled experiments with different levels of speculations, we must find a way to repeatedly generate scenarios that are mostly identical except for their different mixes of agents. This requirement is why we must rely on simulation and not historical data: at any given moment, only one specific market configuration exists and extrapolating on what might have happened for other configurations would not be credible. For this reason, we have adopted the agent-based simulation framework described earlier in the section. This agent-based commodity market simulation is driven primarily by a pre-defined event stream, making it possible to construct a background market trend that is qualitatively identical throughout all simulations (this event-driven framework is described in detail in [1]). Although the background market trends are identical, the resulting price streams still have to be generated through agent interactions, and thus every instance will be a stochastic realization that is a function of the strategies of all market participants.

In our study, three types of agents are introduced. A brief description on their roles is provided here. The implementation details of these agents will be discussed in Section 3.

1. Random agents that makes random bids/asks based on a normal distribution of bid-ask spreads and therefore provides the majority of the liquidity.
2. Trend-following agents that executes certain type of trend following strategies by using options or by using the underlying to replicate such options.
3. A market maker agent that provides additional liquidity to the market. In typical cases, the market maker agent is allowed to make a small profit but is obligated to quote both buy and sell bids. When the market is extremely unstable or one-sided (e.g., during market crashes), the market maker agent might temporarily suspend either buy-side or sell-side quoting, because it is restricted from running up large negative losses.

With these agent roles, we can then easily adjust the level of speculation by controlling the ratio between the random agents and the trend-following agents.

3 Implementation of Agent Strategies

In our study, we are interested in the profit-and-loss dynamics of trend-following agents under different levels of speculation (or equivalently, different levels of liquidity available), which will be

determined by the relative ratio of random agents to trend followers in the simulation.

As noted earlier, we have included three types of agents in our experiments: 1) random agents, 2) trend-following agents, and 3) one market-maker agent. A specific trading strategy has been implemented for each agent type. The primary outputs of these strategies are the bids/asks that will be sent by each agent to the market. A bid is simply a two-tuple that contains price and quantity, where positive quantity stands for buying (establishing *long* positions) and negative quantity stands for selling (establishing *short* positions). In our experiments, all submitted bids/asks are handled in real time by a central continuous double auction (CDA). Through this CDA, transactions are matched continuously as new bids come in and price quotes (both bid and ask quotes) are generated and sent back to all agents. Note that in a CDA, all active and valid bids that are not fully transacted will be stored in an order book, and a transaction happens whenever a buy bid has price that is higher than the lowest sell price, or a sell bid has price that is lower than the highest buy price. This implies that partial order matching is allowed as in real markets. The bid (or ask) quote of this CDA is simply defined as the highest (or lowest) price of all standing buy (or sell) bids.

Unlike most of the past research in market microstructure, there is no *informed trader* in our experiments. In our design, random agents resembling the behavior of retail investors create liquidity around current price quote, and trend-following agents resembling the behavior of professional investors simply follow a fixed set of indicators and rules when generating their trades. As such, besides those trend-following agents who mechanically trade according to formulas, all other agents believe that the initial market price is indicative of certain fair asset value consistent with the marginal cost of production of the traded commodity, and they will more or less demonstrate *mean-reverting* behaviors.

3.1 Trend-Following Agents

The trend-following agent tries to replicate the lookback option straddles by trading the underlying asset. Every 5-10 seconds, it will compute a delta based on the Black-Scholes formula:

$$d_1 = \frac{\ln(s/x) + (r + \sigma^2/2)t}{\sigma\sqrt{t}},$$

$$\Delta_{\text{call}} = \Phi(d_1), \quad \Delta_{\text{put}} = \Phi(d_1) - 1,$$

where s is the current price of the traded commodity, x is the strike price, r is the risk-free interest rate, t is the times in years until expiration (we assume that the simulation horizon is equivalent to one year for simplicity; the assumption is valid because the final payoff is independent of time to expiration and volatility), σ is the volatility, and $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. By using the delta from an option straddle to determine trading quantity, the desired position in period i will be $\Delta_i = (\Delta_{\text{call}} + \Delta_{\text{put}})$, i.e., $2\Phi(d_1) - 1$, which lies within the range of $[-1, +1]$ times the trading limit. In our experiments, the strike price is set to be the initial price, and the risk-free interest rate is assumed to be 1% p.a. The minimum volatility is set at 1% initially (purely to avoid nonsensical outputs during the initial period of trading), and is updated continuously by using an exponentially weighted moving average (EWMA) function. To update the volatility by EWMA, we will first compute the periodic returns by taking the natural log of the ratio of prices in the current period over the previous period, $u_i = \ln(S_i/S_{i-1})$. We then update the volatility by incorporating the periodic return:

$$\sigma_n^2 = \lambda\sigma_{n-1}^2 + (1 - \lambda)u_{n-1}^2.$$

Throughout our experiments, we set λ to 0.9, which is a standard choice that can produce numerically stable outputs. With the formula above, we can obtain Δ_i for every period i , and $(\Delta_i - Q)$, where Q stands for the current position, will be the differential quantity that this agent should adjust throughout trading (positive quantity stands for buying, negative quantity stands for selling). However, the agent should not submit its bid unless the bid is consistent with recent market trends, which is obtained by a simple linear regression with a minimum threshold of ± 0.03 (in other words, the agent will do nothing if the absolute value of the slope of the price trend is less than 0.03; this is necessary to avoid the agent “flip-flopping” based on insignificant trends). Our simple trading rule states that the sign of the proposed bid should be the same as the sign of the slope of the recent market trend. If this rule is not satisfied, the agent will not submit any bid.

To avoid trend-following agents becoming artificially homogeneous (which may defeat the purpose of having an agent-based model), we introduce two additional features, as follows:

1. Each agent is given a randomly generated horizon on how far back into the past should it go in collecting market prices for the execution of the linear regression. The horizon parameter is uniformly distributed between 20 and 60 price observations.
2. The quantity differential, $(\Delta_i - Q)$, cannot be executed by the agent if its value is too large; In each time period, the agent will compute a quantity that is uniformly generated between 50 and 250 and execute an order based on the smaller of $(\Delta_i - Q)$ and the random variable. Doing so ensures that orders are “broken up” into more manageable “lots”, which is consistent with real-world market practice.

