

Federal Reserve Bank of New York
Staff Reports

Liquidity and Volatility in the U.S. Treasury Market

Giang Nguyen
Robert Engle
Michael Fleming
Eric Ghysels

Staff Report No. 590
December 2012
Revised November 2018



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Liquidity and Volatility in the U.S. Treasury Market

Giang Nguyen, Robert Engle, Michael Fleming, and Eric Ghysels

Federal Reserve Bank of New York Staff Reports, no. 590

December 2012; revised November 2018

JEL classification: C58, G01, G12

Abstract

We model the joint dynamics of intraday liquidity, volume, and volatility in the U.S. Treasury market, especially through the 2007-09 financial crisis and around important economic announcements. Using various specifications based on Bauwens and Giot's (2000) Log-ACD(1,1) model, we find that liquidity, volume, and volatility are highly persistent, with volatility having a lower short-term persistence than the other two. Market liquidity and volume are important to explaining volatility dynamics but not vice versa. In addition, market dynamics change during the financial crisis, with all variables exhibiting increased responsiveness to their most recent realizations. Our models also reveal different market dynamics around announcements. Finally, we introduce new measures of liquidity risk that are useful for continually monitoring liquidity conditions and the risk of liquidity stress in the market.

Key words: liquidity, Treasury market, limit order book, financial crisis, volatility, announcement

Nguyen: Pennsylvania State University (email: giang.nguyen@psu.edu). Engle: New York University (email: rengle@stern.nyu.edu). Fleming: Federal Reserve Bank of New York (email: michael.fleming@ny.frb.org). Ghysels: University of North Carolina at Chapel Hill (email: eghysels@unc.edu). This paper was previously circulated under the title "Liquidity, Volatility, and Flights to Safety in the U.S. Treasury Market: Evidence from a New Class of Dynamic Order Book Models." The authors benefitted greatly from discussions with Michael Aguilar, Saraswata Chaudhuri, Wayne Ferson, Corey Garriott, Nikolaus Hautsch, Jean Helwege, Jonathan Hill, Pete Kyle, Albert Menkveld, Adam Reed, and Duane Seppi. The authors also thank Brian Gibbons, Casidhe Horan Troyer, Collin Jones, Neel Krishnan, Weiling Liu, Francisco Ruela, and Ron Yang for excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

The importance of understanding market liquidity and volatility dynamics in the U.S. Treasury securities market stems from the market's many vital roles. Treasury securities are commonly used to price and hedge positions in other fixed-income securities and to speculate on the course of interest rates. The securities' creditworthiness and liquidity also make them a key instrument of monetary policy and a crucial source of collateral for financing other positions. These same attributes make Treasury securities a key store of value, especially during times of crisis.

Recent trends in liquidity in this key market make understanding its dynamics ever more important. As we show, market liquidity worsened markedly during the 2007–09 financial crisis, with sharply higher price volatility and markedly reduced depth. After recovering strongly in the years after the crisis, liquidity then deteriorated abruptly again during the mid-2013 taper tantrum, when the market came to expect the Federal Reserve to soon start reducing its large-scale asset purchases. Liquidity then worsened again following the October 15, 2014 flash rally, when Treasury yields traded in a wide range before closing the day largely unchanged. The flash rally, and subsequent study, are a powerful reminder that emerging trends in market structure might have effects on market liquidity that are not yet well understood (see Joint Staff Report, 2015).

In this paper, we model the dynamics of liquidity and volatility and how they interact in a unified econometric framework designed to accommodate empirical regularities unique to the Treasury market. With the availability of intraday data on the limit order book of Treasury securities in the interdealer market, the model can be cast in high frequency time intervals, and can accordingly convey rich and insightful information about the micro behavior of liquidity and volatility. Events with important implications for volatility and liquidity, such as the flash rally, might be missed by lower frequency models. We use our model to examine market dynamics during the period surrounding the 2007-09 financial crisis, and around important economic announcements. With a 10-year sample of high frequency limit order

book data spanning completely the financial crisis and the years that follow, we are able to isolate the effects of the crisis on the dynamics of these important variables.

Our econometric framework is built upon the large literature on modeling high frequency financial time series. It is based on the multiplicative error model (MEM), first developed in Engle and Russell (1998) and formally introduced in Engle (2002), to model non-negative valued variables. Conveniently, our variables of interest – market depth, trading volume, and volatility – are all non-negative variables and such non-negativity is implicitly guaranteed under a MEM structure. This is a major departure from the empirical literature on limit order book modeling, which largely relies on the standard linear Gaussian framework.¹ Clearly, linear error models cannot guarantee non-negativity of all possible outcomes for the modeled variables. This issue is particularly acute if we are interested in not only the conditional moments of a variable but the entire conditional distribution. The framework we introduce allows for meaningful density predictions. From here, we can compute the tails of liquidity and volatility outcomes. One important use of our framework is for constructing measures of liquidity risk that can be used to continually monitor market liquidity conditions for trading and risk management purposes.

Our econometric model also includes important features to accommodate zero or low values of depth, volume, and volatility. These are not uncommon occurrences in this market. For example, the literature shows (e.g., Fleming and Remolona, 1999; Fleming and Piazzesi, 2005), and we confirm that depth in the Treasury market tends to disappear prior to economic news announcements. We also document a large liquidity drop in the fall of 2008. In addition, it is inevitable to have zero volatility or volume during quiet times of a trading day when we measure these variables at high frequencies. The multiplicative error structure we suggest can

¹For example, Ahn, Bae, and Chan (2001) use a regression framework to study the dynamics of the number of limit orders posted. Likewise, Næs and Skjeltorp (2006) regress trade size and number of transactions on volatility to document the existence of a volume-volatility relationship in the Norwegian equity market. Härdle, Hautsch, and Mihoci (2012) propose a dynamic semiparametric factor approach to modelling liquidity supply, combining non-parametric factor decomposition for the order curve’s spatial structure with a vector autoregression (VAR) for time variations of factor loadings. Other studies similar in their use of VAR include Danielsson and Payne (2012) and Hautsch and Huang (2012), among others. Rinaldo (2004) uses an ordered probit regression framework to analyze how the state of the limit order book affects order submission strategy.

accommodate zero or small values of a variable with a more realistic probability distribution than is possible in the class of linear error models. Furthermore, we can rely on a mixture probability distribution as in Hautsch, Malec, and Schienle (2013), augmenting a positive-valued continuous probability distribution with a point mass at zero, to more accurately reflect the distribution observed in the data and facilitate model estimation. Lastly, we adopt Bauwens and Giot (2000)'s Log-ACD₁ formulation to model the log conditional expectation, a key advantage of which is to ensure the non-negativity of the conditional mean, especially if we want to include covariates to explain the dynamics of the modeled variables.

Based on our model, we find that market liquidity, trading activity, and volatility are highly persistent, but that there is an important difference in the dynamics of volatility versus depth and volume. We characterize the short-term persistence as the coefficient on the most recent realization and the long-term persistence as the coefficient on the past conditional expectation (or its log). Both depth and volume have higher short-term persistence, but lower long-term persistence, than volatility. Nevertheless, all three variables exhibit increased short-term persistence during crisis times. There are also important changes in the dynamics of these variables around economic announcements. Volatility's short-term persistence increases both before and after announcements, whereas that of trading volume increases only after announcements, reflecting continued trading on the news. In addition, volatility appears to increase upon the release of negative news (based on the direction and the surprise component of price change upon the news release). Another important finding is that depth and volume are important to the dynamics of volatility, whereas their own dynamics seem much less affected by volatility.

Our model allows us to construct two measures of liquidity risk, with one indicating the probability of an extreme illiquidity shock, and the other showing the level of illiquidity that occurs with a 5% probability (which we refer to as the "liquidity-at-risk" measure due to its similarity in spirit to the concept of value-at-risk). Our liquidity risk measures point to a spike in liquidity risk in late 2008 at the height of the financial crisis. The measures

are relatively subdued after the crisis, but do point to some mildly increased liquidity risk episodes in subsequent years. Importantly, our model shows that extreme illiquidity events are much more highly correlated with extremely high volatility episodes than the usual liquidity-volatility correlation would suggest.

Our study contributes to the extensive literature on price formation and liquidity in the U.S. Treasury market. This strand of literature includes Fleming and Remolona (1999), Balduzzi, Elton, and Green (2001), Huang, Cai, and Wang (2002), Fleming (2003), Brandt and Kavajecz (2004), Green (2004), Fleming and Piazzesi (2005), Goldreich, Hanke, and Nath (2005), Mizrach and Neely (2007), Pasquariello and Vega (2007), Fleming, Mizrach, and Nguyen (2018), Jiang, Lo, and Verdelhan (2011), and many others. However, most of the extant studies use data prior to the 2008 crisis period, leaving market dynamics during the crisis – the most serious to hit the global economy since the Great Depression – less documented.²

Our study is also related to papers that document asset pricing anomalies during the financial crisis. Fleckenstein, Longstaff, and Lustig (2014) find that a significant mispricing arose between Treasury bonds and inflation-swapped TIPS issues with replicating cash flows. Musto, Nini, and Schwarz (2018) detail a large and systematic mispricing during the crisis between notes and bonds with identical cash flows. Hu, Pan, and Wang (2013) show that “noise” in Treasury security prices rose sharply during the crisis. Our paper also reveals unusual market behavior during the crisis, but by directly assessing market liquidity. Moreover, while the previously mentioned pricing anomalies are shown to have arisen largely among less traded Treasury securities, our study identifies liquidity declines in the most actively traded on-the-run Treasury securities.

The paper is organized as follows. We first provide in Section 2 a descriptive analysis of important liquidity and volatility trends in the Treasury market. In Section 3, we present our econometric framework designed to accommodate empirical regularities specific to this

²One exception, written after an earlier draft of this paper, is Adrian, Fleming, Shachar, and Vogt (2017), which documents the evolution of Treasury liquidity from 2006 to 2016.

market. We then discuss the empirical dynamics of Treasury market liquidity and volatility as estimated by the model in Section 4. Section 5 contains our liquidity risk measures, a direct and important application of the model. Finally, Section 6 concludes.

2 The U.S. Treasury securities market

U.S. Treasury securities are debt instruments sold by the U.S. government through public auctions and subsequently traded in the secondary market. Our paper focuses on the most liquid segment of the secondary market, which is the electronic interdealer market for on-the-run Treasury securities. On-the-run securities account for most trading activity in the market, and their liquidity in the interdealer market can be observed and measured with available high frequency data.³ The liquidity of the dealer-to-customer market is more opaque because of data limitations, but likely tracks the liquidity in the interdealer market.⁴

2.1 Data

The data for our study covers the period from 2006 to 2015, and comes from the BrokerTec platform.⁵ The ten-year sample period covers completely the 2007–09 financial crisis and important events in the post crisis period, as well as a period before the crisis. We date the start of the crisis to August 9, 2007, when BNP Paribas announced that it could not value assets in three of its investment funds.⁶ There is no clear ending date to the crisis, so we mark the end with the NBER’s end-of-recession date of June 2009.

³Fleming (2016) estimates that trading in on-the-run securities accounts for roughly 85% of total trading volume across nominal Treasury securities.

⁴Weekly dealer statistics published by the Federal Reserve Bank of New York (FR2004 data) indicate that interdealer trading accounts for slightly less than half of all trading activity, with dealer-to-customer trading accounting for the remainder.

⁵Fleming, Mizrach, and Nguyen (2018) analyze the microstructure of the BrokerTec platform and report that it accounts for 60% of electronic interdealer trading for each of the on-the-run 2-, 5- and 10-year notes, with eSpeed accounting for the rest.

⁶Note that the Business Cycle Dating Committee of the National Bureau of Economic Research (NBER) considers December 2007 as the start of the recession (<http://www.nber.org/cycles.html>). However, for the purpose of modelling liquidity in the Treasury market, the earlier date at which the crisis started in financial markets is more appropriate.

We reconstruct the tick-by-tick limit order book from BrokerTec order message data. For our modeling, we work with five-minute snapshots, extracted from the tick data at the end of each five-minute interval during the New York trading hours of 7:00 to 17:30 Eastern Time.⁷ We drop 104 days from the sample, consisting of 75 days with an early market close and 29 days with unusable raw data files.⁸ Our final sample has 2,395 trading days, each with 126 five-minute intervals, for a total of 301,770 five-minute observations. From this data, we construct important liquidity and volatility variables as described below.

2.1.1 Market liquidity

We measure market liquidity by market depth, the quantity available for trading at various price levels. This is an ex-ante measure of liquidity, and a suitable measure of liquidity especially for this market. The bid-ask spread — an often used measure of liquidity — is often constrained by the minimum tick size and “stuck” at one tick.⁹ More nuanced fluctuations in market liquidity show up in the depth measure as liquidity providers adjust the quantities they are willing to buy or sell. For descriptive statistics purpose, we report depth at the best tier, as well as the second through the fifth best tiers, separately for the bid side and ask side. For model estimation, we use the average of the bid depth and ask depth at the best tier, noting that all depth variables exhibit similar variation.

On the liquidity demand side, we consider trading volume, an ex-post measure of liquidity. It provides complementary information to gauge the level of market liquidity, especially when liquidity providers are reluctant to post a full liquidity supply schedule ex-ante, which is often the case in volatile times. We include both measures in our model: market depth as a

⁷The five-minute frequency reflects a reasonable compromise between the need to understand market dynamics at the highest frequency possible, and a sufficiently long interval for meaningful movement in the modeled variables. For example, realized volatility can be zero during quiet times of a trading day. Therefore, sampling the data at too high a frequency results in having a variable dominated by zeroes.

