UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

ESSAYS IN MARKET MICROSTRUCTURE AND INFORMATION NETWORKS

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

JEFFREY R. BLACK
Norman, Oklahoma
2017

ESSAYS IN MARKET MICROSTRUCTURE AND INFORMATION NETWORKS

A DISSERTATION APPROVED FOR THE
MICHAEL F. PRICE COLLEGE OF BUSINESS

BY

____
Dr. Pradeep Yadav, Chair

____
Dr. Tor-Erik Bakke

____
Dr. Scott Linn

____
Dr. Duane Stock

____
Dr. Wayne Thomas

I dedicate this dissertation to my wife, Alicia Black; her unwavering love and support has made this work possible. I also dedicate this dissertation to my parents, Bryan and Tammy Black, who have always prioritized their lives to give me their utmost support, in this and all other ventures. Finally, I'd like to recognize a few others for keeping me grounded during my doctoral studies: my daughters, Audrey and Jacey, my brother Chris and his family, my wife's family in Oklahoma, as well as numerous other friends and family members. You all helped me immeasurably.

Acknowledgements

I'm extremely grateful to Pradeep Yadav for his guidance and mentorship throughout the doctoral program, and for chairing my dissertation committee. I'm appreciative to Tor-Erik Bakke, Seth Hoelscher, Scott Linn, Hamed Mahmudi, Duane Stock, and Pradeep Yadav for the opportunity to collaborate on research over the past five years. I would also like to thank Tor-Erik Bakke, Scott Linn, Duane Stock and Wayne Thomas for their commitment and advice as members of my dissertation committee. Finally, I'd like to thank Sriram Villupuram for sparking my interest in a PhD years ago, and helping to make it a reality.

Table of Contents

List of Tables

List of Figures

Abstract

This dissertation is a collection of three essays which study the impact of make-take fees on market efficiency, the relevance of three dimensions of liquidity on bond yields, and the value of information flow through the network of corporate directors.

In Chapter 1, I investigate the causal link between exchange-subsidized liquidity, in the form of make-take fees, and market efficiency. Using an exogenous experiment performed by NASDAQ in 2015, I employ difference-in-differences analysis on a matched sample and find that a decrease in take fee and make rebate levels leads to greater absolute pricing error and larger variance of mispricing. This stems from widened bid-ask spreads and decreased informed trading by retail investors. These findings demonstrate that make-take fees are beneficial for market efficiency.

There are three important dimensions of liquidity: trading costs, depth, and resiliency. Chapter 2 investigate the relevance of each of these three dimensions of liquidity – separately and in conjunction – for the pricing of corporate bonds. Unlike previous studies, this sample allows us to cleanly separate the default and non-default components of yield spreads. We find that each of the above three dimensions of liquidity are priced factors. Overall, in our sample, a one standard deviation change in trading costs, resiliency, and depth measures lead to a change in non-default spreads of 5.00 basis points, 2.27 basis points, and 1.27 basis points, respectively. We also find that both bond-specific and market-wide dimensions of liquidity are priced in non-default spreads. Finally, we find that there does exist in some periods a small residual non-default yield spread that is consistent with an additional "flight-to-extreme-liquidity" premium

reflecting investor preference for assets that enable quickest possible disengagement from the market when necessary.

Chapter 3 investigates the value of information flow through the network of corporate directors. More connected directors may have better information and more influence, which can increase firm value. However, directors with larger networks may also spread misleading or value-decreasing management practices. To identify the effect of director networks on firm value, we use the sudden deaths of well-connected directors as a shock to the director networks of interlocked directors. By looking at the announcement returns and using a difference-in-differences methodology, we find that this negative shock to director networks reduces firm value. This evidence suggests that director networks are valuable.

Chapter 1: The Impact of Make-Take Fees on Market Efficiency

1. Introduction

Informational efficiency in financial markets is of paramount interest to financial economists because efficient security prices result in a Pareto optimal allocation of capital, which contributes to economic growth. Over the past two decades, the informational efficiency of prices has increased dramatically due to market improvements such as decimalization, tick-size reduction, and increased institutional trading (Chordia et al., 2008, 2011). Over that same time period, stock markets in the US have become completely electronic. What's more, participation has become completely voluntary. Striking the old specialist model has resulted in a lack of affirmative obligations for traders to stand ready to make a market. As a result, exchanges have made numerous changes in order to cater to and incentivize liquidity providers. One such change has been the introduction of the "maker-taker" pricing model of market access fees, in which traders providing liquidity receive a rebate, and those consuming liquidity pay a fee. These make-take fees, largely propagated by Reg NMS (Regulation National Market System) in 2005, have quickly become one of the most debated aspects of market design (SEC, 2016)

Using the NASDAQ Access Fee Experiment in 2015 as a source of exogenous variation in make-take fees, I examine the resulting change in market efficiency. Utilizing propensity score matching and difference-in-differences (DiD) regressions, I find that an exogenous decrease in make-take fees causes a deterioration in market efficiency, namely an increase in absolute pricing error and an increase in the variance of mispricing. I propose that this effect propagates through the information channel by demonstrating that bid-ask spreads widen and fewer informative trades are executed. This suggests that make-take fees are beneficial for market efficiency.

That is, market makers widen their bid-ask spreads to compensate for the marginal loss of exchange rebates, as shown by Brolley and Malinova (2013), Malinova and Park (2015), and Anand et al. (2016), which increases transaction costs (particularly for traders without direct market access[1]). This increased cost of trading, in turn, discourages informed traders from trading if they indeed prefer the immediacy of marketable orders over the uncertainty of limit orders. This leads to a marginal decrease in the amount of information transmitted through trading. The decrease in information dissemination ultimately leads to a deterioration in market efficiency.

This finding is novel because it is not clear a priori that this should be the case. While Malinova and Park (2015) show that make-take fees improve liquidity by narrowing bid-ask spreads, it is not clear that this would automatically result in more efficient markets. While Chordia et al. (2008) finds a positive correlation between liquidity and market efficiency, Tetlock (2008) demonstrates that increased liquidity has an ambiguous effect on market efficiency. This is because lower transaction costs make trading more attractive, not just for informed traders, but also for noise traders, and the increase of noise trading could lead to more pricing error.

[1] Many traders, including retail and numerous institutional investors, do not have direct market access, but rather trade through a broker. This is an important point raised by Brolley and Malinova (2013) because, in practice, these traders have the bid and ask prices from the exchange passed on to them, but not the volume-based market access fee. They instead pay a flat fee. While this flat fee will also change in the long run, the authors show that in equilibrium, the higher flat fee does not offset the reduced bid-ask spread, and thus overall transaction costs reduce.

Separately, the lower cost of trading (lower take fees) on a public exchange like NASDAQ may encourage some orders – which would otherwise be routed to dark pools – to be routed to “lit” exchanges.[2] This would actually have a positive effect on information transmission, as more trades would be executed in the public eye, ultimately improving market efficiency.

In practice, it is quite possible that when the make-take fees are altered both the liquidity effect and the volume effect simultaneously impact the informational efficiency of a stock’s price, leading to conflicting ex ante hypotheses. In addition, Skjeltorp, Sojli, and Tham (2013) as well as Chung and Hrazdil (2010) conclude that the effect of make-take fees depends largely on the degree of adverse selection occurring in the market, further increasing the ambiguity of the theoretical impact of make-take fees on market efficiency.

In 2005, in hopes of further alleviating market fragmentation and creating a more cohesive national market for securities, the SEC (U.S. Securities and Exchange Commission) introduced Reg NMS, which established rules for market data, order protection (price priority), and market access fees, which were limited to $0.0030 per share traded (or hereafter referred to as 30¢ per 100 shares traded). The “maker-taker” pricing model of market access fees has since developed due to increased competition for trading volume between stock exchanges. In this model, traders with direct market access (DMA) – namely brokers and broker-dealers – are charged a “take” fee when removing liquidity from the market via marketable orders and given a “make” rebate when providing liquidity through limit orders. For example, for the majority of orders and

[2] In fact, this is what NASDAQ hypothesizes in its Dec. 12, 2014 SEC filing (NASDAQ, 2014).

stocks in 2015, NASDAQ had a 30/29 fee structure in place. This meant that for every 100 shares, a trader would be assessed a 30¢ fee for consuming liquidity and credited a 29¢ rebate for providing liquidity. Meanwhile, NASDAQ itself would keep the 1¢ difference.

In February of 2015, NASDAQ – in hopes of increasing its market share (NASDAQ, 2014) – experimentally lowered its fee structure from 30/29 to 5/4 on 14 stocks for a four month period.

To answer the empirical question “How do make-take fees affect market efficiency?” I exploit the changes to make-take fees during the 2015 NASDAQ Access Fee Experiment. Using propensity score matching on pre-shock variables to alleviate concerns of selection bias, I create a control sample for a baseline comparison to the 14 treated stocks before, during, and after the NASDAQ Access Fee Experiment using difference-in-differences (DiD) regression specifications. Following Menkveld (2013), Brogaard et al. (2014), and Fotak et al. (2014), I use a Kalman filter to estimate both the latent pricing error variable every minute as well as the latent variance of pricing error innovations parameter on a daily basis, while controlling for the bid-ask bounce. I find that when the NASDAQ reduces its access fee structure the treated stocks suffer an increase in mean absolute pricing error, as well as an increase in the variance of pricing error innovations, vis-à-vis the control stocks; pricing error of treated stocks increased to 2.51% from 2.34%, while pricing error of control stocks increased from 0.54% to 0.56%. This effect is in addition to, and cannot be explained simply by, widened bid-ask spreads during the experiment.

Furthermore, I show that during the access fee experiment, NASDAQ was less likely to possess the quote at the national best bid and/or the national best offer. I show that before and after the access fee experiment, the national best bid (ask) quote is on the NASDAQ 42.0% (42.26%) of the time. For stocks in the treatment group during the experiment, NASDAQ possessed the national best bid (ask) 19.9% (19.9%) less relative to the control stocks. I also find that the time NASDAQ possessed either/both of the NBBO (National Best Bid and Offer) quotes decreased during the access fee experiment, suggesting wider bid-ask spreads resulting from reduced make-take fees, at least on the NASDAQ. This is consistent with Malinova and Park (2015) who document that bid-ask spreads narrow with higher make-take fees.

Chordia et al. (2008) suggest that one reason liquidity and market efficiency may be associated is the increased incorporation of private information into market prices during more liquid regimes. To test this in regards to the access fee experiment, I examine the changes in adverse selection costs and find that for treatment group stocks, market makers lost less capital to informed traders during the lower make-take fee regime. This suggests that less private information was being incorporated into prices during this time, possibly explaining the steep increase in mispricing.