3.2 Random Agents

The random agent is designed to provide the bulk of liquidity to the market. Its bidding interval is every 5-10 seconds, and it generates buy or sell bids according to a mean-reverting process: every 10 ticks in positive (negative) price difference increase the selling (buying) probability by 0.1, until reaching the upper bound at 0.8. The quantity of the bid is uniformly generated between 50 and 250 and the prices of the bids follow a normal distribution. The mean of the normal distribution is the average of bid and ask quotes, and the standard deviation is the volatility of the price sequence, which is updated using the same EWMA updating procedure employed by the trend-following agent.

3.3 Market Maker Agent

The market maker (MM) agent controls the underlying market movement and provides additional liquidity when liquidity dries up in the market. It achieves this by maintaining a bid list containing multiple bid points. For example, it can provide sell offers of Q_s units for all prices up to D_s ticks above a_0 at an interval of δ , and buy offers of Q_b units for all prices up to D_b ticks below b_0 at an interval of δ (all mentioned variables are configurable parameters). If there are no other agents in the system, this bid will result in ask and bid quotes of a_0 and b_0 , respectively. To control this bid list (i.e. list of bid points), we introduce an event-based mechanism that contains both bullish and bearish events. A bullish event will cause a positive shift of all bid points, resulting in an increase in both bid and ask quotes. On the other hand, a bearish event will cause a negative shift of all bid points, resulting in a decrease of both bid and ask quotes. For both bullish and bearish events, a strength parameter is provided so that the MM agent can determine how bullish (bearish) the response should be. The realization of the impact of event sequence in our experiments is plotted in Figure 1, where the general trend is bearish despite having ups and downs in the market.




Figure 1: A sample realization of the event sequence that is employed in all our experiments. Every point denotes an event. Despite the ups and downs, the general trend of this sequence is bearish.

The MM agent is assumed to be monopolistic in our trading simulation, and it will try to maximize its profits or minimize potential losses. The design of the MM agent can be quite sophisticated: for example, Das and Magdon-Ismail [2] have recently proposed a MM model that is based on Markov decision processes. We have chosen not to complicate the simulation with any such sophisticated design, since our primary focus is on liquidity issues. The MM agent is designed to expect a minimal profit margin of m for all its trades; however, if 10% of the absolute price differential (between initial and current prices) is greater than m , 10% of the absolute price differential will be used as the margin instead. Intuitively, this provides the MM agent with an easy way to find a balance between profit making and providing a service to the market.

Based on the design described above, one can easily compute that the initial total buying and selling volumes provided by the MM agent are $(D_b Q_b)$ and $(D_s Q_s)$, respectively. With properly set D_b and D_s , the MM agent can handle normal trading scenarios effectively while still maintaining reasonable profits. Unfortunately, when the market becomes too one-sided, the volume on either the buys or sells will be consumed quickly, and the MM agent will no longer be able to fulfill its role as a “gap filler” of the bid list. To handle this type of extreme movements, we insert an exponential schedule to allow the MM agent to inject more liquidity as the gap between current and initial prices increases. The function is in the form of $(D_x Q_x) \exp(\beta d)$, where $(D_x Q_x)$ denotes either initial buy-side or sell-side liquidity supply, β is the parameter controlling how strong the reaction should be for the MM agent to respond to price deviation, and d is the deviation (in percentage) of price from its initial level. In our experiments, β is defined such that when the price deviation is close to 100%, the MM agent can provide up to half of the total market liquidity. Doing so will be consistent with our assumption that the opening price of the market reflects certain fair asset value of the traded commodity based on its marginal cost of production.

4 Deviation from Theoretical Assumptions

The theoretical trend-following strategy implemented by using lookback straddles has been described in considerable details by Fung and Hsieh [3]. Fung and Hsieh use such a strategy to construct the so-called “Primitive Trend Following Strategy” (PTFS). PTFS is a key factor in explaining the returns of Commodity Trading Advisors (CTAs). However, PTFS is subject to a number of practical limitations:

1. PTFS tends to work more effectively on commodities, currencies, interest rates and bonds, but has not worked particularly well with stocks.
2. The replication works only based on the assumption that the required Black-Scholes options exist, which may or may not be the case for exchange-traded options, while the bid-ask spreads for over-the-counter options may be too expensive for any practical applications.
3. Many exchange-traded options are not European style options, or may have discrete strike points that do not match with the at-the-money strike that we intend to roll to.
4. The horizon of the options used is often determined by the availability of liquidity in the option calendar. In our experiments, we have assumed that the option expiration will be one year, simply out of convenience because the payoffs are not dependent on the time to expiration. Other “real-life” case studies tend to avoid options with expiration longer than one quarter to avoid illiquid options.
5. PTFS assumes that the manager has no issue in coming up with any upfront cash outlay. That may or may not be the case.

To make our simulations more realistic when constructing trend followers in our simulations, we have relaxed the assumptions of Fung and Hsieh by:

1. Not constantly rolling the strikes of the straddles;
2. Using the underlying to replicate the options, thereby circumventing any issues related to the liquidity of the options; and
3. Using regression betas to “smooth” the trends to avoid transaction costs from excessive buying and selling in violent markets with unclear directions.

The analyses in this section have taken reasonable steps to relax the original assumptions of Fung and Hsieh. To ensure that the results obtained from the agent-based simulation are reasonable, we attempt five different implementations of the trend-following trading strategy. For all these implementations, it is assumed that there will be sufficient market liquidity, and that the speculators’ actions have negligible effect on the market price. The differences among these implementations are based on how well the trader can replicate the theoretical lookback straddle strategy. As expected, when market becomes more and more dominated by speculators, results from simulations with direct participation by trend-followers will become increasingly worse than these five theoretical bounds. Since we are assuming for now that these implementations have no impact on market prices, we can use any “realized” price stream as inputs and construct the respective profit-and-loss curves for all five implementations. These *off-line* analyses can thus be used for comparison and calibration in the computational experiments that we plan to carry out.