⁸Days with an early bond market close, as recommended by the Securities Industry and Financial Markets Association, typically occur before holidays or holiday weekends and are collected from <http://www.sifma.org/uploadedfiles/research/statistics/statisticsfiles/misc-us-historical-holiday-market-recommendations-sifma.pdf>.

⁹Fleming, Mizrach, and Nguyen (2018) show that the bid-ask spread is one tick about 94% of the time.

proxy for liquidity supply and trading volume for liquidity demand. Naturally, one might expect that the two are related, because trading consumes liquidity supplied by the book, but it can also stimulate additional liquidity supply (e.g., Biais, Hillion, and Spatt, 1995). Yet, the model by Johnson (2008) postulates that volume is determined by the flux of entering and exiting agents, which should be independent of the total level of market participation that determines the level of market liquidity. In modeling both liquidity supply and trading volume, we seek to provide additional evidence for the linkage between the supply of and demand for liquidity.

2.1.2 Market volatility

An analysis of market liquidity would be incomplete without consideration of market volatility, which increases pick-off risk and inventory risk for liquidity providers and might affect their liquidity provision strategies. Conversely, the lack of liquidity supply makes it harder for the market to absorb shocks, thereby increasing volatility. Furthermore, as implied by Kyle (1985)'s model, trading conveys information. As a result, trading provides the means for price discovery and accordingly is expected to be associated with increased volatility. Our inclusion of both liquidity supply (depth) and liquidity demand (trading volume) in a system together with volatility seeks to shed light on the relative importance of possible drivers of volatility: the lack of liquidity supply to absorb shocks, and the price discovery induced by trading activity.

To measure market volatility at the five-minute frequency, we choose the noise-robust estimator of volatility based on the popular pre-averaging method introduced by Jacod, Li, Mykland, Podolskij, and Vetter (2009). We apply the commonly used tent-shaped weighting function $g(x) = \min[x, (1 - x)]$ on one-second log returns to construct pre-averaged returns. Following Hautsch and Podolskij (2013), we choose the pre-averaging parameter $\theta = 0.4$. This is on the lower end of the range recommended by Hautsch and Podolskij (2013), because the Treasury securities under study are highly liquid and the return sampling frequency of

one-second is high (adopting too high a θ results in oversmoothing and consequently negative bias in the estimator.) The optimal bandwidth is therefore $\lceil \theta\sqrt{n} \rceil = 7$, where $n = 300$ corresponding to 300 seconds per each five-minute interval. The pre-averaging is important for estimating realized volatility using ultra high frequency data because it helps diminish the influence of microstructure noise. The pre-averaging realized variance estimator is then the sum of squared pre-averaged returns over each five-minute interval. We then take the square root of the realized variance to obtain the pre-averaged realized volatility, annualize it by a factor of $\sqrt{12 \times 24 \times 250}$, and denote it by $\sigma_{i,t}$ throughout the paper’s empirical analysis.

2.1.3 Descriptive statistics

We report some summary statistics of the limit order book variables in Table 1. At the inside tier, the average depth for the 2-year note is nearly \$400 million, roughly eight to ten times greater than that for the 5- and 10-year ones. Average depth at the second through fifth tiers is about four to six times greater than that at the inside tier, indicating that there is typically more depth standing at each of the tiers behind the market than at the inside tier. Furthermore, the inside depth is more volatile relative to its level and has lower auto-correlation as compared to the next four tiers, likely because it is more directly affected by the trading process. With regard to trading volume, the 5- and 10-year notes are more actively traded than the 2-year note, averaging \$277 and \$241 million respectively every five minutes compared to \$197 million, despite having lower market depth than the latter. Average annualized volatility is lowest for the 2-year note (1.6%) and highest for the 10-year note (8.4%), which is what one would expect given the difference in their durations. Finally, the statistics indicate that both liquidity and volatility variables have a positive probability mass at zero, which calls for special care in our econometric modeling of their dynamics. We will discuss this point in detail later in the paper.

2.2 Market behavior through the crisis

We plot the time series behavior of market depth, volume, and volatility for the 2-, 5-, and 10-year notes in Figure 1. For easier comparison across securities, we standardize these series by their levels at the beginning of 2006 when the sample starts. The first four plots collectively show that market depth starts to decline sharply in mid 2007, providing evidence of mounting pressures in the Treasury market at the onset of the crisis. Liquidity drops sharply again (from an already low level) in the fall of 2008 following Lehman Brothers' bankruptcy. Depth then bottoms out towards the end of 2008, before slowly improving in the three years thereafter. There is then a significant upward movement in market depth in 2012 and the early part of 2013, before dropping sharply again during the taper tantrum in mid 2013 and after the flash rally in October 2014. The similarity of patterns observed at both the inside tier and outer tiers indicate that the time series fluctuations occur across price levels in the limit order book.¹⁰

On the liquidity demand side, there is also a major slide in trading volume after Lehman's failure, although trading volume is strong prior to that point despite the decline in market depth. This suggests that the demand for trading and liquidity still existed, and was fulfilled, up to that point, but that liquidity providers were cautious in posting ex-ante liquidity supply. However, at the peak of the crisis in the fall of 2008, both market depth *and* trading activity plunge altogether, signaling serious liquidity stress. After that, trading activity in the 2-year continues to dwindle, but trading volume in the longer maturities rebounds and even exceeds the crisis peak. The divergent behavior of the 2-year note coincides with the zero interest rate environment in the post crisis period, during which volatility and trading demand for the 2-year remain low even though the supply of liquidity is high. That is, the low level of trading activity in the 2-year note after the crisis is likely a reflection of the lessened appeal of the note to traders, rather than a reflection of the difficulty to transact.

¹⁰We do not go beyond the 5th tier because the best five tiers collect the majority of the limit order book's total quantity, on average accounting for 64%, 58%, and 52% of total market depth in the 2-, 5-, and 10-year notes respectively.

Finally, we contrast the variation in liquidity with that in volatility shown in the last plot. Accompanying the drop in market depth in the crisis, price volatility rises and reaches its peak in late 2008 following Lehman’s failure. Volatility declines gradually from this peak, but experiences further spells of volatility in 2010 and 2011 around the time of the Euro-area sovereign debt crisis, in May 2013 during the taper tantrum, and in late 2014 around the flash rally.

2.3 Market behavior throughout the trading day

We plot in Figure 2 the average variation in market liquidity and volatility over the course of the trading day. For each variable, we compute the full-sample average for each five-minute interval, and plot these averages against the corresponding intervals. To facilitate cross-security comparison, we standardize intraday averages by the average for the first interval of the day, i.e., [7:00-7:05).

The figure shows that depth exhibits an inverted U-shape pattern over the trading day. Depth builds up in the morning, reaches its peak shortly before noon and starts to decline after 15:00. There are also dips in depth shortly before 8:30, 10:00, 13:00, and to a lesser extent 14:00. Turning to volume and volatility, the intraday patterns show an increase in the morning, rising to the day’s highest level at 8:30. There are then subsequent spikes at 10:00, 11:00, 13:00, 14:00, and 15:00, after which volume and volatility both trail off. These intraday patterns are consistent with previous studies, such as Fleming and Remolona (1997) and Fleming and Remolona (1999), and demonstrate a departure from the well-documented U-shaped patterns in equity markets (although the Treasury market patterns are U-shaped between 8:30 and 15:00).

The intraday patterns are largely explained by various announcements. The most important scheduled macroeconomic announcements are released at 8:30 and 10:00, results of the Federal Reserve’s open market operations are released right after the operations close at 11:00, Treasury auction results are released within minutes of the 13:00 auction close,

and most announcements by the Federal Open Market Committee (FOMC) in the sample period were made at 14:00 or 14:15. The behavior around 15:00 coincides with the close of open outcry futures trading in Chicago and the pricing of fixed income indices (which likely spurs increased trading demand by investment managers who are seeking to rebalance their portfolios and/or minimize tracking errors relative to the indices).

To ascertain the importance of the announcements, we perform an event study analysis of depth and volatility around scheduled macroeconomic announcements, Treasury auctions, and FOMC announcements. We confirm that the aforementioned drops in market depth and corresponding spikes in volatility and trading volume at 8:30, 10:00, 13:00, and 14:00 can be attributed to these events,¹¹ consistent with the literature (e.g., Fleming and Remolona, 1999; Fleming and Piazzesi, 2005). We thus verify that it is necessary to account for announcement effects in our modeling, as well as the deterministic diurnal patterns in the limit order book variables.

3 Econometric framework

The characteristics of the Treasury market limit order book variables guide the building of our econometric framework. In this section, we discuss modeling considerations, specify our proposed framework, and describe estimation details.

3.1 Modeling considerations

The framework requires choices to be made with respect to each of its four building blocks: 1) the error structure, 2) the distribution for the errors, 3) the model for the conditional expectation, and 4) whether and how to model variables jointly.

¹¹The announcements considered are listed in Appendix A. The results are not presented here to save space, but are available from the authors.

3.1.1 Error structure

The first choice is concerned with whether the error structure is linear or multiplicative. Market liquidity and volatility variables are non-negative variables. Engle (2002) discusses key difficulties of the standard, widely used linear Gaussian framework in modeling non-negative variables, especially the challenge of ensuring non-negativity of the conditional mean and in describing closely their empirical dynamics. A large literature, starting with Engle and Russell (1998)'s ACD model, is devoted to developing non-linear models that can parsimoniously capture important dynamics of high frequency financial series, especially those that take on non-negative values. The MEM is one important class of such models (see Engle, 2002, for an in-depth discussion of the advantages of MEM models). With the error term being a multiplier of the conditional mean and from a positive-valued distribution, the non-negativity of the modeled variables is implicitly guaranteed. This is especially important if we are interested in not only the conditional mean, but also the complete conditional distribution of the variables of interest. It is indeed an important objective of our paper to build high frequency measures of liquidity risk that can be used in a real time setting to monitor liquidity risk in the Treasury market. A multiplicative error structure is therefore highly appropriate for our purpose.

3.1.2 Distribution of errors

Secondly, both liquidity and volatility variables have a non-trivial probability mass at zero. Zero values can happen at quiet times of day, especially for volatility or trading volume. This issue is exacerbated the higher the data sampling frequency. The conventional MEM formulation relies on continuous probability distributions with non-negative support, but such distributions cannot accommodate the point mass at zero. This calls for the use of a zero-augmented distribution as suggested in Hautsch, Malec, and Schienle (2013).¹² In

¹²An alternative, and simple to implement, method of dealing with observed zeros is proposed in Sucarrat and Escribano (2017), which involves replacing zeros with conditional expectation estimates. This method works well when the zeros reflect missing data or measurement errors. In our empirical analysis, zero

this paper, we construct the error distribution by mixing the point mass at zero with the log-normal distribution. Allen, Chan, McAleer, and Peiris (2008) prove consistency and asymptotic normality of QMLE for the Log-ACD model with log-normal distribution. Using Monte Carlo analysis, they demonstrate the desirable finite sample properties of the QMLE based on the log-normal distribution in comparison to several other widely used distributions (e.g., exponential, Weibull, and generalized gamma). Accordingly, QMLE estimates using the log-normal are the most robust.

3.1.3 Conditional expectation model

The third component of the framework is a model for the conditional expectation (or some transformation of the conditional expectation). We choose to work with the Log-ACD₁ model developed by Bauwens and Giot (2000) – a model for the log conditional expectation. The most important advantage of this model is to guarantee the non-negativity of the conditional mean, given that the modeled variable is non-negative, without the need to impose restrictions on parameter estimates. This advantage is especially important when we want to add other explanatory variables of interest to the model to test for their effects. Including additional explanatory variables would make the non-negativity restriction in a typical MEM formulation particularly challenging. In fact, this challenge is one of the key motivations for the class of log-GARCH-X models, as discussed in, e.g., Francq and Sucarrat (2017).

3.1.4 Contemporaneous dependence structure

Lastly, the summary statistics and figures demonstrate that liquidity and volatility are interdependent. This necessitates a multivariate modeling approach in which the interdependencies among the variables are explicitly accounted for. In linear models, this could be done either by assuming a trivariate Gaussian distribution or by a Cholesky factorization

observations are economically valid. For example, prior to importance announcements, traders withdraw their limit orders, leaving the limit order book empty. In such instances, market depth is truly and validly zero. We therefore choose to work with the Hautsch, Malec, and Schienle (2013) method, which treats these zeros as part of the distribution.

of the covariance matrix implying some causal ordering. In multiplicative error models, we face a similar choice in how to account for the contemporaneous correlation structure.

The first approach of adopting a multivariate distribution for the error vector is much more challenging than is a linear setup. Cipollini, Engle, and Gallo (2017) point out the lack of probability density functions for multivariate positive valued random variables. To pursue this route, they suggest the use of copula functions to jointly estimate the parameters of the scale factors and of the correlations of the error processes. However, this approach is especially challenging for our setting due to the fact that the variables we model can take on valid values of zero, necessitating a mixture distribution for each of the error processes, which in turn complicates the estimation of copulas governing the multivariate dependence structure.

