The Impact of High Frequency Market Makers Upon Market Liquidity: Evidence from Exchange Outages*

Matthew Harding†
Department of Economics
Stanford University

Paul Ma‡
Department of Economics
Stanford University

December 2010

Abstract

We identify the presence of high frequency arbitrageurs in the US treasury market through intraday exchange outages. Complementing our identification strategy, we find that order cancelation behavior also changed during the outage, consistent with arbitrageurs' profit maximization motives. Our estimates suggest that arbitrageurs represent approximately 69 to 94% of the quote depth in the spot treasury market. In addition, their presence seems to have large effects for the bid-ask spread of the 30-year treasury bond, which is the most illiquid product within its class. The paper provides more clarity for the conditions under which high frequency trading provides liquidity.

JEL: C33, G00, G12
Keywords: liquidity, treasury, bonds, natural experiment, arbitrage, outages, hedge funds, high frequency trading, order submission strategy, algorithmic trading

*We thank Douglas Bernheim, Peter Hansen, Jerry Hausman, Yesol Hoh, Seema Jayachandran, Alejandro Molnar, John Taylor, seminar participants from the St.Louis Fed Applied Time Series Workshop, and the Shanghai Econometric Society World Congress for their helpful comments and suggestions. Harding acknowledges support from the Presidential Fund for Innovation in International Studies and the UPS Fellowship. Ma acknowledges support from the Carl M. and Carolyn Franklin Fellowship.

†Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305; Phone: (650) 723-4116; Fax: (650) 725-5702; Email: mch@stanford.edu

‡Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305; Phone: (650) 644-8827; Fax: (650) 725-5702; Email: paulma@stanford.edu

There has been a recent surge of interest in understanding how the trading activities of sophisticated market participants affect the health of financial markets. The collapse of LTCM in 1998 convincingly illustrated both the limits of arbitrage (Shleifer and Vishny, 1997) as well as the issue of contagion amongst hedge funds trading similar strategies (Chan, Getmansky, Haas, Lo, and Center, 2005). More recently, on May 6th, 2010 the US equity market experienced a large, sudden, and unexpected price drop which has since been called the "flash crash". While the causes of the crash are still under review, preliminary reports from the CFTC and the SEC have called for further research into understanding the activities of high frequency traders.[1]. In this article, we shed light on the relationship between high frequency arbitrageurs and market liquidity via a natural experiment in the US treasury market.

We exploit the sharp timing of intraday outages- periods in which an unforseen technical glitch suspends operations in one exchange to proxy for the operational presence of a particular set of trading strategies, namely statistical arbitrage within a high frequency setting. The premise is that certain arbitrage strategies require simultaneous access to multiple securities which happen to be exclusively traded on different exchanges. When one of these exchanges unexpectedly shuts down from an outage, arbitrageurs are temporarily prevented from executing their strategies and halt their trading. Because these trading strategies were effectively arbitraging between the two assets and therefore providing liquidity in one asset conditional on the availability of liquidity in the other asset, the cessation of the strategy leads to a natural decline in the availability of liquidity in products traded within the unaffected exchange. We use standard difference in difference techniques to measure the change in liquidity as a consequence of the intraday outage.

Crucial to the interpretation of the findings is our assumption that the outage can

[1]Preliminary Findings Regarding the Market Events of May 6, 2010. Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.

be used as a proxy for the absence of these types of arbitrageurs. To further support that claim, we also provide evidence of strategic order submission and cancelation behavior which is complementary to arbitrageurs' profit maximization since it increases their probability of execution conditional on an arbitrage opportunity. We apply a method for detecting such behavior and show that it is only present during the non-outage period.

The market we focus on is the entire US fixed income treasury market, one of the largest and most liquid markets in the world. The fact that there are no designated market makers within these markets[2] suggests that the presence of arbitrageurs may have a significant effect upon liquidity. Figures 1 and 2 demonstrate graphically the sudden withdrawal of quote depth from the market at the precise moment of the outage while the point estimates indicate that between 67 to 94% of the quote depth of the spot US treasury market was sustained by these arbitrageurs. In addition, we find that their absence during the outage triggers a 18% increase in the bid ask spread of the 30 year treasury bond, one of the most illiquid assets amongst the US treasuries, amounting to a difference of 132 millions in bid ask cost fees. Robustness checks show that every other measurement of liquidity was affected by the outage, from lowered trading volume, lowered trading frequency, and increased duration between trades.

Our paper contributes to the literature in several ways. It complements recent work by Hendershott, Jones, and Menkveld (????) who show that a switch in market technology that facilitated algorithmic trading at the NYSE led to liquidity improvements for large cap stocks. We pinpoint statistical arbitrage, a strategy which can be implemented through algorithmic trading as one of the mechanisms behind the supply of liquidity. The paper also contributes to the growing evidence which links the presence of arbitrageurs

[2]Specifically the US spot treasury market operated by Cantor Fitzgerald and Brokertec, the only two exchanges in the world which trade US spot treasury notes, and not other dealer-only markets such as Govpx.

to liquidity provision (Choi, Getmansky, and Tookes, 2009), but focuses specifically on high frequency trading, a matter of increasing concern to regulators. Finally we provide evidence of a new mechanism governing the order submission and cancelation behavior of arbitrageurs which may explain some of the empirical findings in limit order order dynamics in Beber and Caglio (2005) and Ellul, Holden, Jain, and Jennings (2007).

The rest of the paper is structured as follows, the following section describes the institutional details and background concerning the trading strategies and arbitrageurs, section 2 explains the empirical setup and our identification strategy using the outages as well as our evidence regarding the nature of the strategic order process, section 3 summarizes the results on the impact upon liquidity, section 4 discusses alternative mechanisms and section 5 concludes.

I. Background

A. *Statistical Arbitrage*

Arbitrage, in a broad sense, is one of the most common types of strategies hedge funds use to generate their returns. For example in 2000 of the roughly $137 billion assets in hedge funds, over 90% of assets were classified as being used within a relative value or arbitrage strategy portfolio (Gatev, Goetzmann, and Rouwenhorst, 2006). Statistical arbitrage is just a specific example where quantitative models examine relative values within a basket of related securities for possible trading edges. Borrowing from the literature on co-integration, statistical arbitrage takes advantage of the possibility that a linear combination of the prices of two (or more) correlated non-stationary assets would be stationary along some frame of time. If such a relationship is found, then temporary deviations from the stationary synthetic asset are trading signals because these deviations are likely to mean revert. The source of these arbitrageur profits comes from satisfying idiosyncratic demand from uninformed noise traders and it is in a sense, a more sophisticated form of market making. These idiosyncratic demands will temporarily push prices of one or several of the assets away from the equilibrium of the synthetic combination. However since the noise traders were not trading on any fundamental information, the prices should over time revert to the equilibrium(Long, Shleifer, Summers, and Waldmann, 1990). It of course is not without the typical risks of arbitrage, as deviations can persist longer and become larger forcing the arbitrageur's hand to untimingly liquidate. It is also important to note that the capital costs to construct the necessary software and network infrastructure in order to capitalize on these arbitrage opportunities are not insubstantial. Therefore the returns can also be thought of as compensation for the opportunity cost of arbitrage (Grossman and Stiglitz, 1980).

Statistical arbitrage and its more common derivations such as pair trading (which involves only two assets) have been employed in the financial industry as far back as 1985 across various asset classes such as fixed income, currencies, and equities (Derman, 2007). Recent academic work using daily data from 1962-2002 found simple pair trading strategies on US equities to generate excess annual returns of 11% with half of the risk of the S&P 500 (Gatev, Goetzmann, and Rouwenhorst, 2006). They concluded that the source of the excess returns is compensation for the arbitrageur's efforts in enforcing the law of one price. In fixed income markets, where by 2005 there were over $56 billion assets specifically for fixed income arbitrage[3] Duarte, Longstaff, and Yu (2005) found yield curved based strategies(which are essentially pair trading) produced significant alpha over a six year period. It is not surprising that relative value strategies work well in fixed

[3]This comes from Tremont/TASS Asset Flow Reports which is an underestimate because it does not include the proprietary trading arms of major investment banks.

income presumably because cash flows from bonds are easier to calculate than equities.

B. *The US Treasury Market*

The US treasury market consist of both spot and future treasury notes which are traded in public electronic exchanges and cleared in central clearinghouses. Espeed is an electronic exchange that specializes in trading on the run US Treasury notes [4]. It, alongside with rival exchange Brokertec in a 40 to 60 ratio are the two principal electronic exchanges for trading spot treasury securities in the secondary market, supplying roughly 40%, with the remainder being inter-dealer markets (Fleming and Mizrach, 2005).

While spot treasuries can be traded on either Espeed or Brokertec, in order to trade treasury futures, one must trade on the Chicago Board of Trade (CBOT) platform[5], which has had a history of "glitches" resulting in outages, the earliest of which dated back to 1995. In addition to treasury futures, the CBOT also offered agriculture futures such as soybeans, wheat, oat, ethanol, corn, interest rate swaps, equity futures, and precious metals futures. In October 2006 the Chicago Mercentile Exchange (CME), another futures exchange announced that it was merging with CBOT but it was not until 2008 that the CBOT's trading software migrated to CME's trading system. In between 2002 to 2007, there were 10 media reported outages[6] at CBOT which halted electronic trading on average for 90 minutes during the middle of a trading session.

[4] The most common maturities are the 2 year, 3 year, 5year, 10 year, and 30 year notes

[5] ELX, a competing exchange trading treasury futures was introduced in 2009 in response to CME's merger with CBOT.

[6] Total of 8 trading sessions, 2 trading sessions had 2 outages within the same session

C. *Arbitrage in Fixed Income*

Statistical arbitrage in fixed income seeks to buy underpriced assets and sell overpriced assets in a systematic fashion using a combination of spot and future securities. The most basic example would be a pair trading where one simultaneously purchases a spot treasury note and sells the corresponding future. More sophisticated trading combinations exist as well. The crucial context is that in order for arbitrageurs to perform arbitrage between the spot and future markets, they must have access to both electronic exchanges, since they are traded on different exchanges.

