THE HIGH FREQUENCY ECONOMICS OF GOVERNMENT BOND MARKETS

Giang Thi Huong Nguyen

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics.

Chapel Hill
2014

Approved by:

Eric Ghysels

Michael Fleming

Neville Francis

Jonathan Hill

Adam Reed

Eric Renault

ABSTRACT

Giang Nguyen: The High Frequency Economics of Government Bond Markets
(Under the Direction of Eric Ghysels)

This dissertation is a collection of four essays examining different aspects of government bond markets, with a special focus on the US Treasury securities. The first is a study of the microstructure of BrokerTec, the larger of the two electronic interdealer trading platforms for US Treasury securities, providing institutional background essential for subsequent studies. We characterize empirically market activities and the price discovery process. We show that both limit orders and trades affect prices, and that these effects are greater around monetary policy announcements. Contrary to previous findings pertaining to equity markets, we find that iceberg orders, which allow traders to hide liquidity, are not used frequently, even around volatile times.

The second essay examines a frequently used channel of hidden liquidity – the workup protocol. We ask whether trading activities during workups contain any private information and leave harmful effects on uninformed traders. We find that workup activities account for a significant portion of market liquidity not ex ante observable, but they tend to be less informative than transparent trades. We show that workups are used more often, but contain relatively less information, around volatile times, indicating that workups tend to be used as a channel to guard against adverse price movements, rather than as a channel to hide private information.

In the third essay, we propose a novel model to study jointly the intraday dynamics of liquidity and price risks, two important determinants of bond yields. We show that liquidity declines sharply during the 2008 crisis and on flight–to–safety days, accompanied by increased price volatility. Our model reveals a negative feedback effect between liquidity and volatility, and that each becomes more persistent during the crisis.

The fourth study provides an international perspective by studying the propagation of liquidity and volatility shocks during the 2010-2012 sovereign debt crisis across major euro-area bond markets, namely Belgium, France, Germany, Italy, the Netherlands, and Spain. We show that liquidity is generally the more important source of shocks transmitted across the borders, and this transmission largely originates from Italy and around the Italian crisis.

ACKNOWLEDGEMENTS

Over the last five years, I have received help, support and encouragement from a great number of individuals without whom this dissertation would not have been possible. First and foremost, I would like to thank Dr. Eric Ghysels for always being there to provide guidance and inspiration. I have grown so much as a researcher thanking to his mentoring. I am also indebted to Dr. Michael Fleming, who has been extremely generous with his time and resources, for providing very helpful and detailed feedback on my work. I greatly appreciate insightful comments and suggestions from Dr. Neville Francis, Dr. Jonathan Hill and Dr. Eric Renault during various stages of this dissertation. I would like to thank Dr. Adam Reed for broadening my prospective and offering great career advice. I would like to thank Dr. Helen Tauchen for her resourcefulness and encouragement. I thank Toan Phan, a faculty member and a friend, for his willingness to listen and provide feedback. I thank my friends, including my fellow doctoral students, for their help, support and friendship over the years.

I would like to thank Weiling Liu and Neel Krishnan, formerly at the Federal Bank of New York, for tremendous help with data processing and analysis. I thank Rita Moss, Michele Hayslett and the data team at the UNC Library for acquiring the European bond market data used in the last chapter of this dissertation. I also thank Arthur D'Arcy, Dan Cleaves, and Stuart Wexler from ICAP for clarifying how the BrokerTec platform works, and James Greco from Jefferies for helpful information on the trading of U.S. Treasury securities.

This dissertation is written with the financial support from the UNC Graduate School's Lovick P. Corn Dissertation Fellowship, made possible by a generous gift from the late Mr. Lovick Pierce Corn. Mr. Corn's legacy and generosity will always be remembered. In addition, the financial support from the Kampf Family for my graduate study and research is gratefully acknowledged.

Last but not least, I would like to thank my family for their unconditional love. Words cannot express how grateful I am for everything my husband Anh, my parents, and my mother in law have done to support my academic pursuit. My double bundle of joy, Khuê and Kỳ, are truly a blessing and I thank them everyday for making my life super fun (and super busy, of course).

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
INTRODUCTION

This dissertation is a collection of four essays on the high–frequency economics of government bond markets, with a special focus on the market for securities issued by the U.S. Government. Currently, the total outstanding amount of US Treasury securities exceeds $11 trillion. This market serves many vital roles in the financial system, not only in the U.S. but also around the globe with foreign countries holding nearly half of all Treasuries outstanding. The securities' creditworthiness and sheer liquidity make them a main instrument of monetary policy, a key store of value, a crucial source of collateral for financial transactions and a pricing benchmark for other financial assets.

The need to understand market dynamics in this important market at increasingly high frequencies arises from recent trends in trading. Since the Securities and Exchange Commission (SEC) authorized electronic markets in 1998, speed competition and technological advancements have turned the trading environment into one highly automated by computer algorithms and in which even a few milliseconds can affect the trading outcome (e.g., the Flash Crash of May 2010). The past few years have seen a growing body of literature seeking to further our knowledge of the new high speed trading environment. However, most research addresses equities markets. The Treasury market, despite its vital roles, remains much less studied, in part due to data availability (or lack thereof). The dissertation helps bridge this gap by providing a comprehensive analysis of this market along several dimensions of great interest to market participants, policy makers and academic researchers.

The first essay provides a detailed examination of an electronic trading platform for US Treasury securities, the BrokerTec platform (the larger of the two electronic interdealer trading platforms). Our findings suggest a level of liquidity on the BrokerTec platform that is improving over time and markedly greater than that found by other studies using data from the period before trading in these Treasury securities went electronic.

Importantly, we examine the price impact of not only trades but also of order book activities not previously available for the Treasury market. In fact, given the sheer amount of limit order book activities in comparison to trades, there is a lot to be learned about how these activities affect price dynamics. Furthermore, limit orders are often considered as supplying liquidity and market orders consuming it. Accordingly, our study can delineate the response of price to shocks in liquidity supply from that to shocks in liquidity demand. and show that limit orders as well as trades affect prices. In particular, the price impacts are found to be greater around monetary policy announcements, an important type of information events in this market.

In addition, we shed lights on the use of iceberg orders to hide liquidity in this fixed income market, an analysis not heretofore possible for Treasury securities. Several of our findings are consistent with the equity market evidence, but more importantly, we contribute evidence novel to the U.S. Treasury market. We find that iceberg orders are used much less often than in other markets examined in the literature. Furthermore, contrary to some earlier evidence on iceberg orders in equity markets, we find that Treasury traders are less likely to use iceberg orders when they post more aggressively priced limit orders, or when the market is more volatile, precisely when traders need greater protection.

The puzzling lack of popularity of iceberg orders contrasts sharply with the high usage of the workup protocol, an alternative mechanism to hide liquidity. This is the main subject of the second study. We examine this protocol and ask whether trading activities induced by this protocol, which generally account for more than half of market liquidity, are motivated by private information.

The contribution of this work extends beyond a study of a specific microstructure feature of the U.S. Treasury market. First, the workup mechanism is essentially a dark pool trading mechanism. Our study provides the first set of evidence on dark pool trading in a fixed income market setting. It is therefore a timely addition to the literature on dark pool trading and the current discussion among researchers and policy makers on the effects and implications of dark pool activities on market quality and welfare. We find that volatility tends to generate more workups, but that those workups tend to be less informative, suggesting the value of this dark pool mechanism in protecting traders against adverse price movements. In general, the amount of private information hidden in this Treasury dark pool is quite small, easing concerns that the dark pool could harm less informed traders.

Secondly, our work helps inform the current debate on market design response to high frequency trading. High frequency trading, or computer-driven trading in general, has increased significantly over the last few years – a trend dubbed “rise of the machines” in Chaboud et al. (2013). There is a continual need to devise new market design features to keep up with changing trends in trading, and to understand the implications of those features. The BrokerTec market design with the workup protocol fits neatly into this discussion via an interesting mix of continuous auction (the limit order book) and periodic call auctions (workups). Our empirical results readily provide a glimpse of the implications of such a market feature on price discovery and trading patterns.

With the first two studies providing a micro foundation for our understanding of the trading environment in the US Treasury market, the third study aims to understand the dynamics of liquidity and volatility, two important determinants of Treasury securities values. To this end, we propose a new class of econometric models to capture jointly the dynamics of liquidity and volatility at a high frequency interval (i.e., five-minute). Our models address several interesting questions for the US Treasury market. Is liquidity supply available when it is needed most? How is liquidity supply driven by uncertainty and other market factors, and conversely, does the supply of liquidity have any role in dampening or magnifying volatility in the market? How do the dynamics of the Treasury limit order book differ around the time of economic announcements, through the recent financial crisis, and during flight-to-safety episodes? Most of the previous 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 – much less understood.

We show that liquidity posted in the order book is lower on flight-to-safety days, when liquidity is especially needed. However, a high level of trading activity is also observed on those days, along with an elevated level of price uncertainty. These patterns collectively suggest that liquidity providers monitor the market more closely on these days and refrain from using limit orders to passively supply liquidity to the market. In general, price volatility and liquidity supply at the best price tier are negatively interrelated and each becomes more persistent during the crisis. This dangerous combination provides a great illustration to models of liquidity crashes (for example, Cespa and Foucault (2012)) in that bad shocks to either volatility or liquidity can intensify the negative feedback effect, leading to liquidity crashing while volatility spiking up. Our models also provide consistent evidence with the earlier literature that depth is withdrawn immediately before important economic announcements but then quickly gets refilled once the announcement is released, accompanied by a surge in trading activity and price uncertainty.

If the third essay focuses solely on the interaction between volatility and liquidity within the same bond market, namely the market for US Treasury securities, the final essay provides an international perspective by studying shock transmission across multiple bond markets. Shock propagation during a crisis is a particular concern for policy makers. The euro area sovereign debt crisis provides a valuable opportunity to study bond market linkages and spillover effects. In addition, with some of the bond markets in the area being among the largest in the world, only after the US and Japan markets, the findings of this study complement nicely what we have learned from the US government bond market.

We first measure spillovers via a forecast error variance decomposition of a vector autoregressive model, which captures jointly the dynamics of liquidity and volatility in the government bond markets of Belgium, France, Germany, Italy, the Netherlands, and Spain. The model controls for common trends in sovereign credit risk, financial sector credit risk, funding conditions, aggregate default risk, and proxies for regional and global risk aversion. As in the US Treasury market, liquidity and volatility in these euro area bond markets are closely inter-related, but we show that liquidity is generally the stronger force driving this inter-relationship.

We further show that liquidity is more responsive to macroeconomic developments and also the more important source of shocks transmitted across borders during the euro-area sovereign bond crisis. Our framework permits an assessment of the systemic role of each bond market during the crisis. The evidence consistently points to Italy as the sole net sender of liquidity shocks to other countries in the region. Furthermore, this transmission is greatest around the Italian crisis, as compared to that around the crisis associated with the other smaller sized periphery countries (i.e., Greece, Ireland, Portugal and Spain).

CHAPTER 2
THE MICROSTRUCTURE OF A US TREASURY ECN

2.1 Introduction

Since the early 2000's, trading in the U.S. Treasury securities market has migrated from voice-assisted brokers to fully electronic platforms (Mizrach and Neely (2006)). For the most recently auctioned securities in particular, the transition has been nearly complete, with nearly all interdealer trading now taking place via one of two electronic communications networks, BrokerTec and eSpeed (Barclay et al. (2006)). BrokerTec accounts for about 60% of trading activity (based on comparison with earlier studies using data from eSpeed).

This chapter assesses the microstructure of the U.S. Treasury securities market using tick data from BrokerTec. It is the first paper to closely study a U.S. Treasury market electronic communications network (ECN) and one of the first to analyze any fixed income market ECN.[1] Many previous papers have examined the microstructure of the Treasury market using data from GovPX, which consolidates data from voice-assisted brokers.[2] The migration of bond trading to the electronic platforms (which do not contribute to GovPX) has sharply reduced GovPX coverage of the interdealer market, as noted by Boni and Leach (2004) and others. The breadth of the BrokerTec tick data allows us to provide a comprehensive analysis of the market's microstructure as orders enter and leave the order book, and characterize market liquidity beyond

[1] Campbell and Hendry (2007) examine price discovery in the 10-year note using transactions data from BrokerTec. Mizrach and Neely (2006) estimate bid-ask spreads and market impact using transactions data from eSpeed. Additional studies examine the euro area sovereign debt market using data from MTS (e.g., Cheung et al. (2005), Menkveld et al. (2005), and Beber et al. (2009)). In addition, since the first draft of this chapter, there are several studies that look at different aspects of this market, including Dungey et al. (2013) for trade duration on eSpeed, Engle et al. (2012c) for intraday dynamics of market liquidity and volatility on BrokerTec, and Fleming and Nguyen (2013) for the order flow segmentation induced by the workup protocol on BrokerTec and the informational content of workup and non-workup trades.

[2] Fleming (1997) characterizes intraday liquidity, Fleming and Remolona (1997), Fleming and Remolona (1999), Balduzzi et al. (2001), Huang et al. (2002), and Fleming and Piazzesi (2005) look at announcement effects, Fleming (2002) examines the relationship between issue size and liquidity, Fleming (2003), Brandt and Kavajecz (2004), Green (2004), and Pasquariello and Vega (2007) assess the information content of trades, Goldreich et al. (2005) gauge the relationship between liquidity and value, and Brandt et al. (2007), Campbell and Hendry (2007), and Mizrach and Neely (2008) compare the information content of trades in spot and futures markets.

the inside tier for the first time. This is an important improvement, as the BrokerTec data shows that the inside tier depth is often not greatest in the book, and accounts for only a small fraction of the book's total depth.[3]

In addition, electronic trading facilitates greater speed of order manipulation and execution, permits an increased role for computer-driven and automated trading processes, and enables better market information collection, dissemination and processing. Coupled with the rise in electronic trading is a newly emergent trend in high frequency and/or algorithmic trading, the so-called "rise of the machines" (Chaboud et al. (2013)). Therefore, there is a great interest in understanding this new market structure and its level of trading activity and market liquidity from both academic and practitioner points of view.

Using tick data from 2010 to 2011, we characterize trading activity and liquidity on the BrokerTec platform for the on-the-run 2-, 3-, 5-, 7-, 10-, and 30-year Treasury securities.[4] Our findings suggest a level of liquidity on the BrokerTec platform that is improving over time and markedly greater than that found by earlier studies using data from GovPX. Since BrokerTec's inception, trading activity has grown many folds, e.g., starting at below $5 billion per day in 2001 to between $30-40 billion per day in 2011 for the 5- and 10-year notes. Over the 2010-2011 period, inside bid-ask spreads for maturities of five years or less average less than 1/100th of one percent. An average of over $300 million is available on the platform at the best price on either side for the 2-year note, $80 million for the 3-year note and in the $30 million range for each of the three remaining notes. There are even greater amounts available at the adjacent price tiers. Across the whole book, there is about $2.4 billion on each side for the 2-year note, $700 million for the 3-year note, and around $400 million for the 5- and 10-year notes.

Besides being the first to provide a comprehensive picture of a U.S. Treasury ECN, we make two further contributions. First, while previous studies have assessed price impact using GovPX trade data (e.g. Fleming (2003), Brandt and Kavajecz (2004), and Green (2004)), we examine the price impact of not only trades but also of order book activities not previously available for the Treasury market. In fact, given the sheer amount of limit order book activities in comparison to trades, there is a lot to be learned about how these activities affect price dynamics. Furthermore, limit orders are often considered as supplying liquidity and market orders consuming it. Accordingly, it is important to delineate the response of price to shocks in liquidity supply from that to shocks in liquidity demand.

[3] This fact has also been documented for equity limit order markets (e.g., Biais et al. (1995)).

[4] On-the-run securities are the most recently auctioned securities of a given maturity

We first calculate the permanent price impact of trades following the framework in Hasbrouck (1991a). We then extend this model to include limit order flow, separately for the bid and ask sides. Our work builds upon earlier studies of equity markets that incorporate order book information into the market impact function (e.g., Engle and Patton (2004), and Mizrach (2008)). A recent paper by Hautsch and Huang (2012a) uses a vector error correction model to analyze the dynamics of the limit order book for select NASDAQ stocks, and compute the price impact of orders of different types, sizes, and levels of price aggressiveness. They show that limit orders also have significant market impact.

We find that the price impact of trades on BrokerTec is quite small, but increases in maturity of the securities considered, ranging from 0.006/256th for the 2-year note to 0.450/256th for the 30-year bond per $1 million buyer-initiated volume. Equivalently, it takes about $182 million in signed trading volume to move the price of the 2-year note by 1/256th of one percent of par, whereas the required volume is only $2.2 million to move the price of the 30-year bond by the same amount. Moreover, limit order activities affect prices, and play an especially large role in the price dynamics of longer-dated maturities. Accounting for the impact of limit order activities on trading activities and price dynamics, the price impact of trades is about 9-14% lower for the 2-, 5-, 10-, and 30-year securities, and 26% and 40% lower for the 3- and 7-year notes, respectively. Our analysis also shows that trades and especially limit orders have a larger price impact immediately following Federal Open Market Committee (FOMC) rate decision announcements.

Another contribution lies in our analysis of hidden liquidity in the form of iceberg orders. The ability to enter "iceberg" orders (partially hidden orders) on the BrokerTec platform allows analyses not heretofore possible for Treasury securities.[5] Hidden orders in equity markets have been examined by Harris (1996), Aitken et al. (2001), Hasbrouck and Saar (2002), Anand and Weaver (2004), Tuttle (2006), De Winne and D'Hondt (2007a), De Winne and D'Hondt (2007b), Bessembinder et al. (2009), Pardo and Pascual (2012), and Hautsch and Huang (2012b), among others. We add to this literature by providing the first analysis of iceberg orders in the trading of Treasury securities. In particular, we study traders' order submission decision and explore whether certain order characteristics as well as prevailing market conditions might help predict the likelihood as well as the extent of hidden size of an iceberg order.

Several of our findings are consistent with the equity market evidence. For example, the use of hidden depth increases with order size and the prevailing bid-ask spread, intuitively highlighting the benefit of hidden

[5] Iceberg orders are not used on eSpeed, the other electronic platform for trading U.S. Treasury securities, leaving BrokerTec the only venue to study traders' choice with respect to such hidden orders.

orders as a mechanism to prevent information leakage and mitigate adverse selection risk. Additionally, when there is lower prevailing depth or lower likelihood of future orders whose display size will take precedence over the current hidden size, hidden orders tend to be used more often, as the cost of using them in terms of execution probability is lower.

Perhaps more valuable is our contribution of findings that are novel to the U.S. Treasury market. We find that iceberg orders are used much less often than in other markets examined in the literature. Typically iceberg orders account for less than 2% of order flow in the Treasury market, compared to 18% for stocks on Euronext-Paris (Bessembinder et al. (2009)), and 9% for 30 German blue chip stocks on Deutsche Borse's Xetra platform (Frey and Sandas (2012). Furthermore, contrary to the evidence documented in Bessembinder et al. (2009) that traders are more likely to use iceberg orders when they select a more aggressive limit order price, Treasury traders are generally less likely to do so for quote improving orders, except for the less liquid 7- and 30-year securities.

Another interesting finding of our work is that volatility and hidden order usage are negatively linked. At first blush, the finding seems counter-intuitive, as it suggests that the more volatile the market, the less likely that hidden orders will be used, precisely when traders need greater protection. However, if we place this finding in the context of the Treasury market, in which there exists another mechanism for order exposure management, namely the workup protocol, we can better understand how it could be the case for this market. The workup protocol gives market participants the ability to workup order sizes if and when desired, whereas iceberg orders can be adversely executed when the market is moving so fast that traders cannot cancel soon enough. As documented in Fleming and Nguyen (2013), workups tend to be used more frequently in more volatile times, undermining the popularity of iceberg orders. Likewise, hidden orders are used less often around the release of key macroeconomic reports, FOMC rate decision announcements, and Treasury auctions. These are moments when the market is eagerly waiting for and trading on the newly released announcements, so priority in the order queues seems to be an important consideration.

Overall, our work highlights how the electronic market for trading in U.S. Treasury securities differs from its voice-assisted precedent and from other markets studied in the literature. Comparing with the voice-assisted trading system, the electronic market facilitates a much greater frequency and volume of trades and limit order activities, resulting in greater competition for liquidity provision and thus lower bid-ask spreads and market impact. Comparing with other market setups, the high level of market liquidity and the presence of the more preferred protocol to manage order exposure in this market are likely related to the lower usage of iceberg orders and the seemingly greater importance of execution probability in traders' decisions.

The chapter proceeds as follows. Section 2.2 describes the structure of the interdealer Treasury market. Section 2.3 describes the BrokerTec data, characterizing trading activity and liquidity in the market. Section 2.4 presents the VAR model of returns and trades, and discusses a number of specifications and the resulting estimates of the price impact of trades. In Section 2.5, we add order book information to the model and quantify the price impact of limit orders. Section 2.6 presents our analysis of hidden orders. Section 2.7 concludes.

2.2 Market Structure

The secondary market for U.S. Treasury securities is a multiple dealer, over-the-counter market. The predominant market makers are the primary government securities dealers those dealers with a trading relationship with the Federal Reserve Bank of New York. The dealers trade with the Fed, their customers, and one another. The core of the market is the interdealer broker (IDB) market, which accounts for nearly all interdealer trading. Trading in the IDB market takes place 22-23 hours per day during the week, although we find that slightly over 90% of trading occurs during New York hours, roughly 07:00 to 17:30 Eastern time (comparable with what Fleming (1997) finds using GovPX data).

Until 1999, nearly all trading in the IDB market for U.S. Treasury securities occurred over the phone via voice-assisted brokers. Voice-assisted brokers provide dealers with proprietary electronic screens that post the best bid and offer prices called in by the dealers, along with the associated quantities. Quotes are binding until and unless withdrawn. Dealers execute trades by calling the brokers, who post the resulting trade price and size on their screens. The brokers thus match buyers and sellers, while ensuring anonymity, even after a trade. In compensation for their services, brokers charge a fee.

The migration from voice-assisted to fully electronic trading in the IDB market began in March 1999 when Cantor Fitzgerald introduced its eSpeed electronic trading platform. Cantor spun eSpeed off in a December 1999 public offering. After many ownership changes, eSpeed merged with BGC Partners, an offshoot of the original Cantor Fitzgerald. In 2013, eSpeed was purchased by NASDAQ OMX Group.

In June 2000, BrokerTec Global LLC, a rival electronic trading platform, began operations. BrokerTec had been formed the previous year as a joint venture of seven large fixed income dealers. BrokerTec was acquired in May 2003 by ICAP PLC. Mizrach and Neely (2006) describe the migration to electronic trading in greater detail, and Mizrach and Neely (2011) provide a summary of the evolution of the microstructure in the Treasury market.

2.2.1 The Electronic Platforms

BrokerTec and eSpeed are fully automated electronic trading platforms where buyers are matched to sellers without human intervention. A comparison of BrokerTec trading activity with that of eSpeed reported in Luo (2010) and Dungey et al. (2013) shows that BrokerTec accounts for around 60% of electronic interdealer trading in the on-the-run 2-, 5-, and 10-year notes and slightly above 50% for the 30-year bond.

The brokers provide electronic screens which display the best bid and offer prices and associated quantities. On BrokerTec, for example, a manual trader can see five price tiers and corresponding total size for each tier on each side of the book, plus individual order sizes for the best 10 bids and offers. For computer-based traders, the complete order book information is available. Traders enter limit orders or hit/take existing orders electronically. As with the voice brokers, the electronic brokers ensure trader anonymity, even after a trade, and charge a small fee for their services.

The BrokerTec platform operates as an electronic limit order market. Dealers send in orders that can be aggressive (market orders) or passive (limit orders), but they must all be priced. The minimum order size is $1 million par value. Dealers can enter aggressive orders at a price worse than the current best price. This is typically the case when dealers need to trade a large quantity for which the limit order quantity at the best price is not sufficient. The order will first exhaust all depth, both displayed and hidden, at better price levels until it reaches the originally stated price. Therefore, large aggressive orders can be executed at multiple prices. However, the incidence of market orders walking up or down the book is very small (below 0.5%). This is likely due to the large amount of depth usually available at the best price tier, and the ability to work up volume at a given price point.

The BrokerTec platform allows traders to enter iceberg orders, whereby a trader can choose to show only part of the amount he is willing to trade. As trading takes away the displayed portion of an iceberg order, the next installment of hidden depth equal to the pre-specified display size is then shown. This process continues until trading completely exhausts the iceberg order. It is not possible to enter iceberg orders with zero displayed quantity; that is, limit orders cannot be completely hidden.

The priority of execution of limit orders is based on price, display status and time. That is, limit orders with better prices have higher priority of execution. Displayed limit orders in the same price queue are executed on a first in, first out basis. Once all displayed depth at a particular price level is exhausted, hidden depth at that same price – if there is any – is then shown and executed.

Beside iceberg orders, the electronic brokers have retained the workup feature similar to the expandable limit order protocol of the voice-assisted brokers, but with some important modifications.[6] On BrokerTec, the most important change is that the right-of-first-refusal – previously given to the original parties to the transaction – has been eliminated, giving all market participants immediate access to workups. All trades consummated during a workup are assigned the same aggressive side as the original market order.[7] For a detailed analysis of workup activity in this market, see Fleming and Nguyen (2013).

2.2.2 The Voice-Assisted Brokers: GovPX

Most previous research on the microstructure of the Treasury market has used data from voice-assisted brokers, as reported by GovPX, Inc. GovPX receives market information from IDBs and re-disseminates the information in real time via the internet and data vendors. Information provided includes the best bid and offer prices, the quantity available at those quotes, and trade prices and volumes. In addition to the real-time data, GovPX sells historical tick data, which provides a record of the real-time data feed for use by researchers and others.

When GovPX started operations in June 1991, five major IDBs provided it with data, but Cantor Fitzgerald did not, so that GovPX covered about two-thirds of the interdealer market. Over time, the number of brokers declined due to mergers, and a new non-contributing electronic broker (BrokerTec) was formed. By the end of 2004, GovPX was receiving data from three voice-assisted brokers, but neither eSpeed nor BrokerTec, even though nearly all trading of on-the-run securities had migrated to these fully electronic brokers. After ICAP's purchase of GovPX in January 2005, ICAP's voice brokerage unit was the only brokerage entity reporting through GovPX.

---

[6]Boni and Leach (2004) provided a thorough explanation of this feature in the voice-assisted trading system. This feature allows a Treasury market trader whose order has been executed to have the right-of-first-refusal to trade additional volume at the same price. As a result, the trader might be able to have his market order fulfilled even though the original quoted depth is not sufficient. That is, the quoted depth is expandable.

[7]For a detailed description of the workup process on the BrokerTec platform, see "System and Method for Providing Workup Trading without Exclusive Trading Privileges", U.S. Patent number US8,005,745B1, dated August 23, 2011.

2.3 Data

Our analysis is based on tick data from the BrokerTec platform. The database provides a comprehensive record of every trade and order book change over the BrokerTec system for the on-the-run 2-, 3-, 5-, 7- and 10-year Treasury notes as well as the 30-year Treasury bond. We choose to focus on the period from January 2, 2010 to December 31, 2011. This is the most recent period for which we have available data. It is a sufficiently long sample period for a microstructure study, and relatively distanced from the 2007-2009 financial crisis, so that our analysis can provide an up-to-date characterization of the market's microstructure in a typical trading environment. For market dynamics during the crisis period, see Engle et al. (2012c).

2.3.1 Data Processing

From BrokerTec's detailed record of every trade and order book change, time-stamped to the millisecond, we process the data into two main parts: the trade data and the order book data. The trade data include price, quantity, and whether a trade was seller-initiated (a "hit") or buyer-initiated (a "take"). It should be noted that BrokerTec records the execution of a market order against multiple limit orders, as well as further matches during workups, as separate trade records. We aggregate these multiple trade records that belong to the same workup as one market transaction for the following reasons. First, treating the individual trade records as separate and distinct trades would artificially inflate the serial correlation in both trade initiation and signed trade flow and might compromise econometric modeling and inferences. Furthermore, our aggregation permits a more precise analysis of market order submission and the price impact of market orders, the size of which is better measured by the total volume exchanged during a trade and its associated workup. Our treatment is in line with BrokerTec's workup patent document which states that a workup is conceptually a "single deal extended in time". Nevertheless, the aggregation is not without cost in that it will sometimes overestimate the market order size.

The second part of the data concerns the limit order book, which we recreate from order book changes on a tick-by-tick basis, saving as much of the richness of the data as is practical. Each order book change record specifies the price, quantity change, shown and total quantities for that order, whether the order is a bid or an ask, and the reason for the change. The book can be changed as a result of limit order submission, modification, cancellation or execution against market orders. The order book data provide a view of the Treasury market far more detailed than that provided by GovPX data. In particular, our processed dataset not only tells us the best bid and offer and associated sizes at any given time, but also the depth available outside of the first tier. Moreover, we see the number of individual orders comprising the quantities available at particular prices. In addition, we are able to discern what quantities were visible to market participants at the time and what quantities were hidden.

Over our sample of 500 trading days in 2010 and 2011, BrokerTec intermediated almost \$63 trillion in trading of on-the-run coupon securities, or \$125.6 billion per day. The activity involved nearly 6 million transactions (each comprised of one or more trades), or almost 12,000 per day. Moreover, there were roughly 2.4 billion order book changes at the first five price tiers alone for these securities over our sample period, amounting to over 4.7 million order book ticks per day.

2.3.2 Trends in Trading Activity

To provide a historical perspective of trading activity on the BrokerTec platform since its beginning in the early 2000's, Figure 2.1 plots the average daily trading volume by year for the respective on-the-run coupon securities. As can be observed from the figure, there has been a sharp increase in trading activity over time, especially in the first seven years of the platform's history before the financial crisis intensified in late 2008. For the 10-year note, for example, average daily trading volume grew from \$2.9 billion in 2001 to a level over ten times larger in 2007 and, except for 2009, has remained above \$30 billion since. Another interesting observation is that the 2-year note – which used to be the most actively traded security with an average daily trading volume of nearly \$50 billion in 2008 – has seen lower activity since the crisis as the short rate has stayed at the zero bound. This contrasts with the post-crisis recovery observed in other securities. In 2010 and 2011, the 5-year note is the most actively traded, closely followed by the 10-year note. Trading in the other on-the-run securities is far below the level of the 2-, 5- and 10-year notes, although trading in the 3-year note rose quickly between late 2008 and 2011.

Focusing on the most recent years of 2010 and 2011, Table 2.1 reports average daily trading volume, trading frequency and trade size for each security. The table shows that trading in the 5- and 10-year notes is most frequent, with over 3,000 transactions per day, on average. The 5-year note is the most actively traded in terms of volume, with a daily trading volume exceeding \$36 billion. The 30-year bond is also quite frequently traded with nearly 2,000 transactions per day, but each trade is of much smaller size than that of the other securities, so that its total daily trading volume of nearly \$6 billion is far below the others. On the other hand, the 2-year note has the lowest trading intensity, but the largest average trade size, nearly ten times larger than that of the 30-year bond.

2.3.3 Liquidity Around the Clock

Figure 2.2 plots average BrokerTec trading volume by half-hour interval over the round-the-clock trading day for our six notes and bonds. To make the intraday patterns comparable across securities, we standardize the half-hour volume figures by the total daily volume of the relevant security. The findings are very consistent with what Fleming (1997) finds using GovPX data from 1994, and the patterns are strikingly similar across the six securities. Trading activity is extremely low during Tokyo trading hours (roughly 18:30 or 19:30 the previous day to 03:00 Eastern time), then picks up somewhat during morning trading hours in London. Trading then rises sharply during morning trading hours in New York, peaking between 08:30 and 09:00, and then peaking locally between 10:00 and 10:30. Trading reaches a final local peak between 14:30 and 15:00 and then tapers off by 17:30. This pattern is probably largely explained by scheduled macroeconomic announcements (most of which are made at 08:30 and 10:00), the hours of open outcry Treasury futures trading (08:20 to 15:00), and the pricing of fixed income indices at 15:00.

2.3.4 Spreads

The most basic measure of the bid-ask spread is the quoted spread. The inside quoted spread, s_t, is defined as the gap between the best (lowest) ask price, p_t^a, and the best (highest) bid price, p_t^b, i.e.:

$$s_t = p_t^a - p_t^b.$$

The middle column of Table 2.2 shows the average inside bid-ask spread in multiples of tick size of the relevant security.[8] Spread is generally increasing in maturity, from 1.03 128ths (2.06 256ths) at the 2-year maturity to 2.66 64ths (10.64 256ths) at the 30-year maturity. The 10-year note, however, has a narrower spread than the 7-year note. An interesting feature of the BrokerTec spreads is that they are quite close to the tick size for all of the notes (but not the 30-year bond), suggesting that the minimum tick increment may be constraining. Comparing to earlier studies using GovPX data, BrokerTec spreads are generally narrower.

[8] The tick size for the 2-, 3-, and 5-year securities is 1/128$^{\text{th}}$ of one percent of par and that for the 7-, 10- and 30-year securities is 1/64$^{\text{th}}$ of one percent of par.

Fleming (2003), for example, reports average bid-ask spreads of 0.39 32nds (3.12 256ths) for the 5-year note and 0.78 32nds (6.24 256ths) for the 10-year note, whereas the corresponding BrokerTec spreads are 1.18 128ths (2.36 256ths) and 1.15 64ths (4.60 256ths) respectively for these securities.[9]

Providing new information on how market depth is spaced along the price dimension beyond the inside tier, Table 2.2 shows the average price distance between adjacent price levels up to the fifth level in the book. All of the securities except for the 30-year bond have tightly populated order books at the first five price levels: adjacent depths are roughly one tick apart, although they get slightly wider further away from the inside tier.

To supplement the information provided in Table 2, we show in Figure 2.3 the frequency distributions of inside spreads for the six securities. Immediately apparent from the figure is the high degree of clustering of inside spreads at one tick, except for the 30-year bond whose distribution is more spread out at wider spread levels and peaks at two ticks. In particular, nearly 97% of inside spreads for the 2-year note are 2/256ths, another 3% are 4/256ths, and the negligible remainder is split between 0/256ths and above 4/256ths. Zero spreads, or "locked" markets, are possible, albeit infrequent, because prices exclude the brokerage fee, and because passive limit orders at the same price are not automatically executed against one another.

2.3.5 Market Depth

As a limit order market, liquidity on BrokerTec is supplied by limit orders submitted by market participants. Table 2.3 reports the total visible quantity of limit orders available on average at the best price level, the best five price levels, and across all price levels on each side of the market. Market depth is generally declining in maturity, greatest at the 2-year and lowest at the 30-year segment. At the inside price tier, there is about $300 million available on either side for trading in the 2-year note. It is interesting to observe that while being the most actively traded, the 5- and 10-year notes' market depth is on the lower end, averaging $26-31 million, suggesting a higher replenishment rate of liquidity to meet the high trading activity level. The inside depths reported here greatly exceed average depths on GovPX reported by earlier studies. For the 2-year note, for example, Fleming (2003) reports average depth on GovPX at the first tier of just $25 million (averaging across the bid and ask side).

[9]Note that the prices in both databases do not reflect brokerage fees. Such fees are proprietary, and can vary by customer and with volume, but are unquestionably lower for the electronic brokers than the voice-assisted brokers.

In addition, earlier studies using GovPX data are limited to the inside tier, leaving market liquidity beyond the first tier unknown. As Table 2.3 shows, market liquidity away from the first tier is substantial, several orders of magnitude larger than that available at the inside tier. Collectively across the best five tiers on each side, there is over $1.5 billion market depth for the 2-year note, about $470 million for the 3-year note, in the range of $210-280 million for each of the 5-, 7- and 10-year note, and $28 million for the 30-year bond. The first five tiers account for about 55-79% of total market depth for the notes and 47% of total market depth for the bond. That is, the first five tiers collect a disproportionally large amount of depth, given that there are typically around 16-18 price tiers on each side (slightly higher for the 5- and 10-year notes). The maximum number of price levels on one side during our sample ranges from 43 for the 30-year bond (on the bid side) to 101 for the 2-year note (on the ask side).

While depth in the book concentrates among the best five tiers, the inside tier is not the one with the greatest depth. To learn more about the depth distribution in the book away from the inside tier, we display a depth histogram of the order book in Figure 2.4. The figure illustrates again that order book depth outside the first tier is considerable. A notable feature of the depth distribution patterns is that there is consistently more quantity available at the second and third price tiers (and even fourth and fifth for some securities) than the first. The available quantity generally peaks at the second tier on both the bid and ask sides for the notes, and at the third tier for the bond. Depth then declines monotonically as one moves further away from the inside quotes. Biais et al. (1995) also find depth lower at the first tier than the second tier, but find similar depths at the second through fifth tiers.

2.3.6 Hidden Depth

In addition to information on visible depth at the best five tiers, Figure 2.4 also shows information on hidden depth. Hidden depth is only a small share of total depth at each price tier on average. The first tier has proportionally more hidden depth than other tiers. Among the securities, the 30-year bond has a greater share of depth that is hidden from view.

Next, we examine more closely the extent of hidden depth at the inside tier as well as across all tiers, and report the results in Table 2.4. The column "Full Sample" shows the percentage of hidden depth calculated across all five-minute snapshots over the whole sample period, representing an unconditional estimate of the extent of hidden depth. We then compute the percentages over only those snapshots when there is positive hidden depth (column "Hidden>0"). The percentages of those snapshots with positive hidden depth are reported in column "% of Obs.". These numbers indicate the probability of having hidden depth in the order book at any given time.

Since the results are quite similar between the bid and the ask side (although the numbers on the ask side are slightly lower), we discuss the findings for the bid side only. We find that on the bid side, hidden depth as a share of total depth at the inside tier is roughly in the vicinity of 10% for the 2-, 3-, 5-, and 10-year notes. The extent of hidden depth is particularly low at the 7-year maturity (just under 4%). In contrast, the share of hidden depth for the 30-year bond is far larger, about 23%. In terms of how likely it is for the inside tier to have hidden depth, the 2-year note is at the top with a 45% probability, followed by a nearly 25% probability for the 3-year. When there is hidden depth, the percentage hidden can be quite high, and in the extreme case of the 30-year bond, the average percentage reaches nearly 70%. The 7-year remains at the lower extreme in terms of both the extent and the likelihood of having hidden depth. When analyzing the overall percentages of hidden depth across the whole book, the numbers are much smaller, indicating that depth outside the first tier contains relatively less hidden depth. This finding is consistent with the belief that there is a greater need to hide exposure of orders closer to the market.

2.4 Price Impact of Trades

In this section, we quantify the price impact of trades as the long run cumulative response of price to a unit shock in trades, following Hasbrouck (1991a). This framework allows us to approximate the permanent price impact of trades that incorporates any delayed response and that is not contaminated by transitory effects. Accordingly, it provides a measure for the informational content of trades in this market.

Specifically, we estimate a structural VAR model with five lags for a vector of endogenous variables that consist of return and trade-related variables. We measure returns as changes in the best bid-ask midpoint, i.e., $r_t = m_t - m_{t-1}$, where t indexes transaction time, and m_t is the midpoint prevailing at the end of the t^{th} transaction. We let X_t denote trade-related variables (X_t can be a vector), so that the general structural VAR model is:

$$B_0 \begin{bmatrix} r_t \\ X_t \end{bmatrix} = \sum_{j=1}^{5} B_j \begin{bmatrix} r_{t-j} \\ X_{t-j} \end{bmatrix} + \begin{bmatrix} u_{r,t} \\ u_{X,t} \end{bmatrix},$$

where u_t is the structural innovation vector. The matrix B_0 captures the comtemporaneous effects within the endogenous variable vector. We will explain the chosen direction of contemporaneous effects when we present specific model estimates in subsequent subsections. The model is estimated by Seemingly Unrelated Regressions (SUR).

Based on the estimated dynamics of return and trade-related variables, we then compute the impulse response function (IRF) to a unitary shock in trade, that is,

$$\frac{\partial r_{t+h}}{\partial X_t}.$$

We compute the IRF out to 50 transactions after the shock ($h = 50$).[10] The permanent price impact is approximated by the cumulative return over this horizon.

We consider a number of specifications so as to gain a deeper understanding of how trading affects price dynamics, such as the extent to which trade direction contributes to price impact, both by itself and in conjunction with trade size. We present each specification and the corresponding price impact estimates in turn below.

2.4.1 Baseline Specification

We begin the estimation of market impact with a bivariate VAR of return and order flow q_t. We consider two alternative measures of order flow. The first is the direction of trade initiation x_t with a buy order signed +1 and a sell order signed -1. Trade initiation is recorded in the BrokerTec dataset, so all trades are classified properly. The second is the signed volume $x_t V_t$ where V_t is the actual volume of the t^{th} transaction.

In an ECN like BrokerTec, we can be sure that transactions, as well as the sequence of events associated with each transaction, are recorded in the proper order. That is, a market order arrives, executes against available limit orders on the opposite side, and the order book subsequently updates to reflect the transaction just taking place. This supports the identifying assumption that order flow contemporaneously affects return, but not vice versa. Accordingly, the model specification is:

$$\begin{bmatrix} 1 & -\alpha_{1,2} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} r_t \\ q_t \end{bmatrix} = \sum_{j=1}^{5} B_j \begin{bmatrix} r_{t-j} \\ q_{t-j} \end{bmatrix} + \begin{bmatrix} u_{r,t} \\ u_{q,t} \end{bmatrix}, \tag{2.1}$$

[10] Visual inspection of the IRF indicates that the 50-tick horizon is sufficiently long for the IRF to stabilize.

where B_j are (2×2) matrices. We estimate model (2.1) separately for $q_t = x_t$ and $q_t = x_t V_t$.

The permanent price impact estimates from model (2.1) are reported in Table 2.5. Under the column titled "Trade Direction" is the price response across maturities to a buyer-initiated trade (i.e., computed from the specification with trade initiation), while the column "Signed Trade Volume" shows the response to a $1 million shock in buyer-initiated trade flow (i.e., computed from the specification with signed trade volume). The price impact rises with maturity, except for the 7-year note which has a higher price impact than the 10-year note in the model using signed trade volume. A buy market order results in a permanent price increase, ranging from 0.357/256th for the 2-year note to 2.921/256th for the 30-year bond.

Since transaction size varies across maturities, a better cross securities comparison may be obtained by looking at the price impact per $1 million shock in the order flow of the respective securities. A $1 million increase in buyer-initiated trade flow moves the 2-year note's price by 0.006/256th, or alternatively, it takes about $363 million increase in buyer-initiated transaction volume to move the price by one tick (or 2/256th). The 30-year bond is much less liquid: a $1 million shock in the buyer-initiated order flow permanently increases the price by 0.450/256th, or equivalently, only $8.9 million is needed to move the price by one tick (or 4/256th).

2.4.2 Separate Effects of Trade Direction and Size

In the spirit of Hasbrouck (1991a), we also estimate a specification that incorporates both trade direction and size in order to explore their respective market impact:

$$\begin{bmatrix} 1 & -\alpha_{1,2} & -\alpha_{1,3} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_t \\ x_t \\ xV_t \end{bmatrix} = \sum_{j=1}^{5} B_j \begin{bmatrix} r_{t-j} \\ x_{t-j} \\ xV_{t-j} \end{bmatrix} + \begin{bmatrix} u_{r,t} \\ u_{x,t} \\ u_{xV,t} \end{bmatrix}. \tag{2.2}$$

Based on the model estimates, we compute the permanent price impact of trade direction and the marginal market impact of trade size beyond the minimum size.

We report the results in Table 2.6. The first column shows the price impact of a minimum-sized trade ($1 million), which ranges from 0.271/256th for the 2-year note to 2.378/256th for the 30-year bond. From here, price impact increases directly with trade size. Essentially, this specification disentangles the price impact of trade into two separate components: a "fixed" component due to trade initiation and a "variable" component that scales directly with the volume of the trade. For example, for a \$100 million buyer-initiated transaction in the 10-year note, the buy direction increases price by 1.033/256th, and the \$99 million increment in trade size from the \$1 million minimum increases price by an additional 3.265/256th, for a total price impact of 4.298/256th.

To entertain the possibility that the price impact of trade size does not increase linearly in trade size beyond the minimum size, we explore a further specification that allows for the non-linearity of trade size by incorporating signed trade volume squared in the system, as in Hasbrouck (1991a). Specifically,

$$\begin{bmatrix} 1 & -\alpha_{1,2} & -\alpha_{1,3} & -\alpha_{1,4} \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_t \\ x_t \\ xV_t \\ V_t^2 \end{bmatrix} = \sum_{j=1}^{5} B_j \begin{bmatrix} r_{t-j} \\ x_{t-j} \\ xV_{t-j} \\ V_{t-j}^2 \end{bmatrix} + \begin{bmatrix} u_{r,t} \\ u_{x,t} \\ u_{xV,t} \\ u_{V^2,t} \end{bmatrix}. \tag{2.3}$$

We plot the permanent price impact calculated from this model for varying trade sizes in Figure 2.5. It is clear from the figure that price impact is increasing in trade size. The concavity of the price impact function of the notes is quite mild, almost visually indistinguishable from linearity for the notes, especially the 2-year note whose price impact is already very small. Only the 30-year bond demonstrates a pronounced concavity in the price impact function. Parameter estimates (not shown) reveal that the squared trade size variable has a significant and negative contemporaneous effect on mid-quote return for all notes and bonds, but the magnitude is overwhelmed by the positive effects of trade direction and size. This suggests that a very large trade size is required for the concavity effect of price impact to kick in. We are able to see the concavity of the price impact function for the 30-year bond as trade size in this bond is typically very small (\$3 million).

2.4.3 Asymmetric Effects of Buys and Sells

We extend the baseline specification in equation (2.1) to explore if there is any asymmetry in the price impact between buyer-initiated and seller-initiated trades. Saar (2001), for example, motivates theoretically an asymmetric response to buyer- and seller-initiated block trades. The model we estimate is:

$$\begin{bmatrix} 1 & -\alpha_{1,2} & -\alpha_{1,3} \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_t \\ VB_t \\ VS_t \end{bmatrix} = \sum_{j=1}^{5} B_j \begin{bmatrix} r_{t-j} \\ VB_{t-j} \\ VS_{t-j} \end{bmatrix} + \begin{bmatrix} u_{r,t} \\ u_{VB,t} \\ u_{VS,t} \end{bmatrix}, \tag{2.4}$$

where VB and VS are the buy and sell transaction volume respectively. For buyer-initiated transactions, VB_t is equal to the transaction volume and $VS_t = 0$ (and vice versa for seller-initiated transactions).

The permanent price impact estimates are reported in Table 2.7. The estimates are quite similar in magnitude to the baseline estimates. In addition, there is little evidence to suggest that the market responds asymmetrically to buy versus sell trade initiation.

2.4.4 Asymmetric Effects on the Bid and Ask

Econometric modeling of the order book by Engle and Patton (2004) has stimulated interest in models which allow for a possibly asymmetric price impact on the bid and ask. Escribano and Pascual (2006) provide a detailed review of empirical evidence showing that bid and ask quotes do not adjust symmetrically after a trade. However, most prior evidence of such asymmetry is documented for equity markets. We explore if this asymmetry also prevails in the Treasury market.

We follow Escribano and Pascual (2006)'s generalization of Hasbrouck (1991a)'s structural model. The model allows bid and ask prices to follow separate stochastic processes, but imposes a vector error correction mechanism through the spread. That is, bid and ask prices can follow different dynamics but cannot deviate too much from each other given the spread. Buy and sell volumes are also separated to allow for their price effects to differ. This more flexible specification allows us to explore asymmetries, if any, in the market effects of buyer- versus seller-initiated transactions on bid versus ask prices. Escribano and Pascual (2006)'s generalization leads to the following structural vector error correction representation:

$$\begin{bmatrix} 1 & 0 & -\alpha_{1,3} & -\alpha_{1,4} \\ 0 & 1 & -\alpha_{2,3} & -\alpha_{2,4} \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \triangle p_t^b \\ \triangle p_t^a \\ x_t^b \\ x_t^a \end{bmatrix} = \gamma(L)s_{t-1} + \beta(L) \begin{bmatrix} \triangle p_{t-j}^b \\ \triangle p_{t-j}^a \\ x_{t-j}^b \\ x_{t-j}^a \end{bmatrix} + \begin{bmatrix} u_{p^b,t} \\ u_{p^a,t} \\ u_{x^b,t} \\ u_{x^a,t} \end{bmatrix}, \tag{2.5}$$

where γ is a 4×1 vector of error correction coefficients, s is the bid-ask spread, and $\beta(L)$ are matrices of autoregressive coefficients. The data also support a more parsimonious specification of the AR structure in which bid (ask) prices depend only upon lagged bid (ask) prices, and buys (sells) depend only upon changes in the ask (bid) prices:

$$\beta(L) = \begin{bmatrix} \beta_{1,1}(L) & 0 & \beta_{1,3}(L) & \beta_{1,4}(L) \\ 0 & \beta_{2,2}(L) & \beta_{2,3}(L) & \beta_{2,4}(L) \\ 0 & \beta_{3,2}(L) & \beta_{3,3}(L) & \beta_{3,4}(L) \\ \beta_{4,1}(L) & 0 & \beta_{4,3}(L) & \beta_{4,4}(L) \end{bmatrix}.$$

The system of dynamic structural equations (2.5) is estimated using SUR to allow for the possible correlation among the error terms. Indeed, as shown in Escribano and Pascual (2006), they have common components and therefore cannot be treated as mutually uncorrelated. Based on the model estimates, we compute the impulse response function of bid and ask prices to a unitary shock in buy (or sell) trades by forecasting this dynamic system over a 50-tick horizon following the shock. Before the shock, the system is assumed to be at rest with constant bid and ask prices, no trades and zero spread.

Table 2.8 reports the cumulative impact after 50 transactions following a \$1 million shock to buyer-initiated transaction volume (on the bid and ask prices), and similarly a \$1 million shock to seller-initiated transaction volume (on the bid and ask prices) in that order. The evidence presented in this table again suggests that there is little difference in the impact on bid price versus ask price induced by the same shock to order flow. Likewise, prices appear to respond similarly to a buyer-initiated order flow shock versus a seller-initiated one. In addition, the magnitudes of the price impact estimates remain quite similar to the baseline estimates, further highlighting the lack of asymmetry in the price impact of trades in this market.

2.5 Price Impact of Limit Orders

Given that the order book information is observable by market participants, the decision to place a trade and its size can be influenced by activities in the book. As reported earlier, there are about 4.7 million order book changes in the best five tiers alone, overwhelmingly outnumbering trading activity (about 12,000 transactions per day across the six securities). Theoretically, Boulatov and George (2013) suggest the concept of "informed liquidity provider", that is, informed traders can also be on the supply side, as opposed to the common assumption that informed traders merely consume liquidity. If so, relevant information might also be present in the limit order flow. Empirically, Mizrach (2008) shows that excluding this order book information is likely to overstate the market impact of trades. Hautsch and Huang (2012a) document significant price impact of limit orders for select NASDAQ stocks. We now extend our specification to incorporate information on limit order activities. We first estimate the price impact of trades and limit orders unconditionally. We then examine the price impact of both trades and limit orders following FOMC announcements in order to shed light on how price discovery varies with the information environment.

2.5.1 Price Impact of Limit Orders

We modify our specification (2.4) by adding the visible inside bid and ask net order flow between trades:

$$\begin{bmatrix} 1 & -\alpha_{1,2} & -\alpha_{1,3} & -\alpha_{1,4} & -\alpha_{1,5} \\ 0 & 1 & 0 & -\alpha_{2,4} & -\alpha_{2,5} \\ 0 & 0 & 1 & -\alpha_{3,4} & -\alpha_{3,5} \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_t \\ VB_t \\ VS_t \\ lb_t \\ la_t \end{bmatrix} = \sum_{j=1}^{5} B_j \begin{bmatrix} r_{t-j} \\ VB_{t-j} \\ VS_{t-j} \\ lb_{t-j} \\ la_{t-j} \end{bmatrix} + \begin{bmatrix} u_{r,t} \\ u_{VB,t} \\ u_{VS,t} \\ u_{lb,t} \\ u_{la,t} \end{bmatrix}, \qquad (2.6)$$

where lb, the bid limit order flow, is the volume of limit buy orders submitted to (positive) or cancelled from (negative) the first tier between trades, i.e., between the $(t-1)$ and t transactions. Similarly, la is the ask limit order flow. The measurement timing of the endogenous variables supports the direction of contemporaneous effects from limit order flow to trade flow to return in the above specification.

The measurement of limit order flow variables warrants some further discussion. Because our model already incorporates the effects of trades directly, we explicitly exclude order book changes caused by execution from our limit order flow measures. Specifically, the net flow of limit orders on each side of the market is computed as the difference between the quantity of new order submissions and that of cancellations from the last trade until immediately before the current trade. Our resulting measures of limit order flow account for the non-trade related change in liquidity supply in the market. As a result, our model can capture the dynamic interactions of liquidity demand (trade flow), liquidity supply (limit order flow) and price revisions. The model then allows for delineating the price impact of liquidity supply change from the price impact of liquidity demand change, a novel feature of our empirical exercise.[11]

[11] Hautsch and Huang (2012a) measure the price impact of limit orders by modeling the limit order book as a co-integrating vector comprised of price and depth up to the third level in the limit order book. Our model has a similar spirit in that it also incorporates limit order information in the vector of variables of interest, but focuses on the trading process and the price dynamics as affected by both trading and limit order activities.

We analyze the permanent price impact of trades and limit order activities by computing the cumulative price response to a shock vector that is zero everywhere except for the relevant order flow variable which has a unitary shock. The estimates are reported in Table 2.9. In line with the evidence on the impact of limit orders for equity markets (e.g., Hautsch and Huang (2012a)), our results show that limit order activity also leaves a permanent impact on price, although the impact is of smaller magnitude than that of trades. For example, a $1 million increase in bid limit order volume permanently raises the best bid-ask midpoint by 0.001/256th, 0.014/256th, and 0.035/256th for the 2-, 5- and 10-year notes. This implies that an increase in bid depth of $1.85 billion, $144 million and $113 million is required to raise the best bid-ask midpoint of the respective notes by one tick. In the less liquid 30-year bond, it takes as little as an $11.5 million increase in the bid limit order flow to raise the midpoint by one tick.

Additionally, the price impact of limit orders is higher for longer-maturity securities, both in magnitude as well as in comparison with the corresponding price impact of trades. For the 2-year note, the price impact of limit orders is about one fifth the impact of market orders of equal volume. A similar comparison for the 10-year note shows that limit orders have effects that are roughly two-thirds the corresponding impact of trades. For the 30-year bond, limit orders have almost as large effects as market orders.

The results also show that including information on order book depth affects market impact estimates. In particular, across all securities, trades now show smaller price impact estimates than those estimated from earlier specifications without limit order flow (as in Tables 2.5 and 2.7). These results suggest that ignoring limit order activity overstates the price impact of trades by 9-14% for the 2-, 5-, 10- and 30-year securities. The extent of overestimation is particularly acute for the 3- and 7-year notes: 27% and 40% respectively.

2.5.2 Market Impact Following FOMC Announcements

As our price impact estimates are based on all transactions over our sample period, one question that can naturally arise is whether price impact varies by the information environment. For example, Green (2004) shows that the information content of trades increases following macroeconomic announcements. We explore this question by choosing FOMC announcements as an exemplified information event around which to study the extent to which price impact might differ. FOMC announcements are key information events for the formation of Treasury prices, precipitating high price volatility, high trading volume, and wide bid-ask spreads (Fleming and Piazzesi (2005)).[12] The idea is to quantify price impact using only limit orders and transactions that take place during a short period following FOMC announcements, and compare with that computed during the same time window on non-FOMC days. The choice of the window and the "control" sample is discussed below.

During our sample period, there are 16 announcements following FOMC meetings, three of which occurred at about 12:30, and the rest of which occurred around 14:15. We collect the exact time at which the announcements reached the market, using the timestamp of the first news report appearing in Bloomberg. We focus on the 90-minute intervals after these announcements. In order to avoid the effect of price jumps that typically occur at announcement times without requiring trades, as documented in Fleming and Remolona (1999), we start the post-announcement window two minutes following the actual announcement times. We choose the same time window on the five days preceding and five days following each FOMC announcement to serve as the non-announcement counterpart, effectively controlling for the time-of-day effect and general market conditions.

We estimate model (2.6) using data in the post-announcement window on FOMC days and in the comparable window on non-FOMC days. The corresponding market impact estimates are reported in Panel A and B respectively in Table 2.10. Two important results can be observed. First, market impact is pervasively higher following FOMC announcements. The market impact of a buyer-initiated trade is about 20-40% larger than in the same time interval on non-FOMC days, except for the 10-year note and 30-year bond where the price impact is slightly weaker on FOMC days. The increase in price impact of seller-initiated trades following FOMC announcements varies more widely across securities, ranging from about 5% for the 7-year note to over 100% for the 3-year note.

Secondly, and perhaps more strikingly, limit order flows have a much greater impact after FOMC announcements than during the comparable time window on non-FOMC days. For example, for the 2-year note, the price impact of limit order flow to the bid side is about triple the non-FOMC price impact, and that to the ask side increases nearly five fold. Accordingly, limit order flows become nearly as important as, or occasionally even more important than, trade flows in the price discovery process following FOMC announcements.

[12] Gao and Mizrach (2013) show that price impact in the equity market rose substantially following regularly scheduled Permanent Open Market Operations during the Federal Reserve's first large-scale asset purchase program.

In order to address the concern that our comparable sample, which includes both the five days before and five days after each of the FOMC announcements, may be contaminated by the effect of these announcements, we alternatively include in the non-FOMC sample only the five days preceding each of the announcements. The results are qualitatively similar. Furthermore, we check the robustness of our estimates by varying the starting time of the window between two and five minutes after announcement times, as well as using a shorter window of 60 minutes. These sensitivity analyses all confirm the results obtained above.

2.6 Hidden Orders

The preceding price impact analysis characterizes market liquidity as observable by market participants. However, traders on the BrokerTec platform have the option to submit iceberg, or partially hidden, orders. The hidden quantity represents a source of liquidity that is not known to the marketplace until and unless it is later revealed in trade executions. From the perspective of the hidden order traders, this type of order helps them manage their order exposure, reducing information leakage and front running.

To gain a more complete understanding of this market, we explore the pattern of iceberg order usage on BrokerTec and examine the effects of order characteristics and market conditions on the likelihood and hidden size of iceberg orders. Given the sheer amount of data at the order level, we choose to work with newly submitted orders to the first tier only, and for a subset of trading days in the sample, in order to keep our analysis computationally manageable. Specifically, of the 500 trading days in the sample, we randomly select 100 days for analysis.

2.6.1 Descriptive Analysis of Hidden Orders

Hidden order usage and some basic features of orders, whether completely transparent or partially hidden, are shown in Table 2.11. The number of order submissions per day varies across securities. For the 5-, 7- and 10-year notes, there are over 145,000 orders submitted to the first tier per day, on average. The 2- and 3-year notes have over 60,000 orders per day, and the 30-year bond has the lowest number of orders at 38,000 per day. While order submission activity is quite high, most of the order flow is completely visible – only about 2% or less of orders are iceberg orders. This is much lower than the extent of hidden order usage in other markets. For example, Bessembinder et al. (2009) report that iceberg orders on average account for 18% of order flow for stocks on Euronext-Paris.