In comparison, the second approach similar to a Cholesky factorization in linear models is much more suitable for our specific modeling needs. By explicitly capturing the contemporaneous correlation structure in the conditional mean equations, the errors are orthogonalized and hence the model can be estimated with a separate, univariate probability distribution specified for each error term. This approach allows us to easily accommodate the mixture distributions we specify for the errors. Furthermore, the ability to estimate the model equation-by-equation is particularly important for studies of high-frequency dynamics of market variables where there is often a need to apply the model to high-dimensional multivariate systems (see, e.g., Escribano and Sucarrat, forthcoming, for a discussion of the “curse of dimensionality” and the important advantages of equation-by-equation estimation approaches). One of the main concerns with this approach (in both linear and nonlinear settings) is that it requires an assumption on the recursive ordering of model variables. However, this disadvantage can be satisfactorily remedied through economic intuition. For example, Manganelli (2005) relies on economic reasoning to stipulate the inclusion of contemporaneous other variables in the conditional mean equation for each variable in his model. Similar to Manganelli (2005), we also establish the causal ordering based on our knowledge of the market and relevant

microstructure theories. We will describe this ordering in detail below where we specify our model.

3.2 Model specification

For notation, days are denoted by $t \in \{1, \dots, T\}$, and five-minute intervals within a trading day (from 7:00 to 17:30) are indexed by $i \in \{1, \dots, M\}$, with $M = 126$. The three variables to model are: market depth $d_{i,t}$ (the average of the inside bid and ask depths), trading volume $v_{i,t}$, and volatility $\sigma_{i,t}$.

A generic Log-ACD₁(p,q) model developed by Bauwens and Giot (2000) for a non-negative time series of interest y_i (with the day subscript t suppressed) is:

$$y_i = \exp(\psi_i) \epsilon_i, \quad (1)$$

$$\psi_i = \omega + \sum_{j=1}^p \alpha_j \ln(y_{i-j}) + \sum_{j=1}^q \beta_j \psi_{i-j} + c' \mathbf{z}_{i-1}, \quad (2)$$

where \mathbf{z}_{i-1} are covariates that help explain the dynamics of y_i , and ϵ_i is the multiplicative error having a conditional probability distribution with unit mean and non-negative support. For Log-ACD₁(1,1), a necessary stationary condition is $\alpha + \beta < 1$.

As previously discussed, we choose to capture the contemporaneous error correlation structure explicitly through the conditional mean equations. Thus, we expand the generic Log-ACD₁ formulation as follows, choosing order (1,1) for parsimony:

$$Y_{i,t} = \exp(\Psi_{i,t}) \odot \epsilon_{i,t} \quad (3)$$

$$\Psi_{i,t} = \mathbf{w} + A_0 \ln(\epsilon_{i,t}) + A_1 \ln(Y_{i-1,t}) + B\Psi_{i-1,t} + C\mathbf{z}_{i-1,t},$$

where $Y_{i,t} = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$, \odot is the Hadamard matrix product, $\epsilon_{i,t} = [\epsilon_{i,t}^v, \epsilon_{i,t}^\sigma, \epsilon_{i,t}^d]'$ and $\Psi_{i,t} = [\psi_{i,t}^v, \psi_{i,t}^\sigma, \psi_{i,t}^d]'$, with $\psi_{i,t}^v$, $\psi_{i,t}^\sigma$, and $\psi_{i,t}^d$ denoting the log of the conditional expectations of $v_{i,t}$, $\sigma_{i,t}$, and $d_{i,t}$ respectively.

The ordering of the elements of $Y_{i,t}$ is chosen to reflect the following. By construction, market depth $d_{i,t}$ is measured *at the end* of the interval i and therefore is arguably affected contemporaneously by trading volume $v_{i,t}$ and volatility $\sigma_{i,t}$ that occur *during* the interval i . The ordering between $v_{i,t}$ and $\sigma_{i,t}$, on the other hand, is guided by theory. More specifically, the seminal strategic trading model by Kyle (1985) suggests that informed traders choose the optimal quantity to trade to exploit their information advantage, taking into account the effect their demand would have on price. Kyle's model therefore suggests a causality that runs from trading volume to price, supporting that $v_{i,t}$ have a contemporaneous effect on $\sigma_{i,t}$ but not vice versa. In short, the contemporaneous causal ordering is from $v_{i,t}$ to $\sigma_{i,t}$ to $d_{i,t}$, implying that A_0 is lower triangular with zeros on its diagonal. Thus, the model is to be estimated recursively. We first estimate the equation for trading volume $v_{i,t}$, which does not depend contemporaneously on the other two variables, to obtain the innovations of volume. The contemporaneous volume innovations then enter the second equation for volatility $\sigma_{i,t}$, and subsequently coupled with the contemporaneous volatility innovations, permit the estimation of the equation for liquidity $d_{i,t}$. This way, we are able to estimate the model equation by equation and sidestep the challenge of modeling the copula governing the vector error process.

While our specification is not the usual representation due to the presence of $A_0 \ln(\epsilon_{i,t})$ in the conditional mean $\Psi_{i,t}$, we can rewrite it to the conventional representation as:

$$\begin{aligned}
Y_{i,t} &= \exp(\Psi_{i,t} + A_0 \ln(\epsilon_{i,t})) \odot \epsilon_{i,t} = \exp(\Psi_{i,t}) \odot \tilde{\epsilon}_{i,t} & (4) \\
\tilde{\epsilon}_{i,t} &= \exp(A_0 \ln(\epsilon_{i,t})) \odot \epsilon_{i,t} \\
\Psi_{i,t} &= \mathbf{w} + A_1 \ln(Y_{i-1,t}) + B\Psi_{i-1,t} + C\mathbf{z}_{i-1,t}.
\end{aligned}$$

Taking the logarithmic transformation of $\tilde{\epsilon}_{i,t}$ yields:

$$\ln \tilde{\epsilon}_{i,t} = (A_0 + I) \ln(\epsilon_{i,t}). \tag{5}$$

Without loss of generality, assume that $\ln(\epsilon_{i,t}) \sim N(0, I)$ (see in equation (8) where we introduce a zero-augmented log-normal distribution to address the presence of zeros in the data). Then the covariance structure of $\ln \tilde{\epsilon}_{i,t}$ is $(A_0 + I)(A_0 + I)'$. Recall that A_0 is lower triangular with zeros on its diagonal, so that $(A_0 + I)$ is a typical Cholesky factorization lower triangular matrix with unit diagonal elements.

It is important to point out that the first order multivariate system here can be marginalized to yield univariate representations for the elements of $Y_{i,t}$. To see this, we note that from equation (4), we have $\Psi_{i,t} = (I - BL)^{-1} [\mathbf{w} + A_1 \ln(Y_{i-1,t}) + C\mathbf{z}_{i-1,t}]$. Taking the log transformation of equation (3) yields the linear system:

$$\begin{aligned}
\ln(Y_{i,t}) &= (I - BL)^{-1} [\mathbf{w} + A_1 \ln(Y_{i-1,t}) + C\mathbf{z}_{i-1,t}] + \ln \tilde{\epsilon}_{i,t} \\
&= \mathbf{w} + (A_1 + B) \ln(Y_{i-1,t}) + C\mathbf{z}_{i-1,t} + (I - BL) \ln \tilde{\epsilon}_{i,t} \\
&= \mathbf{w} + (A_1 + B) \ln(Y_{i-1,t}) + C\mathbf{z}_{i-1,t} + (I - BL)(A_0 + I) \ln(\epsilon_{i,t}) \\
&= \mathbf{w} + (A_1 + B) \ln(Y_{i-1,t}) + C\mathbf{z}_{i-1,t} + (I - BL)e_{i,t}
\end{aligned} \tag{6}$$

where $e_{i,t} \sim N(0, (A_0 + I)(A_0 + I)')$. Setting $C = 0$ in the above set of equations yields a constrained VARMA(1,1) model for $\ln(Y_{i,t})$. Zellner and Palm (1974) show that the elements of vector processes can be marginalized to yield univariate ARMA representations. In particular, when applied to the system of equations in (6) we obtain (again setting $C = 0$):

$$\det(I - (A_1 + B)L) \ln(Y_{i,t}) = \tilde{\mathbf{w}} + \text{adj}(I - (A_1 + B)L)(I - BL)e_{i,t} \tag{7}$$

where for any matrix polynomial $D(L)$, $\det(D(L))$ is the scalar determinant finite polynomial in L and $\text{adj}(D(L))$ is the adjugate of $D(L)$.¹³ As noted by Zellner and Palm (1974), and more recently by Hecq, Laurent, and Palm (2016) in the context of multivariate GARCH models, the above implies that univariate models for the elements of $Y_{i,t} = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$ will be finite

¹³We use a generic intercept $\tilde{\mathbf{w}}$ and omit the explicit formula obtained from the product of $\text{adj}(I - (A_1 + B)L)$ with \mathbf{w} .

order ARMA and the parameters of the autoregressive part of each equation will usually be the same. As $Y_{i,t}$ is 3-dimensional and the polynomial $(I - (A_1 + B)L)$ is first degree we can be more specific, namely all marginal univariate processes will involve a third order autoregressive part. We can also say something about the MA term $\text{adj}(I - (A_1 + B)L)(I - BL)e_{i,t}$, namely the cofactor matrix determining $\text{adj}(I - (A_1 + B)L)$ contains quadratic polynomials in L , and therefore the MA terms will have maximal order 3.¹⁴ Therefore, we expect the marginalized univariate representation to be ARMA(3,3). However, in practice, the lag orders will be smaller because of roots cancellation, rank restrictions on the matrix $(A_1 + B)$, sparsity of that same matrix, and finite sample issues in the estimation of the univariate models. Thus, we mainly focus on lower order specifications in our analysis, but recognize that a higher order on the univariate processes might be warranted.

3.3 Estimation

As previously discussed, we choose to work with the lognormal distribution because it is convenient to describe some of the stochastic process features of the vector $Y_{i,t}$ described by the dynamic model appearing in equation (3). The empirical data, however, features many zero entries and we therefore consider the density function of the zero-augmented log-normal distribution for the errors, namely:

$$f(x) = (1 - p^+) \delta(x) + p^+ \frac{1}{xs\sqrt{2\pi}} \exp\left(-\frac{1}{2} \left(\frac{\ln(x) - m}{s}\right)^2\right) \mathbf{1}_{x>0} \quad (8)$$

where p^+ is the probability of strictly positive values, $\delta(x)$ is the probability mass at $x = 0$, and s and m are the scale and location parameters of the log-normal distribution. Requiring that the distribution has unit mean implies a restriction on m , that is, $m = -\frac{s^2}{2} + \ln\left(\frac{1}{p^+}\right)$, which eliminates the need to estimate m .

The model parameters include: \mathbf{w} (3×1 vector of constants), A_0 , A_1 , B (3×3 matrices of

¹⁴This is because the MA term involves the product of quadratic and linear polynomials in L , i.e., $\text{adj}(I - (A_1 + B)L)(I - BL)$.

dynamic coefficients), C ($3 \times K$ matrix of coefficients of covariates \mathbf{z} , if any), and the scale parameter s for the log normal distribution.¹⁵ The model is estimated via (quasi) maximum likelihood, where the density function for each innovation term $\epsilon_{i,t}^v$, $\epsilon_{i,t}^\sigma$, and $\epsilon_{i,t}^d$ is as specified in equation (8), using data for the whole sample. In all specifications, we include a dummy variable for the first observation of each day to absorb any overnight (or weekend) effects.

Lastly, it is important to account for the deterministic intraday patterns documented in Figure 2. Prior work either adds the intraday pattern as model parameters to estimate, or does so non-parametrically through an empirical estimator. We choose the nonparametric route and use the empirical mean for each five-minute interval as the diurnal adjustment factor for that interval. We take special care in accounting for the variation in our modeled variables over time, by doing the diurnal adjustment on a rolling-window basis, with a window size of 250 days that is sufficiently long to estimate the empirical mean reliably, while not too long as to make the adjustment factor stale. That is, for a raw variable $u_{i,t}$, the diurnal adjustment factor for interval i ($i = 1, 2, 3, \dots, 126$) is:

$$s_i = \frac{1}{250} \sum_{j=1}^{250} u_{i,t-j}. \quad (9)$$

Therefore, the diurnally-adjusted series is $\tilde{u}_{i,t} = u_{i,t}/s_i$.

The use of the empirical mean in computing the diurnal adjustment factors is common in the literature (e.g., Bauwens and Giot, 2000). One caveat is that the conditional variances $V(u_{i,t-j})$ and $V(u_{i,t-k})$ might not be the same for $k \neq j$ while the mean gives the same weight to all the u 's, and the mean can be contaminated by large outliers. That our data cleaning process has carefully removed data errors and winzorized outliers at the top 0.01%, and that we use rolling window estimation of the adjustment factors (which makes the assumption of constant conditional variance more reasonable), help ease this concern. While the median has been proposed as a jump-robust estimator of intraday periodicity, we run into the problem

¹⁵We use the empirical probability of zero, so it is not among the estimated parameters.

of having a zero median over some rolling windows, especially with regard to volatility at certain quiet times of day during low volatility periods. We compare the intraday pattern based on the mean with that based on the median over the whole sample period and find them to be very similar (if anything, the mean-based intraday pattern is smoother while the median-based intraday pattern looks slightly jumpy). This provides further assurance to the use of the mean in diurnally adjusting the data.

We estimate the model for all the three Treasury notes. For brevity of exposition, we present estimation results for the 10-year note, but results for the other two securities are largely similar. We show the histogram of the diurnally-adjusted depth, volatility, and volume of the 10-year note in Figure 3 to provide support for our modeling choices laid out above.

