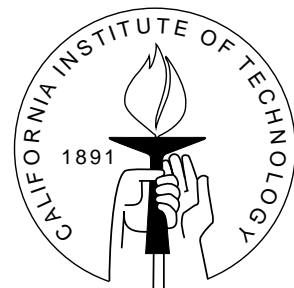# MARKET MICROSTRUCTURE DESIGN AND FLASH CRASHES: A SIMULATION APROACH

Paul Brewer
Economic and Financial Technology Consulting

Jaksa Cvitanic
california Institute of Technology

Charles R. Plott
California Institute of Technology

MARKET MICROSTRUCTURE DESIGN AND FLASH CRASHES: A SIMULATION APPROACH

Paul Brewer, Jaksa Cvitanic, Charles R. Plott*

December 3, 2012

Abstract

We study consequences of regulatory interventions in limit order markets that aim at stabilizing the market after an occurrence of a "flash crash". We use a simulation platform that creates random arrivals of trade orders, that allows us to analyze subtle features of liquidity and price variability under various market structures. The simulations are performed under continuous double-auction microstructure, and under alternatives, including imposing minimum resting times, shutting off trading for a period of time, and switching to call auction mechanisms. We find that the latter is the most effective in restoring the liquidity of the book and recovery of the price level. However, one has to be cautious about possible long term consequences of the intervention on the traders' strategies.

MARKET MICROSTRUCTURE DESIGN AND FLASH CRASHES: A SIMULATION APPROACH

Paul Brewer, Jaksa Cvitanic and Charles Plott

I.   Introduction

In recent years, concern has been expressed about the effect of high-frequency trading in electronic markets. In particular, the infamous ``flash crash" of May 2010 has prompted regulators all over the world to consider restrictions on trading and/or and new market microstructure designs through which trading would take place. Among the objectives is the prevention of sudden large price drops, or, at least, to make them short-lived.

This paper is focused on the consequences of alternative policies for a fixed order flow. That is, the order flow remains constant, the policies are changed and the market responses to a "flash crash" causing event are studied in the context of the alternative policies. The impact of the policies on subsequent liquidity and related volatility is thus studied within a framework in which the order flow itself does not change as a consequence of the policy or subsequent volatility. Uncoordinated order flow has a direct effect on the market and that is our focus. Clearly, order flow can react to the market or the policies themselves and those relationships can be studied, but the coordination of the order flow in response to the market could rely on an additional set of principles.

The proposed interventions we study have a purpose of increased liquidity and a smoothing effect in the times of market instability. The objectives are to mitigate price changes and volatility due to flash crashes and enhance recovery after flash crashes. Specifically, we study: (i) imposing minimum resting times in limit order books (LOB), that is, banning quick cancellations of buy/sell orders when not executed at arrival; (ii) switching to call auction markets instead of the prevalent continuous double auction markets; and (iii) other types of "circuit breakers", that is, interventions in the market with the aim of re-establishing price stability after large moves. We use a simulation approach[1] in which the focus is on the order flow and its interaction with the market micro

---

[1] We perform our analysis using a simulation tool platform that can process, execute, and allocate in the range of 10,000 - 80,000 bids and asks per second. The tool then produces the trade price series that occurs as a result of these bids and asks, their timing, or ordering, and the trade matching rules of the marketplace. A market microstructure rule can be adapted as desired, e.g. specify that a trade occurs whenever the lowest available seller's ask price is met or exceeded by a buyer's bid price, or by contrast specify that the totality of bids and asks determine price as in a call market, or otherwise specify how the price is determined from the bid price(s) and/or the ask price(s). This flexibility is an important

structure, as opposed to strategic behavior or individual decision rules. Buyers and sellers arrive randomly to the market and submit bid or ask orders in a continuous market. Orders are placed in an order book according to price and time priority and stay until executed or cancelled. Price movements and liquidity are products of the order flow. Traders' behavior that produces the flow can be interpreted as being along the lines of the competitive equilibrium model, but the appearance of traders at the market and their preferences when arriving are random, as dictated by background properties of the economy in which they are operating. For example, (the majority of) the traders could be viewed as brokers who are acting on a commission and who submit orders received from their clients. The clients, on the other hand, are motivated by a common background economy but their motivation and timing in response to that background are modeled as random. That is, for the results reported here order flow does not operate through specific, event-coordinated strategies. Thus, while our environment is a continuous market in which buy and sell orders arise from random outcomes, the environment is a flexible framework that can be readily modified, for example to have strategic traders.

We cause a flash crash by a submission of an extremely large order. We then analyze the process through which the large order influences the liquidity of the book and the instability of the transaction prices for a given structure of order flow. The method facilitates a study of the conditions under which the impact of a flash crash might be substantial and of the mechanisms that might serve to mitigate the effects.

While the flash crash is initiated by our deliberate action, it is the type of event that could arise on its own. For example, the continuous time nature of the market invites a randomness of the timing in which uncoordinated and decentralized traders execute orders creating an underlying randomness. That itself can be a cause of a crash, for example as many sell orders happen to appear in the market at the same time. The bottom line is, we study flash crash causes in the simplest possible framework, not necessarily corresponding to the special facts of any historical situation. The purpose is to obtain insight about what could go wrong generically.

---

innovation that allows the study of different market architectures while keeping the order flow controlled as needed for market structure comparisons.

As part of the platform, there is a market simulator using the JavaScript language and a new server-side JavaScript interpreter known as NodeJS (developed by others and released as free software). Unlike other languages, JavaScript's original development as a web-browser language has lead to an asynchronous event-driven execution model derived from requirements to handle events that can fire at known or unknown times or rates. It is thought that an asynchronous model may be more appropriate for programming market mechanisms where the events are order flow related. We use a simulated time approach whereby the computer creates a time stamp for each event in a list and processes the collection of events according to marketplace rules as though the events occurred at the indicated times. This approach allows the simulation of high frequency (HF) flows without large investments or limitations due to computation times that might vary from machine to machine or from year to year. It is faster, which is important because the complexity of simulations is limited by desired waiting time for results.

II.   Background Literature

To our knowledge, there have been very few academic studies of minimum resting times, or call auctions, or other measures of regulating electronic trading. Two recent papers use a simulation approach, as we do. Lee, Cheng, and Koh (2010) consider a market that consists of two types of traders, systematic traders or trend followers, and "zero-intelligence" traders. As the percentage of trend followers (who all apply a similar strategy) increases, the market prices break down and there is a decrease in liquidity. They find evidence that injecting and reducing liquidity by a market maker can both be effective. They suggest that imposing minimum resting times might be a way to control liquidity, and thus, might be helpful. However, they also find that the market maker can accumulate large losses by buying in a one-sided, falling market. Therefore, they conclude that 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 the paper Lee, Cheng, and Koh (2011) the same authors add arbitrageurs and market makers as two additional types of traders. They claim that problems of market instability might be less about high-frequency trading per se, but rather, about the domination of market activities by trading strategies that are responding to a given set of market variables in similar ways. They offer the following conclusions:

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

2.   If there are parallel trading venues, rules to alter the speed of trading may chase away traders to other venues, and may drive liquidity out of the aggregate market. Thus, it is important for parallel trading venues to coordinate their responses to avoid creating unintended domino effects.

3.   Slowing-down trading may lead to potential liquidity withdrawal due to traders' adjustments.

The classic paper Madhavan (1992) compares call auctions with continuous auctions and finds that call auctions lead to more stability and better information aggregation. However, continuous auctions are more popular in practice, and this discrepancy between theory and practice is a puzzle that has not been fully resolved. Coppejans, Domowitz, and Madhavan (2004) find that in electronic limit order markets shocks to liquidity dissipate quickly, indicating a high degree of resiliency, which is in accordance with our results.

Equilibrium models of limit order markets (with continuous auctions) include Parlour (1998), Foucault (1999), Biais, Martimor, and Rochet (2000), Parlour and Seppi (2003), Foucault, Kadan and Kandel (2005), Goettler, Parlour, and Rajan (2005), Back and Baruch (2007), and Biais and Weill. (2009), among others. These papers aim to derive the equilibrium price formation process. Most of these papers do not compare different market designs. In order to cope with the large dimensionality of the state and action spaces of limit order markets, these studies use stylized models with a lot of simplifying assumptions. In contrast to the equilibrium considerations, we study the formation of transaction prices given the distribution of orders, which we take to be exogenous. Dynamic, expected utility maximization models include those of Avellaneda and Stoikov (2008), Cont, Stoikov, and Talreja (2008), Kuhn and Stroh (2009), and Rosu (2009). Those studies assume specific functional forms that govern the traders' preferences. Our approach is more pragmatic: our simulation method works for any possible distribution of orders - equilibrium or otherwise.