There are several additional observations of interest from this table. First, iceberg orders tend to be several orders of magnitude larger than normal orders. For example, an average iceberg order in the 10-year note is about $11.6 million, whereas an average visible order is only $1.4 million. This finding helps reconcile the low percentage of iceberg orders with the much higher percentage of hidden depth residing in the book at any given point in time. That is, iceberg orders, while used sparingly, tend to hide a large quantity, so the hidden depth proportion tends to be higher.

Secondly, iceberg orders tend to be more price-aggressive, as the percentage of orders placed inside the prevailing spread is higher among iceberg orders than completely visible orders. This is especially the case with the 30-year bond with over 53% of iceberg orders that are price-improving. Contributing to this result is the fact that the bid-ask spread for the 30-year bond is typically much wider than the spread for the notes, making it easier to undercut the best price.

Thirdly, the arrival rate of similar limit orders around the time of order submission shows some noticeable differences between visible and iceberg orders. Iceberg orders tend to be used when similar limit orders arrive more slowly. Given the priority rule favoring displayed depth, the hidden part of an iceberg order has lower priority than the displayed depth of future limit orders at the same price point. Accordingly, the higher the expected arrival rate of future competition, the less likely the use of iceberg orders.

2.6.2 Determinants of Hidden Orders

We proceed to analyze factors that might contribute to the likelihood of hidden orders in a multivariate framework. The order-level data from BrokerTec allows us to examine hidden orders directly as they enter the limit order book. This provides for a clean analysis of factors that might be driving the exposure choice. We employ a logistic model to predict the likelihood that an incoming limit order contains some hidden size, building upon the approach in De Winne and D'Hondt (2007b) and Bessembinder et al. (2009). The dependent variable is a binary variable that takes the value of 1 if the order is an iceberg order, and 0 otherwise.

Findings in the literature help guide our selection of explanatory variables. First, in Bessembinder et al. (2009) and references therein, the price aggressiveness of an order has been shown to affect the use of hidden orders. That is, more aggressively priced orders are more likely to be iceberg orders. Since our study is concerned with only the orders coming to the first tier of the book, we measure price aggressiveness by IMP, an indicator variable for whether the order is improving the current best price on the relevant side.

The next explanatory variable is $SIZE$, the total size of the order (logged and standardized by its daily mean and standard deviation). This variable relates to the benefits and costs of order exposure as articulated in Harris (1997). Traders choose to expose their orders in order to attract potential counterparties who otherwise may not have known about the existence of such trading interest. However, exposing their trading interest, especially when such interest is large, can reveal useful information about trading intention and potential future price impact. Such exposure can provide free trading options to other market participants who may then take actions detrimental to the exposing traders. Accordingly, large traders may find it useful to hide part of their orders from the market.

Another possible factor in hidden order usage, as suggested by Moinas (2010), is that informed traders may use hidden orders to mitigate information leakage. As a result, it is predicted that hidden orders are more likely when adverse selection risk is higher. Following Bessembinder et al. (2009), we use the bid-ask spread SPR (measured in basis points of the bid-ask midpoint) as a proxy for the degree of adverse selection and test whether a wider bid-ask spread is associated with a greater probability of an iceberg order and a greater hidden size. Empirically, results reported in Bessembinder et al. (2009) support this prediction for a sample of Euronext-Paris stocks during April 2003.

Furthermore, the state of the order book has been shown to affect the choice of order exposure. Buti and Rindi (2013) argue that uninformed liquidity suppliers use hidden orders to reduce picking-off risk and discourage undercutting in liquidity provision. This has several implications for iceberg order usage concerning the level of depth and price volatility in the order book.

On the one hand, a higher level of prevailing depth on the same side (especially relative to the order size and the prevailing depth on the opposite side) indicates a lower probability of a new limit order being picked off. That is, the greater depth queueing in front of the incoming order provides a greater protection against adverse execution of such order. This in turn reduces the need to use an iceberg order. On the other hand, fully displaying the order size and adding to the already high level of prevailing depth can potentially induce future order submitters to undercut and post price-improving orders instead of joining the current price queue. As a result, higher prevailing depth may lead to a higher probability of hidden order. Therefore, while prevailing depth on the same side, $DSAME$, is argued to be an important determinant of hidden order usage, whether it is negatively or positively linked with hidden order usage is an empirical question.

Moreover, as documented in De Winne and D'Hondt (2007b), the order book imbalance (positive-valued if the book is heavier on the same side as the order) decreases the likelihood of an iceberg order. Thus, it appears that the relative magnitude of same side depth and opposite side depth also matters in explaining hidden order usage. Our model includes prevailing depth on the opposite side of the market, $DOPP$, to explore this conjecture.

Besides the effect of depth, the degree of market volatility increases the risk of adverse execution of limit orders, thereby making hidden orders more useful in helping traders manage their order exposure and reduce the chance of being picked off. We measure this risk by $VOLA$, the prevailing five-minute realized volatility of one-second returns based on the best bid-ask midpoint.

In addition, there is abundant empirical evidence suggesting that the level of trading activity and the rate of limit order arrival in the market helps explain the use and extent of hidden orders (for example, see Bessembinder et al. (2009), De Winne and D'Hondt (2007b), Aitken et al. (2001), and references therein). In particular, Aitken et al. (2001) argue that a higher trading activity level indicates a lower expected time-to-execution of limit orders, thereby reducing the need to hide part of a limit order as a way to mitigate the free trading option inherent in limit orders. If this argument applies, we expect to see a negative link between trading activity, as captured by $NTRANS$ – the number of transactions over the last five minutes – and hidden order usage.

On the contrary, during periods when limit orders on the same side are slow to arrive, the threat of future (displayed) limit orders taking priority over the hidden portion is lessened, resulting in higher likelihood of hidden order usage. The expected arrival rate of similar limit orders is measured by $WAIT$, the average inter-order duration (in seconds) for the last three limit orders on the same side.

Finally, we include several dummy variables to account for potential differences in the order exposure choice around the announcement of important news and overnight trading hours.[13] In particular, $PRENEWS$ and $POSTNEWS$ are indicator variables for whether the order is submitted within the five-minute time window before and after an announcement. $OFFHR$ is an indicator variable that is equal to 1 if the order is submitted outside the New York trading hours of 7:00 to 17:30.

Our model is estimated using data on all order submissions on the 100 randomly selected days. We note that continuous explanatory variables (namely $SIZE$, SPR, $DSAME$, $DOPP$, $VOLA$, $WAIT$, $NTRANS$) are standardized by the corresponding daily mean and standard deviation so that they are comparable across days. Table 2.12 presents the model parameter estimates along with the odds ratios. Since

[13] We consider the same set of key macroeconomic reports, FOMC rate decisions and Treasury auction results as in Engle et al. (2012c)

continuous variables in the model are demeaned and variance-rescaled by daily statistics, the odds ratios correspond to a one standard deviation change in the relevant variable.

As shown by the significantly negative coefficients for IMP for four out of the six securities (the 2-, 3-, 5-, and 10-year notes), price-improving orders are less likely to contain hidden volume, after controlling for other factors believed to affect hidden order usage. This result is opposite to that reported in De Winne and D'Hondt (2007b) and Bessembinder et al. (2009). By achieving price priority with a price-improving order, the order submitter seems to indicate an eagerness for faster execution and the result suggests that the trader prefers to display the full order size so as not to lose the priority to future orders that join the queue at the new price. However, this is not the case for the 7- and 30-year securities, which are traded much less actively, and for which orders are hence less likely to lose priority to future orders.

Consistent with the prior literature, we find that large orders are more likely to be partially hidden, because posting a large order may give away information to the market as suggested by Harris (1997). A one standard deviation increase in the logged order size is associated with more than twice the odds of the order having hidden size. Furthermore, comparing the odds ratios across explanatory variables shows that order size is uniformly the most important driver of the hidden order decision.

With regard to adverse selection risk, we find that hidden orders are more likely when the bid-ask spread is wider for all of the notes, in line with the prediction by Moinas (2010) and the empirical evidence reported in Bessembinder et al. (2009). A one standard deviation increase in the prevailing spread is associated with a roughly 10% increase in the odds of hidden size. The 30-year bond is an exception, in which the bid-ask spread has a negative effect on the likelihood of an iceberg order. That is, when the spread is wide, liquidity providers seem to prefer to submit a fully displayed order rather than an iceberg order, possibly to achieve faster execution in order to earn the spread. As documented earlier, the 30-year bond has a markedly wider spread, on average, and is much less liquid than the notes (i.e., lower trading volume and lower standing depth), indicating a lower degree of competition for liquidity provision. Accordingly, although a wider spread makes it easier for an order to be front-run, the front-running risk is probably not substantial, thereby reducing the need to use iceberg orders. While this finding is somewhat unexpected, the same effect has been documented by De Winne and D'Hondt (2007b) for non-marketable limit orders for the majority of their sample of 40 stocks on Euronext.

The prevailing depth on the same side as the incoming order has a significantly negative effect on the probability of hidden volume across all securities, except for the 2-year note. This finding provides support for the hypothesis that an order is less likely to have hidden volume if the prevailing depth is high, since the plentiful depth on the same side provides a reassuring signal that orders on that side are less subject to being picked-off. In addition, Treasury securities have very tight spreads, providing less scope for undercutting, so that the opposite effect of depth on iceberg order usage relating to undercutting risk is less pronounced. Even for the 30-year bond, whose wider bid-ask spread makes undercutting easier, this risk is evidently not a large concern for traders. However, depth on the opposite side is generally also negatively related with iceberg order usage, suggesting that the likelihood of an iceberg order is lower when the standing order book is deeper.

Interestingly, the consistently negative coefficients on volatility across maturities do not support theoretical predictions of a positive relationship discussed earlier. We suspect that, when prices are moving fast, Treasury traders refrain from using hidden orders altogether since they have an alternative mechanism, namely the workup protocol, to protect themselves from adverse price movements. That mechanism affords them complete control over how much to bid/offer based on changing market conditions, as opposed to making a firm commitment over the total quantity they want to bid/offer, even if part of that commitment is hidden from view. In fact, Fleming and Nguyen (2013) show that workups are utilized more in volatile times.

The rates of limit order arrivals, as captured by the average wait time between recent same-side orders, show expected effects on hidden order usage. Specifically, a longer wait time suggests a slower arrival rate for future orders, and thus, the threat of the order's hidden volume losing priority to future orders is lessened. The positive and significant coefficients on $WAIT$ across securities support this hypothesis. In contrast, there is mixed evidence as to the effect of trading rates, $NTRANS$, on hidden order choice. It is negative for the 3- and 7-year notes, but positive for the 5-, 10- and 30-year securities and insignificant for the 2-year note.

Importantly, there is consistent evidence that hidden orders are used less often around announcement times. This appears to be the period when the market is geared up to receive and then incorporate the news. Accordingly, the lower priority of hidden volume may make hidden orders less attractive around these moments. Lastly, we find that hidden orders are more likely during the overnight trading hours.

2.6.3 Determinants of Hidden Volume

Conditional on the choice to partially hide an order's volume, the next natural step is to explore how hidden volume is determined. For this purpose, we regress the hidden size (logged) of hidden orders on the same set of explanatory variables in the hidden order choice model, and report the results in Table 2.13. To facilitate interpretation, we also report the exponential of parameter estimates so that one can see the effect of each explanatory variable directly on the hidden size, as opposed to logged hidden size.

A few key observations from this table are in order. First, $SIZE$ continues to be a key driver of hidden orders: the larger the size, the larger the hidden volume. A one standard deviation increase in the logged order size is associated with an increase in the hidden size ranging from 47% for the 30-year bond to 140% for the 2-year note. IMP also appears to play an important role in the extent of volume hidden: those hidden orders placed inside of the spread have about 3-10% higher hidden volume than similar hidden orders that are not price-improving. In addition, the $PRENEWS$ and $POSTNEWS$ variables are generally negative, suggesting that hidden orders placed around announcement times tend to have lower hidden volume, in addition to the earlier reported evidence that these hidden orders are less likely around these times. The effects of other variables on hidden volume are less determinative and vary across securities, suggesting that these variables matter more to the choice of hidden orders than the extent of hidden volume once the choice has been made.

2.7 Conclusion

The microstructure of the U.S. Treasury securities market has changed markedly in recent years, with trading activity migrating from voice-assisted brokers to fully electronic brokers. We use tick data from one of these platforms, BrokerTec, to reassess market liquidity. We find that the market is notably more liquid than earlier reports based on GovPX data, and that there has been an increase in liquidity over time, except for the crisis period. In addition, our work offers the first look into market liquidity beyond the inside tier. We show that market liquidity concentrates more heavily at the price tiers immediately behind the market, and that the first five tiers collect over half of the total market liquidity in the order book at any given point in time.

We formally quantify the price impact of trading and limit order book activities. The price impact of trades on BrokerTec is quite small but generally increasing in maturity. Baseline estimates based on the specification with price dynamics and trading activities suggest that it takes \$182 million in signed trading volume to move the price of the 2-year note by 1/256th of one percent of par, but only slightly over \$2 million to move the price of the 30-year bond by the same amount. Accounting for the impact of limit order activities on trading activities and price dynamics, we find that limit order flow itself affects prices, and is especially important in the price dynamics of longer-dated maturities. We also find that part of the price impact of trades initially estimated from the model of trade flow alone can be attributed to limit order activities: including limit order activities reduces the price impact of trades by about 9-40%, on average, for the on-the-run securities. Finally, price impact is larger following FOMC announcements, particularly that of limit orders.

We find that iceberg orders are used sparingly in the Treasury market. However, the hidden portion of these iceberg orders is several orders of magnitude larger than the displayed part, thus occupying a proportionately larger portion of depth residing in the order book at any point in time. We find that the use of hidden depth increases with the order size and the prevailing bid-ask spread, highlighting the benefit of hidden orders as a mechanism to prevent information leakage and mitigate adverse selection risk. Additionally, when there is lower prevailing depth or lower likelihood of future orders whose display size will take precedence over the current hidden size, hidden orders tend to be used more often, as the cost of using them in terms of execution probability is lower. These results are generally in line with the evidence reported for other markets.

However, we also find a number of results in this market that have not been documented elsewhere. Unlike Bessembinder et al. (2009), we find that traders are less likely to use iceberg orders when their orders are price improving, except for the less liquid 7- and 30-year securities. Considering that this market is highly liquid, in which the inside spread is restrained at one tick the majority of the time, opportunities to undercut the market are limited and thus, whenever such opportunities are present, execution probability seems to take on greater importance. The extent of iceberg order usage in this highly liquid market is also much smaller than what has been documented in the literature for equity markets. In this regard, our findings actually support the evidence from equity markets that hidden order usage is lower for the more liquid stocks.

Another interesting departure from both theoretical predictions and previous empirical evidence is that volatility and hidden order usage are negatively linked. At first blush, the finding seems counter-intuitive, as it suggests that the more volatile the market, the less likely that hidden orders will be used, precisely when traders need greater protection. However, if we place this finding in the context of the Treasury market, in which there exists another mechanism for order exposure management, namely the workup protocol, we can better understand how it could be the case for this market. Recall that the workup protocol gives market participants the ability to workup order sizes if and when desired, whereas iceberg orders can be adversely executed when the market is moving so fast that traders cannot cancel soon enough. Empirically, workups tend to be used more frequently in more volatile times, undermining the popularity of iceberg orders. Likewise, hidden orders are used less often around the release of key macroeconomic reports, FOMC rate decision announcements, and Treasury auctions. These are moments when the market is eagerly waiting for and trading on the newly released announcements, so priority in the order queues seems to be an important consideration.

Overall, our study highlights how the electronic market for trading in U.S. Treasury securities differs from its voice-assisted precedent and from other markets studied in the literature. Comparing with the voice-assisted trading system, the electronic market facilitates a much greater frequency and volume of trades and limit order activities, resulting in greater competition for liquidity provision and thus lower bid-ask spreads and market impact. Comparing with other market setups, the high level of market liquidity and the presence of the more preferred workup protocol to manage order exposure in this market are likely related to the lower usage of iceberg orders and the seemingly greater importance of execution probability in traders' decisions.

Table 2.1: Trading Activity

Maturity	Trading Volume	Trade Frequency	Average Trade Size
2-Year	26,354	934	28.2
3-Year	16,204	1,297	12.5
5-Year	36,262	3,052	11.9
7-Year	9,640	1,500	6.4
10-Year	31,462	3,066	10.3
30-Year	5,705	1,921	3.0

The table reports daily averages of trading volume, trade frequency, and average trade size for on-the-run Treasury coupon securities on the BrokerTec platform, for the period 2010-2011. Volume and trade size are reported in millions of dollars. Multiple order matchings (including during workups) associated with the arrival of an aggressive order are aggregated as a single trade.

Table 2.2: Average Bid-Ask Spread and Inter-Tier Price Distance

	Bid 5-4	Bid 4-3	Bid 3-2	Bid 2-1	B-A Spread	Ask 1-2	Ask 2-3	Ask 3-4	Ask 4-5
2-Year	1.02	1.01	1.01	1.00	1.03	1.01	1.01	1.03	1.03
3-Year	1.08	1.04	1.02	1.02	1.12	1.02	1.02	1.03	1.08
5-Year	1.06	1.04	1.03	1.03	1.18	1.03	1.04	1.05	1.09
7-Year	1.13	1.09	1.05	1.04	1.33	1.04	1.04	1.06	1.10
10-Year	1.04	1.03	1.02	1.02	1.15	1.02	1.02	1.04	1.06
30-Year	1.48	1.35	1.33	1.29	2.66	1.32	1.34	1.39	1.49

The table reports the average price distance between adjacent price levels on each side of the book, and the average bid-ask spread. All numbers are in multiples of the tick size of the corresponding security. The tick size for the 2-, 3- and 5-year maturities is $1/128^{th}$ of one percent of par, and that for the 7-, 10- and 30-year maturities is $1/64^{th}$ of one percent of par. The numbers are computed from five-minute snapshots of BrokerTec's limit order book for the respective securities for the hours 07:00-17:30 over the period 2010-2011.

Table 2.3: Limit Order Book Depth

	First Tier		First 5 Tiers		All Tiers	
	Bid	Ask	Bid	Ask	Bid	Ask
2-Year	308	300	1,561	1,538	2,422	2,355
3-Year	82	81	474	467	715	684
5-Year	31	31	278	275	465	443
7-Year	37	36	236	236	301	297
10-Year	26	26	213	211	393	376
30-Year	3	3	28	28	59	59

The table reports average depth on BrokerTec at the first tier, the first five tiers, and across all tiers. The statistics are computed from five-minute snapshots of the limit order book for the respective securities for the hours 07:00-17:30 over the period 2010-2011. Depth is reported in millions of dollars.

Table 2.4: Unconditional and Conditional Percentage of Hidden Depth

	Bid			**Ask**		
	Full Sample	Hidden > 0	% of Obs.	Full Sample	Hidden > 0	% of Obs.
			A: At First Tier			
2-Year	10.2	16.8	45.4	8.2	15.3	39.4
3-Year	9.0	25.1	24.5	7.1	23.4	21.6
5-Year	12.2	38.7	19.7	10.6	37.7	18.1
7-Year	3.9	27.3	10.7	3.4	25.3	10.0
10-Year	11.9	38.6	19.1	10.7	38.5	17.3
30-Year	22.9	68.5	13.8	22.2	69.9	12.7
			B: Across All Tiers			
2-Year	5.4	6.0	85.7	4.1	5.0	77.5
3-Year	5.8	7.7	69.5	4.3	6.7	58.7
5-Year	10.2	11.9	80.2	7.3	9.4	72.7
7-Year	3.1	7.2	39.8	2.7	6.5	38.7
10-Year	10.4	11.6	85.2	8.8	10.2	80.1
30-Year	17.8	21.7	73.8	14.9	19.2	69.2

The table reports the percentage of depth that is hidden at the first tier (panel A) and across all tiers (panel B). Column “Full Sample” shows the percentage of hidden depth based on all observations, while column “Hidden $>$ 0” shows the percentage of hidden depth based on observations with positive hidden depth only. The percentage of observations with positive hidden depth is reported in column “% of Obs’. The statistics are computed from five-minute snapshots of BrokerTec’s limit order book for the respective securities for the hours 07:00-17:30 over the period 2010-2011.

Table 2.5: Baseline Price Impact of Trades

	Trade Direction	Signed Trade Volume
2-Year	0.357	0.006
3-Year	0.486	0.017
5-Year	0.708	0.028
7-Year	1.302	0.078
10-Year	1.340	0.066
30-Year	2.921	0.450

The table reports 50-tick cumulative price impact of trades using a bivariate VAR(5) model of trade and return (based on the best bid-ask midpoint), with two alternative measures for the trade variable: 1) trade direction (1 for buys and -1 for sells), and 2) signed trade volume (positive for buys and negative for sells). Cumulative price impact is in 256ths of one percent of par. Estimation is based on BrokerTec data for the period 2010-2011.

Table 2.6: Separate Price Impact of Trade Direction and Size

	Trade Direction	\$1M Increment
2-Year	0.271	0.003
3-Year	0.390	0.008
5-Year	0.571	0.013
7-Year	1.132	0.031
10-Year	1.033	0.033
30-Year	2.738	0.093

The table reports 50-tick cumulative price impact of trade direction (buy) and size separately, using a trivariate VAR(5) model of return (based on the best bid-ask midpoint), trade direction and signed trade volume. Trade direction is 1 for buys and -1 for sells. Signed trade volume is the volume of trade, signed positive for buys and negative for sells. Cumulative price impact is in 256ths of one percent of par. Estimation is based on BrokerTec data for the period 2010-2011.

Table 2.7: Price Impact of Buyer-Initiated versus Seller-Initiated Trades

	Buy	Sell
2-Year	0.006	-0.005
3-Year	0.017	-0.016
5-Year	0.028	-0.027
7-Year	0.081	-0.075
10-Year	0.066	-0.065
30-Year	0.435	-0.464

The table reports 50-tick cumulative price impact of buyer-initiated versus seller-initiated trades using a VAR(5) model of buy trade volume, sell trade volume and return (based on the best bid-ask midpoint). Cumulative price impact is in 256ths of one percent of par. Estimation is based on BrokerTec data for the period 2010-2011.

Table 2.8: Bid and Ask Price Impact of Buys and Sells (VECM)

	Buy		Sell	
	Bid Price	Ask Price	Bid Price	Ask Price
2-Year	0.005	0.005	-0.005	-0.005
3-Year	0.016	0.016	-0.016	-0.016
5-Year	0.029	0.029	-0.028	-0.028
7-Year	0.081	0.081	-0.077	-0.076
10-Year	0.068	0.068	-0.068	-0.067
30-Year	0.451	0.449	-0.483	-0.485

The table reports 50-tick cumulative bid and ask price impact of buyer-initiated versus seller-initiated trades using a VECM(5) model of bid and ask price revisions, buy trade volume and sell trade volume, with the bid-ask spread as the error correction term. Cumulative price impact is in 256ths of one percent of par. Estimation is based on BrokerTec data for the period 2010-2011.

Table 2.9: Price Impact of Trades and Limit Orders

	Buy Trade	Sell Trade	Bid Limit Order	Ask Limit Order
2-Year	0.0052	-0.0049	0.0011	-0.0011
3-Year	0.0132	-0.0131	0.0078	-0.0053
5-Year	0.0245	-0.0236	0.0139	-0.0119
7-Year	0.0592	-0.0544	0.0285	-0.0275
10-Year	0.0579	-0.0570	0.0353	-0.0341
30-Year	0.3986	-0.4258	0.3491	-0.2801

The table reports 50-tick cumulative price impact of trades and limit orders using a VAR(5) model of buy trade volume, sell trade volume, bid limit order flow, ask limit order flow and return (based on the best bid-ask midpoint). The limit order flow variables are measured as the total volume of limit orders submitted to the inside tier between trades, net of modifications/cancellations. Cumulative price impact is in 256ths of one percent of par. Estimation is based on BrokerTec data for the period 2010-2011.

Table 2.10: Price Impact of Trades After FOMC Announcements

	Buy Trade	Sell Trade	Bid Limit Order	Ask Limit Order
		A: On FOMC Days		
2-Year	0.0049	-0.0069	0.0018	-0.0024
3-Year	0.0153	-0.0264	0.0176	-0.0083
5-Year	0.0258	-0.0283	0.0150	-0.0248
7-Year	0.0749	-0.0592	0.0567	-0.0281
10-Year	0.0504	-0.0702	0.0587	-0.0508
30-Year	0.1925	-0.5242	0.5640	-0.4224
		B: On Non-FOMC Days		
2-Year	0.0037	-0.0042	0.0006	-0.0005
3-Year	0.0123	-0.0121	0.0058	-0.0032
5-Year	0.0210	-0.0237	0.0096	-0.0121
7-Year	0.0512	-0.0564	0.0276	-0.0289
10-Year	0.0541	-0.0576	0.0300	-0.0285
30-Year	0.2134	-0.2732	0.2828	-0.2976

The table reports 50-tick cumulative price impact of trade and limit order flow after 16 scheduled FOMC announcements over the period 2010-2011 (Panel A), and compare with similarly calculated price impact of trade and limit order flow over the same time interval on 5 days preceding and 5 days following these announcements (Panel B). The price impact estimates are based on a VAR(5) model of return (based on the best bid-ask midpoint), buy volume, sell volume, net limit order flow to the inside bid, and net limit order flow to the inside ask. Cumulative price impact is in 256ths of one percent of par. Estimation is based on BrokerTec data for the 90-minute window that begins 2 minutes after the announcement time of each FOMC announcement. The announcement times are collected from Bloomberg.

Table 2.11: Descriptive Statistics of Normal versus Iceberg Orders

	2-Year		3-Year		5-Year	
	Normal	Iceberg	Normal	Iceberg	Normal	Iceberg
Percent Placed Inside Spread	1.18	2.57	2.40	4.58	2.74	10.25
Total Size ($M)	4.4	13.5	2.5	8.8	1.6	8.8
Hidden Size ($M)	0.0	9.4	0.0	6.1	0.0	6.7
Spread (cents/$100)	0.8	0.9	0.9	1.0	0.9	1.0
Same Size Depth ($M)	246	227	65	66	36	34
Opposite Size Depth ($M)	264	210	88	59	32	29
Inter-Order Duration (secs)	2.4	5.0	2.1	5.2	0.9	2.3
Past Realized Vol. (ann.)	0.03	0.03	0.06	0.05	0.09	0.08
Past Trading Rate (#Trades/5min)	8	7	12	10	24	21
Past Trading Volume ($M/5min)	292	228	183	141	369	323
No. of Orders Per Day	62,497	1,244	69,159	1,018	172,152	1,454

	7-Year		10-Year		30-Year	
	Normal	Iceberg	Normal	Iceberg	Normal	Iceberg
Percent Placed Inside Spread	2.04	8.91	2.93	15.98	16.83	53.26
Total Size ($M)	1.5	5.3	1.4	11.6	1.2	7.1
Hidden Size ($M)	0.0	3.7	0.0	9.3	0.0	5.8
Spread (cents/$100)	1.9	2.2	1.8	2.0	4.6	4.8
Same Size Depth ($M)	39	38	29	27	5	3
Opposite Size Depth ($M)	34	34	26	25	3	3
Inter-Order Duration (secs)	1.0	3.0	1.0	3.4	3.9	11.7
Past Realized Vol. (ann.)	0.14	0.13	0.18	0.16	0.30	0.27
Past Trading Rate (#Trades/5min)	13	12	23	20	18	16
Past Trading Volume ($M/5min)	92	89	311	262	57	52
No. of Orders Per Day	145,062	759	151,719	767	38,211	432

The table reports descriptive statistics based on limit orders submitted to the first tier of the order book on the BrokerTec platform on 100 days randomly selected from the 500 trading days spanning the sample period 2010-2011. Past realized volatility, trading rate and trading volume are calculated over the five minute interval before each order submission. Inter-order duration is the prevailing average wait time between orders on the same side, measured in seconds and averaged over the previous three orders on the same side. Past realized volatility is the square root of the past five-minute realized variance computed as the five-minute sum of squared one-second log midquote returns and annualized by a factor of 288×250 (288 five-minute intervals per day and 250 trading days per year).

Table 2.12: Hidden Order Choice Model

	2-Year		3-Year		5-Year		7-Year		10-Year		30-Year	
	Coeff.	Odds Ratio	Coeff.	Odds Ratio	Coeff.	Odds Ratio	Coeff.	Odds Ratio	Coeff.	Odds Ratio	Coeff.	Odds Ratio
CONST	−4.38*		−4.91*		−5.46*		−5.79*		−6.22*		−5.94*	
IMP	−0.26*	0.77	−0.64*	0.53	−0.44*	0.64	0.59*	1.80	−0.20*	0.82	0.72*	2.05
SIZE	0.82*	2.27	0.89*	2.44	0.85*	2.34	0.83*	2.29	0.80*	2.22	0.84*	2.32
SPR	0.09*	1.09	0.17*	1.19	0.14*	1.15	0.11*	1.12	0.10*	1.11	−0.09*	0.91
DSAME	0.01*	1.01	−0.01*	0.99	−0.16*	0.85	−0.21*	0.81	−0.21*	0.81	−0.45*	0.64
DOPP	−0.21*	0.81	−0.29*	0.75	−0.05*	0.95	0.04*	1.04	−0.03*	0.97	0.04*	1.04
VOLA	−0.11*	0.90	−0.20*	0.82	−0.17*	0.84	−0.19*	0.83	−0.11*	0.90	−0.11*	0.90
NTRANS	-0.01	0.99	0.02*	1.02	−0.01*	0.99	0.03*	1.03	−0.10*	0.90	−0.07*	0.93
WAIT	0.04*	1.04	0.01*	1.01	0.04*	1.04	0.03*	1.03	0.05*	1.05	0.05*	1.05
PRENEWS	−0.37*	0.69	−0.37*	0.69	−0.30*	0.74	−0.12*	0.89	−0.71*	0.49	-0.10	0.91
POSTNEWS	−1.15*	0.32	−0.99*	0.37	−1.06*	0.35	−1.25*	0.29	−0.90*	0.41	−0.34*	0.71
OFFHR	0.24*	1.27	0.44*	1.55	0.08*	1.08	−0.35*	0.71	0.20*	1.22	0.72*	2.06
Pseudo R^2	0.14		0.18		0.23		0.17		0.32		0.45	

This table reports the result of the logistic regression: $ICE_i = f(\beta_0 + \beta_1 IMP_i + \beta_2 SIZE_i + \beta_3 SPR_i + \beta_4 DSAME_i + \beta_5 DOPP_i + \beta_6 VOLA_i + \beta_7 NTRANS_i + \beta_8 WAIT_i + \beta_9 PRENEWS_i + \beta_{10} POSTNEWS_i + \beta_{11} OFFHR_i)$, where ICE_i is equal to 1 if the newly submitted limit order i contains some hidden depth, and 0 otherwise. IMP is equal to 1 if the order is placed inside the prevailing spread, and 0 otherwise. $SIZE$ is the total order size (logged). SPR is the prevailing bid-ask spread, measured in basis points of the inside mid quote. $DSAME$ is the prevailing inside depth on the same side as the order, $DOPP$ is the prevailing inside depth on the opposite side of the order. $VOLA$ is the prevailing five-minute realized volatility based on one-second mid-quote log return. $NTRANS$ is the number of transactions over the preceding five minutes. $WAIT$ is the average wait time in seconds between limit order submissions on the same side, averaged over the preceding three same-side limit orders. $PRENEWS$ and $POSTNEWS$ are indicator variables with a value of 1 if the order is within the five-minute window before and after an announcement respectively (see Appendix A for the list of announcements considered). $OFFHR$ is equal to 1 if the order is submitted outside New York trading hours (7:00-17:30 ET). To ensure comparability across days, $SIZE$, SPR, $DSAME$, $DOPP$, $VOLA$, $WAIT$, $NTRANS$ are standardized by the corresponding daily mean and variance. The model is estimated based on limit orders submitted to the first tier of the order book on the BrokerTec platform on 100 days randomly selected from the 500 trading days spanning the sample period 2010-2011. Asterisk $*$ indicates statistical significance at the 5% level.

Table 2.13: Determinants of Hidden Size Conditional on Hidden Order Choice

	2-Year		3-Year		5-Year		7-Year		10-Year		30-Year	
	β	e^β	β	e^β	β	e^β	β	e^β	β	e^β	β	e^β
CONST	0.400*	1.49	−0.076*	0.93	−0.456*	0.63	−0.399*	0.67	−0.537*	0.58	−0.593*	0.55
IMP	0.041*	1.04	0.096*	1.10	0.083*	1.09	0.086*	1.09	0.042*	1.04	0.026*	1.03
SIZE	0.875*	2.40	0.744*	2.10	0.576*	1.78	0.507*	1.66	0.489*	1.63	0.388*	1.47
SPR	−0.003*	1.00	−0.005*	1.00	-0.002	1.00	0.021*	1.02	-0.003	1.00	−0.007*	0.99
DSAME	0.005*	1.00	−0.018*	0.98	−0.022*	0.98	−0.036*	0.96	0.001	1.00	0.019*	1.02
DOPP	−0.002*	1.00	0.009*	1.01	0.000	1.00	−0.017*	0.98	−0.013*	0.99	-0.002	1.00
VOLA	0.008*	1.01	0.009*	1.01	−0.006*	0.99	-0.003	1.00	0.018*	1.02	0.015*	1.02
NTRANS	−0.006*	0.99	−0.010*	0.99	−0.005*	0.99	−0.019*	0.98	−0.022*	0.98	0.007*	1.01
WAIT	0.001	1.00	0.001	1.00	0.003*	1.00	0.004*	1.00	0.003*	1.00	0.001	1.00
PRENEWS	-0.019	0.98	−0.049*	0.95	0.006	1.01	−0.070*	0.93	−0.046*	0.95	0.055*	1.06
POSTNEWS	−0.046*	0.96	−0.027*	0.97	-0.000	1.00	-0.014	0.99	−0.052*	0.95	0.015	1.02
OFFHR	−0.033*	0.97	0.011*	1.01	-0.005	1.00	−0.047*	0.95	0.003	1.00	0.027*	1.03
R^2	0.901		0.874		0.912		0.855		0.907		0.889	

This table reports the result of the regression: $HSIZE_i = \beta_0 + \beta_1 IMP_i + \beta_2 SIZE_i + \beta_3 SPR_i + \beta_4 DSAME_i + \beta_5 DOPP_i + \beta_6 VOLA_i + \beta_7 TRADE_i + \beta_8 WAIT_i + \beta_9 PRENEWS_i + \beta_{10} POSTNEWS_i + \beta_{11} OFFHR_i$, where $HSIZE_i$ is the hidden size of order i (logged). IMP is equal to 1 if the order is placed inside the prevailing spread, and 0 otherwise. $SIZE$ is the total order size (logged). SPR is the prevailing bid-ask spread, measured in basis points of the inside mid quote. $DSAME$ is the prevailing inside depth on the same side as the order, $DOPP$ is the prevailing inside depth on the opposite side of the order. $VOLA$ is the prevailing five-minute realized volatility based on one-second mid-quote log return. $NTRANS$ is the number of transactions over the preceding five minutes. $WAIT$ is the average wait time in seconds between limit order submissions on the same side, averaged over the preceding three same-side limit orders. $PRENEWS$ and $POSTNEWS$ are indicator variables with a value of 1 if the order is within the five-minute window before and after an announcement respectively (see Appendix A for the list of announcements considered). $OFFHR$ is equal to 1 if the order is submitted outside New York trading hours (7:00-17:30 ET). To ensure comparability across days, $SIZE$, SPR, $DSAME$, $DOPP$, $VOLA$, $WAIT$, $NTRANS$ are standardized by the corresponding daily mean and variance. The model is estimated based on limit orders submitted to the first tier of the order book on the BrokerTec platform on 100 days randomly selected from the 500 trading days spanning the sample period 2010-2011. Asterisk $*$ indicates statistical significance at the 5% level.

Figure 2.1: Trading Activity Over Time

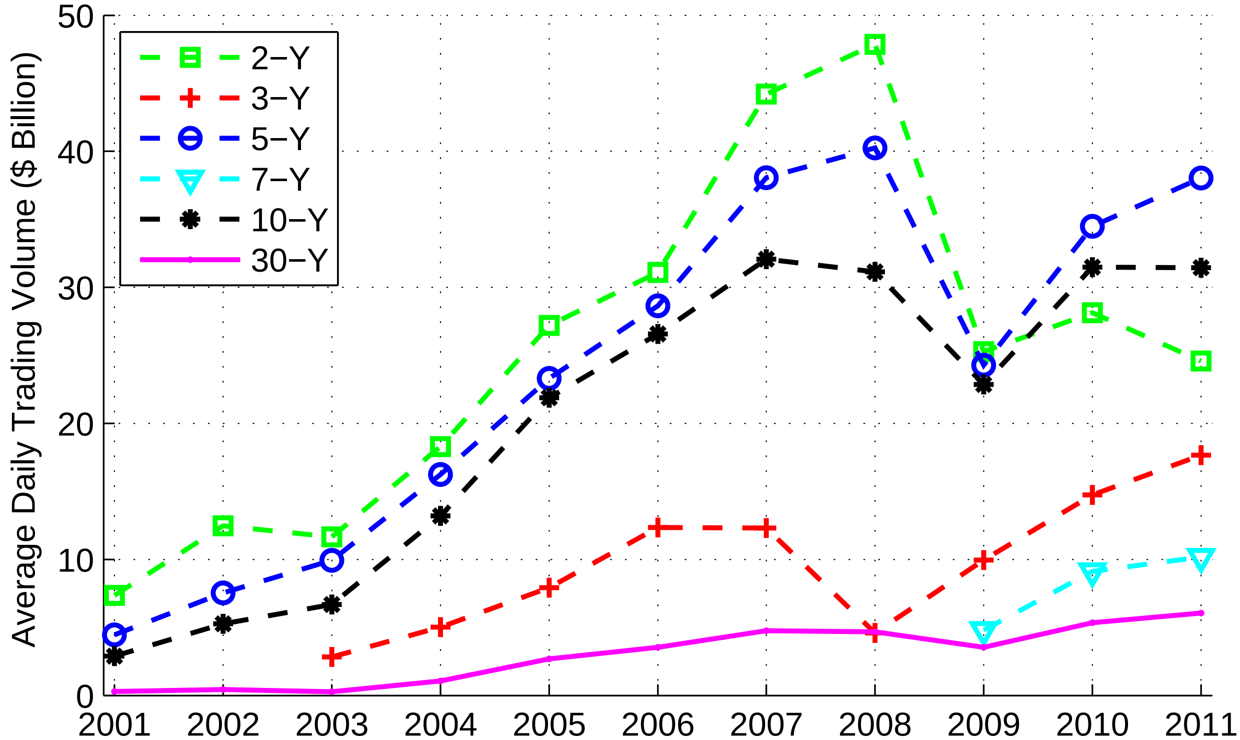

The figure shows average daily trading volume by year in billions of dollars from 2001 through 2011 for on-the-run Treasury coupon securities on the BrokerTec platform. The 2007 and 2008 figures for the 3-year note are based on data through August 2007 and from November 2008 respectively, due to the suspended issuance of this note between August 2007 and November 2008. The 2009 figure for the 7-year note is based on data from February 2009, when this note was reintroduced.

Figure 2.2: Round-the-Clock Trading Activity

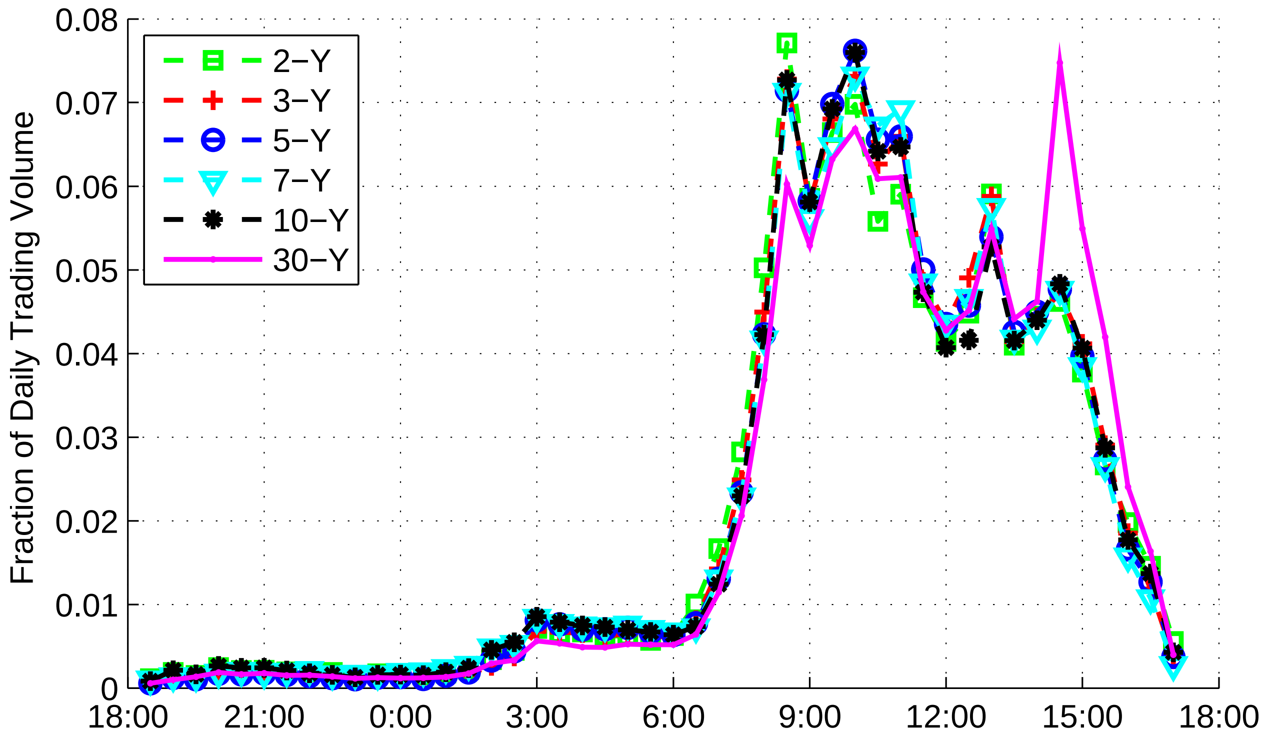

The figure shows the fraction of daily total trading volume by half-hour interval for on-the-run Treasury coupon securities on the BrokerTec platform, based on data for the period 2010-2011. Times are Eastern time and indicate start of half-hour interval.

Figure 2.3: Frequency Distribution of Inside Spread

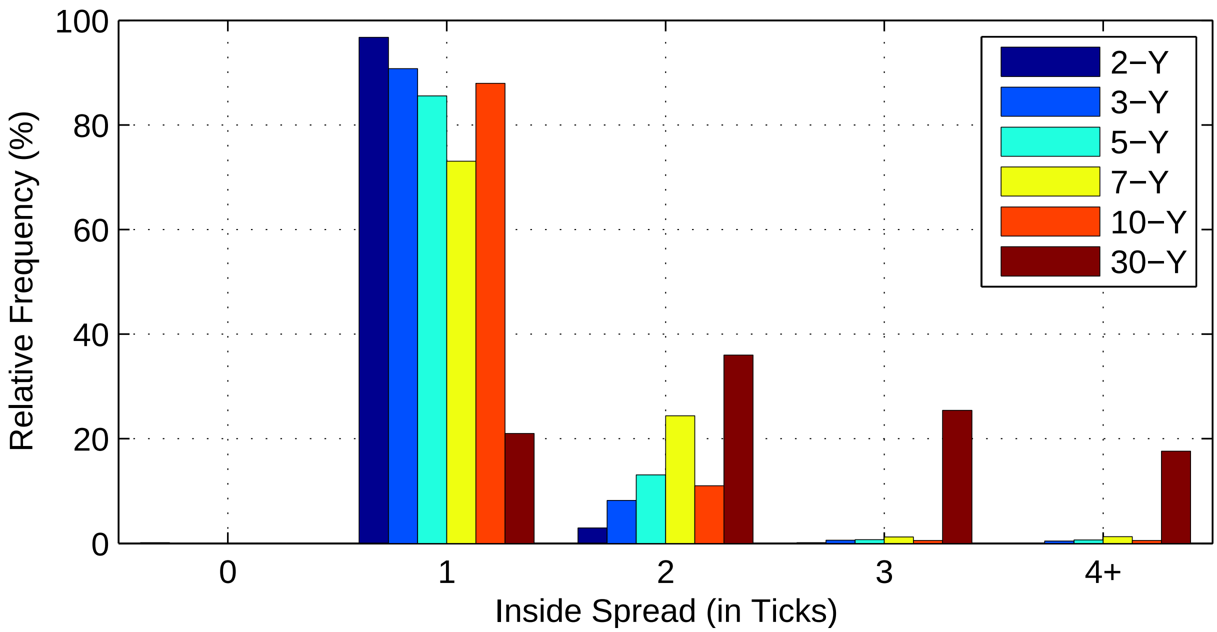

The figure shows the frequency distribution of the inside spread (measured in number of ticks) on BrokerTec. The tick size for the 2-, 3- and 5-year maturities is 1/128th of one percent of par, and that for the 7-, 10- and 30-year maturities is 1/64th of one percent of par. The numbers are computed from five-minute snapshots of BrokerTec's limit order book for the respective securities for the hours 07:00-17:30 over the period 2010-2011.

Figure 2.4: Displayed and Hidden Liquidity at the First Five Tiers

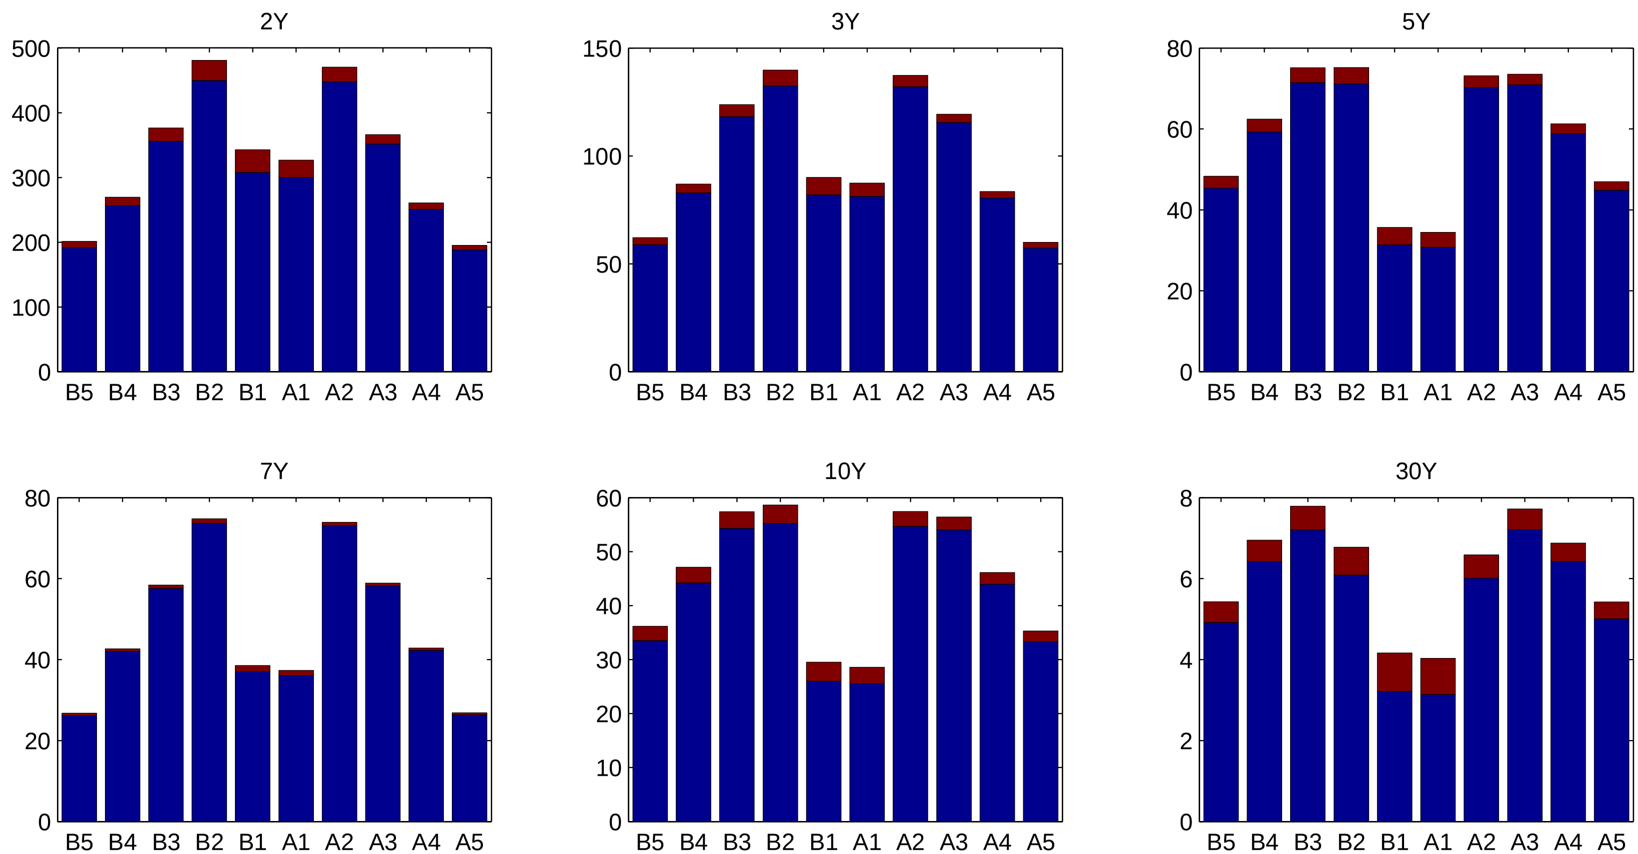

The figure depicts the average displayed depth (blue) and hidden depth (red) by price tier up to the fifth level on each side of the market. The numbers are computed from five-minute snapshots of BrokerTec's limit order book for the respective securities for the hours 07:00-17:30 over the period 2010-2011. Depth is reported in millions of dollars.

Figure 2.5: Non-linear Price Impact of Trade Size

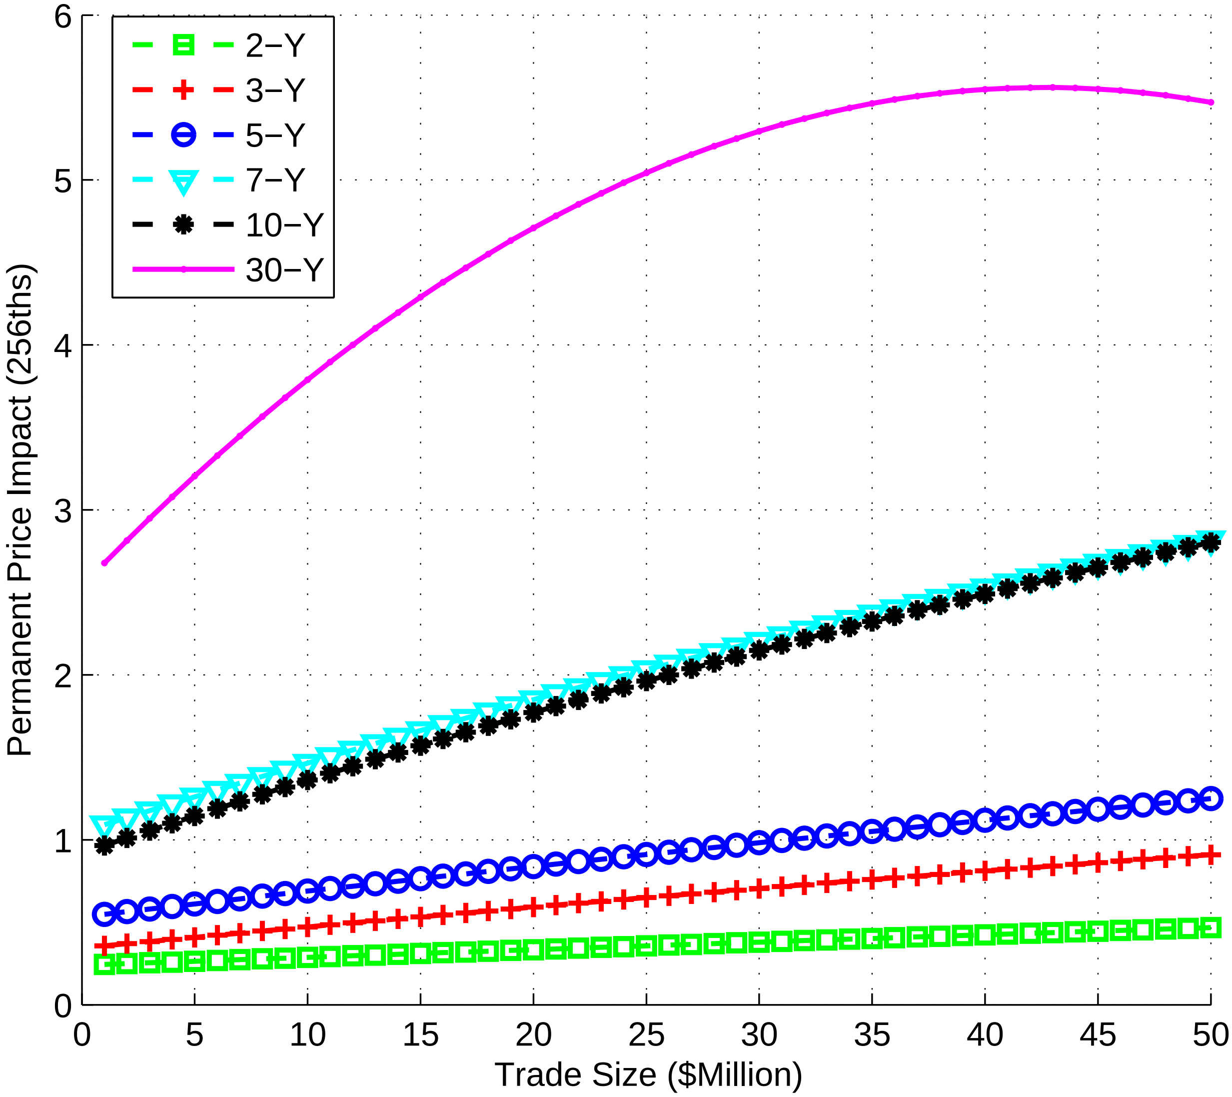

The figure plots the permanent price impact (y-axis) for different buyer-initiated trade sizes (x-axis). The permanent price impact of a given trade size is the cumulative price change (measured in 256ths of one percent of par) over a 50-tick horizon following the trade. This is based on a VAR(5) model of return (based on the best bid-ask midpoint), trade direction, signed trade volume and signed trade volume squared. Trade direction is 1 for buys and -1 for sells. Signed trade volume is the volume of trade, signed positive for buys and negative for sells. Signed trade volume squared is the squared volume of trade, signed positive for buys and negative for sell. Estimation is based on BrokerTec trade data for the period 2010-2011.

CHAPTER 3
DARK POOL TRADING IN THE US TREASURY MARKET

3.1 Introduction

The ability to hide trading intention is important and valuable to market participants. As Harris (1997) writes: "the art of trading lies in knowing when and how to expose trading interest." Many trading venues provide features that enable market participants to manage their exposure, ranging from hidden orders to dark pools.

The BrokerTec platform, one of two interdealer electronic trading platforms for U.S. Treasury securities, is one such place where traders can conceal the extent of their trading interest. One way they can do this is through iceberg orders. An iceberg order is a limit order that displays only a portion of the order quantity, called the display size, with the rest invisible to the market. The hidden quantity is revealed only gradually to the market as the displayed size is fully executed and the next installment becomes visible.

A second way traders can conceal the extent of their trading interest is through the workup process. A workup is a protocol that automatically opens after the execution of each market order. During the workup window, any interested market participants can transact additional volume at the same price established by the initial execution, as long as counter trading interest exists. Workups thus provide traders the option to submit orders of smaller size than desired, and then to increase the size when the workup opportunity opens.

It is intriguing to observe that iceberg orders are used sparingly in this market as compared to workups, even though the former has higher execution priority. On average, less than 5% of transactions involve execution against iceberg orders, whereas volume expansion through the workup protocol happens 49-56% of the time for the on-the-run 2-, 5- and 10-year notes and nearly 40% of the time for the 30-year bond.

The economic significance of the workup protocol is demonstrated not only by its frequent usage, but also by the magnitude of the expanded volume. On average, worked-up trading volume accounts for 43-56% of total daily trading volume, depending on the security being traded. In transactions with a workup, the worked-up dollar volume accounts for over 60% of total transaction volume. Collectively, these statistics show that the workup protocol uncovers a significant amount of market liquidity that is not ex ante observable to market participants. More importantly, the workup feature relates to a salient fact about the Treasury limit order book: market orders rarely walk up or down the book, at least for the on-the-run securities examined in this chapter.

The fact that a significant portion of market liquidity is only revealed during the workup process, while the price is fixed, raises an important question about its role in the price formation process. Does this portion of order flow carry information? If so, how does it compare to the "transparent" part (i.e., the execution of the initiating market order)? Up until now, the literature on price discovery of U.S. Treasury securities, such as Brandt and Kavajecz (2004), Green (2004), Pasquariello and Vega (2007), and others, has been concerned only with the informational role of generic order flow. Implicitly, it is assumed that the trade flow is homogeneous and that the portion transacted in the pre-workup stage has the same impact on price dynamics as that in the workup phase.

The theoretical literature on dark trading suggests that this is not the case. For example, Zhu (2014) argues that the different execution probability of orders in transparent versus dark venues for informed versus uninformed traders likely steers informed traders to the transparent venue and uninformed traders to the dark venue. In contrast, Ye (2012) predicts that informed traders are more likely to hide their information in the dark, thereby reducing price discovery. Although arriving at opposing predictions, both models suggest that the information content of the transparent and dark parts of trading is different as a result of how informed traders optimize their exposure strategy to exploit their information advantage.

It is also not realistic to study the information role of trades in this market without recognizing or appreciating that the workup feature is an integral part of the trading process and a useful device for order exposure management. Accordingly, it can alter traders' optimization outcomes and decisions, as compared to a hypothetical market setup in which this option is not available. For example, aside from choosing between market and limit orders, traders also have choices over the exposure of their trading intention. If it is a limit order, the submitter can 1) display the full order size, 2) submit the full order size but hide part of it from view (i.e., submit an iceberg order), or 3) submit a smaller sized order and wait to expand the order size if and when the order is executed and the workup is open. Likewise, if it is an aggressive order, the available choices are: 1) submit the full sized market order for immediate execution, 2) submit a smaller sized market order and hope to increase the volume in the ensuing workup, or 3) wait for a workup at the right price to trade. The fact that workups are used so frequently in this market speaks for its importance in traders' decision making.