4 Results

4.1 Baseline results: Univariate dynamics

We start with a simple specification to focus our inquiry first on the essential dynamics of each of the modeled variables. That is, we estimate a version of the model specified in equation (3) in which A_0 is zero and both A_1 and B are diagonal, equivalent to estimating univariate Log-ACD₁(1,1) model for each variable. We label this specification as “Baseline” and report the parameter estimates in Table 2. All three variables exhibit a high degree of persistence: the sum of α and β are roughly 0.99 for depth and volatility, and about 0.93 for trading volume. Interestingly, both depth and volume have a much higher degree of short-term persistence (higher α) than volatility, indicating that they are more reactive to prevailing market conditions (or have shorter memory) than volatility. This corroborates our earlier observation that participants in this market have more degree of freedom with respect to quantities than to prices.

4.2 The joint dynamics of depth, volatility, and volume

Next, we estimate two specifications of the model for the joint dynamics of depth, volatility, and volume: 1) without contemporaneous interdependence (i.e., $A_0 = \mathbf{0}$), which we label as “Model 1”, and 2) with contemporaneous interdependence (i.e., $A_0 \neq \mathbf{0}$), which we label as “Model 2”. In both models, A_1 is a full 3×3 matrix of coefficients on lagged variables and B is a diagonal 3×3 matrix.

Model 1: Lagged interdependence

Under the Model 1 specification, the three market variables jointly affect one another with a lag. The parameter estimates for this model are reported in Table 3. The most important finding from this table is that, when accounting for the (lagged) interdependencies among the variables, the dynamics of depth and volume remain largely similar to their baseline dynamics, while those of volatility differ. More specifically, volatility is less persistent once we account for available depth and the level of trading activity in the market. The long-term persistence coefficient (β) is now 0.835, compared to the baseline estimate of 0.985, and the coefficient on lagged volatility (α) turns from 0.005 to -0.011. In other words, the importance of past volatility in predicting future volatility lessens when liquidity supply and demand are also considered. Lower depth and higher trading volume both predict a higher level of volatility in the next five-minute interval, as expected.

Model 2: Contemporaneous interdependence

Model 2 captures the contemporaneous error correlation structure that runs from trading volume to volatility to depth as previously discussed and specified in equation (3). This model addresses the possibility that the interdependencies among the variables occur at a faster speed than a lag of five-minute length would suggest. Table 4 contains the estimates for this model. We continue to observe that the dynamics of depth do not deviate much from earlier specifications, even after taking into account possible contemporaneous effects of volume and

volatility.¹⁶ What is new, and intuitive, in this specification is that contemporaneous shocks in volatility and trading demand both reduce depth. This suggests that liquidity providers in the market react quickly to changes in volatility and trading demand, a finding that resonates well with this Treasury interdealer market being populated by sophisticated participants.

Another interesting result in Table 4 is that the contemporaneous volume shock has a positive, and stronger, impact on volatility than past volume, as reflected in the coefficients on $\ln(\epsilon_{i,t}^v)$ (0.174) and $\ln(v_{i-1,t})$ (0.039) respectively. This suggests that the increase in volatility brought about by trading volume reflects increased price discovery as information gets revealed through trades, and that this price discovery occurs instantaneously with the order flow shock and continues into the next interval. Another source of volatility, due to the lack of liquidity to absorb shocks, is also evident by a negative and significant coefficient of $\ln(d_{i-1,t})$ (-0.188). To see the economic importance of these two sources of volatility, a quick calculation based on the coefficients reported in this table implies that a 25% increase in trading volume is accompanied by a 3.96% immediate increase in volatility and 0.87% increase in volatility in the next five-minute interval, whereas a 25% decrease in market depth at the beginning of the interval predicts roughly a 5.56% increase in subsequent volatility.

Perhaps the most striking result from this table is the significant drop in the persistence of volatility after incorporating the contemporaneous effects of trade flow shock, to 0.207 from 0.835 in Model 1. It appears that volatility clustering is due in large part to the clustering in trade flow, which in turn arises due to information arrival.

Model evaluation

To judge the performance of the three models estimated thus far in capturing the dynamic behavior of depth, volatility, and trading volume, we plot the autocorrelation function (ACF) of the residuals from the models against the ACF of the raw data in Figure 4. Two key observations are in order. First, the residuals exhibit much lower ACFs, indicating that

¹⁶The dynamics of volume do not change between Model 1 and Model 2 because there is no contemporaneous impact of shocks in depth and volatility on trading volume.

the models have done a reasonable job of explaining the dynamics of the modeled variables. Secondly, all three models, when fitted on depth and trading volume, do not result in substantially different ACFs of the residuals. This is consistent with the results discussed above that depth and trading volume dynamics do not seem to be significantly affected by other relevant market variables. However, for volatility, there is a clear improvement when other market variables are included in explaining its dynamics. Taken together, the results indicate that liquidity supply and demand have non-trivial predictive power over volatility, while the reverse is generally not supported by the data.

Another diagnostic exercise, reported in Table 5, is to examine the correlations among the three variables in the raw data and model residuals to assess whether the model specifications satisfactorily account for the joint dynamics of the variables. In the raw data, depth and volatility are negatively correlated with a correlation coefficient of -0.266, while volume and volatility are positively and highly correlated at 0.562. After fitting the three models, the residuals exhibit substantially lower correlation, especially under Model 2 which accounts for contemporaneous dependencies among the variables. The remaining pair, depth and trading volume, has a much lower degree of correlation of 0.081 in the raw data, and so the reduction in the residuals' correlation is much less drastic: the correlation of the residuals under the three modeling choices is in the 0.042–0.048 range. In sum, the much lower correlations observed in the residuals lend support to our multivariate specifications' ability to capture the joint dynamics of depth, volatility, and trading activity.

Higher order representations of marginalized univariate processes

It is important to keep in mind that the univariate representations of the elements of $Y_{i,t}$ can be of higher order than the Log-ACD₁(1,1) estimated here (see Zellner and Palm, 1974, and our detailed discussion in subsection 3.2). If so, restricting the lag order of each equation to (1,1) might result in the persistence coefficient estimate (β) being increased to capture the additional persistence otherwise accrued to older lags. To explore this issue, we extend

all three specifications (i.e., Baseline, Model 1, and Model 2) to additional lags on ψ^v , ψ^σ , and ψ^d up to the maximal order 3, and report the results in Table 6 in Panels A, B, and C respectively. To facilitate comparison, we reproduce the corresponding estimates for the specification with $q = 1$. We also report the log likelihood of each specification at the bottom of each panel.

The table shows that as more lags are included, the persistence spreads out instead of concentrating on lag 1. However, even with longer lag structures, the previous key finding that volatility is less persistent once we account for the effects of depth and trading level in the market continues to hold. As we move from Panel A (no effect of depth and volume) to Panel B (lagged effects of depth and volume) to Panel C (lagged effects of depth and both lagged and contemporaneous effects of trading volume), the persistence coefficients of volatility decline. Another important observation from this table is that the improvement in model log likelihood for volatility going from the Baseline model to Model 1 and then to Model 2 is most noticeable, compared to that for depth and trading volume.

4.3 Dynamics during the crisis

We now turn to test whether the dynamics of liquidity and volatility change during the crisis period. As the time series graphs in Figure 1 show, market depth is severely impaired during the crisis. The lower level of liquidity throughout the crisis is somewhat accounted for in our diurnal adjustment of the data using rolling-window averages. Here, we consider the effect of the crisis on the short-term persistence of the variables. That is, does the market update its expectation of liquidity and volatility based on prevailing market conditions any differently during the crisis as compared to non-crisis times? We implement this test by adding to the representation specified in equation (3) an interaction term of a crisis dummy with the lagged modeled variables as follows:

$$\Psi_{i,t} = \mathbf{w} + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t}, \quad (10)$$

where $\mathbb{1}_{crisis}$ is a dummy variable equal to 1 on days during the crisis period (August 9, 2007 – June 30, 2009) and 0 otherwise. A_0 , A_1 , and B are as in Model 2, and A_c is a 3×3 diagonal matrix capturing the incremental effects of the crisis on the dynamics of the system. If the crisis alters the dynamics of liquidity and volatility, we should find the coefficients on the diagonal of A_c to be significantly different from 0. As reported in Table 7, all coefficients on the diagonal of A_c are significant, and importantly, positive. This indicates that the short-term persistence of liquidity and volatility increases during the crisis, suggesting that during such stressful times the market places a greater weight on current market conditions in updating expectations. We also estimate several other specifications (not reported) and find that the crisis effects remain robust.

4.4 Announcement effects

The prior literature on the Treasury market’s response to economic news, such as Fleming and Remolona (1999), Balduzzi, Elton, and Green (2001), and Green (2004), has documented strong patterns of liquidity and volatility around economic announcements. Our earlier descriptive analysis also shows systematic patterns immediately before, during, and after these announcements, consistent with the evidence in the literature. However, most of the existing evidence pertains to the announcement effects on the level of market liquidity, volatility, and trading volume. We add to this literature by investigating how announcements affect the dynamics of these important market variables. This can provide insights into how market expectations get updated around announcement times.

Similar to the test of crisis effects discussed earlier, we estimate a specification that allows the short-term persistence of the modeled variables to differ before, during, and after announcements:

$$\Psi_{i,t} = \mathbf{w} + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis} + D_1 \mathbb{1}_{pre} + D_2 \mathbb{1}_{news} + D_3 \mathbb{1}_{post}) \ln(\mathbf{Y}_{i-1,t}) + B \Psi_{i-1,t}, \quad (11)$$

where $\mathbb{1}_{pre}$, $\mathbb{1}_{news}$, and $\mathbb{1}_{post}$ are the dummies for the five-minute interval before, during, and after an announcement respectively.¹⁷ As before, in order to keep the model manageable and focus our attention on the effects of announcements on each variable’s short-term persistence, we choose to estimate the model with D_1 , D_2 , and D_3 being diagonal. These coefficients indicate whether the importance of immediate past market information in updating expectations changes around announcement times. We report our estimates of D_1 , D_2 , and D_3 in Table 8. Appendix A contains the list of announcements considered.

The table provides strong evidence that the dynamics of market depth and volatility differ around announcement times. For market depth, for example, there is more intense updating of expectations in response to prevailing market conditions in the periods before and after an announcement. However, in the interval during which an announcement is released, the short-term persistence drops by 0.136, indicating a disconnect in the level of liquidity before and at the announcement. This suggests that, immediately after an announcement, the most recent pre-announcement liquidity level has become stale and less important in shaping up the expected level of liquidity in the next interval. On the other hand, there is always increased volatility clustering before and after announcements. The uncertainty with respect to an upcoming announcement, and the subsequent short period that follows when the market “digests” the announcement, could explain this increase in short-term volatility persistence.

It is interesting to note that increased short-term persistence of trading volume is observed only after, not before or during announcement intervals. This is consistent with market participants continuing to trade on the news into the next five-minute interval after the announcement. In the pre-announcement and announcement intervals, lagged volume reflects pre-news information and thus is less important to volume dynamics after the announcement, explaining the reduced short-term persistence in the pre-announcement and announcement windows.

¹⁷Scheduled announcements are assumed to be released during the interval starting with their release time (so that 8:30 announcements are deemed to have been released during the 8:30-8:35 interval, for example).

Having shown evidence of significant effects of announcements on the dynamics of liquidity and volatility, we now turn to test whether the direction of the announcements matters for these variables. We accomplish this with a simpler specification in which the announcement related dummies enter the log conditional mean equations directly rather than through their interaction with $Y_{i-1,t}$, together with indicators for whether an announcement delivers negative news:

$$\begin{aligned} \Psi_{i,t} = & \mathbf{w} + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t} \\ & + D_1 \mathbb{1}_{pre} + D_2 \mathbb{1}_{news} + D_2^- \mathbb{1}_{news}^- + D_3 \mathbb{1}_{post} + D_3^- \mathbb{1}_{post}^-, \end{aligned} \quad (12)$$

where $\mathbb{1}_{news}^-$ and $\mathbb{1}_{post}^-$ indicate announcement and post-announcement time intervals associated with negative news. A standard approach in the literature is to classify the direction of an announcement based on the deviation of the announced statistics from their professional forecasts. However, it is often the case that multiple statistics are released at the same time and they do not always follow the same direction with respect to the professional forecasts (e.g., some announced statistics are lower, while others are higher, than forecast). Therefore, classifying the aggregate direction of such announcements requires some judgment as to the importance assigned to each statistic in the bundle. We therefore take a different approach and aggregate the direction of an announcement (including announcements in which multiple economic indicators are released) by the resulting price change in the 10-year note over the five-minute interval enclosing the announcement. We denote price declines over such intervals as negative, or “bad news” (but make no judgment as to whether such price declines are good or bad for the economy).

In addition to testing for the importance of the news’ direction, we also test whether the magnitude of an announcement’s surprise matters. To measure the surprise component of each announcement (which again can bundle multiple macroeconomic statistics), we fit an AR(2) model on five-minute returns and use the innovations of returns over announcement

intervals as our measure of announcement surprises, denoted by NS . Thus, the alternative specification is:

$$\begin{aligned} \Psi_{i,t} = & \mathbf{w} + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t} \\ & + D_1 \mathbb{1}_{pre} + (D_2 + D_2^* \times NS) \mathbb{1}_{news} + (D_3 + D_3^* \times NS) \mathbb{1}_{post} \end{aligned} \quad (13)$$

We report the effects of announcement direction and surprise as specified in equations (12) and (13) in Table 9, Panels B and C. Panel A is reserved to show the announcement effects, without consideration to announcement direction nor surprise, on the level of market liquidity and volatility. This is different from the announcement effects on the short-term persistence of the modeled variables reported in Table 8. Here, from Panel A, we can see that the level of depth is lower before, but higher during and after announcements. This is consistent with our descriptive analysis earlier, and reflects the pattern of market participants withdrawing market liquidity immediately before an announcement, but quickly returning once the news has arrived. The pre-announcement drop in market depth is accompanied by increased volatility, which continues into the announcement and especially the post-announcement intervals. Meanwhile, trading volume increases both before and during announcement intervals, but the surge appears to fade already in the post-announcement intervals. Again, we continue to see evidence of trading volume leading volatility.