Of particular relevance to our study, is the science that evolved from a long history of the use of laboratory experimental methods to study the principles that govern market behavior, including price discovery, efficiency, and volatility. Using financial incentives to create markets with controlled parameters, economists have demonstrated that the underlying price discovery process is governed by the law of supply and demand. The original discovery was fundamental (Vernon Smith was awarded a Nobel Prize in Economics) and has been extended to a wide range of economic conditions, parameters, and market institutions (see C.R. Plott and V. L. Smith, 2005). For example, it is well established that the CAPM follows those fundamental principles, see Bossaerts, Plott, and Zame, (2007). The classical settings of experimental markets were generalized by Alton and Plott (2007, 2008),(henceforth "AP (2007, 2008)"), to include the study of markets in which the arrival of traders in the market is stochastic. Our model is based on AP (2007, 2008). The basic supply and demand continuously change according to the randomness of traders' private values.

III.   ECONOMIC ENVIRONMENT

Our study is based on the model of AP (2007, 2008), in which buyers and sellers arrive randomly to the market for units of one asset, according to independent Poisson processes with given arrival rates. Each buyer/seller is assigned

a random reservation value for the trading asset. For example, a buyer with assigned value "x" would not pay more than "x" for one unit, but is willing to pay less. A possible interpretation is that she can sell the asset outside of the market at her reservation value (which was, in fact, the case in the experiments performed by AP (2007, 2008)) . The reservation values are drawn randomly from fixed distributions. In the simulations, we draw the reservation values in the iid fashion mostly from uniform distribution, but also sometimes from normal distribution, and, in some simulations, relative to the previously traded price, conditionally on being profitable to the trader.

In the benchmark case for our simulations, the traders submit their reservation values as limit orders, with no expiry. In these simulations, each order can be viewed as having been tendered by a different buyer or seller, so the number of traders can be viewed in terms of thousands. Even though the arrival of buyers and sellers and their reservation prices are determined at random in the AP (2007, 2008) framework, they show that the concept of Flow Competitive Equilibrium (FCE) can make rough predictions about market price behavior. FCE is defined as the price at which the expected number of buys is equal to the expected number of sells during the period of simulation, and is, thus, a form of a supply-equals-demand concept.

For the majority of the simulations, our reservation values are drawn uniformly from the range of [1,100], leading to an FCE at a price of 50. Given model parameters, prices have high variance around FCE. If we change parameters of the model, for example the frequency of arrivals, the FCE changes, and we see a shift in trading prices, as illustrated in Figure 1. The figure shows the FCE's computed from expected supply and demand at the intersections of the lines at the left and right panel, compared to the median of observed prices, marked by orange arrows in the center panel which displays the observed trading prices.

[FIGURE 1]

The qualitative results we obtain below have been shown to be robust with respect to various parameter values we used. That is, the effect of flash crashes and behavior of the prices thereafter under varying microstructure assumptions are robust with respect to parameter changes. Moreover, the model scales with respect to the speed of order flows and range of order values, so a particular choice of the values for those parameters also does not matter for the nature of the results. Our goal here is not to describe the exact time series properties of the price, but rather,

the mechanisms through which the impact of flash crashes take form, the qualitative impact of flash crashes and the qualitative comparison of policies implemented to reduce those impacts.

IV.   Flash Crash Effects

In this section, we consider the risks associated with a flash crash occurring under the benchmark market structure. As explained above, we perform simulations of limit order markets with traders who have private values drawn randomly from given distributions (one for buyers and one for sellers) and who arrive to the market at random times. We add to this market one large order to cause a flash crash, and study the properties of the order book after the crash. The following is the list of the general conclusions we obtain, with references to the related figures.

- Flash crashes can be caused by events and practices that destroy liquidity, and they create subsequent volatility (Figures 2, 3, and 5-11)

- Liquidity-destroying practices include large market orders that execute immediately (Figures 5-12), and short lived orders, including high frequency orders (Figures 2, 3, and 10, 11)

- The impact of short lived orders on the market depends on the proportion of traders using short lived orders (Figures 2, 3 and 10, 11)

- Short lived orders close in price to supply/demand equilibrium have little effect (Figure 4)

- The severity of the impact of large orders on the market can be increased by the proportion of short lived orders (Figures 2, 3, 10, 11), and change of orders frequency in the subsequent order flow (Figures 8, 9)

- When there are no reactions in subsequent order flow to a large order, recovery in price occurs quickly (Figures 5, 7, 10, 11), but weakened order books allow the price to revisit low values (Figures 5, 7, 10)

- Intervention in markets impacted by large orders should focus on rebuilding liquidity as measured by the order books, since healthy order books limit the range of subsequent prices.

By introducing a large market sell order, we cause an immediate crash by removing liquidity in the book. In our benchmark simulations, we observe that prices tend to recover quickly, as they do in most of variations on the model, except they do not recover if there is a shift in the market parameters (for example a decrease in the frequency of buy orders after the crash). After a crash, the order book contains fewer orders, which makes the

market vulnerable to increased variance and further low prices. The robustness of a recovery cannot be judged from transaction prices alone, because it is the orders in the book, or lack of book orders, that create the potential for renewed weakness.

In most of our simulations, there are equal number of buyers and sellers. The initial order flow has a rate of 100 buy orders/sec plus 100 sell orders/second with values similar to the left panel of Figure 1.

Short lived orders. In this extension of the benchmark, each trader initially submits orders to the market as a Good-Till-Cancelled (GTC) orders (no expiry). However, each 10 seconds, 10% of buyers and sellers begin placing a 0.001 sec expiration limit on their orders. This 0.001 sec expiration has virtually the same effect in our model as fill-or-kill orders. The order is only available in the book for a brief moment, and usually shorter than the expected arrival of the next order. At the end of 100 sec, all buyers and sellers are sending in orders on short expiration. As a consequence, trade slows considerably and most trades occur at the extreme prices created by leftover GTC orders. This is illustrated in the transactions prices reported in Figure 2 showing that as the transaction rate slows, prices tend to spread out towards the extreme values of 0 or 100 and often take these extreme prices, only occasionally hitting a price in the middle.

[FIGURE 2]

This happens because, when traders switch to shorter order expiry, the bid/ask spread expands as the liquidity near the equilibrium gets removed. In other words, the orders that have a longer life build up the book, and if a substantial part of the orders are changed from having a long life to a short life two things happen. The long life orders are removed from the book through trades (which would happen under any circumstance), but the short life orders that missed a trade do not remain in the book to create liquidity. If the short life orders are in an identifiable range, one can see only a small buildup of orders at the prices in that range. Outside the range one can see larger queues in the book. So, looking at a snapshot of the book one sees a price with a big queue in the book and even bigger queues as one gets further from equilibrium. However, between these two big queues there are only small queues. It looks like the "hook'em horns" (index and small fingers extended and two middle fingers not extended). The extended fingers are the size of the order book on either side of the price range of the short life orders. These big queues of orders provide the liquidity that keeps the market inside the horns. The hook'em horns phenomena

with the horns that do the hooking (corral the trades) become wider as the range of the short term orders becomes wider. In short, changing the life-span of orders has a transforming effect on the market.

Figure 3 bins the prices in a simulation in which all buyers and sellers use the same expiration time for orders. The height of the bar shows the number of times a particular trading price was observed. The expiration time is varied from 1ms to 10 sec. These all use the AP framework of traders submitting randomly price orders with reservation values iid uniform on [1,100]. The Poisson arrival rate for orders from each of the buyers and sellers is 200/sec total. In terms of total number of trades, it appears that even fairly short expiration times can produce trading, however the trading initially spans the entire range [1,100] and is characterized as two-party trade rather than trade mediated by a market. In the figure, we see that longer expiration times lead to book formation and limit the domain of prices. Prices become less spread out, and more concentrated around the equilibrium price, and there is much more trading going on.

[FIGURE 3]

High frequency orders in narrow range. In Figure 4, we obtain different conclusions in another simulation, where we have very fast fill-or-kill, fixed-value (rather than random from an interval) orders submitted to buy at a price of 48 and sell at a price of 52. This does not have a significant effect on the price formation, other than increased number of trades around that value. This is a stylized model of HF traders trying to make money by fast submissions and cancellations of specific bid and ask orders (so-called "sniping"). The observation that in this case there is no extreme effect on prices is in agreement with theoretical results of Cvitanic and Kirilenko (2010).

[FIGURE 4]

Large orders. Figure 5 shows the trading prices where a large sell order causes a flash crash, but the recovery can be fast. If the traders do not change their reservation values and otherwise behave as in AP (2007), a singular large order for 1000 units affects the prices only temporarily - the market can be surprisingly robust. The trading prices after the event are even a little higher than before the event. At 100 orders per second normal order flow, suddenly clearing the buy book of 1000 units, one would think, should take at least 10 sec to "recover". The data show only a difficult to notice temporary effect. There are a few revisits of the low price values immediately after the event at time=150.0 sec, but the price data otherwise look a lot like the data preceding the event.