The workup protocol, by providing traders with an option to expand order size beyond the initially submitted level, can be beneficial to informed traders in multiple ways. For example, an informed trader may choose to submit a limit order of less than the intended quantity to minimize information free-riding by others, knowing that he has the opportunity to increase the size when the order is aggressed against (i.e., executed by a market order). These are the "informed liquidity providers" as discussed in Boulatov and George (2013). Alternatively, an impatient informed trader with a large trading interest may submit a market order small enough to execute at the best price, and then search for further counter trading interest during the subsequent workup at the same price, without having to walk deeper into the book.

The informativeness of the workup order flow, however, also depends on the actions of the uninformed. If they behave in the sense of Zhu (2014), or if they are reactive traders who act upon the lead of others as described in Harris (1997), the uninformed are more likely to trade in the workup stage, thereby reducing the informativeness of workup order flow. Likewise, the uninformed can participate in the market on the liquidity provision side, and thus volume expansion of limit orders during workups is not necessarily information motivated.

Given the mixed guidance from theory, the informativeness of workup order flow in practice is an open question. Our work in this chapter aims to address this question and thereby provide a more complete picture of the price formation process of U.S. Treasury securities. We do this by separating order flow into the trade initiation (or "transparent") part and the workup (or "dark") part, and quantifying how these respective components contribute to price discovery.

Based on transaction data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities from the BrokerTec platform over the period 2006-2011, we find that workup order flow is informative, albeit less informative than trade initiation order flow. Workup order flow is most informative for the 2-year note, explaining about 17% of the total variation of the efficient price update. The 5- and 10-year counterparts explain between 7-8% of the variation of the efficient price. It is only 1%, however, for the 30-year bond, indicating that traders in this maturity segment do not opt for the workup as a channel to exploit their information advantage. Price discovery in the 30-year segment is most attributable to public information, which accounts for over 80% of the variation of the efficient price update. Our analysis also illustrates the importance of recognizing the information segmentation of order flow due to the workup protocol. We show that the impact of actively initiating a trade and the share of trade-related information are underestimated if one does not consider the segmentation.

Except for the 2-year note where the workup order flow is slightly more informative, our information share analysis generally supports the predictions of Zhu (2014) in that the transparent part of order flow is more informative than the dark part. Our evidence is also consistent with Comerton-Forde and Putnins (2013) who study dark trading and price discovery of stocks traded on the Australian Stock Exchange and find that the order flow that migrates to the dark is less informed, but not completely uninformed. Intuitively, informed traders with a short-lived information advantage may choose to initiate a trade and realize their information advantage quickly, since the potential of not finding a counter-party during the workup can make the workup option costly (e.g., forgone information value). This is strongly supported by our result that the transparent (or "lit") part of the order flow becomes relatively more informationally important on high volatility days, when adverse price movements could fasten the expiry of information.

Although less informative than the trade initiation flow, the workup process is responsible for the discovery of a significant portion of market liquidity and plays a non-trivial role in price discovery. It is therefore important to understand what factors might predict the use of a workup following a market order execution and the extent of the volume expansion. We employ a logistic regression model for the probability of a workup as a function of hypothesized determinants, including prevailing order book depth on the same side, spread, depth on the opposite side, price volatility and a set of control variables. We employ a Tobit model censored at the zero lower bound to capture the effects of the same set of explanatory variables on workup volume, since volume is zero for non-workup transactions and strictly positive for those with a workup. We find that, in general, a workup is more likely and expands greater volume when the market is deep, upon the discovery of hidden orders, and around the times of high trading intensity, volatility and workup activity. Outside New York trading hours, workups occur less frequently and discover less volume.

Despite workups being an unusual and economically important trading feature in the U.S. Treasury market, academic research specifically on the workup process and its implications on market participants' trading strategies is limited. The paper closest to ours is Boni and Leach (2004) who investigate the workup protocol using GovPX data for October 1997. At that time, interdealer trading was still conducted over a network of voice-assisted interdealer brokers (IDB). GovPX collected and disseminated market data from five such IDBs. Each of the IDBs maintained its own limit order book, facilitated trades and mediated quantity negotiations beyond the quoted depth. Boni and Leach hence characterize this market as one in which limit orders are “expandable”.[1] They find that this expandability option helps limit order traders reduce costs associated with information leakage and adverse execution of stale limit orders.

Dungey et al. (2013) account for the workup feature in their model of trading intensity on eSpeed, the other electronic trading platform. They find that the duration of a workup significantly fastens the arrival of the next market order. They interpret this result in light of Easley and O’Hara (1992)’s theory that market participants infer the presence of informed traders by the time between trading events. They thus suggest that workups provide information to the market.

Adding to this literature, our work is the first to consider the segmented order flow due to the workup protocol and formally quantify the informativeness of the dark trade flow in comparison with that of the transparent trade flow in a well defined microstructure model of the dynamics of price and order flow. We therefore provide a more complete picture of how the trading process with the embedded workup feature – a distinctive microstructure feature of secondary trading in U.S. Treasury securities – affects the price formation process of these securities.

Furthermore, we extend the model used by Boni and Leach by exploring and controlling for a wider range of potential determinants of workups as suggested by the literature on dark pool trading and hidden liquidity. The extension is also valuable because the workup protocol as prevailing on BrokerTec differs considerably from its historical precedent as studied in Boni and Leach, making it important to reevaluate the workup protocol under its current structure.

Our contribution goes beyond a study of a specific microstructure feature of the U.S. Treasury market in two important ways. First, this work is a timely addition to the literature on dark pool trading. There is currently an active discussion among researchers and policy makers on the effects and implications of dark pool activities on market quality and welfare (see for example Degryse et al. (2014)). On the one hand, the existence of undisplayed liquidity compromises pre-trade transparency and can potentially harm less informed

[1] BrokerTec’s workup protocol differs from that of voice-assisted brokers in several ways. First, quantity negotiation on BrokerTec is governed by a set of precise rules stipulating the window of opportunity for workup trades, replacing the role of human brokers in going back and forth between counterparties working up the size of a trade. Second, as explained by Boni and Leach, with the voice-assisted brokers, when a limit order on an IDB’s book is aggressed by a market order, the IDB gives the limit order’s submitter the right-of-first-refusal to provide additional liquidity, even when there are other limit orders at the same price in the book. This exclusivity was completely eliminated on BrokerTec in early 2006, making workups immediately open to all market participants following the original trade(s) on a first come, first served basis (source: “System and Method for Providing Workup Trading without Exclusive Trading Privileges”, patent number US 8,005,745 B1, dated August 23, 2011). Furthermore, the expanded volume can come from either the aggressive or passive side during a workup so that the workup is no longer confined to expanding only limit orders. Recently, BrokerTec instituted a new rule that allows for a workup to terminate prematurely if there is sufficient trading interest at a better price point (source: “System and Method for Providing Workup Trading”, patent number US 7,831,504 B1, dated November 9, 2010).

traders. On the other hand, supporters of dark pool trading mechanisms point to greater market liquidity and better execution quality for trades, especially for large trades that can be executed without causing negative price impact. Complicating matters, dark pools come in many forms. It is therefore important to understand these various types of dark pools in different market setups and how their specific operationalization might affect trading behavior and patterns.

The workup process in the Treasury market resembles a crossing network, a common form of dark pool.[2] As in a crossing network, workup trades are matched on a first come, first served basis at a reference price derived from the initial marketable order execution. While a crossing network is a common form of dark pool used in many equity trading venues, not much is known about this form of dark pool in a fixed income market setting. We find that volatility tends to generate more workups, but that those workups tend to be less informative, suggesting the value of this crossing network in protecting traders against adverse price movements. In general, the amount of private information hidden in this Treasury dark pool is quite small, easing concerns that the dark pool could harm less informed traders.

The second direction in which our findings might be valuable is in the area of market design response to high frequency trading. High frequency trading, or computer-driven trading in general, has increased significantly over the last few years – a trend dubbed "rise of the machines" in Chaboud et al. (2013). This has spurred debate on whether the competition for speed has resulted in socially wasteful investments in trading technology. Budish et al. (2014) argue that a continuous limit order book is a flawed market design in the face of increased high frequency trading. They propose frequent batch auctions as a way to eliminate technical arbitrage opportunities prevailing at very high frequency and to slow down the arms race. Whether this is a feasible proposal is outside the scope of our work, but the important takeaway from this discussion is that there is a continual need to devise new market design features to keep up with changing trends in trading, and to understand the implications of those features.

The BrokerTec market setup fits right into this discussion via its unique design. Its market structure features an interesting mix of continuous auction (the limit order book) and periodic call auctions (workups).

[2] Buti et al. (2011b) characterize dark pools as having "limited or no pre-trade transparency, anonymity and derivative pricing." The workup process has precisely these characteristics. First, the workup process enables execution of additional trading interest not observed by market participants before each transaction. Second, all trades through interdealer brokers in the Treasury market are anonymous. During the allowable workup time window, market participants can send in orders, which are then matched by the system. Any unmatched volume is held in the system waiting for subsequent counter trading interest in the workup. Third, the price for these workup executions derives from the execution of the initial market order that triggers the workup. We thank Joel Hasbrouck for this insight.

These periodic call auctions open and close after every marketable order execution, allowing for the discovery of additional liquidity. This can potentially help slow down activities of high frequency traders and give other traders an opportunity to trade at a given price point. Viewed from this angle, the workup protocol might be one possible design response to high frequency traders and the arms race discussed above. Our empirical results readily provide a glimpse of the implications of such a market feature on price discovery. Even though price is fixed during workups, activities during the workup window do contribute to price discovery, mostly by trades that expand the aggressive side. Nevertheless, the extent of this contribution is small in comparison to that of initiating trades that lead to workups.

The chapter is organized as follows. Section 3.2 describes the workup process in detail, discusses the data used for the analysis and presents key stylized facts concerning workups. Section 3.3 presents a microstructure model for the dynamics of segmented order flow and price, and analyzes the price impact and informativeness of the respective components of order flow. In Section 3.4, we model and discuss the dependence of workup probability and workup volume on market condition variables. Finally, Section 3.5 summarizes our key empirical findings and concludes the chapter.

3.2 The Workup Process

3.2.1 Market Overview

This chapter focuses on the interdealer trading segment of the secondary market for U.S. government securities.[3] Trading in this segment, especially in the on-the-run securities, occurs mostly on two electronic trading platforms, BrokerTec and eSpeed (Barclay et al. (2006)). Comparing BrokerTec trading statistics with those reported by other studies using eSpeed data (e.g., Dungey et al. (2013) and Luo (2010)), we observe a greater market share for BrokerTec across all four securities considered.[4] Dunne et al. (2011) compare price

[3] In this interdealer market, the majority of participants are Treasury securities dealers, some of whom are primary dealers with obligations to participate in Treasury securities auctions. We use the term "dealers" interchangeably with "traders" and "market participants", even though there are other participants in the market such as hedge funds.

[4] Comparison of BrokerTec and eSpeed activity for the same sample periods and trading hours shows that the market share of BrokerTec ranges between 57% to slightly over 60% for the 2-, 5- and 10-year notes. The market share in the 30-year bond is slightly lower, but it is still over 50%.

discovery on the two platforms using non-contiguous data for 2002, 2004 and 2005 and conclude that more price discovery takes place in the more active but less transparent BrokerTec platform.[5]

Both BrokerTec and eSpeed operate as electronic limit order markets with no designated market maker. Liquidity supply comes from the limit order book, which is a collection of limit orders at various price levels submitted by market participants. Execution of orders follows the price and time priority rule. Our empirical analysis is based on BrokerTec data and our discussion of how the market works is for the BrokerTec platform.

In submitting limit orders, market participants can choose to display order size either partially (iceberg orders) or completely. If the former, the rest of the order size is not observable by other market participants. As the displayed portion is exhausted through trading, the next installment of the order quantity becomes displayed. This process continues until the total order quantity is completely executed. It is noted that the hidden portion of an iceberg order takes precedence over the displayed part of orders queuing behind it at the same price level.[6]

Traders demanding liquidity can send in market orders. Market orders must be priced. That is, beside the order quantity and whether it is a purchase or sale, traders must specify a price.[7] When a market order arrives, it is matched with one or more limit orders standing on the opposite side at that price (or better), starting with the displayed depth before executing against any hidden depth. For example, consider a market order to *buy* \$100 million at a price of 25580 when there is \$30 million available at the best ask price of 25578 and another \$100 million at 25580.[8] The first \$30 million of the order will be matched with limit sell orders at 25578. Assuming there is no hidden depth at that price, the remaining \$70 million will be executed at 25580. If a limit sell order is an iceberg order, then upon the execution of the displayed portion, the next portion becomes visible and available for execution. Continuing with the above example, assume that there is \$15 million hidden depth at the best ask price of 25578. The market order will be executed as follows: \$45 million at 25578 (\$30 million displayed + \$15 million hidden), and \$55 million at 25580.

---

[5]These authors note that eSpeed does not have hidden orders, whereas BrokerTec allows such orders and is thus considered as having less pre-trade transparency.

[6]This is different from other market setups in which the hidden part of an iceberg order goes to the end of the queue when it becomes visible.

[7]In essence, they are marketable limit orders. In this chapter, we use the terms “market orders”, “marketable orders” and “aggressive orders” interchangeably.

[8]In the BrokerTec database, prices are stored in 256th’s of one percent of par value. The prices used in the example translate to 99.9140625 (25578/256) and 99.921875 (25580/256).

The execution of a market order is just the beginning of a transaction. Once all possible matches have been made (against displayed and hidden depth in the book, if any), the market then enters into a workup process during which additional volume at the same price can be transacted, until there is no further trading interest. Using the example above, once the original buy order execution of $100 million is complete, the workup protocol opens at the last price of 25580. As described in one of BrokerTec's patent documents relating to the workup protocol, the whole process is conceptually "a single deal extended in time".[9] The protocol is discussed in greater detail in the next subsections.

Finally, it is worth noting that BrokerTec charges a fee for order execution, and that this fee is trader-specific and not order-type specific.[10] That is, a trader is charged the same fee whether his order is a limit or a market order, and whether his order is executed in the pre-workup or workup stage. Therefore, for each trader, the fee is not a consideration when it comes to the choice of order type and exposure. However, traders with different levels of trading activity might be subject to different fee structures.

3.2.2 The Workup Process

The workup is a distinctive feature of trading in U.S. Treasury securities. The process automatically opens after each market order execution, giving all market participants the chance to transact additional quantity at the last price. The ability to transact additional quantity during the workup process thus enables traders to submit orders of smaller size than their desired quantity, and then expand the quantity during the workup. The workup protocol therefore offers a higher degree of control over if and when to trade the additional needed quantity, whereas iceberg orders are subject to the risk of being adversely executed before the order submitters have a chance to modify or cancel. However, the cost for the traders hoping to expand volume in a workup is that the incremental quantity they expect to transact may not materialize if counter trading interest is lacking. Thus, for those who need immediate execution, non-execution risk can be a major deterrent to using workups.

Historically, the workup process consisted of two distinct phases: 1) the private phase, which gave an exclusive right of first refusal to the original parties to the transaction, and 2) the public phase, which is open to all other market participants. However, in 2006, the private phase was replaced by a public phase, making

---

[9]"System and Method for Providing Workup Trading", U.S. Patent No. 7,831,504 B1, dated November 9, 2010.

[10]The exact fees are proprietary and we do not have information on these fees.

the workup a double-public process. As a result, a transaction will progress straight to the first public workup phase after all possible matches with the limit order book have been completed. During this phase, additional trading interest can come from either side of the market, and these extra trades are conducted on a first come, first served basis.

The first public workup phase is open for a pre-specified duration (it was 4 seconds from early 2006 until July 2011 when the duration was shortened to 3 seconds). If there is no trading interest, the workup process automatically expires at the end of this duration. However, if and when a new execution occurs during this time window, the second public phase commences and a new duration opens up. It is then re-settable each time a new execution occurs. This protocol allows the workup to last as long as there is trading interest at the same price point, or to terminate after a predetermined time period if no such interest exists so that the market can move forward.

All trades during a workup – triggered by the initial execution of a market order – are executed at the same price as that of the original market order. As the extra liquidity discovered during the workup process at a given price is not known to the market *ex ante*, the workup process can be likened to a dark pool trading mechanism. More precisely, it resembles a crossing network, in which the last price serves as the reference price for the execution of additional trades during the workup.

We treat the whole process from the initial execution to the expiration of the ensuing workup as a single transaction. Each transaction can involve multiple trades or order matches. For example, a market order can execute against multiple (smaller sized) limit orders. Each of these executions is recorded separately in the database and is referred to as a trade, an order match, or an execution. We refer to those trades (or matches) that happen before the workup as pre-workup trades. Other interchangeable terms for "pre-workup trades" include "transparent trades", "lit trades", "normal trades", and "non-workup trades". Those that occur during the workup phase are referred to as workup trades or "dark trades".

3.2.3 The BrokerTec Data

The tick data from BrokerTec contains records of all market activity, from limit order submission, cancellation, and modification, to matching with incoming market orders, time-stamped to the millisecond. We extract information on all trading activity from this raw, comprehensive database. There is a flag each time a market order arrives. Once automatic order matching with the limit order book completes, another flag indicates the commencement of the workup phase. Finally, when the workup expires, it is also flagged in the database. As a result, we are able to identify the complete sequence of activities pertaining to each transaction.

The trade direction of the original market order, i.e., whether the aggressive side is a buy or a sell, is also recorded, thus providing unambiguous signing of all pre-workup trades. The signing of trades executed during the workup process warrants some further discussion. BrokerTec considers these trades as an extension (time-wise and volume-wise) of the original execution. Hence, workup trades occur at the same price and follow the same trade direction as that of the original execution, even though they can arise from either side of the market. For example, if the original aggressive side is a buy, then the buy side remains the aggressive side in the workup, and the sell side is the passive side. Therefore, there is no confusion as to the signing of workup trades.

After identifying the sequence of activities for each transaction, we aggregate information for the transaction, separately for the pre-workup and workup phases. In particular, we count the number of trades as well as the total dollar volume exchanged in each respective phase. Furthermore, if there is execution against the displayed portion of an iceberg order resulting in the exposure of new depth, we mark the transaction as involving execution against an iceberg order.

The transaction data is then combined with the limit order book snapshots prevailing just before and after each transaction. The limit order book is reconstructed from the raw BrokerTec data by cumulating changes to the order book from the beginning of each trading day. Combining the limit order book data with the transaction data provides a complete picture of the market at each transaction, facilitating our empirical analysis of factors that are related to workup activities.

Our sample covers the period from 2006 to 2011. We focus the study on the on-the-run 2-, 5-, 10- and 30-year fixed principal securities. The on-the-run 3- and 7-year notes are excluded from our analysis due to discontinuity in issuance. Issuance of the 3-year note was suspended between May 2007 and November 2008. Issuance of the 7-year note was suspended between April 1993 and February 2009. We do not have access to comparable tick data for any other Treasury securities.

3.2.4 Univariate Analysis of Workup Activities

3.2.4.1 Trading and Workup Activities

An overview of market trading activity is presented in Table 3.5. Panel A shows the average daily total trading volume and number of transactions. The 2-, 5- and 10-year notes have comparable trading volume in the \$30-35 billion range. This far exceeds the average of \$5 billion in daily trading volume for the 30-year bond. The number of transactions per day varies across securities, from the low 1,000's range for the 2- and 30-year securities to over 2,600 for the 5- and 10-year notes. It follows that trading in the 2-year note tends to occur in much larger size than is the case for other securities.

We find that market participants utilize the workup protocol in 49-56% of transactions for the notes, but only 39% of transactions for the 30-year bond. The workup share of dollar volume happens to be similar to the share of transactions with workups, ranging from 48-56% for the notes, but only 43% for the bond.

To complement the sample average statistics on workups, we further examine the time series trend of the use of workups in Figure 3.1. There is a modest increase in the use of workups from early 2006 until late 2008, when the workup shares of transactions and order flow drop before partially bouncing back in early 2009. The patterns have been fairly stable since then, except for a sharp decline in workups for the 2-year note, to about 40% in the second half of 2011. Also evident from the figure is that workups in the 30-year bond happen less frequently that they do in the notes, but expand proportionally greater volume when they do occur.

Table 3.5, Panel A also reports the probability of transacting against an iceberg order for comparison, and illustrates that workups are used much more frequently in this market. The chance of hitting/taking an iceberg order is only around 4%, which is less than one tenth the probability of having a workup. Additionally, the workup protocol, in providing traders with the opportunity to expand transaction volume at a given price, likely contributes to the finding that transactions in these securities almost never execute at multiple prices, beside the fact that the market is often very deep relative to the size of most market orders.

Further details at the transaction level, with and without workups, are presented in Panels B and C respectively. We discuss first the transactions with workups (Panel B). This panel shows that the 2-year note has the largest dollar volume per transaction when there is a workup, at about \$42 million. This is more than double the size of a transaction in the 5- or 10-year notes and about eight fold the size of a typical transaction in the 30-year bond. The number of trades per transaction is just below 10 for the notes and about 4 for the 30-year bond. Roughly two thirds of trades occur in the workup phase.

Panel C shows that transactions without workups tend to be much smaller in size, in terms of both dollar volume and trade count, than those with workups. Moreover, transactions without workups tend to be somewhat smaller than even just the pre-workup portions of transactions with workups. For example, the 2-year note's average transaction size without a workup is about $12 million, compared with a $16 million pre-workup size and $42 million total size for transactions with a workup. This is consistent with Harris (1997)'s reasoning that small traders are usually not concerned with exposure issues, and importantly, the small size is of little interest to other traders. Small trades can also be absorbed more easily by outstanding limit orders. Consequently, small market orders are more likely to be executed without a workup.

Finally, it is useful to compare workups on BrokerTec with those on eSpeed as reported in Dungey et al. (2013) for the period from January 2006 to October 2006. BrokerTec's greater market share in terms of total trading volume masks the fact that trading is slightly less frequent on BrokerTec, but that an average transaction has a much greater size.[11] The likelihood of workups is a few percentage points higher on BrokerTec than on eSpeed. However, BrokerTec workups typically discover a slightly smaller proportion of transaction volume. Accordingly, the overall share of workup volume does not differ greatly between the two platforms.

3.2.4.2 Intradaily Pattern of Workup Usage

Figure 3.2 plots the probability of workup over the course of a typical trading day, from 18:30 of the previous day to 17:30 of the current day (Eastern Time – ET).[12] The figure shows that workups are most active between 8:30 and 15:00. Outside of New York hours, workup activity is markedly lower, with a mild increase occurring around the start of London trading at 3:00.

The lower workup usage in the overnight hours may be related to the low overall level of activity during the overnight hours (e.g., Fleming (1997), Fleming et al. (2014)). The workup protocol requires more active

[11] Our comparison shows that an average transaction in the 2-, 5- or 10-year note is over 40% larger on BrokerTec than on eSpeed, while that in the 30-year bond is about 14% larger.

[12] Fleming (1997) provides a description of the global trading day in U.S. Treasury securities. It starts at 8:30 local time in Tokyo, which is 18:30 EST (or 19:30 EDT) the previous day (Japan has not adopted daylight saving time). Trading then passes on to London at 8:00 local time, i.e., 3:00 ET. New York trading then starts at 7:30 and continues until 17:30. Statistics for the hour from 18:30-19:30 of the previous day are based on the periods over which the U.S. is on standard time.

monitoring of market activity and exercise of judgment on the part of traders, which are less likely to occur during these hours. Moreover, workups during the off hours are less likely to be filled due to lower market participation and hence lower chance of meeting counter trading interest.

3.2.4.3 Workups and Order Flow

There are further interesting stylized facts relating to trading and workups. Table 3.2 reports several pairwise correlations of interest. First, the signed order flow imbalance, measured by net order flow as a percentage of total order flow for each day, is only weakly related with workup usage, with the absolute correlation coefficient under 0.05 for three out of four securities. However, the absolute order imbalance shows a much stronger correlation with the use of workups: except for the 30-year bond, the correlation coefficient is in the negative 0.2-0.3 range. That is, we tend to see relatively more workup activities on days when the market is balanced than when it is one-sided, whereas the direction of the imbalance does not matter much. This observation can be interpreted in light of Sarkar and Schwartz (2009)'s notion of market sidedness as an indication of asymmetric information, as informed traders tend to collect on one side of the market. If so, they are more likely to initiate trades to exploit their information advantage quickly, as opposed to trade in workups or post expandable limit orders.

Secondly, workups tend to be used relatively more frequently on more volatile days. This is illustrated by the strong positive correlation coefficients across the four securities, ranging from 0.26 for the 30-year bond to 0.54 for the 2- and 10-year notes. Finally, we also observe positive first order auto-correlation in workup activities, consistent with a liquidity externality effect of workup trades as predicted by Buti et al. (2011a). Specifically, increased workup activities imply that it is relatively easier to find counter trading interest in workups, thereby increasing the execution probability, and hence attractiveness, of workup orders.

3.2.4.4 Direction of Workup Volume Expansion

Given the current workup setup on BrokerTec, any trader from either side can join an open workup. Accordingly, volume can be expanded from either the limit order book side, or the aggressive side of the transaction. It is informative to examine the direction of volume expansion in a workup because it is ultimately linked to the degree of pre-trade transparency with respect to liquidity: the level of available liquidity market participants can see before a trade versus what actually shows up in the trade. Moreover, workup volume expansion from the aggressive side suggests the extent to which other traders follow the lead of the initial aggressive trader. It provides an indication for the amount of inactive trading interest which gets revealed only when someone else has initiated a trade. Most importantly, the ability to work up volume on either side of a transaction is one of the key features that differentiates BrokerTec's workup protocol from its voice-assisted predecessor.[13]

Figure 3.3 provides an analysis of the mix of workups with respect to the direction of workup volume expansion. We classify workups into three categories: 1) expanding volume on the aggressive side only, 2) expanding volume on both sides, and 3) expanding volume on the passive side (or both). Specifically, if a transaction's total volume is not greater than the available depth, all of the workup trades must have come from the aggressive side (category 1). If, instead, a transaction's total volume is greater than the available depth, the limit order book must have been expanded during the workup. Whether or not the aggressive side is also expanded can be determined by examining the pre-workup volume. If the pre-workup volume is less than the available depth, it is clear that the workup also expands the aggressive side (category 2). However, if the pre-workup trades completely wipe out the available depth, it is less clear whether the aggressive side is also expanded during the workup, although the passive side is certainly expanded (category 3). Accordingly, the sum of categories 1 and 2 provides a lower bound for the fraction of workups that expand the aggressive side, whereas the sum of categories 2 and 3 equals the percentage of workups that involve expansion on the passive side.

As can be seen from the figure, there is a cross-maturity variation in the direction of workup volume expansion. For the 2-year note, the majority of workups (at least 73%) expands the aggressive side, including the 53% of workups that expand only the aggressive side. Workups that expand the passive side occur 47% of the time. On the other end of the maturity spectrum, for the 30-year bond, workups mostly expand the limit order book (78%). Instances where only the aggressive side is expanded account for only 22% of workups. The 5- and 10-year notes are in the middle, with nearly 40% of workups expanding the aggressive side only, 23% expanding both sides, and another nearly 40% expanding the passive side or both.

These statistics show that aggressive workups are common for the notes (especially the 2-year) but not for the 30-year bond. In addition, workup trades often come from both sides within a given transaction,

---

[13]This is also where BrokerTec's workup protocol differs from eSpeed's. As described in one of eSpeed's patent documents, a market order needs to be sufficiently large to exhaust all displayed passive orders at the best price in order to trigger a workup, during which the initial parties to the trade are granted the right of first refusal (source: "Systems and Methods for Trading", patent application publication number US 2004/0210512 A1, dated October 21, 2004). That is, small-sized market orders do not trigger workups and thus the volume expansion on only the aggressive side is not possible under eSpeed's workup protocol.

accounting for 20% or more of workups in the notes. Taken together, the evidence demonstrates how the workup protocol on BrokerTec's electronic trading platform differs markedly from the earlier workup protocol described in Boni and Leach (2004).

In order to see if the workup mix is sensitive to different times of day and market conditions, we also analyze the direction of workup volume expansion over different trading hours, on volatile days versus tranquil days, and on days with extreme net order inflow versus net order outflow. In general, the patterns are similar and thus, for brevity, not reported. One notable finding from our sensitivity analysis is that traders in the notes tend to expand limit orders more often on extremely volatile days. Together with the evidence documented earlier that workups are used more frequently on volatile days, this finding is consistent with Boni and Leach (2004)'s conclusion that limit order expandability is helpful to limit order traders as it helps them reduce pick-off risk and information leakage associated with posting large limit orders during volatile times.

3.3 Informational Value of Workup Trades

We proceed to specify a microstructure model for the dynamics of price and order flow, built upon the general framework described in Hasbrouck (2007). The notable feature of our model is that it accounts explicitly for the segmentation of order flow due to the workup feature, as theory suggests that the transparent and dark components of order flow likely have different information values. From this model, we derive a structural VAR representation in (irregular) trade time to be estimated using the data. We then discuss the empirical implementation and findings of the model.

3.3.1 A Microstructure Model of Price and Trade

Let t index the t^{th} market order. We distinguish events occurring during the pre-workup and workup phases of the t^{th} transaction by the subscripts t^- and t^+ respectively. P_{t^-} denotes the best bid ask midpoint (logged) observed as of the t^{th} transaction, and m_{t^-} the unobservable efficient price. Let LT_{t^-} be the signed volume of pre-workup (or "lit") trading, and (DT_{t^+}) the signed volume of workup (or "dark") trading. Both volume variables are positive if the t^{th} transaction is a buy, and negative if it is a sell.

The basic building blocks of the model are:

$$
\begin{aligned}
m_{t^-} &= m_{t^- - 1} + w_{t^-} && (3.1)\\
P_{t^-} &= m_{t^-} + cLT_{t^-} && (3.2)\\
LT_{t^-} &= v_{1,t^-} + \beta_1 v_{1,t^- - 1} && (3.3)\\
w_{t^-} &= u_t^- + \lambda_1 v_{1,t^-} + \lambda_2 v_{2,t^+ - 1} && (3.4)\\
DT_{t^+} &= v_{2,t^+} + \beta_2 v_{2,t^+ - 1} + \alpha_1 v_{1,t^-} + \alpha_2 u_t^- && (3.5)
\end{aligned}
$$

where the efficient price m_{t^-} is specified to follow a random walk as in equation (3.1). w_{t^-} is the efficient price increment and the subscript t^- indicates that the price updating takes place with the execution of the t^{th} market order, but before the workup begins. The observed price P_{t^-}, as expressed in equation 3.2, consists of the permanent component m_{t^-} as well as a component reflecting trading frictions (cLT_{t^-}). Since workup trades are conducted at the price determined in the pre-workup trading round, P_{t^-} depends contemporaneously on the lit trade flow LT_{t^-} but not on the dark trade flow DT_{t^+}.

To allow for the positive auto-correlation of transaction sign as predicted by theory (for example, Parlour (1998)) and observed in the data, a MA(1) model is specified for the lit trade flow as in equation (3.3), where v_{1,t^-} is a white noise process and captures the pre-workup trade innovation. Likewise, equation (3.5) for the dark trade flow DT_{t^+} has an MA(1) error structure with the error term v_{2,t^+}. However, it also includes the contemporaneous effect of the innovation in lit trade flow v_{1,t^-} that precedes and initiates the workup process, as well as public information that arrives at the time of the trade u_{t^-}.

Equation (3.4) models the efficient price increment w_{t^-} as consisting of both a public information component u_{t^-} that is unrelated to trade and a trade-related private information component. The latter component consists of non-public information inferred from the lit trade flow, as well as the lagged dark trade flow. While workup trades have no immediate price implication as they are executed at an established price, traders can observe the workup trade flows after each transaction and update their belief about the fundamental security value in subsequent transactions. Therefore the dark trading innovation enters the efficient price increment equation with a lag. Finally, the model's innovation terms, namely u_{t^-}, v_{1,t^-} and v_{2,t^+} are assumed to be uncorrelated.

From this point, we simplify the notation by suppressing the plus and minus superscripts of t. With this setup, we can derive a VMA(2) for $Y_t \equiv \begin{bmatrix} LT_t & \triangle P_t & DT_t \end{bmatrix}^T$ with an error vector ϵ_t, where ϵ_t relates to the model's exogenous variables through the following expression: $\epsilon_t = B \begin{bmatrix} v_{1,t} & u_t & v_{2,t} \end{bmatrix}^T$, with:

$$B = \begin{bmatrix} 1 & 0 & 0 \\ \lambda_1 + c & 1 & 0 \\ \alpha_1 & \alpha_2 & 1 \end{bmatrix}. \tag{3.6}$$

Assuming invertibility condition, a VAR representation (of infinite order) exists for Y_t with the error vector ϵ_t and a covariance matrix $\Omega \equiv \text{Var}(\epsilon_t)$. The matrix B accordingly captures the contemporaneous dynamic structure of the model. It is a lower triangular matrix, reflecting our key assumption with respect to the causal ordering in the model. Specifically, the ordering goes from pre-workup trades to price update and finally to workup trades (which also corresponds to the way we intentionally stack up the vector Y_t). Price revision following the pre-workup trade variable reflects the commonly adopted assumption in the literature that traders watch order flow to update their beliefs about the fundamental value of a security. That the pre-workup trade variable and price revision precede the workup trade variable in the ordering is natural given the way the workup process works: a market order (i.e., pre-workup, or originating, trade) must arrive and execute against standing limit orders before the workup process opens at the established price point.

Formulated this way, the model implies that the price revision incorporates two sources of information: 1) public information unrelated to trades (u_t), and 2) private information inferred from the contemporaneous trade flow innovation ($v_{1,t}$) and the previous workup trade flow innovation ($v_{2,t-1}$). The role of private and public information in the process of price formation in the U.S. Treasury bond market has been well studied in the literature (e.g., Pasquariello and Vega (2007)). Our model goes one step further by delineating the sources of private information and quantifying the informational importance of workup trade activities separately from the information content of initiating a market order. For comparison, we also estimate a standard model of the price impact of trades with only the generic transaction volume, which we refer to as the "bivariate" model (as opposed to our "trivariate" model).

3.3.2 Permanent Price Impact of Trades

The VAR representation discussed in the previous section can be fit to the data and the permanent price impact of the respective components of order flow can be evaluated. For empirical implementation, we estimate a structural VAR(5) model. Given the assumed ordering discussed earlier, the structural dynamics (i.e., the matrix B as well as the structural variance $\sigma_u^2, \sigma_{v_1}^2$, and $\sigma_{v_2}^2$) can be fully identified.

The long-run cumulative response of price provides a measure of the permanent price impact which is attributable to information and not transitory liquidity effects (see Hasbrouck (1991b)). In other words, it corresponds to the increment in the efficient price w_t:

$$\mathbb{E}\left[\triangle P_t + \triangle P_{t+1} + ...|\epsilon_t\right] = \Psi_{\infty,P}\epsilon_t \tag{3.7}$$

We approximate $\Psi_{\infty,P}$ by computing the cumulative impulse response function (IRF) out to a sufficiently long horizon over which the price response has stabilized and any transitory effects have washed out. As is standard in the literature, we compute the IRF for price from the estimated VAR model by forecasting the system recursively forward to the chosen horizon, assuming that the system is initially at rest, i.e., all variables are set to 0 except for the shocked variable. Inspection of the path of the estimated IRFs indicates that a horizon of 50 transactions provides a reasonable approximation of the permanent price impact $\Psi_{\infty,P}$. The price impact is measured with units in hundredths of a percent of par value (basis points), which is equivalent to cents per \$100 par value. The model is estimated separately for each day in our sample.

Figure 3.4 plots the average cumulative change in price up to 25 transactions following an initial \$1 billion shock in trade volume initiated from the buy side. The figure shows clearly that the transparent part of order flow has a much greater price impact than the dark part for most securities. The only exception is in the 2-year note, where the impact of the pre-workup and workup trade flows is comparable. The figure also shows that the cumulative price response largely settles by the fifth transaction.

The estimated permanent price impact per \$1 billion shock is reported in Table 3.3. The table shows the mean and the 95% range of the time series of the daily price impact estimates separately for pre-workup trade flow (under the "Lit Trades" column) and workup trade flow (under the "Dark Trades" column). The mean impact is monotonically increasing in maturity, and this pattern applies to both the pre-workup and workup trade flow. At the shorter end, the 2-year note price increases by merely 3.7 bps if the trading volume during the pre-workup phase increases by \$1 billion. In sharp contrast, the same shock, if it occurs in the

30-year bond, induces a permanent increase of nearly 400 bps, about a hundred times larger. With respect to workup trade flow, the differential in the price impact between the two maturity ends is not as extreme: 51 bps for the 30-year versus 3 bps for the 2-year. In between, the 5- and 10-year notes exhibit a more moderate difference: the price impact of the latter is slightly more than twice that of the former. The ranges of price impact estimates for lit and dark trade flow also generally respect the ordering by maturity just discussed, except for the 30-year bond where a much wider range is observed.

Finally, it is useful to look at the variation over time of the price impact to gain an understanding of how market liquidity has evolved. Figure 3.5 plots the 20-day moving average of price impact over the sample period. Considering first the price impact of initiating a market order, one can see a significant increase during the crisis period (from August 2007 to June 2009), with the sharpest increases (to about four to eight times larger than the pre-crisis level) occurring in late 2008. This is consistent with the patterns of market depth and bid-ask spreads, other measures of market liquidity, shown in Engle et al. (2012a), in suggesting that the market was markedly less liquid during the crisis.

The price impact of workup trades also varies significantly over time, albeit less so than the price impact of pre-workup trades, with a mild increase during the crisis period. There is thus roughly a doubling of price impact from the pre-crisis level to the peak of the crisis for the 5-, 10-, and 30-year securities, versus a four- to six-fold increase for the pre-workup trades. For the 2-year note, there is roughly a quadrupling of price impact at the peak of the crisis for the workup trades, versus a roughly eight-fold increase for the pre-workup trades. The differential response during the crisis means that the 2-year note's price impact estimates for pre-workup and workup trades, which are similar for the pre-crisis period, separate out during the crisis. Taken together, the evidence indicates that initiating a trade produces a greater impact than waiting to trade the same quantity during a workup, and that this gap is more pronounced during times of crisis.

3.3.3 Information Content of Workup Trades

To evaluate the informational value of workup trades, we follow the information share framework as introduced in Hasbrouck (1991b) and applied widely in subsequent studies of price discovery. Conceptually, the information share of a variable measures the extent to which its variation contributes to the variance of the efficient price update w_t. From equation (3.7), this variance can be approximated by:

$$\tilde{\sigma}_w^2 = \Psi_{h,P} \Omega \Psi_{h,P}^T \tag{3.8}$$

Given the structure of the system, it is easy to show that the right-hand side of equation (3.8) is a linear combination of $\sigma_u^2, \sigma_{v_1}^2$, and $\sigma_{v_2}^2$. Each of these terms can then be expressed as a percentage of σ_w^2 and is referred to as the "Hasbrouck information share" of the relevant variable. Specifically, the percentage attributable to σ_u^2 indicates the extent to which public information drives the variation in the efficient price update, whereas those attributable to $\sigma_{v_1}^2$ and $\sigma_{v_2}^2$ quantify the contribution of non-public information revealed through the trade flows during the pre-workup and workup phases respectively. The information share statistics thus allow us to disentangle the information structure and determine the degree of private information being conveyed in workups in comparison to that conveyed through the normal/visible trade flows.

The information share estimates are reported in Table 3.4. We observe that the informativeness of the lit trade flow is quite consistent across all four securities, ranging on average between 15% and 19%. At the 95% upper bound, this part of order flow explains about 30% or more of the total variation in the efficient price innovations for each of the 2-, 10-, and 30-year securities, and about 26% for the 5-year note.

In contrast, there is a much wider range for the informational value of workup trades across maturities. On the one hand, the dark trade flow of the 2-year note drives about 17% of the variation in the efficient price – slightly higher than the contribution of the lit trade flow. On the other hand, there is almost no private information revealed by the workup trade flow for the 30-year bond (1%). Even the 95% upper bound for the bond is only about 5%. In between, the 5- and 10-year notes are quite similar in terms of workup trade informativeness, with average contributions of 7 and 8% respectively, and corresponding 95% upper bounds of 18% and 21%.

Despite the importance of trade flow, the table also shows that public information is nonetheless the main driver of the variation in the efficient price. For the 5- and 10-year notes, the average contribution of public information to the price discovery process is between 73-77%, with a 95% range of roughly 60-90%. The 30-year bond has a slightly higher public information share, with a mean of 82% and a 95% range between 69% and 94%. The 2-year note shows a slightly lower public information share, averaging 67% and ranging between 42% and 88%. That is, trade flow is most informative at the short maturity end and least informative at the long maturity end. Moreover, the breakdown between lit and dark trades shows that this overall differential between public information and trade-related information is explained mainly by the differential in the informativeness of workups.

The time series of the information shares, presented in Figure 3.6, show that the informativeness of the lit trade flow appears rather stable over time, with a slight increase toward the end of the sample period. On the other hand, the informational role of the dark trade flow changes more appreciably over time, most notably among the notes. The information share of workup trades trended down through much of 2007 and 2008, before rebounding in 2009. The information share of workup trades for the 2-year note settled at a new higher level after 2009, whereas the share remained similar or lower than pre-crisis levels for the 5- and 10-year notes, respectively.

3.3.4 Information Structure on Special Days

We now analyze the information structure on days of special interest. We specifically look at days with important announcements, days when the market is highly volatile, and days when the market experiences extreme buying pressure – an indicator of a possible flight-to-safety. Table 3.5 documents this analysis. Under each security, there are three columns: Lit Trades, Dark Trades and Public Information. Different from Table 3.4 where we report the raw information shares of lit and dark trades, in this table, the respective shares are standardized by the total trade-related information share, i.e., these two columns add up to 100%. This makes it easier to see the relative informational importance of lit versus dark trade flow. The private versus public information split can be gauged by examining the public information share reported in the third column for each security.

3.3.4.1 Announcement Days

In Table 3.5, Panel A, we compare non-announcement days to days with announcements of: 1) FOMC rate decisions, 2) important macroeconomic releases, and 3) auction results.[14] These announcements have been shown to be important to Treasury price formation (see Fleming and Remolona (1997), Balduzzi et al. (2001), Green (2004), Pasquariello and Vega (2007) and references therein). For each of these announcement types, we compare the relative informativeness of the lit and dark trade flow on announcement days with that estimated on days when none of these three announcement types occurs.

Interestingly, there is no major change in the private information structure on announcement days, as compared to non-announcement days. That is, the mix of information content of lit and dark trade flow

[14]See Appendix A for the list of announcements considered.

remains quite similar across different announcement types (including no announcement). However, the trade flow collectively has relatively less information value on FOMC and macroeconomic announcement days. This result is intuitive, because there is a greater amount of public information arriving on these announcement days which can move prices without requiring trades, as shown in Fleming and Remolona (1999) and other studies.

3.3.4.2 Volatile Days

Table 3.5, Panel B shows a comparison of the information structure on highly volatile days against days with low volatility (based on the 95th and 5th percentiles of the volatility distribution). We first focus on the private information mix. Consistently across all four securities, the pre-workup trade flow – the lit part – is relatively more informative on high volatility days. It is helpful to tie this result to an earlier stylized fact that workups are used more on volatile days, and, particularly for the notes, more often expand the quoted depth. We interpret these results collectively as indicating that: 1) information is short-lived in volatile times, necessitating fast execution and 2) the increased incidence of quoted depth expansion reflects how liquidity providers (not necessarily informed) use the workup option to guard against adverse execution of their orders.

It is also interesting to see that public information takes on a greater role in price discovery when the market is highly volatile, as compared to when the market is tranquil. That trades are relatively less informational when price is highly volatile is to be expected, because the variance of the efficient price update is a linear combination of the variances of return and the two order flow variables. When price fluctuates greatly, this variability dominates the variance of the efficient price update, leaving a lesser role for trade-related information in the price formation process. An intuitive way to think about this result is that noisier public information makes it harder for market participants to interpret trade flow patterns and discern value-relevant information.

3.3.4.3 Days with Extreme Net Order Flow

In Table 3.5, Panel C, we compare the information structure on days with high net inflows and high net outflows (based on the 95th and 5th percentiles of the distribution of net order flow). Net order flow, if positive, suggests a possible flight-to-safety into Treasury securities (see Beber et al. (2009)), whereas strongly negative net order flow suggests a flight out of Treasury securities. The results show that the 2- and 10-year notes do not exhibit a statistically significant change in the information structure between flights into and out of Treasuries. In contrast, the lit order flow of the 5- and 30-year securities becomes relatively more informative on days with high flows into the market, compared to flows out of the market. However, the shift is fairly small in magnitude. Furthermore, most securities show a similar public information share between high inflow and high outflow days. Overall, the nature of the flows in the market does not seem to alter substantially the information structure and workup characteristics.

3.3.5 Comparison with Standard Model of Price Impact of Trades

Our analysis in the previous section illustrates that delineating the trade flow into the pre-workup and workup components permits a more complete understanding of how the different layers of the trading process convey non-public information and affect price dynamics. One of the key findings is that trade flow is not homogeneous. A $1 million trade initiated in the pre-workup stage generally results in a greater price impact and carries more information than when the same trade occurs in the workup stage.

As a result, if we model only the trade volume variable without considering its respective components, we may underestimate the price impact of a market order, since the lower impact of the workup component pulls down the estimate for the whole transaction size. In addition, omitting the possible endogenous interaction of workup and pre-workup activity might underestimate the overall informativeness of order flow. To formally see this, we estimate a bivariate VAR(5) of trade flow and return, and compute the permanent price impact as well as the information share of transaction volume. The comparison to the trivariate results is provided in Table 3.6.

Panel A illustrates that for the 5-, 10- and 30-year securities the price impact of a lumped-together (or "generic") trade estimated from the bivariate model is much smaller than the price impact of a market order of the same size estimated from the trivariate model (about half the magnitude). At the same time, Panel B shows that the price impact of a generic trade is higher than the price impact of a trade occurring during a workup.

For the 2-year note, the estimated price impact of a generic trade is not only lower than the estimated price impact of a market order (Panel A), but also lower than the price impact of a workup trade (Panel B). As discussed earlier, workup activity in the 2-year note is generally as informative as pre-workup trading activity. Failure of the bivariate model to capture the endogenous dynamics between workups and trade initiation featured in our trivariate model results in a lower price impact estimate than that of workups for the 2-year note.

In addition, as shown in Panel C, the bivariate model attributes less information value to order flow. Our tests of the hypothesis that the information share of trades in the model of segmented order flow is not higher than that implied by the bivariate model are rejected for three of the four securities considered. This is because the bivariate model does not capture and attribute adequately the different contributions and variations in the respective components of the overall order flow. More importantly, as discussed above, the dynamic interaction between pre-workup and workup order flow is absent in the bivariate model, implying a lower information role of order flow than is the case when this dynamic interaction is taken into account.

Economically, it is important to recognize that the workup option is an integral part of the trading process in the Treasury interdealer market. It is undoubtedly factored into the trading decisions of dealers, since they can choose to trade immediately by submitting a market order, or wait to trade in a workup. Factors such as liquidity need, degree of impatience and/or possession of short- versus long-lived information might contribute to the segmentation of order flow, as dealers balance faster execution with higher price impact. Treating this market as one where such a workup option is not available and trade flow is homogeneous may give rise to a less than accurate characterization of the trading process and how trading affects price dynamics.

3.3.6 Is Direction of Workup Expansion Informationally Relevant?

As discussed earlier, the workup protocol on the BrokerTec electronic platform differs from the voice-assisted protocol described in Boni and Leach (2004) in that workup volume can originate from either side, as opposed to just expanding the limit order book. The analysis performed up to this point has considered all workup volume to be equal, but additional insight may be gained by examining whether the direction of volume expansion during a workup matters to price discovery. Our results show that this is indeed informationally relevant.

To proceed, we estimate an expanded VAR model in which the workup trade flow (DT) is replaced by three workup trade flow types (DT_1, DT_2, and DT_3). These are workups that expand volume on: 1) the aggressive side (DT_1), 2) both sides (DT_2), and 3) the passive side (DT_3).[15] The vector of endogenous variables is now $Y_t \equiv \begin{bmatrix} LT_t & \triangle P_t & DT_{1,t} & DT_{2,t} & DT_{3,t} \end{bmatrix}^T$. To check whether the relative importance of each of these workup trade flow types is sensitive to the VAR variable ordering, we also report the results based on an alternative ordering in which the different workup trade flow categories are reversed,

[15]See subsection 3.2.4.4 for a detailed description of this classification.

i.e., $Y_t \equiv \begin{bmatrix} LT_t & \triangle P_t & DT_{3,t} & DT_{2,t} & DT_{1,t} \end{bmatrix}^T$. From the estimated VAR, we compute the information share for each of the variables in the system as described earlier, and report them in Table 3.7. For brevity, the information shares of lit trades and public information are not shown as they are quite similar to the previously reported results.

The results indicate that workup volume coming from the aggressive side contributes significantly more to price discovery. For example, the information share of workups that expand the aggressive side averages 13% for the 2-year note, compared with the 6% information share of workups that expand both sides and the 1% information share of workups that expand the passive side. Interestingly, this pattern holds even for the 30-year bond, where workups mainly expand the passive side. A comparison of the two orderings shows that the results are not ordering sensitive.

To see whether the information contribution of each type of workup order flow is commensurate with its share of volume (shown in Figure 3.3), we rescale the three information shares so that they add up to 100%. The rescaled numbers indicate the relative contribution of each workup type to the total informativeness of workup order flow. We use the mean information shares based on the first variable ordering for this calculation, but the result is similar for the other ordering. The relative information contribution of aggressive workups is 63.9%, 64.1%, 64.1%, and 64.2% respectively for the 2-, 5-, 10- and 30-year securities. These percentages are consistently higher than the volume share of aggressive workups, which are 53%, 38%, 29% and 22%. That is, aggressive workups are disproportionately more informative than the other two workup types.

3.4 Determinants of Workup Trades

As the previous section shows, trading activity that takes place during the workup stage has a non-trivial role in the price discovery process. Additionally, workups take place in more than half of transactions and account for a large share of volume transacted in this market. Collectively, these findings provide a motivation for our subsequent analysis exploring the determinants of the workup option and the extent of volume transacted during this phase. Being able to predict the likelihood and extent of a workup upon the arrival of the next market order, based on prevailing market conditions, can be valuable to market participants in making trading decisions.

In order to identify workup determinants, it is important to understand workup benefits and costs. The most natural cost of waiting to transact in a workup is the risk of non-execution and perhaps the loss of private information advantage, since counter trading interest may not exist in a workup. Therefore, variables that correlate with non-execution risk or the perishability of private information are expected to be negatively associated with workup usage and workup volume.

On the other hand, the obvious benefit of the workup protocol is that traders have more flexibility with what to do with their trading intention, including not doing anything at all if the market moves unfavorably. This provides an important advantage over iceberg orders, since the hidden part of an iceberg order may get executed adversely before the trader has a chance to modify or cancel. Furthermore, the ability to expand volume during workups can be valuable to those traders with a large trading interest. By submitting an initial small sized order, those traders can avoid causing adverse price impact that could have resulted had they submitted the full-sized large order altogether.

The use of workups therefore reflects a trade-off among non-execution risk, increased control over one's trading activities, and the ability to avoid adverse price impact. We thus model the probability of workup (i.e., whether or not a transaction has a workup), as well as the magnitude of the workup volume, with the following explanatory variables capturing this trade-off:

- *DepthSame*: prevailing inside depth on the same side of the transaction (logged).
- *DepthOpp*: prevailing inside depth on the opposite side of the transaction (logged).
- *PretradeSpr*: prevailing relative spread in basis points $\left(10,000\frac{P_A - P_B}{(P_A+P_B)/2}\right)$.
- *MoSize*: pre-workup volume of the transaction (i.e., the volume transacted before the workup start) (logged).
- *HdRevealed*: whether trading activities during the pre-workup stage have revealed any iceberg orders.
- *AveDurLast5*: average transaction duration (in seconds) in the last five minute interval (logged).
- *Vola5Min*: volatility as measured by the high low range of the logged mid-quote over the last five minute interval, capturing the level of volatility immediately before the transaction.
- *PctWkup5Min*: percentage of transactions with a workup in the last five minute interval, to control for the possible liquidity externality of workup activities as predicted by Buti et al. (2011a)'s model.

- *PctWkupV5Min*: percentage of volume expanded during workups (conditional on workup usage) in the last five minute interval. This is another control for the liquidity externality.

- *Tokyo trading hour dummy*: equals 1 if the transaction starts during the period from 18:30 EST (or 19:30 EDT) the previous day to 3:00 ET.

- *London trading hour dummy*: equals 1 if the transaction starts during the period from 3:00 ET to 7:30 ET.

We employ a logistic regression model for the probability of workup, in which the dependent variable equals 1 for those transactions with workup, and 0 otherwise. For the extent of volume expansion during a workup, we estimate a Tobit model in which the dependent variable is the workup volume, and those transactions with no workup are censored at zero. The model estimates are presented in Table 3.8. Given the large number of observations, most of the coefficient estimates are significant at the 5% level. Only those coefficients that are not significant are marked with an asterisk. We discuss each determinant below.

First, the prevailing depth on the same side is positively related with both the probability of workup and the magnitude of workup volume. This supports the argument that a higher level of depth, indicative of longer time to execution for the marginal limit order, might encourage traders with trading interest on the same side to opt for the immediate execution opportunity offered by the workup. This finding is enhanced by the negative relationship between prevailing spread and the likelihood of workup.[16] A tighter spread (the spread is often 1 tick in this market) makes it harder to post limit orders inside the spread, while simultaneously reducing the cost to trade at the workup price (i.e., the forgone spread). Thus, the choice of immediate execution becomes more attractive, despite it being at a worse price than that of a limit order price. Both of these findings provide empirical support for Buti et al. (2011a)'s model of dark pool trading strategies in limit order markets.

Our finding concerning the effect of depth on the opposite side provides some insight into what matters more to traders when the market is shallow on the opposite side. Theoretically, the effect of opposite side depth on the likelihood and extent of a workup can go either way. On the one hand, the model by Buti et al. (2011a) shows that lower depth on the opposite side to absorb incoming orders can result in more adverse price impact for incoming trades. If so, the workup protocol can be valuable as it gives traders an option to

[16]Note that once the workup choice is made, the effect of prevailing spread on the extent of volume expansion during the workup is mixed: negative for the 2- and 10-year securities, but positive for the 5- and 30-year securities.

start with a smaller sized order and expand the size in a workup without bearing significant market impact. On the other hand, lower depth on the opposite side may be a sign that trading interest on that side is lacking. This can reduce the execution probability of trades in a workup, resulting in a lower likelihood of successful matching during the workup window. Even if a workup does occur, the workup quantity is likely to be lower.

Our empirical evidence of a positive relationship for most securities supports the latter argument; that is, non-execution risk appears to be a more important consideration than the adverse price impact concern. The only exception is the 2-year note, for which we observe a negative effect of opposite side depth on the likelihood of a workup. Recall that the transaction size in the 2-year note is often much larger than that for other securities. Accordingly, the adverse price impact associated with the lack of opposite standing depth might become a more important concern, thereby encouraging greater usage of workups. As for workup quantity, the effect of opposite side depth is also positive for all securities, providing additional support for the non-execution risk hypothesis.

Next, the initial size of a transaction is positively associated with the likelihood of workup (as shown by the positive coefficients for pre-workup volume across securities, except for the 30-year bond). This provides direct empirical support of Harris (1997)'s argument that there might be inactive traders in the market who only take action based on the actions of others. A larger volume transacted during the pre-workup phase is more likely to ignite interest from otherwise inactive traders. Another possible explanation is that large initial volume is perceived by the market to be associated with large liquidity demand. This may induce the expansion of the quoted depth during the workup beyond the level observed just before the trade – an idea that finds empirical support in Boni and Leach (2004). Interestingly, once a workup is taking place, the additional volume transacted may increase or decrease with the pre-workup volume depending on the security. For example, it is positive for the 2- and 30-year securities but negative for the 5- and 10-year notes.

Another aspect of pre-workup trading – the revelation of hidden depth – can also predict higher workup usage and volume expansion. The revelation informs market participants that there is a hidden liquidity pool in addition to the initially observed depth and that workup trades have a greater chance of being filled/absorbed.

We further find that price volatility, measured over the 5-minute time window leading to each transaction, is positively related with workup activities. This is consistent with Boni and Leach (2004)'s finding using GovPX data under a protocol in which workups expand the limit order book only. Intuitively, when the market is volatile, the risk of adverse execution of limit orders increases, thereby motivating a greater reliance on workups because the protocol allows traders greater control over when and how much to trade, or even not to trade at all.

Moreover, since the workup protocol can be likened to a crossing network, we can also borrow theoretical insights from that literature for a better understanding of our empirical volatility finding. Ye (2012) suggests a linkage between security value uncertainty and the choice of trading in the crossing network as opposed to the transparent exchange. Specifically, uncertainty increases both the price impact of trades on the exchange and the non-execution probability in the crossing network, but the net effect is that the crossing network is comparatively more beneficial for the informed traders. As a result, Ye predicts that crossing network usage should increase in value uncertainty. Our finding concerning the effect of volatility is generally in line with this prediction, as well as empirical evidence in Ready (2012) for a cross section of the 500 largest NASDAQ stocks from 2005 to 2007.

Our result also shows that the speed of trading in the market significantly increases the likelihood of a workup, as well as the magnitude of worked-up volume. In light of Easley and O'Hara (1992) and Dufour and Engle (2000), high trading intensity is likely reflective of information arrival, and thus, inactive trading interest can be activated and revealed in a workup following the lead of market order traders. Furthermore, the positive coefficients for the prevailing level of workup activity support Buti et al. (2011a)'s argument that dark pool liquidity begets dark pool liquidity, as a higher level of workup activity signals an increased chance of finding counter-party trading interest and successful execution of workup orders.

Finally, the probability of workup and extent of worked up volume are both significantly lower outside New York trading hours, even after having controlled for the level of trading activity through the previously discussed covariates. This seems to be consistent with the hypothesis that workups are used less in the overnight hours when there are fewer traders in the market. There is simply a lower chance of meeting with a counter-party in a workup, or being able to ignite inactive trading interest, when there are not many traders at their desks.

3.5 Conclusion

This chapter studies the workup protocol, a distinctive and frequently used trading feature in the U.S. Treasury securities market. Given its importance in discovering a large portion of market liquidity, we examine its role in the price formation process, and distinguish it from the information value of non-workup trades that initiate workups. We find that workup trade flow generally contains less information than its transparent counterpart, but that its role is not trivial. In addition, it is the aggressive side workups that account for most of the information value of workup trade flow. Workups that expand the limit order book as described in Boni and Leach (2004) are far less informative.

Furthermore, we find that workups occur more frequently around volatile times, when the incidence of workups expanding the pre-trade limit order book also increases for all three notes, suggesting that the workup protocol is helpful to limit order traders in managing their trading interest. Additionally, workups are more likely when the market is more liquid (e.g., greater market depth and tighter bid-ask spreads) or trading more active. Interestingly, lit order flow becomes more informationally relevant on highly volatile days, supporting the belief that traders with better information are more likely to initiate trades and exploit their information before adverse price movements can render the information less valuable. Taken together, the evidence seems to suggest that workups are used more as a channel for liquidity providers to guard against adverse price movements, than as a channel to hide private information.