Interestingly, as Panel B shows, evidence for asymmetric effects of news is lacking, as indicated by the insignificance of most D_2^- and D_3^- coefficients implying that whether an announcement is negative does not have significant effects on the level of market depth or trading activity. The only significant coefficient here is D_2^- of volatility, indicating that bad news seems to increase volatility upon release.

The inclusion of announcement surprises in the model, shown in Panel C, appears more fruitful, at least for market depth and volatility. We continue to find that bad news (negative surprise) is associated with increased volatility, although under this specification the increase

occurs slightly later, in the post announcement interval. For depth, the speed of liquidity returning to the market after announcements, after having been withdrawn right before announcements, seems to depend on the news surprise. A negative surprise hastens the return of liquidity immediately within the announcement interval, whereas a positive surprise is associated with lower depth during the announcement interval but significantly higher depth in the post-announcement window. It is puzzling to find that the effects of announcements on trading volume are not sensitive to the surprise component of the announcements.

5 Application: liquidity risk at high frequency

As noted earlier, one important motivation for our MEM-based modeling framework is that the framework provides for unambiguously non-negative predictions of the variables of interest. This allows us to study the entire distribution of all possible outcomes and construct measures of liquidity risk based on tail outcomes. These measures can be constructed at a high frequency, and therefore can be a valuable tool for market participants in monitoring market liquidity conditions and assessing the risk of sudden liquidity stress for their optimal trading and market making on an almost real time basis.

We introduce two liquidity risk measures. The first measure, which we term “liquidity-at-risk” and denote by $LR1$, is similar in concept to the value-at-risk (VaR) measure. Specifically, for each five-minute interval, based on the deterministic time-of-day level s_i and the log conditional expectation of liquidity $\psi_{i,t}^d$ obtained from a chosen model, we compute the level of market depth that would result if there were an illiquidity shock (equal to the fifth percentile of the distribution of liquidity shocks, denoted by ϵ^*):

$$LR1_{i,t} = s_i \exp(\psi_{i,t}^d) \epsilon^*. \tag{14}$$

We construct this liquidity risk measure based on model parameters estimated from the full data sample, but note that in practice, one can easily estimate the model parameters

using historical data up to the time of modeling. Based on the estimated log conditional expectation model, we can construct the conditional distribution of liquidity based on two possible choices for the error distribution. The first is the log normal distribution assumed in the estimation (whose scale parameter is among the model parameter estimates). The second is an empirical distribution based on the recent history of liquidity shocks (we use 126-interval \times 20-day history). The two alternatives give two separate *LR1* measures, but they are quite close to each other.

To formally compare the quality of the two measures, we report in Table 10 the hit rate *Hit%*, the percentage of times the actual depth level falls below *LR1*, together with the associated p-value of the simple t-test for $H_0 : Hit = 5\%$. We perform this test for the *LR1* measures computed from three alternative log conditional mean models: (i) Baseline (reported in Table 2), (ii) Model 1 (reported in Table 3), and (iii) Model 2 (reported in Table 4). Panel A is for the *LR1* measure based on the empirical distribution of liquidity shocks. Panel B is for the *LR1* measure based on the log-normal distribution whose scale parameter is estimated for each model. Comparing between panel A and panel B, we can see that the *LR1* measure based on the empirical distribution provides an accurate hit ratio, whereas the measure based on the assumed Lognormal distribution does not perform as well on this unconditional test. Comparing across columns, it seems that all three log conditional mean models deliver quite similar results, indicating that the unconditional property of the *LR1* measure is not highly sensitive to the modeling choice.¹⁸

We also entertain a second measure of liquidity risk, denoted by *LR2*, which is the probability of an unexpected illiquidity shock in the next interval. That is, given the prevailing market conditions, the measure indicates the probability that market liquidity falls

¹⁸It is important to point out that the desirable conditional property of no autocorrelation in the hit sequence computed from *LR1* can be back-tested using well-established tests, most notably the dynamic quantile tests developed by Engle and Manganelli (2004). We performed these tests (untabulated) on our *LR1* measure and found that the null hypothesis of no autocorrelation is rejected at standard significance levels. That said, rejection of the null is not surprising given our very large sample size ($T = 301,770$) and may speak more to significance levels for tests using large datasets than to the suitability of our model.

below a low level L^* (to be chosen):

$$\text{LR2}_{i,t} = \text{Prob} \left(\epsilon_{i,t} \leq \frac{L^*}{s_i \exp(\psi_{i,t}^d)} \right). \quad (15)$$

To illustrate this feature, we select the threshold L^* to be equal to the first percentile of the distribution of five-minute market depth over the whole sample period. For the 10-year Treasury note, this threshold is equal to \$2 million in the bid-ask average inside depth.

We plot in Figures 5 and 6 the time series of liquidity risk as measured by $LR1$ and $LR2$ respectively (based on the parameter estimates of Model 2 presented in subsection 4.2). Our liquidity measures are at the five-minute frequency, but for seeing the trend in liquidity risk over time, we aggregate them to the daily level by computing the daily averages. Both measures point to a spike in liquidity risk in late 2008. For example, the chance of observing a liquidity level lower than the historical first percentile elevates to nearly 30% percent for the 10-year note. At the same time, the liquidity-at-risk level approaches zero. There are some additional, albeit more modest, episodes of heightened liquidity risk in 2012 and 2013. Finally, after 2014, the probability of a large liquidity shock is low, even around the flash rally in October 2014.

Another way to utilize the model is to examine the conditional distribution of liquidity and volatility and assess how their tail behaviors correlate. As previously discussed, the ability of our models to make non-negative forecasts becomes even more important if we are interested in not only the mean but also other quantiles of the distribution. Our models can hence deliver useful insights as to how often extreme illiquidity is accompanied by an extremely high volatility level. To illustrate, we rely on our estimates of the conditional density to compute the 5th percentile of liquidity and 95th percentile of volatility, representing the “5% worst case scenarios.” We then examine the correlation of these worst-case scenarios in liquidity with those in volatility in comparison with the usual correlation of liquidity and volatility in the raw data. The correlation between extremely low liquidity and extremely

high volatility is -0.431. Comparing to the correlation of these variables in the raw data, as well as in the residuals of various models (see Table 5), it is clear that extreme illiquidity and high volatility are much more correlated than their correlation in normal times suggests.

6 Conclusion

We study the dynamics of market liquidity, trading activity, and volatility in the Treasury market before, during, and after the 2007–09 financial crisis. We develop an econometric framework that is particularly suited to addressing empirical features of this important market. The framework is based on a multiplicative error structure, which offers important advantages that are highly suited to the high frequency modeling of the Treasury limit order book for which zero or low values of depth, volume and volatility are not uncommon. The distribution of the errors is a zero-augmented log-normal distribution capable of capturing zero occurrences. The conditional expectation is modeled with Bauwens and Giot (2000)’s Log-ACD₁ formulation, which can conveniently accommodate additional covariates without violating the non-negativity constraint on the conditional mean. With these features, our model integrates liquidity and volatility dynamics into a unified framework from which their dynamic interactions can be studied.

Our models provide several important results. First, market liquidity, trading activity, and volatility are highly persistent, with the first two variables exhibiting higher responsiveness to their own short-term fluctuations (higher short-term persistence) than volatility. They also appear to have non trivial explanatory power over volatility dynamics, whereas their own dynamics seem much less affected by volatility. Secondly, during crisis times, the dynamics of all three variables change in a similar fashion, with an increase in the short-term persistence reflecting a more active updating of new information. Furthermore, we find evidence of a significant change in the dynamics of liquidity and volatility around economic announcements. Volatility’s short-term persistence increases both before and after announcements, whereas

that of trading volume increases only after announcements reflecting continued trading on the news. Simpler specifications focusing on the level of liquidity and volatility around announcements (assuming non-changing dynamics) show that volatility tends to be higher following bad news.

Finally, an important advantage of our modeling framework is that we are able to develop useful measures of liquidity risk for market monitoring and study the tail behavior of the conditional distribution of liquidity and volatility on an almost real time basis. Our liquidity risk measures point to a spike in liquidity risk in late 2008, coinciding with the depth of the financial crisis, followed by several episodes of mildly increased liquidity risk in 2012 and 2013. After 2014, the risk of an extreme illiquidity event largely subsides. Importantly, we show that extreme illiquidity events are much more highly correlated with extremely high volatility episodes than the usual liquidity-volatility correlation would suggest.

References

- Adrian, T., M. Fleming, O. Shachar, and E. Vogt, 2017, “Market Liquidity After the Financial Crisis,” *Annual Review of Financial Economics*, 9, 43–83.
- Ahn, H. J., K. H. Bae, and K. Chan, 2001, “Limit Orders, Depth and Volatility: Evidence from the Stock Exchange of Hong Kong,” *Journal of Finance*, 56, 767–788.
- Allen, D., F. Chan, M. McAleer, and S. Peiris, 2008, “Finite sample properties of the QMLE for the Log-ACD model: Application to Australian stocks,” *Journal of Econometrics*, 147, 163–185.
- Balduzzi, P., E. J. Elton, and T. C. Green, 2001, “Economic News and Bond Prices: Evidence from the U.S. Treasury Market,” *Journal of Financial and Quantitative Analysis*, 36, 523–543.
- Bauwens, L., and P. Giot, 2000, “The Logarithmic ACD Model: An Application to the Bid-Ask Quote Process of Three NYSE Stocks,” *Annales d’Economie et de Statistique*, 60, 117–149.
- Biais, B., P. Hillion, and C. Spatt, 1995, “An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse,” *Journal of Finance*, 50, 1655–1689.
- Brandt, M., and K. Kavajecz, 2004, “Price Discovery in the U.S. Treasury Market: the Impact of Order Flow and Liquidity on the Yield Curve,” *Journal of Finance*, 59, 2623–2654.
- Cipollini, F., R. Engle, and G. Gallo, 2017, “Copula-based vMEM specifications versus alternatives: the case of trading activity,” Università di Firenze and New York University Working Paper.
- Danielsson, J., and R. Payne, 2012, “Liquidity Determination in an Order Driven Market,” *European Journal of Finance*, 18, 799–821.
- Engle, R., 2002, “New Frontiers for ARCH Models,” *Journal of Applied Econometrics*, 17, 425–446.
- Engle, R., and S. Manganelli, 2004, “CAViaR: Conditional autoregressive value at risk by regression quantiles,” *Journal of Business and Economic Statistics*, 22, 367–381.
- Engle, R., and J. Russell, 1998, “Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data,” *Econometrica*, 66, 1127–1162.
- Escribano, A., and G. Sucarrat, forthcoming, “Equation-by-equation estimation of multivariate periodic electricity price volatility,” *Energy Economics*.
- Fleckenstein, M., F. A. Longstaff, and H. Lustig, 2014, “The Tips-Treasury Puzzle,” *Journal of Finance*, 69, 2151–2197.
- Fleming, M., 2003, “Measuring Treasury Market Liquidity,” *Federal Reserve Bank of New York Economic Policy Review*, 9, 83–108.

- Fleming, M., 2016, “Is Treasury Market Liquidity Becoming More Concentrated?,” Liberty Street Economics Blog, February 2016.
- Fleming, M. J., B. Mizrach, and G. Nguyen, 2018, “The microstructure of a US Treasury ECN: The BrokerTec platform,” *Journal of Financial Markets*, 40, 2–22.
- Fleming, M. J., and M. Piazzesi, 2005, “Monetary Policy Tick-by-Tick,” Working Paper, Federal Reserve Bank of New York.
- Fleming, M. J., and E. M. Remolona, 1997, “What Moves the Bond Market?,” *Federal Reserve Bank of New York Economic Policy Review*, 3, 31–50.
- Fleming, M. J., and E. M. Remolona, 1999, “Price Formation and Liquidity in the U.S. Treasury Market: The Response to Public Information,” *Journal of Finance*, 54, 1901–1915.
- Francq, C., and G. Sucarrat, 2017, “An equation-by-equation estimator of a multivariate log-GARCH-X model of financial returns,” *Journal of Multivariate Analysis*, 153, 16–32.
- Goldreich, D., B. Hanke, and P. Nath, 2005, “The Price of Future Liquidity: Time-Varying Liquidity in the U.S. Treasury Market,” *Review of Finance*, 9, 1–32.
- Green, T. C., 2004, “Economic News and the Impact of Trading on Bond Prices,” *Journal of Finance*, 59, 1201–1233.
- Härdle, W. K., N. Hautsch, and A. Mihoci, 2012, “Modelling and forecasting liquidity supply using semiparametric factor dynamics,” *Journal of Empirical Finance*, 19, 610–625.
- Hautsch, N., and R. Huang, 2012, “The market impact of a limit order,” *Journal of Economic Dynamics and Control*, 36, 501–522.
- Hautsch, N., P. Malec, and M. Schienle, 2013, “Capturing the zero: A new class of zero-augmented distributions and multiplicative error processes,” *Journal of Financial Econometrics*, 12, 89–121.
- Hautsch, N., and M. Podolskij, 2013, “Preaveraging-based estimation of quadratic variation in the presence of noise and jumps: Theory, implementation, and empirical evidence,” *Journal of Business and Economic Statistics*, 31, 165–183.
- Hecq, A., S. Laurent, and F. C. Palm, 2016, “On the univariate representation of BEKK models with common factors,” *Journal of Time Series Econometrics*, 8, 91–113.
- Hu, G. X., J. Pan, and J. Wang, 2013, “Noise as information for illiquidity,” *Journal of Finance*, 68, 2341–2382.
- Huang, R. D., J. Cai, and X. Wang, 2002, “Information-Based Trading in the Treasury Note Interdealer Broker Market,” *Journal of Financial Intermediation*, 11, 269–296.
- Jacod, J., Y. Li, P. Mykland, M. Podolskij, and M. Vetter, 2009, “Microstructure noise in the continuous case: the pre-averaging approach,” *Stochastic Processes and Their Applications*, 119, 2249–2276.