Type A – Trend-Followers Replicated by Lookback Straddles. As pointed out by Goldman, Sosin and Gatto [4], the lookback straddle can be replicated by dynamically rolling the strikes of standard Black-Scholes straddles over the life of the option. Fung and Hsieh [3] provides an in-depth discussion and analysis based on this classical methodology. We will not attempt to repeat those discussions, except to provide a basic summary for the convenience of those readers who may be unfamiliar with the topic.

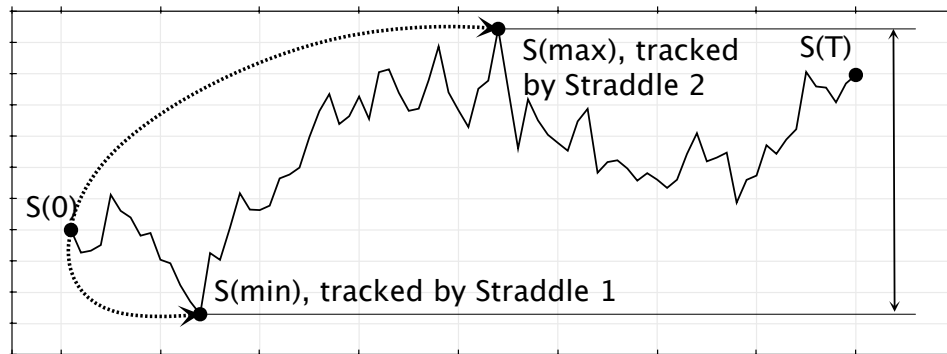


Figure 2: An illustration of the Type A strategy.

Generally speaking, lookback straddles can be thought of as the theoretical best-case scenario. In the perfect world, where there should be ample liquidity and no transaction cost, the Type A strategy can be perfectly implemented with the following payoff (as depicted in Figure 2):

$$|S_{\max} - S_{\min}| - \sum P_{\text{op}},$$

where S_{\max} and S_{\min} are the highest and the lowest observed prices during the investment horizon, and $\sum P_{\text{op}}$ is the total cost spent on rolling options. This payoff is achieved by continuously rolling the strike price of one straddle up whenever a new market high is achieved, and rolling the strike price of the other straddle down whenever a new market low is reached.¹ The transaction costs involved in any such continuous rolling make it very difficult to achieve this type of theoretical payoff in practice, since the required out-of-the-money or in-the-money options may not exist or liquidly traded; even if they do, the bid-ask spreads on the options may be too expensive to make this strategy practical. As we have explained in the footnote, we need to sell at the bid price and buy at the ask price continuously; a wide bid-ask spread implies an expensive “rolling” operation.

Type B – Stock Replication of Lookback Straddle. This is a variation of the Type A trader, in that the underlying commodity is bought/sold by delta hedging to replicate the lookback straddle. Given that such purchases/sales will not meet the continuous and frictionless market assumption of Black-Scholes, we expect that a Type B trader will experience performance deterioration as compared to a Type A trader.

Type C – Static Black-Scholes Straddle. The Static Black-Scholes Straddle is a buy-and-hold strategy that simply purchases one at-the-money straddle at the beginning of the simulation

¹In options trading, the rolling of the strike price means that we sell the existing option and buy a new option at a different strike price. The cost of such a transaction is primarily the bid-ask spread from the simultaneous sales and purchases of the two options, since we have to sell at the bid and buy at the ask, and thus the difference in bid and ask prices (i.e. the *spread*) will be the additional cost we need to pay to “roll” the strike price up or down.




Figure 3: An illustration of the Type C strategy.

period. The payoff for the Type C implementation is (as depicted in Figure 3):

$$|S_t - S_0| - \sum P_{op},$$

where S_t is the price at option maturity, S_0 is the initial price, and $\sum P_{op}$ is the premium paid for options. The intuition is that, when traders do not have perfect foresights to pick the tops and the bottoms in any market, while the continuous rolling of Black-Scholes options may be too expensive, one feasible alternative is to simply purchase a static straddle and hold it to maturity. Such a trader’s belief is that the beginning and the end of the investment horizon will provide a reasonable approximation of the predominant market trend during that period.

Type D – Stock Replication of Static Black-Scholes Straddle. This is a variation of the Type C trader, in that the underlying commodity is bought/sold by delta hedging to replicate the Black-Scholes straddle. Given that such purchases/sales will not meet the continuous and frictionless market assumption of Black-Scholes, we expect that a Type D trader will experience performance deterioration as compared to a Type C trader. One can also think of Type D stock replication as a simplistic, tick-by-tick form of trend following, since a Type D trader will be buying or selling based on a uptick or a downtick, respectively.

Type E – Simplistic Regression-Beta Trend Following. The final type of trader uses a simple regression (say over 30 price observations) to determine whether the market is trending upward or downward. If the regression beta is larger than a specific threshold, then the strategy will purchase/sell the amount suggested by the delta in Type D trader.

By applying practical tweaks to the theoretical strategy of a Type A trader, we expect to see different levels of deterioration of trading performance from Type B to Type E. By feeding these five implementations with a price stream generated by human participants (detail of which will be discussed in the next section), we can indeed observe such deterioration (see Figure 4).

5 Percentage of “Smart Traders” versus “Random Traders”

It should be noted that the simulation results from the previous section are hypothetical profits and losses computed based on an already “realized” price stream, and that none of these hypothetical agents participated in the trading simulation directly. The price stream as shown in Figure 4 is taken from a simulation involving human participants. Unusually interesting insights can be




Figure 4: P&L of different trader types (Type A-E) showing increasing performance deterioration. u_2 , u_{14} , and u_{36} are P&L streams of selected human participants.

obtained by analyzing the participants’ trading patterns and performances; in fact, such insights have motivated our formal analysis framework. We will present our findings from this simulation before reporting our main result.