Our findings provide important implications for research into the price discovery of U.S. Treasury securities. Consistent with theory, we document that the different layers of order flow have different information content. Intuitively, given the option of trading in a workup, a trader who chooses to initiate a trade (as opposed to wait for a workup) conveys a stronger signal to the market than otherwise would be the case in a hypothetical market setup where such a workup option does not exist. Therefore, the act of initiating a trade should contribute more to information discovery than the act of trading in a workup. We show that, without considering this segmentation, the price impact so estimated can underestimate the impact of initiating a trade and the share of non-public information flow.

Beyond the literature on price discovery in financial markets, our research adds to two important areas of research, namely dark pool trading and security market design. The workup protocol in essence is a dark pool mechanism and provides a valuable opportunity for examining how such a mechanism operates in a fixed income market setting. We show that in the market for U.S. Treasury securities, dark pool trades are only mildly informative and that they tend to occur more often at more volatile times, highlighting the benefit of this mechanism in protecting traders against large price swings. While equity dark pools have recently come into the spotlight for the potential of compromising market quality and fairness, our evidence indicates that this is not a major concern for this fixed income market, one that is populated mostly by sophisticated market participants (i.e., government securities dealers).

With respect to security market design, the workup protocol presents an interesting case study of a continuous limit order market combined with periodic call auctions. This is a timely contribution to the current discussion on the market design response to the trend in high frequency trading. With increasing high frequency trading activity across markets, continuous limit order market design has shown certain limitations (e.g., encouraging an arms race in trading technology). Naturally, these limitations invite further research into alternative market design features and necessitate an understanding of possible implications of such features. In this direction, our work readily offers empirical implications on trading patterns, exposure choice and price discovery in a continuous limit order market enhanced with periodic auctions.

Table 3.1: Summary Statistics of Trading and Workup Activities

	2-Year	5-Year	10-Year	30-Year
PANEL A: DAY-LEVEL STATISTICS				
Volume ($B)	33.5	34.0	29.3	4.7
Pre-workup %	51.8	43.6	45.5	57.3
Workup %	48.2	56.4	54.5	42.7
Number of Transactions	1,224	2,679	2,642	1,464
% with Workup	49.0	56.2	55.2	39.1
% with Iceberg Order Match	4.2	4.3	4.4	3.9
% Executed at Multiple Prices	0.0	0.2	0.2	0.5
PANEL B: TRANSACTION-LEVEL STATISTICS (WITH WORKUP)				
Transaction Size ($M)	41.8	18.6	16.4	5.4
Pre-workup	15.6	6.0	5.5	1.9
Workup	26.2	12.7	10.9	3.5
Number of Trades	9.9	8.7	8.6	3.9
Pre-workup	3.2	2.9	3.0	1.4
Workup	6.7	5.8	5.6	2.5
PANEL C: TRANSACTION-LEVEL STATISTICS (WITHOUT WORKUP)				
Transaction Size ($M)	11.9	4.6	4.2	1.7
Number of Trades	2.4	2.1	2.2	1.3
PANEL D: SAMPLE SIZE				
Number of Transactions	1,836,812	4,017,905	3,946,216	2,197,471
Number of Trading Days	1,501	1,501	1,494	1,501

This table provides summary statistics of trading activity in the on-the-run 2-, 5-, 10-, and 30-year Treasury securities on the BrokerTec platform. The sample period is 2006-2011. A transaction refers to a complete sequence of order executions that starts with the arrival of a market order and ends when the workup initiated by the original market order completes. A trade refers to a single paired order matching. There is no data available for the 10-year note on seven days during the sample period (August 3-7, 10-11, 2009). Numbers reported in Panel A are daily averages. Numbers in Panels B and C are averages across all transactions with and without workups respectively.

Table 3.2: Correlations of Workup and Order Flow Variables

	2-Year	5-Year	10-Year	30-Year
Daily Signed Order Imbalance & Workup Usage	0.048	0.169	0.043	0.043
Daily Absolute Order Imbalance & Workup Usage	-0.297	-0.266	-0.206	-0.081
Daily Volatility & Workup Usage	0.541	0.380	0.540	0.256
Workup Autocorrelation	0.110	0.084	0.089	0.063
Workup Volume Autocorrelation	0.098	0.120	0.136	0.114

This table shows correlations of workup and trading variables for the on-the-run 2-, 5-, 10-, and 30-year Treasury securities on the BrokerTec platform. The sample period is 2006-2011. A transaction refers to a complete sequence of order executions that starts with the arrival of a market order and ends when the workup initiated by the original market order completes. Daily signed order imbalance is buy volume minus sell volume, standardized by the day's total trading volume. Daily absolute order imbalance is the absolute order imbalance standardized by the day's total trading volume. Daily volatility is the average five-minute realized volatility of the bid-ask midpoint (logged) for each day. Workup usage is the percentage of transactions with workups for each day. The workup and workup volume autocorrelation coefficients are computed based on transaction-level data.

Table 3.3: Permanent Price Impact of Segmented Order Flow

		Lit Trades	Dark Trades
	Mean	3.70	3.19
2-Year	95% Range Lower Bound	0.76	0.88
	95% Range Upper Bound	13.59	8.57
	Mean	21.88	7.54
5-Year	95% Range Lower Bound	5.23	2.59
	95% Range Upper Bound	59.42	15.44
	Mean	48.16	18.04
10-Year	95% Range Lower Bound	11.02	6.24
	95% Range Upper Bound	129.36	34.00
	Mean	397.97	50.60
30-Year	95% Range Lower Bound	113.59	-61.24
	95% Range Upper Bound	938.46	203.19

This table reports the permanent price impact (in basis points per $1 billion buyer-initiated volume) of pre-workup trades ("Lit Trades") versus workup trades ("Dark Trades"). The estimates derive from a VAR(5) model of pre-workup trade flow, return and workup trade flow. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. Observations outside the [7:00-17:30] time window are excluded. The model is estimated separately for each day. The mean and 95% range are computed from the time series of daily price impact estimates.

Table 3.4: Share of Trade and Non-Trade Related Information

		Trade Related Information		Public
		Lit Trades	Dark Trades	Information
2-Year	Mean	15.23	17.48	67.28
	95% Range Lower Bound	3.52	3.01	41.72
	95% Range Upper Bound	30.67	38.67	87.73
5-Year	Mean	16.09	6.72	77.19
	95% Range Lower Bound	6.07	0.54	63.19
	95% Range Upper Bound	26.35	17.82	91.08
10-Year	Mean	18.52	8.21	73.28
	95% Range Lower Bound	7.83	0.67	59.11
	95% Range Upper Bound	30.10	21.19	86.53
30-Year	Mean	16.66	1.05	82.29
	95% Range Lower Bound	5.66	0.00	68.69
	95% Range Upper Bound	29.95	5.13	93.73

This table reports the information share (%) of pre-workup trades ("Lit Trades"), workup trades ("Dark Trades"), and non-trade-related information ("Public Information"). The estimates derive from a VAR(5) model of pre-workup trade flow, return and workup trade flow. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. Observations outside the [7:00-17:30] time window are excluded. The model is estimated separately for each day. The mean and 95% range are computed from the time series of daily information share estimates.

Table 3.5: Information Structure on Days with Announcements and Extreme Market Movements

	2-Year			5-Year			10-Year			30-Year		
	Lit Trades	Dark Trades	Public Info	Lit Trades	Dark Trades	Public Info	Lit Trades	Dark Trades	Public Info	Lit Trades	Dark Trades	Public Info
PANEL A: ANNOUNCEMENT VERSUS NON-ANNOUNCEMENT DAYS												
FOMC Announcements	47.66	52.34	72.58*	72.51	27.49	84.19*	71.27	28.73	78.94*	94.38	5.62	86.34*
Macro Announcements	47.82	52.18	67.93*	72.22	27.78	78.30*	70.84	29.16	74.15*	93.60	6.40	83.05*
Auction Days	47.03	52.97	66.89	70.79	29.21	77.76*	69.95	30.05	73.01	95.38	4.62	82.36
Non-Announcement Days	49.30	50.70	66.18	71.56	28.44	75.29	69.72	30.28	71.83	93.84	6.16	81.02
PANEL B: DAYS WITH EXTREME VOLATILITY												
High Volatility Days	62.30*	37.70*	78.94*	83.51*	16.49*	84.28*	86.18*	13.82*	78.10*	95.70*	4.30*	83.66
Low Volatility Days	48.47	51.53	67.48	72.66	27.34	77.40	71.25	28.75	73.56	93.97	6.03	82.45
PANEL C: DAYS WITH EXTREME ORDER IMBALANCES												
High Inflow Days	50.11	49.89	66.15	76.64*	23.36*	76.53	68.72	31.28	73.03	95.69*	4.31*	84.01*
High Outflow Days	48.30	51.70	67.21	72.09	27.91	77.14	70.67	29.33	73.28	93.69	6.31	82.29

This table reports the relative informativeness of pre-workup order flow ("Lit Trades") versus workup order flow ("Dark Trades") and the share of public information ("Public Info") on: A) days with announcements versus non-announcement days; B) high volatility days versus low volatility days; and C) high inflow days versus high outflow days. The thresholds for high and low volatility, and similarly for high inflow and high outflow, are the 95th and 5th percentiles of the distributions for volatility and net order flow (Buy Volume minus Sell Volume) respectively. The relative informativeness of each component of the order flow is measured by its information share as a percentage of the total trade-related information share. The public information share is 100% minus the total trade-related information share. The information shares are computed from a VAR(5) model of pre-workup trade flow, return and workup trade flow. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. Observations outside the [7:00-17:30] time window are excluded. An asterisk (*) indicates significantly different informativeness at the 5% level.

Table 3.6: Informational Content of Segmented versus Generic Order Flow

	2-Year	5-Year	10-Year	30-Year
PANEL A: PRICE IMPACT OF $1 BILLION PRE-WORKUP VOLUME				
Model with Segmented Trade Flow	3.70	21.88	48.16	397.97
Model with Generic Trade Flow	2.74	11.34	25.94	181.00
p-value of paired sample t-test (right tail)	<0.001	<0.001	<0.001	<0.001
PANEL B: PRICE IMPACT OF $1 BILLION WORKUP VOLUME				
Model with Segmented Trade Flow	3.19	7.54	18.04	50.60
Model with Generic Trade Flow	2.74	11.34	25.94	181.00
p-value of paired sample t-test (left tail)	1.000	<0.001	<0.001	<0.001
PANEL C: INFORMATION SHARE OF TRADES				
Model with Segmented Trade Flow**	32.72%	22.81%	26.72%	17.71%
Model with Generic Trade Flow	26.64%	22.76%	26.39%	14.45%
p-value of paired sample t-test (right tail)	<0.001	0.201	<0.001	<0.001

This table compares the price impact and informational content of order flow estimated by our trivariate VAR model, which considers separately the pre-workup and workup order flow, with those estimated by a standard bivariate VAR model, which considers the generic order flow without segmentation. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. ** This is the combined information share of pre-workup and workup trades.

Table 3.7: Information Share of Workup Trades by How Workup Volume Arises

		Ordering 1			Ordering 2		
		Mean	95% LB	95% UB	Mean	95% LB	95% UB
2-Year	Aggressive	12.96	2.18	31.18	13.47	2.25	31.97
	Both	6.20	0.52	16.31	5.88	0.39	15.84
	Passive	1.13	0.00	5.96	0.95	0.00	4.98
5-Year	Aggressive	6.55	1.29	14.41	6.81	1.35	14.87
	Both	3.09	0.16	8.72	2.99	0.13	8.68
	Passive	0.58	0.00	3.05	0.43	0.00	2.61
10-Year	Aggressive	7.72	1.59	16.30	8.05	1.66	16.92
	Both	3.73	0.23	10.47	3.56	0.21	10.18
	Passive	0.59	0.00	3.10	0.42	0.00	2.40
30-Year	Aggressive	1.97	0.01	7.10	2.00	0.01	7.13
	Both	0.67	0.00	3.19	0.64	0.00	3.13
	Passive	0.43	0.00	2.42	0.42	0.00	2.42

This table reports the % information share of workups classified by how workups expand volume: 1) the aggressive side only, 2) both sides, and 3) the passive side. The estimates derive from a VAR(5) model of pre-workup trade flow, return and three categories of workup trade flow. For brevity, information shares of pre-workup trades and public information flow are not reported. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. Observations outside the [7:00-17:30] time window are excluded. Ordering 1 columns show the information share based on a variable ordering of Aggressive, Both, and Passive. Ordering 2 columns show the information share based on a variable ordering of Passive, Both, and Aggressive.

Table 3.8: Determinants of Workups

	2-Year		5-Year		10-Year		30-Year	
	Probability	Quantity	Probability	Quantity	Probability	Quantity	Probability	Quantity
Intercept	-0.037	0.650	-0.031	0.556	0.104	0.556	-0.413	0.670
Pretrade Depth – Same Side	0.131	0.175	0.145	0.164	0.147	0.167	0.049	0.065
Pretrade Depth – Opposite Side	-0.101	0.277	0.081	0.357	0.046	0.365	0.522	0.219
Pretrade Spread	-0.217	-0.067	-0.240	0.008	-0.172	-0.008	-0.009	0.010
Pre-workup Volume (logged)	0.095	0.079	0.081	-0.009	0.087	-0.035	-0.193	0.019
Hidden Depth Revealed	0.156	0.303	0.419	0.414	0.338	0.374	0.992	0.323
Volatility Last5Mins	0.084	0.066	0.042	0.024	0.030	0.013	0.007	0.002
Ave Duration Last5Mins (logged)	-0.174	-0.082	-0.264	-0.094	-0.265	-0.109	-0.207	-0.067
Workup Probability Last5Mins	0.496	0.052	0.641	0.161	0.586	0.162	0.406	0.163
Workup Volume Share Last5Mins	0.176	0.023	0.066	0.016	0.115	0.007 *	0.115	-0.013
Tokyo Trading Hour Dummy	-0.369	-0.131	-0.185	-0.225	-0.308	-0.193	-0.502	-0.119
London Trading Hour Dummy	-0.278	-0.205	-0.134	-0.221	-0.168	-0.237	-0.191	-0.103
Max-rescaled R-squared	0.063		0.076		0.083		0.079	

This table reports results of a logistic regression for whether or not a transaction has a workup (under columns titled “Probability”), and a Tobit model for the dollar volume transacted during the workup phase (under columns titled “Quantity”). The models are estimated using BrokerTec data for on-the-run 2-, 5-, 10-, and 30-year Treasury securities over the period 2006-2011. Tokyo trading hour dummy is equal to 1 for the period from 18:30 EST (or 19:30 EDT) the previous day to 3:00 ET. London trading hour dummy is equal to 1 for the period from 3:00 ET to 7:30 ET. Note: an asterisk (*) indicates *in*significance at the 5% level.

Figure 3.1: Workup Activity over Time

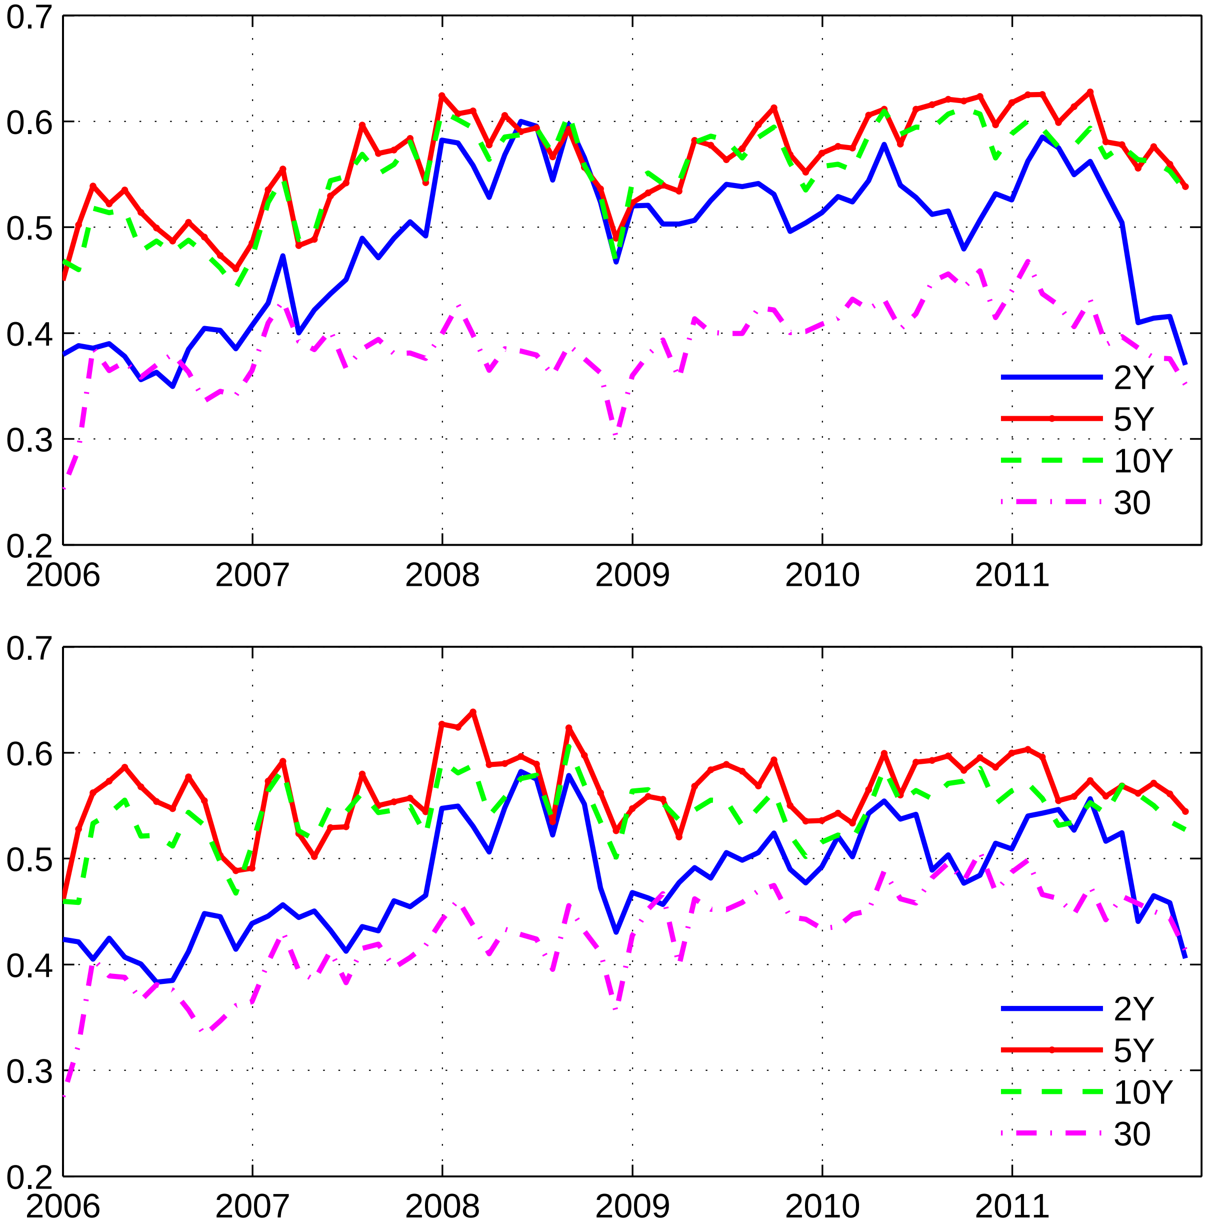

This figure shows the monthly share of transactions with workups (upper plot) and monthly share of volume transacted in workups (lower plot). The numbers are first calculated daily for the on-the-run 2-, 5-, 10- and 30-year Treasury securities on the BrokerTec platform and then averaged across days by month. The sample period is 2006-2011.

Figure 3.2: Intraday Pattern of Workup Probability

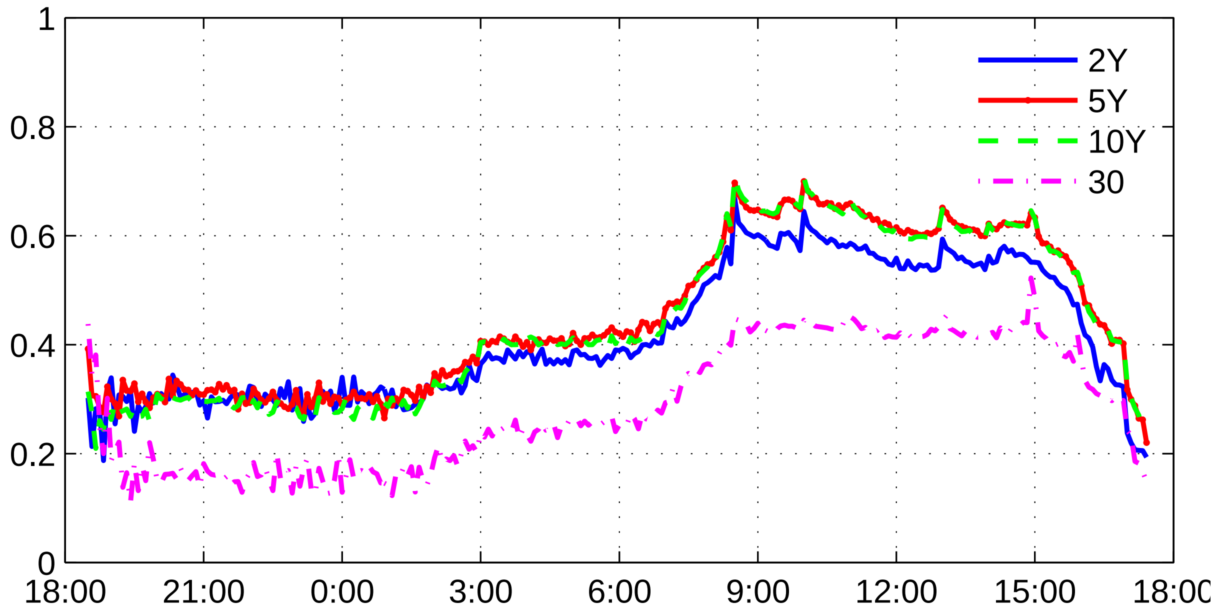

This figure shows the pattern of workup usage over the global trading day (Eastern Time). The plot starts at 18:30 of the previous day and ends at 17:30 of the current day. The numbers are first calculated for a given interval and day for the on-the-run 2-, 5-, 10- and 30-year Treasury securities on the BrokerTec platform and then averaged across days. The sample period is 2006-2011.

Figure 3.3: Which Side Do Workups Expand?

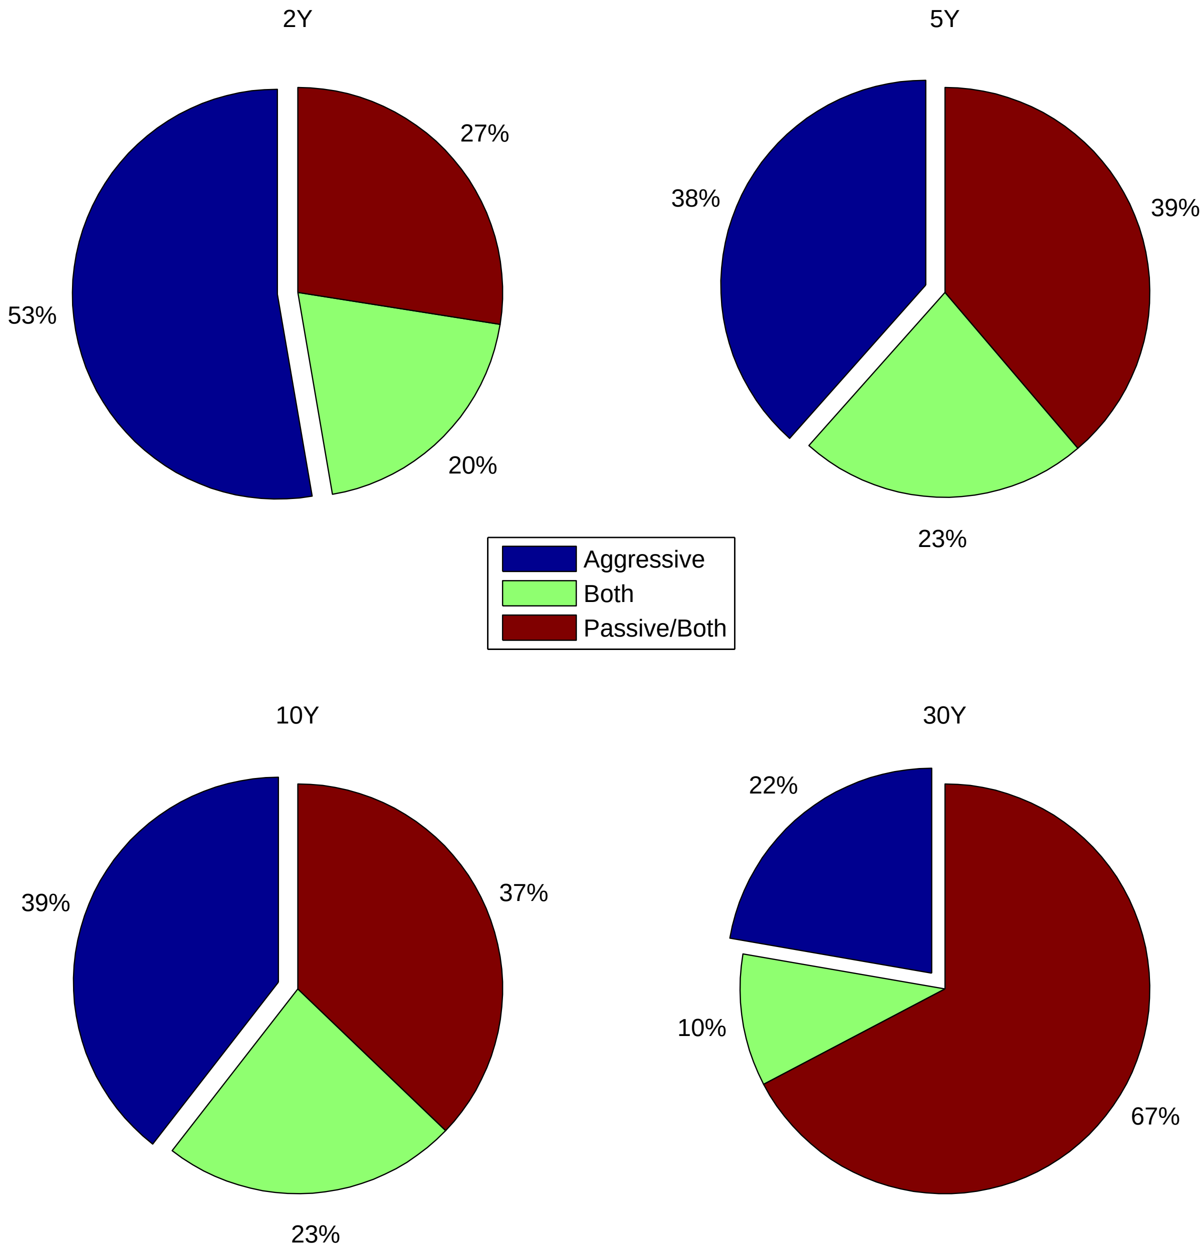

This figure shows the percentages of workups that expand volume on: 1) the aggressive side only, 2) both sides, and 3) the passive side. A workup expands only the aggressive side if the total transaction volume (pre-workup and workup volume combined) is not greater than the depth posted in the limit order book immediately before the transaction. A workup expands both sides if the pre-workup volume is less than the posted depth, but the total transaction volume exceeds the posted depth. A workup expands the passive side if the pre-workup trades exhaust the posted depth. This expansion of the passive side includes instances where the aggressive side is also expanded during the workup. The percentages are first calculated daily for the on-the-run 2-, 5-, 10-, and 30-year Treasury securities on the BrokerTec platform and then averaged across days. The sample period is 2006-2011.

Figure 3.4: Cumulative Impulse Response of Price to Trade

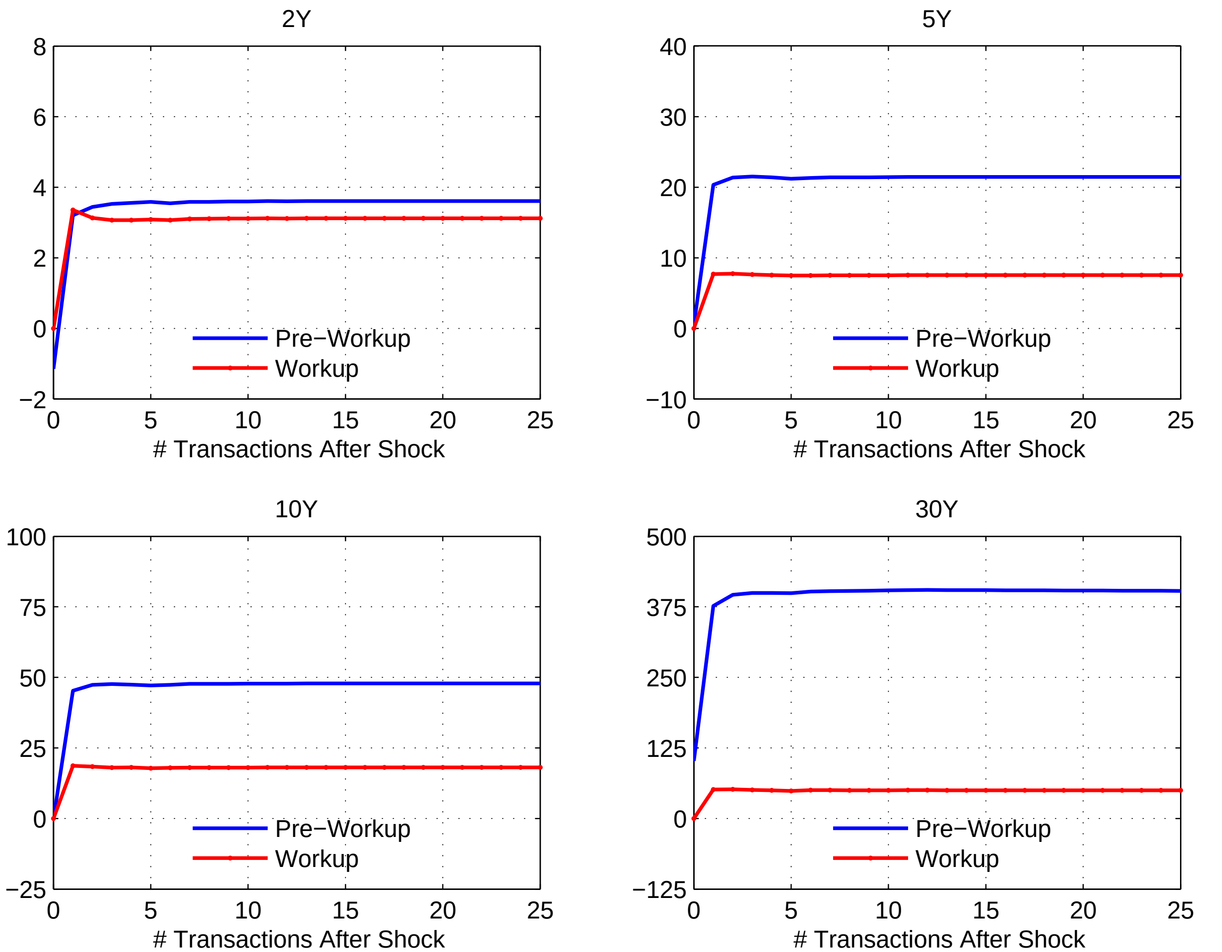

This figure plots the cumulative midpoint return (in basis points) in response to a $1 billion shock to pre-workup and workup trading volume respectively, based on a VAR(5) model of return and segmented order flow. Estimation is first done daily based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities and then averaged across days. The sample period is 2006-2011. Observations outside the [7:00-17:30] time window are excluded from model estimation.

Figure 3.5: Permanent Price Impact of Trade

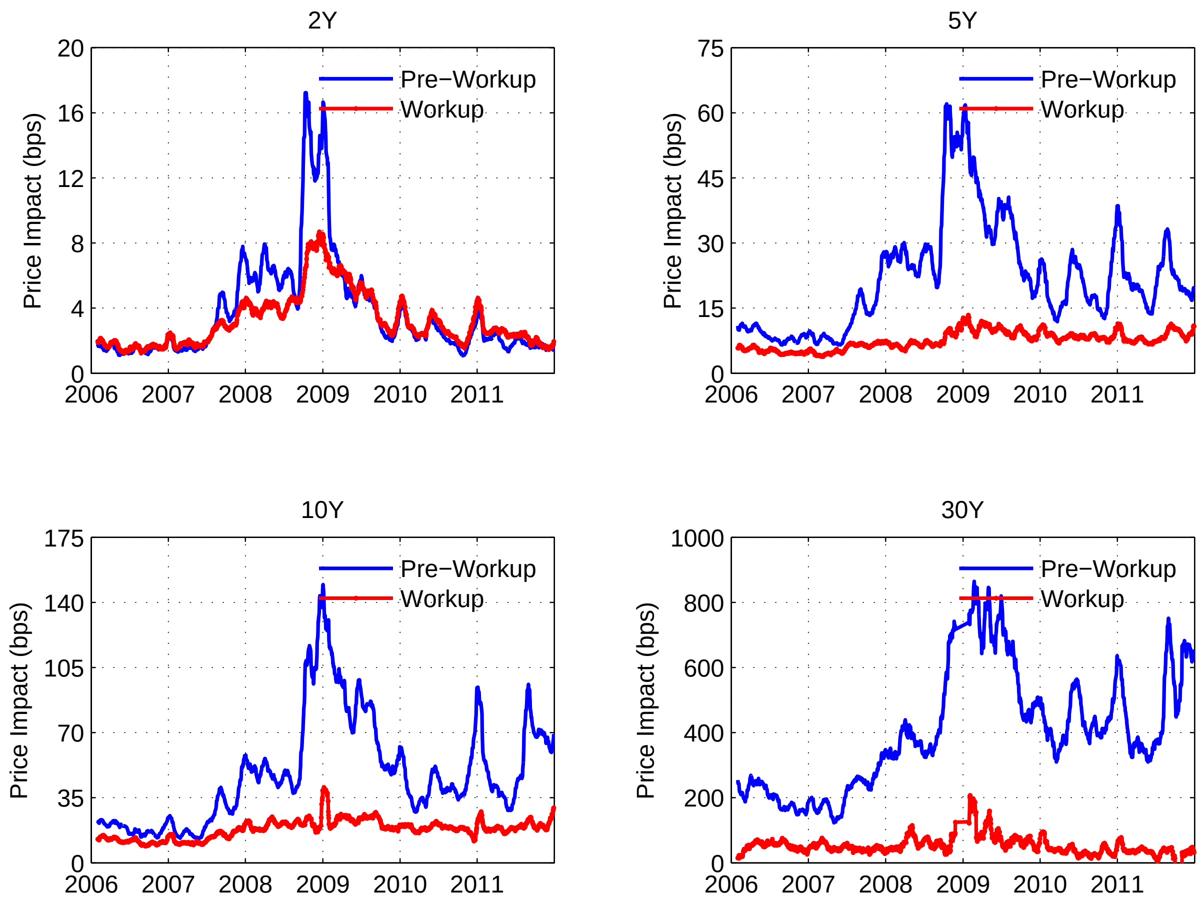

This figure plots the 20-day moving average of the price impact of $1 billion buyer-initiated volume transacted during pre-workup versus workup phases. The price impact measures are first computed daily from a VAR(5) model of return and trade flows, and then averaged over rolling 20-day intervals. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. Observations outside the [7:00-17:30] time window are excluded from model estimation.

Figure 3.6: Information Share of Pre-Workup and Workup Order Flow

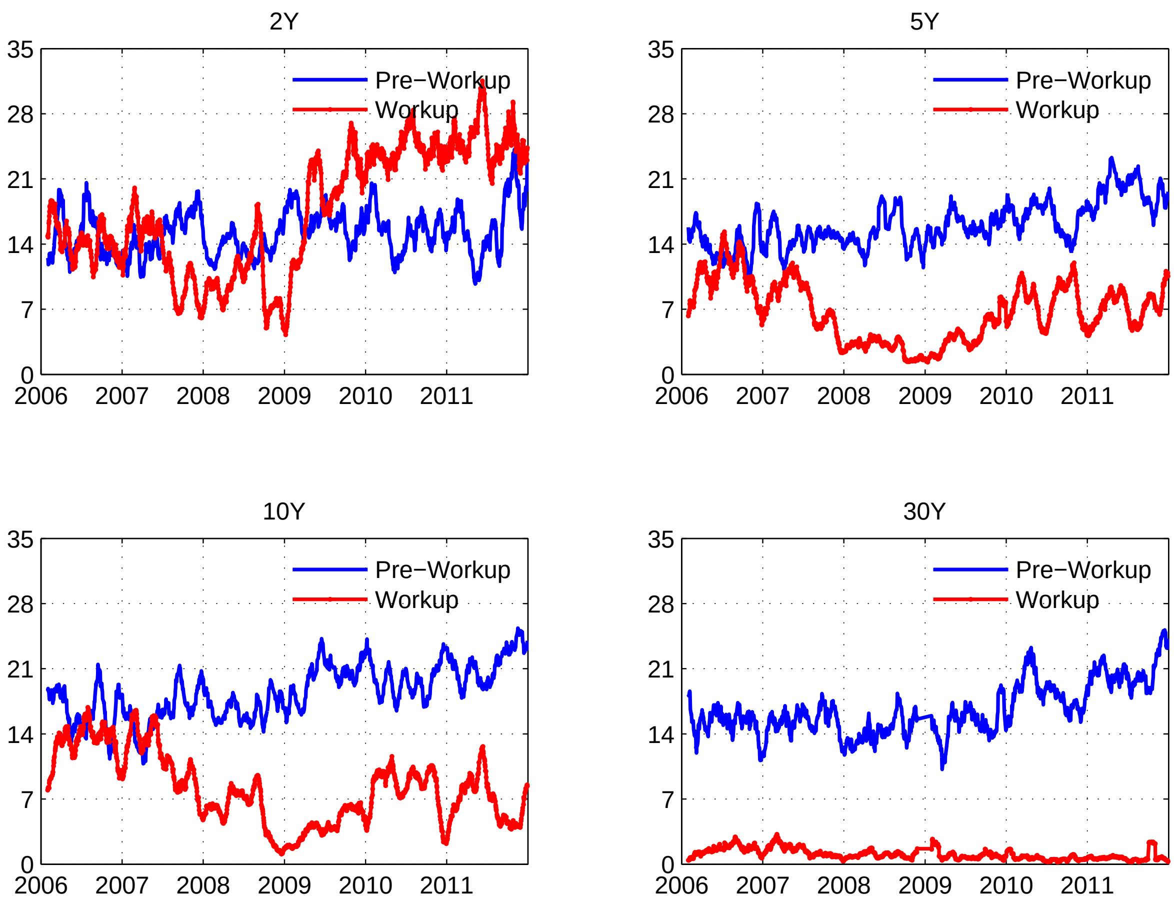

This figure plots the 20-day moving average of the information share of pre-workup versus workup order flow, using Hasbrouck (1991b)'s information share approach. The information share measures are first computed daily from a VAR(5) model of return and trade flows, and then averaged over rolling 20-day intervals. Estimation is based on BrokerTec data for the on-the-run 2-, 5-, 10- and 30-year Treasury securities over the period 2006-2011. Observations outside the [7:00-17:30] time window are excluded from model estimation.

CHAPTER 4
INTRADAY DYNAMICS OF VOLATILITY AND LIQUIDITY IN THE US TREASURY MARKET

4.1 Introduction

Interest in the dynamics of market liquidity and volatility in the U.S. Treasury securities market stems from the market's many vital roles. Because of their liquidity, 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.

The flight-to-liquidity premium in Treasury bond prices documented by Longstaff (2004) is a good example of how the plentiful liquidity in the Treasury market is valued by investors. This poses several interesting questions for the U.S. Treasury market. Is liquidity supply available when it is needed most? How is liquidity supply driven by uncertainty and other market factors, and conversely, does the supply of liquidity have any role in dampening or magnifying volatility in the market? How do the dynamics of the Treasury limit order book differ during flight-to-safety episodes?

A dynamic model for liquidity and its interrelation with volatility is highly useful for addressing these questions and exploring other microstructure issues of interest. 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 in this market.

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 et al. (2001), Huang et al. (2002), Fleming (2003), Brandt and Kavajecz (2004), Green (2004), Fleming and Piazzesi (2005), Goldreich et al. (2005), Mizrach and Neely (2007), Pasquariello and Vega (2007), Fleming and Mizrach (2009), Jiang et al. (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.

Being a safe haven for investors, the role of the Treasury market during flight-to-safety episodes is particularly important. An active literature studying the flight-to-safety phenomenon has provided (1) a number of theoretical models (see for example, Vayanos (2004) and Brunnermeier and Pedersen (2009)) and (2) related empirical evidence (see for example, Longstaff (2004), Goyenko and Sarkissian (2008), Baele et al. (2010), Baur and Lucey (2009), Beber et al. (2009), Bansal et al. (2010), and Baele et al. (2012)). While these studies provide great insights into the potential determinants of flight-to-safety episodes, such as the elevated level of risk, the changing risk aversion of investors, the tightening of margin requirements, and so on, little attention has been paid to how the destination of such flights – the Treasury market – is affected by the actions of those investors seeking safe haven. Our work aims to fill this gap by documenting the behavior of liquidity and volatility during such episodes and by providing an econometric model to isolate the effect of flights to safety on this benchmark market.

Our study is related to papers that have documented asset pricing anomalies that arose during the financial crisis. Fleckenstein et al. (2010) show that a significant mispricing arose during the crisis between Treasury bonds and inflation-swapped TIPS issues with replicating cash flows. Musto et al. (2011) document a large and systematic mispricing during the crisis between notes and bonds with identical cash flows. Hu et al. (2011) show that "noise" in Treasury security prices rose sharply during the crisis. Our work also documents the 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 Treasury securities.

Our study is also relevant for the general market microstructure literature on price discovery. Price incorporates news and converges to fundamental value through the trading process. The availability of liquidity is critical to that process and therefore modelling the evolution of liquidity can complement and further our knowledge on the dynamics of asset prices. Although equity limit order books have been studied extensively, studies on Treasury limit order books remain scant in comparison. There is no a priori reason to expect empirical findings from equity markets to hold up in the Treasury market. Informed trading is important in equity markets, but less so in the Treasury market, which is driven more by macroeconomic conditions and, in particular, monetary policy decisions and macroeconomic data releases. As a result, the dynamics of price and liquidity in this market could potentially differ from that documented for equity markets.

While market liquidity can be measured in many ways, we focus on market depth – a direct measure of the quantity of securities available for purchase and sale. Henceforth, we will use the two terms liquidity and depth interchangeably.[1] It is useful to note that the distinction between liquidity supply and demand may not be clear in limit order markets. In their comprehensive survey of the limit order book literature, Parlour and Seppi (2008) describe how investors with a demand for liquidity may choose to post aggressive limit orders rather than market orders. Such limit orders have the flavor of both supply and demand. Therefore, while the liquidity available in the order book (on both sides) is often considered representative of the liquidity supply, some of this liquidity could potentially come from demanders of liquidity who happen to have more patience to wait for their orders to be executed at better prices than those who submit market orders for immediate execution. In this study, we adopt the traditional approach of considering limit orders as the supply side and market orders as the demand side.

We propose a new joint model of liquidity and volatility based on the multiplicative error model (hereafter "MEM") formally introduced in Engle (2002). There are important features of the U.S. Treasury market that make this modelling choice superior to the standard linear Gaussian framework adopted in many previous empirical models of the limit order book.[2] First, it has been shown that depth tends to disappear prior to economic news announcements (e.g., Fleming and Remolona (1999) and Fleming and Piazzesi (2005)). This study also documents a large liquidity drop in the fall of 2008. Therefore, the model must be able to accommodate zero or small values of depth with a reasonable probability mass. Secondly, both market depth and volatility are nonnegative variables, but under a linear Gaussian framework their predictions are not guaranteed to be nonnegative. Even if log transformation is used to avoid the nonnegativity issue, researchers run into the problem of exact zero values at which the logged depth or volatility is not defined. The log linear framework is also problematic for predicting small values of depth as these are implicitly treated as

[1] Market depths (in plural) refer to depth at multiple price tiers, while market depth (singular) refers to the depth at a particular price tier.

[2] For example, Ahn et al. (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 et al. (2009) propose a dynamic semiparametric factor approach to modelling liquidity supply, combining nonparametric factor decomposition for the order curve's spatial structure with VAR for time variations of factor loadings. Other studies similar in their use of VAR include Danielsson and Payne (2010) and Hautsch and Huang (2009), among others. Ranaldo (2004) uses an ordered probit regression framework to analyze how the state of the limit order book affects order submission strategy.

extreme events whereas empirical evidence tells us that small values of depth are not uncommon. The new class of models we suggest can easily handle spells of near-zero (and positive valued) liquidity. Lastly, our MEM-based model enjoys the benefit of modelling directly the variables of interest, not their log-transformed counterparts, which can be convenient in interpretation and forecasting.

The key insight from our model choice is that we make empirical limit order book models look much like asset price volatility models. This has several advantages. First, we can readily borrow many specifications and modelling strategies from the vast volatility literature. For example, we can study the effect of news via so called news impact curves, see e.g., Engle and Ng (1993). Second, we can easily study the interactions of volatility and limit order book dynamics within a well understood and unified framework. Third, nonnegativity of depth – and obviously volatility – is guaranteed within the MEM specification. This rules out nonsensical predictions and therefore addresses many of the issues discussed in the previous paragraph.

The cross-fertilization of insights from the volatility literature to that of limit order books goes beyond modelling strategies – it also pertains to measurement. In the past decade, the notion of so-called realized (price) volatility has been extensively studied (see the recent survey by Barndorff-Nielsen and Shephard (2007)). We introduce the notion of limit order book depth realized volatility – which measures the variability of liquidity using high frequency data. Namely, our modelling strategy consists of taking five-minute snapshots of the book as well as measuring one-second changes in the book. The latter allows us to compute for every five-minute interval the realized quadratic variation at all levels of the limit order book.[3] This provides us with a measure of liquidity risk, similar to the quadratic variation measure widely used in the volatility literature to characterize price risk. Thanks to this new measure of realized depth volatility, we can study the impact of liquidity uncertainty on the level of liquidity. Needless to say, the realized depth volatility is obviously also a nonnegative process. Hence, our modelling strategy is perfectly suited to include this new measure as well.

Using limit order book data for the 2-, 5- and 10-year U.S. Treasury notes over the period from 2006 to the end of the second quarter of 2010, our class of models identifies several key findings. First, the order book exhibits clustering in all three variables of interest: depth, price volatility, and liquidity volatility (except liquidity volatility at the first tier). More importantly, there is a negative feedback loop between market depth and price uncertainty at the inside bid and ask. For other price levels however, depths tend to lower

[3] Technically speaking the limit order book, unlike high frequency returns, is not a martingale difference sequence. We provide several robustness checks with respect to the drift specification.

subsequent price uncertainty but the reverse effect is not present. Submitting orders to the inside queues subjects dealers to adverse execution risk and thus, when price is volatile, it is easy to see why dealers tend to cut back on supplying liquidity at the best bid and ask. There is room for modification or cancellation of orders behind the market when price moves unfavorably, so price volatility is less of a concern in the decision to supply depth at the outer tiers. In contrast to the negative interaction between liquidity and price uncertainty, we find that an increased level of depth uncertainty tends to bring out more depth, and this is consistent across all price levels. This evidence seems to suggest that liquidity supply tends to increase when it is more valuable to the marketplace, consistent with the findings in Biais et al. (1995).

Examining the dynamics of depth and volatility during the crisis period, we find that both become more persistent during the crisis. This dangerous combination provides a great illustration to models of liquidity crashes (for example, Cespa and Foucault (2012)) in that bad shocks to either volatility or liquidity can intensify the negative feedback effect, leading to liquidity crashing while volatility spiking up. Our models also provide consistent evidence with the earlier literature that depth is withdrawn immediately before important economic announcements but then quickly gets refilled once the announcement is released, accompanied by a surge in trading activity and price uncertainty. Furthermore, the news impact curve – a concept standard in the volatility literature but novel in a limit order book context – shows evidence of an asymmetric response of market depth to negative price changes, whereas price volatility does not seem to discriminate between price increases and decreases. Price volatility instead appears sensitive to the magnitude of the value change only. The fact that many dealers take part on both sides of the market, and large price moves may be indicative of important events around which divergences of opinion often rise, could explain this behavior.

Our analysis of the Treasury market during flights to safety contributes new evidence to the discussions of this phenomenon. In particular, the ex ante liquidity supply, namely the limit order book, is substantially lower on flight days – those days when liquidity is especially needed. However, a high level of trading activity is also observed on those days, along with an elevated level of price uncertainty. These patterns collectively suggest that liquidity providers monitor the market more closely on these days and refrain from using limit orders to passively supply liquidity to the market.

The chapter is organized as follows. Section 4.2 presents stylized facts on trading, liquidity and volatility in the U.S. Treasury market. These stylized facts provide the motivation for our modelling approach based on the multiplicative error framework, which we discuss in Section 4.3. Practical issues with model estimation and the measurement of volatility are also covered in this section. In Section 4.4, we present and discuss the empirical dynamics of Treasury liquidity and volatility as estimated by our proposed class of models. We then provide an analysis of the Treasury market during flight-to-safety episodes in Section 4.5. Finally, Section 4.6 concludes.

4.2 The U.S. Treasury Market – Some Stylized Facts

U.S. Treasury securities are debt instruments sold by the U.S. government through public auctions and subsequently traded in the secondary market. The secondary market is structured as a multi-dealer, over-the-counter market, in which the dealers trade with their customers, the Federal Reserve Bank of New York, and one another. Interdealer trading prior to 1999 was based on a network of voice-assisted brokers. Fully electronic trading started in 1999 with the introduction of the eSpeed platform, followed by the BrokerTec platform in 2000. Mizrach and Neely (2006) estimate that the BrokerTec platform accounts for about 61 percent of all interdealer trading activity.

There are no clearly defined trading hours for this market. Instead, trading spans 22-23 hours per day during the week, commencing around the start of the trading day in Tokyo and fading off with the end of the trading day in New York. During these hours, dealers send in their orders, have their orders executed, or modify or cancel existing orders. Each order specifies the quantity and price, whether it is for purchase or sale, and whether the order is aggressive.[4] Limit orders, when submitted, are queued in the order book according to the price and time priority rules until executed or cancelled. Although trading spans almost the entire day, trading outside of the New York trading hours is sparse, so we limit our analysis to between 7:00 and 17:00 Eastern time.

4.2.1 Data Description

Our analysis is based on order book data from the BrokerTec platform. All order messages sent to this platform are captured and time-stamped to the millisecond. The order book snapshot data is constructed by accumulating these order changes at the corresponding price tiers from the beginning of the trading day. This results in a tick-by-tick dataset with market depths measured in millions of dollars (par value), and prices

[4] Aggressive, or market, orders are executed immediately against best available limit orders on the opposite side of the market. Passive, or limit, orders are queued in the limit order book at the corresponding price level. All orders, whether aggressive or passive, need to specify both quantity and price. Best limit orders on opposite sides with the same price are not automatically executed.

reported in 256ths of a point, where a point equals one percent of the par value. We focus our attention on the on-the-run 2-, 5- and 10-year notes, as these are the most actively traded Treasury securities. The 2- and 5-year notes are newly issued every month, while the 10-year note comes out every quarter with reopenings in the following month and – since November 2008 – two months.[5]

Our sample period is January 2, 2006 to June 30, 2010, and thus covers the financial markets crisis of 2008-2009, 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 (see Boyd (2007)).[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.

For our empirical analysis, we choose to work with the five-minute snapshot data, supplemented by the one-second snapshot data needed for the computation of the realized volatility measures The five-minute snapshot data are extracted from the tick data described above by taking the last observation of each five-minute interval between 7:00 and 17:00. This results in 120 observations per day, except for those days with an early market close. For such days, we discard data after the recommended closing time.[7] The one-second snapshot data is extracted from the tick data in the same fashion. Although the data is available tick-by-tick, our choice of the five-minute interval is to avoid data errors (e.g., erroneous order messages that enter and exit the book in split seconds) and microstructure noise inherent in ultra-high frequencies. The interval is also long enough for sufficient movements in the book, so that meaningful predictions can be made.

We focus our analysis on the best five price tiers on each side of the market. First, we know from Biais et al. (1995) that liquidity is not concentrated at the inside tier. Second, the five-minute sampling implies that depth at different tiers is relevant for the future evolution of the limit order book. We choose to look at five tiers, which mirrors what market participants can see.

---

[5]We apply the following filters to the data. Since we are exclusively looking at recently issued on-the-run securities, and Treasury securities are issued at a price close to par value, then prices have to be in the vicinity of 25,600 (par value). We adopt a range of 20,000 - 30,000 to remove outliers, a filter that is narrow enough to remove obvious price errors, but conservative enough to still capture valid but extreme prices. For depth variables, we adopt a filter of $10 billion of par value for each price tier, which is roughly one-third the size of the typical issue in our dataset.

[6]Note 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 the money markets is more appropriate.

[7]Information on recommended early closes in the bond market is from the Securities Industry and Financial Markets Association and is posted here: `http://www.sifma.org/uploadedfiles/research/statistics/statisticsfiles/misc-us-historical-holiday-market-recommendations-sifma.pdf`.

4.2.2 Some Stylized Facts

In this section, we document a number of stylized facts pertaining to the 2-, 5- and 10-year notes, with particular focus on their time series behavior through the most recent financial crisis, as well as their patterns around macroeconomic announcements.

Market Depths Decline Sharply During Crisis

As shown in Table 4.1, prior to the crisis, the average depth at the best price tier for the 2-year note is over \$400 million. It plummets to roughly one fifth this level during the crisis period. Coming out of the crisis, the market recovers somewhat, but remains far below the pre-crisis level. Similar trends can be observed for the 5- and 10-year notes: the pre-crisis average depth at the best price tier is about \$72 million, before dropping to a level slightly above \$20 million during the crisis. However, unlike the partial recovery observed with the 2-year note, liquidity does not seem to improve much after the crisis for the 5- and 10-year notes. Note that Treasury issue sizes have steadily increased over the sample period. So the decline in liquidity we observe during the crisis is not attributable to a declining issue size.

Another observation of interest is that average market depth is highest at the second tier and gradually declines over the subsequent tiers. This is consistent with the finding of Fleming and Mizrach (2009) using BrokerTec data from January 2005 to February 2006. That liquidity is not concentrated at the inside tier is also documented by Biais et al. (1995) using order book data for stocks on the Paris Bourse. They attribute this finding to the fact that trading consumes liquidity at the front line.

To see the trend in Treasury market depth throughout the crisis period more closely, we graph daily averages of depth in Figure 4.1 for the 2-year note (solid line). Graphs on the left are for the inside ask and on the right for the inside bid. Whether on the bid or ask side, 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 in the fall of 2008 following Lehman Brothers' bankruptcy. Depth then bottoms out towards the end of 2008 and then improves fairly steadily from there. Depths at other tiers as well as the total depth across the best five price tiers have the same time series pattern. The 5- and 10-year notes exhibit very similar trends (not shown), although their recovery following the crisis is not as strong as what is observed with their 2-year counterpart.

Spread Is Tight but Widens Significantly in Late 2008

The bid-ask spread is another useful indicator that supplements our characterization of liquidity in this market. An analysis of the inside spread shows that, for all three securities considered, the spread is quite tight, with an average of slightly above one tick.[8] The tight spread around the small minimum allowable price movement suggests a very liquid market. Its time series behavior - with the inside spread widening significantly and consistently across the three notes in late 2008 – is also consistent with the pattern in market depths that points to a large liquidity drop at the height of the crisis. The 5-year note exhibits the most extreme peak in this period, followed by the 10-year. The spread returns to a level of just over one tick again in 2009 (see Figure 4.2).

Trading Volume Does Not Drop Until Late 2008

In contrast to the drop in market depth – a measure of the ex ante liquidity supply – that happens right at the beginning of the crisis in August 2007, actual trading volume is on the rise during the first half of the crisis (see graphs (a) and (b) in Figure 4.1, where trading volume is depicted by the dash line). It is only after Lehman's failure that we observe a major slide in trading volume which continues largely until the end of 2008, when both market depth *and* trading activity seem to almost vanish altogether. From that point on, they improve and move together.

Volatility Shoots Up During Crisis

As can be seen from Table 4.1, price volatility roughly doubles during the crisis for all three securities. However, only the 2-year note's volatility returns close to its pre-crisis level, whereas volatilities for the two longer term notes remain higher than they were before the crisis. This evidence again supports the finding that the 2-year note is the most liquid and resilient among the three securities considered. Across the best five tiers in the book, price volatility is slightly higher for the outer tiers.

Graphs (c) and (d) in Figure 4.1 provide a closer look into the time series trends in the inside tier's price volatility. In the second quarter of 2007, price volatility is already rising and shoots to a new elevated level on August 9, 2007 – when the crisis is widely believed to start. It keeps increasing and reaches its peak around the time of Lehmans bankruptcy, after which it gradually declines and almost reverts back to its pre-crisis

[8] The tick size for the 2- and 5-year notes is one 128th of a point. It is one 64th of a point for the 10-year.

level. Price volatility at the inside tier closely resembles the trend in trading activity (as shown in graphs (a) and (b) in the same figure). Active trading intensifies the price discovery process and it is thus not surprising to see trading volume and price volatility moving closely together. For brevity, we do not show in this figure the price volatility at other tiers as they exhibit similar time series trends as just described.

For the depth volatility variable, Table 4.1 shows that the inside tier stands out as having significantly higher depth volatility than the outer tiers. This is a natural result since liquidity at the first tier inherits an additional source of randomness from trading. This randomness does not pertain to order execution alone. The pick-off risk inherent in posting orders at the inside tier may require more intense order management and modification activities. We also observe an increase in the volatility of depth during the crisis period, but at a more moderate rate than that of price uncertainty, and a weak post-crisis recovery. Since the depth volatility measure has been standardized by the depth level, it looks reasonably comparable across the three securities.

Liquidity and Volatility Exhibit Clear Intradaily Patterns

Figure 4.3 shows the intraday patterns of market liquidity and volatility measures at the first tier over five-minute intervals.[9] Since the patterns are quite consistent across securities, and across different price tiers for the same security, we show here only the 2-year note and only the first tier.

Depth in the book builds up in the morning, reaches its peak shortly before noon and gradually declines from there, especially after 15:00. There are major dips in depth shortly before 8:30, 10:00 and 13:00. Trading is most active in the morning hours and shows distinct jumps immediately after the drops in market depth described above. There is also a mild peak at 15:00, after which trading diminishes. The peak at 15:00 coincides with the pricing of fixed income indices and hence likely reflects increased trading demand by investment managers who are seeking to rebalance their portfolios, while minimizing tracking errors relative to the indices.

Price volatility is also generally higher in the morning, and fades off toward the end of the day. This is very different from the well-documented U-shape pattern of volatility in equity markets: high around opening and toward market closing. Instead, in the Treasury market, price volatility closely tracks the pattern of trading activity, which peaks in the morning and falls off gradually after 15:00.

