Review of Futures Markets

Volume 19 Special Edition

Research funding for this Special Edition was made possible by a generous donation from the Clearing Corporation Charitable Foundation to the Institute for Financial Markets.

IFM THE INSTITUTE FOR FINANCIAL MARKETS

In Cooperation with The IFM

WOULD PRICE LIMITS HAVE MADE ANY DIFFERENCE TO THE "FLASH CRASH" ON MAY 6, 2010?

Bernard Lee, Shih-fen Cheng, and Annie Koh*

On May 6, 2010, the U.S. equity markets experienced a brief but highly unusual drop in prices across a number of stocks and indices. The Dow Jones Industrial Average (see Figure 1) fell by approximately 9% in a matter of minutes, and several stocks were traded down sharply before recovering a short time later. The authors contend that the events of May 6, 2010 exhibit patterns consistent with the type of "flash crash" observed in their earlier study (2010). This paper describes the results of nine different simulations created by using a large-scale computer model to reconstruct the critical elements of the market events of May 6, 2010. The resulting price distribution provides a reasonable resemblance to the descriptive statistics of the second-by-second prices of S&P500 E-mini futures from 2:30 to 3:00 p.m. on May 6, 2010. This type of simulation avoids "over-fitting" historical data, and can therefore provide regulators with deeper insights on the possible drivers of the "flash crash," as well as what type of policy responses may work or may not work under comparable market circumstances in the future. Our results also lead to a natural question for policy makers: If certain prescriptive measures such as position limits have a low probability of meeting their policy objectives on a day like May 6, will there be any other more effective counter measures without unintended consequences?

---

*Bernard Lee (the corresponding author) is a visiting associate professor (Practice) and Deputy Director, Sim Kee Boon Institute for Financial Economics, at Singapore Management University. E-mail: bernardlee@smu.edu.sg.

Shih-fen Cheng is an assistant professor in the School of Information Systems at Singapore Management University. E-mail: sfcheng@smu.edu.sg.

Annie Koh is an associate professor of finance and Dean, Office of Executive and Professional Education; Academic Director, International Trading Institute, at Singapore Management University. E-mail: anniekoh@smu.edu.sg.

Acknowledgements: The authors thank Ms. Zhan Yingbo of the Quantitative Finance Program at the National University of Singapore for her able research assistance. In addition, the authors gratefully acknowledge the generous financial support by the Institute for Financial Markets, as well as the opportunity to present our findings at the 2011 Asia Pacific Futures Research Symposium. We also want to thank the Monetary Authority of Singapore and the Singapore Exchange for inviting us to discuss our work at a Regulators' Meeting in conjunction with the 2010 Futures Industry Association Asia Conference. Finally, the authors thank both the Chicago Mercantile Exchange and the Singapore Exchange for providing certain primary sources of data. Any additional errors and omissions remain our own.

Keywords: flash crash, May 6, 2010, cascading market failures, complex systems
JEL Classification: G15, G17, G18

There are many publicly-available accounts of the market events of May 6, 2010. We will not attempt to repeat those accounts here. We will aim to provide a relatively straightforward summary, for the purpose of setting the proper context of our simulation analysis. Given that we are simply summarizing basic facts for the convenience of our audience, we would like to acknowledge the relevant sources all at once, including the "Joint CFTC-SEC Preliminary Report" and its corresponding "Final Report" (CFTC 2010a,b), as well as a research report published by the CME Group shortly after the May 6, 2010 incident (CME Group 2010). In addition, we have benefited from primary sources of data provided by the CME Group as well as the SGX.[1]

The trading day of May 6, 2010 started with unsettling political and economic news due to the European debt crisis. Just one day before, the Greek government's debt crisis boiled over into violence on the street of Athens. These factors had weighed on global markets before U.S. trading hours, and the U.S. equity market was down in early trading. At around 2:30 p.m. (all times are shown in Eastern Standard Time), the overall decline suddenly accelerated, after a rush of sell orders. Within a few minutes, both the S&P 500 Index and its June 2010 E-mini futures dropped by more than 5% (shown in Figure 2).

Staff of the Commodity Futures Trading Commission (CFTC) conducted a post-mortem analysis of the top 10 largest longs and shorts. Those analyses suggest that, in most cases, traders with the largest longs and shorts in fact traded on both sides of the market. In other words, there was no obvious one-sided "squeezing" of the market. The CME order books on futures also showed that there were many more sell orders than buy orders from 2:30 to 2:45 p.m. However, the volume of E-mini futures surged to eight times that of SPDRs (after adjustments) between 2:45 and 2:50 p.m. To most traders, this was a clear indication that the futures market was driving the cash market, not the other way around.

The bid-ask of the June 2010 E-mini S&P 500 futures widened considerably at about 2:45 p.m., triggering CME's Globex stop logic functionality. The stop logic functionality aims to prevent the triggering of stop-loss orders that would have resulted in transactions at price levels below the contract's "no-bust range," leading to an avalanche of price declines due to order-book imbalances. This functionality put the market in a "reserve" state when orders could be entered, modified, or cancelled but not concluded. It was, in fact, triggered in the E-mini market at 2:45:28 p.m. for five seconds, precisely when the E-mini contract hit its low of the day. Since futures were not traded during these five seconds, the linkages between the cash and the futures markets would have broken down despite that, in theory, U.S. stock futures that are traded on the CME are supposed to be coordinated with cash equity trading on the New York Stock Exchange (NYSE).

The majority of the single-name stocks had declines consistent with the 5% decline in June 2010 E-mini S&P 500, which traded at its low of 1056 by 2:34:28 p.m. However, three stocks — namely, Proctor and Gamble (PG), 3M (MMM),

1. The authors gratefully acknowledge the help from John Labuszewski of the CME Group as well as that of Sutat Chew from the Singapore Exchange.

Figure 1. DJIA, Dollar-Yen, Nikki and Gold Prices on May 6, 2010.

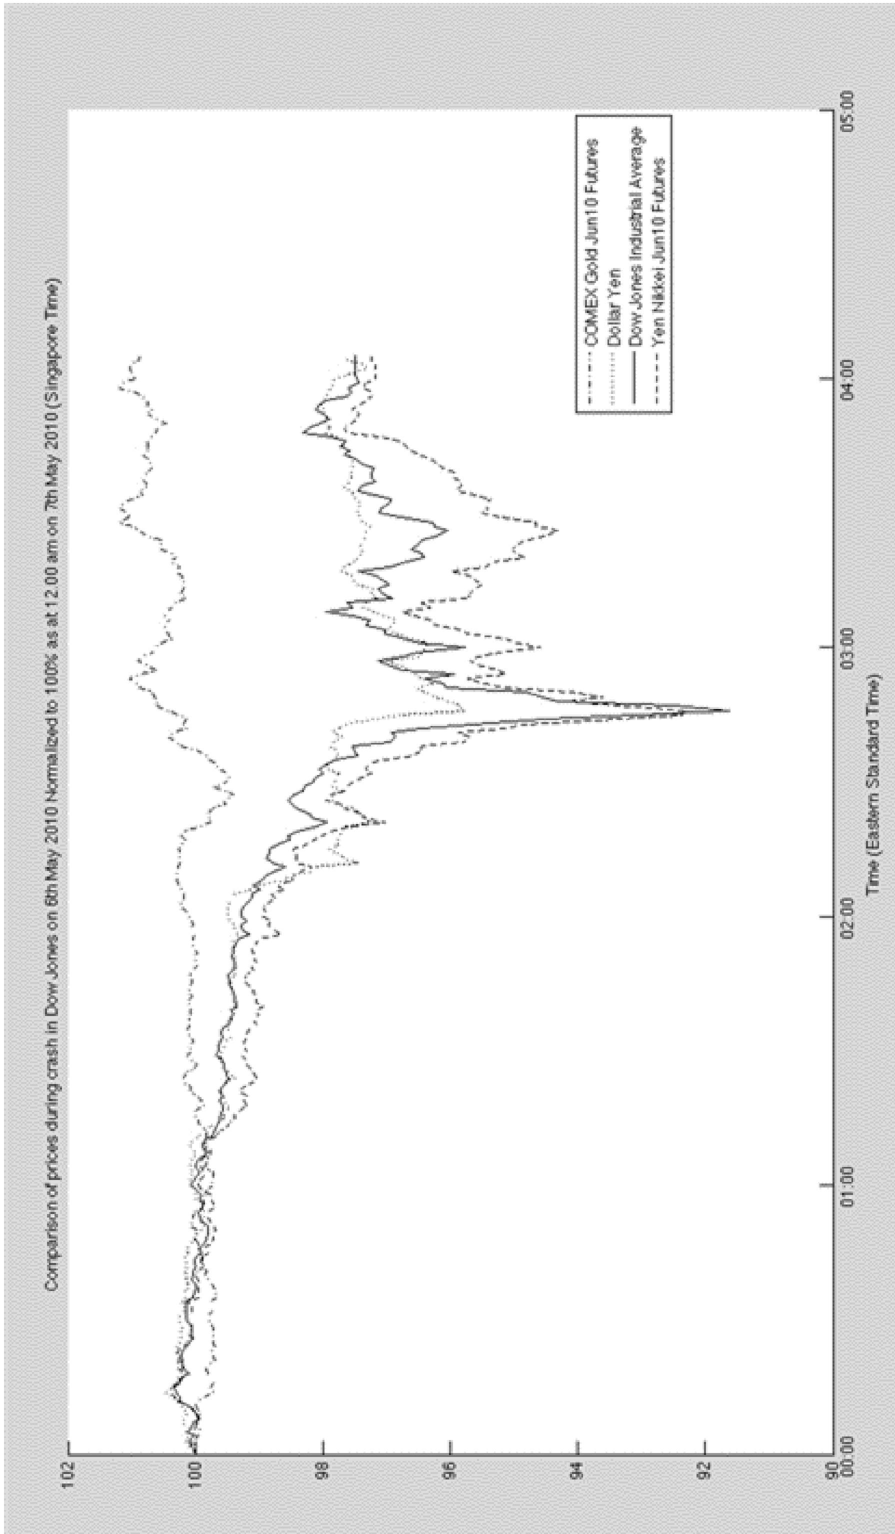

(Data courtesy of SGX.)

Figure 2. June 2010 E-mini futures on S&P 500 vs. SPDRs.

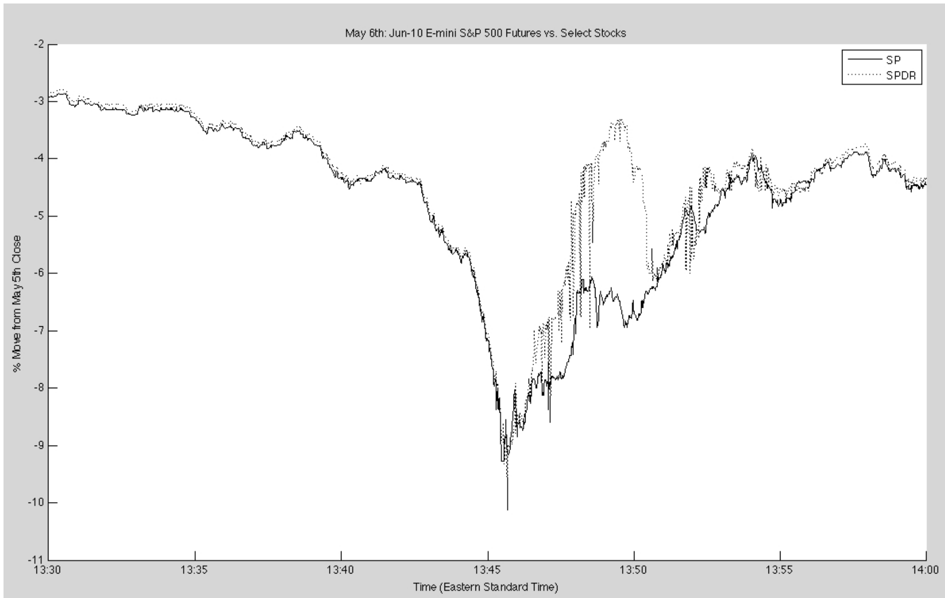

(Data courtesy of CME Group.)

and Accenture (ACN) — continued to decline even as the E-mini S&P 500 contract hit its low and then began to reverse upward (see Figure 3). These three stocks hit their Liquidity Replenishment Points (LRPs) at 2:45:52 p.m., 2:50:36 p.m., and 2:46:10 p.m., respectively, while their lowest trading prices of \$39.37, \$67.98, and \$0.01 were reported at 2:47:15 p.m., 2:45:47 p.m., and 2:47:54 p.m., respectively.

Eventually, Nasdaq announced that it would bust all trades that were more than 60% off the market. Of the U.S.-listed securities with declines of 60% or more away from the 2:40 p.m. transaction prices (resulting in busted trades), approximately 70% were ETFs. This observation suggested that ETFs as an asset class were affected more than any other categories of securities. One hypothesis is that ETF might have been widely used by investors as inexpensive short hedges and in placing stop-loss market orders.

Several hypotheses were raised by the "CFTC-SEC Preliminary Report to the Joint Advisory Committee on Emerging Regulatory Issues" as to what might have caused the trading experience of May 6, 2010:

> 1. *Disparate trading venues in the United States*; this is also known as "market fragmentation." It refers to the fact that multiple exchanges, alternative trading systems, and private matching networks (dark pools) run by broker-dealers all trade the same stocks in the United States simultaneously. While the overall liquidity may appear substantial, whenever there is a liquidity problem faced by one of the many trading venues containing a fraction of the total liquidity, the manner in which that venue reacts to the problem may initiate an overall chain reaction. Such a chain

reaction may not have happened at all if the total liquidity for each stock can be consolidated into a single trading venue.

2. "Liquidity Replenishment Points" (LRPs) at the NYSE and similar practices. Whenever an LRP is triggered, the NYSE will go into a "go slow" mode and pause momentarily to allow liquidity to enter the market. This may have exacerbated the problem, in that automated trading orders are most likely rerouted to other possible trading venues, resulting in a net loss of trading liquidity at the primary market. This may also have the effect of triggering similar cautionary procedures in parallel trading venues, driving liquidity further from the market.

3. "Self-Help remedy." Two exchanges declared "self help" against NYSE Arca in the minutes prior to 2:40 p.m., after NYSE Arca repeatedly failed to provide a response to incoming orders within one second. Such declarations free the declaring exchanges from their obligations to route unmatched orders to the affected exchange, resulting in additional loss of trading liquidity. For instance, a high bid and a low ask on the same stock appearing on two different exchanges, which could have been matched if there was rerouting, would fail to be matched under such circumstances.

4. Stop loss market orders. Some market participants left sell orders much lower than current prices as market orders to sell, primarily as a stop-loss precaution. Those orders were not expected to be executed. In a fast-falling market, these stop-loss market orders might have triggered a chain reaction of automated selling orders, and the sellers would have limited time to reconsider those orders. Typically, such orders would be left by institutional investors, and the quantity involved could be quite substantial as compared to the existing liquidity for a particular stock.