5.1 Simulation with Human Participants

The original simulation was conducted for a group of undergraduate students attending an oil trading workshop. Prior to the simulation session, students were exposed to 8 weeks of intensive instructions focusing on different aspects of crude oil trading.

During the simulation, participants played speculators and traded for profits. The targeted commodity was the Western Texas Intermediate (WTI) crude oil contract and each trader could trade up to a maximum of 2,000,000 barrels (bbls). A trader could take either long or short positions, but once he reached his trading limit, he could take no further positions in the same direction (however, he could still trade to reduce his positions). To encourage active participation, we explicitly asked each participant to transact at least 100,000 bbls in each trade, and at least 2,000,000 bbls during the entire simulation. The contract size was 1,000 bbls.

After the simulation, we computed the P&L for all participants. Only 8 out of 39 participants (around 20%) managed to make profits; surprisingly, almost 60% of all profits were earned by the top 2 traders. After examining the trading patterns of these two traders, we find out that both of them followed some sort of trend following strategies (one of them is illustrated in Figure 5 (a)). The P&L streams of these two users are also plotted in Figure 4 (denoted as u_{14} and u_{36}). We can see that their performances are similar to the theoretical upper bound of trend followers. These human trend followers were able to succeed in the simulation, primarily because most other participants traded randomly (e.g., the trading pattern of the worst trader is illustrated in Figure 5

(b), which is not noticeably different from a totally random computer agent; the P&L stream of this trader is denoted as u_2 in Figure 4). Such seemingly random trading activities create ample liquidity for the trend-following traders to generate profits. After all, it is not possible for these human trend followers to generate profits without other traders who are willing to take the other side of their trades.

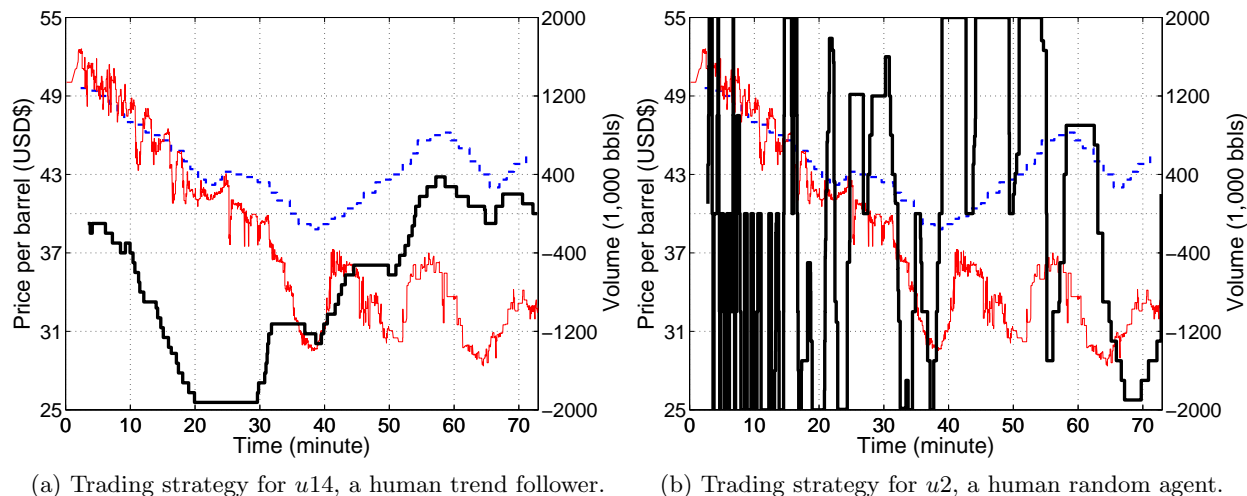


Figure 5: The realized trading strategies: thick line is trader’s position balance; thin line is real market price; dotted line is estimated market price dynamics.

This observation provides the primary motivation for us to carry out these elaborate computational simulations with different levels of speculative activities. It is natural to ask whether the success of these “smart traders” are indeed related to the level of trading liquidity available; if so, then one should expect a potential collapse of their performances once liquidity becomes increasingly scarce.

5.2 Replacing Human Participants with “Smart Traders”

In order to study the impact of liquidity levels, it is necessary to perform the simulations with computer agents participating directly in the simulation game, and to gradually “remove” liquidity from the market simulation. In this sub-section, we ask the following question: What if the human participants in the previous sub-section are replaced by computerized agents consisted of both systematic trend followers and random agents?

The following control experiments all use 50 computer agents and 1 market maker. In each of the experiments, we vary the percentage of “smart traders” versus “random traders”. The “smart traders” use the trend-following algorithm of Type E traders in the previous section, but with randomized parameters (under uniform distributions) in:

1. Forecasting horizon;
2. Quantity per bid; and
3. Bidding interval.

The configuration of each experiment is listed in Table 1. The resulting price streams from the control set of simulations can be found in Figure 6. The histograms of the price return distributions

Simulation	Random Agents	Trend Followers	Market Maker	% of Trend Followers
1	40	10	1	20%
2	30	20	1	40%
3	20	30	1	60%
4	10	40	1	80%

Table 1: The configuration for the four simulations.

from the price streams can be found in Figure 7. For comparison, the histogram of the price return distribution of West Texas Intermediate (WTI), the leading crude oil spot contract traded in the US, can be found in Figure 8. Given the well-known technical difficulties in calibrating non-normal price return distributions, it is surprising how the shapes of these distributions are visually similar to distributions from real-life commodity prices.² The key statistics of those price streams are listed in Table 2, while the key statistics of the WTI spot prices and the experiment with human participants are listed in the last two rows of the same table for the ease of making comparisons. Do note that in the control set of simulations and the experiment with human participants, market returns are computed by sampling the price streams every 10 seconds, thus zero returns are much more common than it should have been if the experiment can be allowed to run for hours if not days. To eliminate this artificial effect caused by the sampling interval, zero returns are removed when computing the statistics shown in Table 2.