- Jiang, G., I. Lo, and A. Verdelhan, 2011, “Information Shocks, Liquidity Shocks, Jumps, and Price Discovery: Evidence from the U.S. Treasury Market,” *Journal of Financial and Quantitative Analysis*, 46, 527–551.
- Johnson, T. C., 2008, “Volume, Liquidity and Liquidity Risk,” *Journal of Financial Economics*, 87, 388–417.
- Joint Staff Report, 2015, “The U.S. Treasury Market on October 15, 2014,” U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, U.S. Securities and Exchange Commission, and U.S. Commodity Futures Trading Commission.
- Kyle, A. S., 1985, “Continuous Auctions and Insider Trading,” *Econometrica*, 53, 1315–1336.
- Manganelli, S., 2005, “Duration, Volume and Volatility Impact of Trades,” *Journal of Financial Markets*, 8, 377–399.
- Mizrach, B., and C. Neely, 2007, “The Microstructure of the U.S. Treasury Market,” Federal Reserve Bank of St. Louis Working Paper No. 2007-052A.
- Musto, D., G. Nini, and K. Schwarz, 2018, “Notes on Bonds: Illiquidity Feedback During the Financial Crisis,” *Review of Financial Studies*, 31, 2983—3018.
- Næs, R., and J. Skjeltorp, 2006, “Order Book Characteristics and the Volume–Volatility Relation: Empirical Evidence from a Limit Order Market,” *Journal of Financial Markets*, 9, 408–432.
- Pasquariello, P., and C. Vega, 2007, “Informed and Strategic Order Flow in the Bond Markets,” *Review of Financial Studies*, 20, 1975—2019.
- Ranaldo, A., 2004, “Order Aggressiveness in Limit Order Book Markets,” *Journal of Financial Markets*, 7, 53–74.
- Sucarrat, G., and A. Escibano, 2017, “Estimation of log-GARCH models in the presence of zero returns,” *European Journal of Finance*, 24, 809–827.
- Zellner, A., and F. C. Palm, 1974, “Time series analysis and simultaneous equation econometric models,” *Journal of Econometrics*, 2, 17–54.

Table 1: Descriptive statistics of limit order book variables

This table provides descriptive statistics of limit order book variables at the five-minute frequency for the 2-, 5-, and 10-year notes. Market depth at *Bid 1* and *Bid 2-5* reflects the quantity of displayed depth (in \$ million of par) at the best bid and collectively at the second through fifth best bid respectively. Market depth at *Ask 1* and *Ask 2-5* are similarly defined for the ask side. *Volume* is the total quantity (in \$ million of par) traded at five-minute frequency. *Volatility* is five-minute realized volatility (annualized), computed from pre-averaged one-second log returns. Data is from BrokerTec and covers the 2006-2015 period.

	Market Depth				Volume	Volatility
	Bid 1	Bid 2-5	Ask 1	Ask 2-5		
Panel A: 2-Year Note						
Mean	402	1547	386	1514	197	0.016
Std dev	435	1017	411	1001	302	0.019
Min	0	0	0	0	0	0.000
Q1	93	721	90	697	4	0.000
Median	265	1396	256	1363	79	0.015
Q3	566	2282	545	2252	273	0.029
Max	12196	17808	11160	11477	6178	1.481
Auto corr (1)	0.777	0.965	0.820	0.966	0.509	0.639
Panel B: 5-Year Note						
Mean	48	308	48	306	277	0.044
Std dev	48	206	50	207	274	0.031
Min	0	0	0	0	0	0.000
Q1	13	153	13	152	95	0.027
Median	35	276	35	274	207	0.040
Q3	67	420	67	419	373	0.056
Max	1019	2532	7383	5620	7218	3.099
Auto corr (1)	0.484	0.944	0.445	0.940	0.660	0.608
Panel C: 10-Year Note						
Mean	40	247	39	245	241	0.084
Std dev	40	159	40	160	235	0.047
Min	0	0	0	0	0	0.000
Q1	12	133	12	132	87	0.057
Median	29	219	29	217	181	0.081
Q3	55	331	54	328	322	0.106
Max	6843	2273	6481	5308	5911	3.468
Auto corr (1)	0.394	0.936	0.389	0.928	0.669	0.650

Table 2: Parameter estimates of baseline Log-ACD₁(1,1) model

This table reports the parameter estimates for the baseline Log-ACD₁(1,1) model for depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note. Each variable, generically denoted by $y_{i,t}$, is modeled as: $y_{i,t} = \exp(\psi_{i,t}) \epsilon_{i,t}$, with $\psi_{i,t} = \omega + \alpha \ln(y_{i-1,t}) + \beta \psi_{i-1,t}$. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. Significance at the 1%, 5%, and 10% levels are marked with ***, **, and * respectively.

	$\psi_{i,t}^d$		$\psi_{i,t}^\sigma$		$\psi_{i,t}^v$	
	Est.	s.e.	Est.	s.e.	Est.	s.e.
ω	0.012***	0.0001	0.002***	0.0000	0.036***	0.0005
α	0.125***	0.0005	0.005***	0.0000	0.110***	0.0010
β	0.866***	0.0005	0.985***	0.0001	0.821***	0.0011

Table 3: Joint dynamics of depth, volatility, and volume: lagged effects

This table reports the parameter estimates for the following multivariate model of depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note: $\mathbf{Y}_t = \exp(\Psi_{i,t}) \odot \epsilon_{i,t}$, with $\Psi_{i,t} = \omega + A_1 \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t}$, where $\mathbf{Y}_t = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$. A_1 is a full 3×3 matrix and B is a diagonal 3×3 matrix. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. Significance at the 1%, 5%, and 10% levels are marked with ***, **, and * respectively.

		$\psi_{i,t}^d$	$\psi_{i,t}^\sigma$	$\psi_{i,t}^v$
ω	Constant	0.012***	0.007***	0.034***
A_1	$\ln(d_{i-1,t})$	0.124***	-0.058***	-0.003***
	$\ln(\sigma_{i-1,t})$	0.001***	-0.011***	-0.018***
	$\ln(v_{i-1,t})$	-0.003***	0.044***	0.121***
B	$\psi_{i-1,t}^d$	0.868***		
	$\psi_{i-1,t}^\sigma$		0.835***	
	$\psi_{i-1,t}^v$			0.823***

Table 4: Joint dynamics of depth, volatility, and volume: contemporaneous effects

This table reports the parameter estimates for the following multivariate model of depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note: $\mathbf{Y}_t = \exp(\Psi_{i,t}) \odot \epsilon_{i,t}$, with $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + A_1 \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t}$, where $\mathbf{Y}_t = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$. A_0 is a lower triangular 3×3 matrix with zeros on its diagonal, A_1 is a full 3×3 matrix and B is a diagonal 3×3 matrix. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. Significance at the 1%, 5%, and 10% levels are marked with ***, **, and * respectively.

		$\psi_{i,t}^d$	$\psi_{i,t}^\sigma$	$\psi_{i,t}^v$
ω	Constant	0.011***	0.061***	0.034***
A_0	$\ln(\epsilon_{i,t}^d)$			
	$\ln(\epsilon_{i,t}^\sigma)$	-0.011***		
	$\ln(\epsilon_{i,t}^v)$	-0.005***	0.174***	
A_1	$\ln(d_{i-1,t})$	0.132***	-0.188***	-0.003***
	$\ln(\sigma_{i-1,t})$	0.021***	0.026***	-0.018***
	$\ln(v_{i-1,t})$	-0.004***	0.039***	0.121***
B	$\psi_{i-1,t}^d$	0.865***		
	$\psi_{i-1,t}^\sigma$		0.207***	
	$\psi_{i-1,t}^v$			0.823***

Table 5: Pearson correlations of depth, volatility, and volume before and after model fitting

This table reports the pairwise (Pearson) correlations among depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note and that among their residuals after fitting them with three dynamic models: Baseline (reported in Table 2), Model 1 (reported in Table 3), and Model 2 (reported in Table 4). Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period.

	Residuals			
	Raw Data	Baseline	Model 1	Model 2
$\rho(d_{i,t}, \sigma_{i,t})$	-0.266	-0.045	-0.062	-0.029
$\rho(d_{i,t}, v_{i,t})$	0.081	0.043	0.042	0.048
$\rho(\sigma_{i,t}, v_{i,t})$	0.562	0.372	0.450	0.147

Table 6: Higher order specifications of marginalized univariate processes

This table reports the parameter estimates for higher order univariate representations of each element in the multivariate model of depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note, specified as: $\mathbf{Y}_t = \exp(\Psi_{i,t}) \odot \epsilon_{i,t}$, with $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + A_1 \ln(\mathbf{Y}_{i-1,t}) + \sum_{j=1}^q B_j \Psi_{t-j}$, where $\mathbf{Y}_t = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$. A_0 is a lower triangular 3×3 matrix with zeros on its diagonal, A_1 is a full 3×3 matrix and B 's are diagonal 3×3 matrices. The last row of each panel reports the log likelihood for the given model. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. All estimates are significant at 1% levels except one marked with *ns*.

	$\psi_{i,t}^d$			$\psi_{i,t}^\sigma$			$\psi_{i,t}^v$		
	$q = 1$	$q = 2$	$q = 3$	$q = 1$	$q = 2$	$q = 3$	$q = 1$	$q = 2$	$q = 3$
Panel A: Baseline model ($A_0 = \mathbf{0}$, $A_1 = \text{diagonal}$)									
ω	0.012	0.016	0.017	0.002	0.004	0.006	0.036	0.043	0.044
α	0.125	0.163	0.175	0.005	0.010	0.014	0.110	0.131	0.136
β_1	0.866	0.434	0.504	0.985	0.180	0.195	0.821	0.506	0.569
β_2		0.392	0.033		0.793	0.125		0.284	0.012
β_3			0.278			0.642			0.204
llk ($\times 10^5$)	-1.4188	-1.4112	-1.4071	-1.8711	-1.8686	-1.8666	-2.7540	-2.7497	-2.7473
Panel B: Model 1 ($A_0 = \mathbf{0}$, $A_1 \neq \text{diagonal}$)									
ω	0.012	0.015	0.016	0.007	0.009	0.009	0.034	0.040	0.042
α	0.124	0.162	0.174	-0.011	-0.014	-0.015	0.121	0.146	0.152
β_1	0.868	0.435	0.505	0.835	0.453	0.509	0.823	0.498	0.561
β_2		0.394	0.033		0.346	0.048		0.293	0.014
β_3			0.279			0.236			0.212
llk ($\times 10^5$)	-1.4182	-1.4107	-1.4065	-1.6790	-1.6741	-1.6717	-2.7468	-2.7419	-2.7392
Panel C: Model 2 ($A_0 \neq \mathbf{0}$, $A_1 \neq \text{diagonal}$)									
ω	0.011	0.014	0.014	0.061	0.059	0.059	0.034	0.040	0.042
α	0.132	0.167	0.177	0.026	0.016	0.011	0.121	0.146	0.152
β_1	0.865	0.472	0.536	0.207	0.049	0.000 ^{ns}	0.823	0.498	0.561
β_2		0.356	0.020		0.276	0.198		0.293	0.014
β_3			0.262			0.187			0.212
llk ($\times 10^5$)	-1.4120	-1.4051	-1.4011	-1.5197	-1.4763	-1.4557	-2.7469	-2.7419	-2.7392

Table 7: Effects of the crisis on the dynamics of depth, volatility, and volume

This table reports the parameter estimates for the following multivariate model of depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note: $\mathbf{Y}_t = \exp(\Psi_{i,t}) \odot \epsilon_{i,t}$, with $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B \Psi_{i-1,t}$, where $\mathbf{Y}_t = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$ and $\mathbb{1}_{crisis}$ is a dummy variable equal to 1 on days during the crisis period (August 9, 2007 to June 30, 2009). A_0 is a lower triangular 3×3 matrix with zeros on its diagonal, A_1 is a full 3×3 matrix, and both B and A_c are diagonal 3×3 matrices. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. Significance at the 1%, 5%, and 10% levels are marked with ***, **, and * respectively.