[FIGURE 5]

In addition to the price, we also consider **the age of paired orders traded**. Changes in the range of price movements are associated with increased age of paired orders traded. More precisely, when the change in model parameters results in orders going deeper into the order book to find trading partners, as orders that exist deep in the order book have typically been there longer, the age of paired orders increases. Thus, the age of paired orders is an indication of how long orders have been in the book before the price reached them, and is a measure of what is happening in the market. If the age is getting older, it means that structural changes are taking place even though they might not be easily visible in price patterns. In a mature market, the orders near the equilibrium price trade with each other. They are not old orders. The old orders build up in the book away from the equilibrium and provide liquidity for orders that are away from the equilibrium. This liquidity is also a cushion that keeps prices near the equilibrium and keeps variance low. A movement of trading old orders in the book reflects a process of removing the liquidity (price stability) that the accumulated orders in the book provide.

Still using the above described simulation, Figure 6 shows the order age data, with a big spike for the large sell order event at T=150. A few additional old orders are matched about a second later and then a few more old orders from either side of the market are matched later.

[FIGURE 6]

Large market order: a "hammer". In a variation on the above, the flash crash occurs also due to a very large sell at market order, that we call a "hammer". Subsequent recovery can be measured by the shape of the market "jaws" (the number of orders in the book at various price points). Figure 7 shows the buy and sell order books just before and immediately after the simulated crash. For instance, bid2 is the 2nd best or 2nd highest bid price, bid5 is the 5th highest bid price, and bid100 is the 100th highest bid price. Similarly, ask2 would be the 2nd best or 2nd lowest asking price. In our simulation, the hammer hits at 150 seconds, at which point Figure 7 shows that the lower jaw of bid prices drops and the upper jaw of ask prices begins to jut out. This takes place because the big sell order removes the orders from the buy book leaving little or no liquidity on the down side. The liquidity on both sides was accumulated during the market maturing process. Once removed, it takes time to build up again. During that build up time, the market will exhibit increased downside variability. Low prices normally occur only at the very

10

beginning of trade when there are few buy orders in the book. Only after the crash at T=150 do low prices reappear, and mostly in the first second or so after the crash, although a few occur later.

[FIGURE 7]

We also performed simulations in which we compare the effect of the hammer in the case in which the orders come from uniform distribution to that in which they come from normal distribution. There is not much difference, because the normally distributed orders can be thought of as a one-to-one transformation of uniform orders (via the inverse normal distribution function). Moreover, in another simulation we have traders receive a value signal uniformly in [0,100] and bid or ask submitted to the market will be the last price plus a draw from a standard normal distribution, conditionally on the resulting price being profitable compared to the value signal. The crash is still self-correcting, and otherwise not too different qualitatively from the cases above.

We conclude that, in general, there are underlying processes of recovery that can be seen to be somewhat independent of the order generation process. After a crash, the buy book contains fewer orders. Prices from random trading can continue to touch the lows generated by the crash, until the book rebuilds.

Big sell order with slowdown of buy flow. If a big sell order flash crash creates a slowdown in the order flow from buyers (while the range of buy orders is unchanged) the market will experience slow recovery and vulnerability to higher variance and additional flash crashes, as illustrated in Figures 8 and 9. Before the large sell order the Poisson rate parameters of buyers and sellers are 10 orders per second. After the large sell order, the buyers' rates fall by 90% and recover on an exponential slope. The buy book will not rebuild as fast because buy orders are slowed down. Furthermore, and unlike the other cases, the relative slowdown in buy orders will allow the ask book to build an equilibrium at much lower prices than is the case without buyer slowdown. These patterns of prices occur because the equilibrium is shifted after the crash.

[FIGURE 8]

Figure 9 shows the age of orders for the scenario with a reduced flow of buyers after the crash. Recall that the crash at T=150 is typically associated with a spike in the order age because the large sell order necessarily matches old orders in the book. What is different here is the persistence of new sell orders matching old buy orders in the book,

11

e.g. at T=160-240 and this occurs because the flow of buyers has reduced. As more buyers return to the market and place orders, the buy order book rebuilds with fresher orders.

[FIGURE 9]

V.   **Mechanisms to alleviate flash crash effects**

We consider these mechanisms to mitigate the negative effects of flash crashes: (a) introducing minimum resting times; (b) switching to call auction market mechanism; and (c) shutting off trading for a period of time. They all help reduce instabilities in the market, but especially helpful, in our simulations, are the introduction of the call auction mechanism.

Minimum resting times. This mechanism involves requiring the traders to leave their orders in the limit order book at least for a minimum required duration, as to prevent them to cancel immediately the trades that are not executed at arrival.

Call auction mechanism. Most prevalent mechanism for trading in today's markets is a continuous (double) auction: as soon as a buy order (for example) arrives that is larger than the minimum sell order resting in the book, the trade is executed between those two orders, at the price value of the resting order. An alternative mechanism, often used at a beginning of a trading day, is a continuous call auction: all orders arriving during time intervals of specified length are collected, after which pairing of buy orders and sell orders is performed in a way that identifies a single price for each batch of orders that maximizes the quantity that can be traded at a single price. This price also maximizes the classical gains from trade or trading surplus. Procedurally, a simple call auction is carried out by priority-sorting the asks (from low price to high price), priority-sorting the bids (from high price to low price), and finding the intersection. In this case, sub-options include switching to call auction only temporarily after a flash crash, or having periodic call auctions.

Shutting off trading. This option stops the trading for a period of time. Sub-options include:

1. A one-time call market (Catch and Release): The policy is to suspend trading and collect the order flow while trading is suspended. The orders are collected in the book. When the market reopens the

accumulated orders are treated as a call. Contracts are identified and executed. Unexecuted contracts remain on the books for liquidity when continuous trading opens after the call.

2. A series of temporary call markets. The policy is to replace the continuous market with call markets for a short period of time.

3. Stop trading, clear books and resume trading. The policy is similar to starting the market fresh. The market is delayed for a period of time; the orders that would have arrived are simply thrown away. The market is then reopened for trading with the books building from an empty slate.

4. Stop trading and take no orders, keep books unchanged. The policy here is to stop trading as before, take no trades and facilitate no trade execution, keep all orders on the books and then resume trading.

We now report on simulations in which the above mechanisms were applied.

Minimum resting times. Differing order expiration times influence both the vulnerability of the market to a flash crash as well as the subsequent healing and buildup of liquidity and consequent stabilization, as we have seen above. Longer order expiration times mean large buildup in the book at various price points. This buildup in the book creates liquidity that reduces price variance in the market. We thus conclude that requiring minimum resting times may be helpful in preventing instabilities in the market. However, this conclusion is based on very specific conditions under which we performed simulations, and it is not necessarily the case that quick cancellations by high frequency traders cause instabilities, as seen also in the experiment described above, in which high frequency traders constantly submit orders at the same value.

In Figures 10 and 11, we present another pair of simulations, the first of which is our benchmark setup with a hammer, while in the second; we have 50% of the traders having very short lifetimes of 0.001 seconds. In each of the figures, the trade prices are shown, and the lower portion of the display zooms in on the detail around the crash at T=150. The latter corresponds to a market in which half of the traders are high-frequency traders who quickly cancel their orders, while the former can be thought of the same market with a ban on cancelations. We see some evidence that the liquidity is lower, and the bid-ask spread and the price variance are larger when cancelations are allowed. There is also evidence, but less clear, that the effect of the hammer is smaller and the recovery faster when cancelations are banned.

13

[FIGURE 10]

[FIGURE 11]

Call auction mechanism. We now study the behavior of call markets (CM) mechanism as a potential mechanism for trading after a flash crash. As we shall see, this results in lower variance and a more rapid recovery from a flash crash. This improved healing of the book with the CM occurs because more aggregation of orders before trading reduces the noisiness and range of the trading price, which further enhances the aggregation of orders away from the trading price into the books - orders which otherwise would become noisy trades and disappear from the book.

In the following simulation, a call market replaces the continuous double auction for the entirety of the simulation. The call market has orders accumulating for a fixed period of time, and then computes the market clearing price. Trades are executed for those orders matched with counterparties, and orders that do not trade are retained to be included in the next call.

In common with previous simulations, the following properties are maintained:

- Poisson order arrival rate of 100/sec for buyers as a group, and sellers as a group
- orders are randomly priced iid uniform on [1,100]
- Large sell order at $T = 150$

In the simulation, the call market price is determined, and the resulting trades settled, every 1 second. In Figure 12, the black line is the trade price. The red line shows the price at which 100 units are in the buy order book. Notice that from the parameters the maximum possible number of trades is 100/second (except when the large order arrives), so the black line (trade price) is always above the red price except when the large order (the hammer) arrives at $T = 150$.

In terms of average price behavior, the call market looks much the same as the continuous market with a book. The nature of the call is to smooth the transaction data, but in terms of a moving average the call market and the continuous market with a book appear similar.