[9]While previous studies have examined the intraday patterns of bid-ask spreads and price volatility in this market (e.g., Fleming and Remolona (1997), Fleming and Remolona (1999)), we believe our intraday analysis of depth and depth volatility is novel.

The volatility of liquidity exhibits a less clear intraday pattern than its price volatility counterpart, but also peaks at the 8:30 and 10:00 time marks. We see volatility spikes slightly lagging liquidity drops around the key time marks above, and in addition around 14:15.

Depths Disappear Immediately Before Announcements

The spikes at certain times documented in the intraday patterns of liquidity and volatility coincide with the release times of major economic announcements: 1) macroeconomic announcements released at 8:30, 2) macroeconomic announcements released at 10:00, 3) announcements of Treasury auction results shortly after 13:00, and 4) announcements of the Federal Open Market Committee's rate policy decision around 14:15 ("FOMC announcements").[10] For a complete list of major announcements, their frequency and time of release, see Appendix A.

To differentiate market behavior around these announcements, we separate days with each of the above news categories from days with no major news and examine the intraday patterns of liquidity and volatility on the news versus no-news days, as in earlier studies.[11] As evident in graphs (a)-(d) of Figure 4.4, in the short time window before an announcement, market depths largely disappear, especially depth at the first tier, but then immediately return to the book after the announcement has been released. This finding is consistent with the evidence documented in Fleming and Remolona (1999) and Fleming and Piazzesi (2005) that dealers often withdraw quotes before announcements due to inventory risk concerns.

An important observation is that on the days with FOMC rate decision announcements, the order book thins out rather gradually, starting from shortly before noon until reaching the minimum just before 14:15. The order book then refills in the next half hour or so and converges to its no-news day level. This pattern differs from that for other announcements, for which limit orders are cleared from the book just shortly before the announcement time. The anticipation leading up to FOMC announcements suggests that market

[10] Announcements after scheduled meetings, which occur eight times per year, were made at about 14:15 during our sample period. Announcements after unscheduled meetings, of which there were two in 2008, do not have a standard announcement time.

[11] We define "no news" days as days without any of the major announcements as listed in Appendix A. For the news days, we separately examine announcements released at a particular time (e.g., 8:30), but include days with announcements released at other times (e.g., 10:00). We therefore observe announcement patterns around the other release times (e.g., around 10:00 on 8:30 announcement days), although typically not as strong as those associated with the release time being examined.

participants consider monetary policy decision announcements so important to the market that they refrain from taking positions in the order book well before the announcement comes out.[12]

Trading and Volatility Jump Immediately After Announcements

Graphs (e)-(h) of Figure 4.4 document the intraday patterns in trading activity around major announcements, and graphs (i)-(l) the patterns in short-term price volatility. Trading activity and price volatility both spike in the five-minute window following announcements. The market is bustling during this short time window with the limit order book filling up, trading demand surging and price discovery intensifying. The market then gradually works its way back to the no-news day pattern in the next hour or so, reflecting the time it takes for disagreement over an announcement's implications to be resolved. These graphs once again show that trading activity and price volatility are highly related.

4.3 A New Class of Dynamic Limit Order Book Models

The evidence presented in the preceding section shows that Treasury market depth can sometimes have zero or low values (e.g., at the peak of the crisis, or immediately before economic announcements). The average frequency of low values of depth, i.e., depth being equal to 1 (the minimum order size on BrokerTec) or 0, equals 4% for the inside price tier (bid and ask) for both the 5- and 10-year notes across our full sample, and is naturally much higher on certain days. Likewise, realized volatility at the five-minute frequency is often zero. For the 2-year note, the realized volatility of price equals zero for 22% of the sample observations, regardless of tier, and is again much higher on numerous days.

In our search for a modelling framework that can accommodate zero or low values with a realistic probability distribution and integrate the dynamics of market depth and volatility in one, the multiplicative error model is particularly fitting. In this section, we start with a description of a general MEM formulation proposed in Engle (2002) and explain how this model choice is novel for the Treasury limit order book. We then specify the details of our model, as well as our measurement of volatility. Lastly, we discuss practical issues with the model estimation.

[12]Fleming and Piazzesi (2005) find that uncertainty about the exact announcement time leads to a protracted reduction in liquidity. However, this only explains the decline about five minutes before and not two hours before such events.

4.3.1 Multiplicative Error Model for nonnegative Valued Processes

The general formulation of an MEM model is as follows. Let X_t be a nonnegative time series of interest. Its dynamics is modelled as:

$$
\begin{aligned}
X_t &= \mu_t \epsilon_t, && (4.1)\\
\epsilon_t | \Im_{t-1} &\sim D(1, \psi), && (4.2)\\
\mu_t &= \omega + \sum_{i=1}^{p} \alpha_i x_{t-i} + \sum_{j=1}^{q} \beta_j \mu_{t-j} + c' z_{t-1}, && (4.3)
\end{aligned}
$$

where $\Im_{t-1}$ presents the information set at time $t-1$, ϵ_t is the multiplicative error with a conditional distribution D having unit mean and defined on nonnegative support, and z_t are weakly exogenous variables. The persistence of X_t is captured by $\sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{q} \beta_j$. The model can be estimated with the exponential quasi log likelihood function

$$
\ln \mathcal{L}(\mathbf{X}; \theta) = -\sum_{t=1}^{T} \left[\ln \mu_t + \frac{X_t}{\mu_t} \right]. \tag{4.4}
$$

The asymptotic properties of the QML estimator have been established in Engle (2002). Hautsch (2012) notes that Newey and West (1987) standard errors are robust not only against distributional misspecification but also against dynamic misspecification in the MEM errors.

The general vector MEM is specified similarly. Let $\mathbf{X}_t$ be a K-dimensional process with nonnegative components. The dynamics of $\mathbf{X}_t$ are specified as follows:

$$
\begin{aligned}
\mathbf{X}_t &= \boldsymbol{\mu}_t \odot \boldsymbol{\epsilon}_t, && (4.5)\\
\boldsymbol{\epsilon}_t | \Im_{t-1} &\sim D(\mathbf{1}, \boldsymbol{\Sigma}), && (4.6)\\
\boldsymbol{\mu}_t &= \boldsymbol{\omega} + \sum_{i=1}^{p} \mathbf{A}_i \mathbf{X}_{t-i} + \sum_{j=1}^{q} \mathbf{B}_j \boldsymbol{\mu}_{t-j}, && (4.7)
\end{aligned}
$$

where $\odot$ indicates the element-by-element product, $\boldsymbol{\mu}_t$, $\boldsymbol{\epsilon}_t$, $\boldsymbol{\omega}$ are Kx1 vectors and $\boldsymbol{\Sigma}$, $\mathbf{A}_i$, $\mathbf{B}_j$ are KxK matrices.

To explain our technical contribution, it is worth elaborating on the key difficulties of the standard linear Gaussian framework – usually adopted among current limit order book models – in modeling nonnegative variables like market depth. In the simplest form, such a framework specifies the dynamics of a variable X_t

as $X_t = \mu_t + \epsilon_t$. As discussed in Engle (2002), the requirement that the conditional mean is positive means that the corresponding error term has to be no more negative than the mean to ensure the nonnegativity of X_t. Accordingly, the range of the error term changes with every observation, presenting a difficulty to estimation. Second, even if $\log(X_t)$ is used to avoid the nonnegativity issue, researchers run into the problem of exact zero values at which $\log(X_t)$ is not defined. Therefore, taking the log is not a solution when zeros are valid observations. It is also not a solution when small values are common, as the log-linear model would imply an extreme event probability to these values.

With its multiplicative error structure, the MEM formulation ensures the nonnegativity of X_t as long as the conditional error distribution has a unit mean and nonnegative support, for which there are many possible candidates. The structure allows us to model market depth and volatility directly, and assigns reasonable probability to low values of these variables. This is important for the Treasury market because, as we saw earlier, the order book thinned out substantially during the crisis period and low values of depth are common immediately prior to important economic announcements. Furthermore, during quiet times, high frequency realized volatility is often zero. Therefore, if we are to model liquidity and volatility in one framework, that framework should be able to accommodate zero or very small values of the dependent variables with appropriate probability distributions.

Additionally, the GARCH-type nature of the multiplicative error model allows us to effectively capture the persistence of market depths. The persistence in the limit order book at intraday frequencies has been documented in the prior literature, see for example Biais et al. (1995). Intuitively, depths queue at different price levels in the book waiting to be executed by coming trades. Over a short time interval, say five minutes, we do not expect these queues to vary substantially, especially in the outer tiers, which are not reached until trades or order cancellations exhaust liquidity at the first tier.

The ability of multiplicative error models to capture the nonnegativeness and persistence of a dynamic process gives this class of models an important place in the finance literature, since many financial series possess these properties. Important applications of this framework in finance include the modeling of conditional trade duration (see Engle and Russell (1998)), volatility, trading volume and intensities (see Manganelli (2000)), volatility, average trade sizes, trading costs and number of trades (see Hautsch and Jeleskovic (2008)), and absolute returns, daily range and realized volatility (see Engle and Gallo (2006)). However, despite it being a valuable tool for modeling nonnegative valued processes, we have not seen an application of this framework among limit order book models.

The closest work to ours, in terms of modeling framework, is by Hautsch and Jeleskovic (2008). That paper provides a review of the technique and applies it to the modeling of the dynamics of volatility and trade characteristics using order book data from the Australian Stock Exchange, but not the evolution of limit order depths – our modeling object of interest. On the other hand, while the paper by Russell and Kim (2010) models precisely this object, it adopts a different approach. Specifically, for both the buy and sell side, they model the total market depth on the given side and combine it with an estimated distribution of the depth across price levels. Their model therefore never predicts zero depth at any individual price level, a scenario that could plausibly happen in this market as previously discussed. Our newly introduced class of order book models based on the multiplicative error structure allows us to model the evolution separately for each price tier, be able to predict economically sensible possibilities, and uncover interesting insights into the dynamics of liquidity in this important market.

4.3.2 Model Specification

We specify a joint MEM model of order (1,1) for three variables – market depth, price volatility, and depth volatility – as formulated in equations 4.5 - 4.7. The model is estimated separately for each of the best five price levels on both sides of the market, resulting in 10 systems of equations in total. Following Engle and Gallo (2006), we assume a diagonal variance covariance matrix for the error terms, acknowledging that there can be a loss of efficiency if this assumption is false. We capture the possible interdependence among the variables by allowing a fully parameterized coefficient matrix $\mathbf{A}$, and restrict matrix $\mathbf{B}$ to be diagonal. We then estimate the model equation-by-equation using exponential quasi log likelihood function as specified in equation 4.4 and compute Newey-West standard errors for our estimates. To avoid the overnight effect, we reinitialize the conditional mean of each variable at the beginning of each day, using the average over the 7:00 - 7:55 period of that day, and estimate the model using data from 8:00 through 17:00.

To fix notation, for each of the five price tiers on each side of the market, the vector $\mathbf{X}$ consists of depth (D, or $X^{(1)}$), realized volatility of price (RVP, or $X^{(2)}$) and realized volatility of depth (RVD, or $X^{(3)}$). The dynamics of $\mathbf{X}$ are modelled at a five-minute frequency, indexed by t, as:[13]

[13]More precisely, the time interval t is the product of day d and time of day j where $d = 1, 2, ..., D$ with D being the total number of days in the sample, and $j = 1, 2, ..., J$ with J being the total number of five-minute intervals in a day. In this study, $J = 120$ for a typical trading day, and $J < 120$ for those days when the bond market closes early.

$$
\begin{aligned}
\mu_t^{(1)} &= (\omega_1 + \alpha_1 X_{t-1}^{(1)} + \beta_1 \mu_{t-1}^{(1)}) + \gamma_2 X_{t-1}^{(2)} + \gamma_3 X_{t-1}^{(3)} + \theta' \mathbf{Z}_{t-1}, \\
\mu_t^{(2)} &= (\omega_2 + \alpha_2 X_{t-1}^{(2)} + \beta_2 \mu_{t-1}^{(2)}) + \gamma_1 X_{t-1}^{(1)} + \gamma_3 X_{t-1}^{(3)} + \theta' \mathbf{Z}_{t-1}, \\
\mu_t^{(3)} &= (\omega_3 + \alpha_3 X_{t-1}^{(3)} + \beta_3 \mu_{t-1}^{(3)}) + \gamma_1 X_{t-1}^{(1)} + \gamma_2 X_{t-1}^{(2)} + \theta' \mathbf{Z}_{t-1}.
\end{aligned} \tag{4.8}
$$

In the above equations, $\mathbf{Z}$ allows for other potential explanatory variables to enter the dynamics of $\mathbf{X}$. This enables a wide range of specifications designed to explore crisis effects, announcement effects, effects of price changes and any possible asymmetry between positive and negative changes via the so called news impact curve, the role of liquidity demand, and last but not least, effects of flights to safety on Treasury liquidity and volatility.

4.3.3 Measurement of Volatility

In this study, we use two volatility measures. One is the volatility of price, and the other is the volatility of depth. Each of the measures is described below.

Volatility of Price

To measure price uncertainty, we use the realized volatility of price for each of the five price levels on both sides of the market (RVP). Beside its simplicity in computation, the main advantage of this measure of volatility, as discussed by Andersen et al. (1999), is that it is effectively error- and model-free. For each five-minute interval and for each price tier, the realized volatility is computed as the square root of the sum of the squared second-to-second price changes:

$$RVP_t = \sqrt{\sum_{k_t=1}^{300} (\triangle P_{k_t})^2}, \tag{4.9}$$

where $\triangle P_{k_t}$ denotes the one-second price change at second k_t of the five-minute interval t. Since prices are reported in 256ths of one percent of par, this volatility of price inherits the same unit of measurement.

We note that this is a measure of total volatility comprised of both transitory and permanent components. We use this measure as we focus on the short-term dynamic interactions of volatility and liquidity in order to predict their evolution throughout a trading day at high frequency. It is important to note that this measure of volatility is based on intraday prices and not yields. The mapping of price changes to yield changes differs over time due to changes in coupon rates and time to maturity. At an intra-daily frequency, such as the five-minute frequency used in this study, the difference between the two methods of calculating volatility should be negligible.

Beyond the intraday boundary, however, the difference between volatility computed from prices and that from yields can magnify. Additionally, the trend in volatility over the sample period documented earlier could change if volatility were computed differently. For comparison, we plot in Figure 4.5 the monthly average of our five-minute price volatility over the sample period, together with a measure of volatility based on daily yields.[14] The latter volatility measure is computed for each month as the square root of the realized variance for the month. The realized variance is the sum of all squared daily yield (absolute) changes in that month.

As Figure 4.5 demonstrates, our measure of price volatility closely tracks the volatility of yields, except that it is more variable than the latter given that it is based on higher frequency data. They both document an elevation in uncertainty during the crisis period, peaking around the time of the Lehman Brothers bankruptcy.

Volatility of Depth

To measure liquidity uncertainty, we introduce the notion of realized depth volatility for each of the five price levels on both sides of the market (RVD), which we compute in a similar way to RVP:

$$RVD_t = \sqrt{\sum_{k_t=1}^{300} (\triangle D_{k_t})^2}, \tag{4.10}$$

where $\triangle D_{k_t}$ denotes the one-second depth change at second k_t of the five-minute interval t. The realized volatility computation requires that the one-second depth change series be a martingale difference sequence. To check the robustness of our RVD measure with respect to this assumption, we use an autoregressive specification at various lag lengths to estimate the time-varying drift of depth changes and then compute RVD as the five-minute sum of squared residuals of one-second depth changes.[15] The resulting RVD measures are very similar to the initial RVD computed from raw one-second depth changes, providing support for its use in our subsequent analysis. We also examine the auto-correlation function of the one-second depth changes

[14]Daily on-the-run Treasury yields for the 2-, 5- and 10-year maturities are from Bloomberg.

[15]We also experimented with various other specifications, such as including a constant drift or allowing for linear dependence via ARMA models.

and find them to exhibit no significant serial correlation. Finally, depth in the limit order book can change due to trade execution, limit order submissions, modifications and cancellations. Therefore, our RVD measure captures the extent of liquidity supply and demand imbalances in the order book.

Our realized volatility of depth measure complements well the realized price volatility in enhancing our understanding of the various sources of uncertainty in the market. On the one hand, the volatility of liquidity is closely related with the volatility of price. When the book is fluctuating actively, the resulting temporary order imbalances induce increased short-run price volatility (see Handa and Schwartz (1996)). This is particularly likely to happen with a thin book. For example, a market order can create a large imbalance causing the price to change, or sweep more than one tier's depth causing the price queues to move forward. In that case, we would expect the volatility of depth and volatility of price to move together and leave similar effects on liquidity supply. On the other hand, for a deep book, it is possible for the depth to change without an accompanying change in price. Accordingly, examining the volatility of depth and whether it helps to predict the subsequent level of liquidity provides understanding not possible with the volatility of price alone.

For our empirical model estimation, we standardize our liquidity volatility measure by the corresponding market depth at the beginning of each five-minute interval to remove the scale effect, making this effectively a measure of volatility relative to the size of the limit order book.

4.3.4 Diurnal Pattern Adjustment

As shown earlier, our variables of interest, namely depth, volatility variables and trading volume, exhibit clear diurnal patterns, making it necessary to remove such seasonality before the models can be estimated. We choose a non-parametric method to adjust for this intraday pattern by dividing the relevant variable at a given time interval by the average for that time interval, essentially assuming a multiplicative seasonality effect. For example, market depth observed at the end of the 9:15-9:20 interval on a given day is adjusted by the level of depth typically observed at that time.

We are careful to account for the different levels of depths and volatility across the pre-crisis, crisis and post-crisis periods by using three sets of intraday averages corresponding to the three sub-sample periods, as opposed to just one set of overall sample averages. In addition, we compute these sets of intraday averages from days without any important public announcements concerning FOMC rate decisions, macroeconomic conditions or Treasury auctions to avoid distorting the typical intraday seasonality with markedly different patterns observed on days with such announcements.

For depth variables, we favor the median as the average statistic in the seasonality adjustment to avoid possible distortion caused by extreme depth values on the seasonally-adjusted depths. For all other variables, we use the mean instead of the median as the latter can be zero at certain quiet intervals of the day.

4.4 Empirical Analysis

In this section, we present and discuss the dynamics of depth and volatility as revealed by our MEM-based class of models. Given the qualitative similarity in findings and for brevity, we report here the results for the 2-year note only.

4.4.1 The Baseline Model

We first estimate the baseline model as specified in equation (4.8) with no covariates and present the results in Table 4.5. The table has three panels: the top one for liquidity dynamics, the middle for price volatility dynamics, and the bottom for depth volatility dynamics. We flag those coefficients that are not significantly different from 0 based on the 5% significance threshold with an asterisk $*$. We discuss each panel in turn.

Liquidity Dynamics

As expected, order book depth exhibits a high level of persistence, as represented by the sum $\alpha + \beta$ being close to 1. In addition, depths at the best bid and ask prices are negatively impacted by price volatility, whereas depths at outer tiers tend to increase with volatility. This may reflect the unwillingness of market participants to supply depth at the first tier for fear that their orders will be adversely executed in a volatile market. Yet, at the same time, higher volatility may help increase the probability of execution for limit orders at outer tiers, making the option inherent in these orders more valuable for the limit order traders. Therefore, during volatile times, depth could move away from the first tier and toward the outer tiers.

The negative linkage between price volatility and depth at the first tier is also consistent with the evidence documented by many previous studies, including Næs and Skjeltorp (2006), that an increased level of trading is often associated with moments of high price volatility. More active trading could deplete liquidity in the book that is not subsequently replenished fast enough, especially if potential liquidity suppliers hesitate to supply liquidity to a volatile market. The shrinkage of liquidity following a rise in volatility is also in line with a proposition put forth by Caballero and Kurlat (2008) that when asset price volatility rises, the risk of illiquidity rises.

The variability of order book depth also shows a significant effect on the subsequent level of depth. Following intervals of active movements in the book, liquidity supply becomes more plentiful. This is intuitive, whether one interprets the volatility of depths as representative of supply uncertainty (e.g., that associated with increased order modifications/cancellations), or as a sign of a strong demand for liquidity (e.g., that associated with increased trading activity). In the former interpretation, the uncertainty of supply may increase the payoffs to those who can actually provide liquidity, thus encouraging them to supply more. Similarly, if it is the strong demand for trading that induces frequent movements of liquidity, the demand would increase compensation for liquidity provision and subsequently invite more depth. This finding supports Biais et al. (1995)'s conclusion that more liquidity is supplied when it is valuable to the marketplace.

Price Volatility Dynamics

We see evidence of price volatility clustering, although not to the same extent as the clustering of depths. More importantly, lower depth at any tier predicts a subsequent increase in price volatility, consistent with Parlour and Seppi (2008)'s assessment that prices are more volatile in thin markets, as the lack of liquidity hinders the price discovery process, causing more uncertainty about the security value. This evidence is also consistent with the idea that depth is withdrawn in advance of expected price changes (e.g., macroeconomic announcements). That is, causality may be reversed, with expected volatility leading to lower depth.

The liquidity-volatility feedback loop at the inside price tier on both sides of the market isolates the effect of adverse execution risk on dealers' liquidity supply decision. This dynamic interaction could help explain episodes of liquidity and volatility feeding on each other and exacerbating a bad shock that could originate from either the liquidity side or the volatility side. The liquidity drop and the heightened volatility during the recent financial crisis as shown in Figure 4.1 is a good example. Our evidence provides empirical support for the theory of liquidity crashes put forth by Cespa and Foucault (2012), although the focus of their theory is on the liquidity-volatility feedback among multiple assets. Here we provide evidence that this theory is also at work for one asset.

Let us turn to the role of depth uncertainty on price uncertainty. The different sign of the effect for the first tier, as compared to the rest of the book, deserves some close attention. Increased volatility of depth at the first tier predicts lower price volatility, while the opposite is true for other tiers. Apparently, depth at the first tier changes in large part due to trading, but depths at the other tiers change mostly by order addition, modification, or cancellation. So in the former case, if we associate the high volatility of depth with active trading in the market, it is understandable how price volatility can be subsequently reduced. The latter case, however, is a better reflection of supply uncertainty on the part of liquidity providers, which arguably means they are also uncertain about security value, explaining the accompanied increase in price volatility.

Depth Volatility Dynamics

The volatility of liquidity also exhibits clustering at price levels behind the market. At the inside tier, the persistence is quite low, adding to our discussion earlier that the liquidity volatility measure at this tier seems to be closely related with trading activity in the market. Next, increasing liquidity level can predict a subsequent decrease in liquidity volatility relative to the size of the book, but we note that this effect is typically small at the first tier. Finally, price uncertainty is positively related with liquidity uncertainty, although we believe the mechanism is again different between the first tier and the other tiers. At the inside bid and ask, high price uncertainty is associated with increased trading that induces greater variability of depth. For outer price tiers, the fluctuation of depths is often the result of limit order submission, modification, or cancellation activities, which also tend to intensify when price is volatile.

4.4.2 Announcement Effects

Prior literature on the Treasury market response to economic news, such as Fleming and Remolona (1999), Balduzzi et al. (2001), and Green (2004), has documented strong patterns of liquidity and volatility around economic announcements. The analysis of intraday patterns of depths and volatilities in this study further confirms that depths largely disappear from the order book immediately prior to an economic announcement, but quickly return thereafter. Both trading activity and volatility are high following an announcement. To formally test these patterns, we estimate an MEM model with pre-announcement and announcement time dummies for each of the three dependent variables indexed by i (where $i = 1, 2, 3$ corresponding to depth, price volatility, depth volatility respectively). The specification is as follows:

$$\mu_t^{(i)} = \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j \neq i} \gamma_j X_{t-1}^{(j)} + \theta_1 \text{prenews2}_{t-1} + \theta_2 \text{prenews1}_{t-1} + \theta_3 \text{news}_{t-1}, \quad (4.11)$$

where $prenews2$ is the dummy for the second-to-last five-minute interval before an announcement, $prenews1$ is the dummy for the five-minute interval before an announcement, and $news$ is the dummy for the five-minute interval containing the announcement release. Since these dummies enter the equations with a lag, θ_1, θ_2 and θ_3 in fact capture the pre-announcement, announcement and post-announcement effects respectively on the liquidity and volatility variables.

As shown in Table 4.6, market depths decline significantly across the top five price levels in the five-minute interval before an announcement, but then increase by a larger magnitude in the announcement interval, reflecting the fast refilling of order book depth once an announcement is released. Depths in the post-announcement five-minute interval continue to show some further increase but the effects start to subside. On the other hand, both measures of volatility are already at elevated levels in the pre-announcement interval, consistent with the withdrawal of orders from the book causing greater fluctuation of prices and depths. The announcement interval witnesses the peak in volatilities, especially the volatility of depth, that come about with the refilling of limit orders and the surge in trading activities following the news arrival. In the five-minute interval following an announcement, both depth and price volatilities are significantly lower, suggesting that the most intense price discovery and order book activities happen within a very short time window. Depth and price volatilities remain high in the next thirty minutes or so, as compared to non-announcement days, but the peak has passed.

4.4.3 Dynamics During Crisis

To explore whether any of the above dynamics change during the recent crisis, we estimate a specification that incorporates a crisis period dummy, allowing it to have both intercept and interactive effects. For i = $1, 2, 3$ corresponding to depth, price volatility, and depth volatility respectively, the specification is as follows:

$$\begin{aligned} \mu_t^{(i)} &= \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j \neq i} \gamma_j X_{t-1}^{(j)} + \theta_1 \mathrm{DC}_{t-1} + \theta_2 \mathrm{prenews2}_{t-1} \\ &\quad + \theta_3 \mathrm{prenews1}_{t-1} + \theta_4 \mathrm{news}_{t-1} + \theta_5 \mathrm{DC}_{t-1} X_{t-1}^{(i)} + \theta_6 \mathrm{DC}_{t-1} \mu_{t-1}^{(i)}, \end{aligned} \tag{4.12}$$

where DC is the dummy for the crisis period, defined to be from August 9, 2007 through June 30, 2009, $prenews2$ is the dummy for the second-to-last five-minute interval before an announcement, $prenews1$ is the dummy for the five-minute interval before an announcement, and $news$ is the dummy for the five-minute interval containing the announcement release. With this specification, θ_1 is the estimate for the level effect of the crisis, while θ_5 and θ_6 collectively show the effect of the crisis on the persistence of each dependent variable. The pre-announcement, announcement and post-announcement effects are captured by θ_2, θ_3 and θ_4 respectively, as in the earlier specification with announcement effects (Equation (4.11)).

Table 4.7 shows the model estimates. The pre-announcement, announcement and post-announcement effects remain qualitatively similar to those previously documented. Depth is lower during the crisis period, consistent with evidence presented up to this point. A surprising observation is that price volatility is also lower, contradicting the model-free descriptive analysis performed earlier. We suspect that outliers in price volatility during the crisis period may skew our seasonality adjustment factors upward (as these are based on the mean), resulting in lower than expected diurnally-adjusted price volatility.

The key observation from this table is that depth and price volatility tend to be more persistent during the crisis period. The immediate implication of the higher degree of persistence is that bad shocks to these variables take longer to fade away. Considering the negative depth-price volatility feedback loop documented earlier, this finding illustrates how this negative feedback effect can intensify in a crisis.

4.4.4 Limit Order Book Dynamics and the News Impact Curve

In this section, we build a model that captures any asymmetric response of liquidity, price volatility and liquidity volatility to the changing value of the securities, in the spirit of the news impact curve technique introduced by Nelson (1991) for GARCH models. This original framework is designed to allow the conditional variance process of a given asset's returns to respond asymmetrically to positive and negative price changes. The question we want to address is whether the movement in the best bid-ask midpoint has any bearing on the dynamics of the order book. The data shows that the distance between any two adjacent price levels in the order book is almost always one tick, so the movement in the best bid-ask midpoint is a good indicator of the overall ups and downs of order book prices.

We specify the news impact curve ("NIC") as:

$$\text{NIC}_t = \theta_1 |Ret_t| + \theta_2 |Ret_t| \mathbf{1}_{Ret<0}, \tag{4.13}$$

where $Ret_t^{S,i}$ is defined as the five-minute return (annualized log return) of the best bid-ask midpoint. This functional form for the NIC particularly suits our needs as the NIC will enter the dynamics of depths and volatilities as a positive covariate. With this specification, the coefficient θ_1 captures the effect of changing price on the dynamics of order book liquidity and volatility, while the coefficient θ_2, if significantly different from zero, will indicate an asymmetric response of liquidity and volatility to negative price changes.

We start with a simple specification that has only the NIC, shown in Table 4.8. We note that the results on other variables of the model are qualitatively the same as those obtained with the baseline specification. Key findings, such as the negative liquidity-volatility feedback loop at the first tier, and the association of greater depths with subsequently lower volatility of both depth and price, remain. We therefore focus our discussion in this section on the NIC coefficients.

The results show that depth generally responds to price movements, although the response is not uniform across all price tiers. The asymmetric coefficient θ_2 is also mostly significant, implying that negative price movements impact limit order book depth differently. Second, with regard to price volatility, it seems that large price movements predict a subsequent increase in price volatility, regardless of the direction. There is no evidence for an asymmetric response of price volatility to positive versus negative price changes. We hypothesize that many dealers submit orders on both sides of the market, i.e., perform a market-making function, so the direction of the price change does not have much of an impact on the price uncertainty, only the magnitude does. Large price swings may be indicative of important news that intuitively could result in an increased divergence of opinions among market participants. Lastly, concerning depth volatility, we document a significant news impact curve function (both the magnitude and asymmetric effects) only up to the second or third tier. Beyond that, value changes do not seem to matter.

We also estimate a specification with the NIC controlling for announcement effects. Although not reported here, the results are consistent with our basic findings.

4.4.5 Effect of Liquidity Demand

In this section, we examine the effect of liquidity demand, as indicated by the volume of market orders, on the limit order book. First and foremost, trading provides the means for price discovery and accordingly is expected to affect price volatility. Trading consumes depth in the book, can stimulate additional liquidity supply (e.g., Biais et al. (1995)), or can change the distribution of depths across price levels. Typically we would expect a negative impact due to the consumption effect, but the stimulus effect might also come into play. Nevertheless, either effect would have the same impact on the volatility of depth.

As shown in Table 4.9, the estimated effect of trading volume (θ_3) on subsequent market depth is positive, especially for tiers at or near the market, suggesting that liquidity is supplied at a faster rate than it is consumed and supporting to some extent the hypothesis that trading demand might help attract additional liquidity supply. As expected, the variability of depth increases with trading volume, but this effect shows up at the first tier only.

In addition, trading activity does not significantly affect subsequent price volatility at the first tier, but rather increases price volatility only at other tiers. This result still holds after we control for announcement effects. The empirical facts documented earlier, with first tier price volatility and trading volume exhibiting strikingly similar patterns over time and over the course of a trading day, can be relied upon in interpreting the model estimates. Apparently, trading and price volatility at the inside tier are contemporaneously related so that once we control for lagged volatility, lagged trading volume has little incremental explanatory power. Beyond the first tier, however, trading volume still shows its relevance in predicting subsequent price volatility, indicating that the effect of trading activity travels to the outer tiers' price volatility with a lag.

4.5 Liquidity and Volatility During Flights to Safety

We now turn to an analysis of the Treasury market during flight-to-safety episodes. Prior research mainly focuses on understanding the motives of such flights, e.g., whether investors are seeking the high quality and/or high liquidity of Treasury securities. Evidence seems to favor the liquidity motive for flights. For example, Beber et al. (2009) show that euro-area bond investors chase liquidity rather than credit quality during times of market stress. Likewise, Longstaff (2004) documents a liquidity premium in Treasury securities, as large of 15% of their values. However, the question of whether the sought-after liquidity is actually there when it is needed the most remains. Addressing this question is the main objective of this section.

4.5.1 Identification of Flights to Safety

We first describe how flights to safety are identified. It is widely observed, and agreed, that a flight to safety occurs when investors withdraw in droves from risky asset markets and move to safe/liquid asset markets. A common flight to safety is the flow out of equity markets and into the Treasury market. Such a flight is often accompanied by an extreme negative equity return concurrent with an extreme positive Treasury return. This is the basis for the identification of flight-to-safety episodes proposed by Baele et al. (2012) - an approach that we adopt in this study.

Specifically, let r_T and r_E be the daily return on the relevant on-the-run Treasury note and the S&P 500 index respectively.[16] σ_T and σ_E correspond to their sample return volatility. The flight-to-safety dummy, FTS is defined as:

$$FTS = \mathbf{1}_{r_E < -\kappa\sigma_E} * \mathbf{1}_{r_T > \kappa\sigma_T} \tag{4.14}$$

where κ is the parameter for the severity of the flight.

We examine three different levels of κ, i.e., 1, 1.5, 2 – which we will refer to as "light", "moderate" and "severe" flights. Appendix B shows the dates of these flights, as well as the total count of flights for each security and each severity level. With 1,124 trading days in the sample period from January 2006 through the second quarter of 2010, the light, moderate and severe flights occur on approximately 6%, 2-3% and 0.8-1.2% of the days respectively. The last quarter of 2008 contains a disproportionately large number of flights. In particular, the majority of the severe flights happen in the aftermath of the Lehman bankruptcy, especially on September 15, 29 and October 6 of 2008 when severe flights to safety occur in all three securities. The identification of moderate FTS episodes in the 2- and 5-year notes also picks up August 9, 2007, which marks the beginning of the crisis.

It is worth noting that using the returns on the 5- and 10-year notes helps identify slightly more FTS episodes than when the 2-year note returns are used. Light and moderate flights tend to happen most frequently with the 5-year note, whereas for severe flights, the 10-year note returns pick up the highest frequency of such episodes. One possible explanation is that the Fed lowered rates to the zero bound, thereby anchored the 2-year, creating a more stable pattern for the shorter maturity. Indeed, flights to the 2-year note seem to occur mostly in the earlier period of the sample which is less affected by the zero bound.

---

[16] We use the daily on-the-run "dirty" prices for the 2-, 5- and 10-year notes from Bloomberg in computing the relevant daily Treasury return for the identification of flights. It is worth noting that there is a discontinuity in price of any given on-the-run Treasury note series on those days when a security goes off-the-run and a new issue becomes on-the-run. To take care of these discontinuities, we first compute daily returns separately for each issue, and then splice together the returns over only the on-the-run period of each issue. The notes are on-the-run from the day following their auction through the day of the following auction (for that maturity).

4.5.2 How Different Are Flight-to-Safety Days?

We compare order book depth, trading volume, price volatility and depth volatility on FTS days with non-FTS days and find them to be markedly different. The comparison across the three Treasury notes sheds light onto which one is affected by these episodes the most. For the discussion below, we use moderate FTS days, i.e., days when the Treasury return exceeds 1.5 times its sample volatility and the S&P 500 return falls below -1.5 times its sample standard deviation.

In terms of trading volume, the average daily figure for the 2-year note on FTS days is over $57 billion whereas the number on non-FTS days is only nearly $34 billion. On a typical non-FTS day, selling pressure dominates buying pressure, with a net selling volume of $652 million. This reverses to an average net buying volume of $945 million on FTS days. Likewise, the number of trades is almost double the typical non-FTS level, and the reversal in the net number of trades is more dramatic than is the case with the volume of trades.[17] The average trade size, being 4.35 on FTS days as compared to 4.65 on other days, suggests that market participants submit smaller sized orders than usual on FTS days.

The 5- and 10-year notes also have higher trading volume on FTS compared to non-FTS days, but to a lesser extent than the 2-year note. As with the 2-year note, we observe net selling pressure in the 5- and 10-year notes on non-FTS days, about $465 and $321 million respectively. However, both notes show changes on FTS days that are less dramatic than those for the 2-year note. For the 10-year note, net volume reverses (as with the 2-year note) so that there is net buying averaging $338 million on FTS days. For the 5-year note, net selling is weaker on FTS days, averaging $222 million, but net volume remains negative.

Differences between FTS and non-FTS days for the 5- and 10-year notes are more striking when we look at the number of trades as opposed to trading volume, as is the case with the 2-year note. That is, the overall increase in activity is greater on FTS days when looking at the number of trades, and the net number of trades flips from negative (i.e., net selling pressure) on non-FTS days, to positive (i.e., net buying pressure) on FTS days for the 5- as well as the 10-year note.

These pieces of evidence together suggest that there is a higher level of trading demand on FTS days, particularly on the buy side, as we would expect to be the case when the Treasury market absorbs the flow of investors fleeing risky asset markets. The differences are especially striking when looking at the number of

[17] We note that the number of trades is not the same as the number of market orders. When a market order is executed against multiple limit orders, the system records each of them as one trade. Thus a market order can correspond to multiple trades. Nevertheless, the number of trades and volume of trades collectively provide a complete picture of trading activity in the market.

trades, suggesting that order sizes are smaller on FTS days. The 2-year note appears to be especially affected by flights to safety.

Even more striking evidence from our analysis is that market depth, representing the ex ante liquidity supply (or the willingness to provide liquidity), is much lower on FTS days, as shown in Figure 4.6. In particular, the order book for the 5- and 10-year notes on FTS days thins out to a greater extent than that for the 2-year note. This lack of willingness to post limit orders spreads over all price tiers, and not just the first tier where depth could naturally be lower if the trading rate exceeded the limit order submission rate.

We in turn examine the behavior of volatility to understand better how the book on FTS days could thin out so much. Both volatility measures are much higher on FTS days, especially price volatility. The evidence seems supportive of our conjecture that, despite the strong demand for liquidity, dealers become more conservative with their liquidity supply on FTS days when the market is highly volatile. They provide less depth to reduce adverse execution risk and allow themselves increased flexibility to respond to the highly volatile market conditions on these days.

4.5.3 Liquidity and Volatility Dynamics and Flights to Safety

We now turn to our econometric framework to explore whether the dynamics of liquidity and volatility change on FTS days. In particular, we allow for the FTS dummy to affect both the level as well as the persistence of liquidity and volatility. For this section, we use a univariate MEM specification for each of the liquidity and volatility variables in order to expose the FTS effect cleanly on each variable's dynamics. Specifically, each variable's conditional mean equation takes the following form:

$$\mu_t = \omega + \alpha X_{t-1} + \beta \mu_{t-1} + c\, FTS_t + c_\alpha FTS_t X_{t-1} + c_\beta FTS_t \mu_{t-1}.$$

Under this specification, the coefficient c captures the level effect of FTS, while c_α and c_β captures the change in the dynamics of the modelled variable. We estimate this equation for market depth (presented in Table 4.3) and price volatility (presented in Table 4.4) for each security and each of the three levels of FTS severity corresponding to three thresholds adopted for κ (1, 1.5, 2).

In general, flights seem to have more of an effect on the dynamics than the level, given that the depth series have been diurnally adjusted to remove intraday regularities. c_α is the coefficient on the interaction between the relevant FTS dummy and the past market depth, hence it captures the marginal impact of news on market depth on FTS days. c_α is mostly negative and significant, suggesting that on those days when flights occur, the immediate past realization of market depth has lower predictive value for depth at the next interval. On the other hand, c_β tend to be somewhat positive on these days, essentially indicating that market participants place greater importance on the historical path of market depth than on just its most recent realization in predicting the next level of market depth. Therefore, while the persistence of depths on FTS days does not significantly differ from that on non-FTS days, the composition has changed that gives less weight to the impact of news on market depth. With respect to the level effects, the 2-year note shows a somewhat lower depth level on FTS days, especially at the inside bid and ask, whereas the evidence is more mixed with the other two notes.

We now describe the effects of flights on price volatility. Consistent with evidence presented earlier, the level of price uncertainty is significantly higher on flight-to-safety days, although statistical significance seems to diminish with the severity of flights – a direct consequence of a much lower number of observations with the most severe flights to safety. Examining the effects of flights on the dynamics of volatility via the estimates for c_α and c_β, we can see that the news impact coefficient c_α is usually not significantly different from that on non-flight days, except for the positive news impact on volatility at some outer price tiers. The other coefficient c_β often has opposite sign to c_α. Similar to the evidence with depth, we also find that while volatility does not exhibit a significantly different level of persistence on flight-to-safety days, there has been a shift in the relative importance of the news impact coefficient on volatility especially at some of the outer price tiers.

4.6 Conclusion

In this study, we propose a new class of dynamic order book models for the purpose of exploring the micro dynamics of depth, price volatility and depth volatility in the interdealer market for the 2-, 5- and 10-year U.S. Treasury notes. Our models are based on the multiplicative error framework introduced by Engle (2002). This class of models offers important advantages that are highly suited to the modeling of the Treasure limit order book. Zero or very low values of market depths are not uncommon, particularly around economic announcements and during the crisis. Likewise, intraday volatility is often zero during quiet times of the day. The MEM guarantees the prediction of nonnegative depths and volatility measures and allows us to integrate liquidity and volatility into a unified framework from which their dynamic interactions can be studied. It also goes beyond the log linear framework in that it allows for more flexible and realistic probability distributions. Additionally, by modeling the limit order book in a similar fashion to asset price volatility models, we can capitalize on the vast literature in the latter to tailor our model specifications in ways that can capture the dynamics between liquidity and volatility as closely as possible.

In addition to the novel use of the MEM framework to model the dynamics of the limit order book, we also introduce the notion of realized volatility of depth, which is parallel in concept to realized volatility of price. Furthermore, apart from testing market microstructure hypotheses, our proposed class of models can be used for the purpose of forecasting liquidity and managing liquidity risk.

Our empirical analysis examines market depth and volatility around economic announcements, through the crisis and during flight-to-safety episodes. Consistent with earlier studies on the impact of economic announcements, we document an important stylized fact that depths tend to disappear before announcements but return shortly thereafter, together with a surge in trading activity and a jump in price volatility which takes an hour or so to fade away. We offer additional facts about the Treasury market over the crisis not previously documented, that is, the order book thins out substantially over the crisis, coupled with an elevated level of price volatility, although trading activity does not decline substantially until the second half of the crisis.

Our models' key finding is that price volatility and depth at the first price tier exhibit a negative relationship, which runs in both directions. This negative feedback effect becomes more pernicious during the crisis when both of these variables are evidently more persistent. This helps explain spells of liquidity deteriorating as volatility increases, and conversely, liquidity improving as volatility decreases, especially observed during the crisis. Last but not least, our study of the Treasury market during flights to safety shows that market depth is substantially lower despite the higher demand for trading on these days. The inflows of trading interest and the accompanied rise in price uncertainty may necessitate greater market monitoring and reduce dealers' incentive to supply liquidity via limit orders.

Table 4.1: Summary Statistics of Depth and Volatility

		Depth($ million)						Price Volatility						Depth Volatility					
		Before Crisis		Crisis		After Crisis		Before Crisis		Crisis		After Crisis		Before Crisis		Crisis		After Crisis	
		Mean	StD	Mean	StD	Mean	StD	Mean	StD	Mean	StD	Mean	StD	Mean	StD	Mean	StD	Mean	StD
2Y	Ask5	287.0	131.0	69.8	56.2	180.9	106.9	2.9	3.0	6.8	4.2	3.4	2.9	0.9	12.9	2.9	10.1	1.1	9.8
	Ask4	338.7	150.8	86.7	65.5	226.5	123.3	2.9	2.9	6.6	3.6	3.4	2.8	1.1	14.3	2.8	9.5	1.2	14.1
	Ask3	473.7	219.5	116.8	84.2	284.4	143.5	2.9	2.8	6.5	3.3	3.4	2.7	1.7	19.7	2.9	12.0	1.1	15.5
	Ask2	647.2	296.9	143.0	105.4	325.5	190.9	2.8	2.8	6.5	3.3	3.4	2.7	1.2	14.8	2.6	8.6	1.7	20.6
	Ask1	399.4	285.8	80.9	81.6	187.1	143.9	2.8	2.8	6.5	3.2	3.4	2.7	13.0	69.2	16.6	43.5	11.5	46.9
	Bid1	418.2	292.8	82.4	84.1	189.3	137.6	2.8	2.8	6.4	3.2	3.4	2.7	13.1	69.7	16.7	46.7	10.7	44.4
	Bid2	676.4	302.2	146.0	107.6	327.7	165.0	2.8	2.8	6.5	3.2	3.4	2.7	1.0	11.3	2.5	8.4	1.1	6.8
	Bid3	498.0	225.5	120.7	88.8	293.0	145.2	2.9	2.8	6.5	3.3	3.4	2.7	1.6	18.3	2.8	9.1	0.9	6.8
	Bid4	355.1	155.9	89.2	67.8	233.2	125.3	2.9	2.9	6.6	3.6	3.4	2.8	1.0	8.9	2.7	8.8	0.9	4.9
	Bid5	298.0	134.0	72.3	59.1	187.0	110.4	2.9	3.0	6.9	4.4	3.4	2.9	0.8	8.4	2.8	8.5	0.9	4.2
5Y	Ask5	67.2	38.5	21.4	19.8	50.7	32.9	6.2	6.1	14.1	14.3	9.1	7.1	2.1	9.5	5.5	10.0	2.4	5.9
	Ask4	83.9	45.4	26.1	22.2	62.4	36.9	6.0	4.9	13.0	9.8	8.8	5.2	2.6	11.4	5.6	9.8	2.0	5.5
	Ask3	139.4	75.8	36.4	29.5	70.3	40.0	5.9	4.1	12.1	6.6	8.7	4.3	3.3	17.4	5.0	9.8	1.9	5.0
	Ask2	164.6	85.2	41.4	33.8	65.7	38.8	5.8	3.6	11.6	4.9	8.6	3.8	2.6	13.5	5.1	9.9	3.1	7.7
	Ask1	72.6	57.9	21.4	22.9	28.0	25.5	5.8	3.3	11.4	4.3	8.5	3.5	15.0	36.5	20.2	36.7	19.2	37.5
	Bid1	73.0	58.0	21.7	23.9	28.5	25.9	5.8	3.3	11.4	4.3	8.5	3.6	15.1	37.3	20.3	35.9	18.8	36.8
	Bid2	165.1	83.8	42.0	34.4	66.8	39.4	5.8	3.5	11.5	4.8	8.6	3.7	2.6	13.6	5.1	9.7	3.0	7.0
	Bid3	140.6	75.4	36.9	30.2	71.1	40.2	5.9	4.0	12.0	6.4	8.6	4.1	3.3	17.6	5.1	10.4	1.9	5.1
	Bid4	86.6	46.2	26.1	22.1	62.8	36.8	6.1	5.1	13.0	10.0	8.7	5.0	2.5	10.6	5.8	10.0	2.0	4.6
	Bid5	69.7	39.3	21.3	19.8	51.3	32.7	6.2	5.8	14.0	13.3	8.9	6.2	2.0	8.7	5.8	10.4	2.4	5.7
10Y	Ask5	67.3	34.8	22.7	17.5	42.8	23.8	11.4	8.1	22.7	18.5	16.4	9.1	1.3	4.1	3.7	6.6	2.2	4.0
	Ask4	80.2	39.3	26.9	19.7	53.4	27.0	11.3	7.3	21.5	14.0	16.3	7.8	1.6	5.2	3.7	6.1	1.8	3.3
	Ask3	127.3	67.2	37.0	27.5	61.2	30.3	11.2	6.8	20.4	9.9	16.1	7.1	2.9	10.1	3.7	6.7	1.7	3.7
	Ask2	166.3	76.3	45.2	34.7	56.6	30.1	11.1	6.4	19.8	8.0	16.1	6.8	2.1	9.7	3.5	7.3	2.6	5.3
	Ask1	71.8	55.4	22.1	23.6	24.3	20.7	11.1	6.3	19.5	7.2	16.0	6.6	15.2	36.8	17.1	29.2	15.1	27.5
	Bid1	72.5	56.2	22.6	24.6	24.7	21.2	11.1	6.3	19.4	7.1	16.0	6.5	15.7	37.5	17.4	29.9	14.8	25.5
	Bid2	167.0	75.2	46.1	36.3	56.9	29.6	11.2	6.3	19.6	7.5	16.0	6.6	2.0	9.0	3.6	7.6	2.6	5.9
	Bid3	127.0	65.5	37.4	28.1	61.4	30.4	11.2	6.5	20.0	8.8	16.1	6.9	2.8	9.0	3.8	7.3	1.6	3.6
	Bid4	81.4	39.5	27.5	20.3	53.6	27.2	11.3	7.0	21.1	12.7	16.2	7.7	1.5	4.3	3.7	6.7	1.8	3.5
	Bid5	69.4	35.5	23.1	17.9	42.9	24.1	11.4	7.6	22.4	17.3	16.4	8.7	1.3	3.9	3.8	7.5	2.3	4.0

Price volatility is the five-minute realized volatility of price, in 256ths of one percent of par. Depth volatility is five-minute realized volatility of depth, standardized by the depth level. The statistics are reported for three periods: 1) Before Crisis (January 2006-August 8, 2007), 2) Crisis (August 9, 2007-June 2009), and 3) After Crisis (July 2009-June 2010).

Table 4.2: Average Daily Trading Volume and Number of Trades on FTS and non-FTS days

	2-Year Treasury Note		5-Year Treasury Note		10-Year Treasury Note	
	non-FTS	FTS	non-FTS	FTS	non-FTS	FTS
Trading Volume ($M of Par)	33,775	57,425	30,627	39,946	25,971	31,440
Buyer-Initiated Volume	16,561	29,185	15,081	19,862	12,825	15,889
Seller-Initiated Volume	17,214	28,240	15,546	20,084	13,146	15,551
Net Volume (Buy - Sell)	-652	945	-465	-222	-321	338
Number of Trades	7,264	13,190	12,532	18,419	12,041	17,602
Number of Buy Trades	3,580	6,767	6,194	9,267	5,963	8,934
Number of Sell Trades	3,683	6,423	6,338	9,151	6,078	8,668
Net Number of Trades (Buy - Sell)	-103	344	-144	116	-116	266

This table shows the average daily trading volume and number of trades on days with a flight to safety ("FTS") and days without such an episode ("non-FTS"), using BrokerTec trade data for the 2006-2010Q2 period. Flights are identified by a large positive return on the Treasury note and a large negative return on the S&P500 index, based on a 1.5 standard deviation threshold. Daily trading volume is the total volume exchanged during the 7:00-17:00 time period. Similarly, daily number of trades is the total number of order executions during the same time period.

Table 4.3: Liquidity Dynamics With Flight-to-Safety Effect

	2-Year Treasury Note			5-Year Treasury Note			10-Year Treasury Note		
	c	c_α	c_β	c	c_α	c_β	c	c_α	c_β
Panel A: Flight–to–Safety Dummy Based on $\kappa = 1$									
Ask5	−0.003*	-0.119	0.098				0.000*	-0.076	0.049
Ask4	-0.009	-0.101	0.093	0.003*	-0.067	0.054	−0.002*	-0.086	0.065
Ask3	−0.004*	-0.072	0.060	−0.001*	−0.022*	0.008*	0.001*	−0.039*	0.014*
Ask2	0.003*	−0.035*	0.009*	0.014	−0.037*	−0.009*	0.004*	−0.046*	0.020*
Ask1	-0.021	-0.045	0.035						
Bid1	-0.014	-0.043	0.005*	0.006*	-0.037	−0.025*	−0.004*	−0.022*	-0.029
Bid2	−0.001*	-0.087	0.067	0.004*	-0.048	0.018*	0.001*	-0.046	0.022*
Bid3	−0.002*	-0.068	0.052	0.003*	−0.038*	0.004*	0.006*	−0.038*	0.005*
Bid4	−0.004*	-0.060	0.047*	−0.000*	-0.077	0.043*	0.002*	-0.068	0.041*
Bid5	-0.007	-0.084	0.073	−0.000*	-0.077	0.043	−0.005*	-0.083	0.067
Panel B: Flight–to–Safety Dummy Based on $\kappa = 1.5$									
Ask5	0.002*	-0.096	0.066*	-0.012	-0.086	0.067	0.004*	-0.080	0.034*
Ask4	−0.006*	-0.085	0.077*	−0.004*	-0.076	0.042*	−0.005*	-0.109	0.082
Ask3	−0.008*	-0.087	0.085	0.004*	−0.025*	−0.007*	0.003*	−0.063*	0.023*
Ask2	0.014	0.037*	−0.074*	0.013	-0.060	0.007*	0.010*	−0.062*	0.013*
Ask1	−0.007*	−0.005*	−0.006*	−0.001*	-0.041	−0.020*	0.007*	−0.033*	-0.055
Bid1	-0.020	-0.091	0.057	0.010*	−0.030*	-0.049	−0.002*	-0.036	−0.045*
Bid2	−0.003*	-0.090	0.070*	0.012	−0.042*	−0.010*	0.004*	-0.068	0.027*
Bid3	0.002*	−0.066*	0.040*	−0.006*	-0.075	0.048*			
Bid4	−0.005*	−0.066*	0.057*	−0.005*	-0.069	0.036*	−0.001*	-0.101	0.064
Bid5	−0.006*	-0.120	0.108	−0.004*	-0.083	0.048*	−0.006*	-0.117	0.088
Panel C: Flight–to–Safety Dummy Based on $\kappa = 2$									
Ask5	0.006*	-0.155	0.086*	−0.010*	−0.083*	0.035*	0.011*	−0.046*	−0.038*
Ask4	-0.010	-0.164	0.143	0.004*	-0.163	0.071*	−0.006*	-0.101	0.066*
Ask3	−0.009*	-0.103	0.091*	−0.000*	−0.020*	−0.028*	0.013*	−0.050*	−0.038*
Ask2	0.009*	−0.061*	−0.016*	0.027	−0.088*	−0.050*	0.016*	−0.096*	0.015*
Ask1	-0.031	-0.066	0.032*	0.031*	-0.130	−0.078*	0.070*	-0.074	−0.168*
Bid1	-0.037	-0.113	0.090	0.173	−0.068*	-0.412	0.008*	-0.108	−0.027*
Bid2	−0.004*	-0.116	0.083*	0.044	−0.043*	-0.147	0.007*	-0.100	0.034*
Bid3	0.003*	−0.100*	0.055*	0.007*	−0.088*	−0.000*	0.011*	−0.066*	−0.009*
Bid4	−0.001*	−0.001*	−0.031*	0.009*	-0.146	0.035*	0.009*	−0.065*	−0.011*
Bid5	−0.009*	-0.154	0.134	0.029	−0.072*	−0.091*	−0.002*	-0.106	0.052*

This table shows estimates for the model of market depth with a flight-to-safety effect for the 2-, 5- and 10-year Treasury notes: $\mu_t = \omega + \alpha X_{t-1} + \beta\mu_{t-1} + c\,FTS_t + c_\alpha FTS_t X_{t-1} + c_\beta FTS_t \mu_{t-1}$. FTS is the flight-to-safety dummy, and equal to 1 if the S&P return falls below $-\kappa$ times its sample standard deviation while the Treasury note return exceeds κ times the latter's sample standard deviation. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level.

Table 4.4: Volatility Dynamics With Flight-to-Safety Effect

	2-Year Treasury Note			5-Year Treasury Note			10-Year Treasury Note		
	c	c_α	c_β	c	c_α	c_β	c	c_α	c_β
Panel A: Flight-to-Safety Dummy Based on $\kappa = 1$									
Ask5	0.098	0.023*	−0.018*	0.194	0.074	-0.172	−0.018*	0.077	0.002*
Ask4	0.110	0.019*	−0.019*	0.076	0.035*	−0.046*	0.012*	0.068	−0.022*
Ask3	0.123	0.006*	−0.013*	0.059	0.024*	−0.015*	0.036*	0.045*	−0.018*
Ask2	0.140	0.022*	−0.037*	0.081	0.060*	−0.065*	0.019*	0.058*	−0.021*
Ask1	0.117	0.009*	−0.015*	0.117	0.050*	−0.090*	0.082	0.039*	−0.053*
Bid1	0.165	0.019*	−0.061*	0.105	0.049*	−0.080*	0.124	0.005*	−0.056*
Bid2				0.111	0.096	-0.130	0.128	0.046*	-0.103
Bid3	0.145	0.016*	−0.044*	0.056*	0.005*	−0.001*	0.100	0.004*	−0.038*
Bid4	0.125	0.021*	−0.037*	0.029*	0.035*	−0.014*	0.063*	0.072	−0.077*
Bid5	0.097	0.021*	−0.027*	0.026*	0.080	−0.058*	0.048*	0.087	-0.074
Panel B: Flight-to-Safety Dummy Based on $\kappa = 1.5$									
Ask5	0.052*	0.002*	0.046*	0.161	0.007*	−0.058*	−0.008*	0.115	−0.020*
Ask4	0.102	0.021*	0.009*	0.070	−0.039*	0.050*	0.024*	0.061*	0.001*
Ask3	0.160	0.038*	−0.042*	0.101	−0.011*	0.004*	0.039*	0.017*	0.034*
Ask2	0.123	0.040*	−0.022*	0.082*	0.079*	−0.074*	−0.006*	0.049*	0.027*
Ask1	0.151	0.020*	−0.025*	0.106*	0.026*	−0.044*	0.110*	0.024*	−0.034*
Bid1	0.152	0.009*	−0.022*	0.137	0.060*	−0.096*	0.220	−0.007*	−0.088*
Bid2	0.174	0.020*	−0.043*	0.110*	0.157	-0.161	0.227	0.046*	−0.146*
Bid3	0.174	−0.013*	−0.016*	0.052*	−0.037*	0.060*	0.206	0.004*	−0.087*
Bid4	0.091	0.016*	0.006*	0.050*	0.084	−0.064*	0.060*	0.086*	−0.067*
Bid5	0.065*	0.024*	0.003*	0.040*	0.112	−0.090*	0.054*	0.101	−0.067*
Panel C: Flight-to-Safety Dummy Based on $\kappa = 2$									
Ask5	0.057*	0.026*	0.048*	0.241	−0.010*	−0.050*	−0.037*	0.115*	0.011*
Ask4	0.086*	0.040*	0.023*	0.074*	−0.085*	0.117	−0.036*	0.225	−0.103*
Ask3	0.201	0.092*	−0.082*	0.157*	−0.026*	0.017*	0.053*	0.027*	0.031*
Ask2	0.139*	0.121*	−0.079*	0.057*	0.062*	−0.017*	−0.049*	0.100*	0.022*
Ask1	0.165*	0.054*	−0.038*	0.078*	−0.020*	0.033*	0.099*	0.029*	−0.011*
Bid1	0.137*	0.004*	0.013*	0.242	0.024*	−0.113*	0.317	−0.006*	−0.137*
Bid2	0.139*	0.014*	0.003*	0.261	0.197	-0.269	0.373*	0.015*	−0.195*
Bid3	0.177*	−0.010*	0.007*	0.039*	−0.060*	0.119	0.327	−0.003*	−0.147*
Bid4	0.083*	0.050*	0.000*	0.033*	0.086*	−0.034*	0.064*	0.052*	−0.026*
Bid5	0.039*	0.024*	0.034*	0.008*	0.104*	−0.035*	0.082*	0.075*	−0.048*

This table shows estimates for the model of realized price volatility with a flight-to-safety effect for the 2-, 5- and 10-year Treasury notes: $\mu_t = \omega + \alpha X_{t-1} + \beta\mu_{t-1} + c\,FTS_t + c_\alpha FTS_t X_{t-1} + c_\beta FTS_t \mu_{t-1}$. FTS is the flight-to-safety dummy and equal to 1 if the S&P500 return falls below $-\kappa$ times its sample standard deviation while the Treasury note return exceeds κ times the latter's sample standard deviation. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level.