A simple pair trading strategy entails placing limit orders, which is an offer of liquidity, on both securities and have them executed only when certain market conditions are met, for example the co-integration linear combination deviates sufficiently from some known mean. In the scenario where an outage knocks out one of the exchanges, pair trading stops mechanically because arbitrageurs no longer have the ability to arbitrage. Furthermore, arbitrageurs must cancel their previous limit orders in the unaffected exchange, once again, because those orders were only there for the purpose of fulfilling a potential arbitrage opportunity, which is no longer possible. Consequently liquidity supplied from arbitrageurs in the unaffected exchange must decrease.

To give credence to the co-integration relationship between the spot and future treasury notes, figure 5 plots the price series for both the spot and future 10 year treasury note[7] on 4 different days where outages occurred. The two series are extremely correlated, and an augmented Dickey-Fuller test of the difference in the two price series rejects the null of a unit root.

[7] Note that during the outage, we still have pricing data from the open outcry market, the line in red, and we see that prices are still very correlated even during the outage, a sign that the markets were still operating relatively efficiently despite the loss of access to the electronic market.

II. Data and Empirical Strategy

A. *Data*

We focus our study in the time period between 2006 and 2007. It is not a particularly selective sample period because after 2007 the CBOT merged with CME, and switched software platforms, which was the common fault for the crashes and since there have been no reported outages to our knowledge. While there were outages prior to 2006 at CBOT, for example one occurred in March 2005, Espeed did not begin to collect detailed limit order book level data until 2006, which explains our starting date. We have available to us message level data regarding the state of the limit order book from Espeed between 2006 and 2007 for all on the run treasury maturities. While Espeed treasury data has not been used commonly in the past, it is comparable to Brokertec data used by other researchers[8]. To complement our spot treasury data, we also have high frequency trade and quote level data from the Chicago Board of Trade on the corresponding US treasury futures between 2006 and 2007. Finally, we obtain records of outages directly from searching historical media sources; Table 1 contains summary statistics of the reported outages, of the six reported outages in the study period, the average outage lasted 95 minutes.

B. *Liquidity Measurements*

There is no universal variable that captures every dimensional quality of liquidity. Some of the most common candidates in the US treasury literature include the bid ask spread, the quote size, trade size, and the frequency of trading (Amihud and Mendelson, 1991) .

[8]See Fleming and Mizrach (2005) as a comparable example of the Brokertec data.

The bid ask spread[9] which is the difference between the best ask price and the best bid price is the most common measurement (Fleming and Sarkar, 1999). It does not however fully capture the essence of liquidity because it ignores information regarding the quantity available for trade at each price- the quote size, which proxies for how "deep"[10] the market is. Other common proxies for liquidity include the trade size and the duration between trades. Finally, we also examine the unique number of trading accounts with orders at each price point. This serves as a proxy for the number of market participants for each particular security[11]. Table 2b presents high frequency summary statistics of all the liquidity measurements used in this paper for the entire sample period.

As a preview of the effects of the outage, Figure 1 plots the total quote depth [12] over time for six different outage days for the 10 year spot treasury note. The two dashed vertical bars denote the beginning and the end of the reported outage. There is a sharp discontinuous drop in the quote depth at the exact time of the outage in the spot market. We also see that once the exchange recovers from the outage, the quote depth steadily increases back up, but less than what it was originally possibly due to intraday seasonality and users queueing to log back into the trading system platform at the Chicago Board of Trade. Figure 2 plots the total number of market participants in the market in the same ten year spot treasury note and the effect of the outage graphically is nearly identical, suggesting that it wasn't just one large market participant who left, but rather a large population of smaller traders.

Table 2a reports summary statistics of key liquidity measurements measured at the one minute level for the 2, 5, 10, and 30 year treasury note before, during, and after

[9]In the estimation procedure, the bid ask spread is first normalized in terms of minimum tick increments.

[10]The electronic limit order book also does not display "hidden liquidity" which are comprised of possible orders which are not present for strategic reasons.

[11]This is available retrospectively in the historical data in Espeed.

[12]sum of quote size of best five ask prices and best five bid prices

an outage. The before and after periods were defined as a 90 minute window, the same as the average duration of an outage. As an example, the ten year treasury note had a average volume of 64 per minute just prior to the outage, which decreased to 20 during the outage, and came back up to 28 after the outage. Quote depth decreased from 1056 to 322 during the outage, and increased back to 724 after the outage. The patterns for duration between trade, trading frequency, number of market participants, and average order size are similar. The exception seems to be the bid ask spread[13] which did not seem to significantly change in the 2, 5, or 10 year notes, but did increase by 44% in the 30 year treasury note from the outage, from an average of 1.69 ticks to 2.44 ticks, and subsequently decreasing to 1.74 after the outage. To more formally quantify the effect of the outage upon liquidity in the spot market we estimate separately for each maturity note the following empirical specification:

$$Y_{t,d} = \beta_0 + \beta_1 O_{t,d} + \lambda_t + \gamma_d + \epsilon_{t,d} \tag{1}$$

where Y is the liquidity measurement outcome of interest at time t for day d. O is an indicator which equals 1 during an outage and 0 otherwise. λ_t are hourly fixed effects to control for intraday seasonality,γ_d are monthly fixed effects, and $\epsilon_{t,d}$ are random disturbances. Time t is measured at the minute level, so each the liquidity outcome variables are measured as end of each particular minute from the raw tick level quote and trade data. The parameter of interest is β_1 and the identification comes from our argument that the outages are a consequence of an exogenous process unrelated to the error term.

[13]The bid ask spread is normalized in terms of ticks for easy comparison across different maturities.

C. *Identification Strategy*

In order to estimate the effect of the arbitrageurs (practicing the form of arbitrage as described) upon market liquidity, one would ideally like to have a random on/off switch for the presence of the arbitrageurs. One could in principal use differential observance of holidays by different exchanges as a candidate, but that becomes subject to the criticism of pre-anticipation by market participants and to the criticism of local average effects since holiday days are generally different from non-holiday days. Using the outages at CBOT is a more favorable option because it is both unanticipated, high frequency in the sense that one could observe the shock intraday, and finally conditionally random as we argue in a later section, and therefore orthogonal to the actions of market participants and market activity. The identification strategy is most similar in spirit to a combination of Hendershott and Jones (2005) and Chaboud, Chernenko, and Wright (2008).

Figure 3 plots intraday volume of all the CBOT treasury futures using original transaction level data from the Chicago Board of Trade for a particular outage day. In each case one observes a volume gap in trading during the reported outage period and serves as a confirmation that an outage actually occurred [14]. For each reported outage, we researched media sources to discover if any reasons were given for the failure, and in all cases they were either due to "glitches" in software or unknown. While other exchanges also have software failures of various natures from time to time, for example the London Stock Exchange went offline briefly in 2008, and the US flash crash in May 2010 may have been related to a delay in the NYSE's quote system[15], the Chicago Board of Trade seemed especially prone, with a history going back as far as 1995. In fact one of the arguments the CME management team made to shareholders during the merger talks

[14] We confirmed these gaps for a representative sample of the other products traded in the CBOT e.g. agricultures, metals, and currencies

[15] `http://www.nanex.net/20100506/FlashCrashAnalysis_CompleteText.html`

was that they would be moving to a more stable exchange software platform.

Detailed information and summary statistics concerning the timing of the outages are listed in Table 1. The average duration of an outage is roughly 95 minutes and they generally happen in the morning period[16] When an outage actually occurs, the network connection between traders and the exchange immediately severs and all outstanding orders are canceled[17]. As the CBOT diagnoses the problem they may give advance notice to its trading members on when it expects to be back online. We as the econometrician however do not observe this information. After the problem is solved the exchange reopens with a reconnection queue to prevent overloading. Therefore market participants return more slowly than they exited into the exchange's system.

Note that in an outage all electronic securities offered at the CBOT simultaneously became unavailable. The CBOT also operates a floor exchange trading the same products via open outcry. The floor operation was not interrupted during the electronic outage. Therefore the resulting drop in liquidity cannot be due to the lack of a futures market per se, but instead more specifically the lack of an electronic futures market. In 2007 electronic volume of treasury futures constituted 94% of the total volume. More specifically, we are only identifying the effect of the entire product space of the CBOT being unavailable, and not specifically treasury futures. Besides the interest rate futures, CBOT also trades agricultural, equities, and precious metals futures. However these other products have statistically low correlation with interest rates in general and do not have the same level of economic relationship that the spot and future treasuries share, and would therefore unlikely be part of an arbitrage strategy involving the spot treasury products. In

[16]The electronic trading hours for the treasury futures is 22 hours a day from 6:00 pm to 4:00 pm central time.

[17]If by chance an order was completed just prior to the outage, the trade confirmation message may be lost. Traders would then typically have to contact the exchange in order to verify their current position from the exchange's internal system.

addition, because all the future interest rate products at the CBOT became unavailable simultaneously, we are only able to estimate the combined effect of their absence upon the liquidity of each individual spot treasury note.

D. *Are Outages Exogenous?*

The central identification assumption is that the outages are exogenously generated. The nature of the outages from all media accounts seem to point to fundamental deficiencies in the underlying trading platform of the CBOT. This suggests that during particularly high periods of message traffic, the exchange simply overloads and shuts down. The average volume traded for the 10 year treasury future, the most commonly traded interest rate future, just prior to each of the outages was 834,000. In comparison in 2006 the average volume traded for an entire session was 810,000 contracts. The 834,000 number is significantly higher than the average volume of 530,000 traded in the same pre crash period but using historical data from every trading session in 2006. Mapping the volume traded prior to the crash against the unconditional volume distribution shows that the average crash day was in the top 90% of trading days in terms of volume activity. If we use price volatility prior to the crash instead and compare it against the historical volatility from 2006 we find that crash days were in the top 79% of trading days in terms of volatility.

The suggestive evidence indicate that the probability of crashing increases with volume and volatility, hence it may be important to use similarly high volume or volatility non-crash days as a control in the regression exercise. Restricting the sample to trading days which are in the top 10% in volume relative to the historical distribution from 2006 to

2007 [18], we find that the probability of an outage is 12%[19].