Finally, I consider the effect of the access fee experiment on trade volumes. While NASDAQ reports that they lost 1.5% market share on the treated stocks (relative to control stocks) during the low make-take fee regime, I find that NASDAQ lost 3.4% of the market share relative to the control group in this study. Delving further into the changes in volume, I find that volume on other exchanges actually increased during the experiment for treated stocks by 5.8% more than control stocks. However, the NASDAQ lost 11.5% of its volume during the experiment for treated stocks, relative to control stocks. This may suggest that the reduction in make-take fees actually did entice dark pool volume on to lit exchanges, however, brokers still routed marketable orders to the exchanges with higher rebates when possible. Clearly, more data is needed to make conclusions in this area. The potential SEC-proposed market-wide access fee pilot program may allow further research to be conducted along this vein.

Overall, the empirical analysis suggests that an exogenous decrease in make-take fees is detrimental to market efficiency. Taken in conjunction with the recent literature's claims that make-take fees are beneficial to market liquidity, it would not be altogether outlandish to conclude that make-take fees, while still highly debated, are actively improving market quality.

The remainder of the chapter is organized as follows. In Section 2, I give a brief review of the relevant literature. In Section 3, I describe the NASDAQ Access Fee Experiment, as well as describe the data, matching process, and research methods. In Section 4, I complete the empirical analysis of the NASDAQ Access Fee Experiment. Finally, Section 5 contains my concluding remarks.

2. Literature Review

Research in the area of make-take fees is still in its relative infancy. While analyzing new features of equity markets in the 21st century, Angel et al. (2011) recommend that the SEC either require access fees to pass through to end-users, stipulate that fees be included in the order protection (price priority) rule, or simply ban access fees outright.

They cite the increased agency costs between brokers and clients that arise from the maker-taker pricing model.

Since this recommendation, several theoretical works have analyzed aspects of market access fees. Colliard and Foucault (2012) emphasize the importance of distinguishing net fee and the breakdown between take fees and make rebates. They further show that an increase in net fee can either increase or decrease volume, based on several parameters. Empirically, Cardella et al. (2015) find that an exchange's trading volume is decreasing in its net access fee. Hence, one very important aspect of this paper is that the net fee is held constant in the NASDAQ Access Fee Experiment.

Skjeltorp et al. (2013) posit that make-take fees actually create a positive liquidity externality unless adverse selection is sufficiently high, in which case make-take fees may actually cause a negative liquidity externality because market makers are averse to trading opposite informed traders.

After recognizing that the choice between market and limit orders arises from a trader's inherent value of speed, Foucault et al. (2013) endogenize the demand for speed, and propose a model which shows that the breakdown of make and take fees becomes economically significant when the minimum tick size restricts bid-ask spread adjustment. Since markets now permit trades to be executed up to 4 decimal places, and access fees are in the 3 to 4 decimal range, this is of less concern, at least in US markets.

However, Brolley and Malinova (2013) show that because make-take fees are not, in practice, passed through brokers to end-use traders[3], make-take fees should improve

[3] Rather than make-take fees, brokers typically charge a flat trade fee to its customers, which include retail traders, as well as other institutional traders not structured as brokers with direct market access. While these flat fees may increase in the long run, Brolley and Malinova (2013) show that in equilibrium, the increased fee does not offset the narrowed bid-ask spread, therefore overall transaction costs reduce.

market quality in numerous aspects - lowering transaction costs, increasing trading volume, and improving welfare. These theoretical predictions are confirmed by Malinova and Park (2015). They use a change in make-take fees on the Toronto Stock Exchange to show that raw bid-ask spreads improve, but after adjusting for the fees, total transaction costs remain unaffected. However, since access fees are not passed to non-DMA traders, overall liquidity improves.

Similarly, Anand et al. (2016) find that overall execution costs for liquidity demanders decline following the introduction of the make-take fee structure in options markets, consistent with increased quote competition. Lutat (2010), on the other hand, finds that spreads aren't affected by make-take fees, but depth at the best bid and ask quotes improves. Also related, Battalio et al. (2016) document a negative relationship between limit order execution and rebate/fee levels.

This study builds upon Malinova and Park (2015) and Anand et al. (2016) by showing that not only are bid-ask spreads improved by higher make-take fees, but that make-take fees also reduce mispricing. I show that this occurs through the information channel. Interestingly, Malinova and Park (2015) find that adverse selection costs actually decline as make-take fees increase while I show that adverse selection costs decline as make-take fees *decrease*.

This paper is also related to the vein of literature which relates liquidity to market efficiency. Particularly, Chordia et al. (2008) find a positive correlation between liquidity and market efficiency. While they hypothesize that this could be because liquidity stimulates greater arbitrage activity, enhancing market efficiency, no causal evidence is explored. Chung and Hrazdil (2010) reinforce these findings, also documenting that the liquidity-efficiency relationship is amplified when adverse selection spread is higher (more informative trading).

While this paper does not provide direct evidence of a causal relationship between liquidity and informational efficiency due to confounding economic mechanisms, the results of this paper are consistent with both Chordia et al. (2008) and Chung and Hrazdil (2010), as I find informational efficiency is improved by make-take fees because transaction costs are reduced and trades become more informative.

This study also increases our understanding of the liquidity-efficiency relationship because even during a decrease in make-take fee level, and subsequent widening of the bid-ask spread, I document an increase in trading volume, suggesting that liquidity is defined by more than just transaction costs. Additionally, whereas Chordia et al. (2008) define market efficiency as the inverse of short-horizon return predictability from order flows, I define market efficiency following Fama and French (1988) and others, by removing the random walk component of intraday stock prices (specifically NBBO midpoints) to find mispricing, and the variance thereof.

3. Sample and Methodology

3.1. NASDAQ Access Fee Experiment

In November of 2014, NASDAQ announced its intention to change its make-take fee structure for select stocks in order to analyze the changes' effect on market share, displayed liquidity, effective spreads, and volatility. In NASDAQ's filing with the SEC the exchange stated that it believed take fees had grown to a level which was discouraging certain traders from directing their trades to one of the 14 "lit" exchanges, opting instead to trade in dark pools. NASDAQ hypothesized that by reducing its take fees and make rebates that it would be able to increase its market share. In the SEC filing, the exchange requested permission to experimentally change its market access fee structure to charge a $0.0005 fee per share to remove liquidity (from $0.0030), and to credit a $0.0004 rebate per share to add displayed liquidity (from $0.0029)[4] (NASDAQ, 2014).

Late in 2014, NASDAQ announced the 14 stocks included in its access fee experiment: American Airlines (AAL), Micron Technology (MU), FireEye (FEYE), GoPro (GPRO), Groupon (GRPN), Sirius XM (SIRI), Zynga (ZNGA), Bank of America (BAC), General Electric (GE), Kinder Morgan (KMI), Rite Aid (RAD), Transocean (RIG), Sprint (S), and Twitter (TWTR). The NASDAQ Access Fee Experiment commenced on February 2, 2015 and was set to run for a term of four months, though wording seemed to indicate that an extended timeframe would be possible if it was deemed valuable later on.

Throughout the course of the experiment, NASDAQ reported on various aspects of the market for these stocks. The exchange had seen a small uptick in liquidity consumption (marketable orders), but it was not offset by the major losses in liquidity provision (executed limit orders), "time at the inside" of the NBBO, and market share. Therefore, NASDAQ elected to cease the experiment after the initial four month term.

[4] There were further alterations to the fee structure which included rebates for non-displayed liquidity, non-displayed midpoint liquidity, and several other obscure order types, but in general, the make-take fee structure was reduced by 25¢ per 100 shares. Also important to the validity of the results in this study, the net fee remained 1¢.

3.2. Data and Research Design

The majority of the data used in this study is collected from NYSE TAQ (Trade and Quotation). The data represents the consolidated tape, which covers virtually all trades and quotes on the 14 U.S. public stock exchanges. The NASDAQ experiment ran for four months, from February through May of 2015. In order to obtain a baseline sample outside of the experiment period, I collect TAQ data on all stocks spanning October 2014 through September 2015. For each day, I calculate the total volume, volume on the NASDAQ, price, NASDAQ market share, percentage bid-ask spread, dollar bid-ask spread, adverse selection cost, the percentage of time NASDAQ spent on the inside of the NBBO, and multiple mispricing measures.

Specifically, I calculate volume as the sum of the trade size (in shares) for all trades on every exchange. Similarly, NASDAQ volume is the sum of the shares traded on the NASDAQ exchange. I divide the NASDAQ volume by the total volume to measure NASDAQ's market share. In later regressions, use the natural logarithm of both of the volume variables to address the skewness of the distribution of volumes (because volume has a lower bound of 0, volume is positively skewed). Actual transaction prices are often executed at the bid or ask price. Therefore, to eliminate the bid-ask bounce, I use a volume-weighted average of the NBBO midpoint of each trade as a proxy for stock price. Similarly, I calculate the daily percentage bid ask spread as the mean of the NBBO bid-ask spread scaled by the midpoint taken after each new quote. I multiply this by the daily midpoint to calculate the dollar bid-ask spread[5].

[5] A differentiation between percentage and dollar bid-ask spreads is important because the make-take fees are "volume-based," rather than "value-based," which means the fees and rebates are assessed on a dollars-per-share basis, rather than a percent-of-value basis, therefore the fee structure will affect the bid-ask spreads, and ultimately the market efficiency, of low and high priced shares differentially.

In order to estimate the level of informed trading, I calculate adverse selection costs, which represent the money that market makers lose to informed traders on average. In order to calculate this, I first sign the trades using the Lee and Ready (1991) algorithm. I calculate adverse selection costs ASC_k as

$$ASC_k = \frac{1}{T}\sum_{t=1}^{T}\frac{d_t(m_{t+k} - m_t)}{m_t}, \quad (1)$$

where m_t is the NBBO midpoint at trade t, d_t equals 1 for a buy and -1 for a sell (according to the Lee and Ready (1991) algorithm), T is the number of trades in a given day, and k is the number of minutes after the initial trade. I calculate the adverse selection spread using a k of 1, 15, 30, and 60 minutes.

In order to calculate the amount of time the NASDAQ has a quote on the inside of the NBBO, I first create two binary variables at the quote level: one which equals 1 when the national best bid is located on the NASDAQ exchange (and 0 otherwise), and one which equals 1 when the national best ask is located on the NASDAQ exchange. From these, I create two more quote-level binary variables: one which equals 1 when the NASDAQ has the best bid *and* offer (represented mathematically as $bestbid \times bestask$), and another which equals 1 when the NASDAQ has *either* the best bid or ask (represented mathematically as $max(bestbid,bestask)$). Next, I calculate the daily averages of these four binary variables (Best Bid, Best Ask, Best Both, and Best Either) to find the amount of quote-time the NASDAQ is at the inside of the NBBO.