5. Short sales and stub quotes. Short sales against stub quotes accounted for more than 70% of the busted trades between 2:45 and 2:50 p.m. and approached a staggering 90% between 2:50 and 2:55 p.m. The fact that stub quotes were never intended to be executed, and that there would be limited (if any) upside to short selling against near-zero bids, suggests that at least some of these short sales were placed in a somewhat automated manner, since it would be unlikely for any experienced human trader to execute such orders.

In Lee, Cheng, and Koh (2010), the authors constructed a simulated market with multiple types of computer agents, including a market maker, systematic traders (deploying several varieties of trend-following strategies, which are among the most common techniques deployed by hedge funds), and "retail-like" investors who place randomized bids and asks in the market in a mean-reverting manner. Unlike traditional market simulations, the evolution of asset prices is the direct result of how these agents are trading against each other as in real markets, and there are no a priori

Figure 3. June 2010 E-mini Futures on S&P 500 vs. PG, MMM, and ACN.

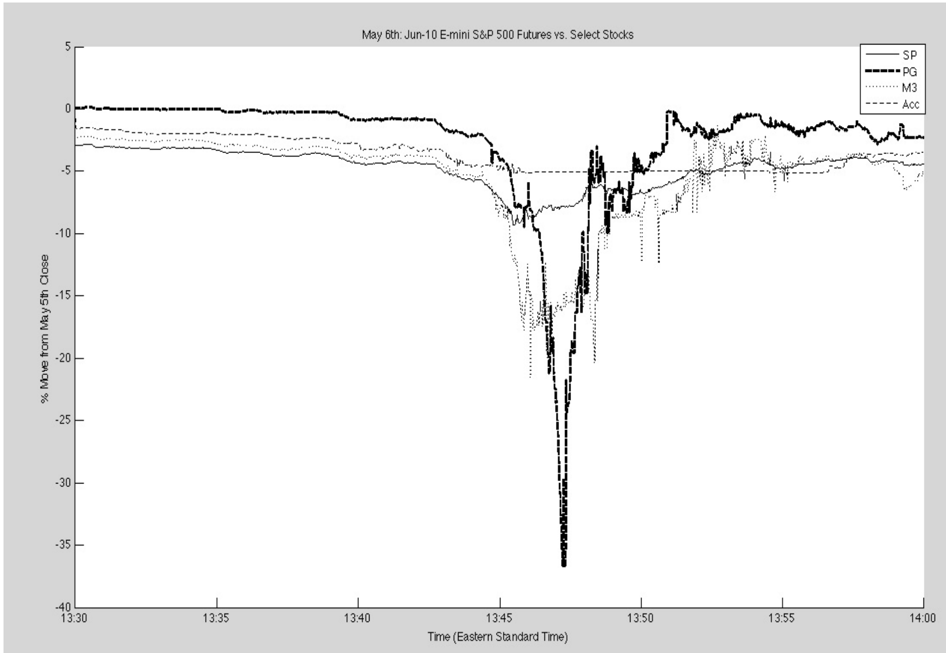

(Data courtesy of CME Group.)

assumptions on asset price distributions. While market simulation is hardly new, the academic contributions of our work are the following:

(i) We provide a convincing description of market dynamics based on the structure of the market and the type of participants.

(ii) The resulting price distribution provides a reasonable resemblance of the descriptive statistics of certain commodity markets.

(iii) Yet the simulation does not contain so many degrees of freedom that it essentially "over-fits" historical data, resulting in limited predictive power and insights.

The key findings from our earlier study include the following:

1. In theory, trend-following is a trading strategy that can be replicated by lookback straddles, which is a traditional "long gamma" strategy. The theoretical strategy is supposed to have unlimited upside but limited downside, much like any option. However, most option pricing theories work under the unrealistic assumptions of infinite liquidity and zero transaction costs. What we have observed is that, as we deliberately withdraw liquidity from the market, the profit-and-loss profiles of the trading strategies will deviate further and further away from the theoretical bounds derived based on option theories.

2. As the percentage of systematic traders in the market exceeds a certain threshold (between 60% and 80%) relative to the total number of market participants, the bids and offers in the market will concentrate on only one side of the market, especially during extreme market movements. Market prices will begin to behave erratically, leading to the eventual breakdown of the market.

3. Finally, any attempt to restore market liquidity by changing the "rules of the game" in the middle of trading is unlikely to produce the desired outcome. The process for market agents to adjust to any new set of rules, as well as subsequently reversing to the original state of the market, appears to cause more problems than it solves by creating significant liquidity disruptions to the market.

The goal of this paper is to determine if the findings from the earlier paper can be used to understand and assess potential regulatory responses, such as those listed in the "Joint CFTC-SEC Preliminary and Final Report." In particular, the authors contend that the events of May 6, 2010, show a pattern consistent with the type of "flash crash" observed in our earlier study. While some commentators assigned blame to high-frequency trading, our analysis was unable to identify a direct link to high-frequency trading per se. Rather, the likely causes are the domination of market activities by trading strategies that are responding to the same set of market variables in similar ways, as well as various pre-existing schemes that modify the "rules of the game" in the middle of trading, that results in a significant withdrawal of liquidity during extreme market movements. In addition, certain micro-structural safety mechanisms in the market, such as the uneven triggering of circuit breakers by the cash equity, futures, and ETF markets at different times, may have exacerbated the problem.

Furthermore, the triggering of the Liquidity Replenishment Points at the New York Stock Exchange (NYSE), commonly known as "go slow" mode, might have further driven liquidity out of the market when it was needed the most. Only when certain stocks reached "stupid cheap" levels, other investors seized the opportunity to buy and market prices began restoring to levels consistent with fundamental valuations. Moreover, the subsequent cancelling of trades by the NYSE has created a significant worry for market participants (market makers in particular) who can potentially step in to provide much-needed liquidity in similar episodes in the future.

To achieve our objectives, we have constructed nine different simulations in this study, in an attempt to recreate various market conditions for the cascading effects leading to the type of flash crash seen on May 6. Those results allow us to study the potential effects of:

- imposing position limits by traders.
- changing from continuous time auctions to discrete time auctions.
- imposing price limits during a major market dislocation, with different trigger levels.

Figure 4. A Sample Market Structure that Agents Need to Understand.

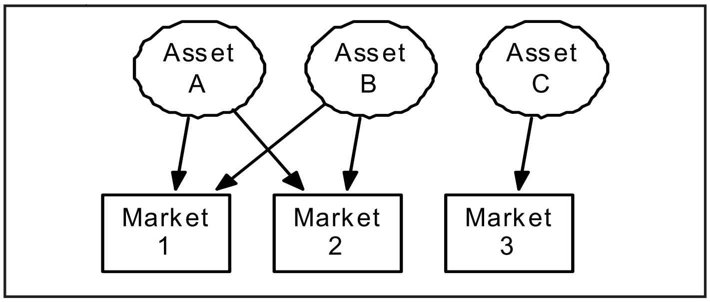

I. DESIGNING THE SIMULATION PLATFORM

It has been widely speculated that the Flash Crash on May 6, 2010 was caused primarily by two factors: (a) trading venues with different and often inconsistent rules of operations and (b) complex dependency among multiple assets (e.g., among index tracking ETFs and its component stocks). The first factor contributes to the congestion of orders when trading venues are slowing down unevenly, while the second factor contributes to the contagion of instability from one asset to other related assets. In order to reconstruct the market conditions leading to the Flash Crash and to evaluate policies that could help preventing similar incidents, we have developed a realistic microscopic financial simulation even though, to the best of our knowledge, no financial simulator can reproduce faithful replications of both features.

The simulation platform utilized in this paper is derived from the model first introduced in Cheng (2007), and used subsequently for analyzing extreme market conditions in Lee et al. (2010). In the following subsections, we will briefly describe the enhancements necessary for the simulation platform to model the two features mentioned above.

A. Multiple Trading Venues

With any sufficiently generic market engine, introducing multiple trading venues is relatively straightforward: The engine can simply create additional markets according to rules as specified by the user. However, the key challenge of having multiple trading venues is not about creating additional markets but avoiding operational bottlenecks. More specifically, we need to address how we can design a conceptual structure that is understandable by software agents and come up with a reasonable price discovery process under multiple trading venues.

For the software agents that we plan to introduce to the system, they need to recognize the relationship(s) among multiple markets. For example, for the case where a particular asset A is traded simultaneously in two markets, an agent needs to understand that buying and selling A in both markets will directly affect the position of A. In other words, agents in the simulation will need to load a conceptual mapping like the one illustrated in Figure 4. In our simulation design, we allow structural information to be defined compactly and all agents are required to load this same structural information at the starting-up phase. Once such mapping is loaded, an

Figure 5. Introducing Dependencies to Assets.

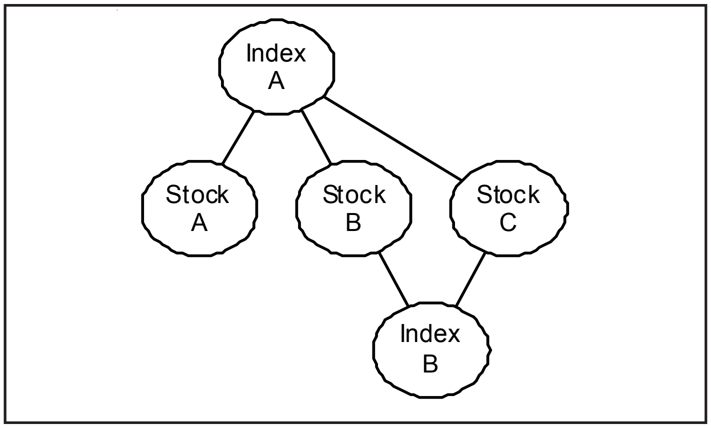

agent will then be able to keep an aggregated view on position balances through the linkages between markets.

Another important issue that needs to be addressed when introducing multiple trading venues is how prices of the same asset are synchronized across different markets. Take asset A in Figure 4 as an example: An agent intending to establish a long position in asset A needs to decide which market to trade in, since markets 1 and 2 are running independently and may have different prices. Agents certainly may have their own logic in deciding which market to go for; however, to simplify agent design and to emulate real-world trading rules, we assume that all bids and offers submitted by agents will go through a mechanism similar to the National Best Bid and Offer (NBBO) rule implemented in the U.S .stock market. In other words, when picking which market to trade in, an agent will simply pick the market with lowest ask prices (from all markets) when buying and the market with highest bid price when selling. Our assumption is that the updates on best ask/bid prices from all markets will be instantaneous without delay.

The framework presented above will allow us to design arbitrary market structures that suit our needs.

B. Complex Asset Dependency

Another important feature that we want to introduce is to allow assets to be related to each other. For example, the trading price of an index future should be dependent on the prices of all stock components this index future tracks. By allowing such dependencies, we are effectively linking up independent assets. An example of such dependency is illustrated in Figure 5.

Prices of linked assets cannot be directly synchronized, since prices of all assets still need to be determined by the market. Therefore, we need to go through a market mechanism to synchronize these asset prices. In order to achieve such synchronization, we introduce a special agent class called the "Arbitrageur." Arbitrageurs understand the relationship between assets, and they will trade whenever market prices are significantly out-of-sync.

Taking Index A in Figure 5 as an example: By assuming that Stocks A, B, and C are equally weighted in Index A, we can design the Arbitrageur using the following rules to eliminate any out-of-sync prices:

- If $Bid_{IndexA} \geq (1 + a)\{Ask_{StockA} + Ask_{StockB} + Ask_{StockC}\}$, then the arbitrageur should buy the basket of three stocks and sell the index.

- If $(1 + a)Ask_{IndexA} \leq \{Bid_{StockA} + Bid_{StockB} + Bid_{StockC}\}$ then the arbitrageur should buy the index and sell the basket of three stocks.

The parameter a is introduced to account for market frictions like delays or transaction costs. Arbitrageur will constantly review its holding, and whenever any of the following conditions is met, the Arbitrageur will liquidate its positions:

(1) If the price discrepancy disappears, that is, $Mid_{Index} \approx Mid_{StockA} + Mid_{StockB} + Mid_{StockC}$. The tolerance for being "sufficiently close" for liquidation can be adjusted empirically based on the bid-ask spreads shown in the tradable assets.

(2) If a perfect arbitrage is unsuccessful because of market slippage, we will implement a stop-loss rule to "reverse out" from any yet-to-be completed arbitrage trade based on a time trigger. This will happen when say only three out of the four legs of the arbitrage trade can be executed at the intended prices. This is an important feature to be included in any type of "flood to the gate" scenario, when one or more legs of an arbitrage trade is moving away from its intended price and the Arbitrageur has no choice but to unwind the trade.

(3) If, instead of convergence, an arbitrage trade diverges and creates losses instead of profits, the Arbitrageur will automatically "reverse out" from the arbitrage trade to prevent any run-away negative P&L. This is consistent with real-world practices and is another important feature to be included in any type of "flood to the gate" scenario. The trigger for stop loss is set to 5% initially and will be adjusted empirically based on the actual price behavior shown in the tradable assets.

The above rules for the Arbitrageur can be easily generalized to include an arbitrary number of assets and uneven weights.

II. SIMULATION DESIGN

A. Current Study

As mentioned earlier, we have conducted nine different simulations in this study, in an attempt to recreate various market conditions for the cascading effects leading to the type of flash crash seen on May 6. Those results allow us to study the potential effects of imposing position limits by traders, changing from continuous time auctions to discrete time auctions, and imposing price limits during a major market dislocation, with different trigger levels.

Specifically, there are the "deltas" from one simulation to the next in the current study:

> *Simulation 1 → Simulation 2:* Compressing the action-reaction time from the "go slow" mode in exchange 1 to the "go slow" mode in exchange 2, in order to pinpoint the potential triggering conditions leading to cascading effects. The purpose is to illustrate how market micro-structural issues can make a significant difference to market stability.
>
> *Simulation 2 → Simulation 3*: Imposing position limits by trader, instead of typical position limits by symbols (i.e., per stock trading on each individual exchange).
>
> *Simulation 3 → Simulation 4*: Changing the clearing mechanism from continuous time auction to discrete time auction, which would have negate any trade execution advantages of high-frequency, algorithm-based trading.
>
> *Simulation 3 → Simulation 5*: Simulation 5 is a variant of Simulation 3, in which quotes are not updated during the slowdown.
>
> *Simulation 3 → Simulation 6*: Simulation 6 is a variant of Simulation 3, in which price limits are imposed when prices have dropped by more than 40%, respectively, when compared to the base prices that are sampled from the last done prices every 60 seconds.
>
> *Simulation 6 → Simulation 7*: The trigger level above is set to 30% instead.
>
> *Simulation 7 → Simulation 8*: The trigger level above is set to 20% instead.
>
> *Simulation 8 → Simulation 9*: The trigger level above is set to 10% instead.

B. Technical Descriptions of Market Agents

For each stock, there are two markets in which it can be traded, with one market being roughly twice as large as another market (in terms of initially-available liquidity). Each stock is serviced by a Market Maker (MM) that is willing to provide liquidity by earning a small fee; the Index market, on the other hand, is not serviced by any MM. Besides the Market Maker, there are also Zero Intelligence (ZI) (or "random") agents, Trend Following (TF) agents, and Arbitrageur (AA) agents, with the latter having been described in detail in Section IB. Both ZI and TF agents are allowed to trade every stock available; however, only ZI agents are allowed to trade the Index. When trading in the Index market, ZI agents are designed to understand the linkage between index and its stock components. Whenever there are sufficiently large gaps between prices of index and component stocks, the AA agent will be performing arbitraging trades as described in Section IB and pulling the Index back to its fair value in the process. Non-convergence in the Index market is allowed and is one critical element of the market that we intend to model.