Table 4.5: Liquidity and Volatility Dynamics: Baseline Estimates for 2-Year Treasury Note

	Ask5	Ask4	Ask3	Ask2	Ask1	Bid1	Bid2	Bid3	Bid4	Bid5
DEPTH ($X^{(1)}$)										
ω_1	0.005	0.000	0.006	0.006	0.015	0.014	0.001*	0.007	0.009	0.013
α_1	0.341	0.314	0.286	0.291	0.111	0.114	0.286	0.293	0.330	0.348
β_1	0.651	0.679	0.705	0.701	0.877	0.874	0.708	0.698	0.659	0.638
γ_2	0.004	0.008	0.001*	0.002	-0.005	-0.004	0.003	−0.000*	−0.001*	−0.001*
γ_3	0.001*	-0.000	0.002	0.001	0.003	0.002	0.002	0.003	0.003	0.003
RVP ($X^{(2)}$)										
ω_2	0.216	0.215	0.220	0.209	0.154	0.202	0.237	0.230	0.223	0.199
α_2	0.239	0.234	0.230	0.229	0.209	0.234	0.241	0.240	0.241	0.241
β_2	0.600	0.605	0.609	0.610	0.678	0.613	0.585	0.590	0.591	0.599
γ_1	-0.032	-0.033	-0.038	-0.029	−0.023*	-0.023	-0.042	-0.039	-0.036	-0.022
γ_3	0.001	0.001	0.001	0.002	−0.002*	-0.002	0.002	0.001	0.001	0.002
RVD ($X^{(3)}$)										
ω_3	0.364	0.468	0.634	0.633	0.561	0.431	0.525*	0.429	0.418	0.334
α_3	0.685	0.505	0.350	0.520	0.012	0.023	0.644*	0.435	0.486	0.696
β_3	0.215	0.143	0.137	0.023	0.018	0.072	0.074*	0.257	0.245	0.221
γ_1	−0.064*	−0.076*	-0.145	-0.077	-0.062	-0.033	−0.076*	−0.087*	−0.083*	−0.036*
γ_2	0.131*	0.252	0.357	0.411	0.874	0.974	0.216*	0.236	0.140	0.044*

This table shows estimates for the baseline model with no covariates for the 2-year Treasury note. The model is estimated jointly for depth ($X^{(1)}$), price volatility ($X^{(2)}$) and depth volatility ($X^{(3)}$). Each conditional mean equation is specified as: $\mu_t^{(i)} = \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j \neq i} \gamma_j X_{t-1}^{(j)}$. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level, based on Newey West robust standard errors.

Table 4.6: Liquidity and Volatility Dynamics With Announcement Effects for 2-Year Treasury Note

	Ask5	Ask4	Ask3	Ask2	Ask1	Bid1	Bid2	Bid3	Bid4	Bid5
DEPTH ($X^{(1)}$)										
θ_1	-0.147	-0.079	-0.056	-0.095	-0.157	-0.114	-0.090	-0.099	-0.059	-0.092
θ_2	0.123	0.139	0.259	0.280	0.200	0.077	0.287	0.304	0.303	0.315
θ_3	0.088	0.143	0.014*	0.005*	0.051*	0.128	-0.001*	0.035*	0.002*	0.256
RVP ($X^{(2)}$)										
θ_1	1.175	0.967	0.674	0.592	0.282	0.519	0.598	0.721	0.907	1.123
θ_2	0.543	1.350	0.862	0.635	0.838	0.678	0.626	0.568	0.533	0.560
θ_3	-0.917	-0.272	-0.165*	-0.672	-0.728	-0.587	-0.661	-0.666	-0.744	-0.891
RVD ($X^{(3)}$)										
θ_1	3.265	0.468	0.224	0.222	0.056	-0.195	0.405	0.191	0.490	3.320
θ_2	31.622	38.731	41.688	55.011	13.568	11.733	40.250	33.868	25.131	3.944
θ_3	-5.611	-13.968	-11.939	-11.615	-0.631	-1.407	-10.738	-15.354	-9.844	-0.620

This table shows estimates of the announcement effects for the 2-year Treasury note. The model is estimated jointly for depth ($X^{(1)}$), price volatility ($X^{(2)}$) and depth volatility ($X^{(3)}$). Each conditional mean equation is specified as: $\mu_t^{(i)} = \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j\neq i} \gamma_j X_{t-1}^{(j)} + \theta' \mathbf{Z}_{t-1}$, where $\mathbf{Z}$ represents other covariates. $\mathbf{Z} = (prenews2, prenews1, news)$ where $prenews2$ is the dummy for the second to last five-minute interval before announcements, $prenews1$ is the dummy for the five-minute interval before announcements and $news$ is the dummy for the five-minute interval containing announcement time. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level.

Table 4.7: Liquidity and Volatility Dynamics With Crisis Period and Announcement Effects for 2-Year Treasury Note

	Ask5	Ask4	Ask3	Ask2	Ask1	Bid1	Bid2	Bid3	Bid4	Bid5
DEPTH ($X^{(1)}$)										
θ_1	-0.040*	-0.030*	0.007	-0.003*	-0.035	-0.026	-0.004*	-0.008*	-0.012*	-0.035
θ_2	-0.033	-0.039	-0.052	-0.097	-0.136	-0.091	-0.095	-0.106	-0.032	-0.088
θ_3	-0.009*	-0.008*	0.271	0.305	0.212	0.076	0.306	0.331	0.257	0.137
θ_4	0.023*	0.055	-0.006*	-0.028*	0.020*	0.105	-0.033*	0.005*	0.172	0.090
θ_5	0.091	-0.009*	-0.124	-0.138	-0.111	-0.108	-0.170	-0.149	0.014*	0.056
θ_6	-0.027*	0.045*	0.116	0.142	0.147	0.135	0.173	0.156	0.010*	-0.016
RVP ($X^{(2)}$)										
θ_1	-0.122	0.006	-0.039	-0.146	-0.144	-0.151	-0.136	-0.146	-0.140	-0.133
θ_2	1.137	0.270	1.275	0.554	0.441	0.471	0.565	0.678	0.863	1.077
θ_3	0.514	1.331	2.015	0.606	0.651	0.634	0.603	0.543	0.515	0.540
θ_4	-0.885	0.709	-0.099*	-0.695	-0.578	-0.608	-0.687	-0.680	-0.733	-0.847
θ_5	0.017*	-0.176	0.226	-0.028	-0.031	-0.043	-0.026*	-0.028	-0.005*	0.009*
θ_6	0.077	0.101	0.120	0.146	0.146	0.164	0.135	0.143	0.113	0.091
RVD ($X^{(3)}$)										
θ_1	-0.123	-0.045	0.000*	-0.085	0.000*	-0.093	-0.045	-0.099	0.006	-0.137
θ_2	0.729	0.450	0.211	0.230	0.062	-0.178	0.401	0.192	0.485	0.978
θ_3	22.863	38.497	41.866	53.526	12.800	11.137	39.401	33.544	24.597	18.279
θ_4	-8.537	-13.682	-11.575	-10.592	-0.530	-0.907	-9.791	-15.777	-9.594	-6.085
θ_5	0.020	-0.097	-0.101	-0.119	-0.034	-0.026	-0.108	0.036	-0.098	-0.135
θ_6	0.021	0.016	0.006	0.009	-0.056	-0.057	0.010	0.010	0.014	0.155

This table shows the estimated crisis and announcement effects for the 2-year Treasury note from a joint model for depth ($X^{(1)}$), price volatility ($X^{(2)}$) and depth volatility ($X^{(3)}$). Each conditional mean equation is specified as: $\mu_t^{(i)} = \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j \neq i} \gamma_j X_{t-1}^{(j)} + \theta' \mathbf{Z}_{t-1}$, where $\mathbf{Z} = (DC, prenews2, prenews1, news, DC \times X^{(i)}, DC \times \mu^{(i)})$. DC is the dummy for the crisis period (August 9, 2007 - June 30, 2009). $prenews2$ is the dummy for the second to last five-minute interval before announcements, $prenews1$ is the dummy for the five-minute interval before announcements and $news$ is the dummy for the five-minute interval containing announcement time. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level.

Table 4.8: Liquidity and Volatility Dynamics With News Impact Curve for 2-Year Treasury Note

	Ask5	Ask4	Ask3	Ask2	Ask1	Bid1	Bid2	Bid3	Bid4	Bid5
DEPTH ($X^{(1)}$)										
ω_1	0.000	0.000	0.006	0.004	0.016	0.013	0.001*	0.007	0.007	0.012
α_1	0.240	0.306	0.287	0.291	0.111	0.114	0.287	0.296	0.333	0.350
β_1	0.752	0.689	0.704	0.701	0.877	0.874	0.708	0.696	0.658	0.638
γ_2	0.004	0.007	0.001*	0.001*	-0.005	-0.004	0.002	−0.001*	−0.002*	−0.002*
γ_3	0.001	0.000	0.002	0.001	0.003	0.002	0.002	0.003	0.003	0.003
θ_1	0.003	-0.000	0.002	0.002	-0.001	0.001*	0.000*	-0.001	0.000*	−0.000*
θ_2	-0.003	0.000	-0.003	-0.002	0.001	−0.000*	0.000*	0.003	0.002	0.002
RVP ($X^{(2)}$)										
ω_2	0.219	0.220	0.225	0.213	0.216	0.206	0.242	0.235	0.227	0.202
α_2	0.230	0.225	0.220	0.218	0.221	0.222	0.229	0.229	0.231	0.233
β_2	0.593	0.598	0.601	0.602	0.599	0.605	0.578	0.581	0.584	0.594
γ_1	-0.033	-0.034	-0.039	-0.029	-0.027	-0.024	-0.043	-0.040	-0.037	-0.023
γ_3	0.001	0.001	0.001	0.002	-0.002	-0.002	0.002	0.001	0.001	0.001
θ_1	0.007	0.007	0.007	0.008	0.010	0.008	0.010	0.010	0.009	0.006
θ_2	0.001*	0.002*	0.002	0.001*	−0.000*	0.003	−0.000*	−0.001*	−0.001*	0.001*
RVD ($X^{(3)}$)										
ω_3	0.365	0.469	0.615	0.656	0.455	0.422	0.525	0.442	0.420	0.334
α_3	0.682	0.502	0.348	0.525	0.020	0.037	0.643	0.430	0.483	0.692
β_3	0.219	0.144	0.144	0.000	0.090	0.066	0.064	0.242	0.245*	0.221
γ_1	−0.065*	−0.076*	-0.141	-0.079	-0.051	-0.032	-0.075	−0.087*	−0.083*	−0.036*
γ_2	0.144*	0.255	0.303	0.339	0.833	0.836	0.155	0.232*	0.143*	0.040*
θ_1	0.007*	0.015*	0.068	0.058	-0.064	0.181	0.042	−0.015*	−0.015*	−0.014*
θ_2	−0.034*	−0.033*	-0.070	-0.021	0.155	-0.206	0.002*	0.048*	0.028*	0.035*

This table shows the news impact curve estimates for the 2-year Treasury note. The model is estimated jointly for depth ($X^{(1)}$), price volatility ($X^{(2)}$) and depth volatility ($X^{(3)}$). Each conditional mean equation is specified as: $\mu_t^{(i)} = \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j \neq i} \gamma_j X_{t-1}^{(j)} + \theta' \mathbf{Z}_{t-1}$, where $\mathbf{Z} = (|Ret|, |Ret| * \mathbf{1}_{Ret<0})$ is the news impact curve and a function of the five-minute return Ret. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level.

Table 4.9: Liquidity and Volatility Dynamics With Trading Volume Effects for 2-Year Treasury Note

	Ask5	Ask4	Ask3	Ask2	Ask1	Bid1	Bid2	Bid3	Bid4	Bid5
DEPTH ($X^{(1)}$)										
ω_1	0.015	0.000	0.007	0.007	0.020	0.019	0.003*	0.009	0.012	0.013
α_1	0.189	0.253	0.287	0.292	0.112	0.116	0.288	0.295	0.330	0.348
β_1	0.797	0.744	0.703	0.698	0.872	0.867	0.705	0.693	0.653	0.635
γ_2	−0.005*	0.003*	−0.001*	0.000*	-0.010	-0.010	0.001*	-0.004	-0.004	-0.003
γ_3	0.002*	0.001*	0.002	0.001	0.003	0.002	0.002	0.003	0.002	0.003
θ	0.001*	−0.001*	0.002	0.002	0.004	0.004	0.002	0.004	0.005	0.004
RVP ($X^{(2)}$)										
ω_2	0.213	0.213	0.221	0.207	0.193	0.100	0.372	0.227	0.221	0.193
α_2	0.207	0.201	0.193	0.191	0.181	0.100	0.277	0.202	0.205	0.211
β_2	0.597	0.601	0.606	0.608	0.636	0.800	0.427	0.588	0.591	0.599
γ_1	-0.037	-0.038	-0.046	−0.035*	−0.029*	−0.022*	-0.067	-0.043	-0.042	-0.023
γ_3	0.000*	0.001	0.001	0.001*	−0.002*	−0.001*	0.001	0.001	0.001	0.002
θ	0.036	0.038	0.041	0.040*	0.033*	0.029*	0.024	0.039	0.036	0.031
RVD ($X^{(3)}$)										
ω_3	0.374	0.483	0.615	0.656	0.430	0.352	0.531	0.439	0.419	0.338
α_3	0.687	0.510	0.341	0.521	0.014	0.024	0.645	0.439	0.486*	0.697
β_3	0.213*	0.140	0.158	0.019	0.030	0.061	0.075	0.252*	0.245*	0.220*
γ_1	−0.065*	−0.077*	−0.151*	-0.080	-0.047	−0.026*	−0.076*	−0.088*	−0.083*	−0.037*
γ_2	0.135*	0.261	0.343	0.426	0.804	0.885	0.225	0.246*	0.140*	0.048*
θ	−0.012*	−0.020*	0.014*	-0.027	0.160	0.167	−0.014*	−0.014*	−0.000*	−0.007*

This table shows estimates for the model with trading volume for the 2-year Treasury note. The model is estimated jointly for depth ($X^{(1)}$), price volatility ($X^{(2)}$) and depth volatility ($X^{(3)}$). Each conditional mean equation is specified as: $\mu_t^{(i)} = \omega_i + \alpha_i X_{t-1}^{(i)} + \beta_i \mu_{t-1}^{(i)} + \sum_{j\neq i} \gamma_j X_{t-1}^{(j)} + \theta' \mathbf{Z}_{t-1}$, where $\mathbf{Z} = qV$, the volume of trading initiated by the opposite side. Estimation is based on five-minute snapshots of BrokerTec limit order book over the period 2006-2010Q2. (*) denotes insignificance at the 5% level.

Figure 4.1: Daily Liquidity and Volatility at First Price Tier of 2-Year Treasury Note

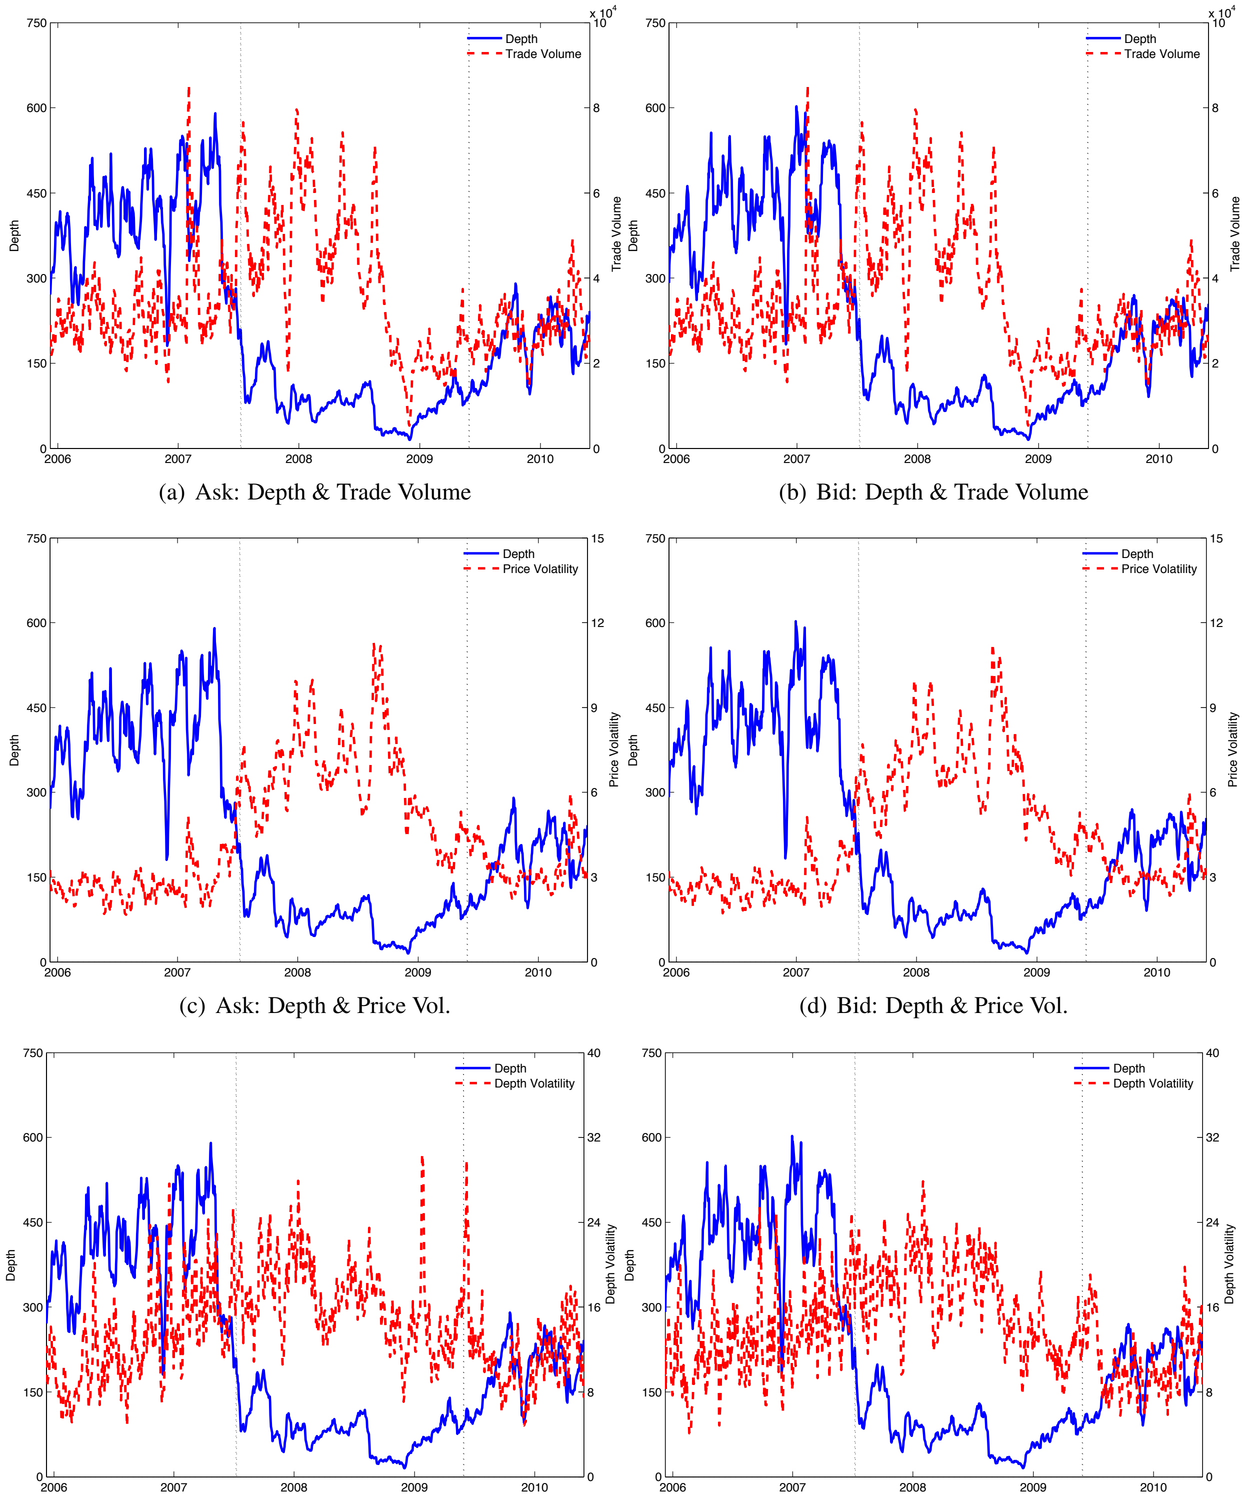

(a) Ask: Depth & Trade Volume

(b) Bid: Depth & Trade Volume

(c) Ask: Depth & Price Vol.

(d) Bid: Depth & Price Vol.

(e) Ask: Depth & Depth Vol.

(f) Bid: Depth & Depth Vol.

This figure shows the 2-year Treasury note's daily average market depth, total trading volume, price volatility and depth volatility at the first price tier, using BrokerTec order book data over the period 2006-2010Q2. Two vertical dotted lines mark the beginning (August 9, 2007) and ending (June 30, 2009) of the crisis. The series are smoothed using a 5-day moving average for better viewing of the trend.

Figure 4.2: Daily Bid-Ask Spread at First Price Tier

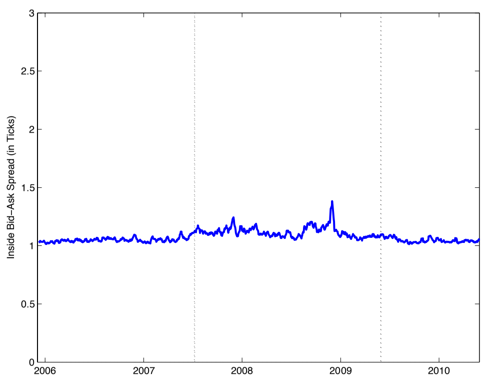

(a) 2-Year Treasury

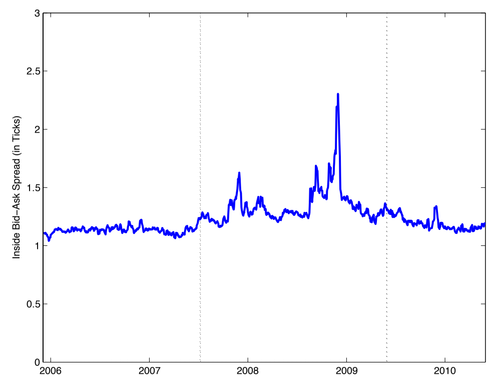

(b) 5-Year Treasury

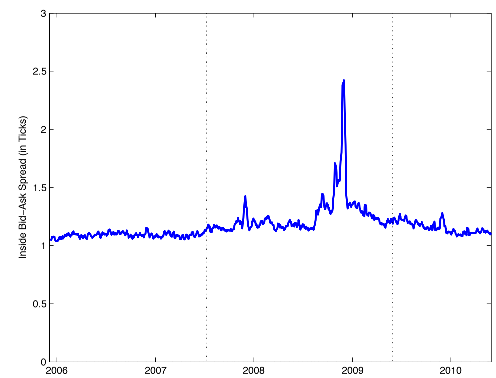

(c) 10-Year Treasury

This figure shows the daily average bid-ask spread at the first price tier for the 2-, 5- and 10-year Treasury notes, using BrokerTec order book data over the period 2006-010Q2. The spread is standardized by the relevant tick size, i.e. 1/128th of a point for the 2- and 5-year notes, and 1/64th of a point for the 10-year note. Two vertical dotted lines mark the beginning (August 9, 2007) and ending (June 30, 2009) of the crisis. The series are smoothed using a 5-day moving average for better viewing of the trend.

Figure 4.3: Intraday Patterns of Liquidity and Volatility at First Price Tier of 2-Year Treasury Note

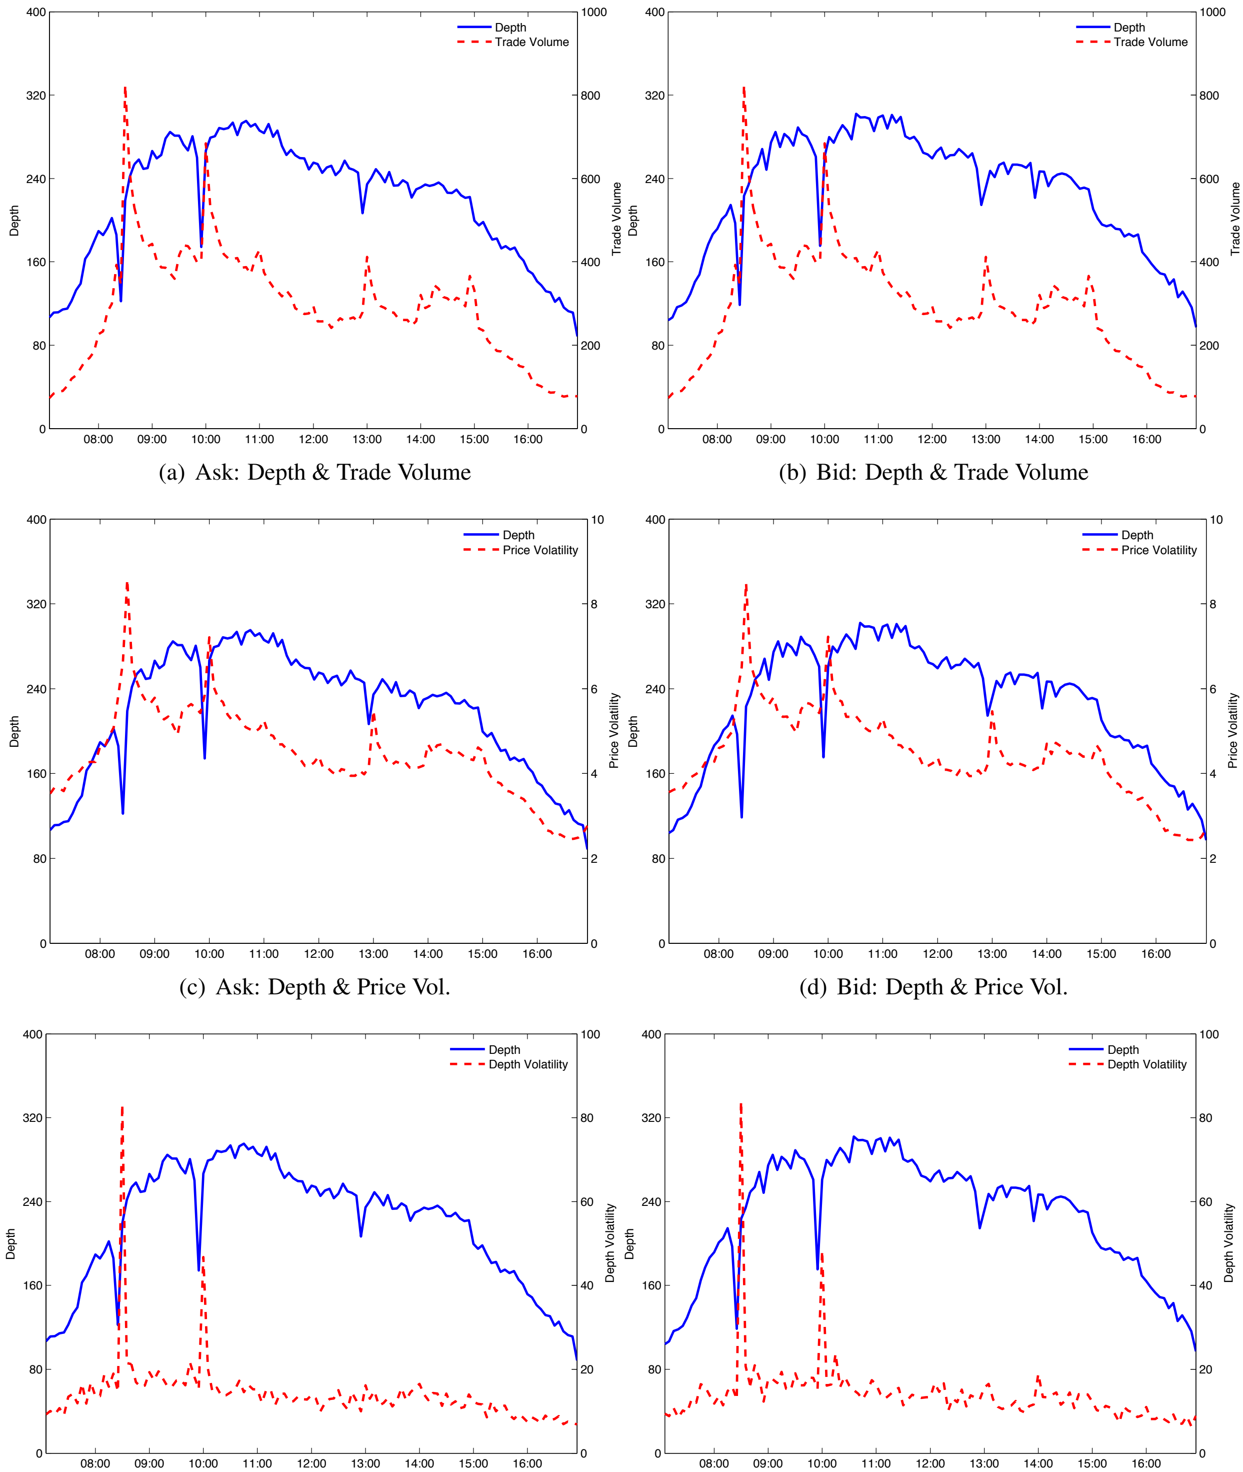

(a) Ask: Depth & Trade Volume

(b) Bid: Depth & Trade Volume

(c) Ask: Depth & Price Vol.

(d) Bid: Depth & Price Vol.

(e) Ask: Depth & Depth Vol.

(f) Bid: Depth & Depth Vol.

This figure shows the 2-year Treasury note's intraday patterns of market depth, price volatility and depth volatility at the first price tier, using BrokerTec order book data over the period 2006-2010Q2.

Figure 4.4: 2-Year Treasury Note's Depth, Trading Volume and Price Volatility on Days With and Without Key Announcements

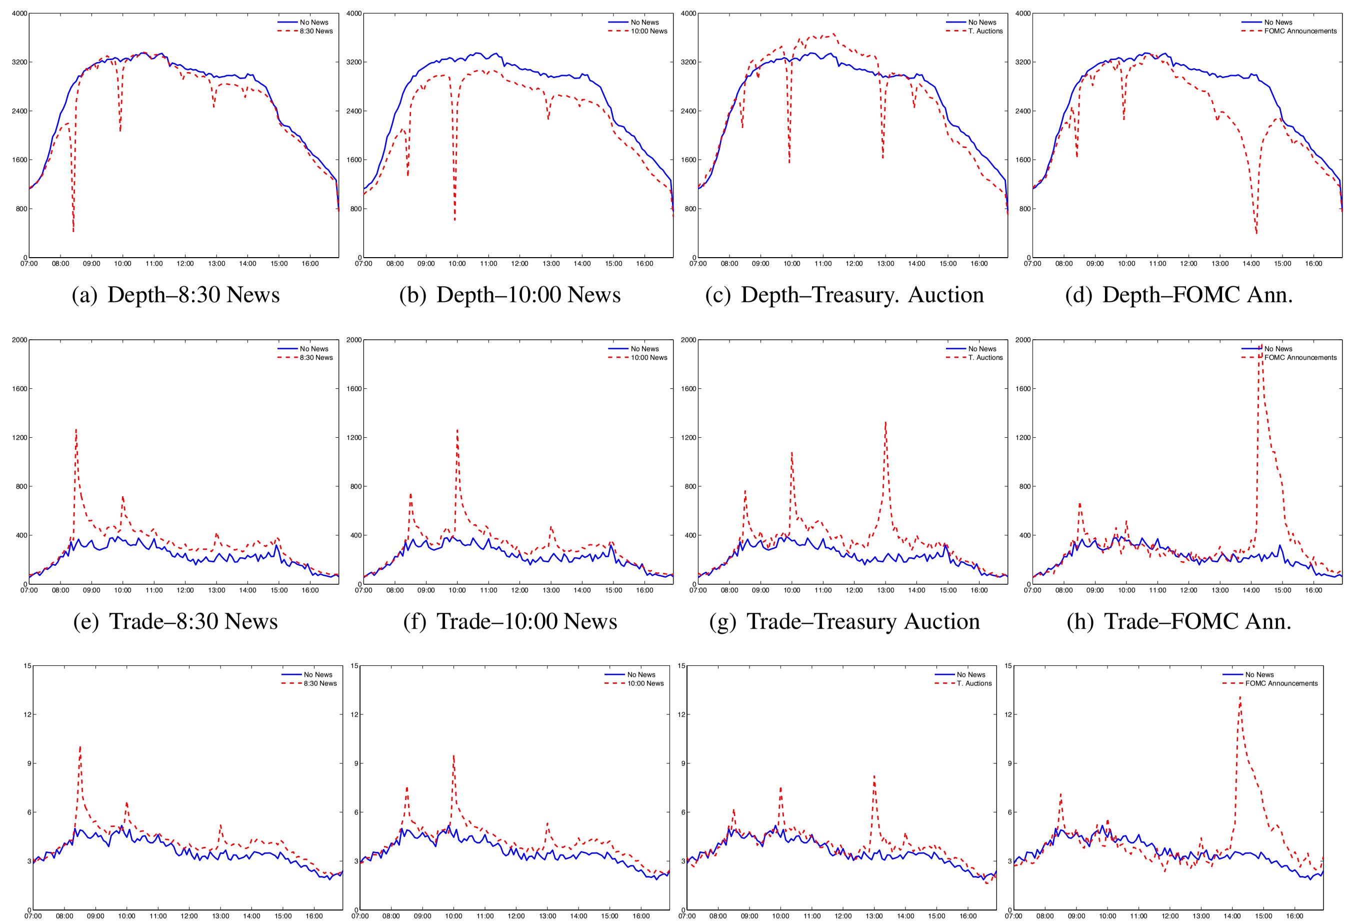

(a) Depth–8:30 News (b) Depth–10:00 News (c) Depth–Treasury. Auction (d) Depth–FOMC Ann.

(e) Trade–8:30 News (f) Trade–10:00 News (g) Trade–Treasury Auction (h) Trade–FOMC Ann.

(i) Price Vol.– 8:30 News (j) Price Vol.–10:00 News (k) Price Vol.–Treasury Auction (l) Price Vol.–FOMC Ann.

Intradaily pattern of depth, trading volume and price volatility on announcement days and non-announcement days, using BrokerTec data over the period 2006-2010Q2. Depth is the total market depth ($ million) at the best five bid and ask prices. Trade is the five-minute total volume of trades ($ million). Price volatility is the five-minute realized volatility of the best bid ask mid-point. Treasury auction results are announced shortly after 13:00 on auction days, and FOMC rate decision announcements are typically made around 14:15.

Figure 4.5: Comparison of Realized Volatility Measures: Price Changes or Yield Changes?

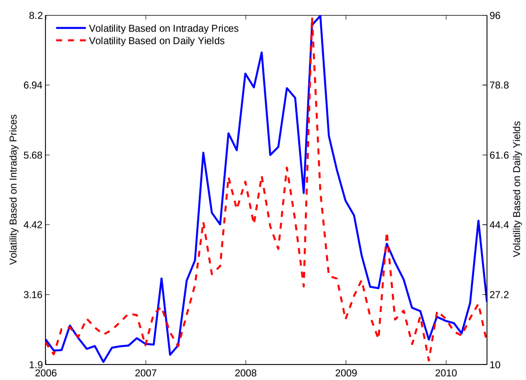

(a) 2-Year Treasury Note

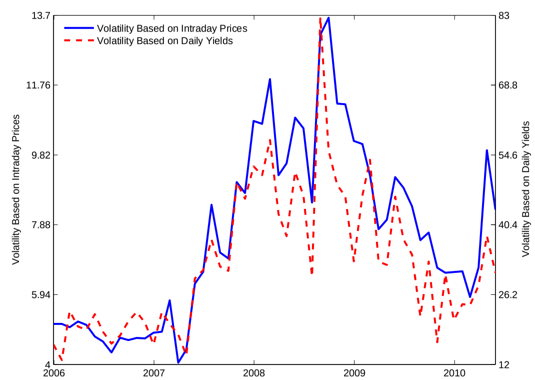

(b) 5-Year Treasury Note

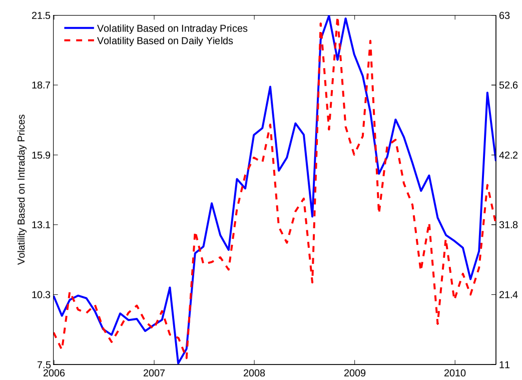

(c) 10-Year Treasury Note

This figure plots realized volatility computed from intraday price changes (‘RV of price’) compared with that computed from daily yield changes (‘RV of yield’) over the period 2006-2010Q2. Specifically, the RV of price is computed for each five-minute interval as the square root of the sum of squared second-by-second price changes. The figure shows the monthly average of this five-minute volatility measure using the left scale where one unit equals one 256th of one percent of par value. The RV of yield for each month is the square root of the realized variance of yield, the latter of which is computed as the sum of squared daily yield changes over the month. It is shown on the right scale where one unit equals one basis point of yield. Daily on-the-run Treasury yields are from Bloomberg.

Figure 4.6: Liquidity and Volatility on FTS and non-FTS days

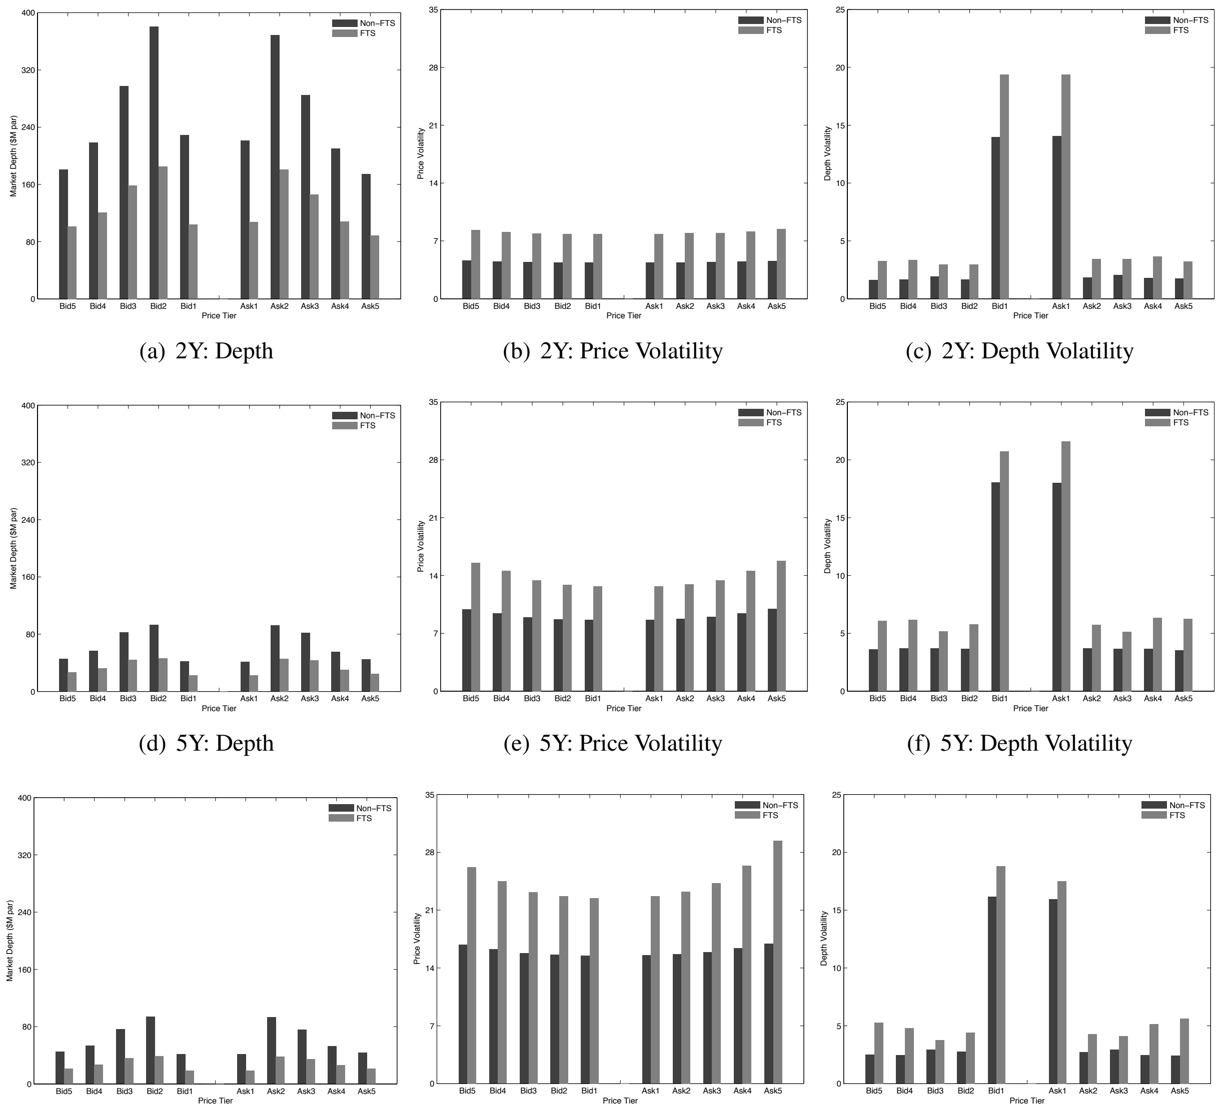

(a) 2Y: Depth (b) 2Y: Price Volatility (c) 2Y: Depth Volatility

(d) 5Y: Depth (e) 5Y: Price Volatility (f) 5Y: Depth Volatility

(g) 10Y: Depth (h) 10Y: Price Volatility (i) 10Y: Depth Volatility

This figure shows average market depth, price volatility and depth volatility at a five-minute frequency, on days with a flight to safety ("FTS" – light colored bars) and days without such an episode ("nonFTS" – dark colored bars), using BrokerTec order book data over the period 2006-2010Q2. Flights are identified by a large positive return on the Treasury note and a large negative return on the S&P500 index, based on a 1.5 standard deviation threshold.

CHAPTER 5

VOLATILITY AND LIQUIDITY SPILLOVERS DURING THE EURO AREA SOVEREIGN DEBT CRISIS

5.1 Introduction

Financial market linkages across countries are important to study, especially in times of crisis. While these linkages facilitate better capital flows across borders, they also make it easier for shocks in one market to spread to others in times of market stress. The recent euro area sovereign debt crisis provides an interesting setting to study the issue of financial market connectedness and cross-country spillovers. Countries in this area share the same currency (the euro), are subject to the same monetary policy making body (the European Central Bank, henceforth ECB), and operate in a highly coordinated fiscal and economic environment. Although the crisis started from imprudent fiscal behavior in Greece, contagion exacerbated the crisis in other countries and spreaded instability across financial markets, as the ECB's Vice President Constancio once remarked (see Constancio (2012)). This source of contagion risk is the main focus of the paper.

Specifically, we examine the extent of shock transmission across euro area sovereign bond markets during the crisis period of 2010-2012 to answer the following questions. Which source of shocks, liquidity or volatility, is more prominent during this period? Which country, or countries, are the main "exporters" of shocks that threaten the stability of the region? The answers to these questions are not obvious, because the crisis did not arise from a single country. Instead, problems of different natures showed up in different member countries at different times and deepened the crisis as it progressed.[1] In this paper, we consider contagion as extreme spillovers of shocks in the sense of Allen and Gale (2000), and measure spillovers in the spirit of Diebold and Yilmaz (2014).

We focus on sovereign bond markets in our study of shock transmission during the eurozone crisis for several reasons. First, government bond markets are among the key financial markets affected by a crisis of

[1]For example, the root cause of the crisis in Greece is fiscal imprudence, whereas Italy experiences a combination of fiscal and political difficulties. For Spain, it is mainly a banking crisis.

sovereign credit risk, because bond yields reflect investors' assessment of a given sovereign's creditworthiness. Secondly, government bonds are an important asset class for investment and hedging purposes. Their liquidity and volatility conditions therefore have important implications for portfolio and risk management. Importantly, government bond markets serve as an important channel of monetary policy transmission. Sovereign bonds are among the most important classes of collateral that can be pledged at the ECB for access to liquidity facilities. Therefore, they can have a direct impact on the level of aggregate funding liquidity and credit intermediation.[2] The ECB's active and direct interventions to address the malfunctioning of the secondary market during the crisis speak volume about the importance of these instruments.

We consider jointly volatility and liquidity spillovers because the volatility-liquidity relation during market crises is an important concern for market participants and policy makers. That is, the potential feedback effect between illiquidity and high volatility can be detrimental. A volatile market increases inventory risk for liquidity providers, and thus can discourage them from supplying liquidity to the market. Conversely, without sufficient liquidity supply to absorb trading demands, price can deviate significantly from the underlying value, resulting in higher price volatility.

From a methodological standpoint, modeling jointly the dynamics of volatility and liquidity in multiple markets permits more insights. Beyond quantifying the extent of shock transmission across countries and identifying systemically important country(ies), we can learn more. In particular, we can delineate the relative magnitude of spillovers through volatility versus liquidity channels and track them over time. Therefore, we can characterize empirically which source of shock tends to propagate more intensely across borders, or how spillovers evolve with market events. In addition, within each bond market, the framework allows us to analyze quantitatively the feedback relationship between volatility and liquidity, and determine if and when one might dominate the other. The vast literature on asset pricing has shown liquidity and volatility to be important pricing factors. Thus, an understanding of how these factors dynamically respond to each other can shed more light on the dynamics of bond yields.

Using daily measures of volatility and liquidity constructed from intraday data on bonds issued by Belgium, France, Germany, Italy, the Netherlands, and Spain, we find that the shock transmission across borders during the crisis appears more liquidity-driven than volatility-driven. Moreover, Italy is the main exporter of both volatility and liquidity shocks to the others, demonstrating its systemic role among euro area

[2] Indeed, Corradin and Rodriguez-Moreno (2014) document a monetary funding premium embedded in euro-denominated sovereign bonds during the 2008-2013 period.

bond markets. By virtue of being the largest bond market and the third largest economy in the eurozone, the dynamics of the Italian bond market can affect other countries in a significant way that none of the other markets can replicate. In addition, we show that this transmission occurs more strongly through the liquidity channel.

The paper contributes to a large literature on spillovers and contagion during the European sovereign debt crisis. Claeys and Vasicek (2014) examine and measure bilateral spillovers using euro area bond yields and find that spillover has increased substantially since 2007. However, the increase in spillover is mostly due to larger sized shocks, and not to a significant change in the transmission mechanism. De Santis (2014) shows that spillover effects from Greece are among the major forces underlying developments in euro area government bond yields in recent years. De Grauwe and Ji (2013) show that the surge in government bond yields during the 2010-2011 crisis is disconnected from underlying fiscal developments, except for Greece. Instead, the surge reflects a negative self-fulfilling market sentiment – one of the possible channels for the propagation of negative shocks.

Mink and de Haan (2013) suggest another channel for contagion, the "wake up call" hypothesis, whereby the Greek crisis prompts the market to scrutinize more closely the viability of other periphery countries. The authors also find that only bailout news relating to the Greek crisis affects European banks' returns, leading them to conclude that the market does not worry about a Greek default per se but relies on bailout news to discern how the European authorities combat the financial crisis. Arghyrou and Kontonikas (2012) document the changing composition of the sources of contagion and examine the role of shifting country-specific market expectations. Bai et al. (2012) study the possible contagion channels during the crisis and conclude that the contagion is predominantly through the fundamental credit risk channel and not through the liquidity risk channel.

Previous papers have measured spillovers among key market variables in the euro area during the crisis. For example, Calice et al. (2013) model the spillover between liquidity and credit risk for countries in the area. Alter and Beyer (2014) measure the spillovers among sovereigns, among banks, from sovereigns to banks, and from banks to sovereigns during the crisis. Based on a generalized impulse response analysis ("GIRF") derived from a VAR model using sovereign and bank CDS spreads, the authors find that spillovers increase prior to key events and policy interventions. In particular, the sovereign-bank spillovers trend upwards during periods of distress, suggesting intensifying feedback loops between euro area banks and sovereigns. Adopting a similar modeling approach based on a factor-augmented VAR model and a GIRF-based forecast error variance decomposition, Claeys and Vasicek (2014) focus on the transmission of shocks across yields of 16 European Union countries.

Surprisingly, despite the numerous studies on the European sovereign debt crisis, not much work is devoted to the issue of volatility and liquidity spillovers across euro area sovereign bond markets. Instead, the literature generally places a greater emphasis on the sovereign CDS markets, the interactions between liquidity and credit risks, or the effects of news or key events on euro area government bond yields to examine spillovers in event study analyses. Even among studies that rely on econometric methods to quantify spillovers as responses to shocks, the focus is mostly on yields (e.g., Claeys and Vasicek (2014)). We argue that, by examining bond market volatility and liquidity – two important determinants of bond yields – we can learn about bond market linkages at a more fundamental level than that by looking at comovement in yields only.

There are a variety of methods proposed in the literature for the testing and measurement of spillovers and contagion. Forbes and Rigobon (2002) focus on cross-market correlations in times of crisis (for a review of work viewing contagion as the effect of contemporaneous movements across countries, see Dungey et al. (2005, 2006)). Some recent papers have relied on copula-based approaches to model the tail dependence among variables of interest (e.g., Rodriguez (2007)), on the basis that structural breaks in tail dependence are symptomatic of contagion. Another branch of the literature measures contagion and spillovers as dynamic responses to shocks. Engle et al. (2012b) model volatilities of eight East Asian countries using a multiplicative error model, and measure volatility spillovers via impulse responses to shocks emanating from each country.

In the linear class of models, Diebold and Yilmaz (2014, 2009) develop spillover measures based on a forecast error variance decomposition of VAR models. Diebold and Yilmaz (2014) propose to estimate the chosen VAR model on a rolling basis to obtain the time-varying dynamics of spillover measures. Parsimony and the simplicity in computing time-varying spillover estimates are appealing features of this VAR-based framework. Parsimony is an important consideration when one has a large number of variables to model. Accordingly, we follow this approach but with an important deviation. That is, instead of assuming a certain variable ordering for the variance decomposition (as in Diebold and Yilmaz (2009)), or using an ordering-invariant GIRF-based approach (as in Diebold and Yilmaz (2014)), we estimate the spillover measures for a large number of permutations of variable orderings and take the average. This way, we avoid making ex-ante assumptions about the causal ordering of variables, while at the same time sidestep the issue of extreme identification assumptions implicit in the GIRF-based approach.[3]

The paper is organized as follows. Section 5.2 describes the data, explains the measurement of bond market volatility and liquidity, and provides background information on the European sovereign debt crisis. Section 5.3 presents the general framework to estimate liquidity and volatility dynamics, from which to construct spillover measures. In Section 5.4, we provide a full-sample analysis of volatility and liquidity spillovers among the six bond markets considered. We also discuss the effects of variables capturing market conditions on liquidity and volatility dynamics. Next, we show in Section 5.5 the multiple layers of spillovers and how they evolve over the 2010-2012 crisis period. Section 5.6 concludes the paper.

5.2 Data

5.2.1 European Sovereign Bond Markets

The main dataset for this analysis is the intraday trade and quote data from MTS. The MTS is the largest electronic interdealer trading platform for European fixed income securities, including government, quasi-government and covered bonds.[4] The sample period is 2010-2012, which covers almost completely the euro area sovereign debt crisis.[5]

While data are available for all euro area countries, we focus on the six largest and most active markets, namely Belgium, France, Germany, Italy, the Netherlands and Spain.[6] These are the major economies in the area, contributing roughly 87% to the total euro area's GDP. As of the end of 2012, over 5,515 billion euros of these sovereign bonds are outstanding, making up 90% of the euro area government bond market size. Furthermore, trading in bonds issued by these six countries, especially Italy, accounts for 94% of total trading

[3]For a critique of the generalized impulse response function in VAR models, see Kim (2012).

[4]Cheung et al. (2005) and Darbha and Dufour (2013) describe in detail the microstructure of the MTS platform.

[5]4/1/2011 is excluded from the sample due to missing quote data. The earliest quote across all bonds on this date occurred at 13:30.

[6]While the origin of the crisis can be traced back to Greece, Ireland and Portugal, there is almost no trading activity in Greek, Irish and Portuguese bonds in the later part of the sample period. We are primarily interested in spillovers that occur through financial market linkages. The inactivity in a given market means that this linkage is not available to transmit shocks to or receive from other countries. As a result, these countries are not included in the analysis.

activity on MTS during the 2010-2012 period. Given their sizes and economic importance, these countries provide a quite complete picture of the euro area for the analysis.[7]

5.2.1.1 Bond Market Volatility

We compute realized variance from intraday quoted prices of benchmark coupon bonds issued by the six sovereigns. Trading prices are not usable for this purpose since most bonds are traded only a few times a day, except for Italian bonds. On the other hand, the quoted prices are binding and are available throughout the trading day, so they present a good alternative to trading prices.

Specifically, the daily realized variance (RV) of a given bond is calculated as:

$$RV_t = \sum_{i=1}^{N} (\triangle Y_{t,i})^2, \tag{5.1}$$

where $\triangle Y_{t,i} = p_{t,i} - p_{t,i-1}$ is the change in the log mid-quote for the interval i on day t, and N is the number of intraday intervals per day. The trading hours on MTS are from 8:00 to 17:30 Central European Time.

In computing RV_t, we need to choose a sampling frequency, i.e., the time interval between $p_{t,i-1}$ and $p_{t,i}$. As discussed in Andersen et al. (2013), a five-minute sampling frequency is often adequate as it offers a reasonable balance between retaining the richness of the data and keeping the microstructure noise component in check. To compensate for the information loss from not using the highest available frequency of data (one-second frequency), we sub-sample the estimator at the one-minute interval. This involves picking a different starting minute during the first five minutes of the day to sample five–minute returns. This gives five five-minute return series. We compute daily RV based on each of the five series and finally take the average of the five computed RV's.[8]

From daily realized variance of individual bonds, we compute the aggregate realized variance for a given bond market by averaging across all bonds issued by the same sovereign. Bonds with less than one year to maturity are excluded. We report several descriptive statistics of realized volatility (square root of realized variance) by country in Panel A of Table 5.1. The mean annualized volatility is between 6-7% for Germany, France and the Netherlands, compared with the 9-10% volatility in Belgium and Italy and 13% in Spain. The

---

[7]For an overview of the euro area government bond markets, and key information about the six selected markets, see Appendix C

[8]This type of estimator has been advocated in Andersen et al. (2011), Ghysels et al. (2010) and Ghysels and Sinko (2011). These papers show that this RV estimator performs on par with, and often better than the more complex RV estimators.

high mean volatility of the latter three countries appears driven by some extreme values. If we look at the median, these countries appear to be much more similar in terms of volatility: all except Spain are in the 6-7% range. Other distributional statistics confirm that the Spanish bond market is the most volatile.

5.2.1.2 *Market Liquidity*

We measure market liquidity for each individual bond on a given day by the time-weighted relative bid-ask spread, expressed as a fraction of the mid-quote (in basis points):

$$Liq = \frac{1}{T_K - T_1} \sum_{k=1}^{K-1} (T_{k+1} - T_k) \frac{(A_k - B_k)}{0.5 \times (A_k + B_k)} \times 10000, \tag{5.2}$$

where T_k is the time of the k^{th} quote update, A_k and B_k are the best ask and bid quotes at time T_k, and K is the total number of quote updates of the day. A higher bid-ask spread indicates a lower level of market liquidity.

A number of studies have found the bid-ask spread to be the most important measure of bond market liquidity (e.g., Fleming (2003); Beber et al. (2009); and Favero et al. (2010)). In addition, since MTS market share is uneven across countries and there are multiple venues in which bonds can be traded, quantity-based measures of liquidity using MTS data might not be comparable across the countries.[9] The bid-ask spread, on the other hand, is a liquidity measure that is based on prices and not quantities. Because we can expect MTS prices to be largely consistent with those prevailing at other trading venues, the bid-ask spread is arguably least subject to the issue of uneven coverage discussed above. To measure the aggregate liquidity level for a given country, we average the daily liquidity measures across all bonds issued by that country, excluding bonds with less than one year to maturity.

Descriptive statistics of liquidity are presented in Panel B of Table 5.1. The German and Dutch bond markets have quite similar level of market liquidity. On average, the bid-ask spread as a fraction of the mid-quote is about 13 bps (mean) or 12 bps (median). Over the sample period, this measure of liquidity never exceeds 60 bps for either market. On the other end, the relative bid-ask spread of Spanish bonds is

[9] It is particularly noted that trading activity in German government bonds is quite low on MTS, because these bonds are also traded actively elsewhere, such as German stock exchanges, international electronic trading platforms and also in the over-the-counter (OTC) market. The total trading volume in German Treasuries on MTS in 2012 is 52.7 billion euros, which is roughly 1% of the total market trading volume of nearly 5,400 billion euros (`http://www.deutsche-finanzagentur.de/en/institutional/secondary-market/`).

nearly six times larger, with a mean and median of 73 bps and 59 bps respectively, and a maximum of 487 bps. Nevertheless, this is still lower than the largest illiquidity reading of the Italian market: 667 bps! The liquidity distribution is highly skewed, especially for Belgium, Italy and Spain.

5.2.2 Market Variables

Beside bond market data, we also collect data capturing market developments during the sample period. Most of the data items are obtained from Datastream, unless otherwise indicated. We use these variables to control for market wide factors that might affect volatility and liquidity across bond markets.

- Daily sovereign credit default swap (CDS) spreads: These are obtained from Bloomberg. We use the spreads for the 5-year maturity, which is widely considered to be the most liquid segment of the sovereign CDS market. Thus, the spreads are least subject to the CDS market's liquidity effects. These CDS spreads provide market-based measures of credit quality of the six countries.

- Daily series of the 5-year iTraxx Financials Senior index: This is a CDS index comprising of the most liquid 25 financial names in Europe (senior subordination). This reflects the aggregate credit risk level of the financial sector. To the extent that banks are a major group of players in the government bond markets, in both market making and trading capacities, their viability is expected to have an effect on bond market liquidity and volatility.

- Daily 3-month Euribor and EOIS rates: The difference between these rates reflects the cost of unsecured borrowing for banks over the risk free rate. It captures the willingness of banks to lend to each other and is thus often used as a measure of funding condition.

- Daily yields on AAA-rated and BAA-rate corporate bond indices: The difference between these two series provides a market-based assessment of the aggregate default risk in the economy.

- Daily series for the VSTOXX: This is the implied volatility of options on STOXX, a major stock market index in the eurozone. This series serves as a proxy for the aggregate risk aversion specific to the eurozone.