E. *Strategic Order Placement and Order Cancelation*

Are outages good proxies for the absence of arbitrageurs? While the mechanisms of arbitrage seem to suggest that arbitrageurs and the arbitrageurs' supply of liquidity would be impacted, there is no prima facie evidence of them actually exiting during the outage. Instead, one can try to identify residual signatures left in the market due to arbitrage activity, and then show that these signatures disappeared during the outage. This would further corroborate the story that it was the arbitrageurs who actually exited. One such signature to be described in detail below is the way in which limit orders are submitted and canceled as there is a strategic aspect which confers an advantage to high speed arbitrageurs. Therefore if our hypothesis regarding the arbitrageurs exiting during the outage is correct, then this particular form of behavior also ought to disappear during the outage.

Given the increasing sophistication of the strategies and the overcrowding of hedge funds into the statistical arbitrage field, profits and losses are often determined by who has the fastest order execution time and who is able to minimize execution risk. Hedge funds are increasingly devoting more capital to more powerful hardware, more efficient coding, or even locating servers closer to the exchange to reduce network congestion[20]. Once again, given these high costs of capital, the arbitrageurs' returns are quite consistent with Grossman and Stiglitz (1980).

As a result of the overcrowding, it is not just a matter of identifying a profitable

[18] Historical volume as measured by the 10 year treasury future at the CBOT during 2006-2007.

[19] Between 2006 and 2007 there were a total of 6 outages and 50 high volume/high volatility days.

[20] There is a recent debate over the ethnics of high frequency due to an arms race in technology: hedge funds are now competing to put their servers closer to exchanges in order to reduce their execution time.

arbitrage, but also having the fastest execution software relative to everyone else. In order to increase the likelihood of a successful execution conditional on an profitable trading opportunity, arbitrageurs can gain an advantage over one another if they somehow happen to be closer to the front of the queue in the limit order book. This is advantageous in exchanges where the order matching engine is based on price time priority, meaning orders closest to the incoming trading price are matched first, and orders within the same price are matched based on which order arrived first. Therefore arbitrageurs in the front of the queue tend to have more frequent opportunities to actually participate in an arbitrage[21]. In order to actually be in closer to the front of the queue, one can design a complementary program which continuously submits orders into all possible pricing points, and selectively canceling them depending on market conditions so that one does not trade by accident. By sending these orders in much earlier than the actual trading opportunity, the arbitrageur can improve his position in the queue so that when an actual opportunity does arise, he is more likely to be in a position to execute it. The phenomenon of submitting orders early in order to get an execution advantage is not new, Biais, Hillion, and Spatt (1995) found such behavior in the limit order book in the Paris Bourse.

The key reason for the existence of an order's option value is because there is no direct cost to sending an order to the exchange or canceling an existing order [22]. The option value creates an incentive for arbitrageurs to blanket a security's limit order book with orders at a variety of price points. However if markets move too quickly, there is the risk that an undesirable order would not be canceled in time, leaving the arbitrageur with an unwanted position which he may have to liquidate at a loss. One can empirically test

[21] This is applicable to arbitrage opportunities where the bid and ask prices do not cross, which are the majority of the cases now due to the nature of high frequency trading.

[22] Some exchanges do have limits on the ratio between "messages" and volume traded to prevent excessive network congestion without any actual trades but we do not know if it binds.

whether a specific form of order placement behavior changed during the outage period relative to normal non-outage periods by using information from the state of the limit order book. If we find that order cancelation behavior did change during the outage and not before or after, then it would be suggestive that it was indeed arbitrageurs who actually withdrew their liquidity from the limit order book, since they were benefitting the most from the option value of the orders.

While the concept of maintaining a favorable queue position is intuitive there could be multiple ways of actually implementing such a strategy, each of which could affect the actual order placement and order cancelation behavior differently. We focus on one of the crucial components of the strategy- the timing of order cancelations, because if implemented incorrectly, it could increase the arbitrageur's exposure to unwanted orders being fulfilled rather than canceled. The decision to cancel an existing order depends on the probability of it being fulfilled while an arbitrage opportunity does not exist. The higher the probability, the more likely one should cancel that particular order. A simple heuristic rule could be that as the trading price approaches the price of the submitted order the order should be more likely to be canceled in the absence of an available arbitrage opportunity.

As an example, denote $P^b_{t,j}$ and $Q^b_{t,j}$ as the jth best bid price and the jth best depth at time t. Now suppose an active sell order arrives and executes against $P^b_{t,1}$ at time t. This signals to the arbitrageur that the risk of the price going down has increased which may lead her to cancel or decrease her order sizes at the second best bid price $P^b_{t,2}$ or lower.

A crude measurement of order cancelation behavior would be to compute the 1 second change in the depth of the second best bid depth in response to an incoming sell order at the best bid price, and symmetrically the 1 second change in the depth of the best ask price depth in response to an incoming buy order at the best ask price.[23]. If programs are indeed behaving this way, then we should expect the second best bid and ask depth to decrease as orders are canceled there in response to an incoming sell or buy order[24]. Our hypothesis however is that this behavior only occurred outside of the outage periods because this option value is only positive in the presence of available arbitrage opportunity which the outage took away. Therefore, we estimate the following empirical specification:

$$Log(Q_{it}) = \beta_0 + \beta_1 * Order_{it} + \alpha_i + \epsilon_{it} \tag{2}$$

where Q_{it} represents either the second best bid or ask depth for treasury note i at time t, and $Order_{it}$ is either an active buy order or active sell order, and α_i represents the fixed effects of each note. Following our hypothesis, we expect that the coefficient β_1 should be negative during non-outage periods and closer to zero during the outage period itself. A negative β_1 signifies order cancelations which causes the depth to decrease. The estimate of the equation from Table 9 shows that an active buy order triggers a -3.7% drop in the market depth offered at the second best ask price[25] just prior to the outage period. However the estimated coefficient drops to -1.1% and is insignificant during the outage period, and increases back to -1.6% after the outage, although this is estimated more noisily. Correspondingly, we see an decrease of -3.5% in the second best bid depth after an active sell order of roughly the same magnitude just prior to the outage. Again, the estimated coefficient drops to -.005% and becomes insignificant during the outage period, and increases back to -2.8% just after the outage.

[23]We chose 1 second to illustrate that only programs would be capable of this behavior. The results do not change if we use sub-second increments

[24]We may also expect similar behavior at bid or ask depth farther away but at a decreasing rate since the risk of execution is lower farther away from the current trading price. We have additional regressions which show that order cancelation also happened at order prices farther away from the second best price and at a decreasing rate.

[25]Holding the best ask price fixed because we only examine scenarios where the buy order doesn't completely deplete the depth at the best ask price.

While Table 9 shows the estimates from a comparison of before and after the outages, there may be intraday seasonalities confounding the timing of the outages. Results from Table 10, which includes hourly and daily fixed effects of a control group of all days in 2006 and 2007, show a precisely estimated 2.5% decrease in the depth as a result of an active trade during non-outage periods and during the outage, insignificant estimate of .5 to 1%. In conjunction, there appears to be some evidence that the outage caused the exit of a class of market participants who used strategic order placement and cancelation strategies, which a statistical arbitrageur would use due to the benefits of improved order execution probability beneficial to arbitrage.

III. Empirical Results

A. *Impact Upon Liquidity Measurements*

Figure 4 plots the average price impact in ticks[26] for a market order of an arbitrary size one minute before and one minute after an outage for the spot market ten year treasury note. It essentially represents the average cost curve for a particular market transaction size. As order size increases, it becomes increasingly costly to transact due to the price impact of the order having to trade through the stacks to reach the quantity desired. The upward shift in the cost curves after the outage represents a depletion in the quote depth due to the outage, which magnifies the price impact of each order size. Reading the figure, just prior to the outage, an order size of 10 units would have cost .05 ticks per unit in excess of the best bid/ask price, and this number jumps to .225 ticks per unit just one minute after the outage, a four fold increase. This is another way of visualizing the effect of the outage using combined information from the bid ask spread as well as

[26]A tick is the minimum price increment and is $156.25 for the ten year treasury note.

the quote depth of the five best bid and ask prices. One should not accept the four fold increase as literal because it does not adjust for hidden liquidity(Harris, 1997), but it is suggestive of the immediacy and the magnitude of the effect the sudden outage had. Figure 5 asks the reverse question of what happens after the exchange recovers from the outage. Just 1 minute prior to the recovery, the average price impact of an order size of 10 units would have caused a .1 tick cost per unit above and beyond the best bid/ask price and goes to zero per unit after the recovery, reflecting the fact that there were at least ten units in depth at the best bid and ask just after the recovery.

Table 3 gives estimates of our baseline empirical specification in equation 1 of the impact of the outage upon the log market depth. Without any controls, the effect of the outage on market depth depending on the maturity year ranges from -39% to -68% in column 1 and highly significantly. When hour and day fixed effects are included to remove intra and interday seasonality, the estimates become larger and range from -83% in the 2 year note to -99.6% in the 30 year note in column 2. To check if outages days were relatively different than non-outage days and therefore cause us to estimate local averages instead of the true average because of the higher likelihood of crashes on high volume or high volatility days, columns 4 and 5 restrict the control sample to only the top 10% of days relative to either historical volume or historical relatively respectively. The estimates from column 4 and 5 qualitatively similar, in general restricting via volatility results in slightly smaller estimates, but both columns are significant. The values from columns 2 and 3 are similar, which seems to rule out the fact that we are estimating a local average, rather than the true average treatment effect. Economically the effects are large and the decreased quote size correspond to price impact graphs from Figure 4 and 5. From column 3, it seems that approximately 67 to 94% of the quote depth in the US treasury market came from arbitrageurs who withdrew their limit orders during the outage period.