The rolling exponentially-weighted moving average volatility, with sample interval of 1 second, of prices from the control set of simulations can be found in Figure 9.

% of TF	Mean	Median	StdDev	Kurtosis	Skewness	Min	Max
20%	-0.0002	0.0011	0.0074	5.7333	-0.7792	-0.0353	0.0266
40%	-0.0013	-0.0012	0.0082	5.2083	-0.8238	-0.0361	0.0290
60%	-0.0012	-0.0012	0.0088	5.6394	-0.6163	-0.0426	0.0324
80%	-0.0042	-0.0046	0.1526	9.3048	-0.0360	-0.7963	0.9163
WTI	0.0027	0.0074	0.0469	3.0465	-0.5323	-0.1923	0.2512
Human-Exp	-0.0012	-0.0012	0.0245	4.1229	-0.2304	-0.1148	0.1134

Table 2: Key statistics of price streams from the control set of simulations (TF stands for Trend Followers), WTI crude oil price return distribution (June 1999 – Jan 2010), and the human experiment.

5.3 Analysis

We will provide selected charts generated by our simulation in the interest of space. The profits and losses of the control set of simulations are plotted in Figures 10, 11, 12 and 13. Analyzing these results suggest the following:

1. In a market with finite liquidity, it is not possible for all players to make money. After all, trading is a zero-sum game. In order for any “smart” trader to generate profits, there must be other market players who are willing to take the other side of their trades.

²Calibration of non-normal distribution is one of the topics discussed in Lee and Lee [7].






Figure 6: Price streams from the control set of simulation. $p\%$ TF refers to the case of having $p\%$ trend followers in the simulation.

2. In general, our trend-following agents generate profits and losses that are upper-bounded by Type A traders (the lookback straddle as proposed by Fung and Hsieh [3]), except in cases where the market becomes extremely unstable and erratic.
3. Furthermore, our trend-following agents generate profits and losses that increasingly underperform those of its counterpart Type E traders (which are theoretical and have not participated in the simulation) as and when liquidity is withdrawn from the market.
4. Simulation 4 represents a breakdown of the market. When there are limited liquidity and extreme price shocks in the market, it may not be possible for trend-followers to execute their strategies by making simultaneous buys and sells in a one-sided market, unless there are liquidity providers who are willing to take the other side of their trades. That will be rare or non-existent by construction.
5. In all cases, trading essentially stops before the end of the trading simulation, after trend-followers have reached their position size limits. While orders are still being matched between the random agents and the market maker, the P&L of the trend followers will be based on static positions from that point forward.

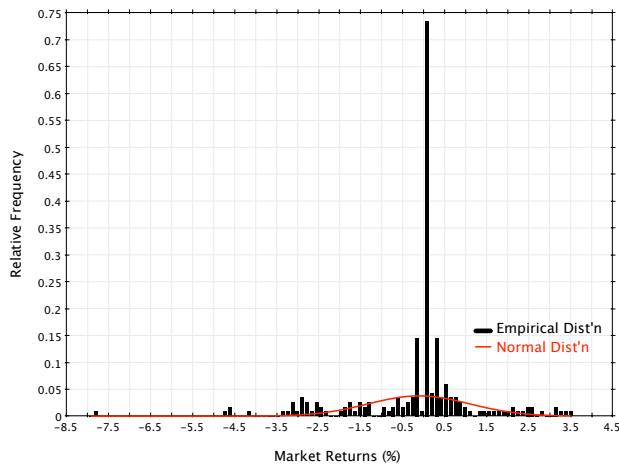
These outcomes suggest that the dynamics generated by our trading game are consistent with the expected experience in real-world markets. Our logical next step is to study the statistical properties of the pricing streams produced by these simulations. Using the profit-and-loss charts and standard statistical analyses on market price streams available in Table 2, we observe the following trends:



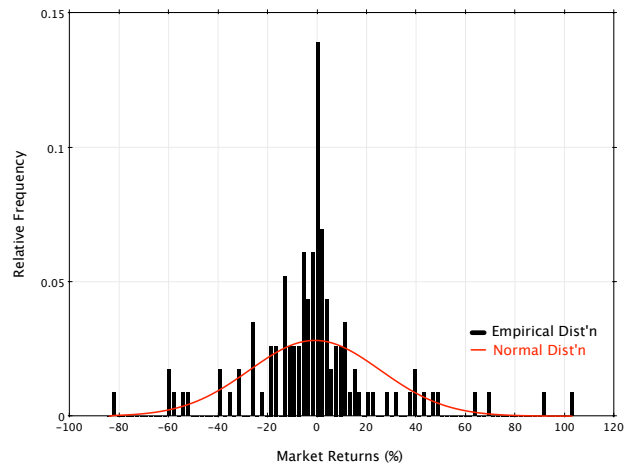
(a) 20% Trend Followers



(b) 40% Trend Followers



(c) 60% Trend Followers



(d) 80% Trend Followers

Figure 7: Return distribution from the control set of simulations.