Later, in order to construct a proper control sample, I collect industry (NAICS and SIC) and listing exchange data from CRSP over the same time frame. All continuous variables were then winsorized at the 1st and 99th percentiles.

3.2.1 Market Efficiency Measures

Because a stock's true fundamental value, and therefore pricing error cannot be directly observed, further assumptions must be made to estimate fundamental value and pricing error. Since at least Fama (1965) the random walk model, or the weak form of market efficiency, has been widely accepted in the financial economics literature. Following Fama and French (1988) and Hasbrouck (1993), I assume that the logarithm of a stock's observed transaction price p_t follows the equation

$$p_t = f_t + s_t \tag{2}$$

where s_t is the pricing error of the stock on day t, and the stock's fundamental value, f_t, follows a random walk with a drift μ, and white noise innovation ε_t,

$$f_t = \mu + f_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim N(0, \sigma_\varepsilon^2). \tag{3}$$

If pricing error is assumed to follow mean-reverting process

$$\Delta s_t = -\alpha s_{t-1} + \phi_t, \qquad \phi_t \sim N(0, \sigma_\phi^2) \tag{4}$$

with mean reversion parameter α and white noise innovation ϕ_t, then combining equations (2), (3), and (4), we get:

$$p_t = \mu + (1-\alpha)p_{t-1} + \alpha f_{t-1} + \theta_t, \qquad \text{where } \theta_t = \varepsilon_t + \phi_t. \tag{5}$$

Following Boehmer and Kelly (2009), Menkveld (2013), Brogaard et al. (2014), and Fotak et al. (2014), I use Kalman filter estimation methodology with the transition equation:

$$\begin{bmatrix} p_t \\ f_t \\ 1 \end{bmatrix} = \begin{bmatrix} 1-\alpha & \alpha & \mu \\ 0 & 1 & \mu \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} p_{t-1} \\ f_{t-1} \\ 1 \end{bmatrix} + \begin{bmatrix} \theta_t \\ \varepsilon_t \\ 0 \end{bmatrix}, \tag{6}$$

and the measurement equation:

$$p_t = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} p_t \\ f_t \\ 1 \end{bmatrix}, \tag{7}$$

Elsewhere in the economics literature, the Kalman filter is used to observe otherwise latent variables, i.e. mispricing, from observable variables, given an assumed structure. In this case, I am removing the random-walk component from stock midpoint prices to observe mispricing. If a time-series was indeed a random walk, I would find no mispricing[6].

Because the actual transaction price tends to "bounce" due to the bid-ask spread, and Malinova and Park (2015) show that make-take fees directly affect the bid-ask spread, I use the log of the midpoint of the bid-ask spread as a proxy for log price, p_t, in this estimation. Omitting the first 5 minutes of trading to eliminate the opening noise, I collect p_t from each stock at every minute from 9:35 to 16:00 over the entire sample. Using BFGS maximum likelihood optimization, I obtain estimates of μ, α, σ_ϕ^2, and σ_ε^2 for every stock, each day. I further calculate the mean absolute pricing error (MAPE) each day by averaging $|s_t|$ over the day. I use σ_ϕ^2 to measure the variance of pricing error innovations on each day. These two variables were also winsorized at the 1st and 99th percentiles.

3.2.2 Matching Procedure

In a laboratory setting, the treatment and control groups would be randomly selected. However, in the NASDAQ experiment, the 14 treated stocks were said to be chosen based on the estimated proportion of "off-exchange" (dark pool) trading (NASDAQ, 2014). We also know that the effect of a make-take fee reduction will differ based on the share price of a stock since make-take fees are based on shares traded, and not dollar value. Therefore, to assuage concerns over selection bias, I created a matched control sample of

[6] For more information on the Kalman filter smoothing-estimation procedure, see chapter 13 of Hamilton (1994).

firms using nearest neighbor propensity score matching. To measure the propensity of being selected into the NASDAQ Access Fee Experiment, I employ a probit model using pre-experiment data (from October 2014 through January 2015).

I omit stocks which have less than 80 trading days in the 4 month window (out of 84 possible), have a missing or unclassifiable industry (SIC 9999), are listed as Financial Vehicles (NAICS 525990), or are missing the outcome variables in the study (pricing error, adverse selection spread, volume, etc.). This results in 7,573 stocks (14 of which are treated). I then take the 4-month average of the variables in the propensity score model to estimate the probit model on a firm level.

Ideally, the proportion of off-exchange volume would be a key component to the matching process, however, since dark pools do not disseminate any comprehensive transaction data, I instead substitute the NASDAQ volume into the model. I also match on dollar bid-ask spread and midpoint price, as well as average MAPE, to attempt to get the control sample close to the treatment sample on these pre-shock characteristics. The results of the probit model, along with marginal effects are displayed in Table 1.

Next, to select the control sample, for each of the 14 treated stocks I eliminate as potential matches untreated stocks which are not listed on either the NASDAQ or NYSE, or with a difference in average price greater than 10 percent. I then select the 5 untreated stocks with the closest propensity score, based on the above probit, allowing for replacement. This creates a control sample of 70 stocks – 63 of which are unique. The means and medians of the pre-shock variables of interest for treatment and control groups are displayed in Table 2.

The sample means are statistically different at the 10% significance level or above for each variable in the table, except price. This is partially due to the large sample sizes reducing the standard errors, as the means are usually economically very similar. However, we do see some economically significant differences in a few key variables, namely the adverse selection costs, MAPE, variance of mispricing, and the dollar bid-ask spread. While this raises a concern that perhaps the control and treatment samples will behave differently, violating the parallel trend assumption of DiD analysis, the difference in means is controlled for by the binary *Treated Dummy* variable in the DiD regressions, therefore, the difference in means are only problematic if they suggest that the treatment stocks and control stocks will behave differently. We also see that the medians are economically similar, suggesting that the statistical differences are driven by a few poorly-matched outliers. Still, to assuage concerns that the results may be driven by these differences in control and treated samples, I create alternative control samples.

4. Empirical Results

4.1 Main Results

To determine the effect of make-take fees on market efficiency, I regress the two (inverse) measures of market efficiency, MAPE and variance of pricing error innovations (σ_{ϕ}^{2}) in panel DiD regressions. I include a binary *Treated Dummy* variable, which equals 1 if the stock is 1 of the 14 in the NASDAQ Access Fee Experiment and 0 if it is in the control stock, a binary *Experiment Dummy* variable, which equals 1 if the lower access fees were in effect (Feb. 2015 – May 2015) and 0 for all other dates, and finally, an interaction of the two, which will produce the regression coefficient representing the difference in differences in the means. To correct the standard errors for autocorrelation and heteroskedasticity between stocks, I use robust standard errors clustered two-ways, by day and stock, as suggested by Pedersen (2009). The results of these regressions are represented in Table 3.

In the first regression in Panel A, we see that there is no statistical difference in MAPE between treatment and control samples before and after the access fee experiment, however, when NASDAQ lowers its access fees to the 5/4 structure, the MAPE increases by 0.17% of the stock price[7]. When I include linear controls for the price, percentage bid-ask spread, and volume, this result holds. Relative to the control stocks, the mean absolute pricing error of treated stocks increased 0.17%, roughly \$0.0422 taken at the mean – far in excess of the \$0.0025 difference in access fees.

When regressing the variance of mispricing innovations in the DiD model, I find that the access fee experiment increased the variance 0.0046 for treated stocks relative to control stocks (an increase of 6.78% in standard deviation). When including controls in the regression model, the DiD effect remains at 0.0049, with strong statistical significance.

It is possible that somehow trader behavior changed after the access fee experiment, even though the make-take fees returned to the previous price structure. To ensure this wasn't driving the DiD results, Panel A of Table 4 includes regressions which exclude observations after May 2015, the end of the NASDAQ experiment. When excluding these observations, the DiD results remain largely unchanged. In fact, the DiD coefficients on

[7] Since p_t is the log of the midpoint price, $|s_t|$, and therefore MAPE, are also measured in the same units.

the MAPE and variance regressions increase relative to the full sample. Therefore, we can conclude that the mispricing increases when make-take fees are reduced.

As suggested by Roberts and Whited (2013), if the mispricing effect is truly driven by the reduction in make-take fees, then when the NASDAQ restored to the original fee structure, the effect of mispricing should reverse. I test this reversal effect in Panel B of Table 4. Indeed, we see that when the fee structure is raised, MAPE decreases 0.08% (0.12% when including control variables). However, we also see that the variance of mispricing increased when the fee structure was reversed, though lacking statistical significance.

Subsequently, I confirm the findings of Malinova and Park (2015) that make-take fees affect bid-ask spreads. However, where their sample involves the entire market switching to a make-take fee structure, this sample involves only one exchange, NASDAQ, changing its make-take fees. Consequently, since the other 13 public exchange did not change their fee structure, the bid-ask spread may not have necessarily altered since other exchanges had make rebates as high as 24¢ and take fees as low as -6¢ (a rebate of 6¢)[8]. Therefore I instead examine the amount of time that NASDAQ quotes are at the inside of the NBBO. I use the same DiD regression framework as in the previous tests, with the constructed *Best Ask*, *Best Bid*, *Best Either*, and *Best Both* proportion variables.

The results of these regressions are displayed in Table 5. It's important to note, that the time at the inside of the NBBO did not differ between treatment and control groups

[8] Inverted "taker-maker" pricing structures are located on the NASDAQ Boston Exchange and Direct Edge EDGA Exchange. This pricing structure has been argued to be an alternative to dark pools since it is cheap (due to the rebate) to execute trades taking liquidity on these exchanges. This structure has also been promulgated as a hot bed for predatory high frequency traders who are essentially being charged a "make" fee by the exchange to have access to the information provided by these fee-sensitive traders.

prior to the experiment, according to the *Treated Dummy* coefficient. But when NASDAQ lowered its fee structure, we see a significant differential reduction in time at the inside of the NBBO for NASDAQ quotes. For treated stocks, NASDAQ spends 19.9%, 19.9%, 27.2%, and 12.4% less time possessing the best bid, best ask, either the best bid or ask, and both the best bid and ask, respectively. As Brolley and Malinova (2013) would posit, the decrease in rebates for limit orders results in less time inside the NBBO as market makers submitting orders on the NASDAQ widen their spreads to compensate. These results hold in Panel B, when I include *price* and *volume* and linear controls. I don't include the bid-ask spread as a control, since it is mechanically related to the NASDAQ's time inside the NBBO.