We have designated separate agents to emulate automatic stop losses and to generate the initial selling pressure in the Index market similar to the rush of sell orders at around 2:30 p.m. on May 6. A group of four agents (known as Bear Market agents) will automatically start piling in sell orders quickly once the major market slows down, to simulate the initial triggering of sell orders by traders who are likely to interpret the "go slow" mode as highly-negative market sentiments. To trigger automatic stop losses as and when the market suffers significant losses, a group of three agents will constantly monitor the stock prices. When asset price drops to below 60% of initial asset price, these agents (known as Stop-Loss agents) will begin placing large amounts of sell orders. For both groups of agents, the amount of sell orders each agent can issue is capped with a predetermined upper bound.

In all of our simulations, we fixed the agent composition at 18 ZI agents, 27 TF agents, and 9 AA agents, in order to represent a market in which there is significant presence of professional traders using algorithm-based techniques as well as those who are looking for arbitrage opportunities.

III. ANALYSIS OF SIMULATION RESULTS

This section contains a detailed analysis of our nine simulations.

A. Simulation Results

We have conducted nine different types of simulations based on a slowdown on market 1 followed by a slowdown in Market 2. In each case, we have plotted out the price history (for Stocks A, B, and C as well as the Index), the rolling exponentially-weighted volatility based on a λ value of 0.9 and the trading volume of each asset in 30-second buckets. The entire simulation lasted 900 seconds, which is comparable to the most active time period of the "flash crash" on May 6, 2010.

1. Simulation 1

The simulation shown in Figure 6 is based on a slowing down of Market 1 from 120 to 360 seconds and then a slowing down of Market 2 from 240 to 480 seconds. In the first case, we can see that prices collapsed, rolling volatilities spiked, and trading volumes picked up during the interval from 120 to 240 seconds and then during the interval from 400 to 600 seconds. This observation is consistent with our earlier research, in that the real problem appears to be caused by changing the "rules of the game" in the middle of trading, instead of the simple domination of the market by any specific type of traders. Since there are no changes to the fundamental demand-and-supply balance during the simulation (except for the initial triggering of selling orders by Bear Market agents), the market will function properly once it is stabilized, but the subsequent reversion to normal speed of clearing once again create an imbalance of demand and supply leading to significant price instabilities. In addition, we observe that, in some cases, price actually hit the value of $1, which is the value of stub quotes left by market-makers.

2. Simulation 2

The simulation shown in Figure 7 is based on a slowing-down of Market 1 from 120 to 240 seconds, and then Market 2 slowed down from 180 to 360 seconds. We are interested in understanding what may happen as and when we push the two slow-down periods closer together, emulating the cascading effects among unstable parallel markets. As expected, we no longer observe two distinct periods of shocks. Even more interesting are the observations that (a) the price-shock periods are compressed; as a result, there really isn't a sufficient time lag for supply and demand conditions in the market to recover from the first price shock before entering the second price shock; (b) prices go through an extended period of instability after the 360th second or the end of the second shock period; and (c) during the time when prices go through an extended period of instability, there continue to be many instances in which the Arbitrageur agents are unable to pull the Index back to its fair value. This is shown in Figure 15. Simulation 2 will be treated as our base scenario for testing other potential policy responses.

3. Simulation 3

The simulation shown in Figure 8 is based on imposing position limits by trader, instead of typical position limits by symbol (i.e., per stock trading on each individual exchange). Although not apparent from the descriptive statistics, the markets in this simulation experienced a significant increase in violent "up and down" shocks, and the price graph clearly shows signs of increased price instability. Readers should note that the type of extreme "up and down" shocks is actually consistent with the type of price movements shown on May 6. Those shocks are not observable with exchange data at the second-by-second level, but the authors have examined internal aggregated client data provided by a broker-dealer at the microsecond level showing exactly that type of extreme "up and down" shocks during the 2:30 to 3:30 p.m. EST period on May 6. The fact that these shocks actually become significantly more pronounced due to the imposition of position limits suggests that position limits are unlikely to have worked as an effective regulatory tool to eliminate "flash crash"-like symptoms.

4. Simulation 4

The simulation shown in Figure 9 is based on changing the clearing mechanism from continuous time auction to discrete time auction, which would have negated any trade execution advantages of high-frequency, algorithm-based trading. The modified clearing mechanism does not mean that the algorithm-based traders cannot execute trades; it only means that certain traders do not have any speed advantage relative to other market players, so they will profit only when they can come up with a fundamentally superior trading strategy that is not based on more timely execution. Based on both the price graphs and the descriptive statistics, it is not obvious that negating the advantages of high-frequency trading can make any significant difference in maintaining market stability.

5. Simulation 5

The simulation shown in Figure 10 is based on Simulation 3, in which quotes are not updated during the slowdown. This simulation is designed to address the following question: Instead of a total and abrupt stoppage — which is generally considered by the market as a blunt and ineffective tool since it simply delays the resolution to any fundamental imbalances in supply and demand — what would have been another alternative to a simple “go slow” mode? The typical “go slow” mode bears a certain degree of resemblance to discrete time auctions, in that primarily the amount of through-put in the clearing process is slowed down. Therefore, it is natural to ask whether stopping the publishing of quotes will make any difference. Based on both the price graphs and the descriptive statistics, it is not obvious that stopping the publishing of quotes could have made any significant difference in maintaining market stability.

6. Simulations 6, 7, 8, and 9

The simulations shown in Figures 11, 12, 13, and 14 are based on Simulation 3, in which price limits are imposed when prices have dropped by more than 40%, 30%, 20%, and 10%, respectively, when compared to the base prices that are sampled from the “last done” prices every 60 seconds. As a result of imposing this new policy, there are significant decreases in the skewness, kurtosis, and maximum drawdown statistics, with more significant improvements as and when the trigger level is lowered. Readers should note that imposing price limits does not address any fundamental supply and demand imbalances. Such imbalance should result in a natural drop in prices until a new market equilibrium is found, instead of any extreme “up and down” shocks, which rarely result in genuine price discovery and the orderly resolution of excessive demands/supplies. Moreover, there are more extreme “up and down” shocks when the price limit trigger is set either too low (40%) or too high (10%) — that may mean that regulators are either intervening too late (thus not providing any relieves) or needlessly (potentially making the situation worse). The ideal trigger level seems to be between 20% and 30%, which is consistent with the intuitive expectations of some market practitioners. Although we started these simulations by modifying Simulation 3, agent-level position limits are not breached in almost all cases, so that in practical terms Simulation 2 should be considered our true base scenario for these four simulations.

Figure 6. Price, Exponentially-Weighted Volatility and Trading Volume.

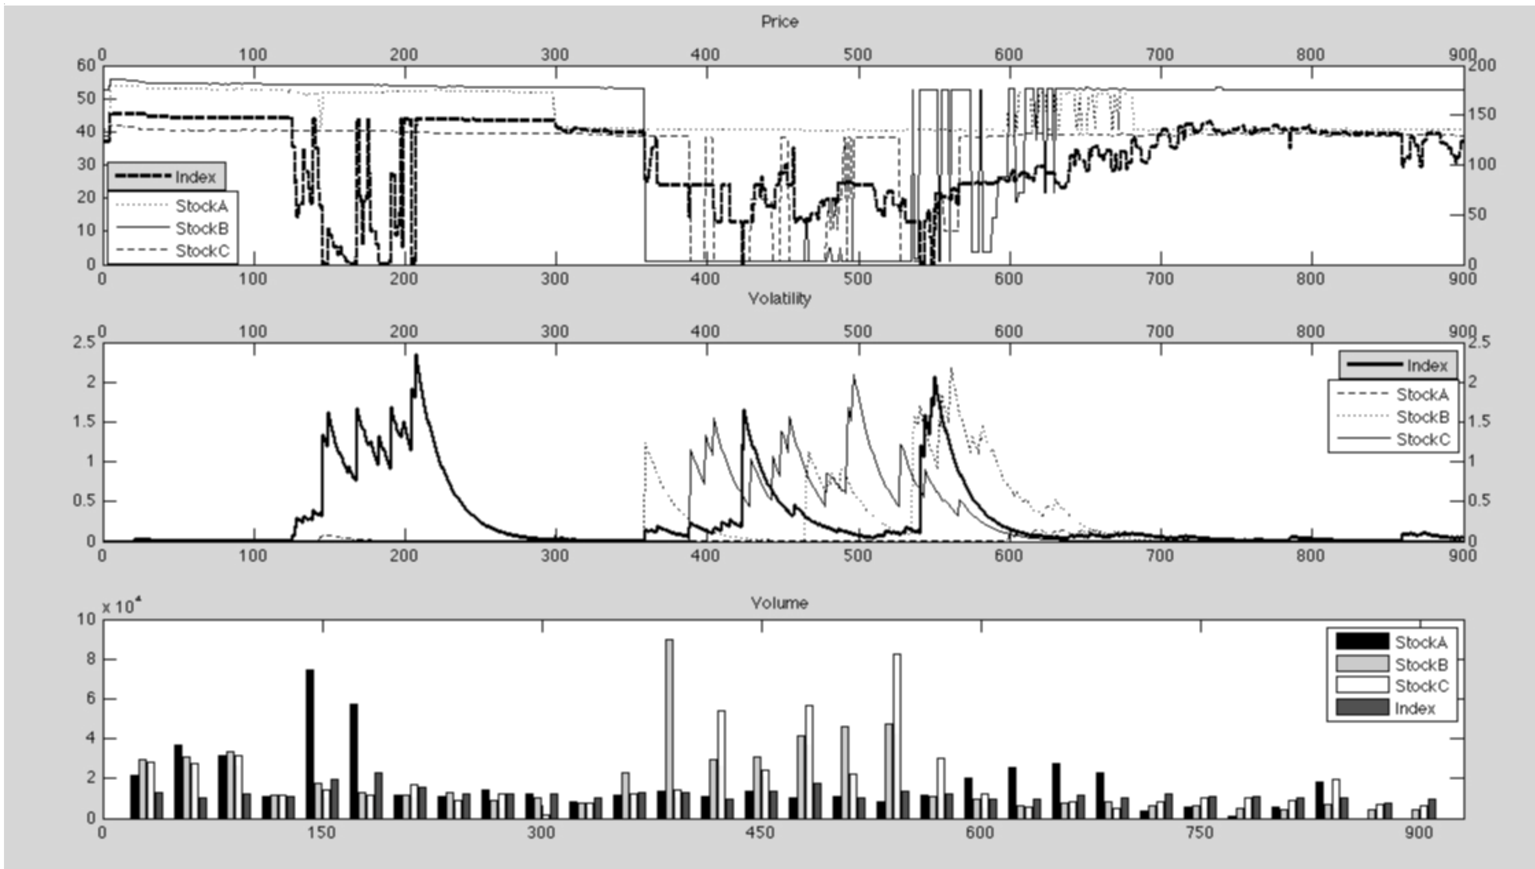

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 1 due to slow-down of Market 1 from 120 seconds to 360 seconds; Market 2 from 240 seconds to 480 seconds. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 7. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 2.

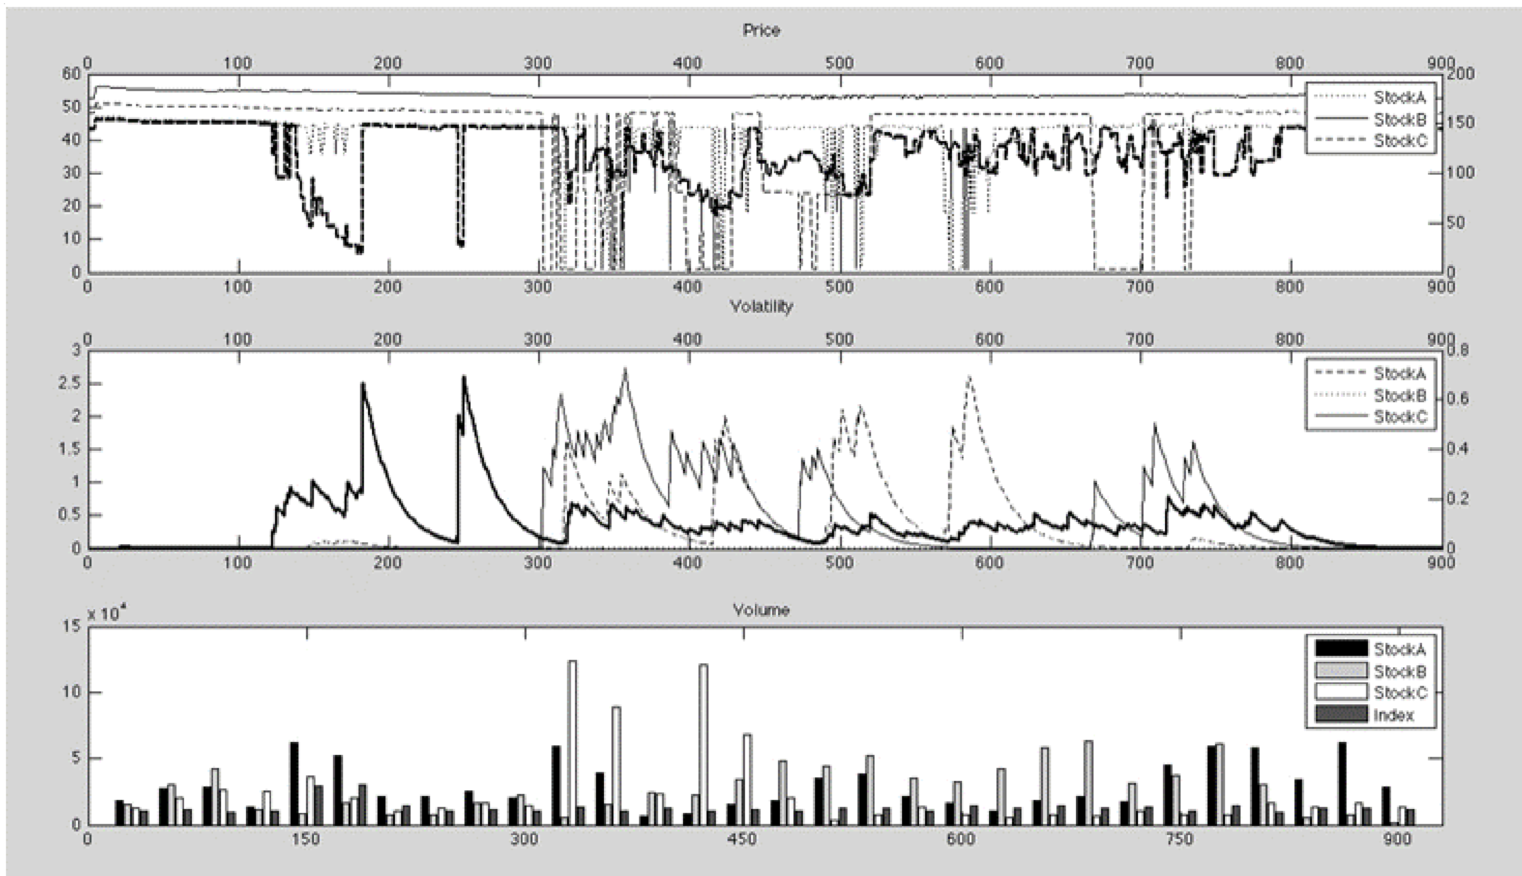

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 2 due to slow-down of Market 1 from 120 seconds to 240 seconds; Market 2 from 180 seconds to 360 seconds. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 8. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 3.