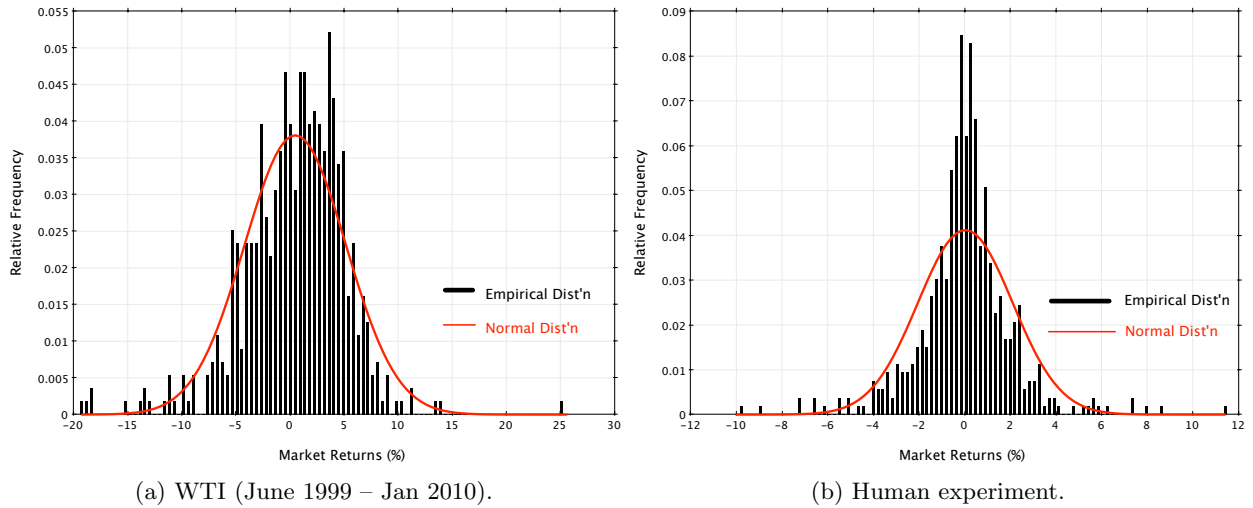


Figure 8: The return distribution for both WTI and the experiment with human participants (based on the price stream illustrated in Figure 4).




Figure 9: Rolling EWMA volatility of price Streams from the control set of simulations (sample interval = 1 second).

1. As and when liquidity is removed from the market, volatility will increase and extreme price shocks become more prominent. This is consistent with empirical observations in the market, in that: 1) commodity futures markets that are increasingly dominated by speculative activities as contracts “roll in” tend to exhibit increasing volatility; and 2) markets such as that of convertible bonds, which is 95% dominated by professional speculators, (i.e. hedge funds and proprietary trading desks, as opposed to long-term investors) show significantly higher volatility than other markets.
2. The kurtosis of the market actually decreases when the market is dominated by more and more trend followers. Nonetheless, all simulation results show significant fat-tail behavior. More theoretical discussion on the modeling of fat-tail behavior in financial markets can be found in Lee and Lee [7].
3. Once the market has gone over a “tipping point” of destabilization (e.g. as in simulation 4), the instability usually does not become significantly worse by further withdrawal of liquidity. The key is that the market must find a new equilibrium of lower (or higher) prices when there is an overwhelming majority of sellers (or buyers). This is consistent with empirical observations by real-life practitioners.

Our next step is to look for evidence as to whether specific techniques of intervention may or may not be effective in a market facing liquidity crises.




Figure 10: P&L of different strategies in simulation 1 with 40 random agents (denoted as RA), 10 trend followers (20% trend followers, denoted as TF) and 1 market maker (denoted as MM).




Figure 11: P&L of different strategies in simulation 2 with 30 RA, 20 TF (40% TF) and 1 MM.




Figure 12: P&L of different strategies in simulation 3 with 20 RA, 30 TF (60% TF) and 1 MM.




Figure 13: P&L of different strategies in simulation 4 with 10 RA, 40 TF (80% TF) and 1 MM.

6 Volatility-Driven Changes to Position Size Limits

In this section, we want to better understand whether it may be effective to stabilize market prices by controlling position size limits. Specifically, a policy proposal was made by the U.S. Commodity Futures Trading Commission to reduce position size limits, in response to what was described as market stabilization by a handful of speculators holding out-sized positions in the commodities market. We seek to ask the more general question: Is there any analytical evidence suggesting that controlling position size limits can be an effective policy tool to stabilize market prices?

In our analysis, a volatility spike is defined as each instance in which the rolling exponentially-weighted moving average (EWMA) volatility exceeds 5%. For completeness, we study both the first case in which the position size limits of the traders are halved, as well as the second case in which the position size limits of the traders are doubled. Any change to the position size limits is reversed as soon as the rolling EWMA volatility drops below 5%.

6.1 Halving Position Size Limits

The resulting price streams from the set of simulations with halving position size limits can be found in Figure 14. The key statistics of those price streams are listed in Table 3. The rolling exponentially-weighted moving average volatility (with sample interval of 1 second) of the price streams can be found in Figure 15.

The profits and losses of the “60% Trend Followers” scenario are plotted in Figure 16. It should be noted that *not* all pricing streams of the “60% Trend Followers” scenarios necessarily trend downward towards the \$10 range, but we have managed to identify one such path. In this experiment, a rapid drop of the price stream has resulted in the trend followers having to cut short positions due to the halving of position size limits. Prices then restore to its previous levels due

to “short covering”. However, as soon as volatility drops below 5% and the original position size limits are restored, the market crashes again. The pattern repeats itself until the market stabilizes at around the \$40 range. In short, we observe that, rather than stabilizing markets, such a policy may result in “thrashing” whereby prices move from one extreme to the other, both at the halving of position size limits and at the restoration of the original limits.

Similarly, in the “80% Trend Follower” scenario, where prices fall because of the sellers’ domination, the policy results in erratic price behavior instead of stabilizing prices each time when there is a sharp volatility spike.

% of TF	Mean	Median	StdDev	Variance	Kurtosis	Skewness	Min	Max
20%	0	0	0.00446	0.0020%	23.398	-1.05	-0.0333	0.0320
40%	-0.0001	0	0.00449	0.0020%	20.718	-0.80	-0.0322	0.0298
60%	-0.0001	0	0.03653	0.1334%	121.999	3.71	-0.5448	0.5521
80%	-0.0005	0	0.08204	0.6731%	23.366	1.06	-0.5445	0.8214

Table 3: Key statistics of price streams from the set of simulations with halving position size limits. TF stands for trend followers.




Figure 14: Price streams after positions size limits are halved based on volatility.