14

The impact of the hammer is a fall in price, as expected. In the call market, full recovery of the "100 units" level of the buy book from the large order takes about 75 seconds, in contrast to 140 seconds for a previous simulation in an ordinary double auction. It seems there is no hope of preventing flash crashes altogether, only of reducing their impact.

[FIGURE 12]

Call markets could be substituted for the continuous market in the event of a flash crash. The next simulation captures the implications of such a policy. This is similar to simulations in the AP (2010) environment previously reported for the ordinary markets with books, but with the following features:

1. Initially there is no call market.

2. The list of orders studied is the same for both treatments for comparison.

3. A "Call Market Treatment" changes the market format to a call market 0.1ms after the large sell order "crashes" the market at T=150.

4. The No Call Market Treatment, using the same exact orders as in item 3 above, does not change the market format.

As previously shown, the red lines show the price in the book at which an aggregate of 100 units are available. Figure 13 demonstrates that the implementation of the call market restores the market to near the FCE price immediately and removes the subsequent residual variance due to the crash. It is also noteworthy that the orders coming in immediately after the crash, combined with the call market, create higher prices than existed immediately before the crash. This is due to the iid uniform random nature of the order prices.

[FIGURE 13]

Figure 14 shows that on average, as was illustrated in the call market analysis, market price variance in terms of aggregated prices is about the same when comparing the call market with the continuous market. The median prices of Figure 14 are by definition not sensitive to the extremes of the trading prices and therefore less sensitive to certain

levels of the order book that act as a bound on prices. In contrast, the extreme prices in Figure 13 are determined primarily by the health of the order books.

[FIGURE 14]

Shutting off trading. Five different types of circuit breaker policies are examined. These can be viewed as temporary treatments as opposed to a complete and lasting change in the market structure. The same set of orders underlies all five simulations.

1.   A one-time call market (Catch and Release)

As shown in Figure 15, the liquidity removed by the hammer is replaced during the trading suspension. When the market opens, the liquidity continues to build with a consequence of subsequent small market variability. More precisely, in the catch and release policy there is substantial buildup of the book, but the variability of the continuous flow of orders has a tendency to use the liquidity far away from the market. Thus, the system experiences additional variability while the liquidity is building. The fact that the liquidity is in constant use slows its buildup. By contrast, the call protects the liquidity far from the market by coordinating and limiting trading to those orders close to the market. Thus liquidity has a cushion near the natural equilibrium price when the call process stops. The cushion is removed by the variability once the market returns.

[FIGURE 15]

2.   A series of temporary call markets.

The liquidity is naturally replaced in the call market because the orders far away from the market accumulate rather than execute. Figure 16 shows the accumulation of liquidity and the subsequent return of the market to before crash conditions and behavior.

[FIGURE 16]

3.   Stop trading, clear books and resume trading

In the cases "Clear" and "Delay", the market is delayed 10 sec, the orders that would have arrived from T=150 to T=160 sec are simply thrown away. The market is then reopened for trading with the books building from an empty slate. The difference is that with "Clear" the asymmetry created in the books, that can create subsequent market

16

variability, is removed. The market buildup of liquidity is balanced because nothing is on either side of the market in the books.

[FIGURE 17]

       4.   Stop trading (10 seconds) and take no orders, keep books unchanged.

The policy here is to stop trading (10 seconds) as before, take no trades and facilitate no trade execution, keep all orders on the books and then resume trading. This maintains existing liquidity, but the asymmetry of liquidity created by the flash crash remains. As observed earlier, it requires a minute or two for the balance of liquidity to be restored and the extra price variability caused by the hammer to be restored.

[FIGURE 18]

       5.   We also provide Figure 19 for the market where nothing is done.

[FIGURE 19]

VI.   Summary

We study simulated limit order markets with buy and sell orders with random values, arriving at random times. A "flash crash" is introduced by means of a very large order. The study is focused on the subsequent effects on the liquidity of the book and the variability of the transaction prices. The simulations are performed under continuous double-auction market microstructure, and under alternative structures, including shutting off trading for a period of time, and/or switching to call auction mechanisms. While all of these help restoring the liquidity of the book and recovery of the price level, the call markets prove to be the most effective for the purpose.

Liquidity in limit order markets is summarized by the following principles:

1. An appropriate measure of liquidity is the depth (the size of the queue) of the order book at price points away from the prices at which trades are taking place. Orders that would otherwise move the market price are stopped at the price points where the depth is sufficiently large.

2. The size of the queues is governed by the arrival rate and the departure rate of limit orders (due to trades or cancellations).

3. Orders at various price points are themselves governed by a random process of uncoordinated decisions made by different bidders and based on different goals and different perceptions of the future.

4. Events that reduce the depth (the queue size) on either side of the book contribute to subsequent price volatility until the natural queuing process allows the liquidity to return. Such events include: (i) ``Flash Floods'' – the accidental arrival at the same time of one or more very large orders with limits considerable off market; and (ii) Liquidity Erosion – asymmetric arrivals on one side of the book that have limits off the market, having the effect of reducing liquidity at several price points off the market.

At a general level, we offer the following conclusions:

1. A flash crash can leave the market eroded.
2. The nature and impact of liquidity erosion is sensitive to the market structure.
3. Exact nature of crash and book erosion depends on the structure of the order flow.

In terms of policy, it may be helpful to require frequent traders to provide liquidity at all times, but requiring minimal resting times may not be very helpful, at least not for the purpose of combating flash crashes. Rather, if there is a significant chance that a sudden fall in prices will have a long-term disruptive effect, switching to call auctions until the prices stabilize may be helpful. On the other hand, in our simulations, most of the time the market stabilizes relatively soon even without outside interventions.

References:

Alton, M. and Charles R. Plott (2007) `` Principles of Continuous Price Determination in An Experimental Environment With Flows Of Random Arrivals And Departures ''. Social Science Working Paper No. 1276 Pasadena: California Institute of Technology.

Alton, M. and Charles R. Plott (2008) "The Dynamics of Price Adjustment in Experimental Random Arrival and Departure Environments". California Institute of Technology, Pasadena

Avellaneda, M. and Stoikov, S. (2008) ``High-frequency trading in a limit order book''. *Quantitative Finance* 8:217-224.

Back, K. and Baruch, S. (2007), ``Working Orders in Limit-Order Markets and Floor Exchanges''. (Forthcoming in the *Journal of Finance*.)

Biais, B., Martimort, D. and Rochet, J.-C. (2000), ``Competing mechanisms in a common value environment''. *Econometrica* 68:799–837.

Biais, B. and Weill, P.-O. (2009), ``Liquidity shocks and order book dynamics''. Working paper.

Bossaerts, P., Charles R. P., Zame W. (2007) "Prices and Allocations in Financial Markets: Theory, Econometrics, and Experiments," *Econometrica,* 75:993-1038.

Cont, R., Stoikov, S. and Talreja, R. (2008), ``A stochastic model for order book dynamics''. Working paper.

Coppejans, M., Domowitz, I, and Madhavan, A. (2004) "Resiliency in an Automated Auction". Unpublished manuscript.

Cvitanic, J. and A. Kirilenko (2010) "High-Frequency Trading and Asset Prices". Working paper.

Foucault, T. (1999), ``Order flow composition and trading costs in a dynamic limit order market''. *Journal of Financial Markets* 2:99–134.

Foucault, T., Kadan, O. and Kandel, E. (2005), ``Limit order book as a market for liquidity''. *Review of Financial Studies* 18:1171-1217.

Goettler, R., Parlour, C., and U. Rajan (2005) "Equilibrium in a dynamic limit order market", *Journal of Finance* 60:2149–2192.

Kuhn, C. and Stroh, M. (2009), ``Optimal portfolios of a small investor in a limit order market – a shadow price approach''. Working paper.

Lee, B., Cheng, S-F., and Koh, A. (2010) "An Analysis of Extreme Price Shocks and Illiquidity Among Systematic Trend Followers." *Review of Futures Markets* 18(4):385-419.

Lee, B., Cheng, S-F., and Koh, A. (2011) Liquidity Withdrawal and the ``Flash Crash'' on May 6, 2010. Working paper.

Madhavan, A. (1992) "Trading Mechanisms in Security Markets". *Journal of Finance*, 47:607-641.

Parlour, C. (1998), `` Price dynamics in limit order markets". *Review of Financial Studies* 11:789–816.

Parlour, A.S. and Seppi, D.J. (2003) "Liquidity-based competition for order flow". *Review of Financial Studies* 16:301-343.

Plott, Charles R., and Vernon L. Smith (2005), "Handbook of Experimental Economics Results Vol. 1", edited by C. Plott and V. L. Smith. Elsevier (2008).