Because market makers are posting limit orders in a way that is widening the bid-ask spread – due to the smaller rebates – one would expect less information to be transmitted by retail and other non-DMA traders due to the marginal increase to transaction costs. Therefore, I examine the changes in adverse selection costs to determine how the make-take fee reduction affected the incorporation of private information into security prices. The DiD analysis is presented in Table 6, without controls in Panel A and with controls in Panel B. When examining the average 1-minute, 15-minute, and 30-minute profits of liquidity takers (ASC_1, ASC_{15}, and ASC_{30}) we see a significant decrease for treated stocks, relative to the control stocks, during the experiment. We also see a decrease in 60-minute profits of liquidity takers, however the difference is not statistically significant. This evidence, in conjunction with prior results, suggests that the widened bid-ask spreads resulting from the make-take fee reduction led to fewer informative trades being made by non-DMA traders, which made markets less informationally efficient.

NASDAQ hypothesized that by lowering their make-take fees, they would garner more volume onto the exchange by luring away “off-exchange” trading to the NASDAQ. This increase in volume could also arguable lead to an increase in market efficiency – because more trades take place on lit exchanges, leading to more incorporation of private information into stock prices. We can see from the above results that this was not the case, but that doesn’t necessarily preclude that the treatment stocks experienced an increase in volume.

In Table 7, I investigate the changes in volume propagated by the access fee experiment. Looking at the DiD interaction coefficient, I find that NASDAQ’s market share reduced by 3.4% relative to the control sample during the experiment – a somewhat higher loss than the 1.5% NASDAQ estimated against its own control sample. Along these lines, I find that the volume on the NASDAQ (for these 14 treated stocks) reduced by 11.5% relative to control stocks, and the volume of the treated stocks actually increased by an average of 5.8% on all other exchanges, in comparison with control stocks. This shows that volume transitioned from the NASDAQ and on to other exchanges due solely to the level of make-take fees.

These results suggests that while NASDAQ did not benefit from the reduction in make-take fees, other exchanges did. This is quite possible if market makers were still apt to submit limit orders to the exchanges with the higher rebates (for example, BATS, with a 25/24 fee structure), and the reduction from 30¢ to 25¢ to execute these orders was enough to entice more traders to take liquidity, whether they would have otherwise not traded or traded in dark pools. While I’ve shown that this increase in volume did not result in an overall increase in trade informativeness, or a decrease in pricing error, without a more complex structural model, the volume effect of make-take fees on market efficiency is not possible to estimate, due to make-take fees' simultaneous effect on the bid-ask spread. However, these results suggest that a market-wide experiment on make-take fees – such as the SEC proposed in the summer of 2016 (SEC, 2016) – may lend valuable data to extend this line of research.

4.2 Robustness Tests

Roberts and Whited (2013) suggest falsification (placebo) tests as one method to further demonstrate that the natural experiment is truly the driving force in variable changes. To perform falsification tests, I create dummy variables for stock-day observations in November, December, and January, and subsequently interact these transactions with the dummy for treated stocks. Using a sample of the four months prior to the NASDAQ pilot, I regress MAPE and the variance of mispricing on these variables in Table 8. Examining the regression coefficients on the interaction terms (of the month and treated dummy variables), we see that there was very little difference in the outcome variables of treated and control stocks between October, November, December, and January. We see a slight increase in MAPE and the variance of mispricing for the treated stocks in November, however the statistical significance is very slight (with p-values between 0.07 – 0.17). When we consider that the pricing error was significantly higher (2.45% vs 2.18%) for treated stocks during the NASDAQ pilot than in November, we can conclude that it is likely not a coincidental time trend which was behind the increase in mispricing, but instead that it was the market access fee pilot.

Due to the statistical differences in treatment and control groups, I create four alternate control samples to demonstrate the robustness of results. To construct each of these samples, I first eliminate as potential matches untreated stocks which have less than 80 trading days in the 4 month window (out of 84 possible), have a missing or unclassifiable industry (SIC 9999), are listed as Financial Vehicles (NAICS 525990), are missing the outcome variables in the study (pricing error, adverse selection spread, volume, etc.), which are not listed on either the NASDAQ or NYSE, or with a difference in average price greater than 10 percent. Next, I match on one variable at a time: MAPE, NASDAQ volume share, dollar bid-ask spread, or the percent of NBBO quotes coming from the NASDAQ.

Next, I repeat the regressions from Table 3 with MAPE as the dependent variable using the four alternate samples. In all four alternate samples, we see that pricing error increased during the access fee pilot, with magnitudes ranging from 0.14% to 0.36%. We also see that in all of the alternate samples, there was no statistical difference in MAPE between treated and control samples. This further assuages any concerns about violating the parallel trend assumption, and demonstrating that lower make-take fees increased pricing error.

5. Concluding Remarks

While market efficiency is a paramount assumption in markets, directly affecting every trader and investor, the effect of market access fee level on market efficiency has not been addressed in the literature until this paper. While *ex ante* predictions range from an increase in efficiency via subsidized liquidity to a decrease in efficiency due to the prohibitive costs of trading using market orders, I find that a decrease in take fees and make rebates causes greater absolute pricing error and larger variance of mispricing, stemming from the widened bid-ask spreads and decreased informed trading by retail investors and other traders without direct market access. This suggests that higher levels of make-take fees lead to greater market efficiency. However, further research may be necessary to document the effect of make-take fees on dark pool trading volume and the market share of lit exchanges.

Chapter 2: The pricing of different dimensions of liquidity: Evidence from government guaranteed bonds[9]

1. Introduction

Bank liabilities are often insured selectively by government programs of different countries.[10] The empirical analysis in this chapter has been made possible by one such program: the U.S. government's Debt Guarantee Program of 2008. In an attempt to stem bank contagion risk during the 2008 financial crisis, the FDIC instituted a program wherein bank-issued bonds were backed by the full faith and credit of the U.S. government, and thus made equivalent in credit quality to U.S. Treasury securities. While these bonds were as safe as Treasuries from a default perspective, they differed significantly from Treasuries, and from each other, in their liquidity. Thus, these bonds impounded a yield spread over comparable Treasuries that was arguably a significant function of liquidity, but independent of any default-related considerations. We use this unique situation to analyze how different dimensions of liquidity affect the pricing of corporate bonds: specifically, bonds issued by banks.

The yield spreads of corporate bonds (relative to Treasuries) have been shown by Elton et al. (2001), among others, to be significantly larger than can be explained by default risk and state taxes. Chief among the factors shown to affect non-default spreads is liquidity. For example, Longstaff et al. (2005) and Dick-Nielsen et al. (2012) show that

[9] This chapter is based on collaborative work with Duane Stock and Pradeep K. Yadav, published in the Journal of Banking and Finance.

[10] A common example is deposit insurance where, in the United States, the insuring agency is the Federal Deposit Insurance Agency (FDIC).

an important dimension of liquidity – the trading cost dimension as measured by the bid-ask spread – is priced in the non-default component of yield spreads. The focus of this chapter is on the relative pricing relevance of different dimensions of liquidity. In this context, the early seminal literature in market microstructure – Garbade, 1982; Kyle, 1985, and Harris, 1990; Harris, 2003 – identifies three main dimensions of liquidity: the trading cost dimension, the tradable quantity or the depth dimension, and the time dimension as manifested in the resiliency in liquidity subsequent to order-flow shocks.[11] In this chapter, our main aim is to investigate whether these three different dimensions of liquidity are priced in government-guaranteed bank bond yields, estimate the relative importance of each of these liquidity dimensions for pricing, and determine the comparative pricing relevance of bond-specific and market-wide dimensions of liquidity.

Unlike previous studies, our sample allows us to cleanly separate the default and non-default components of yield spreads. We are accordingly able to contribute significantly to the extant literature on the pricing of liquidity in fixed income markets in several important ways. We are the first to examine whether the resiliency dimension of liquidity is priced in bond yields. Second, we are also the first to test whether the aforementioned three dimensions of liquidity – trading costs, depth, and resiliency – are priced *in conjunction*, as opposed to being priced separately. Third, an important methodological contribution we make is to use the principles underlying the empirical measure of resiliency developed (for limit-order-book markets) by Kempf et al. (2015) to

[11] Holden et al. (2014) provide an excellent review of the empirical literature on liquidity. Specifically focusing on the three dimensions of liquidity mentioned above, see, for example: (a) Glosten and Milgrom (1985) and Stoll (1989) for the trading cost dimension; (b) Kyle (1985); Glosten and Harris (1988); Hasbrouck (1991), and Kempf and Korn (1999) for the depth dimension; and (c) Foucault et al. (2005), and Kempf et al. (2015) for the resiliency dimension.

define and estimate a new measure for the resiliency of over-the-counter dealer markets (like corporate bond markets). Fourth, we analyze the relative pricing relevance of both bond-specific and market-wide dimensions of liquidity. Finally, we examine whether there exists – after controlling for (bond-specific and market-wide) trading cost, resiliency, and depth dimensions of liquidity – a "residual" non-default yield spread arising from an additional flight-to-extreme-liquidity premium over Treasuries, or a "quality" spread related in some way to the probability of the government guarantee being invoked.

Because the liquidity risk of a security and the default risk of a firm are endogenously related, separating the two is problematic and involves measurement error. Intuitively, and according to Ericcson and Renault (2006), among others, increases in default risk lead to increases in liquidity risk. Interestingly, He and Xiong (2012) and He and Milbradt (2014) theorize that the inverse is also true – that deterioration in liquidity leads to increases in default risk.[12] In this study, we use a sample of bonds in which liquidity risk is exogenously separated from default risk, since the sample bonds do not carry any default risk above that of US Treasury bonds. This allows us to analyze the non-default component of the yield spread (hereafter "non-default spread" or "NDS") without the potential for measurement error induced by using models for the default spread, as has been done in earlier studies. The absence of measurement error in our sample allows us to cleanly and accurately determine the magnitude of the non-default spread, and relate it to different dimensions of liquidity.

[12] This occurs because firms sustain higher losses when rolling over maturing debt, which make strategic default more likely.