6.2 Doubling Position Size Limits

The resulting price streams from the set of simulations with doubling position size limits can be found in Figure 17. The key statistics of those price streams are listed in Table 4. The rolling exponentially-weighted moving average volatility (with sample interval of 1 second) of the price streams can be found in Figure 18.




Figure 15: Rolling EWMA volatility of price streams when position limits are halved (sample interval = 1 second).




Figure 16: P&L of different strategies when position limits are halved (with 20 RA, 30 TF (60% TF), and 1 MM).

In this analysis, the volatility spikes in the “60% Trend Followers” scenario are even more pronounced. In the “80% Trend Followers” scenario, the price stream becomes erratic as trend followers are allowed to increase their ability to take short positions in volatility spikes, resulting in a “death spiral” of prices to near zero levels.

In general, our simulation results have shown that changing the rules of the game in the midst of trading can be quite problematic. For instance, it is difficult to define clear guidelines on how quickly any trader exceeding his/her new limit must cut positions during a volatility spike, and forcing a rapid liquidation of positions can result in even more volatility. Once the initial volatility spike has finally subsided, the restoration of the original position size limits may result in a flood of new orders, thereby destabilizing prices once again. Based on our analysis, controlling position size limits does not appear to be a particularly effective policy tool. It may be more practical to intervene by direct/indirect market intervention, as we will show in the next section.

% of TF	Mean	Median	StdDev	Variance	Kurtosis	Skewness	Min	Max
20%	0	0	0.00453	0.0021%	23.783	-1.04	-0.0419	0.0289
40%	-0.0002	0	0.00529	0.0028%	22.605	-1.63	-0.0426	0.0324
60%	0	0	0.02367	0.0560%	254.081	12.16	-0.1692	0.5124
80%	-0.0034	0	0.12722	1.6185%	170.802	2.61	-1.9308	2.0053

Table 4: Key statistics of price streams from the set of simulations with doubling position size limits. TF stands for trend followers.




Figure 17: Price streams after positions size limits are doubled based on volatility.




Figure 18: Rolling EWMA volatility of price streams when position limits are doubled (sample interval = 1 second).

7 Liquidity Reduction/Injection

In this section, we want to better understand whether it may be effective to stabilize market prices by various forms of liquidity control. Specifically, direct/indirect³ liquidity injection was frequently used as a policy tool by central banks during the most recent financial crisis. We seek to better understand the mechanics of using liquidity injection as a policy tool to stabilize market prices.

In this analysis, a volatility spike is also defined as each instance in which the rolling exponentially-weighted moving average (EWMA) volatility exceeds 5%. For completeness, we study both the first case in which the market maker is forced to reduce the amount of liquidity that it can provide to the market, as well as the second case in which the market maker is injecting liquidity into the market. Any change to the market maker’s ability to provide liquidity is reversed as soon as the rolling EWMA volatility drops below 5%.

7.1 Liquidity Reduction

In this analysis, the market maker’s maximal ability to provide liquidity is reduced from 50,000 contracts to 25,000 contracts whenever exponentially-weighted moving average (EWMA) volatility exceeds 5%. The resulting price streams from the set of simulations under liquidity reduction can be found in Figure 19. The key statistics of those price streams are listed in Table 5. The rolling EWMA volatility (with sample interval of 1 second) of the price streams can be found in Figure 20.

The profits and losses of the “60% Trend Followers” scenario are plotted in Figure 21. Interestingly, by reducing the ability for dominant smart traders to accumulate massive positions, market prices are stabilized. Moreover, the market maker is able to improve its P&L due to the sudden

³Although it is uncommon for central banks to directly intervene in markets outside of FX, interest rate and (on very rare occasions) equity markets, regulators can nonetheless intervene indirectly by controlling the amount of trading liquidity available to active market participants such as hedge funds by adjusting the amount of leverage allowed in the banking system. A detailed discussion of this complex process is beyond the scope of this paper, but can be found in a piece authored by Lee in a forthcoming book on hedge funds and systematic risk edited by Fung, Hsieh and Naik, which will be published by the CFA Research Foundation.

reversals of the massive shorts held by the trend followers (i.e. a classic short squeeze). Thereafter, market prices stabilize and there is no “thrashing” of prices in this case.

In the “80% Trend Follower” scenario, prices simply collapse to near-zero levels when the buying orders from the short squeeze cannot counter the overwhelming selling pressure due to the domination of trend followers. Simply put, a short squeeze can only be applied as a symptomatic cure, but it will not fix any fundamental imbalance in the market.

% of TF	Mean	Median	StdDev	Variance	Kurtosis	Skewness	Min	Max
20%	0	0	0.00431	0.0019%	24.139	-0.69	-0.0283	0.0372
40%	-0.0001	0	0.00461	0.0021%	25.336	-1.00	-0.0365	0.0298
60%	-0.0002	0	0.02080	0.0433%	281.453	7.09	-0.3593	0.3898
80%	-0.0038	0	0.10597	1.1230%	175.755	1.16	-1.8458	1.8458

Table 5: Key statistics of price streams from the set of simulations under liquidity reduction. TF stands for trend followers.




Figure 19: Price streams under liquidity reduction.

7.2 Liquidity Injection

In this analysis, the market maker’s maximum ability to provide liquidity is increased from 50,000 contracts to 100,000 contracts whenever exponentially-weighted moving average (EWMA) volatility exceeds 5%. The resulting price streams from the set of simulations under liquidity injection can be found in Figure 22. The key statistics of those price streams are listed in Table 6. The rolling EWMA volatility (with sample interval of 1 second) of the price streams can be found in Figure 23.




Figure 20: Rolling EWMA volatility of price streams when liquidity is reduced (sample interval = 1 second).




Figure 21: P&L of different strategies when liquidity is reduced (with 20 RA, 30 TF (60% TF), and 1 MM).

The profits and losses of two different “60% Trend Followers” scenarios are plotted in Figure 24 – in the “Run 1” simulation path, a violent price shock is experienced before prices recover and stabilize, while in “Run 2” prices are relatively stable throughout the simulation. The profits and losses of the “80% Trend Follower” scenario are plotted in Figure 25. While there is no obvious reduction of volatility, the market trends downward due to selling pressure, and stabilizes after finding a new (but much lower) equilibrium level.