While quote depth is vital for large order execution, one of the traditional measures of liquidity has been the bid ask spread. Table 4 shows estimates of the impact of the outage upon the log bid ask spread for each bond maturity year. Note that the average bid ask spread is normalized to the security's minimum tick size and ranges from 1.21 in the 2 year treasury note to 2.74 in the 30 year treasury bond. After including additional time fixed effects, column 3 shows that the outage had small insignificant effects for all notes except the the 30 year bond where it increased the bid ask spread by 30%. In column 4 and 5, when we restrict the sample so that the control group comprise of high volume or high volatility days, the results are largely similar. The largest effect seems to be on the 30 year treasury note, which suffered a 13% to 18% increase in the bid ask spread as a result of the outage. It is important to note that bonds with longer maturities are generally more illiquid, for example the summary statistics from Table 2b show that the 30 year bond has largest average bid ask spread amongst all the bonds. The large effects for the 30 year bond can also be attributed to Espeed having a larger share of the 30 year treasury bond market relative to Brokertec. Hence it seems that the arbitrageurs are providing liquidity where it is most needed, but the subtle caveat is that it is only conditional on the availability of the corresponding futures market for arbitrage purposes.

To complete our description of the outage, Table 5 reports similar regressions but with volume traded as the outcome variable measured in levels rather than logs due to a large fraction of zero volume observations. In the main specification of interest, column 3, the outage decreased trading volume ranging from 6 units in the 30 year(μ=18.6) to 47 units in the 2 year(μ=27). We find a different story with order size in Table 6, where the outage seemed to decrease the order size for the 2 year note and 30 year bonds, but increase the order size for the 5 and 10 year notes.

Column 3 of Table 7 reveals that order frequency fell significantly for all four bonds, ranging from one third of a standard deviation for the 30 year bond to over 1 standard deviation for the 5 year note. Finally Table 8 show that the time between trades increased significantly during the outage as traders likely timed their trading to reduce market impact. From Column 3 in Table 8 duration increased between 17 to 37 seconds per minute on average.

It is important to recall that even though the electronic exchange was unavailable at CBOT, traders still had the option of telephoning orders to be filled at CBOT's open outcry market, which was unaffected. Unfortunately intraday floor market volume is to our knowledge unavailable so we would not be able to estimate how much substitution in trading occurred from the electronic to the floor market. However, we do observe the total daily floor volume which for the ten year future increased by 75% during outage days relative to days immediately before and afterĖlectronic traders substituting towards floor trading may be one reason why liquidity in the spot market was not even worse during the outage, but full substitution was unlikely to be possible due to the technical nature of high speed arbitrage, which only the electronic market could provide.

B. *Impact Upon Volatility*

We provide a simple stylized model in the appendix which describes the possible effects upon volatility from a withdrawal of arbitrage traders from one of the two markets. To provide the basic intuition, there are two competing forces affecting the volatility of the spot market. On the one hand, the outage prevents electronic traders from partaking in arbitrage activities hence price deviations from fundamental value should be more likely and volatility may increase. On the other hand, the outage also increased the costs of trading temporarily. Since market participants anticipate the outage to be short term, it may be rational for noise traders to temporarily delay their trading until after the outage.

The subsequent reduction in noise trading would decrease volatility, since noise traders were the source of volatility by model assumptions.

Summary statistics from Table 2 provide annualized volatility for the 2, 5, 10, and 30 year treasury notes for all three periods: before, during and after the outage. Every treasury note decreased in volatility during the outage compared to the pre and post periods. For example the 5 year treasury note's volatility decreased by 33% between the crash and non crash periods[27]. While these results are certainly descriptively interesting , it would be difficult to conclude that high frequency arbitrageurs increase volatility in general or that illiquidity leads to lower volatility(Gillemot, Farmer, and Lillo, 2006). Because the duration of the treatment was so short, we hesitate making a general policy statement regarding the effects of arbitrageurs because of more ambiguous general equilibrium effects as arbitrageurs seek alternative trading strategies.

IV. Discussion

While calculation of welfare gains/losses is an exercise beyond the scope of the current paper, one can give a basic back of the envelope estimate of the welfare transfers from liquidity seekers to liquidity providers. Focusing on the 30 year treasury bond, the average bid ask spread in 2006 was 2.74 ticks where each tick equaled $152. The average volume in the same period was 17,000 bonds per day. Assuming a 50% drop in volume from estimates in Table 5[28], the lower bound on the annual increased cost of trading for the 30 year bond alone was roughly $132 million[29]. Note that this figure would likely increase

[27]This does not take into account intra and interday seasonality.

[28]This is an upper bound on the effect upon volume due to the bid ask spread widening since traders are able to delay their trades only because it was a temporary rather than a permanent change.

[29]The actual calculation is 17,000(daily volume) *.5(reduction in volume due to outage)*250(number of trading days per year)*2.74(average bid ask spread)*.15(estimated increase in bid ask spread, the average of column 3 and 4 in Table 4)*152(the nominal dollars per tick). We are also ignoring hidden

with order size because the calculations are ignoring the outage's negative effects upon quote size.

Outages in the futures market for treasuries may affect liquidities in other markets and through those markets in turn impact the liquidity in the spot market. To the extent that this happens, the outages may not fully serve as a proxy for the operational presence of the arbitrageurs. While we cannot rule out this case, first order effects seem to suggest the presence of the arbitrageurs being responsible, and the evidence from lack of strategic order cancelation only during the outage period seem to support the argument of the exodus of arbitrageurs.

We might also be concerned if exchange outages were somehow correlated. For example the day the London Stock Exchange had an outage, Intercontinental Exchange (ICE) based in the US also had an outage[30]. In such a case, the outage at the CBOT could trigger traders to respond in anticipation of other exchanges also becoming unavailable, and not through the direct absence of the treasury futures. A thorough search of outages in other exchanges revealed no other problems during the same period as the CBOT outages. All evidence seems to point to unique problems within the CBOT's trading system.

V. Non-technical Summary

Given the importance that arbitrage plays in the analysis of security markets, it is surprising that there are relatively few empirical studies of market effects from arbitrageurs and especially those of a high frequency nature. Our paper uses random outages in the

liquidity considerations which could cause us to upward bias the welfare loss

[30]The media reported that the outages were not related.

Chicago Board of Trade's trading platform as a proxy for the operational presence of high frequency arbitrageurs between the spot and future US treasury market. Under a difference-in-difference framework with the outage as the treatment, we find that these arbitrageurs provided up to 67-94% of the quote depth of the US Treasury notes. Their disappearance during the outage also increased the bid ask spread for the longest maturity 30 year bonds by 14-25%. To rule out alternative non-arbitrage mechanisms at play, we also detect strategic order cancelation behavior complementary to arbitrage trading because it increases order execution probabilities and show that it is only present during the non-outage periods. To understand the effects of the outage upon market volatility, we present a simple stylized model which incorporate the effects of arbitrageurs and noise traders who can intertemporally substitute their trading needs when faced with a known temporary shock to liquidity. We find the latter effect to be predominant, as volatility actually declines during the outage period. It must be interpreted with caution however, as the effects we estimate are short term. Finally, it is quite ironic that despite the mass computerization in the financial sector simple outages at the exchange level still occur. Fortunately they provide an excellent source of intraday variation for analyzing the effects of multi-market trading strategies; one direction for future research could be directed at exploiting outages in other exchanges trading other asset classes.

References

Amihud, Y., and H. Mendelson (1991): "Liquidity, Maturity, and the Yields on U.S. Treasury Securities," *The Journal of Finance*, 46(4), 1411–1425.

Beber, A., and C. Caglio (2005): "Order Submission Strategies and Information: Empirical Evidence from the NYSE," *SSRN eLibrary*.

Biais, B., P. Hillion, and C. Spatt (1995): "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *The Journal of Finance*, 50(5), 1655–1689.

Chaboud, A. P., S. V. Chernenko, and J. H. Wright (2008): "Trading activity and macroeconomic announcements in high-frequency exchange rate data," *Journal of the European Economic Association*, 6(2-3), 589–596.

Chan, N. T., M. Getmansky, S. M. Haas, A. W. Lo, and O. C. Center (2005): "Systemic risk and hedge funds," *NBER working paper*.

Choi, D., M. Getmansky, and H. Tookes (2009): "Convertible bond arbitrage, liquidity externalities, and stock prices," *Journal of Financial Economics*, 91(2), 227–251.

Derman, E. (2007): *My life as a quant: reflections on physics and finance*. Wiley.

Duarte, J., F. A. Longstaff, and F. Yu (2005): "Risk and Return in Fixed-Income Arbitrage: Nickels in Front of a Steamroller?," *Rev. Financ. Stud.*, 20(3), 769–811.

Ellul, A., C. W. Holden, P. Jain, and R. Jennings (2007): "Order dynamics: Recent evidence from the NYSE," *Journal of Empirical Finance*, 14(5), 636–661.

Fleming, M., and B. Mizrach (2005): "The microstructure of a US Treasury ECN: The BrokerTec platform," *Unpublished paper, Federal Reserve Bank of New York. Tzucker, R., and M. Islam.*

Fleming, M., and A. Sarkar (1999): "Liquidity in US Treasury spot and futures markets," *Market Liquidity Research Findings and Selected Policy Implications, Bank for International Settlements.*

Gatev, E., W. N. Goetzmann, and K. G. Rouwenhorst (2006): "Pairs trading: Performance of a relative-value arbitrage rule," *Review of Financial Studies*, 19(3), 797.

Gillemot, L., J. D. Farmer, and F. Lillo (2006): "There's more to volatility than volume," *Quantitative Finance*, 6(5), 371–384.

Grossman, S. J., and J. E. Stiglitz (1980): "On the impossibility of informationally efficient markets," *The American Economic Review*, 70(3), 393–408.

Harris, L. (1997): "Order exposure and parasitic traders," *Working Article, University of Southern California.*

Hendershott, T., and C. M. Jones (2005): "Island Goes Dark: Transparency, Fragmentation, and Regulation," *Rev. Financ. Stud.*, 18(3), 743–793.

Hendershott, T., C. M. Jones, and A. J. Menkveld (????): "Does Algorithmic Trading Improve Liquidity?," *Journal of Finance.*

Long, J. B. D., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990): "Noise trader risk in financial markets," *Journal of Political Economy*, 98(4).

Shleifer, A., and R. W. Vishny (1997): "The Limits of Arbitrage," *The Journal of Finance*, 52(1), 35–55.