Rosu, I. (2009), ``A dynamic model of the limit order book". (Forthcoming in *The Review of Financial Studies.)*

FIGURE 1: **Changes in order rates or distribution causes shift in trade prices**
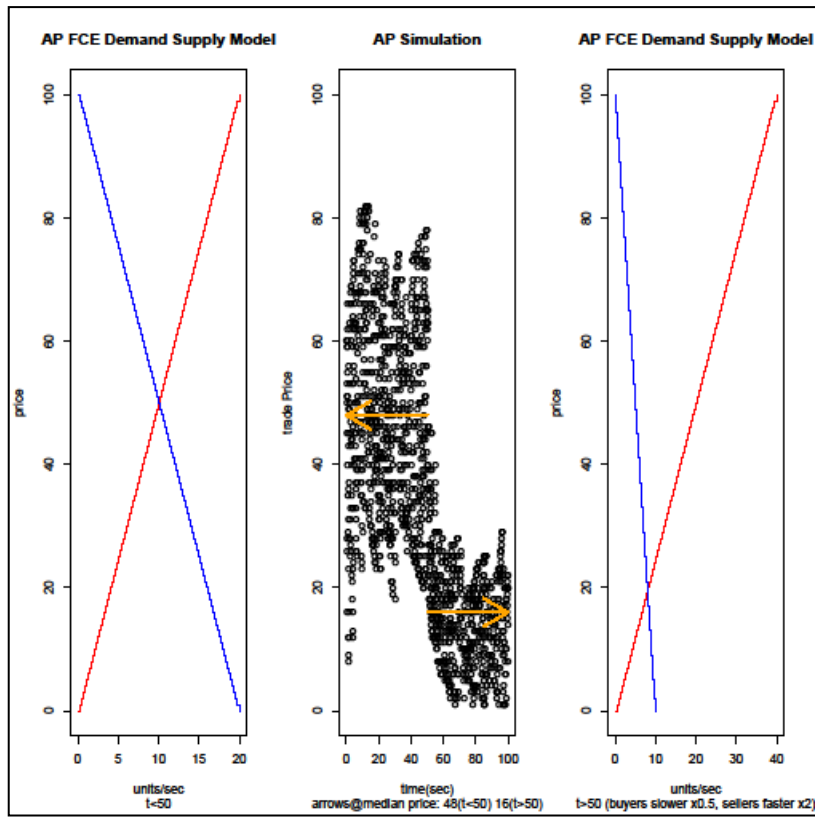
FIGURE 2. **Reducing the percentage of no-expiry orders destabilizes the market**
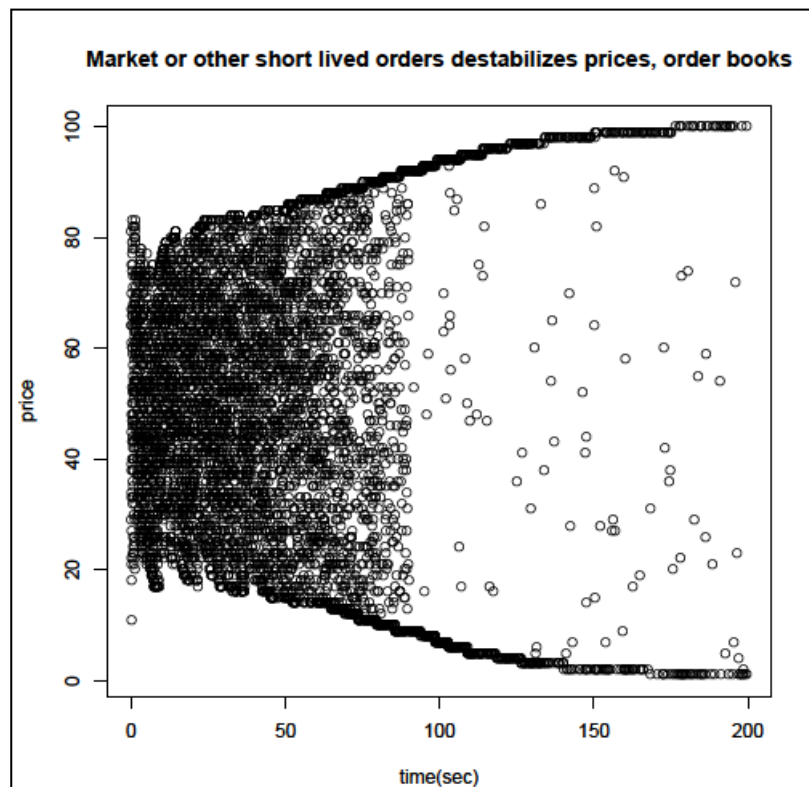
FIGURE 3: **Having longer order expiration times stabilizes the transaction prices**
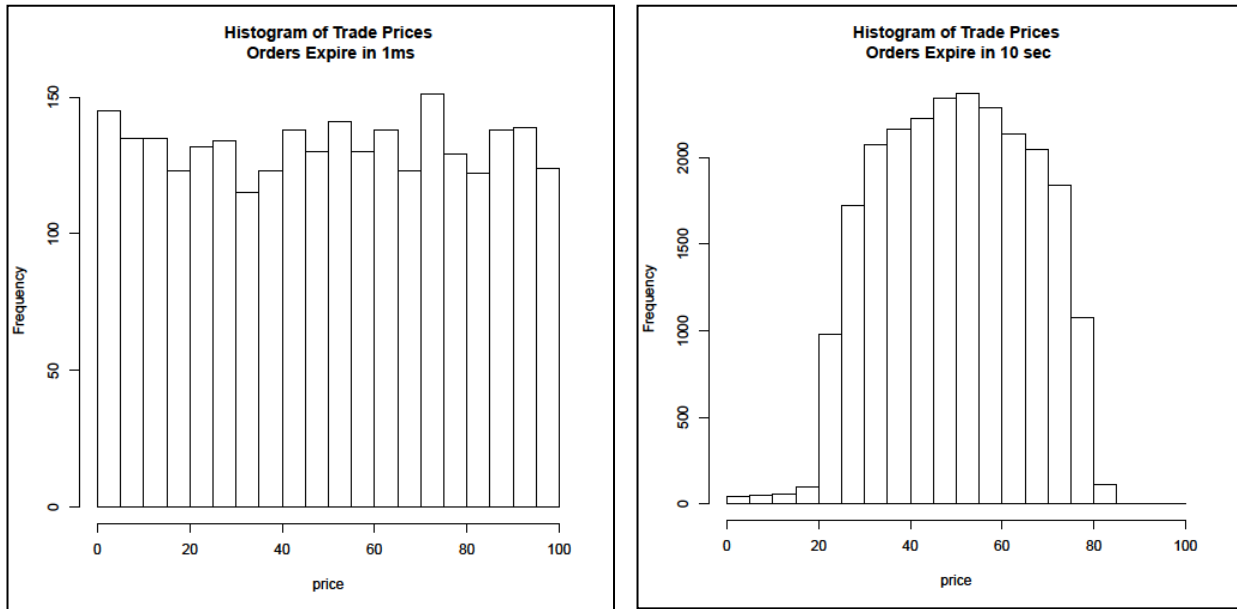
FIGURE 4. **HF short lived orders at fixed values do not destabilize the market**
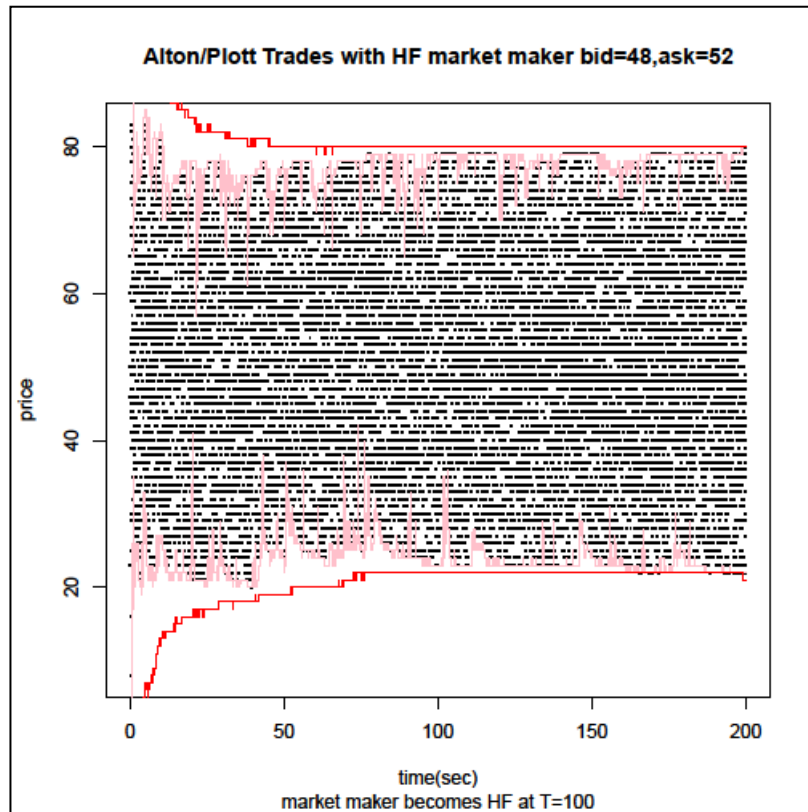
FIGURE 5: A **large sell limit order at 20 submitted at T=150**