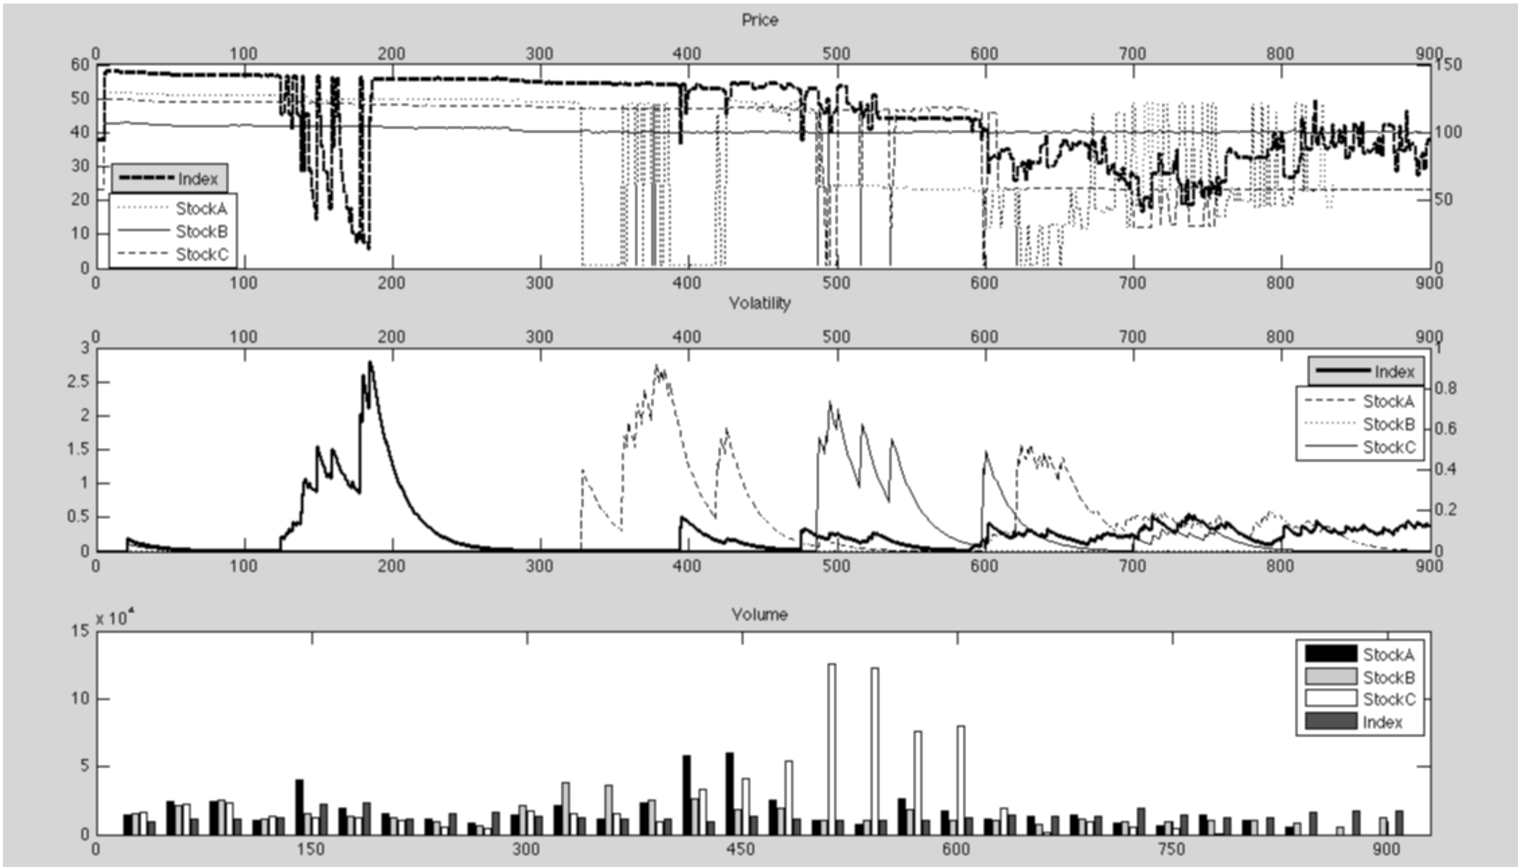

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 3 due to imposing position limits by trader, instead of typical position limits by symbols (i.e., per stock trading on each exchange). Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 9. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 4.

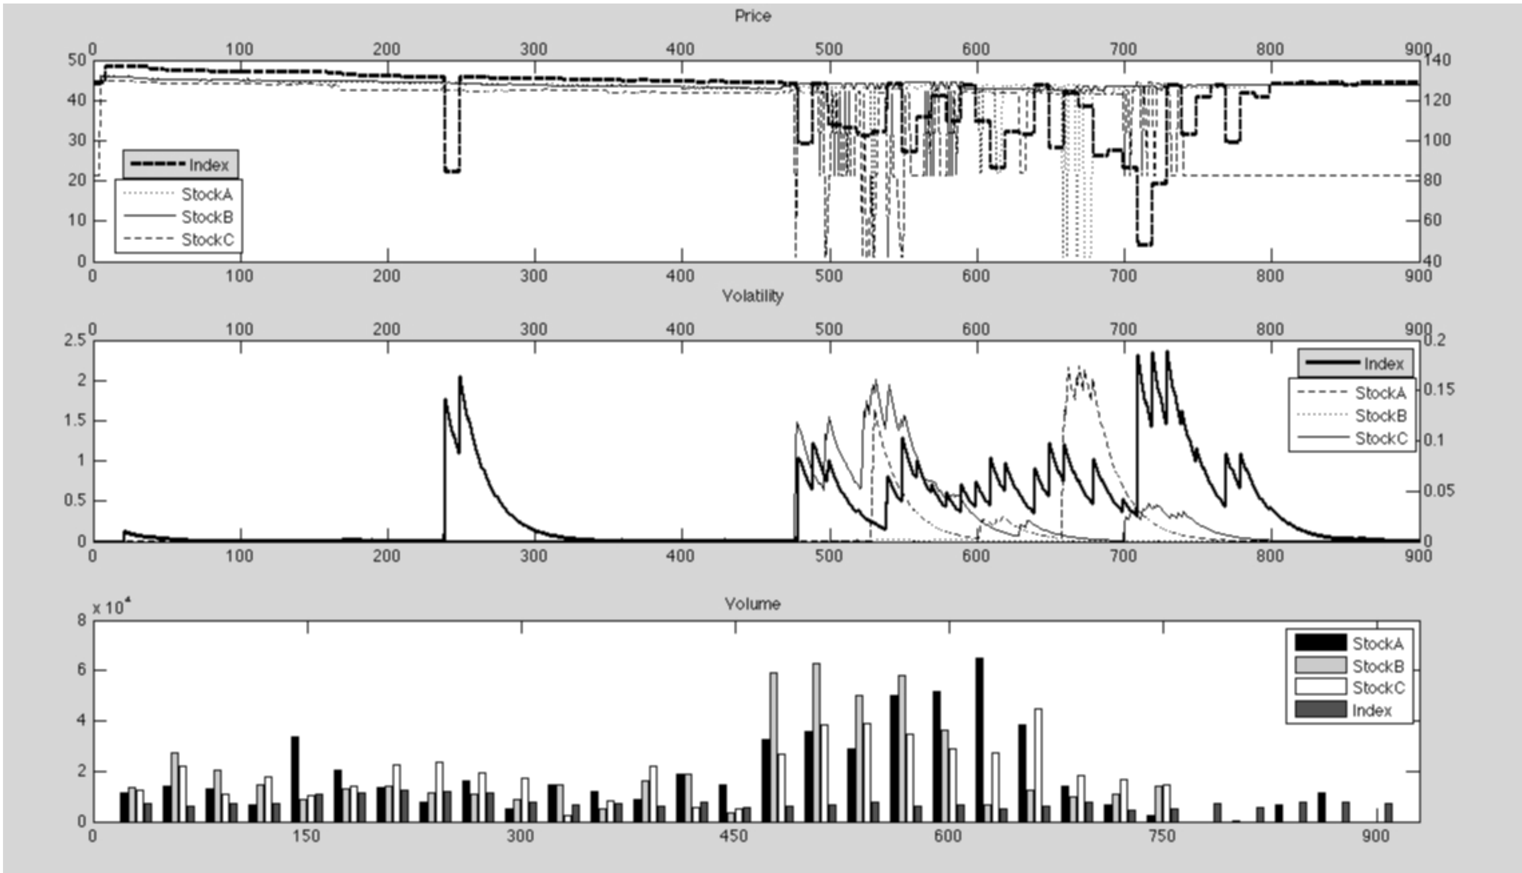

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 4 after changing the clearing mechanism from continuous time auction to discrete time auction, which would have negate any advantages of high-frequency trading. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 10. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 5.

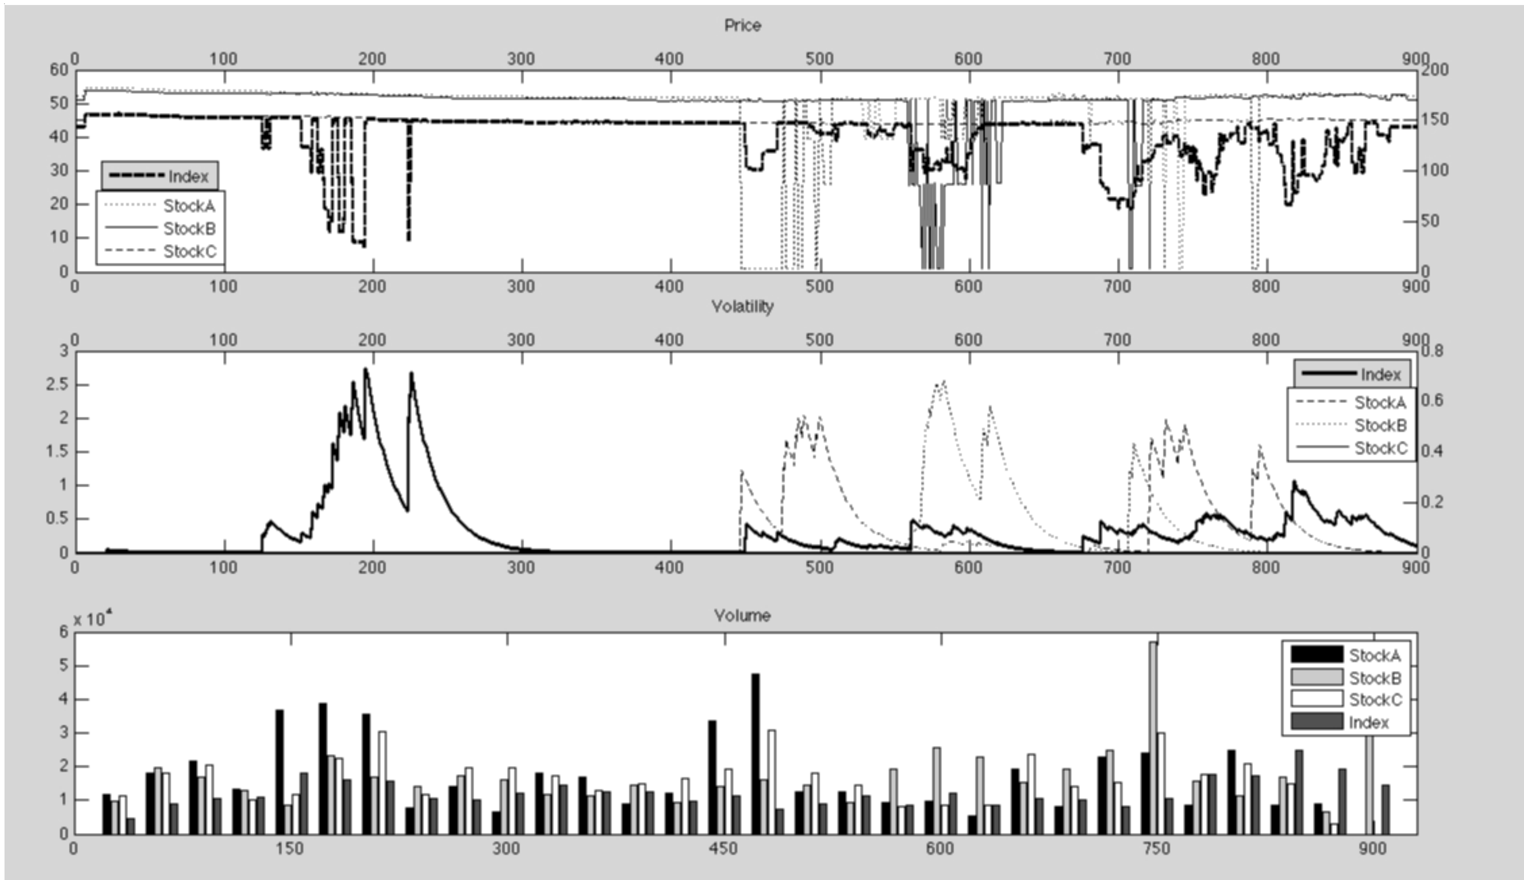

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 5, which is a variant of Simulation 3, but its quotes are not updated during the slowdown. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 11. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 6.