We calculate bid-ask spread to proxy for the trading cost dimension of liquidity following the methods of Hong and Warga (2000). To measure the depth dimension of liquidity we use the Amihud (2002) illiquidity measure, which is a direct measure of the price impact of trading volume, and consistent with Kyle (1985). Finally, we develop a measure for the resiliency dimension of liquidity in OTC dealer markets based on the rate of mean reversion of aggregate dealer inventories – following the conceptual notion of resiliency in Garbade (1982) and the principles underlying the empirical resiliency measure developed (for limit-order book markets) by Kempf et al. (2015). We find that each of the three dimensions of liquidity – trading costs, depth, and resiliency – are priced factors in the non-default spread. We find that the non-default spread is most affected by the trading cost and resiliency dimensions, while the depth dimension has a smaller, but still statistically significant effect. A one standard deviation change in trading costs, resiliency, and depth measures lead to a change in non-default spreads of 5.00 basis points, 2.27 basis points, and 1.27 basis points, respectively. For the case when the individual liquidity dimensions are at their average values, we find that about 80% of the non-default spread (attributable collectively to these three liquidity dimensions) comes from the bid-ask spread, about 17% from resiliency, and a relatively minuscule 3% from the Amihud depth measure. The average non-default spread in our sample period is about 21 basis points. To put this into perspective, the total yield spreads of Aaa and Baa industrial bonds over the same period (which impound both default and non-default risk) are about 90 basis points and 232 basis points respectively.[13]

[13] These total yield spread figures are calculated from 30-year Moody's Baa and Aaa yields less 30-year Treasury bond yields from H.15 releases.

Commonality in liquidity has been examined in several studies (see, for example, Chordia et al., 2000a; Chordia et al., 2000b; Pástor and Stambaugh, 2003; Acharya and Pedersen, 2005; Lin et al., 2011; and Bao et al., 2011). These articles suggest that market-wide liquidity factors may affect the non-default spread more than their idiosyncratic counterparts. In this context, we therefore create indices that measure the trading costs, depth, and resiliency of the Treasury bond market as a whole. We construct a market liquidity index based on the liquidity of Treasuries because the "market" for our government-guaranteed bank bonds is arguably much more comparable to the market for bonds carrying the same credit risk (i.e., the market for Treasuries), rather than the market for other corporate bonds carrying credit risk. We find that each of the three dimensions of market-wide liquidity has significant pricing relevance over our full sample period. When we control for the possibility of different liquidity pricing relationships during the financial crisis (as suggested by Dick-Nielsen et al. (2012) and Friewald et al. (2012)), we find that only the market-wide trading cost dimension is significantly priced during the crisis, in addition to bond-specific resiliency and bond-specific trading costs. However, in the post-crisis subsample, each of the three bond-specific and market-wide liquidity dimensions is significantly priced.

Finally, we find that the overall average *residual* non-default spread over our sample period (after controlling for state taxes and our three dimensions of liquidity) is not significantly different from zero. However, this *residual* non-default yield spread is statistically significant in some periods, albeit small in magnitude. In this context, Longstaff (2004) has earlier investigated government guaranteed Refcorp bonds, and found (like we do for our sample bonds) a non-default spread between these government guaranteed Refcorp bonds and Treasuries, even though they had the same credit risk. Longstaff (2004) concluded that this non-default spread was a "flight-to-liquidity" spread. However, Longstaff (2004) did not incorporate any controls (as we do in this study) for differences in (time-varying measures of) liquidity between Treasuries and his sample of guaranteed bonds, arguing that the differences (for example) in bid-ask spreads are too small in magnitude to explain the large yield spreads of Refcorp bonds. Our results in this study show that most of the Longstaff (2004) "flight-to-liquidity" premium is a liquidity premium directly related to the conventional measures of liquidity – spreads, depth, and resiliency. However, we also find that the non-default spread in some periods, particularly periods of crisis, impounds a tiny additional "flight-to-extreme-liquidity" premium that, in the spirit of the quote of former Federal Reserve Bank Chairman Alan Greenspan cited at the start of Longstaff (2004), reflects a strong investor preference for assets that enable quickest possible disengagement from the market if circumstances make that necessary.[14] Furthermore, we find that the residual non-default yield spread (after controlling for state taxes and our measures of liquidity) is *not* a positive function of issuer default risk, and hence unlikely to represent a "quality spread" arising (for example) because these bonds are guarantees, rather than direct obligations, of the U.S. Treasury. This latter result is also consistent with the indirect evidence in this regard in Longstaff (2004).[15]

---

[14] Longstaff (2004) quotes former Federal Reserve Bank Chairman Alan Greenspan as saying the following on October 7, 1998: *"But what is crucial… is that the individuals who were moving from, let's assume, the illiquid U.S. Treasuries to the liquid on-the-run liquid issues, are basically saying, 'I want out. I don't want to know anything about whether a particular investment is risky or not. I just want to disengage.' And the reason you go into these liquid instruments is that that is the vehicle which enables one to disengage as quickly as possible."*

[15] Longstaff (2004) investigated the government guaranteed bonds of only one entity – Refcorp. Hence, his conclusion in this regard was based on the absence of a positive dependence of the non-default spread on

The remainder of the chapter proceeds as follows. In Section 2, we develop the hypotheses tested in the study. Section 3 describes the sample used for our empirical analysis, including details of the FDIC's Debt Guarantee Program, and the estimation processes we use for the three liquidity dimensions. We report our empirical results in Section 4. Finally, Section 5 contains our concluding remarks.

2. Development of hypotheses

The most researched aspect of liquidity in extant literature is the trading cost dimension, typically estimated by the bid-ask spread of a security. In the bond market, Longstaff et al. (2005) split corporate yield spreads into default and non-default components and find that, among other factors, bid-ask spreads are indeed priced in the non-default component. Dick-Nielsen et al. (2012) also find that bid-ask spreads are priced in the non-default spreads of corporate bonds. We therefore base Hypothesis 1a on those studies.

H1a: *The trading cost dimension is priced in the non-default spread of bank bonds.*

Research has also analyzed the pricing of the depth dimension of liquidity. For equity markets, Brennan and Subrahmanyam (1996) document that an estimate of Kyle's λ – the depth dimension – is a priced risk factor in equities. In the bond market, Dick-Nielsen et al. (2012) find that the Amihud (2002) measure of depth is priced in the non-default spread. These studies motivate Hypothesis 1b.

H1b: *The depth dimension is priced in the non-default spread of bank bonds.*

---

the yield difference between AAA and BBB bonds (proxying for a possible perception of default risk in guaranteed Refcorp bonds).

Kempf et al. (2015) first developed a measure of the resiliency dimension for limit-order-book markets, using Garbade (1982) as the basis for modeling resiliency as the mean reversion of order-flow. Kempf et al. (2015) model time varying liquidity using a mean reverting model, $\Delta L_t = \alpha - \phi L_t - 1 + \varepsilon_t$, where L_t is the level of liquidity at time t. ϕ, the intensity of mean reversion, is their estimate of resiliency in liquidity. Using this measure of resiliency, Obizhaeva and Wang (2013) show that an optimal strategy of trading a given security depends largely on the resiliency of the security. We could not find any research studies on the pricing relevance of resiliency in liquidity – neither for stocks nor for bonds. However, Dong et al. (2010) provide evidence that price resiliency predicts the cross-section of stock returns. Also, Pástor and Stambaugh (2003) show the pricing relevance of an illiquidity measure based on equity return reversals, and hence closely related to price resiliency. We accordingly postulate Hypothesis 1c, and are the first to explore this dimension of bond market liquidity.

H1c: *The resiliency dimension is priced in the non-default spread of bank bonds.*

Commonality in liquidity has been widely explored in the existing microstructure literature, beginning with Chordia et al. (2000a); Chordia et al. (2000b) who show that the bid-ask spreads of securities covary with one another, and that the depths of securities also co-move with one another. In their seminal work, Pástor and Stambaugh (2003) show that a market-wide illiquidity measure is priced in stocks. Similarly, Acharya and Pedersen (2005) demonstrate that a stock's return depends on its relationships with market liquidity. In the bond markets, Lin et al. (2011) show that investors in corporate bonds are compensated for their exposure to general market illiquidity. Moreover, Bao et al.

(2011) show that for high-rated bonds, market illiquidity actually explains more than credit risk. These findings collectively motivate Hypothesis 2.

H2: *The non-default spread varies also with market-wide liquidity dimensions.*

Finally, we turn our attention to the residual component, if any, of the non-default spread that remains after accounting for the non-default spread arising from state-level taxes and the three (aforementioned) dimensions of liquidity we analyze.[16] If the non-default spread is driven entirely by state taxes and these dimensions of liquidity, this residual yield spread should be zero. If there is a significantly positive residual non-default yield spread, it could be a "quality spread" related to the risk of issuer default, arising because government guarantees may be considered inferior to direct government obligations because of possible procedural and time delays when the guarantee is actually invoked.[17] Alternatively, following Longstaff (2004), the residual non-default yield spread could also be a "flight-to-extreme-liquidity premium" related to the fear of future volatility, reflecting investor preference for assets that enable quickest possible disengagement from the market if that becomes necessary – an aspect of liquidity not necessarily fully captured by our time-varying measures of our three dimensions of liquidity. Accordingly, we propose Hypothesis 3a, 3b, and 3c.

H3a: *The residual non-default yield spread that remains after accounting for the trading cost, depth, and resiliency dimensions of liquidity, is zero.*

[16] While the non-default spread has also been explained empirically using variables like maturity, market uncertainty, and certain debt covenants, these factors should affect the value of the bond only through illiquidity or state taxes as a channel.

[17] There could be good reasons for this. For example, in the formation of the Debt Guarantee Program, the FDIC initially claimed it would issue bondholders checks for the full amount of the guaranteed debt within days of a default; however, in the finalized program in November 2008, the FDIC stated that it would continue to make the scheduled payments of the defaulted debt issue (Federal Registrar, 2008).

H3b: *The residual non-default yield spread that remains after accounting for the trading cost, depth, and resiliency dimensions of liquidity, is a "quality spread" related to the risk of issuer default.*

H3c: *The residual non-default yield spread that remains after accounting for the trading cost, depth, and resiliency dimensions of liquidity, is a "flight-to-extreme-liquidity spread" related to the fear of future volatility.*

3. Sample and research design

In order to isolate the non-default spread of bonds, we must control for default risk. To do this, we use a special set of corporate bonds with the same default risk as the US Treasury. This special set of bonds comes out of the financial crisis and Debt Guarantee Program (DGP), in which the FDIC insured bank debt against default with the full faith and credit of the United States government. The FDIC's backing is reflected in the highest possible ratings in the rating system, i.e., AAA ratings, for each of these guaranteed bond issuances, even though this was not necessarily the case for other bonds of the same issuer, with ratings varying all the way down to BB.