		$\psi_{i,t}^d$	$\psi_{i,t}^\sigma$	$\psi_{i,t}^v$
ω	Constant	0.011***	0.057***	0.034***
A_0	$\ln(\epsilon_{i,t}^d)$			
	$\ln(\epsilon_{i,t}^\sigma)$	-0.010***		
	$\ln(\epsilon_{i,t}^v)$	-0.005***	0.171***	
A_1	$\ln(d_{i-1,t})$	0.132***	-0.183***	-0.006***
	$\ln(\sigma_{i-1,t})$	0.020***	0.020***	-0.017***
	$\ln(v_{i-1,t})$	-0.004***	0.035***	0.119***
B	$\psi_{i-1,t}^d$	0.861***		
	$\psi_{i-1,t}^\sigma$		0.214***	
	$\psi_{i-1,t}^v$			0.819***
A_c	Crisis $\times \ln(d_{i-1,t})$	0.007***		
	Crisis $\times \ln(\sigma_{i-1,t})$		0.116***	
	Crisis $\times \ln(v_{i-1,t})$			0.017***

Table 8: Announcement effects on the dynamics of depth, volatility, and volume

This table reports the parameter estimates for the following multivariate model of depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note: $\mathbf{Y}_t = \exp(\Psi_{i,t}) \odot \epsilon_{i,t}$, with $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis} + D_1 \mathbb{1}_{pre} + D_2 \mathbb{1}_{news} + D_3 \mathbb{1}_{post}) \ln(\mathbf{Y}_{i-1,t}) + B \Psi_{i-1,t}$, where $\mathbf{Y}_t = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$, $\mathbb{1}_{crisis}$ is a dummy variable equal to 1 on days during the crisis period (August 9, 2007 to June 30, 2009), and $\mathbb{1}_{pre}$, $\mathbb{1}_{news}$, and $\mathbb{1}_{post}$ are the dummies for the five-minute interval before, during, and after an announcement respectively (the list of announcements is in Appendix A). A_0 is a lower triangular 3×3 matrix with zeros on its diagonal, A_1 is a full 3×3 matrix, and B , A_c , D_1 , D_2 , and D_3 are diagonal 3×3 matrices. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. Significance at the 1%, 5%, and 10% levels are marked with ***, **, and * respectively.

		$\psi_{i,t}^d$	$\psi_{i,t}^\sigma$	$\psi_{i,t}^v$
D_1	Prenews $\times \ln(d_{i-1,t})$	0.136***		
	Prenews $\times \ln(\sigma_{i-1,t})$		0.027**	
	Prenews $\times \ln(v_{i-1,t})$			-0.031*
D_2	News $\times \ln(d_{i-1,t})$	-0.136***		
	News $\times \ln(\sigma_{i-1,t})$		0.013	
	News $\times \ln(v_{i-1,t})$			-0.041***
D_3	Postnews $\times \ln(d_{i-1,t})$	0.019**		
	Postnews $\times \ln(\sigma_{i-1,t})$		0.120***	
	Postnews $\times \ln(v_{i-1,t})$			0.247***

Table 9: Announcement effects on the level of depth, volatility, and volume

This table reports the parameter estimates for three specifications with announcement effects on the dynamics of depth $d_{i,t}$, volatility $\sigma_{i,t}$, and volume $v_{i,t}$ of the 10-year Treasury note. The specification in Panel A is: $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t} + D_1 \mathbb{1}_{pre} + D_2 \mathbb{1}_{news} + D_3 \mathbb{1}_{post}$, where $\mathbf{Y}_t = [v_{i,t}, \sigma_{i,t}, d_{i,t}]'$, $\mathbb{1}_{crisis}$ is a dummy variable equal to 1 on days during the crisis period (August 9, 2007 to June 30, 2009), and $\mathbb{1}_{pre}$, $\mathbb{1}_{news}$, and $\mathbb{1}_{post}$ are the dummies for the five-minute interval before, during, and after an announcement respectively (the list of announcements is in Appendix A). The specification in Panel B is: $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t} + D_1 \mathbb{1}_{pre} + D_2 \mathbb{1}_{news} + D_2^- \mathbb{1}_{news}^- + D_3 \mathbb{1}_{post} + D_3^- \mathbb{1}_{post}^-$, where $\mathbb{1}_{news}^-$ indicates announcement time intervals over which price declines (bad news), and $\mathbb{1}_{post}^-$ indicates the intervals immediately following bad news announcements. The specification in Panel C is: $\Psi_{i,t} = \omega + A_0 \ln(\epsilon_{i,t}) + (A_1 + A_c \mathbb{1}_{crisis}) \ln(\mathbf{Y}_{i-1,t}) + B\Psi_{i-1,t} + D_1 \mathbb{1}_{pre} + (D_2 + D_2^* NS) \mathbb{1}_{news} + (D_3 + D_3^* NS) \mathbb{1}_{post}$, where NS is the news surprise measured by the unexpected price change on the 10-year note over the corresponding announcement time interval (based on an AR(2) model of five-minute returns). A_0 is a lower triangular 3×3 matrix with zeros on its diagonal, A_1 is a full 3×3 matrix, and B , A_c , and all the D 's are diagonal 3×3 matrices. Modeled variables are diurnally adjusted using the 250-day rolling window average by time of day before estimation. Estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period. Significance at the 1%, 5%, and 10% levels are marked with ***, **, and * respectively.

		$\psi_{i,t}^d$	$\psi_{i,t}^\sigma$	$\psi_{i,t}^v$
Panel A: Symmetric Model				
D_1	Prenews	-0.596***	0.027***	0.132***
D_2	News	0.508***	0.014**	0.118***
D_3	Postnews	0.041***	0.035***	-0.051***
Panel B: Asymmetric Model				
D_1	Prenews	-0.596***	0.027***	0.132***
D_2	News	0.492	-0.008	0.062
D_2^-	Bad News	0.037	0.048***	0.127
D_3	Postnews	0.042	0.030***	-0.024
D_3^-	Post Bad News	-0.003	0.013	-0.060
Panel C: News Surprise Model				
D_1	Prenews	-0.596***	0.027***	0.132***
D_2	News	0.508***	0.014**	0.117***
D_2^*	News x Surprise	-0.154***	-0.124	-0.211
D_3	Postnews	0.041***	0.035***	-0.051***
D_3^*	Postnews x Surprise	9.962***	-7.637***	0.810

Table 10: Backtesting Liquidity-at-Risk Measure

This table reports the results of backtesting the liquidity-at-risk measure $LR1$ developed for the 10-year Treasury note per equation (14) in the text. The measure is computed using each of the three dynamic models of market depth $d_{i,t}$: Baseline (reported in Table 2), Model 1 (reported in Table 3), and Model 2 (reported in Table 4). Hit % is the percentage of times the actual liquidity level is below $LR1$. The p-value for the t-test for $H_0: Hit = 0.05$ is reported underneath. Panel A is for the $LR1$ measure based on the empirical distribution of liquidity shocks. Panel B is for the $LR1$ measure based on the log-normal distribution whose scale parameter is estimated for each model. Model estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period.

Statistic	Baseline	Model 1	Model 2
Panel A: Empirical Distribution			
Hit %	4.991	5.003	4.988
pval: Hit = 5%	0.8312	0.9301	0.7794
Panel B: Lognormal Distribution			
Hit %	4.463	4.461	4.464
pval: Hit = 5%	0	0	0

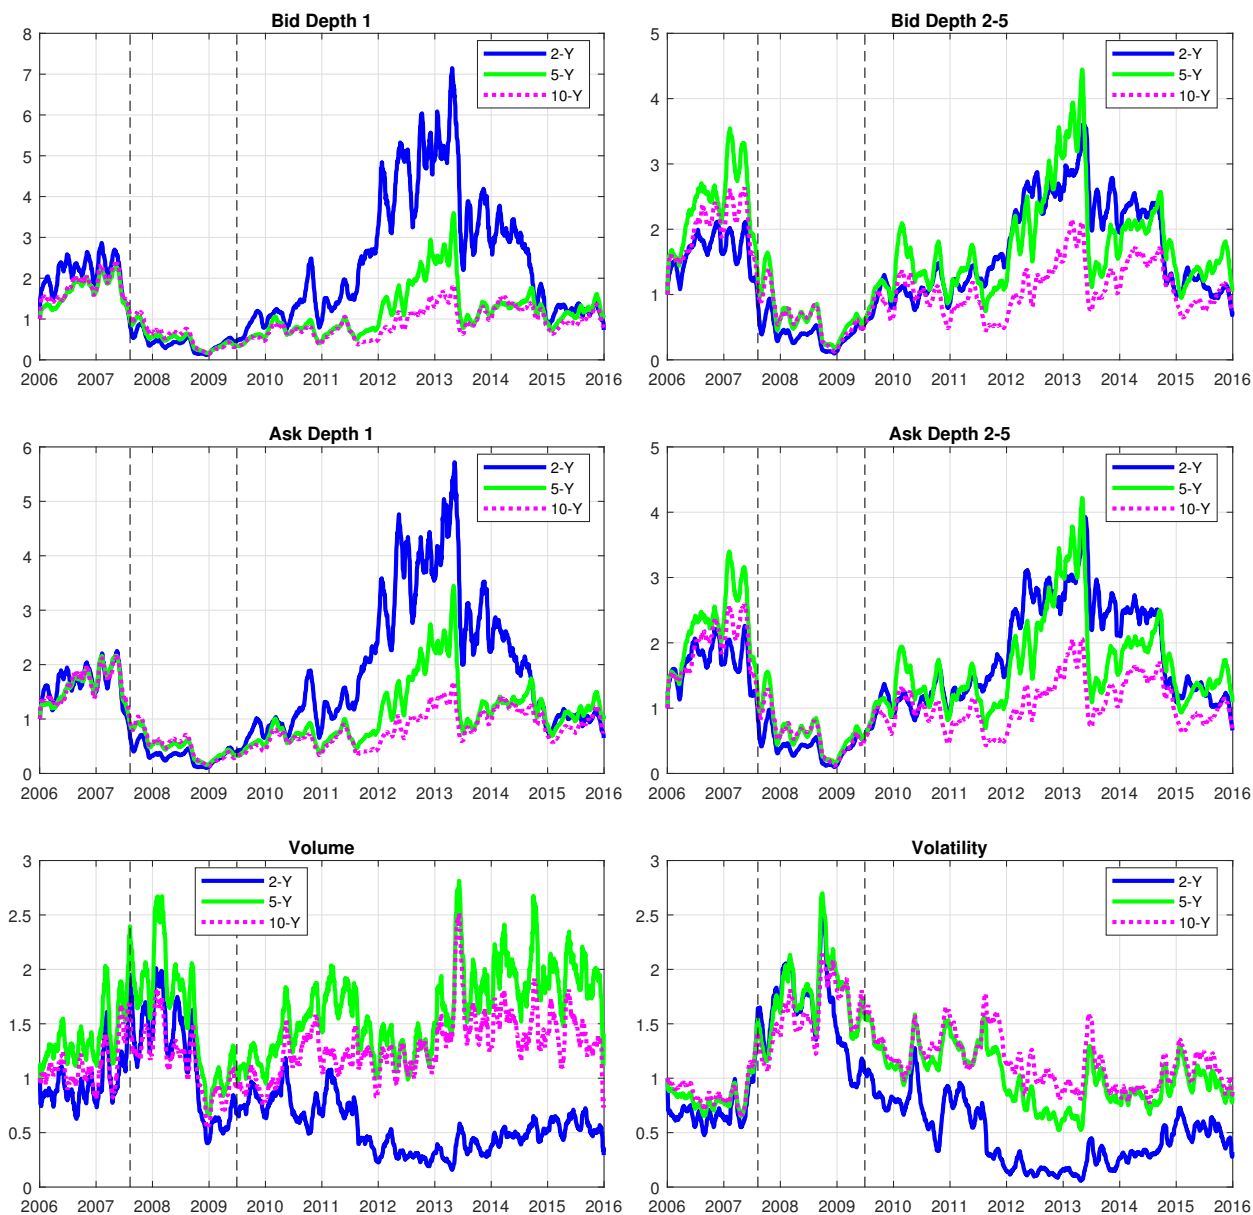


Figure 1: Variation of limit order book variables over time

This figure shows the time series variation of limit order book variables relative to their level at the beginning of 2006 for the 2-, 5-, and 10-year notes. The time series are based on the average of each variable on each day. The daily time series are smoothed using a 21-day moving average window. Vertical dashed lines (from left to right) mark the start (August 9, 2007) and end (June 30, 2009) of the crisis. Data is from BrokerTec and covers the 2006-2015 period.