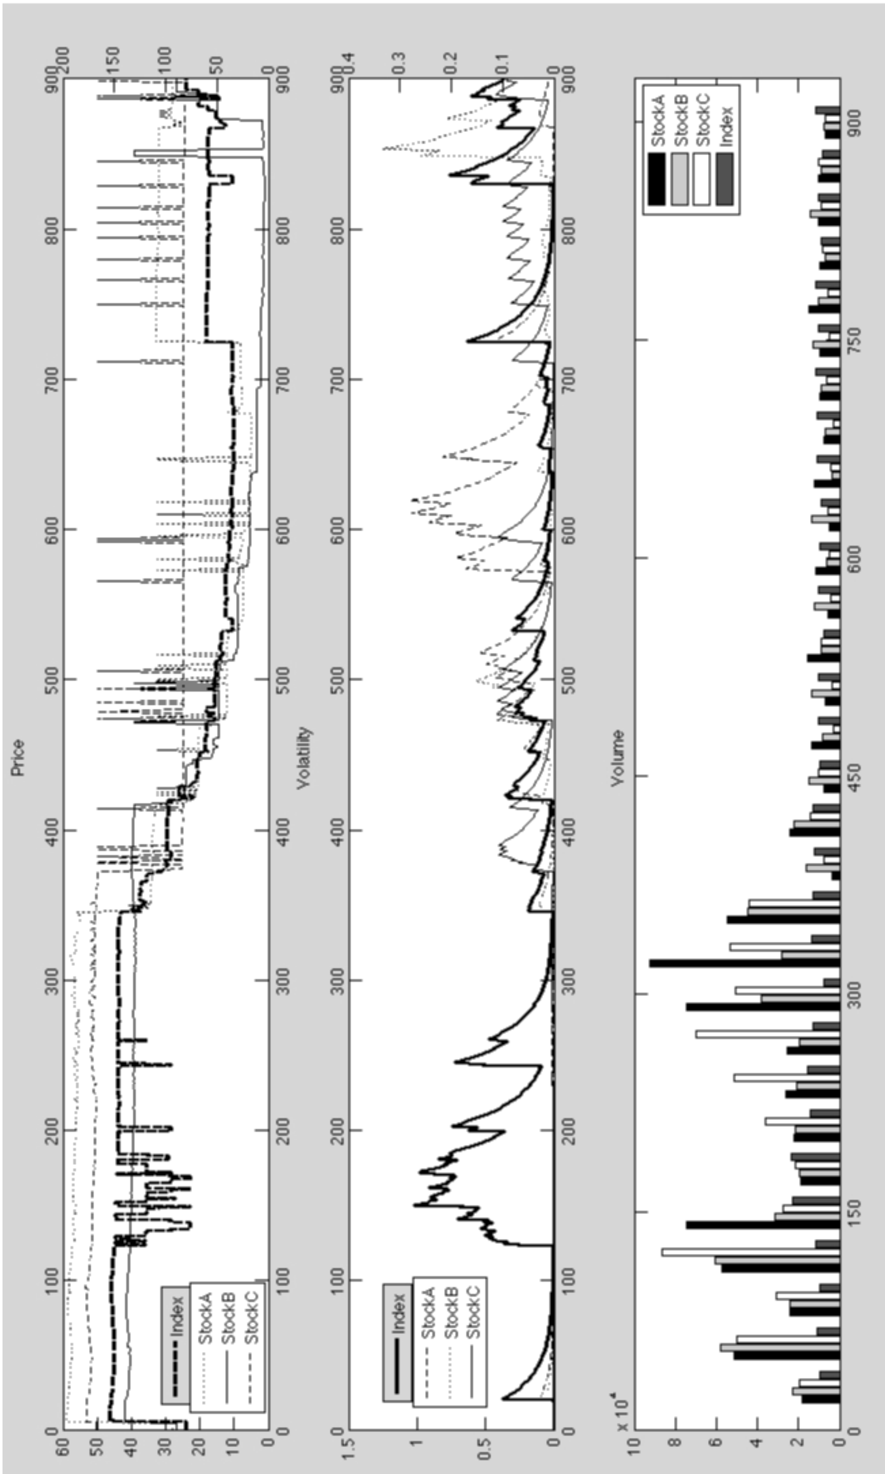

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 6, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 40% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 12. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 7.

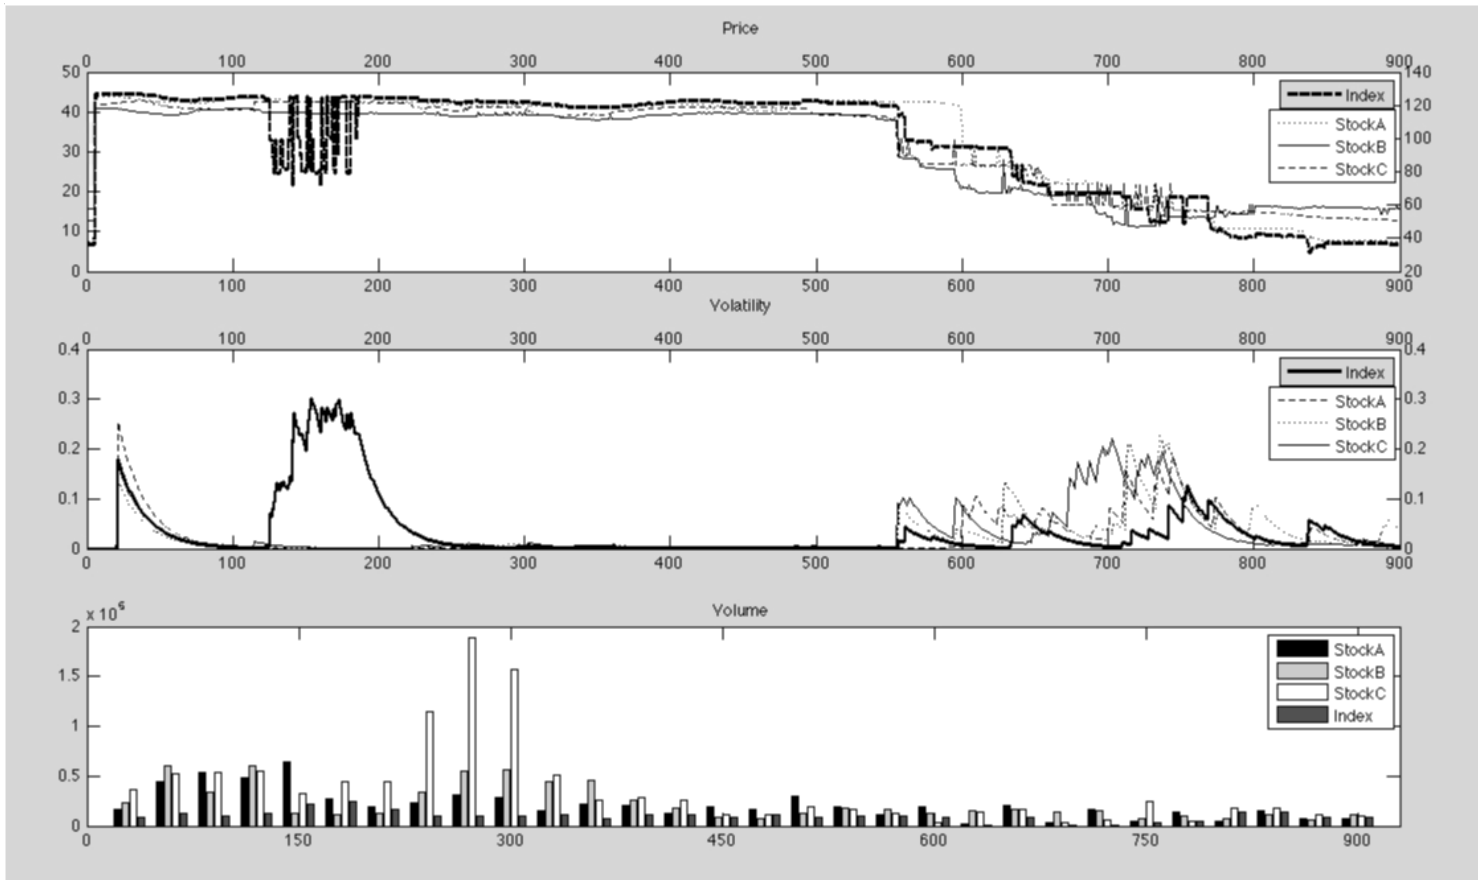

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 7, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 30% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 13. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 8.

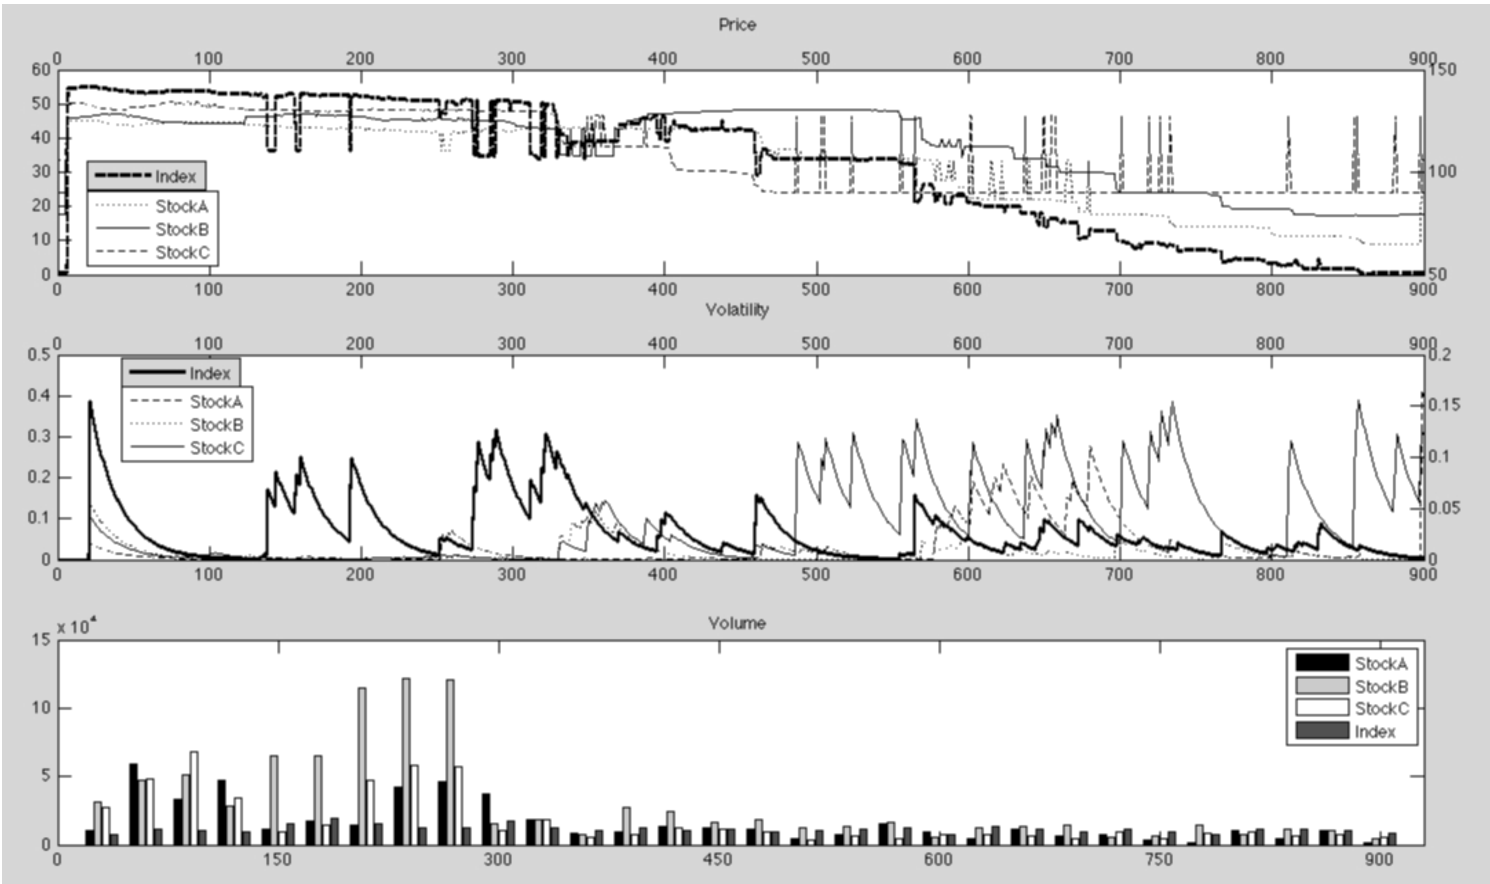

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 8, which is a variant of Simulation 3, but price limits are imposed whenever prices have dropped by more than 20% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 14. Price, Exponentially-Weighted Volatility and Trading Volume in Simulation 9.

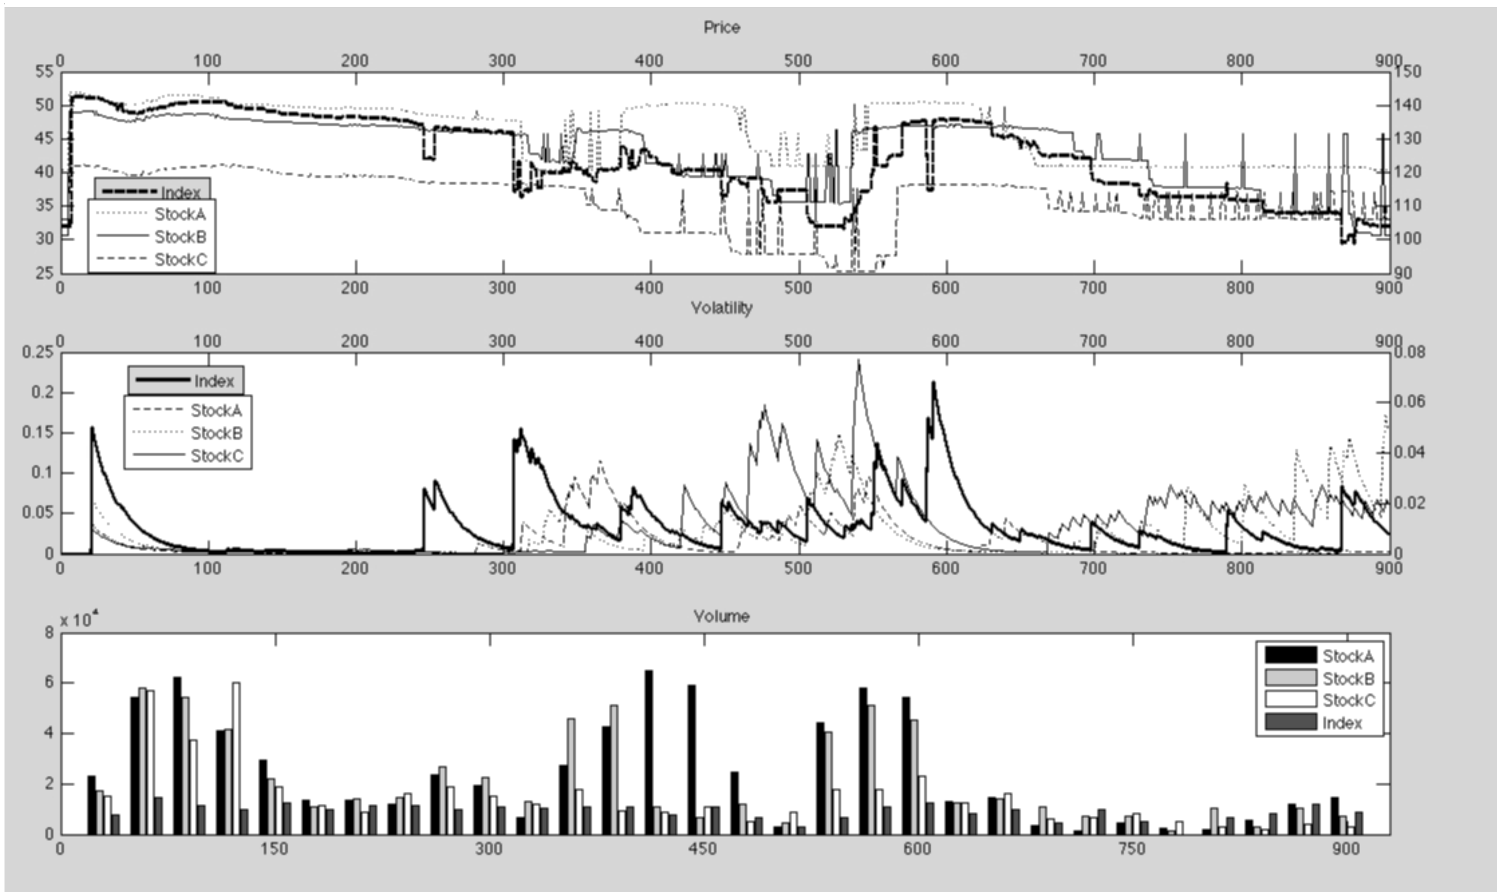

Price, exponentially-weighted volatility and trading volume (in 30-second buckets) in Simulation 9, which is a variant of Simulation 3, but price limits are imposed whenever prices prices have dropped by more than 10% when compared to the average of the last 5 trades. Left axis is for stocks and right axis is for Index. Time axis is in seconds.