These fixed-rate insured bonds provide a very clean setting in which to analyze the yield spreads of corporate debt. This is because these insured bonds should have default risk equal to that of US Treasuries and, therefore, no additional default premium. By subtracting the yields of Treasury debt from the yields of these insured bonds, we can observe the implied non-default component of the yield spread without relying on the kind of measurement-error-inducing models that are used in extant literature.

Transaction-level data for this study comes from the TRACE (Trade Reporting and Compliance Engine) Enhanced dataset. The sample collected from TRACE includes all transactions of DGP bonds with fixed or zero coupons. The program began in October 2008 and continued through December 2012. That is, guaranteed bonds could be issued between October 14, 2008 and October 31, 2009 where the government guarantee on these issuances expired December 31, 2012. In practice, all of the bonds issued under the DGP matured prior to this deadline. Bond-level data for the bonds in the sample was obtained from the Mergent Fixed Investment Securities Database (FISD) and merged by CUSIP. To eliminate erroneous entries in the TRACE data, the transactions are filtered according to the methods outlined by Dick-Nielsen (2009). We also employ the agency filter from Dick-Nielsen (2009) to remove paired agency trades. The data are then processed further using a 10% median filter as described by Friewald et al. (2012). Following Bessembinder et al. (2009), daily yields are obtained by weighting individual trade prices by volume, and finding the yield from the resulting price.

Daily Treasury yields are obtained from the H-15 release data from the Federal Reserve and maturity-adjusted for each observation using linear interpolation, following Dick-Nielsen et al. (2012). The non-default spread is then estimated by subtracting these Treasury yields from the yields of the government-guaranteed bonds. After later merging these non-default spreads with the different measures of liquidity, we are left with 10,122 bond-day observations. To test the aforementioned hypotheses, we calculate proxies for each of the three dimensions of liquidity. The TRACE Enhanced dataset makes this possible by providing non-truncated volumes and a buy/sell indicator.

As a measure of the trading cost dimension, we follow Hong and Warga (2000) and approximate the daily bid-ask spread for each bond by taking the difference between the daily volume-weighted averages of the buy and sell prices. The effective half-spread is then scaled by the midpoint of the average buy and sell prices as follows:

$$Spread_{id} = \frac{\frac{\sum_{D=1} q_{itd} p_{itd}}{\sum_{D=1} q_{itd}} - \frac{\sum_{D=-1} q_{itd} p_{itd}}{\sum_{D=-1} q_{itd}}}{\left(\frac{\sum_{D=1} q_{itd} p_{itd}}{\sum_{D=1} q_{itd}} + \frac{\sum_{D=-1} q_{itd} p_{itd}}{\sum_{D=-1} q_{itd}}\right)}, \tag{8}$$

where q_{itd} is the volume of trade t for bond i on day d, p_{itd} is the price of that trade, and D equals 1 for all public buys and −1 for all public sales.

Similar to Dick-Nielsen et al. (2012) we use the Amihud (2002) illiquidity measure as a proxy for price impact of trades, and thus the depth dimension of liquidity. We estimate the Amihud measure as the following:

$$Amihud_{id} = \frac{100}{T} \times \sum_{t=2}^{T} \frac{abs\left(ln(p_{itd}) - ln\left(p_{i,t-1,d}\right)\right)}{q_{itd}/1{,}000{,}000}, \tag{9}$$

where T represents the number of trades of that particular bond on day d. This measure captures the change in price for a given quantity traded. To the extent that overall quantity traded (rather than signed order flow) represents the order flow in Kyle (1985), this is a theoretically valid measure of depth, and is extensively used as such in recent literature.

The empirical measure of resiliency in liquidity that has been used in the literature is the Kempf et al. (2015) measure for limit order book markets based on the principles outlined by Garbade (1982). In this framework, resiliency in liquidity (i.e., trading cost or depth) is the extent to which distortions in liquidity (trading cost or depth as the case may be) get neutralized within a pre-specified time. Based on this framework, we construct a measure of resiliency for over-the-counter dealer markets, like U.S. corporate bond markets. Since the change in aggregate dealer inventories represents the overall signed order flow in a dealer market, we define resiliency in liquidity as the extent to which distortions in dealers' aggregate inventory get neutralized by the change in inventory within a pre-specified period. Dealers target a given inventory level, and will give attractive prices to buyers and unattractive prices to sellers when they have relatively high inventory levels, and vice versa when their inventory levels are relatively low (see Amihud and Mendelson, 1980; Ho and Stoll, 1981 ; Ho and Stoll, 1983; Hansch et al., 1998). Hence, the stronger the mean reversion in aggregate dealer inventories, the higher the resiliency. Accordingly, to estimate a bond's resiliency, we measure the extent of mean reversion in aggregate dealer inventories; i.e., the relationship between the level of dealer inventory at time t and the change in dealer inventory from time t to time $t+1$. The daily ϕ measure from the following regression is used as our resiliency measure in further analysis:

$$\Delta Inv_{itd} = \alpha_{id} - \phi_{id} Inv_{i,t-1,d} + \varepsilon_{itd}. \tag{10}$$

Consistent with earlier literature (e.g., Naik and Yadav, 2003), we assume that aggregate dealer inventory is zero at the beginning of the sample, and adjust aggregate dealer inventory for each trade over the life of the bond. ϕ_{id}, our measure of resiliency, is a mean reversion parameter, and should theoretically be between 0 and 1, with 0 indicating that dealer inventory is a random walk with no mean reversion, and 1 indicating perfect resiliency, meaning that dealers are always at their target inventory, which eliminates any liquidity-related pressures on prices to deviate from their intrinsic value. Therefore, the higher the value of ϕ, the greater the resiliency in liquidity.

After we estimate the non-default spread, bid-ask spread, Amihud measure, and resiliency measure for each bond-day, we winsorize each of these variables at the 1st and 99th percentiles. Then, to correct for skewness and – more importantly for this study – to improve interpretability of regression coefficients, we take the natural logarithm of the winsorized bid-ask spreads, Amihud measures, and resiliency measures. Finally, in order to test the relationship of these three dimensions independent of the others, we orthogonalize the three liquidity dimension variables by regressing them on the other two, and keeping the residual from these three regressions.[18] Because resiliency decreases as illiquidity increases, we lastly multiply the resiliency value by -1, so that it, as well as the bid-ask spread, the Amihud measure, and the non-default spread, are all increasing with illiquidity.

Market-wide liquidity measures are obtained from GOVPX, which provides trades and quotes for US Treasuries, from 2008 through 2012. For observations in 2008, we limit our Treasury sample to those indicated as "Active," or on-the-run. Similarly, for all other years, we limit our sample to Type 151 and 153 instruments, which are "Active Notes and Bonds" and "Active Treasury Bills," respectively.[19]

The best bid and ask prices for US Treasuries are provided by GOVPX. We first use these values to calculate the inside half-spread, and then average these values for every bond-day to get one bid-ask spread observation per bond-day. For consistency with the TRACE dataset, we estimate our market-wide Amihud illiquidity measure and

[18] For example, these residuals give us the variation in the depth dimension of liquidity while controlling for the bid-ask spread and resiliency, and likewise for the other two dimensions. This is important because the price impact of a trade is affected by more factors than just volume, like the bid-ask spread, for example.

[19] We do this because the GOVPX dataset is split into years 2008 and prior, and 2009 and after, with slightly different variables in the two subsets.

resiliency measure using only on-the-run Treasury trade data. We construct both the Amihud and resiliency measures as we do for the guaranteed bonds above, on a bond-day basis. We then winsorize the bid-ask spread, Amihud, and resiliency variables at the 1 and 99 percentile levels before averaging across days to construct three daily time series. Finally, we take the natural log of these three series to construct *ln(Market Spread)*$_t$, *ln(Market Amihud)*$_t$, and *ln(Market Resiliency)*$_t$. Similar to the individual bond measures, we change the sign of resiliency so that it is increasing in illiquidity.

Following Elton et al. (2001) we use a bond's coupon rate to control for the state tax premium. Due to constitutional law in the United States, state and federal governments cannot tax income from one another. This is most commonly illustrated in municipal bonds, wherein the income is exempt from federal taxation. However, the roles are reversed for Treasury bonds. States cannot tax the income from Treasuries. They can, however, tax the income (coupon payments) from corporate bonds; therefore corporate bonds, even those of equal default and liquidity risk, will have a slight yield spread over Treasuries, due to this “state tax premium”. When analyzing the residual non-default yield spread, we use the daily VIX level (obtained from the CBOE indices database) as well as S&P firm ratings (obtained from Compustat). Descriptive statistics for these measures are in Panel A of Table 10. Panel B contains the correlations of these variables.

4. Empirical results

4.1. Pricing of bond-specific liquidity dimensions

We begin our empirical analysis by testing Hypotheses 1a through 1c: whether the three liquidity dimensions are priced factors in these bonds. We do so using the following regression model:

$$NDS_{id} = \alpha + \beta_1 ln(Spread)_{id} + \beta_2 ln(Amihud)_{id} + \beta_3(-ln(Resil))_{id} + \boldsymbol{\beta}'\boldsymbol{X} + \varepsilon_{id}, \quad (11)$$

where $\boldsymbol{X}$ is a vector of control variables, including coupon and fixed-effects in various specifications. For this model, we use robust standard errors clustered two-ways, by day and bond, as suggested by Pedersen (2009). This corrects the standard errors for autocorrelation within bonds and heteroskedasticity between bonds.

The results of these regressions are reported in Table 11. In Model 1, we use no fixed-effects and find that the coefficient on the log of the bid-ask spread is 0.059. Using the means and standard deviations reported in Table 10, this means that a one standard deviation increase in a bond's bid-ask spread is associated with an increase in yield of 4.94 basis points.[20] Similarly, for a one standard deviation increase in resiliency, yields decrease by 2.07 basis points.[21] Both effects are statistically significant. However, the dependence on the Amihud measure is not statistically significant in this preliminary specification, although it is in the expected direction. Consistent with Elton et al. (2001) we find that state taxes are roughly 4.12 percent on the margin.

[20] We obtain the effect of a one standard deviation change by calculating $\beta \times ln(1 + s/m)$, where β is the calculated regression coefficient using the log transformation, "*s*" is the standard deviation of the non-transformed variable, and "*m*" is the mean of the non-transformed variable.

[21] Recall that the sign of resiliency is changed in the presentation of the regression results.