Large sell order (sell all for P>20) at T=150 effect on market
1 sec before to 1 sec after
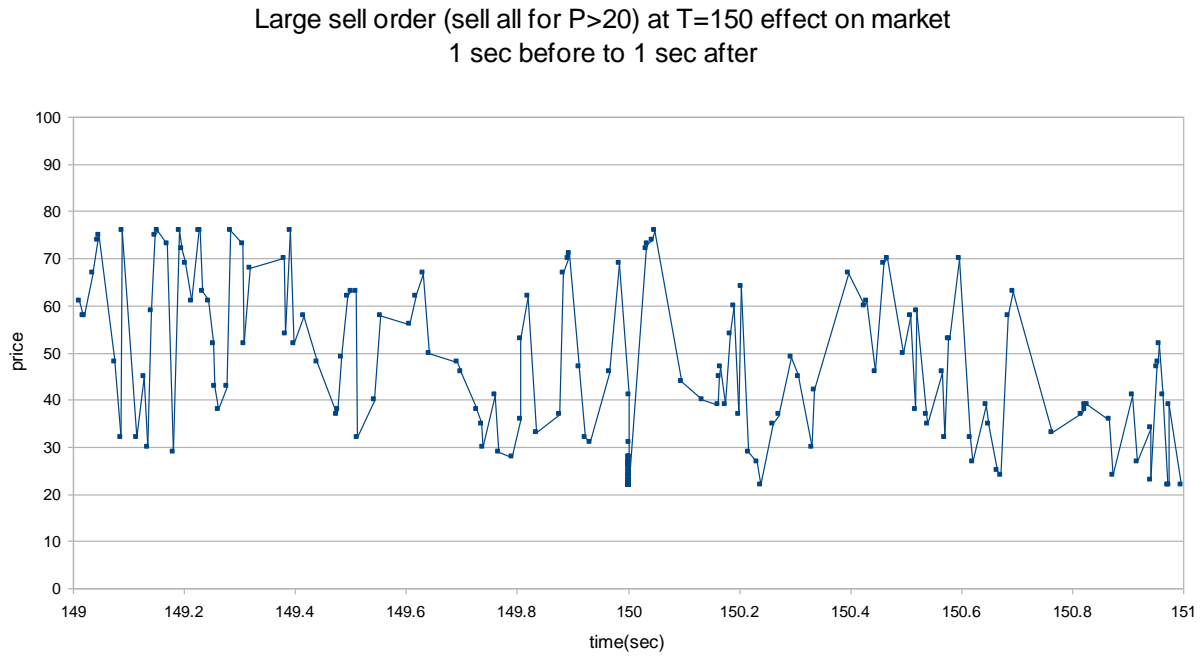
FIGURE 6: **State of the market as represented by age of matched orders**

Age of Matched Orders
Large sell order at T=150
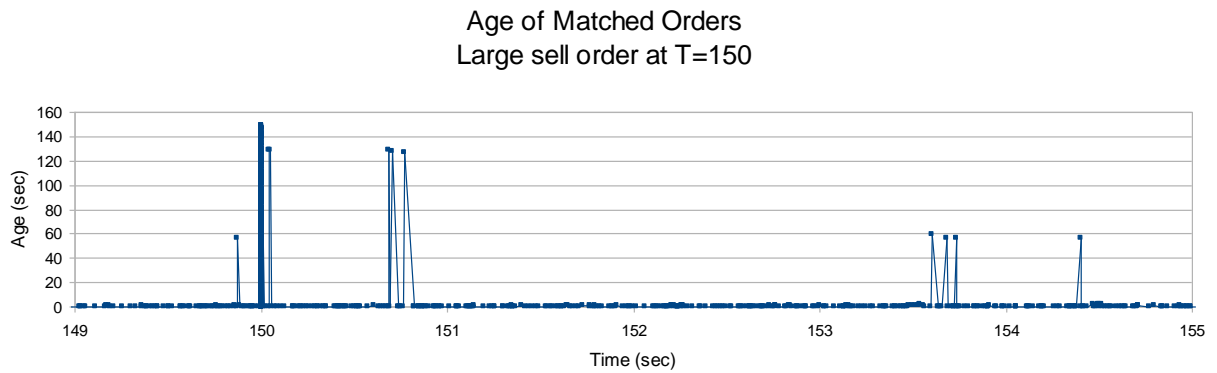
FIGURE 7: **Large sell market order at T=150**

Order prices ~U[1,100]
sell order for 1000@1 at T=150
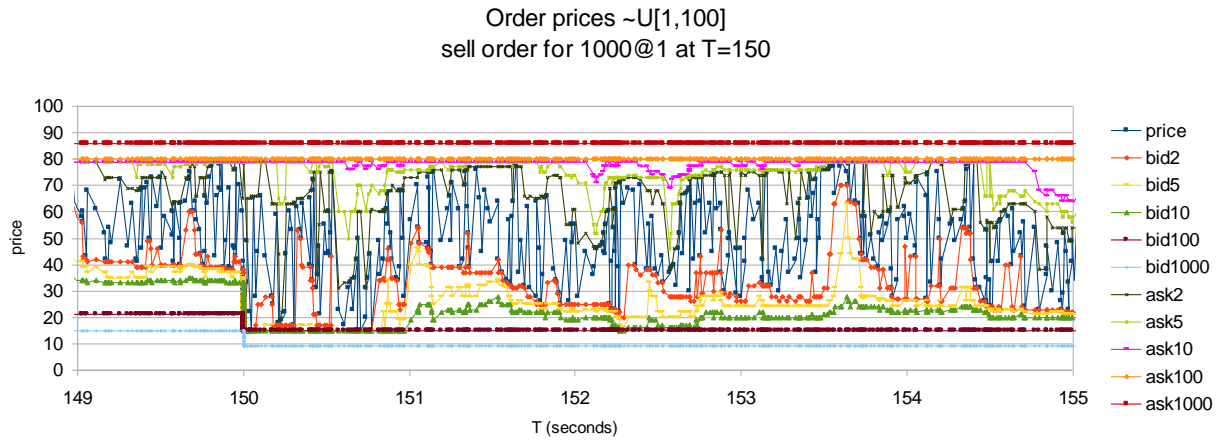
FIGURE 8: **Transaction prices with buy orders less frequent after the crash**

Orders   ~U[1,100]
Large sell order @T=150, reduced flow of buyers

24

FIGURE 9: **Age of matched orders with buy orders less frequent after the crash**