Figure 15. Comparison of Synthetic Fair Value vs. Traded Index Values in Simulation 2.

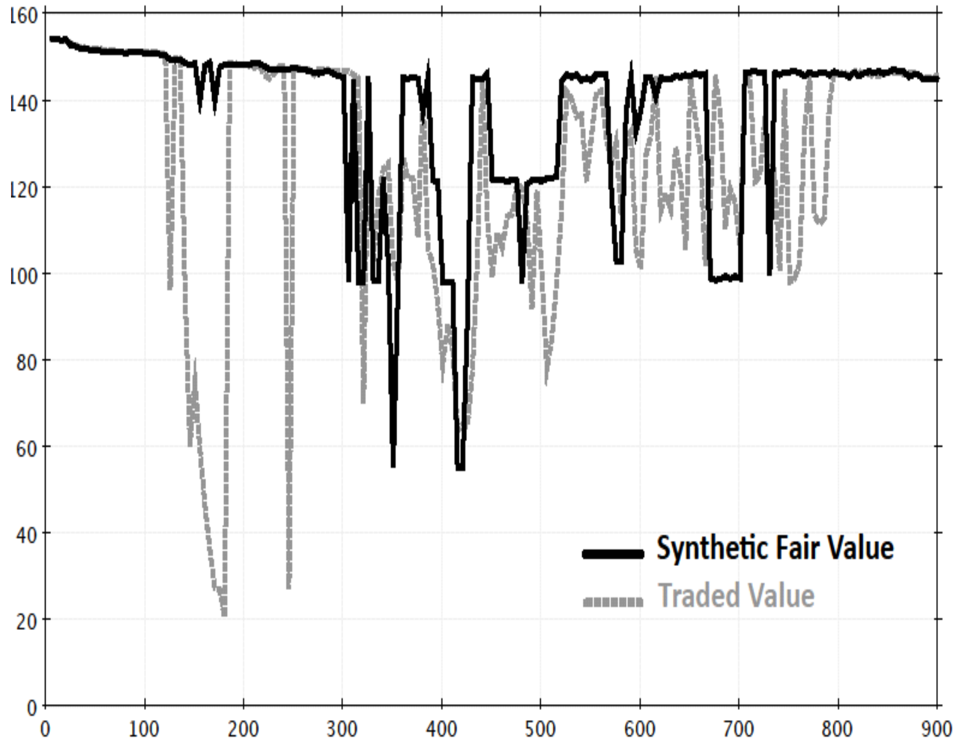

B. Statistical Analysis

The summary statistics below (Tables 1–4) are computed based on second-by-second data using absolute differences in returns on the Index. Because our simulated Index is composed of only 3 stocks instead of 500 securities in the SPX, the difference in base index values means that computing the proportional differences may produce non-comparable (if not non-sensical) results and in particular unreliable skewness statistics. Skewness and kurtosis are scale invariant, and the simulated skewness and kurtosis appear to be "close enough" when compared to those observed from the SPX E-mini futures market on May 6, 2010. Moreover, the minimum and maximum values of the simulations are roughly about 10 times the size of their corresponding standard deviations. That is not reasonable as compared to real-market returns on May 6, 2010 especially those of single-name stocks. (Refer to our earlier study for a further discussion on the challenges and goals in getting "close enough" when matching moments in simulating extreme market movements.)

The comparison is particularly striking when the outputs of these simulations are lined up side by side against typical fat-tail distributions created by a priori mathematical assumptions. Our assessment is that these simulations have produced price distributions with "reasonable resemblance" of the actual evolution of the prices on SPX E-mini futures from 2:30 to 5:00 p.m. EST on May 6, 2010; changing the observation window within the 30-minute time frame does not result in any dramatic changes to the descriptive statistics on the prices of the SPX E-mini futures.

Figure 16. Comparative Return Distributions.

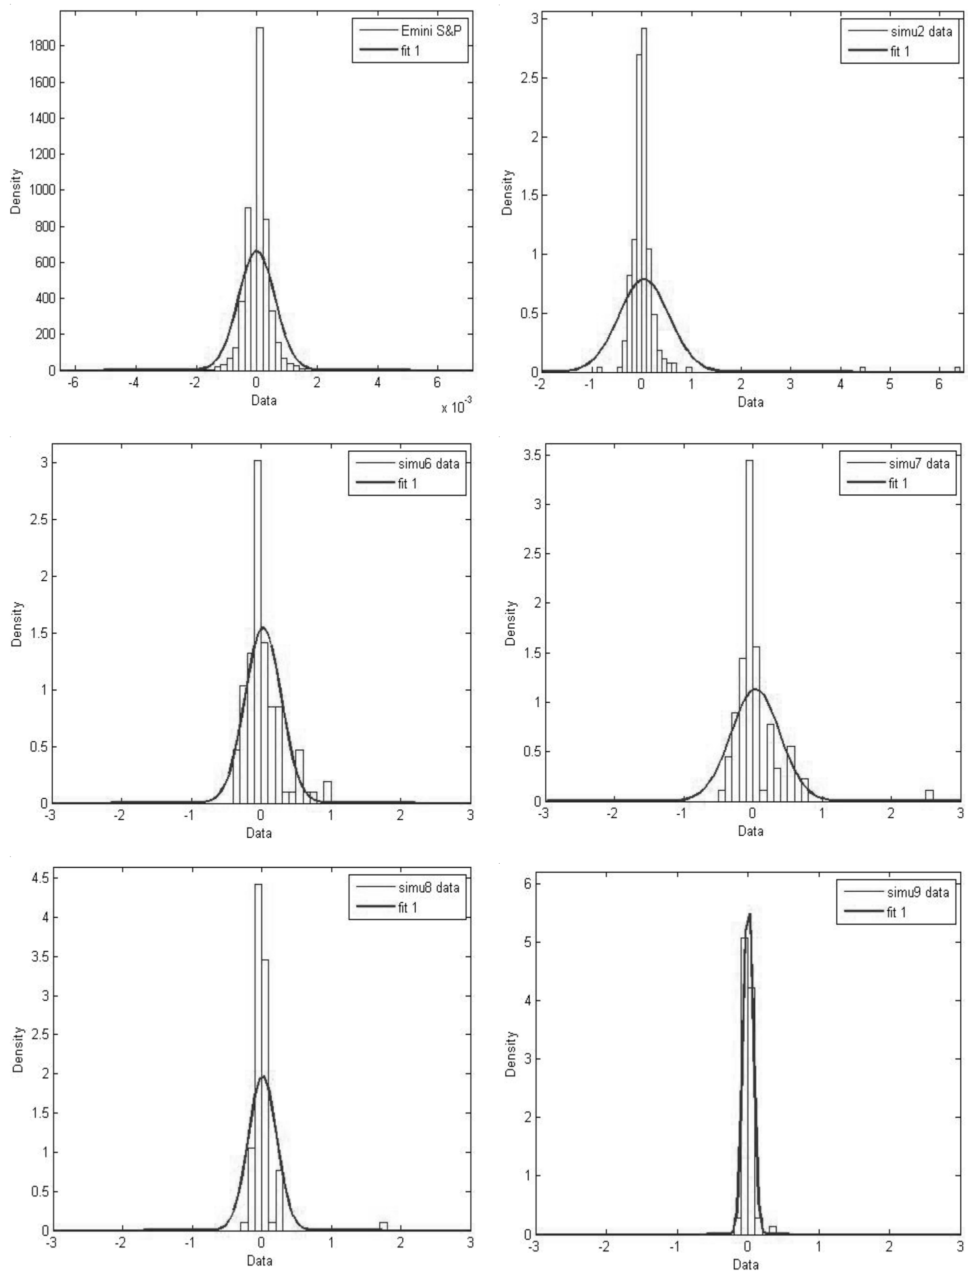

Comparative return distributions based on the SPX E-mini futures as well as the Index from Simulations 2, 6, 7, 8 and 9.

Table 1. Descriptive Statistics on Stock A as well as the SPX E-mini Futures on May 6, 2010.

Stock	A									Emini S&P*	P&G	3M	Accen-ture
Simulation	1	2	3	4	5	6	7	8	9				
Observations	776	745	744	735	705	810	830	785	815	1794	1680	1346	656
Mean	45.64	42.28	35.04	43.34	48.16	35.43	34.06	32.92	46.30	1108.80	60.37	81.74	40.83
Stdev	5.72	8.37	17.95	5.03	12.90	19.86	12.68	12.68	4.40	17.39	3.20	3.46	0.45
Skewness	0.32	-4.13	-0.66	-7.54	-3.14	-0.19	-1.05	-0.71	-0.33	-1.03	-3.58	-1.82	-0.24
Kurtosis	1.13	2.04	4.72	2.81	11.52	1.55	2.38	1.83	1.29	3.25	17.98	5.32	1.83
Min	40.40	1.00	1.00	1.00	1.00	4.90	7.40	9.10	40.30	1056.00	39.37	67.98	40.01
Max	53.90	47.00	52.10	46.10	54.90	58.90	44.00	45.10	52.10	1130.80	62.25	85.49	41.53
Max - Min	13.50	46.00	51.10	45.10	53.90	54.00	36.60	36.00	11.80	74.75	22.88	17.51	1.52
CVaR(95%)	-3.88	-20.09	-28.29	-8.15	-16.43	-10.05	-1.91	-3.65	-2.59	-1.54	-0.69	-1.63	-0.10
MaxDD	0.22	0.98	0.98	0.98	0.98	0.84	0.30	0.48	0.18	0.01	0.17	0.09	0.01
#(DD(>=10%))	13	38	43	7	21	29	13	18	10	0	3	0	0

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.

Table 2. Descriptive Statistics on Stock B as well as the SPX E-mini Futures on May 6, 2010.

Share	B									Emini S&P*	P&G	3M	Accen-ture
Simulation	1	2	3	4	5	6	7	8	9				
Observations	820	804	762	787	729	798	847	807	816	1794	1680	1346	656
Mean	39.15	53.74	40.77	43.97	49.66	23.53	31.05	38.51	43.60	1108.80	60.37	81.74	40.83
Stdev	23.01	0.72	0.85	0.84	8.53	16.98	11.04	10.89	4.84	17.39	3.20	3.46	0.45
Skewness	-1.00	1.03	0.89	0.18	-4.27	-0.18	-0.62	-0.93	-0.96	-1.03	-3.58	-1.82	-0.24
Kurtosis	19.52	3.47	4.34	4.32	21.67	1.19	1.53	2.29	2.66	3.25	17.98	5.32	1.83
Min	1.00	52.70	39.90	42.70	1.00	1.00	11.20	17.30	30.70	1056.00	39.37	67.98	40.01
Max	55.90	56.10	43.00	46.10	54.00	42.00	41.00	48.40	49.20	1130.80	62.25	85.49	41.53
Max - Min	54.90	3.40	3.10	3.40	53.00	41.00	29.80	31.10	18.50	74.75	22.88	17.51	1.52
CVaR(95%)	-11.94	-0.59	-0.33	-1.00	-14.70	-4.79	-1.79	-1.94	-4.25	-1.54	-0.69	-1.63	-0.10
MaxDD	0.98	0.02	0.02	0.04	0.98	0.96	0.39	0.19	0.33	0.01	0.17	0.09	0.01
#(DD(>=10%))	13	0	0	0	17	19	11	8	11	0	3	0	0

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.

Table 3. Descriptive Statistics on Stock C as well as the SPX E-mini Futures on May 6, 2010.

Share	C									Emini S&P*	P&G	3M	Accen-ture
Simulation	1	2	3	4	5	6	7	8	9				
Observations	749	745	702	804	664	858	839	830	812	1794	1680	1346	656
Mean	33.65	40.70	41.79	39.68	44.85	36.66	32.72	35.00	35.96	1108.80	60.37	81.74	40.83
Stdev	13.11	16.24	12.03	8.56	0.89	13.17	11.31	11.35	4.29	17.39	3.20	3.46	0.45
Skewness	-1.87	-1.72	-1.69	-2.70	0.81	0.23	-0.74	0.23	-0.88	-1.03	-3.58	-1.82	-0.24
Kurtosis	1.93	2.33	4.58	2.49	2.70	1.06	1.78	1.19	2.76	3.25	17.98	5.32	1.83
Min	1.00	1.00	1.00	1.00	43.80	24.50	12.90	23.90	25.20	1056.00	39.37	67.98	40.01
Max	41.90	51.10	50.00	45.00	47.30	53.30	42.90	50.50	41.10	1130.80	62.25	85.49	41.53
Max – Min	40.90	50.10	49.00	44.00	3.50	28.80	30.00	26.60	15.90	74.75	22.88	17.51	1.52
CVaR(95%)	-7.82	-20.49	-10.48	-22.44	-0.51	-14.68	-2.79	-11.69	-5.19	-1.54	-0.69	-1.63	-0.10
MaxDD	0.97	0.98	0.98	0.98	0.03	0.51	0.27	0.49	0.33	0.01	0.17	0.09	0.01
#(DD(>=10%))	11	23	13	30	0	24	18	25	32	0	3	0	0

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.

Table 4. Descriptive Statistics on Index as well as the SPX E-mini Futures on May 6, 2010.

Index										Emini S&P*	P&G	3M	Accen-ture
Simulation	1	2	3	4	5	6	7	8	9				
Observations	684	622	637	711	714	749	746	744	752	1794	1680	1346	656
Mean	106.55	123.43	110.85	121.80	133.93	88.62	96.31	101.97	124.41	1108.80	60.37	81.74	40.83
Stdev	40.31	29.40	30.59	16.69	25.91	45.59	33.12	31.76	11.77	17.39	3.20	3.46	0.45
Skewness	-0.86	-1.29	-0.70	-1.94	-1.90	0.28	-0.73	-0.31	-0.15	-1.03	-3.58	-1.82	-0.24
Kurtosis	61.30	2.03	9.87	7.05	6.47	1.40	1.90	1.60	1.74	3.25	17.98	5.32	1.83
Min	0.80	19.48	14.83	48.60	23.39	33.10	31.50	50.50	98.90	1056.00	39.37	67.98	40.01
Max	152.10	154.10	145.10	137.30	156.00	154.10	127.20	141.50	142.50	1130.80	62.25	85.49	41.53
Max - Min	151.30	134.62	130.27	88.70	132.61	121.00	95.70	91.00	43.60	74.75	22.88	17.51	1.52
CVaR(95%)	-41.74	-32.01	-28.45	-10.24	-35.88	-20.99	-17.25	-11.10	-4.49	-1.54	-0.69	-1.63	-0.10
MaxDD	0.99	0.82	0.86	0.44	0.80	0.40	0.43	0.20	0.15	0.01	0.17	0.09	0.01
#(DD(>=10%))	52	60	58	12	43	28	23	12	2	0	3	0	0

*Average of available bid and ask based on second-by-second data from 14:30:00 to 14:59:59 EST on May 6, 2010.