VI. Appendix

To understand how the outage could potentially affect the volatility of the related assets, we present a simple stylized model which incorporates the effect of arbitrageurs and noise traders on prices. Suppose prices for assets A and B have the following time process:

$$P_t^A = P_{t-1}^A + \eta_t + \epsilon_{t,1} - \alpha_1[V_{t-1} - \bar{V}]$$
$$P_t^B = P_{t-1}^B + \eta_t + \epsilon_{t,2} - \alpha_2[V_{t-1} - \bar{V}]$$

Assets A and B follow a random walk with idiosyncratic serially independent disturbances $\epsilon_{t,1}$ and $\epsilon_{t,2}$ with a common disturbance η_t reflecting information pertinent to both assets. $\epsilon_{t,1}$ and $\epsilon_{t,2}$ are assumed to be independent from each other and η_t. This reflects the view that not all noise traders have access to both markets. For example, uninformed retail investors may observe the stock market, but they may be unaware of or do not have access to to the corresponding options market for the equities. The term $V_{t-1} \equiv P_{t-1}^A - P_{t-1}^B$ is the difference in price between A and B. The last term $V_{t-1} - \bar{V}$ denotes the current relative pricing relative to the historical relative pricing, and is a measurement of the degree of mispricing between the two assets. We can interpret α_1 and α_2 as mean reverting factors meant to proxy the action of arbitrageurs who work to equilibrate the relative prices. Note that arbitrage is not assumed to be costless, so the α's also embed the costs to arbitrage.

The volatility of asset A is given by:

$$\begin{aligned}
P_t^A &= P_{t-1}^A + \eta_t + \epsilon_{t,1} - \alpha_1[V_{t-1} - \bar{V}] \\
P_t^A - P_{t-1}^A &= R_t^A = \eta_t + \epsilon_{t,1} - \alpha_1[V_{t-1} - \bar{V}] \\
\sigma_{R^A}^2 &= \sigma_\eta^2 + \sigma_{\epsilon,1}^2 + \alpha_1^2\sigma_V^2
\end{aligned}$$

Asset B has the corresponding volatility:

$$\sigma_{R^B}^2 = \sigma_\eta^2 + \sigma_{\epsilon,2}^2 + \alpha_2^2\sigma_V^2$$

We can also derive the expression for the volatility of the difference in prices A and B, the "basis", σ_V^2 as:

$$\begin{aligned}
P_t^A - P_t^B &= P_{t-1}^A - P_{t-1}^B + \epsilon_{t,1} - \epsilon_{t-2} + (\alpha_2 - \alpha_1)[V_{t-1} - \bar{V}] \\
P_t^A - P_t^B &= V_{t-1} + \epsilon_{t,1} - \epsilon_{t-2} + (\alpha_2 - \alpha_1)[V_{t-1} - \bar{V}] \\
P_t^A - P_t^B &= [1 + \alpha_2 - \alpha_1]V_{t-1} + \epsilon_{t,1} - \epsilon_{t-2} + C \\
\sigma_V^2 &= [1 + \alpha_2 - \alpha_1]^2\sigma_V^2 + \sigma_{\epsilon,1}^2 + \sigma_{\epsilon,2}^2 \\
&= \frac{\sigma_{\epsilon,1}^2 + \sigma_{\epsilon,2}^2}{1 - [1 + \alpha_2 - \alpha_1]^2} \qquad (3)
\end{aligned}$$

Now, suppose the process V_t follows:

$$V_t = V_{t-1} + \epsilon_t - \alpha(V_{t-1} - \bar{V}) \qquad (4)$$

Assuming the innovations of ϵ are independent then from inspection, we see that V_t is just an AR(1) process with parameter $1-\alpha$. Therefore we know the autocovariance function of V_t is:

$$\gamma_j = \frac{(1-\alpha)^j \sigma_\epsilon^2}{1-(1-\alpha)^2} \tag{5}$$

Now we analyze the case where the exchange of one of the assets goes offline unexpectedly. Suppose the exchange for asset B is unavailable. What is the impact upon the volatility of asset A? We assume that α, the degree of reversion is lower during an outage due to the absence of arbitrageurs. Equation 3 then predicts that volatility should increase for the basis. Empirically this is difficult to test because we don't actually observe the price series for the basis. We do however have prices from the open outcry market, and thus we do have a rough measurement of σ_V.

It would be presumptive to assume that σ_ϵ^2, the variance generated from noise traders would remain unchanged during the outage. Suppose that there are always some noise traders in the market and that they face random shocks for liquidity. The objective function of the noise traders is to minimize trading costs given their exogenous demands for liquidity which may not be time sensitive. Specifically facing a shock ξ_t where $\xi_t \sim (0, \sigma_\xi^2)$, and a limited time window to fulfill the transaction, the trader will delay his trades if he expects future liquidity costs to decrease. Hence it's possible that σ_ϵ^2 becomes smaller during the outage. Consequently whether σ_V^2 is higher or lower during the outage is ambiguous.

A similar argument can be made to show that the change in $\sigma_{R^A}^2$ is also ambiguous during an outage, since we expect α_1 to be smaller during the outage, which will increase volatility, while noise traders during the outage would respond by demand less liquidity, and so $\sigma_{\epsilon_1}^2$ would be lower, which would decrease volatility.

$$\begin{aligned}
\sigma_{R^A}^2 &= \sigma_\eta^2 + \sigma_{\epsilon,1}^2 + \alpha_1^2 \sigma_V^2 \\
&= \sigma_\eta^2 + \sigma_{\epsilon,1}^2 + \frac{\alpha_1^2[\sigma_{\epsilon,1}^2 + \sigma_{\epsilon,2}^2]}{1-[1+\alpha_2-\alpha_1]^2}
\end{aligned} \tag{6}$$

In the empirical section we find the average trade size to be smaller during the outage, and the average duration between each trade to be longer. Hence the fact that market volatility decreased during the outage can be rationalized by the noise traders delaying their orders due to an expected increase in future liquidity.

FIGURE 1: Ten Year Spot Treasury Quote Depth During Outage Days

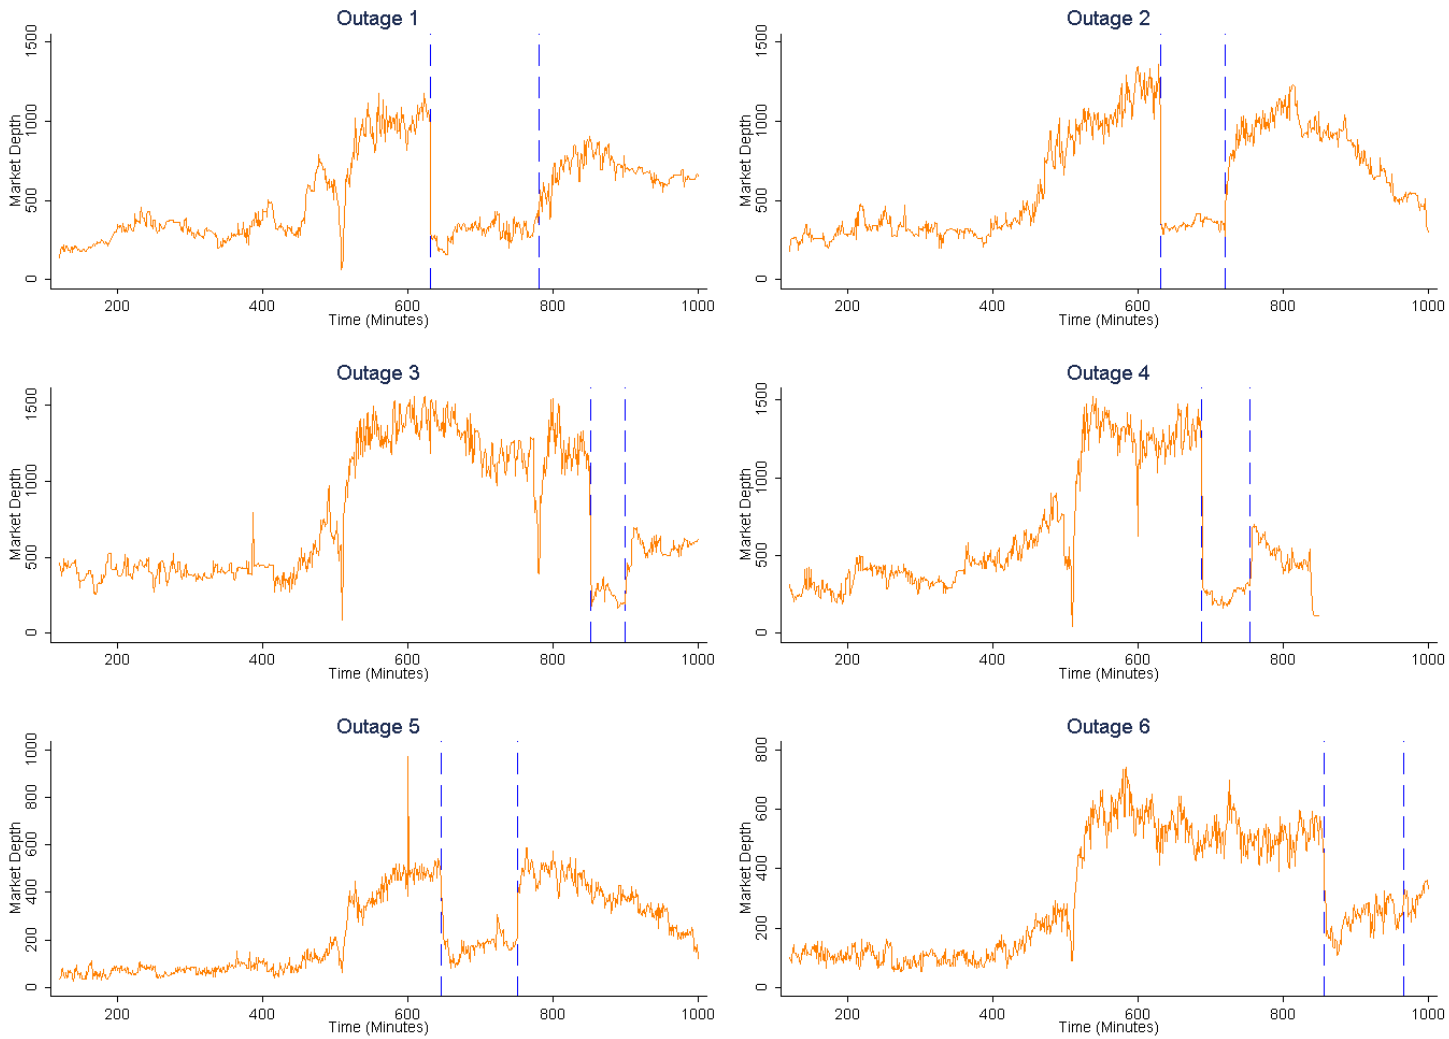

The yellow line measures the market depth, which is defined as the sum of the quantities available for trade at the best bid and best ask prices. Each panel represents a different outage day for the ten year spot treasury note traded on the Espeed exchange.