Age of Matched Order
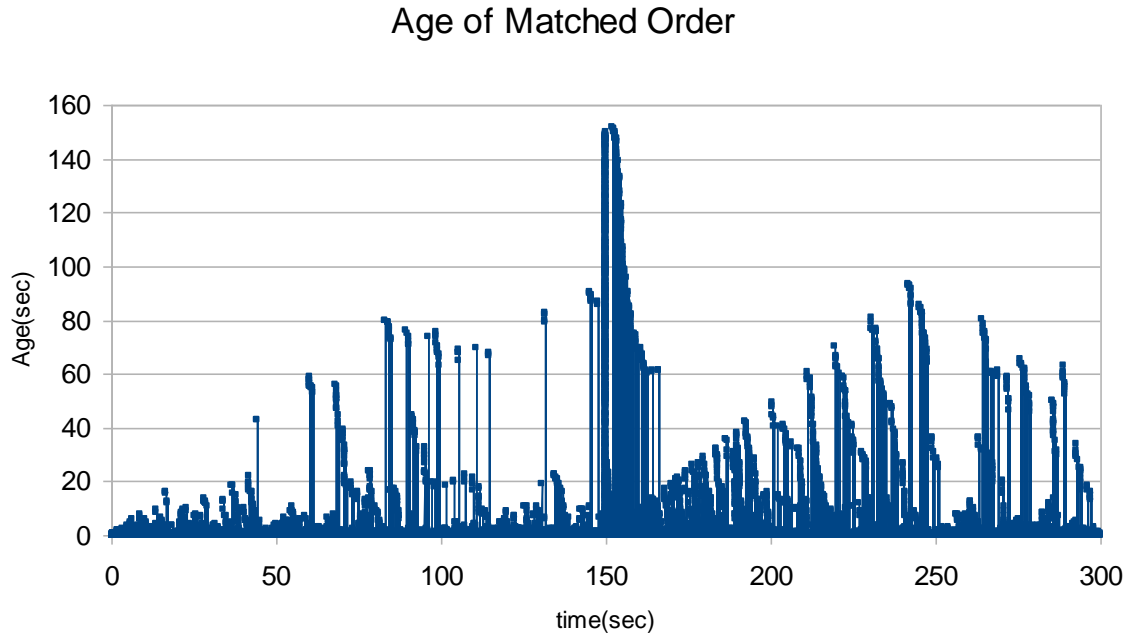
FIGURE 10: Hammer effect without HF traders
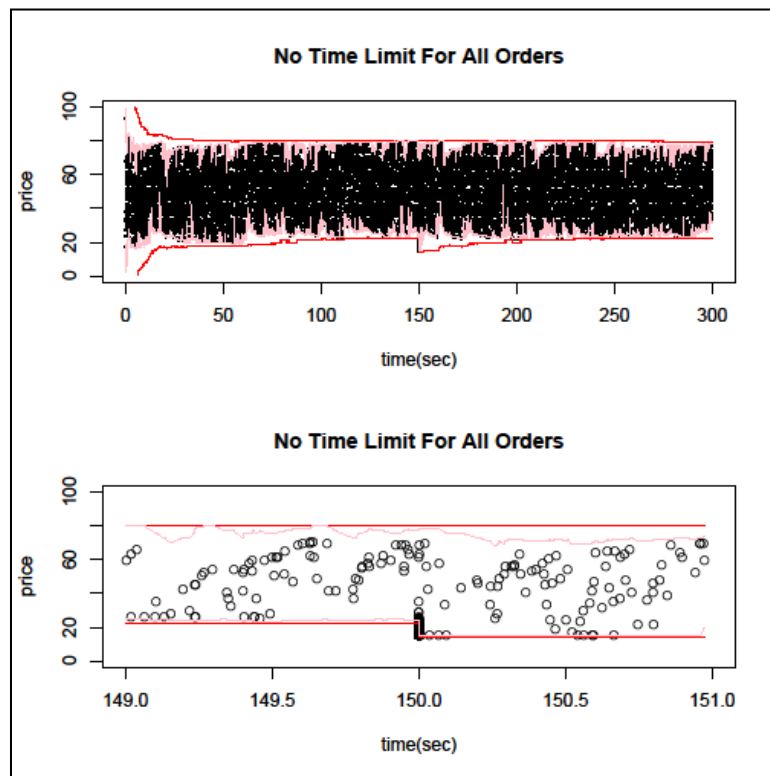
FIGURE 11: **Hammer effect with 50% of HF traders**
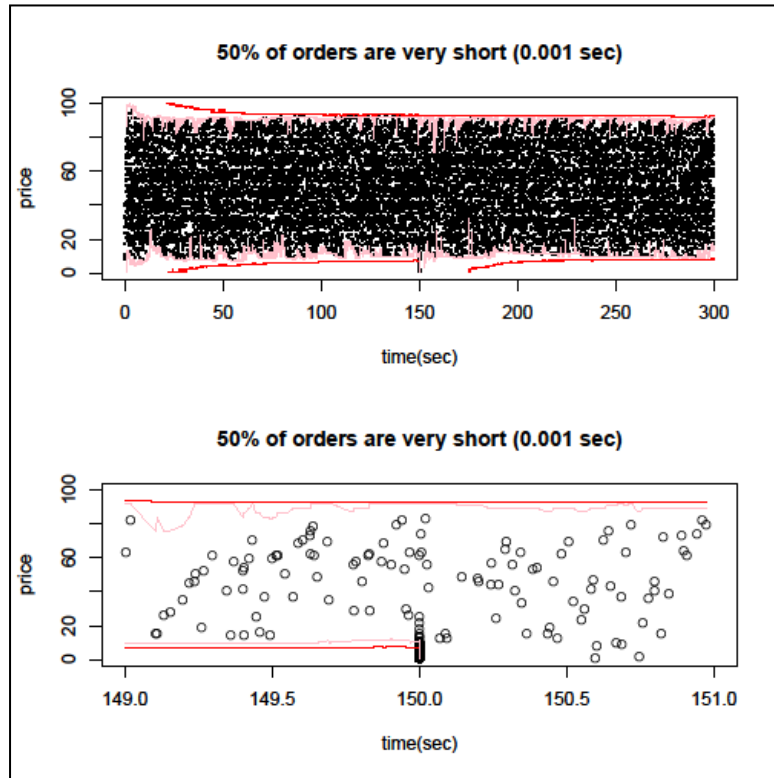
FIGURE 12: **Flash crash with call markets**
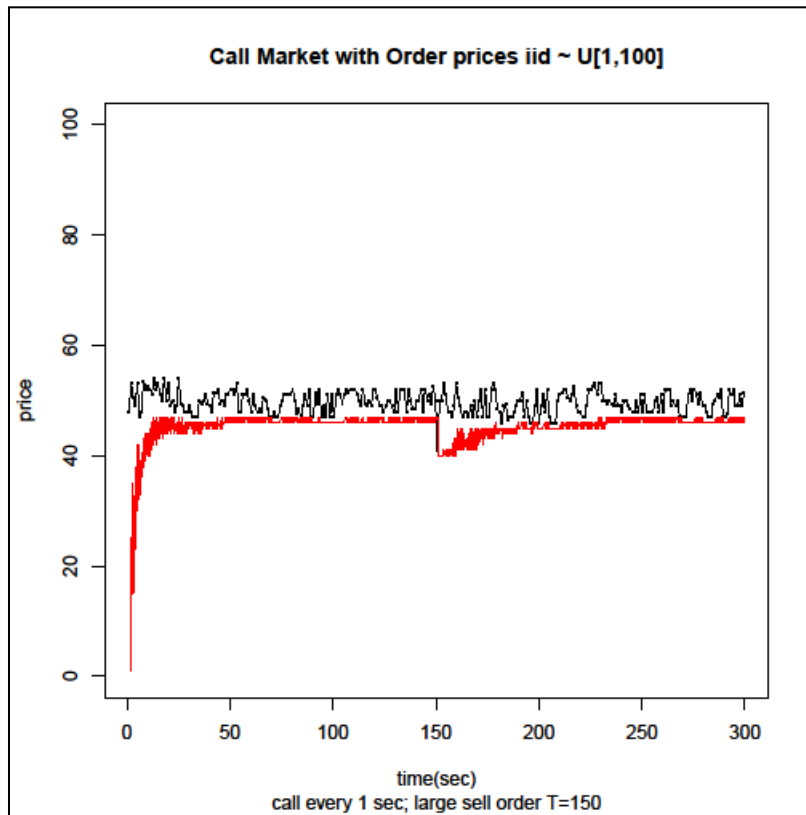
FIGURE 13: **Transaction prices with and without introduction of call market after crash**
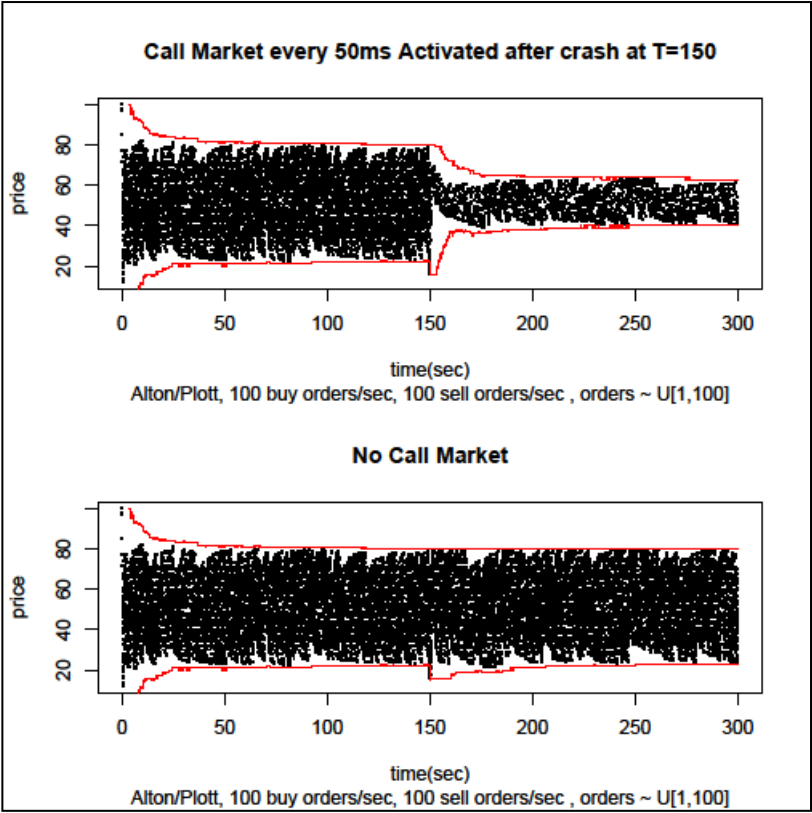
FIGURE 14: **Moving median of 101 prices with/without introduction of call market after crash**
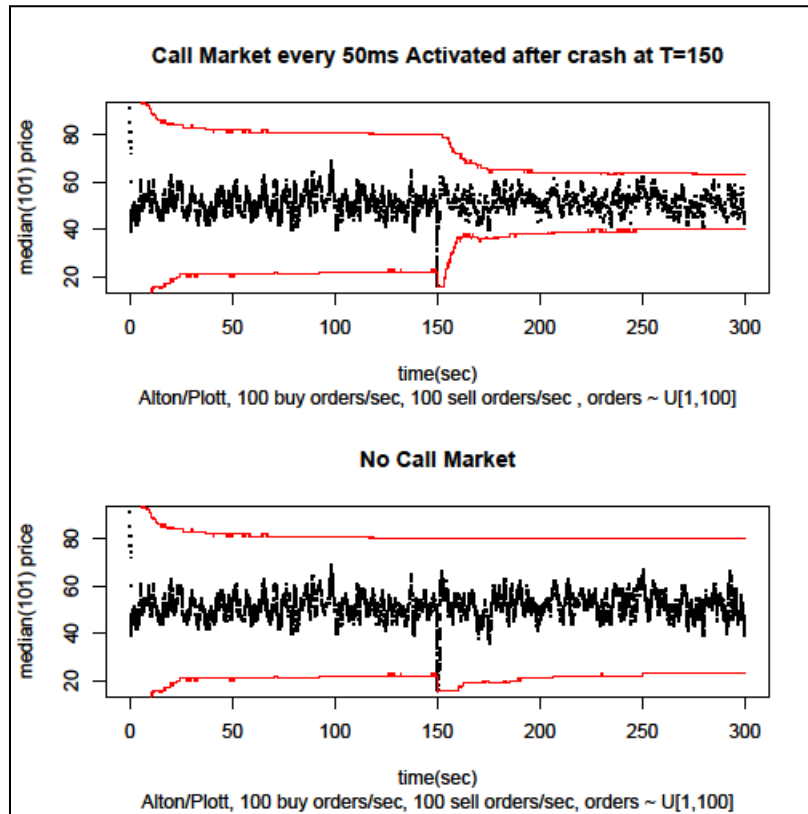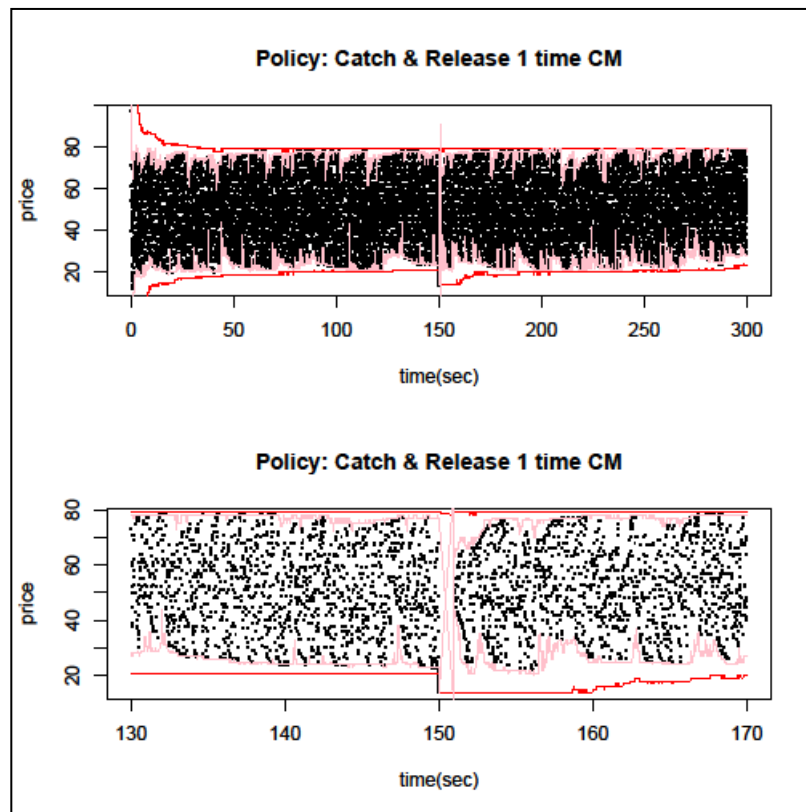