Figure 17. Realized P&L in Simulation 1 for Different Agent Types.

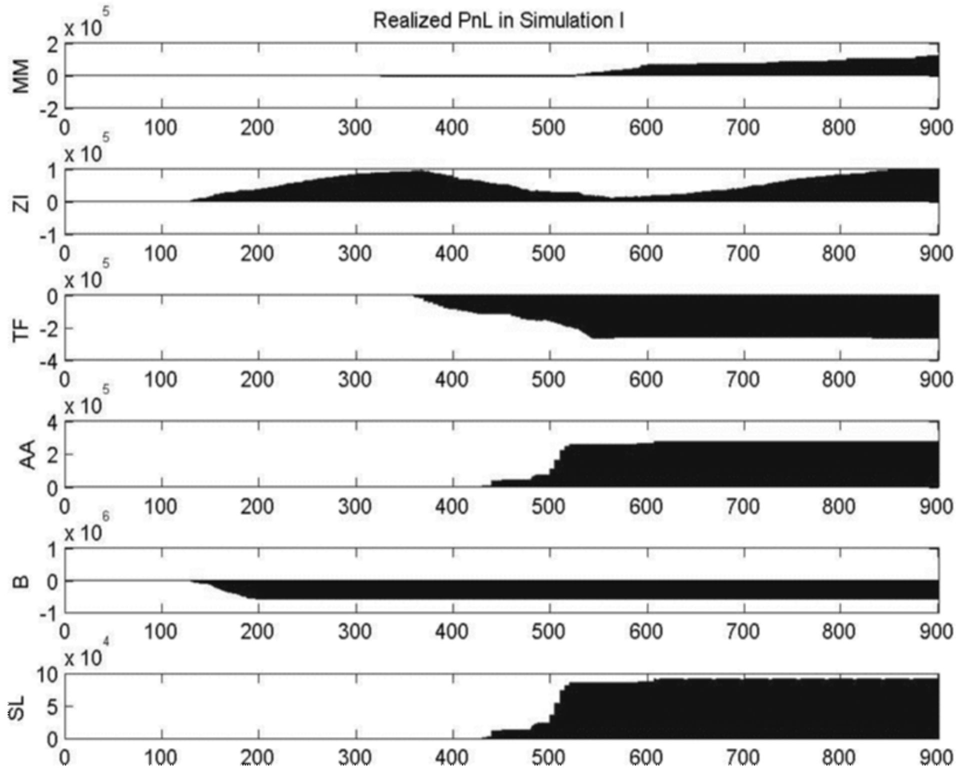

Realized P&L in Simulation 1 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

Figure 16 plots out the comparative return distributions based on the SPX E-mini futures as well as the Index from Simulations 2, 6, 7, 8 and 9.[2] Readers should visually examine the degree of similarity between the return distribution in our base scenario of Simulation 2 and that from the SPX E-mini futures. Not surprisingly, their skewness (-1.29 for Simulation 2 vs. -1.03 for SPX E-mini) and kurtosis (2.03 for Simulation 2 vs. 3.25 for SPX E-mini) statistics are also quite close. This graph also shows how the base scenario evolves under the price limit triggers set at 40%, 30%, 20%, and 10%, with tighter and tighter fits against their corresponding normal distribution curves.

C. Agents P&Ls

We have plotted the realized and unrealized P&Ls for all agent types in Simulations 1 and 2 in Figures 17, 18, 19, and 20. From these base scenarios we make the following observations:

2. To ensure an objective comparison, "zeros" have been deleted from the return distributions, as discussed in Lee et al. 2010.

Figure 18. Unrealized P&L in Simulation 1 for Different Agent Types.

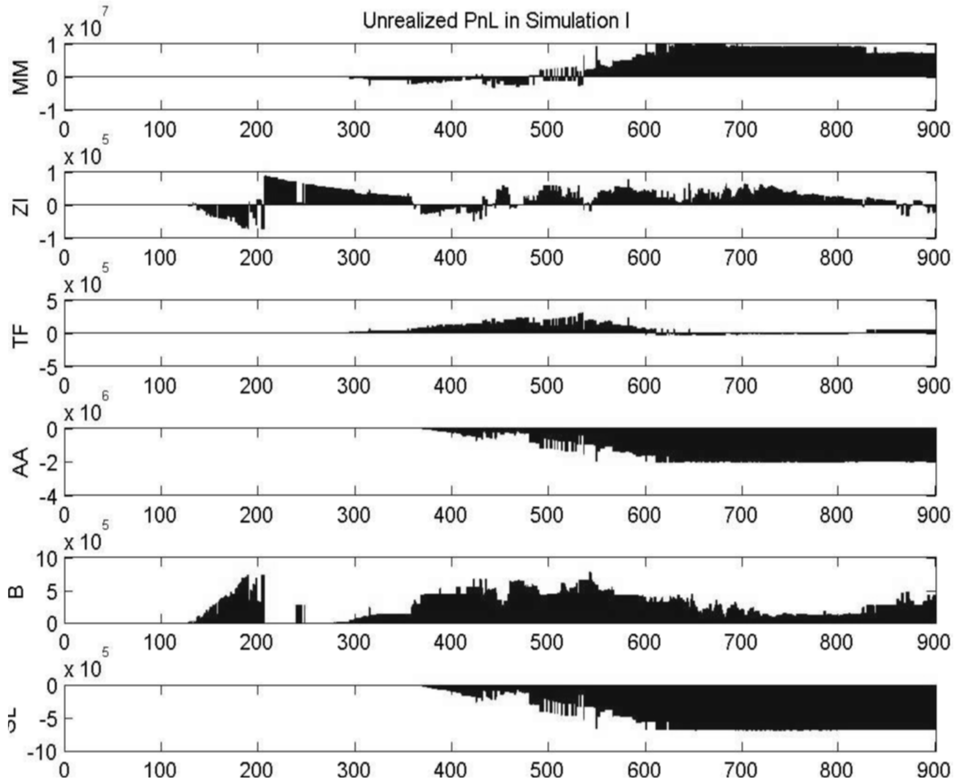

Unrealized P&L in Simulation 1 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

1. In the absence of market interventions, Market Makers almost always make profits by design of their trading algorithms.

2. Neither the ZI (or "random") agents nor the trend follower TF agents are able to make consistent profits.

3. As expected, Arbitrageurs may suffer heavy losses when the Index fails to converge to its fair values.

4. The Bear Market seller may or may not make any profits, depending on the market's recovery path.

5. The Stop-Loss agents will almost always lose money in flash crash by selling at unusually low prices that consequently recover.

If trades are "busted" at a certain level, then the P&Ls of the Market Makers will become uncertain. Doing so is expected to have a highly negative impact on the Market Makers' willingness to participate in the markets during flash crashes.

Figure 19. Realized P&L in Simulation 2 for Different Agent Types.

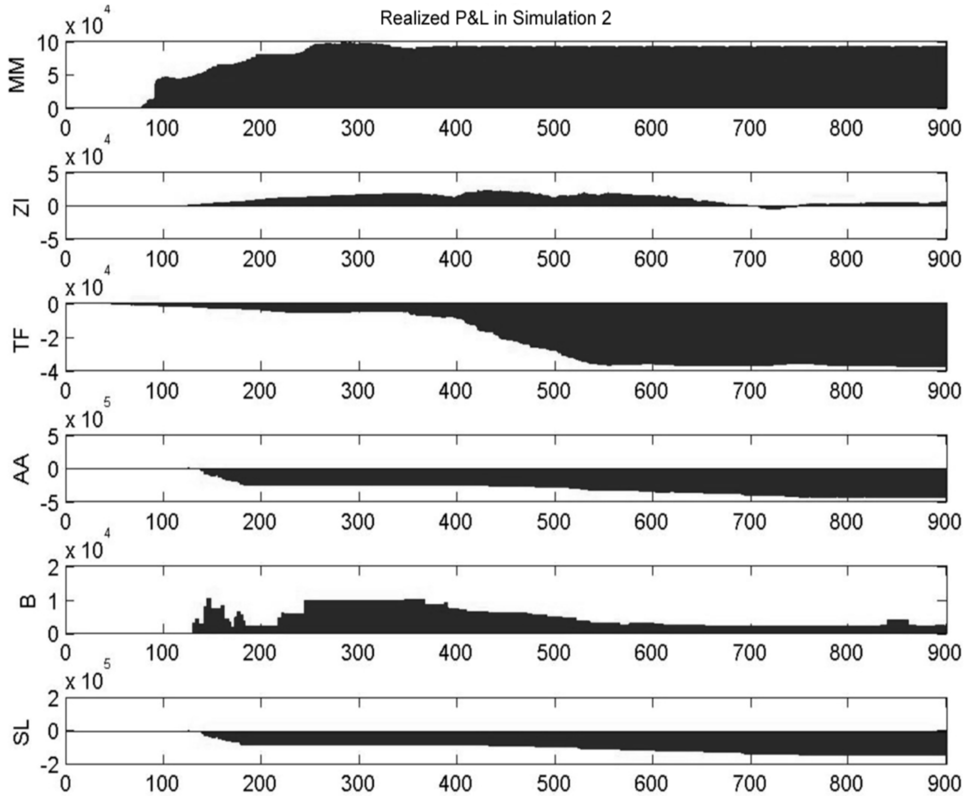

Realized P&L in Simulation 2 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

Without their participation in such markets, the authors contend that (a) it will be even more likely for the market to break down faster when liquidity is withdrawn faster from the market and (b) it will be more difficult for the market to recover from the destabilizing effects of any "flash crash."

In addition, the unrealized P&Ls for all agent types in Simulations 3, 4, 7, and 8 (Figures 21, 22, 23, and 24) show that:

- Both imposing position limits by trader and changing the clearing mechanism from continuous time auction to discrete time auction may be ineffective in terms of eliminating "flash crash"-like symptoms, but these measures do not cause any unexpected changes to the P&L patterns among different types of market players.

- In Simulations 7 and 8 where price limits are imposed, it appears that certain professional traders are able to make profits at the expense of the Market Maker and to some extend the ZI (or "random") agents.

Figure 20. Unrealized P&L in Simulation 2 for Different Agent Types.

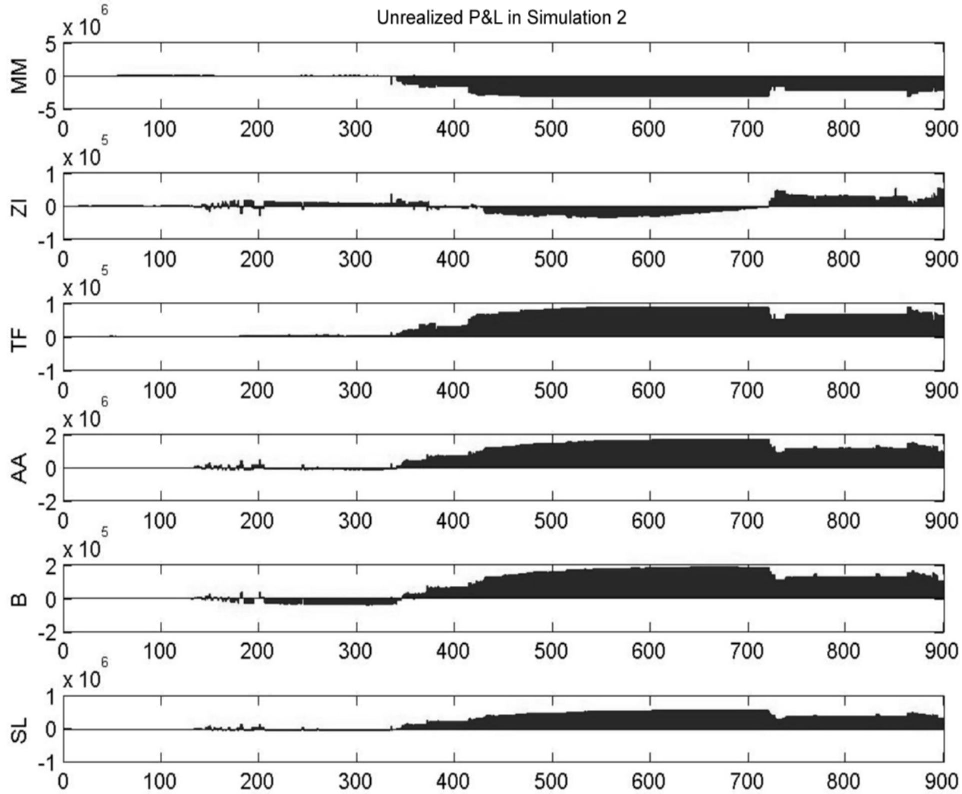

Unrealized P&L in Simulation 2 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

Observation 2 is troubling, but not hugely surprising. When the market knows which direction a particular asset is going to trade because of regulatory intervention, professional traders can usually find ways to take advantage of the anticipated market movements. Market participants who are likely to be on the losing side of their trades will be the retail-like zero intelligence investors who typically deploy unsophisticated trading strategies assuming a fairly even distribution of market ups and downs, or market makers who are obligated to quote under the assumption that bids and asks should be reasonably even and random. From a regulatory viewpoint, imposing price limits can be an effective policy to eliminate "flash crash"-like symptoms, but nonetheless one that may create unintended fairness issues for certain market participants.

1. "Busting" Trade

Finally, we used the base scenario of Simulation 2 to test the potential P&L impacts due to "busting trades" at or below 60% of the opening price of the asset traded:

Figure 21. Unrealized P&L in Simulation 3 for Different Agent Types.