FIGURE 2: Ten Year Spot Treasury Market Participants During Outage Days

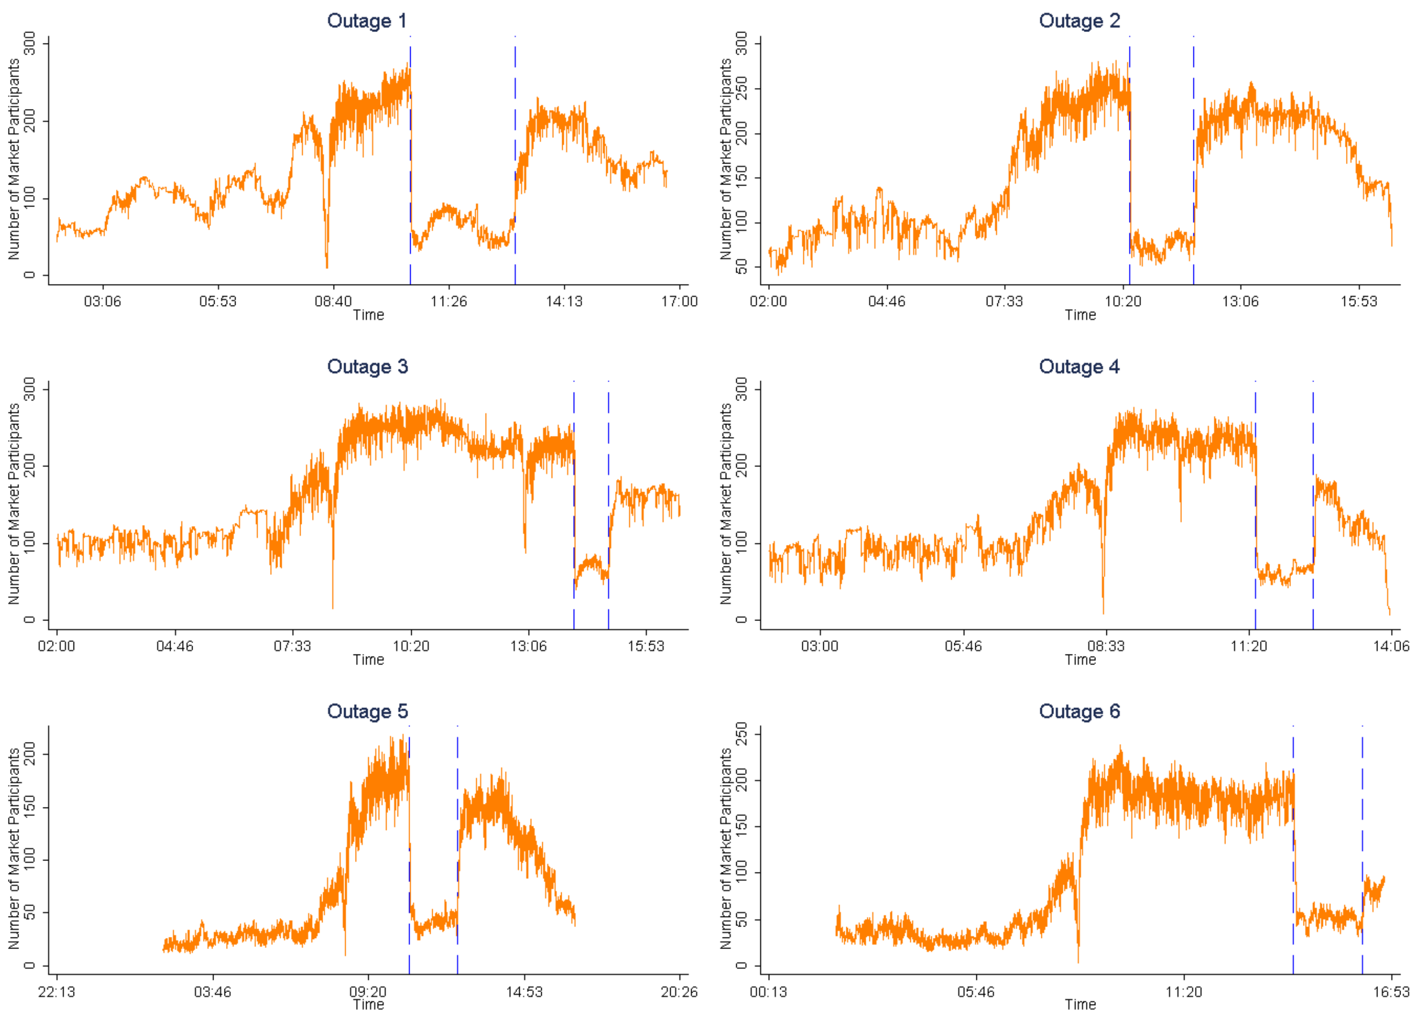

The yellow line measures the total number of participants in the 10 year US spot treasury market on Espeed aggregated from the sum of the participants in the best five ask prices and the best five bid prices. Each panel represents a different outage day. Blue dashed lines represent the timing of the initial crash and the subsequent recovery as reported by the media. See text for more details.

FIGURE 3:

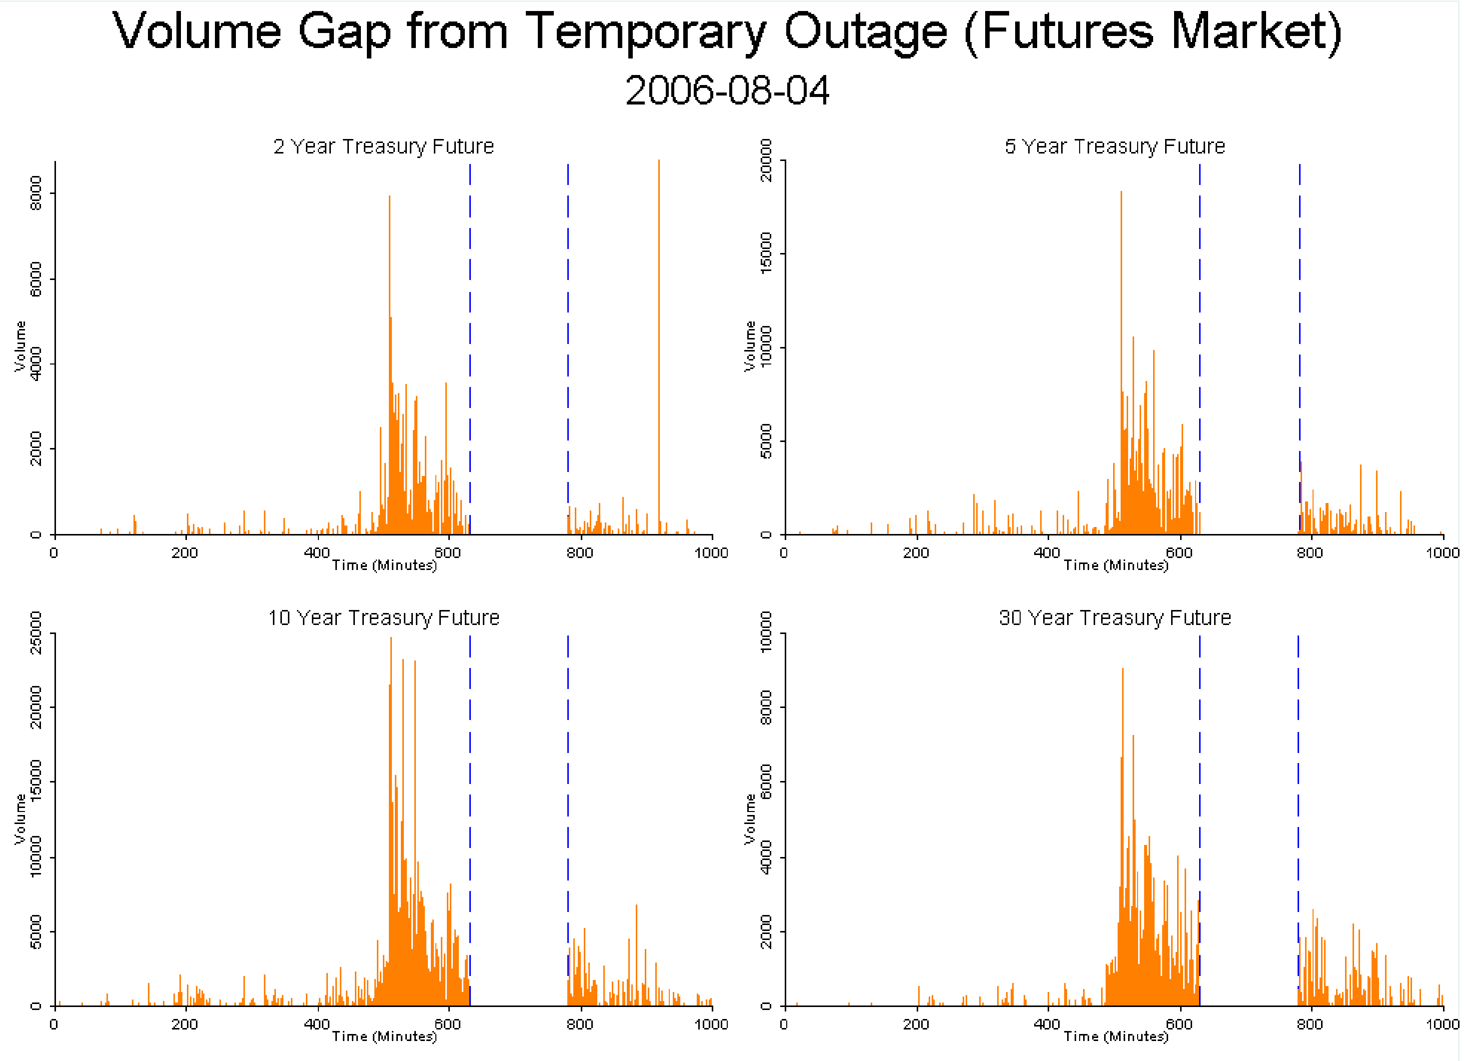

The vertical dashed lines indicates the beginning and end of an outage as reported by news agencies. Each panel represents a different futures maturity year for a representative outage day. Volume is from the electronic market only.