Next, in order to control for time-invariant, bond-specific effects, we use bond fixed effects in Model 2, which allows us to analyze the central research question of this study, i.e., the impact of the time-series variation in the liquidity dimensions of a particular bond on the non-default spread of that bond, while ignoring any variation between bonds. In this model, the effect of the bid-ask spread and resiliency on non-default spreads remains statistically significant and roughly unchanged in economic magnitude. Economically, a one standard deviation change in the bid-ask spread is associated with a 5.00 basis point change, and a one standard deviation change in resiliency is associated with a 2.27 basis point change in the non-default spread. However, the effect of the depth dimension also becomes statistically significant, but the magnitude is still considerably less than the effect of bid-ask spread and resiliency – a one standard deviation increase in the Amihud measure is associated with only about a 1.27 basis point increase in bond yield. Meanwhile, when we take the regressors at the mean and multiply them by their respective coefficients, we find that a non-default spread of 16.3 basis points is attributable to the bid-ask spread, 3.5 basis points is attributable to resiliency, and 0.6 is attributable to the Amihud measure. This means that, for this particular case, about 80% of the non-default spread (attributable collectively to these three liquidity dimensions) comes from the bid-ask spread, about 17% from resiliency, and a relatively minuscule 3% from the Amihud depth measure.[22]

To further analyze the relative impact of these three dimensions of liquidity, we reproduce Model 2 of Table 10 (in unreported regressions) and exclude each of the three orthoganalized liquidity variables, one at a time. We find that the R^2 of the model falls by

[22] Due to the bond fixed effects in this model, the means of the components need not add up exactly to the overall mean of the dependent variable, which is 20.7 basis points.

9.53%, 2.42%, and 0.58% when we exclude the bid-ask spread, resiliency, and Amihud measure variables, respectively. These results are qualitatively consistent with those documented above: the impact of the bid-ask spread is about four times that of resiliency, and that of resiliency is about four or five times that of depth. These relative impacts are also consistent with the univariate correlations in Panel B of Table 10. For example, in the regression we see that the non-default spread is most impacted by bid-ask spreads, then resiliency, and finally the Amihud measure; while in the correlation table, we find that non-default spreads vary most closely with orthoganalized bid-ask spreads (0.292), then orthoganalized resiliency (0.140), and finally the orthoganalized Amihud measure (−0.040). Since R^2 is a goodness-of-fit measure, when we square the correlations, we see that bid-ask spreads (without any controls) account for approximately 8.5% of the variation in bid-ask spreads, resiliency accounts for 2.0%, and influence of depth (Amihud measure) is relatively negligible.

In Model 3, we employ time (day) fixed effects to explore the effect on the non-default spread of the cross-sectional differences in liquidity of different bonds *within a given day* (rather than within bonds over time.) This model controls for day-specific effects that don't change across bonds, in particular, all market-wide variables. The effect it measures is different from Model 2; the coefficients in Model 3 measure the impact of a cross-sectional difference in liquidity between bonds on non-default spread on a particular day. All three dimensions in this model are again statistically significant. A one standard deviation difference (between different bonds) in spreads, resiliency, and depth changes the non-default spread of the bond by about 0.54, 0.17, and 0.63 basis points, respectively. We see that when specifying the model in terms of cross-sectionally examining variation *between* bonds, rather than over time for a particular bond, the depth dimension is surprisingly more important to investors. However, we are more interested in examining the pricing of liquidity for a particular bond over time (which is the bond fixed-effects specification).

Finally, as a check for robustness, we utilize firm-fixed effects in Model 4 to control for any time-invariant effects which affect firms' bonds differentially. In this model, we find results strikingly similar to those of the bond-fixed effect model. This shows that the effect of liquidity of the cost of debt is not firm-dependent. Again, these results indicate that the trading cost dimension of liquidity affects the non-default spread more than the resiliency and depth dimensions.

Overall, these results offer strong evidence in support of Hypothesis 1a, 1b, and 1c – that the trading cost, depth, and resiliency dimensions are each priced factors in the non-default spread of bonds. Furthermore, the trading cost dimension and the resiliency dimension are clearly more important to traders than the depth dimension.

4.2. Pricing of market-wide liquidity dimensions

As discussed previously, market-wide liquidity has been well documented in the literature. Because of this, we test Hypothesis 2, which states that the non-default spread varies also with market-wide liquidity measures. We do this by creating the aforementioned market liquidity variables from Treasury bond data. We then utilize these variables in our analysis of the non-default spread. Rather than estimating multiple liquidity "market models" to estimate the market and idiosyncratic components of liquidity, we opt instead to include both bond-specific and market-wide liquidity measures in the same regression. This parsimonious strategy reduces estimation error by assuming that the effect of market liquidity on bond-specific liquidity is constant over the entire sample.

Model 1 in Table 12 presents results without any fixed effects in the regression specification, while Model 2 presents results with bond fixed effects, which is what is directly relevant for the research question we are investigating. Models 1 and 2 cover the entire sample period. Interestingly, the inclusion of the market-wide liquidity proxies in the regression model does not materially affect the previous bond-specific results. We see that, over the full sample period, even after controlling for market-wide liquidity dimensions, a one standard deviation increase to a bond's bid-ask spread is associated with an increase in non-default spread of 4.20 basis points; a one standard deviation decrease in a bond's resiliency is associated with an increase in non-default spread of 2.03 basis points; and a one standard deviation increase in a bond's Amihud measure is associated with an increase in non-default spread of 1.16 basis points; each of them is statistically significant at the 1% level.[23]

The effects of the market-wide liquidity dimensions are also significant and large in magnitude over the full sample period. Focusing on the more relevant Model 2, we see that, even after controlling for bond-specific liquidity dimensions, a one standard deviation increase in market-wide trading costs, one standard deviation increase in market-wide depth, and one standard deviation decrease in market-wide resiliency is accompanied by an increase in non-default spread of about 0.56, 3.17, and 14.77 basis points, respectively; each of them are statistically significant at the 1% level.

[23] Again, we obtain these figures by calculating $\beta \times ln(1+s/m)$, using the means and standard deviations from the descriptive statistics provided in Table 10 and the regression coefficients provided in Model 2 of Table 12.

Dick-Nielsen et al. (2012), as well as Friewald et al. (2012), show a dichotomy in liquidity pricing between crisis and non-crisis times. In light of this finding, we split our overall sample into crisis and post-crisis subsamples, and present the corresponding results (with bond fixed effects) in Models 3 and 4 respectively. Model 3 includes only the financial crisis period and Model 4 includes only the post-financial-crisis period. We classify transactions from 2008 and 2009 as being within the "crisis" subsample and transactions in 2010 and later as being in the "post-crisis" subsample. The results confirm a strong contrast in the two pricing regimes. We see that during the post-crisis period (Model 4), the pricing relevance of each of the dimensions of both bond-specific and market-wide liquidity remain highly significant, and qualitatively similar to what we have for the overall period. However, during the crisis period (Model 3) the situation is different. Bond-specific trading costs, market-wide trading costs, and bond-specific resiliency are the only liquidity dimensions that remain statistically and economically significant.

4.3. Residual non-default spread

Our results thus far show that the non-default component of the yield spread in our sample of FDIC-guaranteed DGP bonds depends significantly on the three widely accepted dimensions of liquidity – spreads, depth, and resiliency – and also reflect state taxes, as they should, since these bonds are subject to state taxes while U.S. Treasuries are not. In this sub-section, we examine if there is any residual non-default spread that remains unaccounted for after accounting for state taxes and the three dimensions of liquidity we have investigated.

The results of our tests for this residual non-default yield spread are reported in Table 13. In this table, we again regress the non-default spread on the three liquidity dimension proxies – trading costs, depth, and resiliency – and the coupon rate; but the important difference from earlier tables is that the liquidity dimension proxies in the regressions reported in this table have been transformed so that the regression coefficient on the constant term can be interpreted as the remaining magnitude of the non-default spread when the various liquidity dimension variables represent perfect liquidity. We do this by multiplying the liquidity variables by 100, adding 1 and taking the natural logarithm, except that for resiliency, we multiply “1 minus resiliency” by 100, add 1, and then take the logarithm. The intercept provides the residual non-default yield spread since it is the conditional mean of the dependent variable of the regression (the non-default spread) when all of the other variables are zero. This specification allows us to interpret the intercept term as the residual non-default yield spread remaining after controlling for state taxes and the three dimensions of liquidity we analyze, while keeping the distributions of the liquidity variables similar to previous analysis. Therefore, by using these transformed variables, the intercept estimates the mean value of the non-default spread when the bid-ask spread is zero (i.e., perfect liquidity from a trading cost perspective), the Amihud measure is zero (i.e., perfect liquidity from a depth perspective), and the resiliency is 1 – or more accurately “1 minus resiliency” is zero (i.e., perfect liquidity from a resiliency perspective). By including the coupon rate, we also control for the state tax premium.

Model 1 in Table 13 presents the results of running the above regression for the overall sample with only bond-specific liquidity dimensions. We find that the residual non-default yield spread is not significantly different from zero despite a sample of over 10,000 bond-day observations. The intercept term of 0.0097 percent, while not significantly different from zero, represents about 4% of the mean non-default spread.[24]

To further explore the robustness of our conclusion, we employ more specifications and controls. First, since Dick-Nielsen et al. (2012) show that liquidity is priced differently in crisis and non-crisis periods, we run, in Models 2 and 3 respectively, separate regressions for the crisis (2008–2009) and post-crisis (2010–2012) portions of our sample. When we account for potentially different dependence on liquidity measures in different periods, we do find statistically significant residual non-default yield spreads of about 8 basis points in both the crisis sample and the post-crisis sample. We then run the regression for the overall sample but control for market-wide liquidity dimensions in Model 4 of Table 13. Even when we include our three market-wide liquidity dimension proxies, we find, similar to Model 1, no statistically significant residual non-default yield spread for the overall sample. However, when we split the regression sample into the crisis and post-crisis time periods in Models 5 and 6, we again find a residual non-default yield spread of about 8 basis points in each sub-period, though the p-value in the crisis subsample is only 0.13, i.e., not significant at the conventionally used levels of significance.

Overall, our results indicate that, after we control for state taxes and for the trading cost, depth, and resiliency dimensions of liquidity, the residual non-default spread is, on average, zero or minuscule in magnitude; but it may not be appropriate to definitively

[24] It should be noted that this model has an adjusted R^2 of only 0.089. So even though 96.12% of the size of the non-default spread has been statistically accounted for by the state tax premium and the three liquidity dimensions, about 91% of the variation of the non-default spread remains unexplained.

rule it out in its entirety in each period. We accordingly explore two possible reasons for such a residual non-default spread. Since the residual non-default yield spread is not the primary focus of this study, our analysis is largely exploratory, leaving an in-depth investigation of the reasons driving the observed residual non-default yield spread for future research.