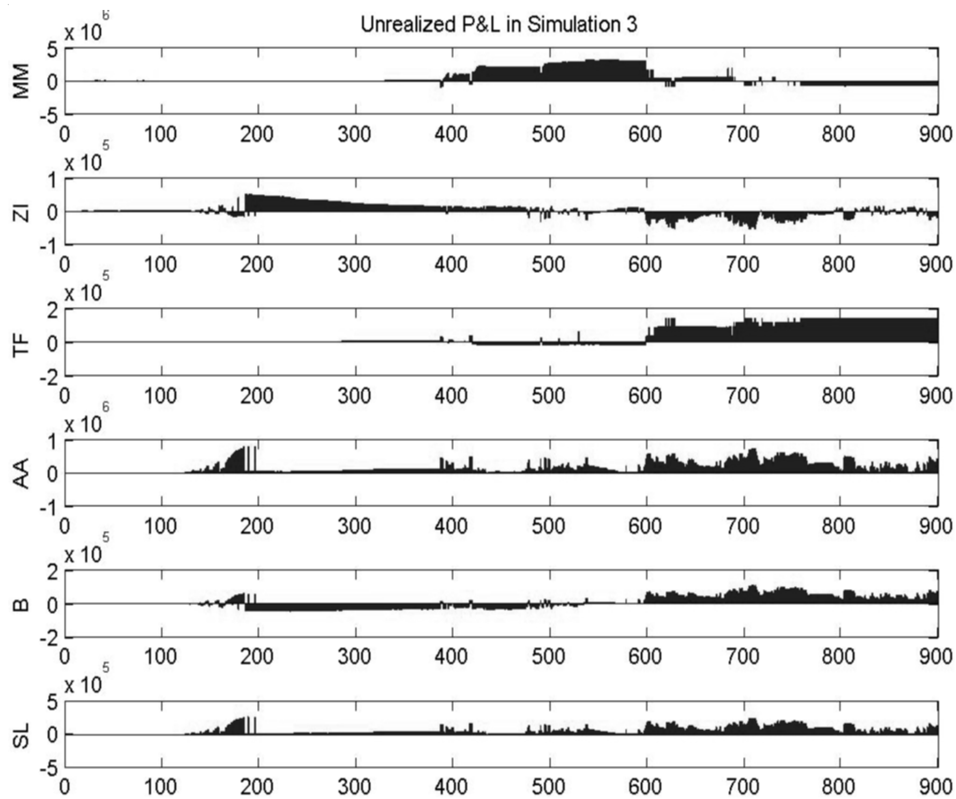

Unrealized P&L in Simulation 3 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

> 1. If a long position is cancelled by the exchange after the trading session, then it is assumed that the agent has to "replace" the position at the asset's closing price, resulting in a negative P&L impact.
>
> 2. If a short position is cancelled by the exchange after the trading session, then it is assumed that the agent has to "replace" the position at the asset's closing price, resulting in a positive P&L impact.

The most interesting observation from Table 5 is that Market Makers and Zero-Intelligence end up bearing most of the impacts. These 2 agent types must quote or place trades based on the simple assumption that the bids and offers are evenly distributed. They are likely to suffer whenever there is a massive market adjustment in any one direction. Exchange officials should be aware of these unintended fairness issues before deploying the blunt tool to "bust" trades.

Figure 22. Unrealized P&L in Simulation 4 for Different Agent Types.

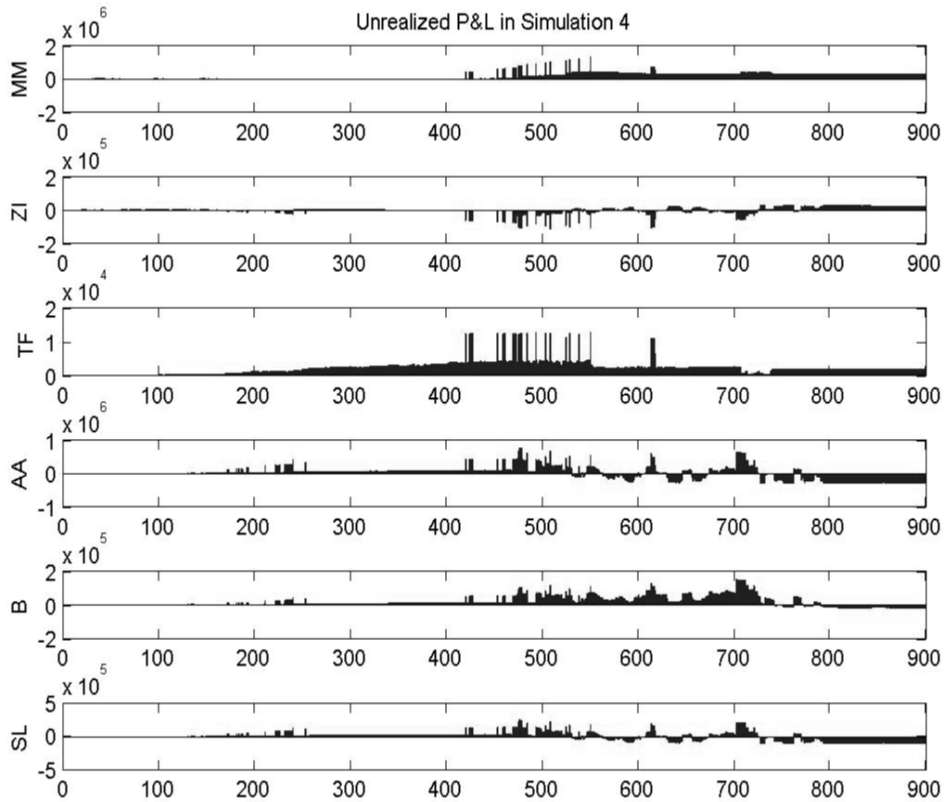

Unrealized P&L in Simulation 4 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents).

5. CONCLUSIONS AND RECOMMENDATIONS

The authors contend that the events of May 6, 2010 exhibit patterns consistent with the type of "flash crash" observed in their earlier study. While some commentators assigned blame on the May 6, 2010 "flash crash" to high-frequency trading, the authors suggest that the issue may be less about high-frequency trading per se, but rather the domination of market activities by trading strategies that are responding to the same set of market variables in similar ways, as well as various pre-existing schemes that modify the "rules of the game" in the middle of trading. The consequent lack of market participants interested in the "other side" of their trades may result in a significant liquidity withdrawal during extreme market movements.

This paper describes an attempt to reconstruct the critical elements of the market events of May 6, 2010 based on the five hypotheses posed initially by the Joint CFTC-SEC Preliminary Report and the corresponding Final Report. The authors contend that the simulated asset price distributions have shown "reasonable

Figure 23. Unrealized P&L in Simulation 7 for Different Agent Types.

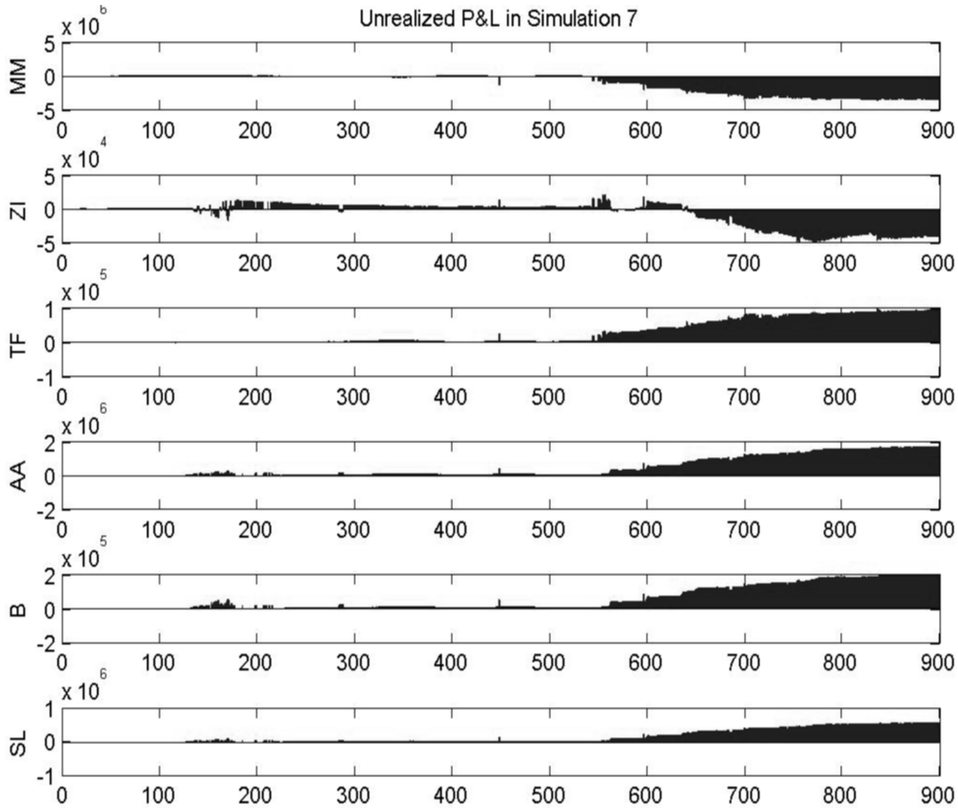

Unrealized P&L in Simulation 7 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents)

resemblance" in descriptive statistics without over-fitting historical data.

Our specific recommendations are:

1. Any scheme to"slow down" trading does not address the fundamental demand and supply imbalance leading to flash crashes, and it may cause more problems than it solves.

2. In a "fragmented" market with parallel trading venues, the "action-reaction" nature of complex exchange rules to alter the speed of trading may initiate a chain reaction that may drive liquidity further out of the aggregate market. Thus, it is important for parallel trading venues to coordinate their responses to avoid creating unintended domino effects.

3. The uneven slowing-down of trading at different trading venues often results in non-convergent fair values, because there is no or limited liquidity to complete one of more "legs" in an arbitrage trade. Arbitrageurs may suffer heavy losses in such markets, resulting in further withdrawal of

Figure 24. Unrealized P&L in Simulation 8 for Different Agent Types.

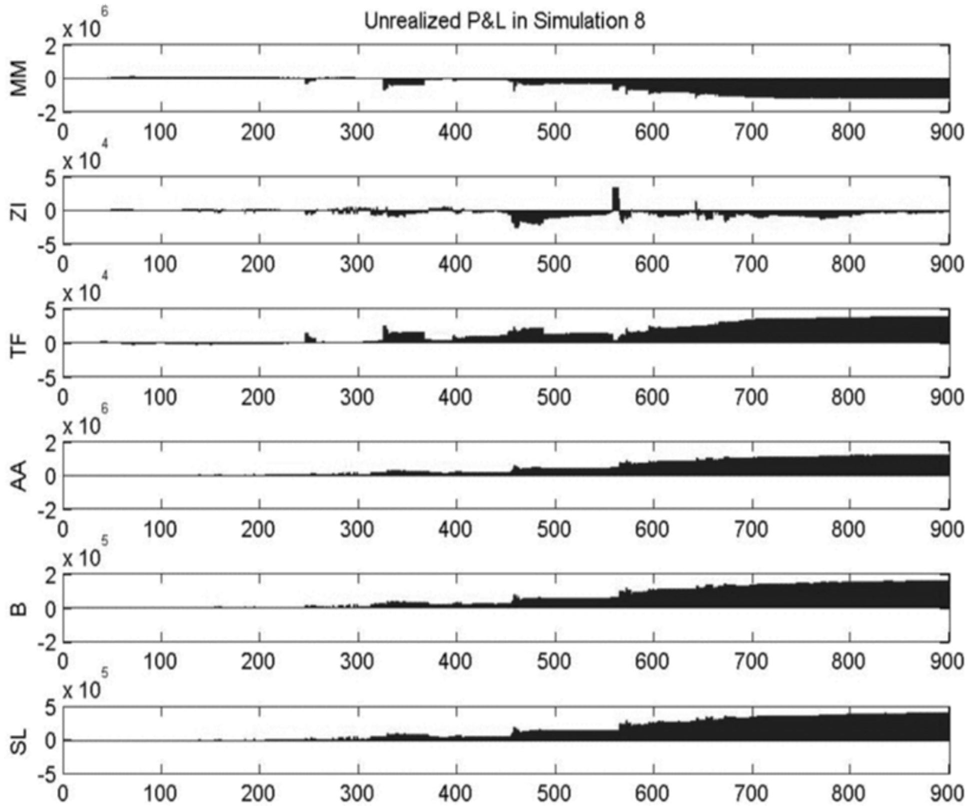

Unrealized P&L in Simulation 8 for different agent types. (MM = Market Maker, ZI = Random Agents; TF = Trend Follower, AA = Arbitrageur, B = Bear-Market Seller; SL = Stop-Loss Agents)

liquidity due to their needs to "reverse out" from loss-making, incomplete arbitrage trades. Thus, it is important for parallel trading venues to coordinate the execution of their responses — in the event that going into a "slow mode" is the correct response, then its execution should be done in parallel by all relevant exchanges to avoid needlessly amplifying the uncertainties faced by market participants.

4. The problem appears to be less about the slowing-down of trading per se. It is about the potential liquidity withdrawal due to the adjustments and chaos as a result of the initial slowing-down, as well as from the subsequent adjustments once the "normal" speed of trading is resumed.

5. "Busting trades" may discourage key participants such as Market Makers from trading in the markets as and when they are most needed. Unless there are clear technical errors involved, busting trades at arbitrary price levels is a blunt instrument that should be used sparingly and with extreme caution.

Table 5. Potential P&L Impacts of Different Agent Types.

Agent Type	Aggregated P&L without busted trades($)	Aggregated P&L with busted trades($)	Delta P&L($)
Market Maker (MM)	8,220,800	2,341.30	-8,218,458.70
Zero-Intelligence (ZI)	1,114,700	228,960.00	-885,740.00
Trend Follower (TF)	-5,930,600	184,590.00	6,115,190.00
Arbitrager (AA)	-132,040	-26,852.00	105,188.00
Bear Market (B)	-1,487,700	-148,520.00	1,339,180.00
Stop Loss (SL)	-1,581,800	-37,224.00	1,544,576.00

Potential P&L impacts of different agent types due to "busting trades" at 60% or below the opening price of each asset.

6. Price limits appear to be more effective than different implementation of positions limit in terms of stabilizing the market during the period of time when the market is finding its new equilibrium due to supply and demand imbalances.

7. Price limits do have limitations. When professional traders are reasonably certain of potential market outcomes, they can normally find ways to make profits based on trading algorithms. That creates fairness issues for unsophisticated retail investors or market makers who are under obligations to quote. Therefore, the deployment of such blunt tools should be a regulatory policy of last resort.

References

Cheng, S-F., 2007, Designing the Market Game for a Commodity Trading Simulation. *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, San Jose, CA, November, 445-449.

Easley, D., Lopez de Prado, M.M. and O'Hara, M., 2010, The Microstructure of the 'Flash Crash': Flow Toxicity, Liquidity Crashes and the Probability of Informed Trading. *The Journal of Portfolio Management*, 37(2), 118-128. Available at SSRN: http://ssrn.com/abstract=1695041.

Findings Regarding the Market Events of May 6, 2010, 2010. *Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*, September 30. Available from http://www.cftc.gov/ucm/groups/public/@otherif/documents/ifdocs/staff-findings050610.pdf.

Kirilenko, A.A., Kyle, A.S., Samadi, M., and Tuzun, T., 2011, The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. Available at SSRN: http://ssrn.com/abstract=1686004.

Lee, B., Cheng, S-F., and Koh, A., An Analysis of Extreme Price Shocks and Illiquidity among Systematic Trend Followers. *Review of Futures Markets*, 18, 385-419.

Preliminary Findings Regarding the Market Events of May 6, 2010, 2010. *Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues*, May 18. Available from http://www.cftc.gov/ucm/groups/public/@otherif/documents/ifdocs/opa-jointreport-sec-051810.pdf.

What Happened on May 6?, 2010. *CME Group Research Report*, May 10. Available from http://www.cmegroup.com/files/May-6-2010-Incident.pdf.