FIGURE 4:

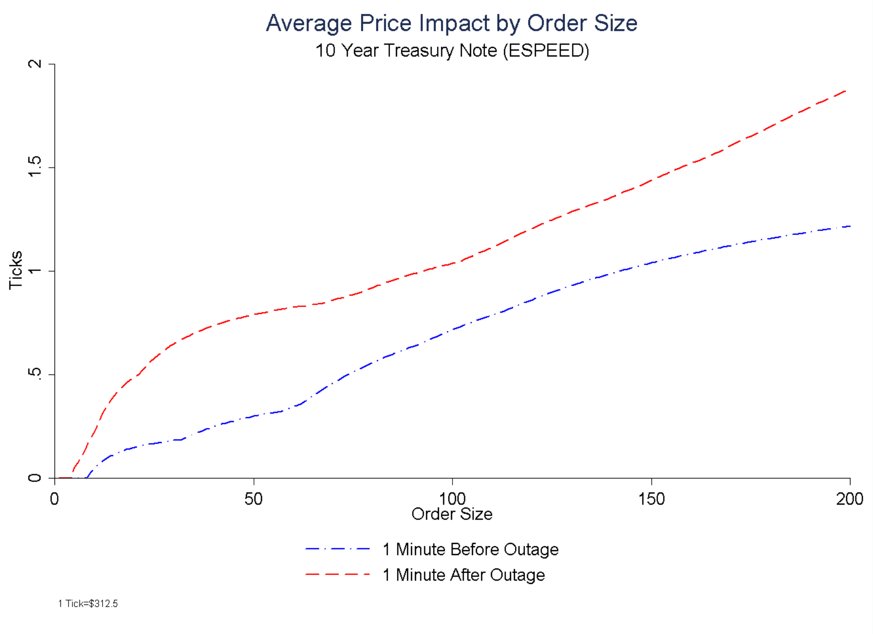

This measures the average price impact of an arbitrary order size above and beyond the current best bid/best before and after an outage.

FIGURE 5:

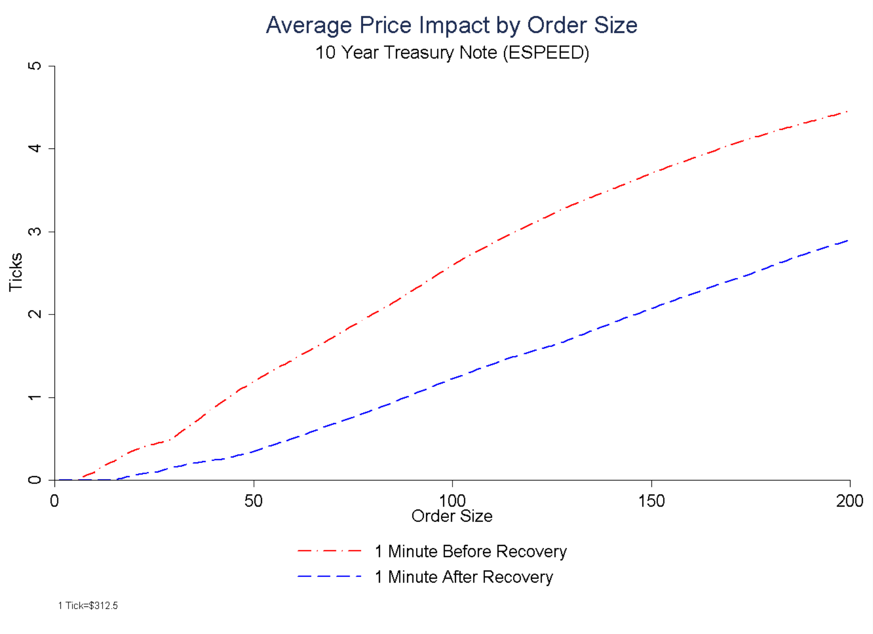

This measures the average price impact of an arbitrary order size above and beyond the current best bid/best before and after an outage recovery.

Table I: CBOT Electronic Exchange Outages

Date	Crash Time	Recovery Time	Duration	Volume	Volume %	Volatility %	VIX
8/4/2006	10:30 AM	1:00 PM	150	756137	97%	97%	14.34
10/3/2006	10:30 AM	12:00 PM	90	445727	70%	46%	12.24
1/11/2007	2:11 PM	3:00 PM	49	1157200	96%	68%	10.87
1/12/2007	11:27 AM	12:35 PM	68	745536	88%	70%	10.15
8/23/2007	10:46 AM	12:30 PM	106	613665	87%	96%	22.62
9/19/2007	2:15 AM	4:05 PM	110	1286518	98%	98%	20.03

Note: All times are based on EST and duration of outages are measured in minutes. Volume traded in the 10 year treasury future in CBOT for the day up until the outage is reported in column 5. Columns 6 and 7 report the percentiles of the volume and volatility up until the outage in relation to historical volume and volatility of the 10 year treasury future from 2006. Volatility is calculated via the realized s.d. of 1 minute observation of the best ask price. VIX is measured at end of day from CBOE.
Source: Bloomberg, Lexis Nexus

Table IIa: Intraday Summary Statistics Around Outage Period

2 Year	**Before**	**During**	**After**	**5 Year**	**Before**	**During**	**After**
Bid Ask Spread	1.06	1.10	1.04	Bid Ask Spread	1.15	1.23	1.11
Market Depth	4025.05	1332.87	2521.22	Market Depth	1029.19	315.72	690.48
Number of Participants	192.60	64.72	136.10	Number of Participants	180.39	42.14	122.19
Volume	79.14	24.01	30.89	Volume	77.53	18.36	34.18
Order Size	6.77	4.87	4.82	Order Size	3.95	3.73	2.93
Order Frequency	8.25	1.97	3.53	Order Frequency	19.96	3.84	8.91
Duration Between Trade	20.08	55.09	48.69	Duration Between Trade	7.55	26.43	20.66
Volatility	0.50%	0.38%	0.39%	Volatility	1.00%	0.66%	0.61%
10 Year	**Before**	**During**	**After**	**30 Year**	**Before**	**During**	**After**
Bid Ask Spread	1.08	1.11	1.09	Bid Ask Spread	1.69	2.44	1.82
Market Depth	1056.41	322.33	723.96	Market Depth	104.45	28.27	61.49
Number of Participants	214.45	61.69	161.90	Number of Participants	52.77	12.86	32.58
Volume	63.90	20.47	28.10	Volume	8.83	2.59	4.58
Order Size	3.19	3.53	2.54	Order Size	1.56	0.95	1.25
Order Frequency	18.49	4.87	8.38	Order Frequency	5.09	1.32	2.35
Duration Between Trade	6.92	22.02	20.50	Duration Between Trade	22.77	78.05	51.18
Volatility	2.94%	2.47%	1.11%	Volatility	2.98%	2.09%	2.01%

Notes: Each summary statistic represents the mean of one minute intraday observations. The different columns denoting before, during, and after correspond to 90 minutes before an outage, during an outage, and 90 minutes after an outage. The bid ask spread is normalized in terms of the minimum tick size (typically $\frac{1}{128}$ or $\frac{1}{64}$). Market depth is measured as the sum of the depth of the best five ask and bid prices. Number of participants is proxied with the number of unique orders in the market. Duration between trade is measured in seconds and is right censored. Volatility is measured using the changes in the best ask price to minimize noise from bid ask bounce.
Source: BG Cantor Market Data and authors' calculations

Table IIb: Summary Statistics of Liquidity Measurements

	Mean	Median	Std. Dev.	Min	Max	N
2 Year Note						
Quote Depth	2397.4	1913	1748.2	2	9491	494962
Bid Ask Spread	1.21	1	0.88	1	68	494962
Volume Traded	27.0	0	69.6	0	1580	494962
Order Size	3.42	0	7.31	0	1000	494962
Order Frequency	2.94	0	6.89	0	217	494962
Duration Between Trades	64.4	60	142.1	0	23130	494962
5 Year Note						
Quote Depth	595.9	480	426.5	2	4528	494717
Bid Ask Spread	1.47	1	1.64	1	124	494717
Volume Traded	24.3	2	47.4	0	1085	494717
Order Size	2.16	1	3.46	0	459	494717
Order Frequency	6.56	1	12.1	0	240	494717
Duration Between Trades	51.6	60	128.1	0	25431	494717
10 Year Note						
Quote Depth	595.0	486	399.0	2	6405	495526
Bid Ask Spread	2.69	2	2.81	1	124	495526
Volume Traded	19.8	2	38.1	0	859	495526
Order Size	1.85	1	2.76	0	250	495526
Order Frequency	6.42	1	11.9	0	240	495526
Duration Between Trades	49.9	60	106.5	0	22379	495526
30 Year Bond						
Quote Depth	65.0	56	41.5	2	659	494676
Bid Ask Spread	2.74	2	3.66	1	216	494676
Volume Traded	3.22	0	7.10	0	263	494676
Order Size	0.72	0	1.19	0	49	494676
Order Frequency	1.91	0	3.86	0	86	494676
Duration Between Trades	67.6	60	184.1	0	22151	494676
Total						
Quote Depth	913.4	437	1277.0	2	9491	1979881
Bid Ask Spread	2.03	2	2.58	1	216	1979881
Volume Traded	18.6	0	47.3	0	1580	1979881
Order Size	2.04	0	4.42	0	1000	1979881
Order Frequency	4.46	0	9.57	0	240	1979881
Duration Between Trades	58.4	60	143.2	0	25431	1979881