FIGURE 15. **One-time call market ("catch and release")**

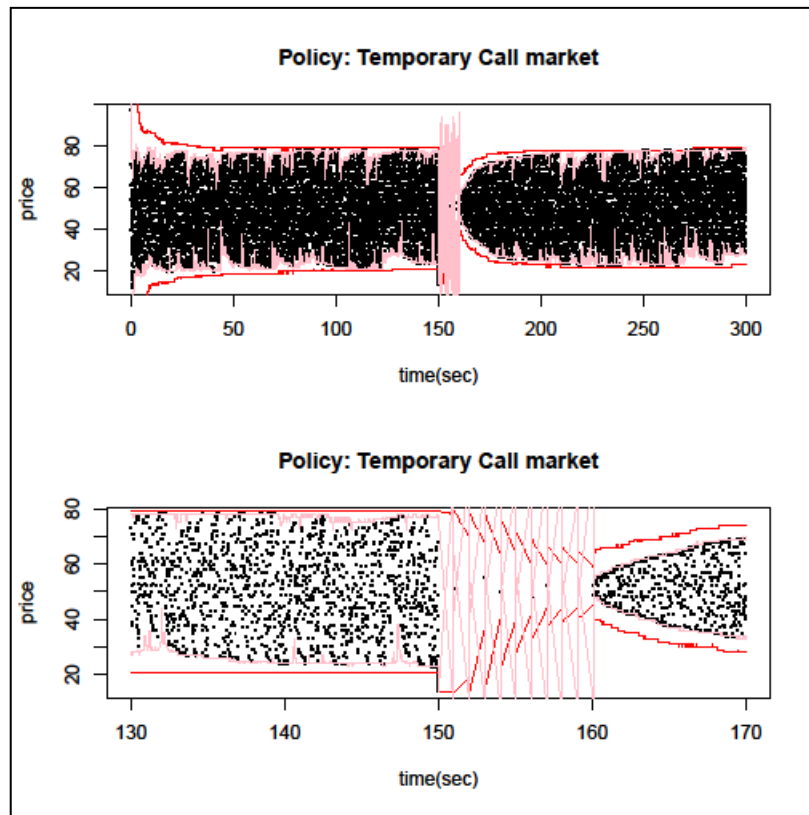FIGURE 16: **Temporary call market then standard trading resumes**

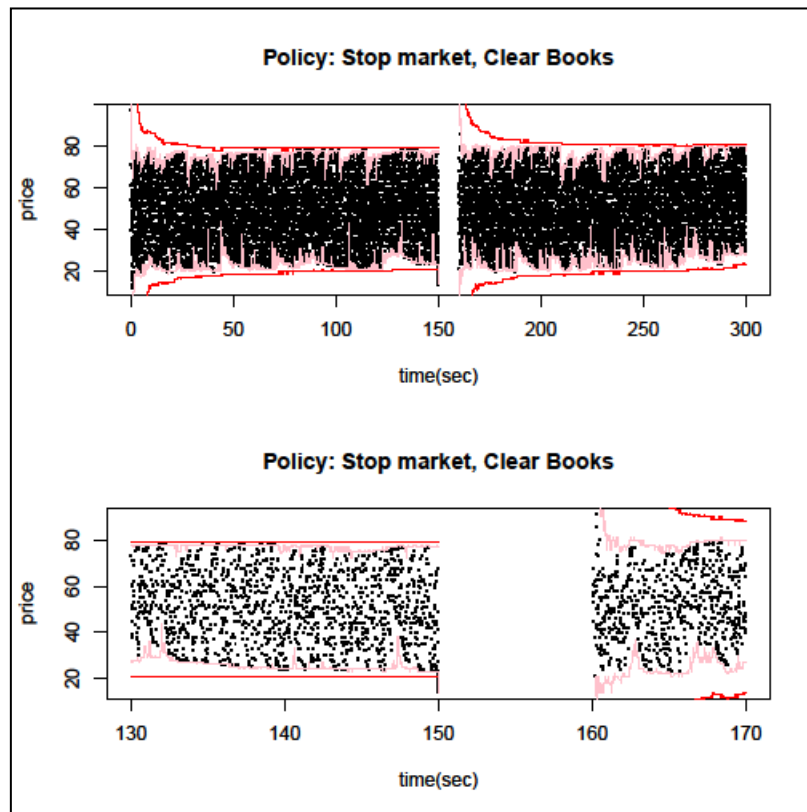FIGURE 17: **Waiting 10 sec with no orders or trades, then clearing books and resuming**

FIGURE 18. **Waiting 10 sec with no orders or trades, keeping books and resuming**

FIGURE 19: **No treatment action performed**