It should be noted that this type of “market rescue” effort may also result in the market maker running up a large negative P&L in any one-side market, as shown in Figure 25. Unless governments are willing to step in to underwrite the large potential losses of market makers, none of them may volunteer to participate in any such market rescue efforts.

% of TF	Mean	Median	StdDev	Variance	Kurtosis	Skewness	Min	Max
20%	0	0	0.00444	0.0020%	25.229	-1.29	-0.0371	0.0311
40%	-0.0001	0	0.00463	0.0021%	26.676	-1.72	-0.0387	0.0319
60% (Run 1)	0	0	0.02367	0.0560%	254.081	12.16	-0.1692	0.5124
60% (Run 2)	-0.0001	0	0.00493	0.0024%	24.040	-1.84	-0.0395	0.0355
80%	-0.0017	0	0.10136	1.0274%	141.775	-3.89	-1.9459	0.9968

Table 6: Key statistics of price streams from the set of simulations under liquidity injection. TF stands for trend followers.




Figure 22: Price streams under liquidity reduction based on volatility.




Figure 23: Rolling EWMA volatility of price streams when liquidity is injected (sample interval = 1 second).




Figure 24: P&L of different strategies when liquidity is injected (with 20 RA, 30 TF (60% TF), and 1 MM).




Figure 25: P&L of different strategies when liquidity is injected (with 10 RA, 40 TF (80% TF), and 1 MM).

8 Conclusions

In our final section, we will conclude by summarizing the analytical and policy implications of our findings.

1. If extreme price shocks can be linked to liquidity, then some of the recent research and policy analysis efforts that focus on predicting the next crisis by using extreme value theory on historical prices and related statistical data may be wrong-headed. As long as there is no constant mix of market players, it is *unlikely* that the next crisis will be driven by liquidity factors manifesting themselves under conditions comparable to those seen in the recent past. So, making any predictions based on a statistical distribution derived from recent historical data may not be terribly meaningful.
2. If one key factor driving extreme price shocks is the proportion of speculative activities (i.e., the portion of speculators versus long-term investors), then enforcing position size limits per se (as recently proposed by the U.S. Commodity Futures Trading Commission) may not be an effective policy tool in mitigating the breakdown of markets. In fact, a direct policy solution is for a liquidity provider to step in whenever the market becomes dangerously unbalanced. However, doing so may require governments that are willing to underwrite large potential losses in any such market rescue efforts.
3. Finally, the potential existence of a tipping point of market destabilization, as observed in our many simulation runs, is a helpful reminder to policy makers that systematic risk may become extremely difficult to contain beyond a certain tipping point. Policy makers may like to reconsider their role to be one of deploying preemptive measures before approaching

the tipping point, rather than looking for suitable corrective measures only after the tipping point has already been breached.

Any stylized model of market contagion should accommodate some or all of the characteristics described above. Such a framework will likely be an extension of the approach similar to that of Kyle and Xiong [5] by increasing the number of tradable assets, the policy options available to the central bank and the types of investors in the financial economy.

A flexible agent-based simulation platform will be a valuable tool to validate any follow-up research, since additional counterparties, such as central banks or additional investor types, can be easily incorporated by introducing new agent types. All in all, this agent-based simulation tool has provided a flexible and easy-to-use suite of tools for policy makers to evaluate the possible outcomes of market policies. Although the tool is not meant to replace classical econometric models, it will supplement classical analyses with additional insights into agent-to-agent and agent-to-policy interactions.

9 Acknowledgements

The authors want to thank Srikanthan Natarajan and Peter Zheng for their able research assistance, as well as Professor William Fung of the London Business School for his kind comments and suggestions. Any errors and omissions remain the sole responsibility of the authors. This work was supported in part by various academic units within the Singapore Management University, including the Provost's Office, International Trading Institute, School of Information Systems, and Wharton-SMU Research Centre under Grant C220/MSS7C008. In addition, the authors acknowledge generous financial support from International Enterprise Singapore as well as from the Finance Graduate Immersion Programme run by the Monetary Authority of Singapore.

References

- [1] S.-F. Cheng and Y. P. Lim. An agent-based commodity trading simulation. In *Twenty-First Annual Conference on Innovative Applications of Artificial Intelligence*, pages 72–78, Pasadena, CA, 2009.
- [2] S. Das and S. Magdon-Ismail. Adapting to a market shock: Optimal sequential market-making. In *Advances in Neural Information Processing Systems 21*, pages 361–368. 2008.
- [3] W. Fung and D. A. Hsieh. The risk in hedge fund strategies: theory and evidence from trend followers. *Review of Financial Studies*, 14(2):313–341, 2001.
- [4] M. B. Goldman, H. B. Sosin, and M. A. Gatto. Path dependent options: “buy at the low, sell at the high”. *The Journal of Finance*, 34(5):1111–1127, 1979.
- [5] A. S. Kyle and W. Xiong. Contagion as a wealth effect. *The Journal of Finance*, 56(4):1401–1440, 2001. 10.1111/0022-1082.00373.
- [6] B. LeBaron. A builder's guide to agent-based financial markets. *Quantitative Finance*, 1:254–261, 2001.
- [7] Y. Lee and B. Lee. How “Sharpe” are funds of funds? In B. Schachter, editor, *Intelligent Hedge Fund Investing*. Risk Books, 2004.

- [8] K. M. Lochner, S.-F. Cheng, and M. P. Wellman. AB3D: A market game platform based on flexible specification of auction mechanisms. Technical report, University of Michigan, 2007.
- [9] L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.
- [10] L. Tesfatsion. Agent-based computational economics: A constructive approach to economic theory. In T. Leigh and L. J. Kenneth, editors, *Handbook of Computational Economics, Volume 2: Agent-based Computational Economics*. Elsevier, 2006.