- Daily series for the VIX: This is the implied volatility of options on the S&P 500 index, obtained from the Chicago Board Options Exchange website. This has been widely used as a proxy for global systemic risk factor. Given that the US market closes after the European market, the VIX series is lagged by one day when it is used as an explanatory variable.

5.2.3 The European Sovereign Debt Crisis

This section describes major developments during the euro area sovereign debt crisis in order to provide a background essential for the subsequent analysis on spillovers in the paper. Figure 5.1 offers an overview of key market developments. The top panel plots the 10-year benchmark yield spreads of five countries relative to Germany's 10-year benchmark yield. The middle panel shows the evolution of sovereign credit risk of the six countries in consideration. The bottom panel documents the ECB's interventions in the bond market.[10] We identify and mark on the plots five event dates around which large spikes in yields and CDS premia occurred.

The first corresponds to May 10, 2010 (vertical line (1)), which marks the beginning of the Greek crisis. On this date, the ECB announced the SMP in an effort to restore bond market's stability. The second system-wide jump in yields occured in August 2011, when the market grew increasingly concerned about Italy's ability to service its public debts and complicated political environment. On August 8, 2011, the ECB reactivated its SMP (vertical line (2)). During that week, the ECB purchased the largest amount in the history of the program.

The political turmoil in Italy continued into the fourth quarter of 2011, ultimately resulting in a change of government on November 26, 2011 (vertical line (3)). As the markets reacted to this news, bond prices were considerably depressed and yields reached a new height. Another episode of yield jumps occurred when Standard and Poor's downgraded the credit ratings of France, Italy and Spain on January 13, 2012 (vertical line (4)). Shortly after this date, Fitch also downgraded the credit ratings of Belgium, Italy and Spain. From this point, the yields on French, Belgian and Dutch bonds declined gradually toward the end of 2012. On the contrary, the Italian and Spanish yields continued to experience one more episode of severe stress in June and July of 2012, when Spain struggled with a banking crisis. The euro group granted financial assistance to

[10]During the European sovereign debt crisis, the ECB implements a number of additional measures complementing its regular liquidity operations. These complementary measures include outright interventions in public and private debt markets via the so-called Securities Market Programme ("SMP") and the Covered Bonds Purchase Programme. Of direct relevance to the sovereign bond markets is the SMP. The only publicly available source of information on the SMP operations is the ECB's weekly financial statements, which state the total amount of purchases over a given week (without further details on the timing and the targeted country/maturity of the operations). We manually compile a weekly series of SMP amounts from these statements. We thank Olivier Vergote at the ECB for providing helpful instructions in compiling this series.

Spain's banking sector on July 20, 2012. However, only until the ECB President's statement on July 26, 2012 of the ECB's commitment to do "whatever it takes to preserve the euro" that the pressure on the Italian and Spanish yields seemed to ease (vertical line (5)).

As shown in the last panel of Figure 5.1, the ECB's outright operations in the secondary markets for sovereign bonds of the periphery countries concentrated in May-June 2010 and August-December 2011, with the two largest purchases undertaken at the peak of the Greek crisis and Italian crisis. The program, which concluded in January 2012, acquired about 218 billion euros. Only until early 2013 that the ECB published the country composition of its SMP bond holdings. Italian bonds account for roughly 47%, followed by Spain (20%), Greece (16%), Portugal (11%) and Ireland (6%). Based on proprietary data on individual SMP purchases, Ghysels et al. (2013) show that these purchases are successful at reducing bond yield volatility in participating countries.

Examining the plots over the whole sample period, we can observe clearly that the pressures on yields and CDS spreads in May 2010 around the Greek crisis appear much less serious than those prevailing around the crises in Italy and Spain in late 2011. These two countries also account for the largest shares of the ECB's bond market intervention program.

Bond market volatilities and liquidity levels closely follow the above events. Figure 5.2 shows the euro area bond market volatility over the sample period from 2010 to 2012. The figure reveals that the German, Dutch and French bond markets are much less volatile than the markets for Italian, Belgian and Spanish bonds. The scale used for plotting volatilities of the countries in the latter group is three times larger than that for the former. There was a clear jump in volatility across markets at the break out of the Greek crisis on May 10, 2010. The markets seemed to calm down shortly afterward, with a brief increase in volatility in Belgium, Italy and Spain in late 2010, around the time of Ireland's bailout package. Nevertheless, the markets remained rather stable until the summer of 2011, when concerns over the spreading of the crisis to Italy and Spain became apparent. The change of government in Italy in November 2011 led to significant volatility increases across the markets. Most notably, Italy and Spain saw volatilities exceeding the 100% p.a. level. Volatility in France and especially Belgium also increased significantly around this time, reaching 45% and over 70% respectively. After these peaks, the markets still experienced one more brief round of heightened volatility in the middle of 2012 as Spain coped with its banking crisis.

The time-series evolution of liquidity follows that of volatility closely. As shown in Figure 5.3, there are clear episodes of extreme illiquidity for the most volatile group of countries (Belgium, Italy and Spain). The first of such episodes occurred at the breakout of the Greek crisis in May 2010, when the bid-ask spread in the Spanish bond market reached nearly 400 basis points for the first time, more than five times larger than the typical level of 73 bps. The second wave of illiquidity occurred around the time of the crisis in Ireland in November 2010. The next and also the most serious spell of market-wide illiquidity came in the summer of 2011 and remained until the end of that year. Notably, the bid-ask spread on Italian bonds shot to an unprecedented level of nearly 700 bps in early December 2011, while that for Spanish, Belgian and French bonds reached nearly 500, 200 and 120 bps respectively. The Spanish bond market experienced another major spike in illiquidity in July 2012, before seeing improvement in liquidity toward the end of 2012. Interestingly, other countries' liquidity did not seem to be significantly affected this time.

It is easy to see that the liquidity patterns correspond closely to episodes of heightened volatility and periods of distress in Greece, Ireland, Italy, and Spain discussed earlier. Formally, the correlation between liquidity and volatility in a given bond market is high, ranging between 0.63 for the Spanish market to 0.79 for the Italian market. This high degree of correlation indicates that modeling only volatilities, or only liquidity measures, will be incomplete. Considering them jointly enables us to disentangle which of the two is the more important source of variation in the market. The next section describes the model and how the various spillover indicators are computed.

5.3 Model

To measure spillovers and assess the degree of connectedness among the major euro area bond markets, we adopt the VAR-based approach by Diebold and Yilmaz (2014). The VAR model is a standard and parsimonious framework to model a large network of interconnected variables. Thus, it is a natural choice for the analysis in this paper, in which we seek to model the dynamic interactions among liquidity and volatility of six bond markets. The availability of high frequency data allows us to directly measure volatility at the daily frequency. This is a great advantage, allowing us to avoid potential mis-specification concerns if volatility were to be estimated parametrically from daily returns. Furthermore, we can model realized volatility jointly with liquidity in one convenient and unified framework. In this section, we describe details of the VAR model used, some key modeling considerations, and how spillovers are measured.

5.3.1 *Modeling Dynamic Interdependencies*

Consider a vector of n variables of interests, denoted by Y_t, that follows a linear vector autoregression:

$$Y_t = C + \sum_{j=1}^{p} A_j Y_{t-j} + \epsilon_t \tag{5.3}$$

In particular, to fix ideas, we take Y_t to be a vector of daily logged realized variances and daily logged relative bid-ask spreads of six bond markets in the euro area, namely Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES). Our objective is two-fold: 1) to disentangle the mutual feedback effects between volatility and liquidity within each bond market ("intra-country spillover") and 2) quantify the effect of spillovers from one market to other markets ("inter-country spillover"). This requires measuring the effects of shocks of one variable on another variable at various horizons, which is commonly achieved by a forecast error variance decomposition.

From the estimated VAR model, the h-step forecast error is:

$$Y_{t+h} - \hat{Y}_{t+h|t} = \epsilon_{t+h} + \Psi_1 \epsilon_{t+h-1} + \Psi_2 \epsilon_{t+h-2} + ... + \Psi_{h-1} \epsilon_{t+1}, \tag{5.4}$$

where Ψ_j is the j-lag MA coefficient matrix in the MA(∞) representation of the VAR model:

$$Y_t = C + \epsilon_t + \Psi_1 \epsilon_{t-1} + \Psi_2 \epsilon_{t-2} + ...$$

The variance of the h-step forecast error in equation (5.4) is:

$$V_h = \Omega + \Psi_1 \Omega \Psi_1' + \Psi_2 \Omega \Psi_2' + ... + \Psi_{h-1} \Omega \Psi_{h-1}', \tag{5.5}$$

where $\Omega = E(\epsilon_t \epsilon_t')$. Orthogonalization of shocks follows Sims (1980)' recursive method, assuming that the j-shock has contemporaneous effects on only variables placed after it in the ordering.[11] That is, $\epsilon_t = Au_t$, where A is a lower triangular matrix capturing contemporaneous effects and u_t is a vector of orthogonalized

[11] Diebold and Yilmaz (2014) use the generalized forecast error variance decomposition to produce spillover measures that are not sensitive to the variable ordering. The limitations of this approach are discussed in Kim (2012). We instead opt for using the standard recursive orthogonalization procedure, but implement an additional step to robustify our spillover measures, which will be discussed in greater details in the subsequent subsection.

residuals. Thus, $\Omega = A\Sigma A'$, where Σ is the variance of u_t. Performing the Cholesky factorization on the matrix Ω, i.e., $\Omega = PP'$ (and thus $P = A\Sigma^{1/2}$), we can decompose the forecast error variance V_h in equation (5.5) into contributions by each of the orthogonalized residuals $V_h^{(j)}$, for $j = 1, ..., n$:

$$V_h = \sum_{j=1}^{n} V_h^{(j)},$$

where

$$V_h^{(j)} = p_j p_j' + \Psi_1 p_j p_j' \Psi_1' + \Psi_2 p_j p_j' \Psi_2' + ... + \Psi_{h-1} p_j p_j' \Psi_{h-1}',$$

and p_j is the j^{th} column of matrix P. The percentage contribution of the j-variable to the forecast error variance of the i-variable is thus:

$$C_{j\to i} = 100 \times \frac{V_h^{(j)}(i,i)}{V_h(i,i)}.$$

The above calculation results in a $n \times n$ matrix of bilateral spillovers ($n = 12$), in which the (i, j) element is $C_{j\to i}$. Intuitively, $C_{j\to i}$ measures the extent to which the j^{th} shock drives the variation in the i^{th} variable at a given horizon h, i.e., the extent of shock transmission from j to i at the horizon h.

5.3.2 Effects of General Market Conditions

To account for the possibility that developments in the macroeconomic environment might explain some of the comovement of volatility and liquidity across the sovereign bond markets in consideration, we augment the VAR model with the following control variables.

First is the common trend in credit quality of these sovereigns. The commonly used proxy for credit quality of a country is the CDS premium on its sovereign debts. A principal component analysis on the CDS premia of the six countries reveals that the first principle component ($SovPC1$) explains over 91% of the total variation. We use $PC1$ as a control variable in the VAR model, instead of the six individual CDS series, realizing that it has captures sufficiently the common trend in sovereign credit risk that might affect the bond markets' liquidity and volatility. The model is also leaner this way.

Next, we include changes in the iTraxx-Financial index to proxy for funding condition ($Fund$). iTraxx-Financial index is an index of CDS on 25 financial firms in Europe. The financial sector plays a key role in intermediating and supplying credit, so their viability has a first order effect on the funding condition in the euro area. Another proxy for funding condition is the cost of credit (CC), as measured by the difference between the 3-month Euribor and the EOIS rate.

To control for aggregate risk factors, we include the following variables. First, we use default spread ($DefSpr$), which is the yield difference between AAA-corporate and BAA-corporate bond indices, to capture the aggregate credit risk level in the economy. Secondly, the implied volatility of options on the STOXX index, $VSTOXX$, serves as a proxy for regional risk aversion. Lastly, we include the VIX to control for movements in global risk factor beyond those pertaining to the euro area.

The model includes both contemporaneous and one-lag effects of these variables to account for possible delay in the bond markets' reaction to these time-varying market conditions. Thus, the model estimated is:

$$Y_t = C + \sum_{j=1}^{p} A_j Y_{t-j} + B_0 X_t + B_1 X_{t-1} + \epsilon_t, \tag{5.6}$$

where $X_t = [SovPC1, Fund, CC, DefSpr, VSTOXX, VIX]'$.

5.3.3 Spillover Measures

After estimating the model and performing the variance decomposition as outlined earlier, we obtain a 12×12 bilateral spillover matrix. From this table, we classify spillovers into two major categories: 1) intra-country spillovers, and 2) inter-country spillovers. To facilitate a clearer description of spillover measures in this section, we introduce a different set of indexing notations. Let c be the country index, with $c \in \{$DE, FR, NL, BE, IT, ES$\}$, L and V denote liquidity and volatility respectively, and an arrow ($\rightarrow$) indicate the direction of variance contribution. We now discuss each spillover category in detail.

Intra-country spillovers reflect the extent to which liquidity shocks are transmitted to volatility, and vice versa, within the same bond market. These intra-country dynamics reveal which source of shocks tends to dominate a given market, and are measured by the elements $C_{c,L \rightarrow c,V}$ and $C_{c,V \rightarrow c,L}$ in the spillover matrix.

On the other hand, inter-country spillovers measure the extent of shock transmission across borders. To assess the systemic importance of each country, we compute the following directional spillover measures capturing the extent of shocks sent and received by each country in the same fashion as Diebold and Yilmaz (2014):

1. Spillovers Received:

$$IN_{c,L} = \sum_{c^* \neq c} \left(C_{c^*,L \to c,L} + C_{c^*,V \to c,L} \right), \tag{5.7}$$

$$IN_{c,V} = \sum_{c^* \neq c} \left(C_{c^*,L \to c,V} + C_{c^*,V \to c,V} \right). \tag{5.8}$$

2. Spillovers Sent:

$$OUT_{c,L} = \sum_{c^* \neq c} \left(C_{c,L \to c^*,L} + C_{c,L \to c^*,V} \right), \tag{5.9}$$

$$OUT_{c,V} = \sum_{c^* \neq c} \left(C_{c,V \to c^*,L} + C_{c,V \to c^*,V} \right). \tag{5.10}$$

3. Net Spillovers Sent:

$$NET_{c,L} = OUT_{c,L} - IN_{c,L}, \tag{5.11}$$

$$NET_{c,V} = OUT_{c,V} - IN_{c,V}. \tag{5.12}$$

Notice that each of the above directional spillover measures is the sum of liquidity-related and volatility-related variance shares. The separation of liquidity variables from volatility variables permits a more granular view about which source of shocks is the main driver of a country's systemic importance.

5.3.4 Measuring Time-Varying Spillovers

Based on the above framework, we can study the time-varying dynamics of spillovers among these countries by estimating the VAR model on a rolling basis and computing the corresponding spillover measures, as suggested in Diebold and Yilmaz (2014). For implementation, we need to choose the length of each rolling window and the rolling frequency.

For the window size, we follow Diebold and Yilmaz (2014) and use a window of 100 days. This choice balances between having reliable model estimates (preferring a longer window) and being able to quickly pick up changes in spillover (preferring a larger number of shorter windows). Different from Diebold and Yilmaz (2014), we choose to roll the estimation at the weekly frequency, from one Friday to the next (or Thursday if Friday is a holiday), instead of the daily frequency. This results in a weekly series of spillover estimates, each of which is based on market dynamics over the 100 days up to the end of the corresponding week. For the 2010-2012 period, the first available spillover estimate corresponds to the week ending on May 28, 2010.

To keep the time-series analysis of spillovers tractable, we compute the following aggregate spillover indices and examine their behavior over time. Note that the time reference (at the weekly frequency) for these indices is suppressed.

1. Total spillover index. This measures the average contribution of shocks from all other variables, or the average extent to which a given variable is driven by others:

$$TotalS = \frac{1}{12}\sum_{j=1}^{12}\left(\sum_{i\neq j} C_{j\rightarrow i}\right). \tag{5.13}$$

2. Intra-country spillover. This is a component of the total spillover index that captures only the spillover between liquidity and volatility of the same bond market. The intra-country spillover is calculated as:

$$IntraS = \frac{1}{12}\sum_{c}\left(C_{c,L\rightarrow c,V} + C_{c,V\rightarrow c,L}\right). \tag{5.14}$$

3. Inter-country spillover. This is the other component of the total spillover index that shows the extent of spillovers received from external countries. It can be calculated as:

$$InterS = TotalS_t - IntraS_t. \tag{5.15}$$

The inter-country spillover index can be further decomposed into four sub-indices based on the type of shocks received from and sent to the other countries: 1) volatility to volatility spillover ("$V2V$"), 2) volatility to liquidity spillover ("$V2L$"), 3) liquidity to liquidity spillover ("$L2L$"), and 4) liquidity to volatility spillover ("$L2V$"):

$$V2V = \frac{1}{6}\sum_{c}\left(\sum_{c^*\neq c} C_{c^*,V\to c,V}\right), \tag{5.16}$$

$$V2L = \frac{1}{6}\sum_{c}\left(\sum_{c^*\neq c} C_{c^*,V\to c,L}\right), \tag{5.17}$$

$$L2L = \frac{1}{6}\sum_{c}\left(\sum_{c^*\neq c} C_{c^*,L\to c,L}\right), \tag{5.18}$$

$$L2V = \frac{1}{6}\sum_{c}\left(\sum_{c^*\neq c} C_{c^*,L\to c,V}\right). \tag{5.19}$$

5.3.5 Modeling Considerations

Lag order The selection of the lag length p is based on the AIC and BIC, which suggest 2 and 1 respectively, using the full data sample. We choose 1 lag to reduce the risk of overfitting, especially when it comes to estimating the model on a shorter rolling window.

Forecasting horizon for forecast error variance decomposition We choose a forecasting horizon of 10 days (or two weeks) ahead. A similar choice is made by Diebold and Yilmaz (2014) who use a 12-day forecasting horizon in their study of volatility spillovers across major financial stocks. Alter and Beyer (2014) use a five-day horizon for their impulse response analysis in studying sovereign and bank CDS spillovers. For risk management, too long a forecasting period can reduce the forecast's usefulness because the forecast error might be too large. On the other hand, if the forecasting horizon is too short, we might not be able to capture the long run effect of a shock.

Variable Ordering The paper computes variance contributions of shocks orthogonalized by Sims (1980)' recursive method, which depend on how variables are ordered in the system. Instead of assuming a specific ordering and letting the results depend on the chosen ordering, we robustify the variance contributions by averaging across multiple permutations of variable ordering. This is similar to the approach suggested by Klößner and Wagner (2012) to overcome the ordering-dependent problem. A common alternative to the ordering issue is to use Pesaran and Shin (1998)'s generalized impulse response approach, as employed in Diebold and Yilmaz (2014). However, as discussed in Kim (2012), this approach relies on extreme identifying assumptions that might provide misleading inferences.

Given that the model has 12 variables, the number of all possible variable orderings is too large (12!), making the variance decomposition of all orderings not computationally practical. This necessitates the choice of a subset of orderings to work with. A logical start is to vary the ordering by country. For each country ordering, the liquidity and volatility variables of a given country are placed together. With six countries, there are 720 country orderings in total. We then vary the ordering of liquidity and volatility within each country. Specifically, we first perform the variance decomposition with liquidity preceding volatility for each country. We then repeat the exercise, but reverse the within-country ordering for volatility to appear first. This results in 1,440 orderings over which the Cholesky decomposition is done. The resulting 12×12 spillover matrix contains the variance contributions averaged across 1,440 orderings.

5.4 Full Sample Analysis of Spillovers

This section provides an analysis of the spillover effects among the six major euro area bond markets for the full sample period of 2010-2012. However, to facilitate a better understanding of the dynamics among variables in the system, we start first with a discussion of the model estimates and the effects of the general market conditions.

5.4.1 Model Estimates

5.4.1.1 Feedback Effects

Estimates of the feedback parameters are provided in Table 5.3, where each column shows the effects of lagged variables on the variable indicated in the column heading. The diagonal of the table contains all autoregressive coefficients. They are significant and positive, providing evidence of liquidity and volatility clustering. Liquidity is particularly persistent in four markets, i.e., NL, BE, IT, and ES, with a coefficient estimate ranging between 0.78 and 0.87, whereas DE and FR show a moderate persistence level. Volatility persistence is lower, mostly below the 0.5 level.

Off-diagonal estimates capture the feedback mechanisms among variables in the system. It is immediately apparent from the table that most of the feedback effects occur due to liquidity. This can be seen from the top half of the table where many of the feedback coefficients of liquidity are significant. On the contrary, in the bottom half of the table where we present the feedback effects of volatility, most coefficients are insignificant.

5.4.2 *Effects of Market Conditions*

The coefficient estimates for explanatory variables capturing market conditions are reported in Table 5.4. We observe that two variables are particularly important. First is $VSTOXX$, the regional risk factor. The coefficient estimate for this factor is positive and significant in all liquidity equations, indicating that an increase in the euro area's aggregate risk level is associated with an immediate increase in illiquidity (or decrease in liquidity) in the bond markets. Interestingly, once we control for this regional risk factor, the global risk factor – proxied by the $USVIX$ – is mostly insignificant, suggesting that liquidity reacts mainly to euro area specific developments. Another noteworthy observation is that $VSTOXX$ does not affect volatility as strongly as it does to liquidity: significantly positive coefficients are seen for Belgium and Spain only. For Germany and the Netherlands, this effect comes with a lag. The results seem to suggest that the eurozone stock market volatility transmits to bond markets rather through the liquidity channel than the volatility channel.

The second important determinant is $iTraxxFin$, a proxy for the aggregate credit risk of the financial sector. A deterioration in the financial sector's viability makes the bond markets more volatile. This is an intuitive result, because the financial sector is a key group of players in the bond markets with respect to market making, trading, and investing activities. Surprisingly, however, it is not a major determinant of liquidity, except in the Italian and Spanish markets, where the financial sector's decreased credit worthiness translates to a decrease in liquidity. It seems that liquidity is mainly driven by the eurozone stock market volatility as discussed above, and to a lesser extent, by the funding cost ($FinCost$) and aggregate credit risk level ($DefSpr$).

Perhaps the most surprising result is that movements in sovereign credit quality have little effects on bond market liquidity and volatility. It appears that sovereign credit risk has been reflected in other aggregate market indicators discussed above and accordingly has low incremental explanatory power. The only significant effects are observed for the German bond market, where an increase in the overall sovereign credit risk of major euro area sovereigns is associated with an increase in liquidity and decrease in volatility. This evidence resonates well with a flight-to-safety hypothesis. Germany maintains its AAA credit rating through the euro area sovereign debt crisis and has traditionally served a safe haven role in the area. Accordingly, when the credit quality of other euro area countries deteriorates, it is reasonable to observe greater liquidity coming to this market, and a lower level of uncertainty with regards to German-issued securities. Nevertheless, this explanation does not help explain the puzzling negative coefficient of sovereign credit risk on Spanish bond market volatility.

Overall, the evidence presented in this subsection highlights an important finding that most aggregate market conditions affect the bond markets through the liquidity channel. Moreover, these effects generally occur contemporaneously rather than with a lag, indicating that the bond markets incorporate new developments quickly.

5.4.2.1 Correlation of Liquidity and Volatility Innovations

In Table 5.5, we show the contemporaneous correlation matrix of the model residuals (ϵ_t). The model residuals exhibit a high degree of correlation. Most correlation coefficients are in the 0.4-0.7 range. Although the residual correlation coefficients are lower than those observed among the original variables (as shown in Table 5.2), the results show a significant degree of commonality of volatility and liquidity innovations among the bond markets. This is consistent with Chordia et al. (2005) who document that liquidity and volatility shocks in bond and stock markets are driven to a large extent by common factors.

5.4.3 Volatility and Liquidity Spillovers

5.4.3.1 Bilateral Spillovers

The full spillover matrix is presented in Table 5.6. The (i, j) element of the matrix indicates the variance contribution by the variable in column j to the variable in row i. We observe several strong cross-country linkages. First are the volatility linkages among Germany, the Netherlands and France, ranging between 9% and 21%. In particular, German bond volatility explains 21% variation in Dutch bond volatility and receives 18% variance contribution in return. Secondly, liquidity of the Italian and Spanish bond markets are also closely connected, explaining respectively 16% and 14% of the variation in the other. Importantly, liquidity shocks emanating from the Italian market are the second most important source of liquidity variation in the other markets beside their own liquidity shocks.

Each number on the diagonal reflects the extent to which a variable is driven by its own shocks, and is evidently the largest source of variation for each of the variables considered. Own shocks account for 31-43% of volatility variation and 25-40% of liquidity variation. The remaining portion of variation is attributable to other variables and is split into own country effects (referred to as “intra-country spillovers”) and inter-country spillovers.

5.4.3.2 Intra-Country Spillovers

Intra-country spillovers reflect the strength of the inter-dependency between volatility and liquidity in a given bond market. These measures are reported in Table 5.7. For Germany, the Netherlands and France, liquidity innovations contribute little to volatility dynamics, and vice versa: all variance shares are below 5.2%, and are relatively balanced between liquidity and volatility. On the other hand, Belgium, Italy and Spain show a markedly greater degree of liquidity-volatility connectedness, ranging between 14% and 20% for liquidity-induced variation and between 8% and 10% for volatility-induced variation. For example, liquidity variation in the Italian market contributes nearly 20% to its volatility variation and receives 10% in return. Importantly, liquidity exerts greater influence on volatility than vice versa, as shown by the net liquidity spillovers between 7% to 10%, indicating that liquidity is the more important source of variation within each of these markets.

5.4.3.3 Inter-Country Spillovers

We quantify cross-country spillovers in Table 5.8. The first panel, labeled “Received From Others”, shows the proportions of variation receive from outside countries. These inward spillovers are split into the sources of shocks, namely “Vol” (the total variance share by external volatility shocks) and “Liq” (the total variance share by external liquidity shocks). Consider first volatility variables. External volatility shocks contribute between 21% to 40% to variation of volatility in a given market, while external liquidity shocks account for another 16-27%. In total, shocks from the outside generally explain at least half of volatility fluctuation in a market. Nevertheless, as shown in the second panel labeled “Sent to Others”, each market also sends out a considerable amount of volatility shocks to others, affecting both their volatility and liquidity variation. In particular, German bond volatility plays a major role in the other markets’ volatilities: its variance contributions sum up to 47%, exceeding the volatility spillovers it receives from others. This makes Germany a net exporter of volatility shocks. Another net exporter of volatility is Italy. The remaining four countries, on the other hand, are net recipients of the volatility transmission.

With respect to liquidity, the magnitude of shocks sent and received across borders is larger that what is observed for volatility shocks. For example, inward spillovers from external liquidity shocks are in the

37-51% range (compared to the 21-40% range for inward volatility spillovers), and outward spillovers of shocks to foreign liquidity are between 24% and 72% (compared to the 17-47% range for outward volatility spillovers). The outward liquidity spillover range would have been 24-48% without Italy. Italy's liquidity fluctuation is the major source of liquidity variation in others, accounting for an aggregate variance share of 72%. On balance, the three crisis-related countries, namely Belgium, Italy and Spain, are sending more liquidity shocks than received, making them net exporters of liquidity shocks. Germany, on the other hand, is the largest net recipient of liquidity shocks from its neighbors. This is not withstanding the observation that liquidity shocks also affect foreign markets' volatility, although this cross-country liquidity-volatility connection tends to be weaker and balanced with the corresponding reciprocal spillover effects.

Overall, the results just presented above clearly demonstrate Italy's systemic role among sovereign bond markets in the euro area. Liquidity shocks from Italy overwhelmingly propagate the major euro area bond markets, in addition to its net export of volatility spillovers. Furthermore, liquidity is generally the more dominant source of spillovers across borders. The strong cross-country liquidity linkages are consistent with evidence of liquidity commonality documented in Chordia et al. (2005).

Our own findings with regard to the effects of market conditions and those documented in Pelizzon et al. (2013) provide some helpful insights to explain the important role of liquidity. These authors find that market liquidity in Italy is significantly driven by external market conditions, including sovereign credit risk. It appears that these factors first drive liquidity fluctuations in the Italian bond market, which in turn affect volatility and propagate to other bond markets through both volatility and liquidity channels, with the latter channel being more prominent.

5.5 The Dynamics of Spillovers

The previous section focuses on identifying the major forces underlying the inter-connection of euro area bond markets over the full sample period. The result shows that, Italy is the major net sender of both liquidity and volatility shocks to the other five countries. To complement the full sample analysis, we examine how spillover effects evolve over time through the crisis. In particular, does the full-sample spillover effects previously documented prevail throughout the whole period, or just show up significantly at some brief moments in time?

We start with a discussion of the trend in the total spillover index, which measures the overall degree of inter-dependency among the bond markets considered. This discussion also describes the two components making up the total spillover index, namely the intra-country spillover and inter-country spillover. We then decompose the inter-country spillover effects by the source and destination of shocks into four sub-indices: 1) liquidity to liquidity, 2) liquidity to volatility, 3) volatility to liquidity, and 4) volatility to volatility. Finally, we analyze the dynamics of spillovers into and out of the Italian bond market, given its systemic importance among bond markets in the euro area.

5.5.1 Trends in Aggregate Spillovers

Figure 5.4 shows the time series plot of the total spillover index, together with its two components reflecting intra-country and inter-country spillover effects. The total spillover index indicates the average extent to which a variable is driven by all other variables in the system. The intra-country spillover measures the average degree of inter-dependency between liquidity and volatility within the same bond market. On the other hand, the inter-country spillover measures the average extent of spillover a country receives from the other countries. We also mark on the figure the five important dates discussed in the data section, namely: (1) May 10, 2010, (2) August 8, 2011, (3) November 16, 2011, (4) January 13, 2012, and (5) July 26, 2012.

The total spillover index exhibits several episodes of increased inter-dependencies among the markets, but it is generally trending down over the sample period. The index starts at roughly 80% in late May 2010, and finishes at about 55% at the end of 2012. Note that the high level at the beginning of the indices reflects the effects of the Greek crisis. It is followed by three more occasions of rising spillovers - evidence of contagion in the area. The first occurs in late November 2010, around the time that Ireland sought financial assistance (November 21, 2010), thereby renewing concerns over a widening crisis. Interestingly, the crisis in Portugal in April 2011 does not seem to affect the overall degree of shock transmission among the countries in consideration.

The second episode of contagion occurs in the summer of 2011, when it was apparent that the crisis could soon spread to two of the larger economies in the area, namely Italy and Spain, prompting the ECB to reactivate the SMP. The total spillover index exceeds the 70% level and stays between 70-80% for the rest of 2011 before declining to around 60% in early 2012. The degree of shock transmission among the six bond markets increases one more time in June and July of 2012 amid the banking crisis in Spain and a growing fear of a possible breakup of the eurozone. However, this upward trend in spillovers is soon reversed following the ECB President's statement to do "whatever it takes to preserve the euro". The markets appear to calm down and spillovers of shocks lessen after this assurance.

The figure also reveals that it is mainly the inter-country spillover component that drives the time series trend in the total spillover index discussed above. The within-country feedback dynamics between liquidity and volatility remain quite stable at around 10% over the sample period. Crisis events appear to affect only the cross-country spillovers. This finding confirms that the increased total spillover around these events is mainly attributable to increased shock transmission across borders (evidence of contagion) and not to increased inter-dependency within the same market.

5.5.2 Inter-Country Shock Transmission: A Decomposition

We now examine the dynamics of inter-country spillovers at a deeper level, which helps decipher the major source of shocks that permeates the euro area bond markets. Figure 5.5 plots the four types of cross-country spillover effects: 1) volatility to volatility ("V2V"), 2) liquidity to volatility ("L2V"), 3) volatility to liquidity ("V2L"), and 4) liquidity to liquidity ("L2L").

Visual inspection of the indices reveals that liquidity shock transmission is the main driver of cross-country spillovers. The L2L index is above all other three indices for almost all of the sample period. Importantly, the contagion relating to the Italian crisis in the summer of 2011 is heavily liquidity-driven, as shown by a dramatic increase in the L2L index.

The transmission of volatility shocks accounts for the second largest portion of cross-country spillovers. Initially, the V2V index appears comparable with the L2L index, and both are fluctuating around the 40% level. But after 2010, the extent of liquidity shock transmission far exceeds that of volatility shock transmission. In other words, the contagion associated with the Greek crisis is driven roughly equally by liquidity and volatility forces, but the contagion associated with the later phase of the crisis can be characterized as liquidity-driven.

The other two cross-type spillover indices (L2V and V2L) appear weaker, generally between 10-30%. They measure how strongly external liquidity shocks drive a bond market's volatility, and similarly, how strongly external volatility shocks drive a bond market's liquidity. Despite being weaker, these cross-type spillovers also increase around times of crisis (e.g., the Greek and Irish crises in 2010, the Italian crisis in the later half of 2011, and the Spanish crisis in the middle of 2012), thereby contributing to the spreading of shocks across borders.

5.5.3 The Dynamics of Italy's Systemic Role

As the full sample analysis has shown, Italy is the main exporter of both volatility and liquidity shocks to other markets. Given its systemic importance, we now provide a close-up analysis of the spillovers into and out of Italy over time. We first look at the extent to which foreign shocks affects Italy's liquidity ("Liquidity Shocks Received") or volatility ("Volatility Shocks Received") and plot them in the negative half in Figure 5.6. In the positive half, we show the extent to which Italian bond market liquidity and volatility drive other markets by the lines "Liquidity Shocks Sent" and "Volatility Shocks Sent" respectively.

Interestingly, external shocks affect Italian bond market's volatility and liquidity rather equally, and this influence is quite stable over the sample period. Roughly 50% of the variation in liquidity or volatility is due to shocks in the other markets.

On the other hand, liquidity and volatility fluctuations in the Italian market have markedly greater effects on the other countries. This is observed for most of the sample period, indicating that Italy is a consistent net exporter of shocks to others. In particular, between August and November of 2011 – the peak of the Italian crisis – the extent of liquidity shock transmitted from Italy to the other markets is about three times the reciprocal effect. Volatility transmission is also twice as large going from Italy than to Italy. This major increase in shock transmission from Italy, especially that of liquidity, is the major force underlying the increase in the overall cross-country liquidity spillovers. We did not observe the same magnitude of shock transmission in the remaining markets, including Spain, even though it is also a country in crisis.

5.6 Conclusion

In this paper, we examine the transmission of liquidity and volatility shocks across six major bond markets in the euro area during the 2010-2012 sovereign credit crisis. We model liquidity and volatility jointly in a linear VAR model with additional variables to control for the common trends in sovereign credit risk, financial sector's credit risk, funding tightness, overall default risk, and regional and global risk factors. Spillovers are measured based on the forecast error variance decomposition of the model. Apart from a full sample analysis of the spillover effects, we also examine their dynamics throughout the crisis period using rolling estimates.

Our main findings can be summarized as follows. First, variables capturing the aggregate market conditions affect the bond markets mainly through the liquidity channel. Furthermore, most of the feedback effects within the system are due to liquidity. In addition, Italy is the main exporter of both liquidity shocks and volatility shocks to others. Germany also contributes net volatility shocks to the system, albeit at a much smaller scale than that by Italy. However, the German bond market receives considerable liquidity spillovers from the other countries, most notably, Italy.

The time series variation of spillovers shows that the degree of shock transmission among the major euro area bond markets increases around major crisis events. Examining multiple layers of spillovers, we find that liquidity is the main source of shocks transmitted across borders. Moreover, the liquidity shock transmission is largely originated from Italy, especially in the summer of 2011 when the crisis over Italy's swelling public debt intensified. Being the largest bond market and third largest economy in the area, the crisis in Italy has important implications for the overall crisis in the eurozone. This is illustrated by the sizable amount of spillovers Italy transmits to other major bond markets in the area.

The results in this paper highlight the importance of liquidity to euro area bond market dynamics. We show that liquidity not only explains a significant portion of market dynamics, but also reveals strong spillover effects – often stronger than volatility spillover effects. This is consistent with the liquidity commonality that has been documented for various asset classes and markets. It also suggests that incorporating liquidity effects into a traditional volatility spillover framework is insightful as it allows for a better understanding of the sources of shocks that get transmitted across borders during a crisis.

Table 5.1: Descriptive Statistics of Volatility and Liquidity

	Mean	Min	P25	Median	P75	Max
			A. Market Volatility (% p.a.)			
Germany	6.01	2.40	4.65	5.52	7.01	18.08
Netherlands	6.37	2.47	4.65	5.69	7.47	19.72
France	6.77	2.35	4.37	5.67	7.63	45.00
Belgium	9.45	2.03	4.39	6.61	10.68	74.17
Italy	9.91	2.50	4.87	6.89	11.00	127.16
Spain	13.15	2.89	6.43	9.67	15.32	149.15
			B. Market Liquidity (bps)			
Germany	13.89	8.78	11.40	12.81	15.12	59.00
Netherlands	13.24	6.17	8.90	11.90	14.98	56.94
France	24.25	10.39	14.70	19.13	27.83	119.04
Belgium	40.33	7.32	21.69	29.46	44.25	267.65
Italy	42.62	11.49	18.96	30.97	47.26	667.43
Spain	72.78	11.38	34.18	58.59	92.76	487.11

This table presents descriptive statistics of volatility and liquidity of six Euro-area bond markets, based on MTS intraday quote data for the period 2010-2012. Volatility is the square root of the annualized realized variance (RV) of each bond market (average of daily RV's of individual bonds issued by that country, excluding bonds with less than one year to maturity). Each bond's daily RV is computed as the daily sum of squared five-minute log mid-quote returns, and annualized by a factor of 250. The liquidity measure for each country is the average of daily liquidity measures across all bonds issued by that country, excluding bonds with less than one year to maturity. Daily liquidity of each individual bond is measured by the time-weighted relative bid-ask spread as a fraction of mid-quote.

Table 5.2: Correlation of Liquidity and Volatility Across Markets

	LiqDE	LiqNL	LiqFR	LiqBE	LiqIT	LiqES	VolDE	VolNL	VolFR	VolBE	VolIT	VolES
LiqDE	1.00											
LiqNL	0.78	1.00										
LiqFR	0.80	0.81	1.00									
LiqBE	0.78	0.66	0.84	1.00								
LiqIT	0.84	0.81	0.93	0.87	1.00							
LiqES	0.66	0.80	0.78	0.64	0.82	1.00						
VolDE	0.65	0.58	0.64	0.65	0.68	0.54	1.00					
VolNL	0.64	0.67	0.69	0.65	0.69	0.60	0.91	1.00				
VolFR	0.70	0.59	0.70	0.77	0.71	0.42	0.79	0.77	1.00			
VolBE	0.65	0.53	0.70	0.86	0.71	0.45	0.66	0.65	0.81	1.00		
VolIT	0.75	0.78	0.86	0.77	0.92	0.79	0.73	0.74	0.72	0.70	1.00	
VolES	0.58	0.72	0.71	0.57	0.74	0.88	0.61	0.64	0.51	0.51	0.82	1.00

The table shows the correlation matrix of liquidity (Liq) and volatility (Vol) of six Euro-area bond markets: Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES). Volatility is the square root of the annualized realized variance (RV) of each bond market (average of daily RV's of individual bonds issued by that country, excluding bonds with less than one year to maturity). Each bond's daily RV is computed as the daily sum of squared five-minute log mid-quote returns, and annualized by a factor of 250. The liquidity measure for each country is the average of daily liquidity measures across all bonds issued by that country, excluding bonds with less than one year to maturity. Daily liquidity of each individual bond is measured by the time-weighted relative bid-ask spread as a fraction of mid-quote. The above correlation coefficients are calculated based on log-transformed variables. The sample period is 2010-2012.

Table 5.3: Model Estimates – Feedback Effects

	Equation											
Lag 1	LiqDE	LiqNL	LiqFR	LiqBE	LiqIT	LiqES	VolDE	VolNL	VolFR	VolBE	VolIT	VolES
LiqDE	0.57***	−0.14**	−0.13**	−0.17**	−0.17**	−0.17**	−0.28*	−0.40**	-0.26	−0.81***	−0.62***	−0.77***
LiqNL	0.07**	0.78***	0.00	-0.04	0.11**	0.14**	0.04	0.09	0.37***	0.09	0.23	0.28
LiqFR	−0.18***	−0.27***	0.52***	−0.19***	−0.24***	−0.29***	−0.29**	−0.38**	−0.56***	-0.28	-0.21	-0.25
LiqBE	0.01	-0.05	0.02	0.83***	0.07	0.14***	0.18**	0.18*	0.27**	0.91***	-0.01	0.25
LiqIT	0.11***	0.27***	0.18***	0.17***	0.87***	0.12*	0.32***	0.40***	0.44***	0.34	0.88***	0.40*
LiqES	0.03	0.05	-0.00	0.08*	0.02	0.79***	0.05	0.09	−0.20*	-0.13	0.04	0.36**
VolDE	0.03	-0.03	−0.06**	0.03	0.05	0.08**	0.55***	0.18**	0.11	0.09	0.11	0.25*
VolNL	-0.02	0.00	0.06**	0.02	-0.01	-0.03	0.06	0.42***	0.04	0.03	0.02	−0.20*
VolFR	0.01	0.04*	-0.00	0.02	-0.03	−0.06**	-0.04	-0.02	0.28***	0.09	−0.12*	-0.06
VolBE	0.00	0.02	0.02*	-0.01	0.02	0.01	-0.04	-0.01	0.00	0.13**	0.06	-0.03
VolIT	0.02	-0.01	0.01	-0.01	0.01	0.03	-0.03	-0.04	-0.03	-0.03	0.39***	0.16**
VolES	-0.01	−0.03**	-0.00	-0.01	-0.00	-0.02	-0.01	-0.03	0.05	0.10	-0.01	0.29***

The table shows estimates of the feedback parameters in the VAR(1) model with market condition variables, as specified in equation 5.6. The model variables are liquidity (Liq) and volatility (Vol) of six Euro-area bond markets: Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES). A bond market's liquidity is the average liquidity level across all coupon bonds in that market, where each bond's liquidity is measured by the daily time-weighted average bid-ask spread as a fraction of mid-quote between 9:00 and 17:00. A bond market's volatility is measured by the logged average of realized variances of all coupon bonds in that market. A bond's daily variance is the daily sum of squared five-minute returns. The sample period is 2010-2012. Note: (*), (**), and (***) indicate significance at the 10%, 5% and 1% levels.

Table 5.4: Model Estimates – Effects of Aggregate Market Conditions

		Equation											
Lag	Variable	LiqDE	LiqNL	LiqFR	LiqBE	LiqIT	LiqES	VolDE	VolNL	VolFR	VolBE	VolIT	VolES
	SovPC1	-0.10^{***}	-0.07	-0.05	-0.01	0.03	-0.03	-0.29^{***}	-0.24^{*}	-0.21	-0.02	0.04	-0.42^{**}
	iTraxxFin	0.24	0.35	0.25	0.42	0.71^{**}	0.66^{*}	1.85^{***}	2.16^{***}	1.55^{*}	1.24	2.61^{***}	2.97^{***}
	FinCost	0.48^{**}	0.42	0.70^{**}	0.31	0.94^{**}	0.39	-0.08	0.67	0.89	0.93	0.78	1.26
L0	DefSpr	0.01	0.04^{***}	0.02^{**}	0.01	0.02^{**}	0.03^{**}	0.02	0.05^{*}	0.02	-0.01	0.05	0.10^{**}
	VSTOXX	0.27^{***}	0.35^{**}	0.43^{***}	0.46^{***}	0.39^{**}	0.48^{***}	0.22	0.04	0.50	1.37^{**}	0.76	1.25^{**}
	USVIX	0.17^{**}	0.18	0.21^{*}	-0.08	0.06	0.04	0.28	0.34	0.37	-0.43	0.04	0.16
	SovPC1	0.01	0.02	-0.02	0.05	0.04	0.00	0.04	0.09	0.11	0.27	-0.07	-0.04
	iTraxxFin	0.05	0.07	0.24	0.03	0.28	0.26	-1.21^{**}	-1.48^{**}	-1.23	-1.04	-0.50	-0.75
	FinCost	-0.32	-0.40	-0.45	0.06	-0.56	-0.38	0.50	-0.33	0.34	0.63	-0.27	-1.78
L1	DefSpr	-0.01	-0.03^{***}	-0.01^{*}	-0.02	-0.02^{*}	-0.01	-0.02	-0.04	-0.04	-0.03	-0.04	-0.06
	VSTOXX	-0.10	0.19	-0.02	0.32	0.01	0.05	0.73^{*}	1.15^{**}	0.66	1.30^{*}	0.86	0.88
	USVIX	-0.05	-0.10	-0.09	-0.11	-0.14	-0.13	0.08	0.12	0.12	0.30	-0.00	-0.20

The table shows the estimated effects (both contemporaneously and at lag 1) of market condition variables in the model specified in equation 5.6. The model variables are liquidity (Liq) and volatility (Vol) of six Euro-area bond markets: Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES). $SovPC1$ is the first principle component extracted from the 5-year sovereign CDS spreads of the six countries in consideration. $iTraxxFin$ is the iTraxx Financial CDS index. $FinCost$ is the difference between the 3-month Euribor and 3-month EOIS rates. $DefSpr$ is the yield differential between AAA-rated and BAA-rated corporate bond indices. $VSTOXX$ is the volatility of the STOXX index. $USVIX$ is the volatility of the S&P500. Except $FinCost$ and $DefSpr$, all other market condition variables are first differenced. The sample period is 2010-2012. Note: (*), (**), and (***) indicate significance at the 10%, 5% and 1% levels.

Table 5.5: Correlation of Residuals

	LiqDE	LiqNL	LiqFR	LiqBE	LiqIT	LiqES	VolDE	VolNL	VolFR	VolBE	VolIT	VolES
LiqDE	1.00											
LiqNL	0.65	1.00										
LiqFR	0.70	0.79	1.00									
LiqBE	0.54	0.65	0.71	1.00								
LiqIT	0.59	0.59	0.67	0.64	1.00							
LiqES	0.55	0.57	0.65	0.64	0.79	1.00						
VolDE	0.42	0.43	0.42	0.34	0.43	0.38	1.00					
VolNL	0.40	0.55	0.49	0.36	0.39	0.38	0.83	1.00				
VolFR	0.42	0.54	0.57	0.47	0.47	0.45	0.71	0.70	1.00			
VolBE	0.37	0.49	0.51	0.70	0.47	0.46	0.45	0.46	0.56	1.00		
VolIT	0.44	0.51	0.56	0.52	0.76	0.65	0.58	0.53	0.58	0.53	1.00	
VolES	0.38	0.44	0.53	0.51	0.62	0.73	0.52	0.48	0.57	0.53	0.72	1.00

The table shows the correlation matrix of residuals from the VAR(1) model specified in equation 5.6. The model variables are liquidity (Liq) and volatility (Vol) of six Euro-area bond markets: Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES). Explanatory variables include the following. $SovPC1$ is the first principle component extracted from the 5-year sovereign CDS spreads of the six countries in consideration. $iTraxxFin$ is the iTraxx Financial CDS index. $FinCost$ is the difference between the 3-month Euribor and 3-month EOIS rates. $DefSpr$ is the yield differential between AAA-rated and BAA-rated corporate bond indices. $VSTOXX$ is the volatility of the STOXX index. $USVIX$ is the volatility of the S&P500. Except $FinCost$ and $DefSpr$, all other market condition variables are first differenced. The sample period is 2010-2012.

Table 5.6: Bilateral Spillover Matrix

	VolDE	VolNL	VolFR	VolBE	VolIT	VolES	LiqDE	LiqNL	LiqFR	LiqBE	LiqIT	LiqES
VolDE	43	18	9	3	5	4	2	2	3	3	5	3
VolNL	21	38	9	3	4	3	2	4	4	4	5	3
VolFR	11	9	39	5	5	5	2	5	5	5	6	3
VolBE	4	3	5	42	4	5	3	4	4	18	5	4
VolIT	6	3	4	4	31	8	3	4	5	5	20	9
VolES	5	3	4	3	11	33	3	3	4	6	11	14
LiqDE	5	2	2	2	5	2	34	8	10	5	15	9
LiqNL	2	3	3	2	4	2	5	37	11	5	15	7
LiqFR	3	3	4	4	6	4	7	10	25	11	15	8
LiqBE	4	2	2	8	4	4	4	5	7	40	11	11
LiqIT	4	2	2	3	10	5	4	5	6	8	36	14
LiqES	4	1	2	3	7	7	4	5	7	11	16	35

The table shows the full bilateral spillover matrix. This is based on the 10-day forecast error variance decomposition of the VAR(1) model with market condition variables, as specified in equation 5.6. Element (i,j) of the matrix indicates the variance contribution by the variable shown in column j's heading to the variable shown in row i's heading. Numbers in each row sum to 100%. The model variables are liquidity (Liq) and volatility (Vol) of the following Euro-area bond markets: Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES). A market's liquidity is the average liquidity level across all coupon bonds in that market, where each bond's liquidity is measured by the daily time-weighted average bid-ask spread as a fraction of mid-quote. A market's volatility is measured by the logged average of realized variances of all coupon bonds in that market. A bond's daily variance is the daily sum of squared five-minute returns. Daily liquidity and volatility measures are computed from the MTS intraday quote data for the 2010-2012 period.

Table 5.7: Intra-Country Spillovers

Country	Liq → Vol	Vol → Liq	Net Spillover
DE	2.5	5.1	-2.6
NL	4.3	3.3	0.9
FR	5.2	3.5	1.7
BE	18.2	7.9	10.3
IT	19.6	10.4	9.2
ES	14.0	7.5	6.5

The table shows the extent to which liquidity drives volatility ("Liq → Vol" column) and conversely, the extent to which volatility drives liquidity ("Liq → Vol" column) within each bond market considered. The last column ("Net Spillover") shows liquidity's net variance contribution to volatility. These figures are based on 10-day forecast error variance decomposition of the VAR(1) model of liquidity and volatility of six bond markets, i.e., Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES), with market condition variables, as specified in equation (5.6). A market's liquidity is the average liquidity level across all coupon bonds in that market, where each bond's liquidity is measured by the daily time-weighted average bid-ask spread as a fraction of mid-quote. A market's volatility is measured by the logged average of realized variances of all coupon bonds in that market. A bond's daily variance is the daily sum of squared five-minute returns. Daily liquidity and volatility measures are computed from the MTS intraday quote data for the 2010-2012 period.

Table 5.8: Inter-Country Spillovers

	Received from Others			Sent to Others			Net Sent to Others		
	Vol	Liq	Total	Vol	Liq	Total	Vol	Liq	Total
VolDE	38.4	16.4	54.8	47.0	16.7	63.7	8.6	0.3	8.9
VolNL	40.0	17.5	57.5	35.4	9.9	45.4	-4.5	-7.5	-12.1
VolFR	34.5	21.4	55.9	30.1	11.7	41.8	-4.4	-9.7	-14.1
VolBE	20.5	19.1	39.6	17.1	14.2	31.3	-3.4	-4.8	-8.3
VolIT	24.1	25.4	49.5	29.4	25.4	54.8	5.2	0.1	5.3
VolES	25.5	27.2	52.6	24.0	17.0	41.0	-1.5	-10.1	-11.6
LiqDE	13.2	47.2	60.4	13.3	24.0	37.4	0.1	-23.2	-23.1
LiqNL	14.3	45.0	59.3	17.5	33.2	50.7	3.2	-11.7	-8.6
LiqFR	20.2	51.0	71.1	19.1	41.6	60.7	-1.1	-9.3	-10.4
LiqBE	15.1	36.9	52.0	22.6	40.4	63.0	7.6	3.5	11.0
LiqIT	16.1	37.6	53.6	31.8	71.6	103.5	15.8	34.0	49.8
LiqES	16.2	41.7	57.8	22.5	48.4	70.9	6.3	6.8	13.1

The table shows the spillovers received from other countries ("Received from Others" panel), the spillovers sent to other countries ("Sent to Others" panel) and the net spillover effects ("Net Sent to Others"). Each panel shows separately the contributions by volatility ("Vol") and liquidity ("Liq") components, and reports the sum of them in the column "Total". These spillover measures are computed as in equations (5.7)-(5.11). They are based on 10-day forecast error variance decomposition of the VAR(1) model of liquidity and volatility of six bond markets, i.e., Germany (DE), France (FR), the Netherlands (NL), Belgium (BE), Italy (IT) and Spain (ES), with market condition variables, as specified in equation (5.6). A market's liquidity is the average liquidity level across all coupon bonds in that market, where each bond's liquidity is measured by the daily time-weighted average bid-ask spread as a fraction of mid-quote. A market's volatility is measured by the logged average of realized variances of all coupon bonds in that market. A bond's daily variance is the daily sum of squared five-minute returns. Daily liquidity and volatility measures are computed from the MTS intraday quote data for the 2010-2012 period.

Figure 5.1: The Euro Area Sovereign Debt Crisis

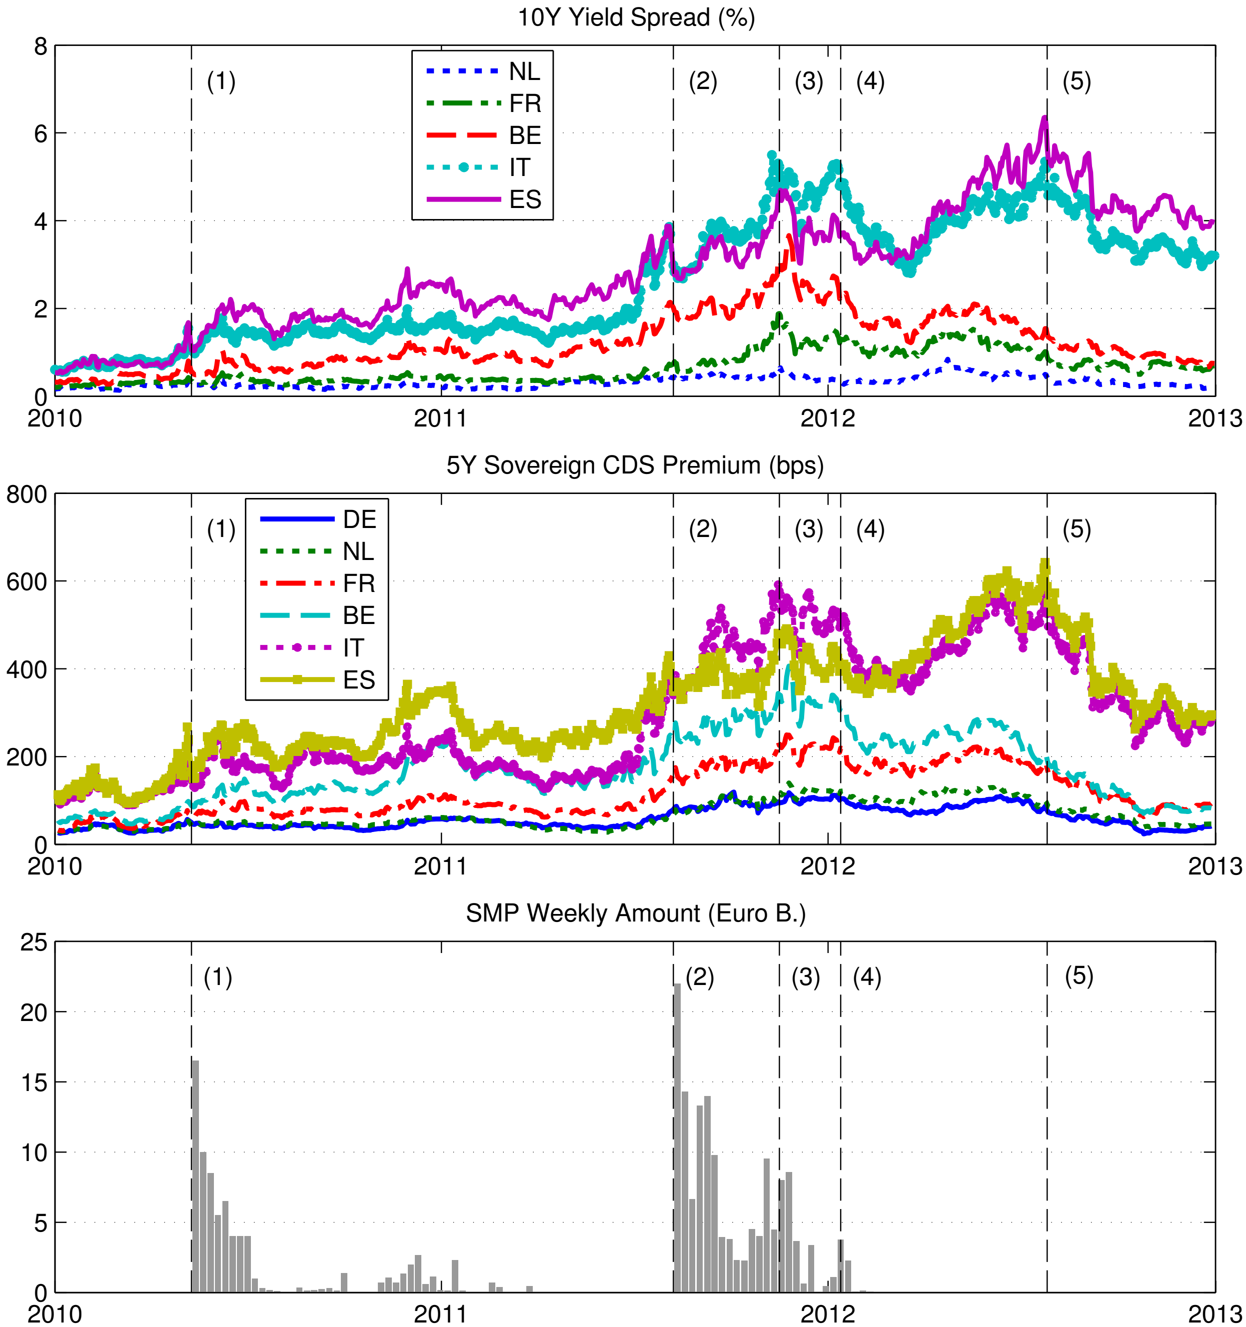

This figure shows the time series variation in: (top panel) 10-year benchmark yield spreads (over German 10-year benchmark yield), (middle panel) sovereign CDS spreads, and (bottom panel) the European Central Bank's Securities Market Programme activity. Vertical lines mark (1) 5/10/2010 (SMP1), (2) 8/8/2011 (SMP2), (3) 11/16/2011 (change of government in Italy), (4) 1/13/2012 (S&P credit rating downgrades for France, Italy and Spain), (5) 7/26/2012 (ECB President's statement to do "whatever it takes to preserve the euro").