Notes: Each summary statistic represents the mean of one minute intraday observations. The bid ask spread is normalized in terms of the minimum tick size (typically $\frac{1}{128}$ or $\frac{1}{64}$). Market depth is measured as the sum of the depth of the best five ask and bid prices. Order size is the average size of a trade. Order frequency is the number of trades within a minute. Duration between trade is measured in seconds and is right censored (time intervals with no trades are replaced with an 60 second duration).
Source: BG Cantor Market Data and authors' calculations

Table III: Impact of Outage upon Log Quote Depth by Maturity Year

	(1)	(2)	(3)	(4)
2 Year Note	-0.442	-0.828	-0.669	-0.456
	[0.285]	[0.278]	[0.232]	[0.265]
5 Year Note	-0.463	-0.990	-0.824	-0.727
	[0.218]	[0.238]	[0.186]	[0.189]
10 Year Note	-0.385	-0.870	-0.730	-0.539
	[0.102]	[0.135]	[0.126]	[0.160]
30 Year Note	-0.683	-0.996	-0.939	-0.790
	[0.0445]	[0.0859]	[0.142]	[0.147]
Observations	495490	495490	55853	54818
Hourly Fixed Effects	N	Y	Y	Y
Monthly Fixed Effects	N	Y	Y	Y
Top 10% Volume	N	N	Y	N
Top 10% Volatility	N	N	N	Y

Notes: Standard errors are in brackets and are clustered at the daily level. The coefficient in each cell measures the impact of the crash upon log market depth. Observations are at the minute level. Each row represents a different bond maturity. Top 10% volume days restricts observations to be on days where the volume exceeded the 90% percentile of the historical distribution from Jan 2006 to December 2007 in addition to outage days. Top 10% volatility days is constructed analogously.

Table IV: Impact of Outage upon Log Bid Ask Spread by Maturity Year

	(1)	(2)	(3)	(4)
2 Year Note	-0.0516	0.00925	0.00415	0.0111
	[0.0165]	[0.0129]	[0.0126]	[0.0154]
5 Year Note	-0.0893	0.0188	-0.00737	-0.0139
	[0.0324]	[0.0407]	[0.0277]	[0.0386]
10 Year Note	-0.111	-0.0280	-0.0384	-0.0774
	[0.0114]	[0.0191]	[0.0134]	[0.0383]
30 Year Note	0.0321	0.249	0.180	0.134
	[0.0524]	[0.0563]	[0.0477]	[0.0561]
Observations	494676	494676	55775	54719
Hourly Fixed Effects	N	Y	Y	Y
Monthly Fixed Effects	N	Y	Y	Y
Top 10% Volume	N	N	Y	N
Top 10% Volatility	N	N	N	Y

Notes: Standard errors are in brackets and are clustered at the daily level. The coefficient in each cell measures the impact of the crash upon log bid ask spread, where bid ask spread is normalized to the minimum tick increment. Observations are at the minute level. Each row represents a different bond maturity. Top 10% volume days restricts observations to be on days where the volume exceeded the 90% percentile of the historical distribution from Jan 2006 to December 2007 in addition to outage days. Top 10% volatility days is constructed analogously.

Table V: Impact of Outage Upon Volume

	(1)	(2)	(3)	(4)
2 Year Note	-3.513	-22.44	-47.00	-40.81
	[2.217]	[4.319]	[9.047]	[5.567]
5 Year Note	-5.821	-27.81	-48.14	-42.63
	[2.325]	[5.546]	[8.418]	[6.672]
10 Year Note	0.683	-15.70	-34.69	-23.51
	[2.618]	[5.148]	[7.817]	[5.210]
30 Year Note	-0.604	-3.292	-5.652	-4.034
	[0.357]	[0.621]	[1.152]	[0.733]
Observations	501000	501000	56000	56000
Hourly Fixed Effects	N	Y	Y	Y
Monthly Fixed Effects	N	Y	Y	Y
Top 10% Volume	N	N	Y	N
Top 10% Volatility	N	N	N	Y

Notes: Standard errors are in brackets and are clustered at the daily level. The coefficient in each cell measures the impact of the crash upon volume traded. Observations are at the minute level. Each row represents a different bond maturity. Top 10% volume days restricts observations to be on days where the volume exceeded the 90% percentile of the historical distribution from Jan 2006 to December 2007 in addition to outage days. Top 10% volatility days is constructed analogously.

Table VI: Impact of Outage Upon Order Size

	(1)	(2)	(3)	(4)
2 Year Note	1.475	-0.280	-1.802	-1.268
	[0.579]	[0.765]	[0.862]	[0.829]
5 Year Note	1.588	0.421	0.249	0.428
	[0.338]	[0.376]	[0.352]	[0.432]
10 Year Note	1.691	0.820	0.523	0.823
	[0.175]	[0.220]	[0.222]	[0.243]
30 Year Note	0.234	-0.295	-0.463	-0.319
	[0.0941]	[0.113]	[0.125]	[0.126]
Observations	501000	501000	56000	56000
Hourly Fixed Effects	N	Y	Y	Y
Monthly Fixed Effects	N	Y	Y	Y
Top 10% Volume	N	N	Y	N
Top 10% Volatility	N	N	N	Y

Notes: Standard errors are in brackets and are clustered at the daily level. The coefficient in each cell measures the impact of the crash upon order size. Observations are at the minute level. Each row represents a different bond maturity. Top 10% volume days restricts observations to be on days where the volume exceeded the 90% percentile of the historical distribution from Jan 2006 to December 2007 in addition to outage days. Top 10% volatility days is constructed analogously.

Table VII: Impact of Outage Upon Order Freq

	(1)	(2)	(3)	(4)
2 Year Note	-0.955	-3.148	-6.188	-6.319
	[0.274]	[0.548]	[0.981]	[0.674]
5 Year Note	-2.681	-8.722	-14.79	-13.93
	[0.355]	[1.352]	[2.129]	[1.544]
10 Year Note	-1.537	-6.845	-12.86	-10.42
	[0.522]	[1.452]	[2.295]	[1.484]
30 Year Note	-0.579	-2.101	-3.537	-2.763
	[0.149]	[0.358]	[0.632]	[0.415]
Observations	501000	501000	56000	56000
Hourly Fixed Effects	N	Y	Y	Y
Monthly Fixed Effects	N	Y	Y	Y
Top 10% Volume	N	N	Y	N
Top 10% Volatility	N	N	N	Y

Notes: Standard errors are in brackets and are clustered at the daily level. The coefficient in each cell measures the impact of the crash upon order frequency. Observations are at the minute level. Each row represents a different bond maturity. Top 10% volume days restricts observations to be on days where the volume exceeded the 90% percentile of the historical distribution from Jan 2006 to December 2007 in addition to outage days. Top 10% volatility days is constructed analogously.

Table VIII: Impact of Outage Upon Duration Between Trades

	(1)	(2)	(3)	(4)
2 Year Note	-6.536	10.70	27.10	26.76
	[4.825]	[4.998]	[4.389]	[4.410]
5 Year Note	-17.49	13.02	20.64	20.03
	[3.717]	[5.503]	[4.740]	[5.320]
10 Year Note	-19.70	7.217	16.70	13.58
	[2.000]	[4.244]	[3.675]	[3.910]
30 Year Note	1.710	25.64	37.01	30.63
	[3.844]	[5.542]	[6.204]	[5.398]
Observations	501000	501000	56000	56000
Hourly Fixed Effects	N	Y	Y	Y
Monthly Fixed Effects	N	Y	Y	Y
Top 10% Volume	N	N	Y	N
Top 10% Volatility	N	N	N	Y

Notes: Standard errors are in brackets and are clustered at the daily level. The coefficient in each cell measures the impact of the crash upon duration between trades. Duration between trade is measured in seconds and is right censored (time intervals with no trades are replaced with an 60 second duration) Observations are at the minute level. Each row represents a different bond maturity. Top 10% volume days restricts observations to be on days where the volume exceeded the 90% percentile of the historical distribution from Jan 2006 to December 2007 in addition to outage days. Top 10% volatility days is constructed analogously.

Table IX: Measuring The Absence of Strategic Order Cancelation Behavior During Outages

	Second Best Ask Depth			Second Best Bid Depth		
	Before Outage	During Outage	After Outage	Before Outage	During Outage	After Outage
Impact of Buy	-0.0369	-0.0107	-0.0166			
Order	[0.0133]	[0.0284]	[0.0168]			
Impact of Sell				-0.0346	-0.00586	-0.0278
Order				[0.0130]	[0.0276]	[0.0183]
Observations	7910	3700	4176	8122	4147	4348
Maturity Year Fixed Effects	Y	Y	Y	Y	Y	Y
Hourly Fixed Effects	Y	Y	Y	Y	Y	Y
Monthly Fixed Effects	Y	Y	Y	Y	Y	Y

Table X: Measuring The Absence of Strategic Order Cancelation Behavior During Outages (Controlling for Seasonality)

	Second Best Ask Depth		Second Best Bid Depth	
	Outage Period	Non-Outage Period	Outage Period	Non-Outage Period
Impact of Buy	-0.0107	-0.0261		
Order	[0.00266]	[0.000851]		
Impact of Sell			-0.00586	-0.0248
Order			[0.00748]	[0.000791]
Observations	3700	2989162	4147	3126032
Maturity Year Fixed Effects	Y	Y	Y	Y
Hourly Fixed Effects	Y	Y	Y	Y
Monthly Fixed Effects	Y	Y	Y	Y

Tables IX and X are both based on equation 2 of the empirical section. The coefficient measures the effect of an incoming sell order upon the quote size at the second best bid price and the effect of an incoming buy order upon the quote size at the second best ask price. A negative coefficient represents order cancelations. In Table IX, the regression is ran for the sample of outage days only, with three periods, before, during, and after an outage, as defined in the paper. In Table X, the regression is ran using all sample days in the period between 2006 and 2007. Bond maturity fixed effects are included in all specifications and standard errors are clustered at the daily level.