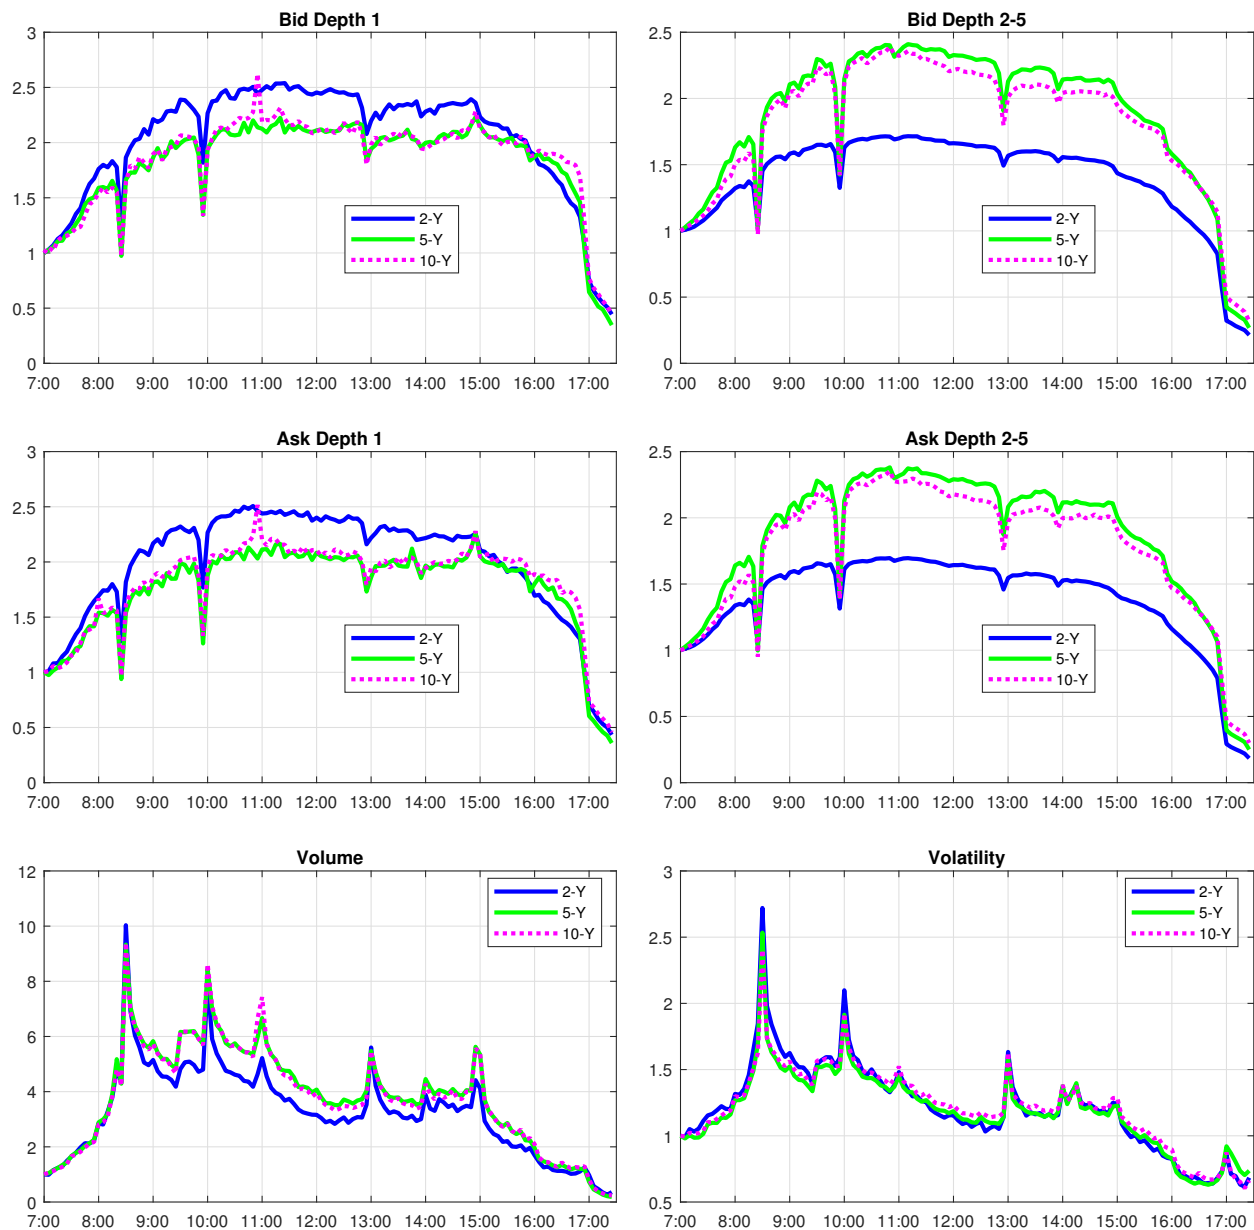


Figure 2: Intraday pattern of limit order book variables

This figure shows the intraday pattern of limit order book variables relative to their respective average level at the beginning of New York trading hours (7:00) for the 2-, 5-, and 10-year notes. For each five-minute interval of the trading day, the average is calculated for that interval across all days in the sample. All intraday averages are then anchored to the average of the first five-minute interval starting at 7:00. Data is from BrokerTec and covers the 2006-2015 period.

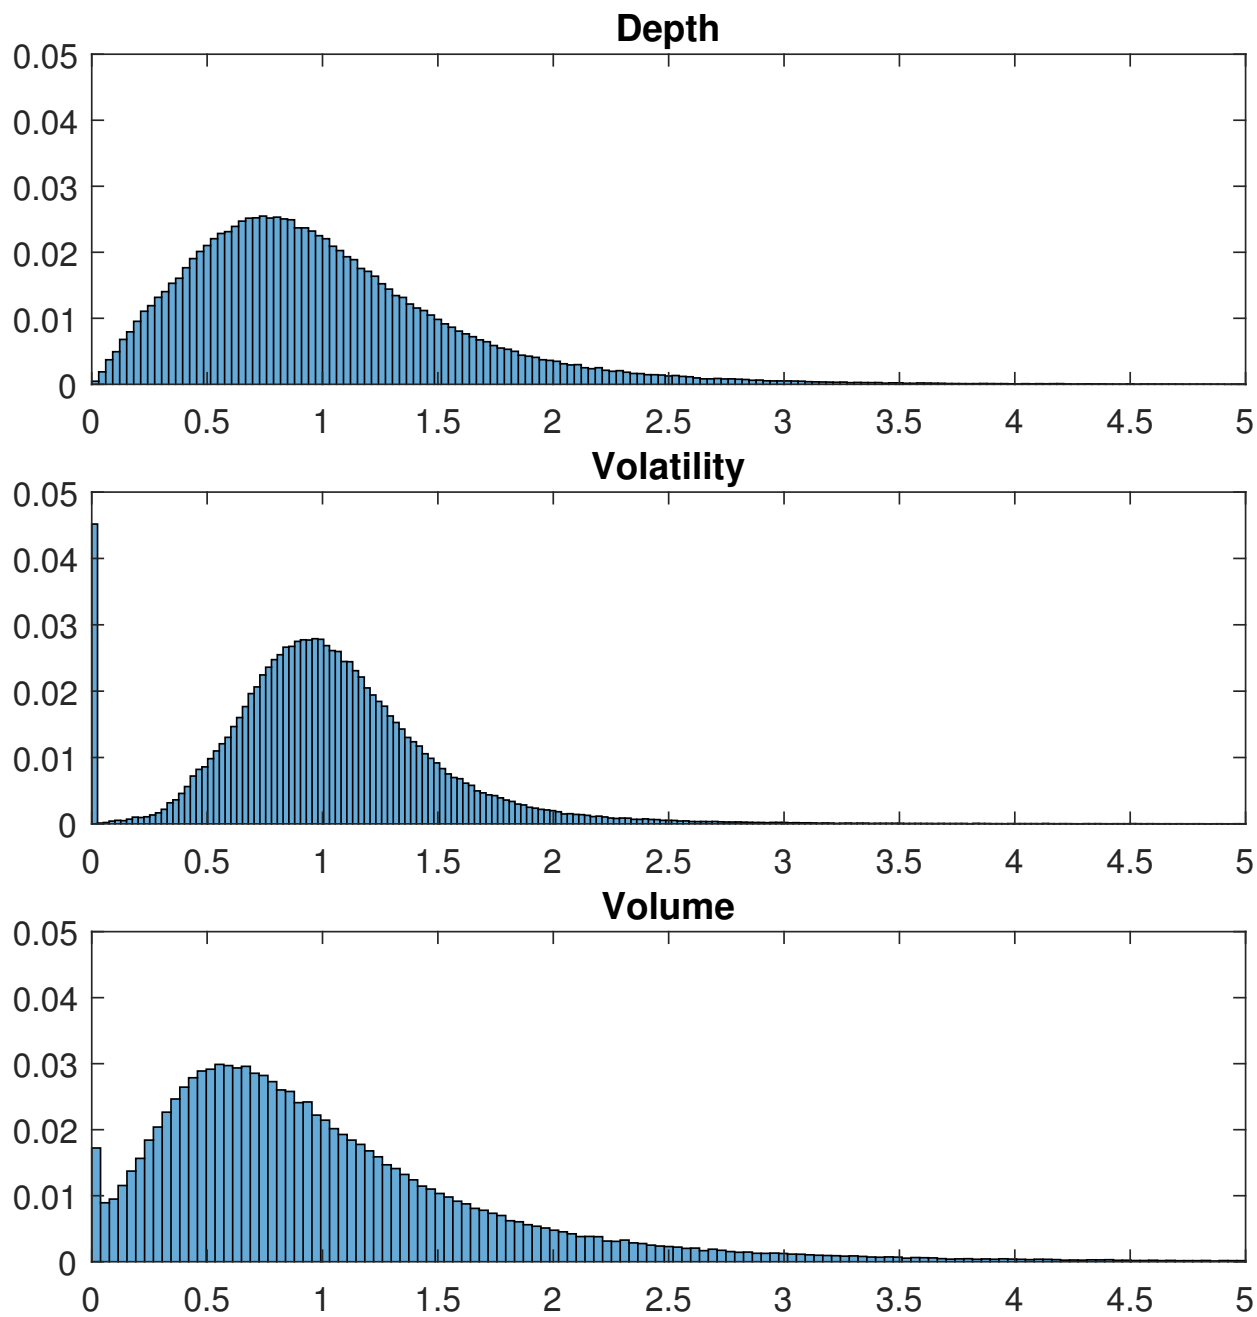


Figure 3: Histogram of diurnally-adjusted data for the 10-year note

This figure presents the histograms of diurnally-adjusted depth, volatility, and trading volume of the 10-year note. The diurnal adjustment factor for each five-minute interval is the rolling average of the given variable for that interval over the preceding 250 days. Data is from BrokerTec and covers the 2006-2015 period.




Figure 4: Autocorrelation functions of raw data versus models' residuals

This figure shows the ACFs of diurnally-adjusted depth, volatility, and volume before model estimation, versus residuals of three models. “eps-spec0” is the residuals based on the baseline model reported in Table 2. “eps-spec1” is the residuals based on Model 1 reported in Table 3. “eps-spec2” is the residuals based on Model 2 reported in Table 4. Model estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period.




Figure 5: Market liquidity at risk

This figure plots the daily average of the liquidity-at-risk measure $LR1$. It is the 5% left tailed market liquidity level, and computed as: $LR1_{i,t} = s_i \exp(\psi_{i,t}^d) \epsilon^*$, where s_i is the diurnal adjustment factor for time interval i , $\psi_{i,t}^d$ is the log conditional expected depth for interval i on day t based on Model 2 reported in Table 4, and ϵ^* is the fifth percentile of the empirical distribution of $\epsilon_{i,t}$ from the model. The measure is then averaged for each day. Model estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period.




Figure 6: Probability of illiquidity shock

This figure plots the daily average probability of low market liquidity, for the low threshold of \$2 million (based on the first percentile of the distribution of market depth over the sample period). This probability corresponds to $P\left(\epsilon_{i,t} < \frac{2}{s_i \exp(\psi_{i,t}^d)}\right)$, where s_i is the diurnal adjustment factor for time interval i , and $\psi_{i,t}^d$ is the log conditional expected depth for interval i on day t based on Model 2 reported in Table 4. This probability is computed based on the empirical distribution of $\epsilon_{i,t}$ from the model. The probability is then averaged for each day. The model estimates are obtained by QMLE using lognormal density for the residuals $\epsilon_{i,t}$ and five-minute snapshot data from the BrokerTec platform over the 2006-2015 period.

A Announcements

We consider three categories of news that are relevant for the Treasury market: 1) macroeconomic announcements, 2) monetary policy announcements (i.e., FOMC rate decision announcements), and 3) Treasury auction results.

A.1 Macroeconomic announcements

Included in our analysis are major economic announcements - those that are classified as “Market Moving” indicators by Bloomberg. Data for each announcement’s release date and time are from Bloomberg. These announcements include:

Time	Announcement	Frequency
8:30	Employment Report	Monthly
8:30	Consumer Price Index	Monthly
8:30	Durable Goods Orders	Monthly
8:30	GDP	Monthly ¹⁹
8:30	Housing Starts	Monthly
8:30	Jobless Claims	Weekly
8:30 ²⁰	Personal Income and Outlays	Monthly
8:30	Producer Price Index	Monthly
8:30	Retail Sales	Monthly
8:30	Trade Balance	Monthly
9:15	Industrial Production	Monthly
10:00	Existing Home Sales	Monthly
10:00	ISM Manufacturing	Monthly
10:00	New Home Sales	Monthly
Varied ²¹	Philadelphia Fed Business Outlook	Monthly

A.2 Monetary policy announcements

Included in our analysis are FOMC rate decision announcements. Such announcements usually occur after regularly scheduled FOMC meetings, of which there are eight per year. From January 2006 to January 2013, announcements after such meetings are made at about

¹⁹GDP is reported for each quarter in three rounds (advance, preliminary, and final) in the first, second, and third month following the reported quarter. Therefore, there is a GDP announcement every month.

²⁰Exceptions are the announcements on December 23, 2014, and February 26, 2016, both of which occur at 10:00.

²¹The announcement time is 12:00 from the start of the sample in 2006 to December 2007. It is 10:00 between January 2008 and October 2015. From November 2015 to present, the announcement time is 8:30.

14:15, excluding eight 12:30 announcements during the 2011-2012 period. From the March 2013 meeting, rate decision announcements are made at about 14:00.

Announcements can also occur after unscheduled meetings. During our sample period, such announcements occurred on January 22, 2008 at 8:20, and October 8, 2008 at 7:00. Announcement times are from Bloomberg.

A.3 Treasury auction result announcements

Included in our analysis are auction result announcements for the 2-, 5-, and 10-year notes. Most auctions for these notes close at 13:00. A few auctions close at other times (e.g., 11:00 and 11:30). We collect auction closing times from Treasury Direct (<http://www.treasurydirect.gov/>). Auction results are announced within minutes of the auction closing time.