Figure 5.2: Time Series Variation in Euro Area Bond Market Volatility

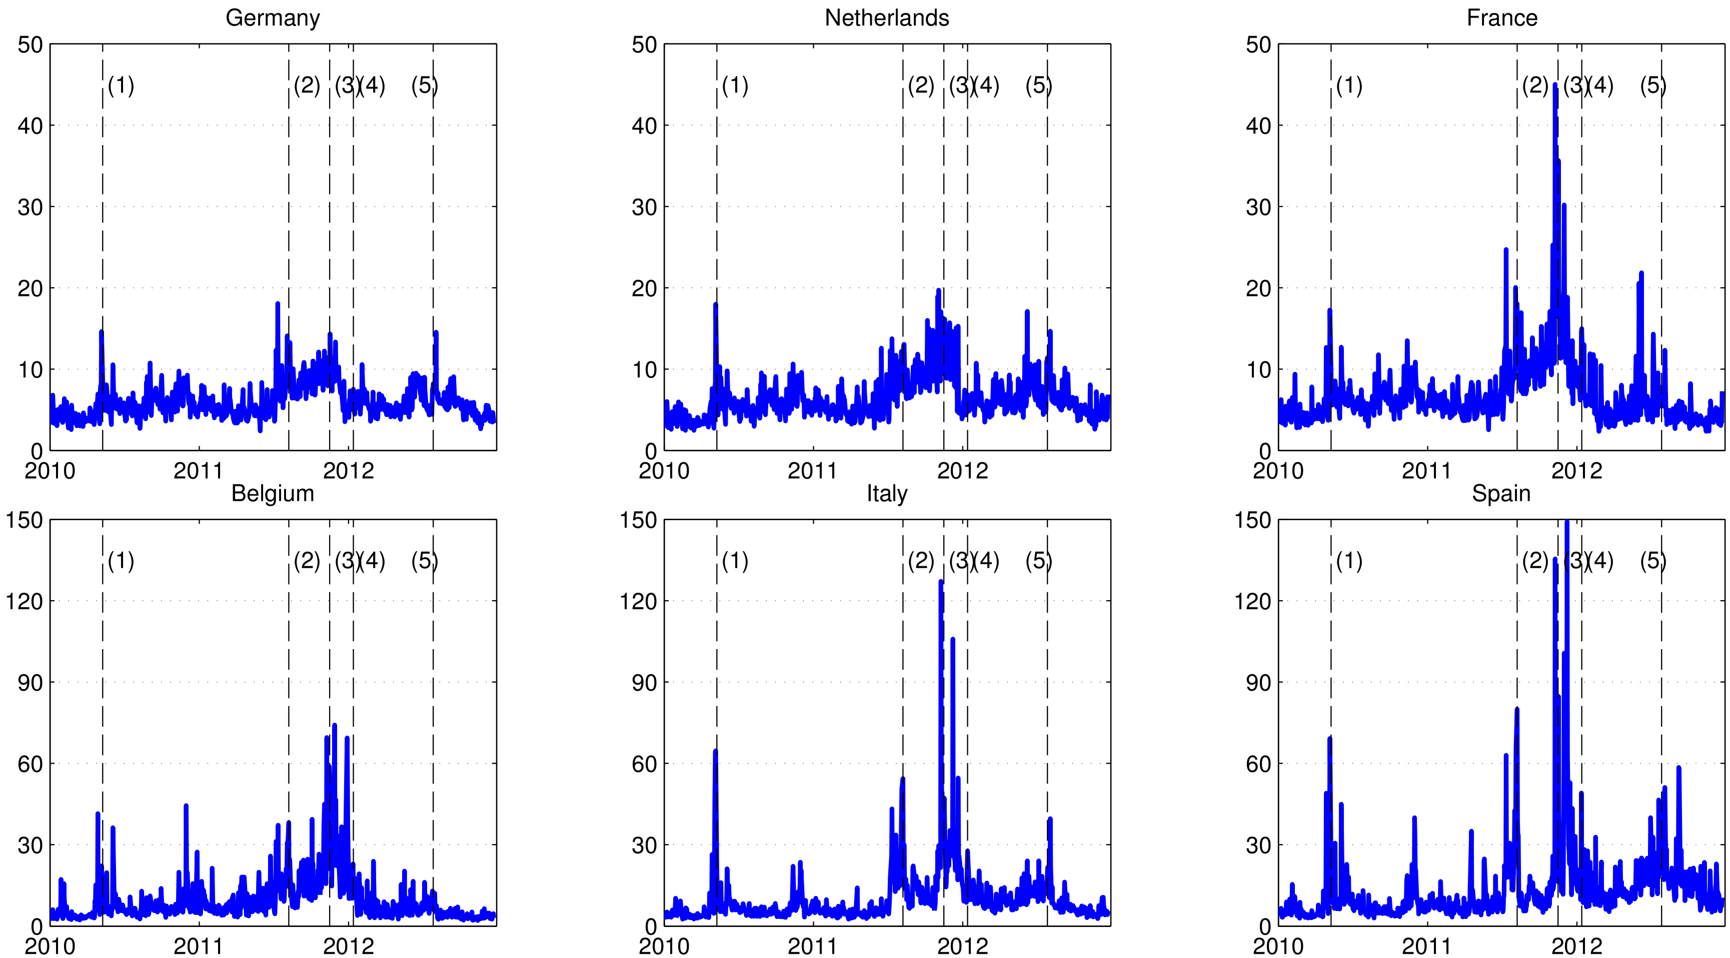

This figure plots daily annualized volatility (in %). Vertical lines mark (1) 5/10/2010 (SMP1), (2) 8/8/2011 (SMP2), (3) 11/16/2011 (change of government in Italy), (4) 1/13/2012 (S&P credit rating downgrades for France, Italy and Spain), (5) 7/26/2012 (ECB President's statement to do "whatever it takes to preserve the euro"). Volatility is the square root of the annualized realized variance (RV) of each country's government bond market. Daily RV of each market is the volume-weighted average of the daily RV's of individual bonds. Each bond's daily RV is computed as the daily sum of squared 5-minute log mid-quote returns, and annualized by a factor of 250. Individual bond RV's are sub-sampled at 1-minute interval. Data source: MTS intraday quote data for the period 2010-2012.

Figure 5.3: Time Series Variation in Euro Area Bond Market Liquidity

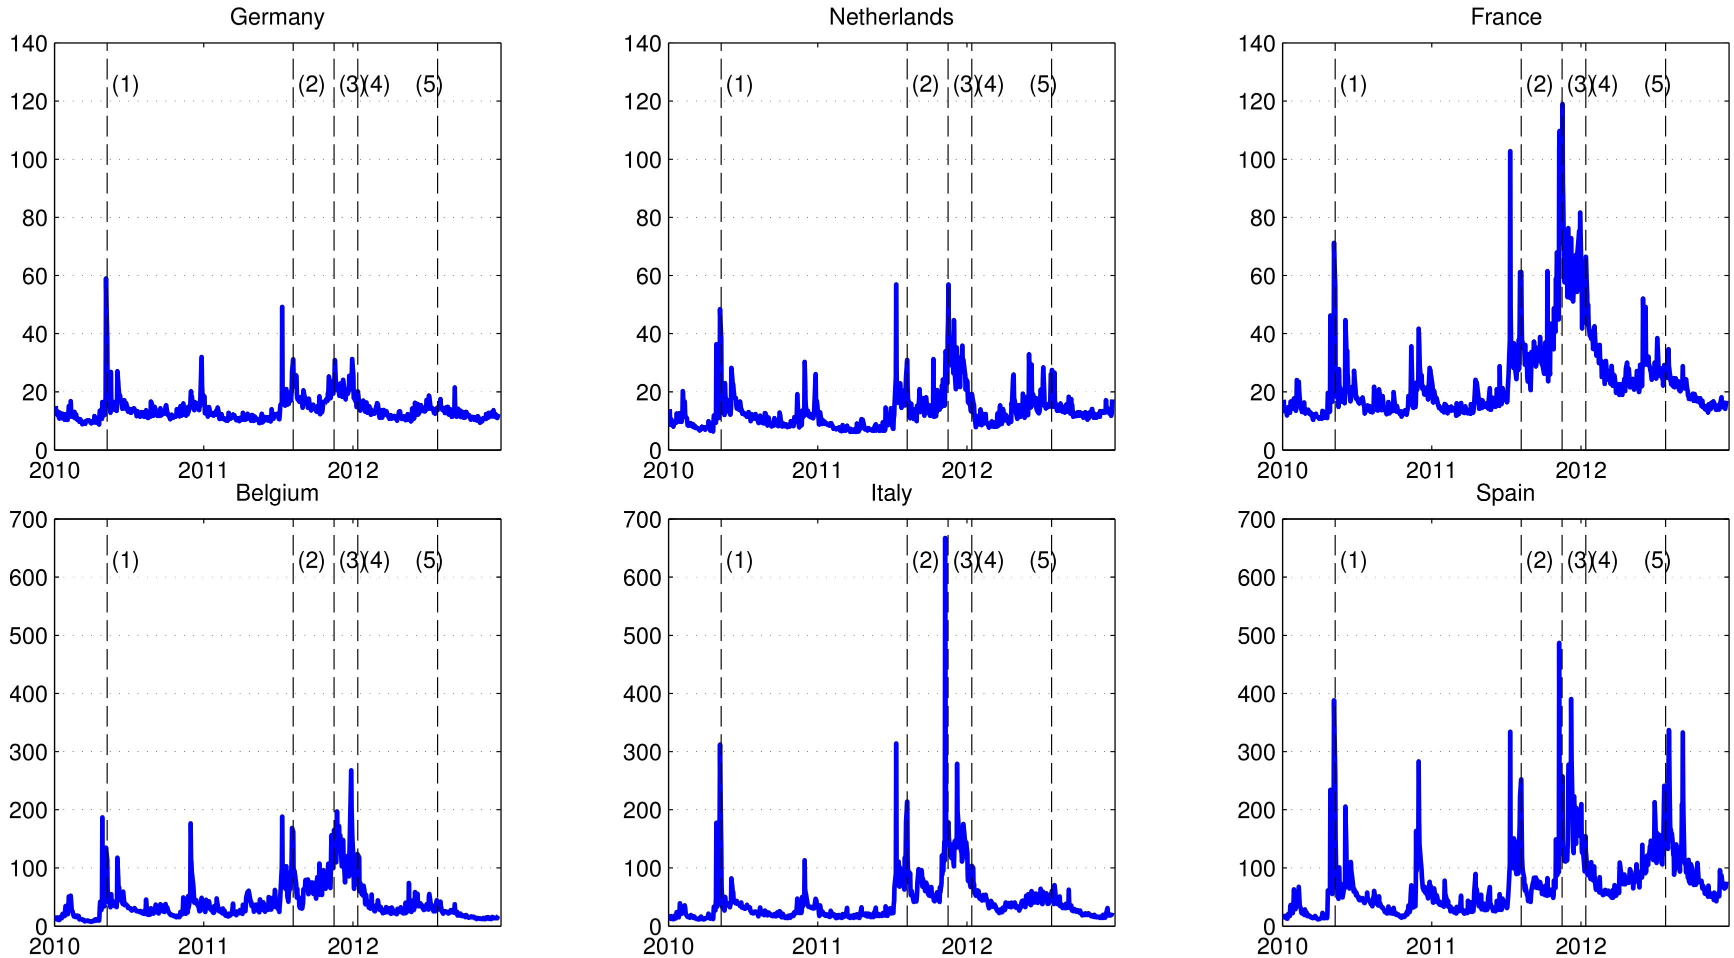

This figure plots daily (il)liquidity measure (in bps). The liquidity measure for each country is the volume-weighted average of the daily liquidity measures of bonds issued by that country. Daily liquidity of each individual bond is measured by the time-weighted percentage bid-ask spread computed from intraday quote data. Vertical lines mark (1) 5/10/2010 (SMP1), (2) 8/8/2011 (SMP2), (3) 11/16/2011 (change of government in Italy), (4) 1/13/2012 (S&P credit rating downgrades for France, Italy and Spain), (5) 7/26/2012 (ECB President's statement to do "whatever it takes to preserve the euro"). Data source: MTS intraday quote data for the period 2010-2012.

Figure 5.4: Dynamics of Total Spillover Among Key European Government Bond Markets

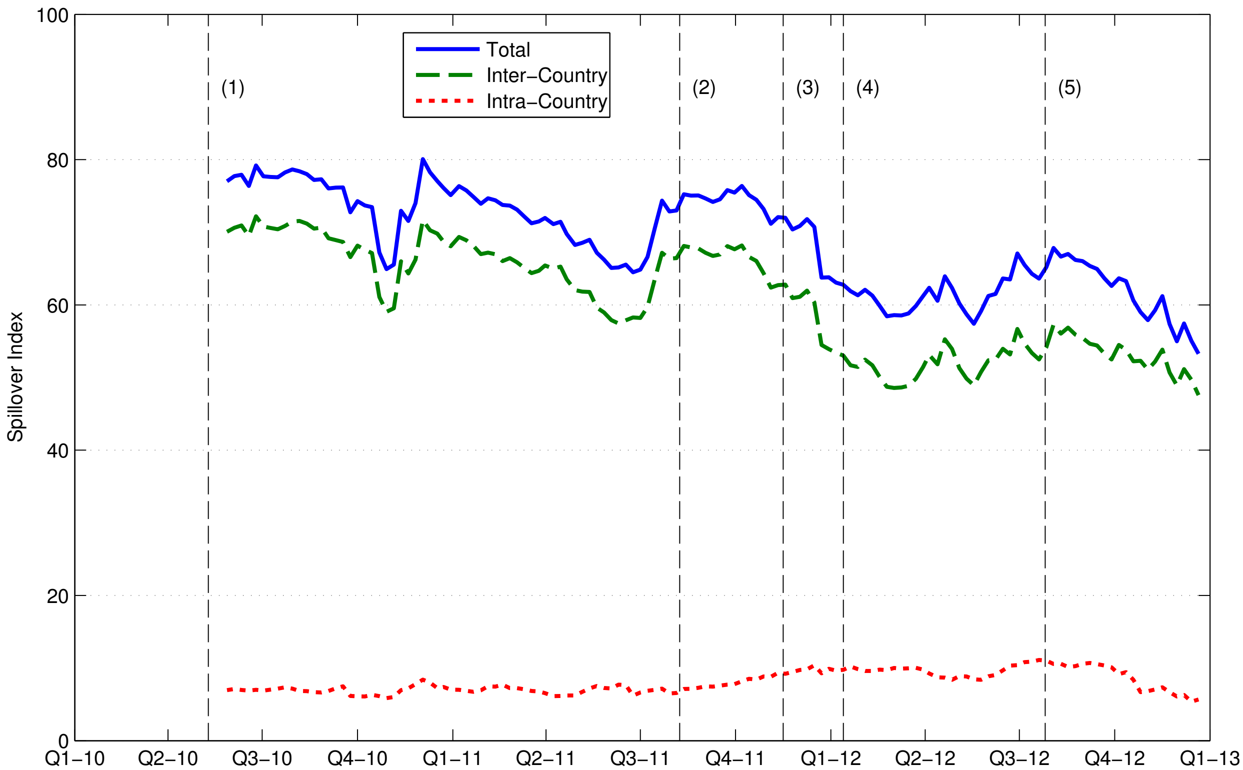

This figure presents the weekly total spillover index and its two components (intra-country and inter-country spillover indices). These spillover indices are based on 10-day forecast error variance decomposition of the VAR(1) model of liquidity and volatility of six bond markets (Germany, France, the Netherlands, Belgium, Italy and Spain) with market condition variables, as specified in equation (5.6). The model is estimated on a weekly rolling basis, with a window size of 100 days. The indices are computed as in equations (5.13)-(5.15). Vertical lines mark (1) 5/10/2010 (SMP1), (2) 8/8/2011 (SMP2), (3) 11/16/2011 (change of government in Italy), (4) 1/13/2012 (S&P credit rating downgrades for France, Italy and Spain), (5) 7/26/2012 (ECB President's statement to do "whatever it takes to preserve the euro"). Daily liquidity and volatility measures are computed from MTS bond market intraday quote data for the period 2010-2012.

Figure 5.5: The Sources of Shocks that Travel Across Borders

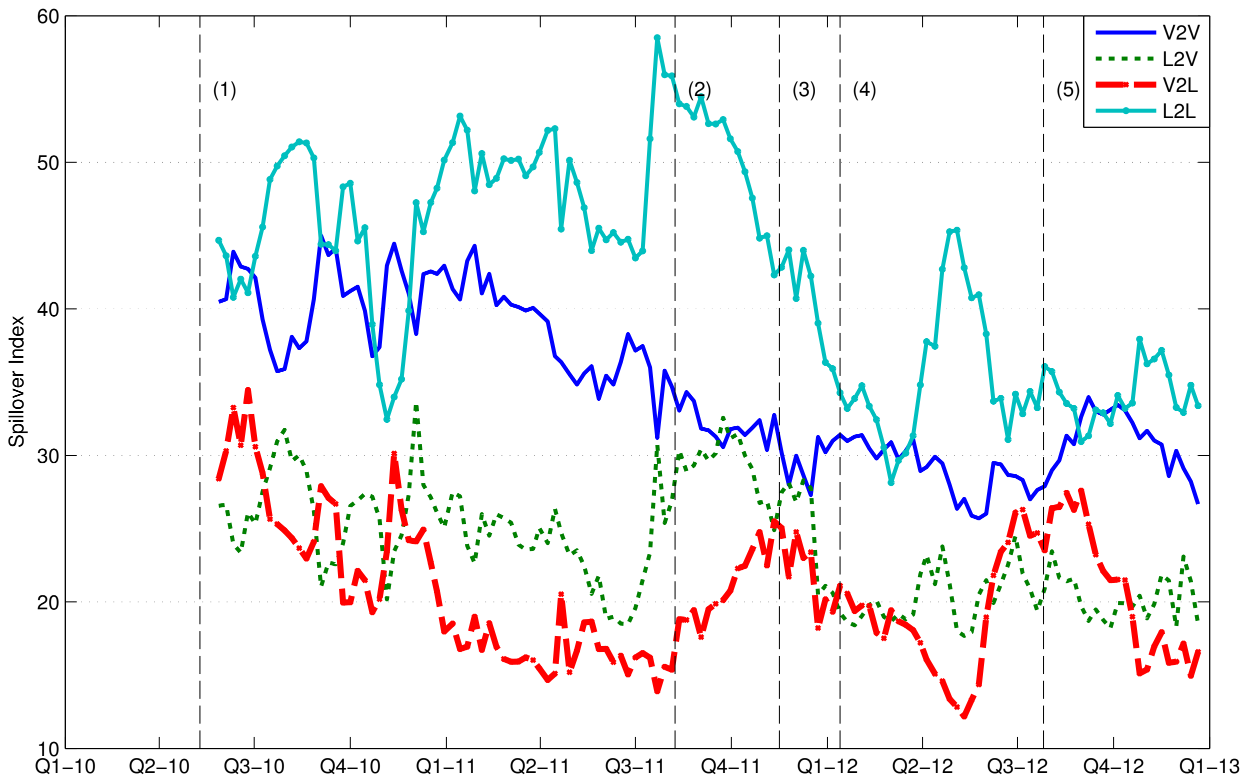

This figure shows four inter-country spillover sub-indices classified by the source of shocks (volatility or liquidity) and destination of spillovers (volatility or liquidity). These spillover indices are based on 10-day forecast error variance decomposition of the VAR(1) model of liquidity and volatility of six bond markets (Germany, France, the Netherlands, Belgium, Italy and Spain) with market condition variables, as specified in equation (5.6). The model is estimated on a weekly rolling basis, with a window size of 100 days. The indices are computed as in equations (5.16)-(5.19). Vertical lines mark (1) 5/10/2010 (SMP1), (2) 8/8/2011 (SMP2), (3) 11/16/2011 (change of government in Italy), (4) 1/13/2012 (S&P credit rating downgrades for France, Italy and Spain), (5) 7/26/2012 (ECB President's statement to do "whatever it takes to preserve the euro"). Daily liquidity and volatility measures are computed from MTS bond market intraday quote data for the period 2010-2012.

Figure 5.6: The Systemic Role of Italy

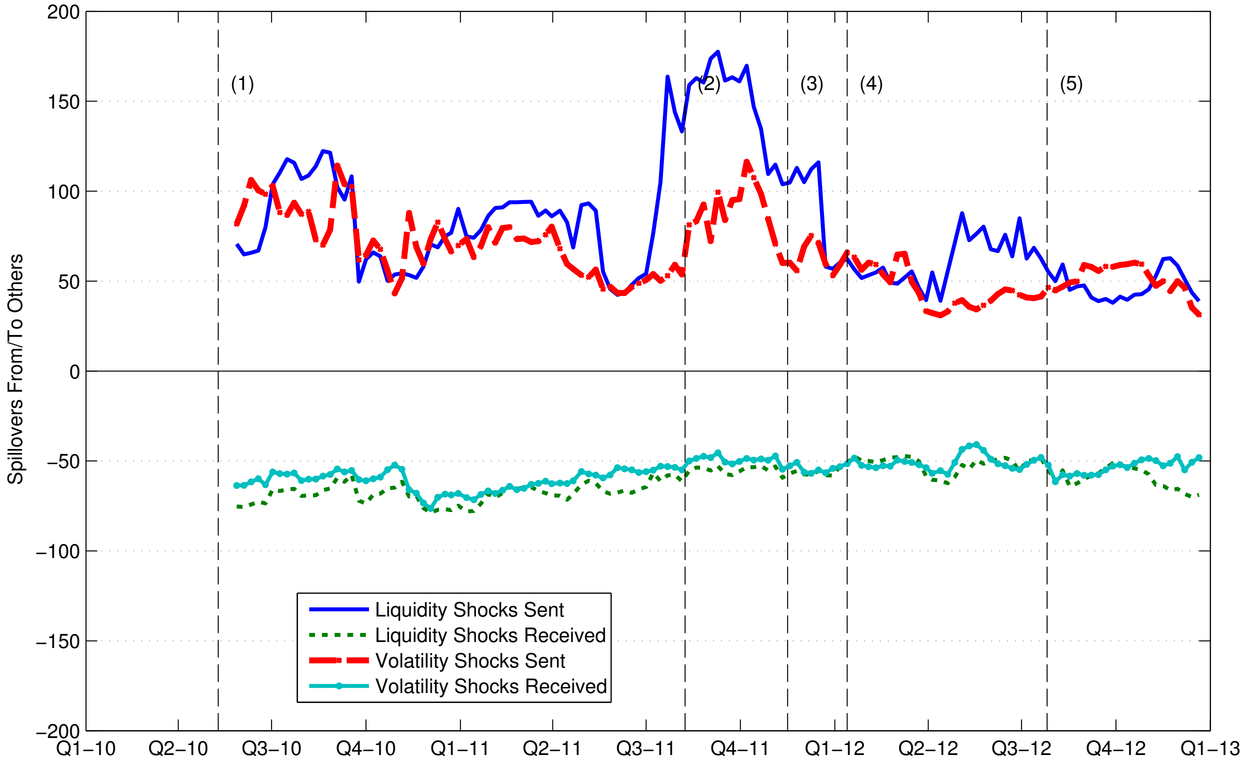

This figure shows the degree of spillovers sent (plotted above the zero line) and spillovers received (plotted below the zero line) by the Italian bond market. The "Liquidity Shocks Sent" line shows the extent to which Italy's bond market liquidity affects the other markets, while the "Liquidity Shocks Received" shows the extent to which it is reciprocally affected by other markets. The "Volatility Shocks Sent" line shows the extent to which Italy's bond market volatility affects the other markets, while the "Volatility Shocks Received" shows the extent to which it is reciprocally affected by other markets. Spillovers are measured based on 10-day forecast error variance decomposition of the VAR(1) model of liquidity and volatility of six bond markets (Germany, France, the Netherlands, Belgium, Italy and Spain) with market condition variables, as specified in equation (5.6). The model is estimated on a weekly rolling basis, with a window size of 100 days. Vertical lines mark (1) 5/10/2010 (SMP1), (2) 8/8/2011 (SMP2), (3) 11/16/2011 (change of government in Italy), (4) 1/13/2012 (S&P credit rating downgrades for France, Italy and Spain), (5) 7/26/2012 (ECB President's statement to do "whatever it takes to preserve the euro"). Daily liquidity and volatility measures are computed from MTS bond market intraday quote data for the period 2010-2012.

APPENDIX A
ECONOMIC ANNOUNCEMENTS

A.1 Macroeconomic Announcements

The macroeconomic announcements we consider are those classified as ”Market Moving” indicators by Bloomberg: 1) Employment Report, 2) Consumer Price Index, 3) Durable Goods Orders, 4) GDP, 5) Housing Starts, 6) Initial Jobless Claims, 7) Personal Income and Outlays, 8) Producer Price Index, 9) Retail Sales, 10) Trade Balance, 11) Industrial Production, 12) Existing Home Sales, 13) ISM Manufacturing, 14) New Home Sales, and 15) Philadelphia Fed Survey.

Time	Announcement	Frequency
8:30	Employment Report	Monthly
8:30	Consumer Price Index (MoM)	Monthly
8:30	Durable Goods Orders	Monthly
8:30	GDP QoQ (Annualized)	Quarterly
8:30	Housing Starts	Monthly
8:30	Initial Jobless Claims	Weekly
8:30	Personal Income and Outlays	Monthly
8:30	Producer Price Index (MoM)	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
10:00	Philadelphia Fed. (after 2008)	Monthly
12:00	Philadelphia Fed. (before 2008)	Monthly

A.2 Monetary Policy Announcements

The monetary policy announcements included in our analysis are FOMC rate decision announcements. Such announcements typically occur after regularly scheduled FOMC meetings, of which there are eight per year. In addition, there were rate changes announced after unscheduled meetings on two occasions during our sample period, on January 22, 2008 and October 8, 2008.

A.3 Treasury Auction Result Announcements

The Treasury auction results we consider are those for the 2-, 5-, 10- and 30-year fixed principal Treasury securities. Auction results are announced shortly after the auction close on auction dates for a given security. The 2- and 5-year notes are newly issued every month. The 10-year note is newly issued every quarter, with reopenings in the following month and – since November 2008 – two months. Starting in May 2009, the 30-year bond is also on a quarterly issuance cycle with two reopenings. For the 2006-2008 period, the 30-year bond was newly issued once a year with irregular reopenings.

APPENDIX B

FLIGHT-TO-SAFETY EPISODES DURING 2006-2010Q2 PERIOD

	Light FTS			**Moderate FTS**			**Severe FTS**		
Date	2-Year	5-Year	10-Year	2-Year	5-Year	10-Year	2-Year	5-Year	10-Year
02/27/2007	X	X	X	X	X	X	X		
03/13/2007	X	X							
07/26/2007	X	X	X						
08/03/2007	X	X	X	X	X				
08/09/2007	X	X	X	X	X				
08/28/2007	X	X							
09/07/2007	X	X	X						
10/19/2007	X	X	X	X					
11/01/2007	X	X	X	X	X	X			
11/07/2007	X	X		X	X				
11/19/2007	X	X	X						
11/21/2007	X	X	X						
11/26/2007	X	X	X						
12/11/2007	X	X	X	X	X	X			
01/04/2008	X								
01/08/2008		X							
01/15/2008	X	X	X						
01/17/2008	X	X	X		X	X			
01/25/2008	X	X	X						
02/05/2008	X	X	X	X	X		X		
02/29/2008	X	X	X	X	X	X			
03/06/2008	X	X	X						
03/14/2008	X	X							
03/19/2008	X	X	X	X	X	X			
04/11/2008	X	X	X						
05/07/2008	X	X							
06/06/2008	X	X	X	X	X	X			
06/11/2008	X	X							
06/26/2008	X	X		X	X				
07/02/2008	X								
07/09/2008	X	X	X						
07/24/2008	X	X	X						
07/28/2008	X	X	X						
08/07/2008	X	X	X						
08/25/2008	X	X	X						
09/04/2008	X	X	X						
09/09/2008	X	X	X	X	X	X	X		
09/15/2008	X	X	X	X	X	X	X	X	X
09/17/2008	X			X			X		
09/29/2008	X	X	X	X	X	X	X	X	X
10/02/2008	X	X	X	X	X	X	X	X	
10/06/2008	X	X	X	X	X	X	X	X	X
10/15/2008	X	X	X	X	X	X	X	X	
10/21/2008	X	X	X		X				

continued on next page

Table B.1 – *continued from previous page*

	Light FTS			**Moderate FTS**			**Severe FTS**		
Date	2-Year	5-Year	10-Year	2-Year	5-Year	10-Year	2-Year	5-Year	10-Year
10/22/2008	X	X	X	X	X	X			X
10/24/2008	X								
11/12/2008	X	X	X		X				
11/14/2008		X	X			X			
11/17/2008			X						
11/19/2008		X	X		X	X		X	X
11/20/2008	X	X	X		X	X			X
12/01/2008	X	X	X		X	X		X	X
12/04/2008	X	X	X			X			
12/09/2008	X	X	X						
12/11/2008	X		X						
12/18/2008		X	X						
01/09/2009	X	X							
01/12/2009			X						
01/14/2009		X	X						
02/10/2009	X	X	X	X	X	X		X	X
02/17/2009	X	X	X		X	X		X	X
02/27/2009	X	X							
03/02/2009	X	X	X	X	X	X		X	X
03/05/2009		X	X		X	X			X
03/30/2009		X							
04/14/2009		X	X						
04/20/2009		X	X			X			
05/11/2009	X	X	X						
06/15/2009			X						
06/22/2009	X	X	X						
07/02/2009		X							
08/17/2009		X	X						
10/01/2009	X	X	X		X	X			
10/30/2009	X	X	X		X	X			
02/04/2010	X	X	X						
04/16/2010		X	X						
04/27/2010	X	X	X						
04/30/2010			X						
05/04/2010		X	X						
05/06/2010	X	X	X		X	X			X
05/14/2010		X	X						
05/20/2010		X	X		X	X			X
06/04/2010	X	X	X		X	X		X	X
06/22/2010			X						
06/29/2010			X						
No. of FTS days	64	74	69	22	33	29	9	11	14

This table presents flights to safety (FTS) over the period 2006-2010Q2. Flights are identified by a large positive return on the Treasury note and a large negative return on the S&P500 index. Light FTS episodes are based on a 1 standard deviation threshold, moderate FTS on a 1.5 standard deviation threshold and severe FTS on a 2 standard deviation threshold.

APPENDIX C

OVERVIEW OF THE EUROPEAN GOVERNMENT BOND MARKET

C.1 Market Overview

The analysis in this paper is based on intraday trade and quote data from MTS. This is the main electronic interdealer (IDB) trading platform for European fixed income securities, including government, quasi-government and covered bonds. MTS consists of multiple domestic platforms and a pan–European trading platform EuroMTS. Only those bonds that have attained the benchmark status are traded on EuroMTS in addition to their domestic platform. Nevertheless, most trading activity in a specific bond tends to occur on the relevant domestic platform (roughly 90% on average). Various estimates indicate that MTS has about 35-40% market share in terms of the daily number of trades across all IDB platforms in Europe, but the market share of MTS in a specific country can vary. Table C.1 below provides an overview of this market.

Table C.1: European Bond Market

		DE	FR	NL	BE	AT	FI	IT	ES	PT	GR	IE
No. Bonds	2006	51	129	33	46	14	11	60	113	20	22	5
	2010	61	186	58	52	16	11	65	140	20	28	13
	2011	63	182	63	67	16	11	74	148	19	26	11
	2012	65	176	68	78	17	13	74	145	20	25	11
Maturity	2006	6.64	9.40	7.36	5.93	7.90	4.04	6.92	10.36	5.34	4.95	6.48
	2010	6.17	10.39	6.65	5.43	8.61	5.17	6.49	9.69	5.62	6.63	5.92
	2011	5.96	10.25	6.95	6.31	8.18	5.98	6.33	9.53	5.42	6.54	6.20
	2012	6.21	10.03	7.76	6.71	7.68	7.06	5.87	9.18	4.59	15.18	5.74
Coupon	2006	3.60	1.41	3.17	1.96	4.19	4.02	3.21	0.95	3.12	4.30	3.47
	2010	2.73	0.94	3.32	1.75	4.01	3.01	3.11	0.89	2.95	4.26	3.89
	2011	2.57	0.93	3.15	1.46	4.02	3.50	3.16	0.89	3.15	4.24	4.19
	2012	2.39	0.87	3.01	1.22	3.88	2.85	3.20	0.89	2.83	2.42	4.46
Age	2006	4.36	8.26	3.83	5.33	4.78	5.40	3.18	5.35	3.71	4.53	4.76
	2010	4.09	6.53	4.74	4.76	5.28	4.84	3.56	5.60	3.62	4.31	2.95
	2011	4.16	7.05	4.81	5.88	5.73	3.51	3.68	5.79	3.86	5.03	4.52
	2012	4.31	7.66	4.79	6.62	5.96	3.83	3.87	6.21	3.85	1.64	5.09
Volume	2006	146	41	161	150	131	286	1462	45	491	207	182
	2010	74	50	196	289	48	85	875	67	247	58	85
	2011	55	69	258	198	73	231	744	47	40	6	22
	2012	46	68	216	242	60	175	501	38	23	1	9
No. Trades	2006	23	5	11	17	15	31	270	5	56	23	21
	2010	11	6	16	30	7	10	153	8	31	15	11
	2011	9	9	22	21	11	27	123	7	6	5	7
	2012	9	10	21	26	11	22	95	7	5	1	3
Ave. Issue (mil. euros)		11,909	11,884	14,118	12,072	10,198	5,658	13,694	12,836	7,315	5,573	8,173

The table reports the monthly average number of bonds by year and country, and average bond characteristics. *Volume* is the average monthly trading volume per bond in million euros. *No. Trades* is the average number of trades per bond. Figures for 2006 are based on data for October, November and December 2006. The sample includes benchmark sovereign bonds only. Bonds with less than 90–day to maturity are excluded. The resulting sample has 801 unique bond codes.

Figure C.1 shows trading activities by country. We observe that the bulk of trading activity occurs in Italian bonds, followed by trading activity in bonds issued by France, Netherlands, Spain, Belgium and Germany. We note in particular that trading activity on German government bonds is quite low. This is because these bonds are also traded actively elsewhere (German stock exchanges, numerous international electronic trading platforms and also in the over-the-counter (OTC) market), in contrast to bonds issued by other countries, which are mostly traded on MTS (see Cheung, de Jong and Rindi, 2005). The total trading volume in German Treasuries on MTS in 2012 is 52.7 billion euros, which is roughly 1% share of the total market trading volume of nearly 5,400 billion euros.[1] Nevertheless, the quality of German government bond prices is no doubt the best among fixed income securities in European capital markets, primarily due to a highly liquid futures market. The futures market is more liquid and larger than the cash market, ensuring fair market prices of German government bonds.

Figure C.2 plots trading activity throughout a trading day for each of the bond markets. Interestingly, the intraday trading pattern is quite similar across countries. The peak in trading activity occurs in the morning between 10:00 and 11:00. After 11:00, trading level remains quite stable through the rest of the trading day, except for a slight increase around 15:00.

C.2 Relevant Statistics for Six Select Countries

The analysis focuses on the six largest bond markets in the Euro-area, both in terms of the amount outstanding as well as daily trading volume. These countries also account for 87% of Euro-area total GDP. Table C.2 provides several important statistics for these countries for background information.

[1] http://www.deutsche-finanzagentur.de/en/institutional/secondary-market/

Table C.2: Country Statistics

Rating Downgrades 2010-2012			
	S&P	Moody's	Fitch
Belgium:	11/25/11	12/16/11	1/27/12
France:	1/13/12		
Italy:	9/11/11	10/5/11	10/7/11
	1/13/12	2/13/12	1/27/12
Spain:	4/28/10	9/30/10	5/28/10
	10/13/11	3/10/11	7/7/11
	1/13/12	10/18/11	1/27/12
	4/26/12	2/13/12	6/7/12

Sovereign Bond Outstanding (million euros)	
Belgium	364,815
France	1,365,452
Germany	1,116,223
Italy	1,638,724
Netherlands	341,759
Spain	688,231
6-Country Total	5,515,204
Euro-area Total	6,152,907
6-Country % Share	89.6%

Country GDP (million euros)	
Belgium	375,852
France	2,032,297
Germany	2,666,400
Italy	1,566,912
Netherlands	599,338
Spain	1,029,002
6-Country Total	8,269,800
Euro-area Total	9,483,173
6-Country % Share	87%

MTS Trading Volume 2010-2012 (million euros)	
Belgium	241,886
France	278,028
Germany	117,851
Italy	1,227,959
Netherlands	283,653
Spain	222,742
6-Country Total	2,372,118
Euro-area Total	2,535,709
6-Country % Share	94%

Figure C.1: Trading Activity By Country

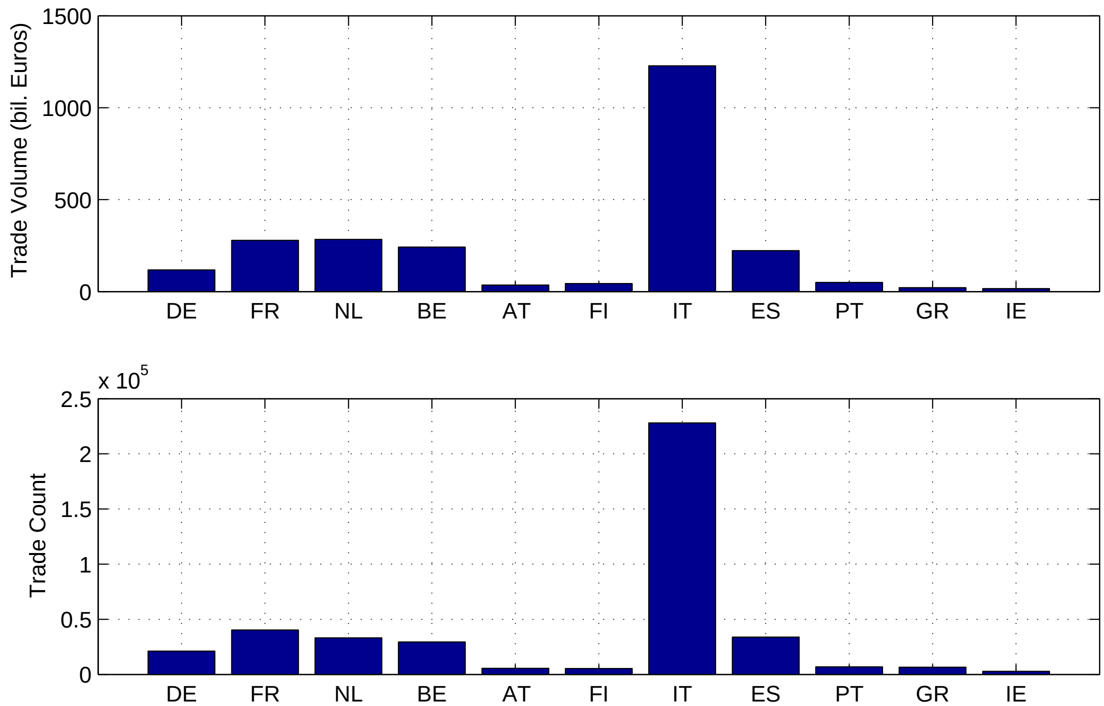

Source: MTS data 2010-2012. Figure shows trading volume by country, and reflects trading in benchmark bonds only.

Figure C.2: Intraday Pattern of Trading Activity

Source: MTS data 2010-2012. Figure shows trading volume by half hour interval throughout a trading day as a fraction of daily trading volume for each country, and reflects trading in benchmark bonds only.

REFERENCES

Ahn, H. J., Bae, K. H., and Chan, K. (2001). Limit orders, depth and volatility: Evidence from the Stock Exchange of Hong Kong. *Journal of Finance*, 56:767–788.

Aitken, M. J., Berkman, H., and Mak, D. (2001). The use of undisclosed limit orders on the australian stock exchange. *Journal of Banking and Finance*, 25:1589–1603.

Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108:1–33.

Alter, A. and Beyer, A. (2014). The dynamics of spillover effects during the european sovereign debt turmoil. *Journal of Banking and Finance*, 42:134–153.

Anand, A. and Weaver, D. G. (2004). Can order exposure be mandated? *Journal of Financial Markets*, 7:405–426.

Andersen, T., Bollerslev, T., Christoffersen, P. F., and Diebold, F. X. (2013). Financial risk measurement for financial risk management. In G. Constantinedes, M. H. and Stulz, R., editors, *Handbook of the Economics of Finance, Volume 2, Part B*, pages 1127–1220. Elsevier.

Andersen, T., Bollerslev, T., Diebold, F., and Labys, P. (1999). Realized volatility and correlation. Manuscript, Northwestern University, Duke University and University of Pennsylvania.

Andersen, T. G., Bollerslev, T., and Meddahi, N. (2011). Market microstructure noise and realized volatility forecasting. *Journal of Econometrics*, 160:220–234.

Arghyrou, M. G. and Kontonikas, A. (2012). The EMU sovereign-debt crisis: Fundamentals, expectations and contagion. *Journal of International Financial Markets, Institutions and Money*, 22:658 – 677.

Baele, L., Bekaert, G., and Inghelbrecht, K. (2010). The determinants of stock and bond return comovements. *Review of Financial Studies*, 23:2374–2428.

Baele, L., Bekaert, G., Inghelbrecht, K., and Wei, M. (2012). Flights to safety. Tilburg University, Columbia University, Ghent University and Federal Reserve Board of Governors Working Paper.

Bai, J., Julliard, C., and Yuan, K. (2012). Eurozone sovereign bond crisis: Liquidity or fundamental contagion. Federal Reserve Bank of New York and London School of Economics Working Paper.

Balduzzi, P., Elton, E. J., and Green, T. C. (2001). Economic news and bond prices: Evidence from the U.S. Treasury market. *Journal of Financial and Quantitative Analysis*, 36:523–543.

Bansal, N., Connolly, R. A., and Stivers, C. (2010). Regime-switching in stock index and treasury futures returns and measures of stock market stress. *Journal of Futures Markets*, 30:753–779.

Barclay, M. J., Hendershott, T., and Kotz, K. (2006). Automation versus intermediation: Evidence from treasuries going off the run. *Journal of Finance*, 61:2395–2414.

Barndorff-Nielsen, O. and Shephard, N. (2007). Variation, jumps, market frictions and high frequency data in financial econometrics. In Blundell, R., Newey, W., and Persson, T., editors, *Advances in Economics and Econometrics. Theory and Applications, Ninth World Congress - Volume 3*, pages 328–372. Cambridge University Press.

Baur, D. G. and Lucey, B. M. (2009). Flights and contagion – an empirical analysis of stock–bond correlations. *Journal of Financial Stability*, 5:339 – 352.

Beber, A., Brandt, M. W., and Kavajecz, K. A. (2009). Flight–to–quality or flight–to–liquidity? evidence from euro–area bond market. *Review of Financial Studies*, 22:925–957.

Bessembinder, H., Panayides, M. A., and Venkataraman, K. (2009). Hidden liquidity: An analysis of order exposure strategies in electronic stock markets. *Journal of Financial Economics*, 94:361–383.

Biais, B., Hillion, P., and Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance*, 50:1655–1689.

Boni, L. and Leach, C. (2004). Expandable limit order markets. *Journal of Financial Markets*, 7:145–185.

Boulatov, A. and George, T. J. (2013). Hidden and displayed liquidity in securities markets with informed liquidity providers. *Review of Financial Studies (Forthcoming)*.

Boyd, S. (2007). BNP Paribas freezes funds as loan losses roil markets. *Bloomberg News*.

Brandt, M. and Kavajecz, K. (2004). Price discovery in the u.s. treasury market: the impact of order flow and liquidity on the yield curve. *Journal of Finance*, 59.

Brandt, M. W., Kavajecz, K. A., and Underwood, S. E. (2007). Price discovery in the treasury futures market. *Journal of Futures Markets*, 27:1021–1051.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22:2201–2238.

Budish, E., Cramton, P., and Shim, J. (2014). The high-frequency trading arms race: Frequent batch auctions as a market design response. University of Chicago and University of Maryland Working Paper.

Buti, S. and Rindi, B. (2013). Undisclosed orders and optimal submission strategies in a limit order market. *Journal of Financial Economics*, 109:797–812.

Buti, S., Rindi, B., and Werner, I. M. (2011a). Dark pool trading strategies. University of Toronto, Bocconi University and Ohio State University Working Paper.

Buti, S., Rindi, B., and Werner, I. M. (2011b). Diving into dark pools. University of Toronto, Bocconi University and Ohio State University Working Paper.

Caballero, R. J. and Kurlat, P. D. (2008). Flight to quality and bailouts: Policy remarks and a literature review. MIT Department of Economics Working Paper No. 08–21.

Calice, G., Chen, J., and Williams, J. (2013). Liquidity spillovers in sovereign bond and cds markets: An analysis of the eurozone sovereign debt crisis. *Journal of Economic Behavior and Organization*, 85:122 – 143.

Campbell, B. and Hendry, S. (2007). Price discovery in canadian and u.s. 10-year government bond markets. Bank of Canada Working Paper 2007-43.

Cespa, G. and Foucault, T. (2012). Illiquidity contagion and liquidity crashes. Working Paper, Caas Business School and HEC Paris.

Chaboud, A., Chiquoine, B., Hjalmarsson, E., and Vega, C. (2013). Rise of the machines: Algorithmic trading in the foreign exchange market. *Journal of Finance (Forthcoming)*.

Cheung, Y., Rindi, B., and de Jong, F. (2005). Trading european sovereign bonds: the microstructure of the mts trading platforms. Working Paper.

Chordia, T., Sarkar, A., and Subrahmanyam, A. (2005). An empirical analysis of stock and bond market liquidity. *Review of Financial Studies*, 18:85–129.

Claeys, P. and Vasicek, B. (2014). Measuring bilateral spillover and testing contagion on sovereign bond markets in europe. European Central Bank Working Paper Series No. 1666.

Comerton-Forde, C. and Putnins, T. (2013). Dark trading and price discovery. University of Melbourne and University of Technology Sydney Working Paper.

Constancio, V. (2012). Contagion and the european debt crisis. *Banque de France Financial Stability Review*, 16:109–121.

Corradin, S. and Rodriguez-Moreno, M. (2014). Limits to arbitrage: Empirical evidence from euro-area sovereign bond markets. European Central Bank and Universidad de Navarra Working Paper.

Danielsson, J. and Payne, R. (2010). Liquidity determination in an order driven market. *European Journal of Finance* (Forthcoming).

Darbha, M. and Dufour, A. (2013). Microstructure of the euro-area government bond market. In Baker, H. K. and Kiymaz, H., editors, *Market Microstructure in Emerging and Developed Markets. Robert W. Kolb Series in Finance*. John Wiley & Sons, Inc.

De Grauwe, P. and Ji, Y. (2013). Self-fulfilling crises in the eurozone: An empirical test. *Journal of International Money and Finance*, 34:15–36.

De Santis, R. A. (2014). The euro-area sovereign debt crisis: Identifying flight-to-liquidity and the spillover mechanisms. *Journal of Empirical Finance*, 26:150–170.

De Winne, R. and D'Hondt, C. (2007a). Hidden liquidity in a pure order-driven market. *Finance Letters*, 5.

De Winne, R. and D'Hondt, C. (2007b). Hide-and-seek in the market: Placing and detecting hidden orders. *Review of Finance*, 11:663–692.

Degryse, H., de Jong, F., and van Kervel, V. (2014). The impact of dark trading and visible fragmentation on market quality. *Review of Finance (Forthcoming)*.

Diebold, F. X. and Yilmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. *Economic Journal*, 119:158–171.

Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, Forthcoming.

Dufour, A. and Engle, R. F. (2000). Time and the price impact of a trade. *Journal of Finance*, 55:2467–2498.

Dungey, M., Fry, R., Gonzlez-Hermosillo, B., and Martin, V. (2006). Contagion in international bond markets during the russian and the ltcm crises. *Journal of Financial Stability*, 2:1–27.

Dungey, M., Fry, R., Gonzlez-Hermosillo, B., and Martin, V. L. (2005). Empirical modelling of contagion: a review of methodologies. *Quantitative Finance*, 5:9–24.

Dungey, M., Henry, O., and McKenzie, M. (2013). Modeling trade duration in u.s. treasury markets. *Quantitative Finance (Forthcoming)*.

Dunne, P., Li, Y., and Sun, Z. (2011). Price discovery in the dual-platform u.s. treasury market. Central Bank of Ireland and Queen's University Belfast Working Paper.

Easley, D. and O'Hara, M. (1992). Time and the process of security price adjustment. *Journal of Finance*, 47:577–605.

Engle, R. (2002). New frontiers for ARCH models. *Journal of Applied Econometrics*, 17:425–446.

Engle, R., Fleming, M., Ghysels, E., and Nguyen, G. (2012a). Liquidity, volatility and flights to safety in the u.s. treasury market: Evidence from a new class of dynamic order book models. Federal Reserve Bank of New York Staff Reports, No. 590, December 2012.

Engle, R. and Gallo, G. (2006). A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics*, 131:3–27.

Engle, R., Gallo, G., and Velucchi, M. (2012b). Volatility spillovers in east asian financial markets: a mem–based approach. *Review of Economics and Statistics*, 94:222–233.

Engle, R. and Russell, J. (1998). Autoregressive Conditional Duration: A new model for irregularly spaced transaction data. *Econometrica*, 66:1127–1162.

Engle, R. F., Fleming, M. J., Ghysels, E., and Nguyen, G. (2012c). Liquidity, volatility, and flights to safety in the u.s. treasury market: Evidence from a new class of dynamic order book models. Federal Reserve Bank of New York Staff Report 590, December 2012.

Engle, R. F. and Ng, V. (1993). Measuring and testing the impact of news on volatility. *Journal of Finance*, 48:1749–1778.

Engle, R. F. and Patton, A. J. (2004). Impacts of trades in an error correction model of quote prices. *Journal of Financial Markets*, 7:1–25.

Escribano, A. and Pascual, R. (2006). Asymmetries in bid and ask responses to innovations in the trading process. *Empirical Economics*, 30:913–946.

Favero, C., Pagano, M., and von Thadden, E.-L. (2010). How does liquidity affect government bond yields. *Journal of Financial and Quantitative Analysis*, 45:107–134.

Fleckenstein, M., Longstaff, F. A., and Lustig, H. (2010). Why does the treasury issue tips? the tips-treasury bond puzzle. *Journal of Finance, Forthcoming*.

Fleming, M. (2003). Measuring treasury market liquidity. *Federal Reserve Bank of New York Economic Policy Review*, 9:83–108.

Fleming, M. J. (1997). The round-the-clock market for u.s. treasury securities. *Federal Reserve Bank of New York Economic Policy Review*, 3:9–32.

Fleming, M. J. (2002). Are larger treasury issued more liquid? evidence from bill reopenings. *Journal of Money, Credit, and Banking*, 34:707–735.

Fleming, M. J. and Mizrach, B. (2009). The microstructure of a U.S. Treasury ECN: The BrokerTec platform. Federal Reserve Bank of New York, Staff Report No. 381.

Fleming, M. J., Mizrach, B., and Nguyen, G. (2014). The microstructure of a U.S. Treasury ECN: The BrokerTec platform. Federal Reserve Bank of New York, Staff Report No. 381.

Fleming, M. J. and Nguyen, G. (2013). Order flow segmentation and the role of dark trading in the price discovery of u.s. treasuries. Federal Reserve Bank of New York Staff Report 624, August 2013.

Fleming, M. J. and Piazzesi, M. (2005). Monetary policy tick–by–tick. Working Paper, Federal Reserve Bank of New York.

Fleming, M. J. and Remolona, E. M. (1997). What moves the bond market? *Federal Reserve Bank of New York Economic Policy Review*, pages 31–50.

Fleming, M. J. and Remolona, E. M. (1999). Price formation and liquidity in the U.S. Treasury market: The response to public information. *Journal of Finance*, 54:1901–1915.

Forbes, K. and Rigobon, R. (2002). No contagion, only interdependence: Measuring stock market co-movements. *Journal of Finance*, 18:2223–2261.

Frey, S. and Sandas, P. (2012). Effective search for liquidity using iceberg orders. Working Paper.

Gao, C. and Mizrach, B. (2013). High frequency trading in the equity markets during u.s. treasury pomo. Rutgers University Working Paper.

Ghysels, E., Idier, J., Manganelli, S., and Vergote, O. (2013). A high frequency assessment of the ecb securities markets programme. University of North Carolina and ECB Working Paper.

Ghysels, E., Mykland, P., and Renault, E. (2010). In-sample asymptotics and across-sample efficiency gains for high frequency data statistics. University of North Carolina and University of Chicago Working Paper.

Ghysels, E. and Sinko, A. (2011). Volatility forecasting and microstructure noise. *Journal of Econometrics*, 160:257–271.

Goldreich, D., Hanke, B., and Nath, P. (2005). The price of future liquidity: Time–varying liquidity in the u.s. treasury market. *Review of Finance*, 9.

Goyenko, R. and Sarkissian, S. (2008). Flight-to-liquidity and global equity returns. McGill University Working Paper.

Green, T. C. (2004). Economic news and the impact of trading on bond prices. *Journal of Finance*, 59:1201–1233.

Handa, P. and Schwartz, R. A. (1996). Limit order trading. *Journal of Finance*, 51:1835–1861.

Härdle, W., Hautsch, N., and Mihoci, A. (2009). Modelling and forecasting liquidity supply using semi-parametric factor dynamics. SFB 649 Discussion Paper 2009-044, Humboldt-Universität zu Berlin, Germany.

Harris, L. E. (1996). Does a large minimum price variation encourage order exposure? University of Southern California Working Paper.

Harris, L. E. (1997). Order exposure and parasitic traders. University of Southern California Working Paper.

Hasbrouck, J. (1991a). Measuring the information content of stock trades. *Journal of Finance*, 46:179–207.

Hasbrouck, J. (1991b). The summary informativeness of stock trades: An econometric analysis. *Review of Financial Studies*, 4:571–595.

Hasbrouck, J. (2007). *Empirical Market Microstructure*. Oxford University Press.

Hasbrouck, J. and Saar, G. (2002). Limit orders and volatility in a hybrid market: The island ecn. Working Paper.

Hautsch, N. (2012). *Econometrics of Financial High–Frequency Data*. Springer.

Hautsch, N. and Huang, R. (2009). The market impact of a limit order. SFB 649 Discussion Paper 2009-051, Humboldt-Universität zu Berlin, Germany.

Hautsch, N. and Huang, R. (2012a). The market impact of a limit order.

Hautsch, N. and Huang, R. (2012b). On the dark side of the market: Identifying and analyzing hidden order placements. Humboldt-Universität zu Berlin Working Paper.

Hautsch, N. and Jeleskovic, V. (2008). Modelling high frequency volatility and liquidity using multiplicative error models. SFB 649 Discussion Paper 2008-047, Humboldt-Universität zu Berlin, Germany.

Hu, X., Pan, J., and Wang, J. (2011). Noise as information for illiquidity. Working Paper, MIT and University of Hong Kong.

Huang, R. D., Cai, J., and Wang, X. (2002). Information-based trading in the treasury note interdealer broker market. *Journal of Financial Intermediation*, 11:269–296.

Jiang, G., Lo, I., and Verdelhan, A. (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.

Kim, H. (2012). Generalized impulse response analysis: General or extreme. Auburn University Working Paper.

Klößner, S. and Wagner, S. (2012). Robustness and computation of spillover measures for financial asset returns and volatilities. Saarland University Working Paper.

Longstaff, F. (2004). The flight–to–liquidity premium in U.S. Treasury bond prices. *Journal of Business*, 77:511–526.

Luo, H. (2010). Profitable opportunities around macroeconomic announcements in the u.s. treasury market. Brock University Thesis.

Manganelli, S. (2000). Volume, volatility and the price impact of trades. PhD Dissertation, UCSD.

Menkveld, A. J., Cheung, Y. C., and de Jong, F. (2005). Euro-area sovereign yield dynamics: The role of order imbalance. ECB Working Paper No. 385, November 2005.

Mink, M. and de Haan, J. (2013). Contagion during the greek sovereign debt crisis. *Journal of International Money and Finance*, 34:102–113.

Mizrach, B. (2008). The next tick on nasdaq. *Quantitative Finance*, 8:19–40.

Mizrach, B. and Neely, C. (2006). The transition to electronic communication networks in the secondary Treasury market. *Federal Reserve Bank of St. Louis Review*, 88:527–541.

Mizrach, B. and Neely, C. (2007). The microstructure of the u.s.treasury market. Federal Reserve Bank of St.Louis Working Paper No.2007-052A.

Mizrach, B. and Neely, C. (2008). Information shares in the u.s. treasury market. *Journal of Banking and Finance*, 32:1221–1233.

Mizrach, B. and Neely, C. (2011). The microstructure of the u.s. treasury market. In Mizrach, B. and Dangerfield, B., editors, *The Encyclopedia of Complexity and System Science*, pages 863–877. Springer-Verlag.

Moinas, S. (2010). Hidden limit orders and liquidity in limit order markets. Toulouse University Working Paper, March 2010.

Musto, D., Nini, G., and Schwarz, K. (2011). Notes on bonds: Liquidity at all costs in the great recession. University of Pennsylvania Working Paper.

Næs, R. and Skjeltorp, J. (2006). Order book characteristics and the volume–volatility relation: Empirical evidence from a limit order market. *Journal of Financial Markets*, 9:408–432.

Nelson, D. B. (1991). Conditional heteroskedasticity in assets returns: A new approach. *Econometrica*, 59:347–370.

Newey, W. K. and West, K. D. (1987). A simple, positive semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55:703–708.

Pardo, A. and Pascual, R. (2012). On the hidden side of liquidity. *The European Journal of Finance*, 18:949–967.

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

Parlour, C. A. and Seppi, D. J. (2008). Limit order markets: A survey. In Thakor, A. V. and Boot, A., editors, *Handbook of Financial Intermediation and Banking*, pages 63–96. Elsevier Science.

Pasquariello, P. and Vega, C. (2007). Informed and strategic order flow in the bond markets. *Review of Financial Studies*, 20.

Pelizzon, L., Subrahmanyam, M. G., Tomio, D., and Uno, J. (2013). Sovereign credit risk, liquidity, and ecb intervention: Deus ex machina? Working Paper.

Pesaran, H. and Shin, Y. (1998). Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58:17–29.

Ranaldo, A. (2004). Order aggressiveness in limit order book markets. *Journal of Financial Markets*, 7:53–74.

Ready, M. J. (2012). Determinants of volume in dark pools. University of Wisconsin-Madison Working Paper.

Rodriguez, J. C. (2007). Measuring financial contagion: A copula approach. *Journal of Empirical Finance*, 14:401–423.

Russell, J. R. and Kim, T. (2010). A new model for limit order book dynamics. In Bollerslev, T., Russell, J. R., and Watson, M. W., editors, *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, pages 354–364. Oxford University Press.

Saar, G. (2001). Price impact asymmetry of block trades: An institutional trading explanation. *Review of Financial Studies*, 14:1153–1181.

Sarkar, A. and Schwartz, R. A. (2009). Market sidedness: Insights into motives for trade initiation. *Journal of Finance*, 64:375–423.

Sims, C. (1980). Macroeconomics and reality. *Econometrica*, 48:1–48.

Tuttle, L. (2006). Hidden orders, trading costs and information. Working Paper.

Vayanos, D. (2004). Flight to quality, flight to liquidity, and the pricing of risk. Working Paper 10327, National Bureau of Economic Research.

Ye, M. (2012). A glimpse into the dark: Price formation, transaction costs, and market share in the crossing network. UIUC Working Paper.

Zhu, H. (2014). Do dark pools harm price discovery? *Review of Financial Studies*, 27:747–789.