First, we note that, although our DGP bonds were backed by the full faith and credit of the United States government, they differed from Treasuries in that they were only guarantees and not direct obligations. Hence, there could potentially exist a "quality spread," reflecting possible procedural and time delays, related arguably to the market-perceived risk of actual issuer default (as it should closely proxy for the probability of the guarantee actually being invoked).

Second, a residual non-default yield spread could also arise because of variables we may have omitted in our regression specifications, or variables that we may have specified in a functional form that did not fully reflect the dependence of the non-default spread. In particular, in the spirit of the Alan Greenspan quote from Longstaff (2004) cited in footnote 3 above, the residual non-default yield spread could potentially be driven, for example, by a "flight-to-extreme-liquidity" premium reflecting a strong investor preference for assets that enable the quickest possible disengagement from the market if circumstances make that necessary.

In light of the previous results, we further analyze the residual non-default yield spread to determine whether or not this yield spread can be driven by a flight-to-extreme-liquidity or by the difference in quality between government guarantees and government obligations. Longstaff (2004) suggests that the yield spread between these bonds and

Treasuries is driven by flight-to-liquidity, which is spawned by a general market fear motivating investors to place their capital in securities which allow them to disengage from the market as easily as possible. We therefore include a proxy for general market fear factor – the VIX level – in the residual non-default yield spread regression specification. We demean the VIX for each regression specification in order to keep the intercept coefficients interpretable. This does not affect the covariance of the non-default spread and the VIX, thus the associated regression coefficients on the VIX are unaffected. Our regression results are in Table 14.

As we see in Table 14, the VIX is positively related to non-default spreads after controlling for liquidity and state taxes, and the dependence is statistically significant. This is consistent with the residual non-default yield spread being indeed driven by this general market fear factor, as Longstaff (2004) suggests. Specifically, we find that a one unit increase in the VIX is associated with a 1.75 basis point increase in residual non-default yield spreads. This effect is increased to 2.22 basis points per unit increase during the crisis period, and reduced to 0.34 basis points per unit increase in the post-crisis period – which is consistent with a flight-to-liquidity premium being more important in times of crisis. We find that, even after controlling for this market fear, the conditional mean of residual non-default yield spreads in the post-crisis period remains at 8 basis points. However, the conditional mean of the residual non-default yield spread in the crisis period rises to about 29 basis points, driven by the VIX level of 31.84 that existed in that period. Over our full sample, we find residual non-default yield spread levels of around 12 basis points.

We finally examine whether this residual non-default yield spread could also be caused by a market perception that these guaranteed bonds are of inferior credit quality to Treasury bonds. Thus, we investigate whether the residual non-default yield spread is a function of market-perceived default risk. To do this, we include issuer credit rating fixed effects in the three regressions modeled in Table 14. These fixed effects are graphically represented in Fig. 1. When looking at the full and post-crisis samples, we find absolutely no evidence that the residual non-default yield spread is a function of market-perceived default risk. Specifically, we show that the residual non-default yield spread does not increase as issuer credit ratings worsen. This is also shown for the crisis subsample, for bonds of all credit ratings, albeit with one single exception. Two bonds, both issued by New York Community Bank, which had a “BBB-” Standard & Poor's credit rating during the crisis sample period – the worst rating of any bond in that period and hence one most likely to default – have much higher non-default spreads than their liquidity and VIX levels would suggest. This could be interpreted as indicating that, during the stressful crisis period, investors became wary of guaranteed bonds with the highest probability of default – possibly due to the uncertainty of guarantee payments in the event of default, or the possible red tape involved in receiving payments – and priced that risk accordingly. Alternatively, these two extreme observations from one particular bank in one particular sub-period could just be outliers. Thus, while we cannot completely rule out the conjecture that the residual non-default yield spread is due to a perceived inferiority of guaranteed bonds to direct obligation bonds, our overall results are not consistent with that view – a conclusion that is consistent with the earlier indirect evidence in Longstaff (2004).

Taken in conjunction, our results indicate that, while most of the Longstaff (2004) "flight-to-liquidity" premium is a liquidity premium arising from the conventional measures of liquidity – spreads, depth, and resiliency – the non-default spread could also impound, particularly in periods of crisis, a tiny additional "flight-to-extreme-liquidity" premium reflecting, as suggested by Longstaff (2004), a strong investor preference for assets that enable quickest possible disengagement from the market if necessary.

4.4. Robustness tests

We document a strong relationship between the non-default spread and each of the three dimensions of liquidity – trading costs, depth, and resiliency. The direction of causality in this relationship should arguably be from liquidity to non-default spreads, since it is difficult to think of a credible economic rationale for higher (lower) yields to *cause* correspondingly lower (higher) levels of liquidity. However, without a shock to bond liquidity that is exogenous to yields, we cannot formally test the causal direction of the relationships we document between non-default spread and liquidity. Instead, we attempt to address this empirically using a changes specification, vector autoregressions, and impulse response functions. All of these suggest that shocks to the non-default spread do not cause changes to the three liquidity dimensions, and point instead to causality from the three liquidity dimensions to the non-default spread.

We begin by analyzing the relationship of daily *changes* in the non-default spread and the liquidity dimensions using the following regression model:

$$\Delta NDS_{id} = \alpha + \beta_1 ln(Spread)_{id} + \beta_2 ln(Amihud)_{id} - \beta_3 (ln(Resil))_{id} + \beta_4 NDS_{i,d-1} + \boldsymbol{\beta' X} + \varepsilon_{id}. \tag{12}$$

Because the non-default spread is arguably an integrated time-series – specifically the sum of a collection of previous shocks to the non-default spread – which is suggested by Longstaff (2004), we include the lagged level of the non-default spread. While this changes the interpretation of the regression coefficients, if we find that the levels of the liquidity dimensions affect the shocks to non-default spreads, it suggests that the liquidity dimensions causally affect non-default spreads, and not the contrary. The results of these regression specifications can be found in Table 15. In Model 1 of Table 15, we see that, as in the levels specification (Table 11), the effect of the bid-ask spread is larger on the non-default spread than the effect of the other two liquidity dimensions. When we control for time-invariant, bond-specific factors by including bond fixed-effects in Model 2 of Table 15, these results still hold. Next, we include market liquidity variables (as in Table 12) in Model 3 and find that while the market variables are significantly positively correlated with shocks to the non-default spreads, the bond-specific liquidity dimensions remain strongly significant. Finally, we split the sample into crisis and post-crisis. We again find a reduced effect of liquidity on the non-default spread during the crisis period. We find that the non-default spread has much less mean reversion during the crisis than in other periods (indicated by a smaller absolute value of the regression coefficient on the lagged NDS level). Thus, the non-default spread could still be a function of liquidity levels but in this specification, the lagged NDS already impounds previously-observed liquidity levels. In the post-crisis subsample, we find that all six liquidity dimensions are significantly priced, and the non-default spread is largely mean-reverting. Once again, we find that the level of market resiliency has a larger effect on the non-default spread than any other dimension. These results largely confirm our earlier analysis and dissuade any concerns that the previous regressions suffered from misspecification.

Typically, in the extant literature, it is assumed that the non-default spread is affected by the contemporaneous level of liquidity. We examine the following vector autoregression of the non-default spread and the liquidity dimension variables, in order to examine whether the lagged level of the liquidity dimensions affects the non-default spread, as well as investigate the reverse causality possibility:

$$\mathbf{V_{id}} = \boldsymbol{\alpha}' + \boldsymbol{\beta}_1'\mathbf{V_{i,d-1}} + \boldsymbol{\beta}'\boldsymbol{X} + \boldsymbol{\varepsilon}_{\mathbf{id}}. \tag{13}$$

where $\mathbf{V}_{id}$ is a vector containing the non-default spread, *ln(Spread), ln(Amihud), and ln(Resiliency)* for bond *i* on day *d*. The lagged liquidity dimensions are excellent proxies for the contemporaneous liquidity dimensions because their exogeneity is difficult to argue – the non-default spread on day *d* cannot affect the level of liquidity on day $d-1$, especially after controlling for the non-default spread in day $d-1$. We also attempt to control for any remaining residual non-default yield spread using the contemporaneous VIX level as variable proxy for the "fear factor". We display the VAR for the crisis subsample in Panel A of Table 16 and the VAR for the post-crisis subsample in Panel B of Table 16. In the crisis subsample we see that all three dimensions of liquidity are priced when we use lagged dimension levels as proxies, confirming that in the crisis liquidity levels and non-default spreads were very persistent. More importantly, we see that the lagged non-default spread has a much smaller statistical effect on the contemporaneous liquidity dimensions than the effect of the lagged liquidity levels on the non-default spread. This goes a long way in dissuading a reverse causality argument, albeit without a properly identified exogenous event. We confirm this when we include contemporaneous variables into the VAR to examine the impulse responses of these four variables. Visual representations of the impulse response functions during the crisis period are provided in

Panel A of Fig. 2. From these impulse responses, we see that the liquidity dimensions have a much smaller response to a one standard deviation shock to the non-default spread than the non-default spread has to a one standard deviation shock to the liquidity dimensions, once again weakening the reverse causality argument. These results hold when we examine the impulse response functions during the post-crisis period in Panel B of Fig. 2. Interestingly, when we examine the VAR during the post-crisis period in Panel B of Table 15, we see that the lagged liquidity dimensions are not significantly priced in the non-default spread after we control for the lagged non-default spread and the VIX level. In conjunction with the changes specifications, this suggests that during the crisis, non-default spreads and liquidity levels were very persistent, but in the calmer, less uncertain environment of the post-crisis period, the non-default spread is more mean reverting and is a function of contemporaneous liquidity levels. Irrespective, overall, these results strongly point towards causality from the three liquidity dimensions to the non-default spread.

5. Concluding remarks

The seminal market microstructure literature – Garbade (1982), Kyle (1985), Harris (1990), and Harris (2003) – identifies three important dimensions of liquidity: trading costs, depth, and resiliency. This is the first study to investigate the relevance of each of these three dimensions of liquidity – separately and in conjunction – for the pricing of corporate bonds, specifically, bank bonds. Unlike previous studies, our sample allows us to cleanly separate the default and non-default components of yield spreads. We find that each of the above three dimensions of liquidity are priced factors in the non-default spread. Both bond-specific and market-wide dimensions of liquidity are priced. The trading cost and resiliency dimensions are relatively more important than the depth dimension as determinants of the level of the non-default spread. Finally, we find that, even after controlling for these three dimensions of liquidity, there does exist in some periods a small residual non-default yield spread that is consistent with an additional "flight-to-extreme-liquidity" premium (related to the fear of future volatility, and consistent with Longstaff (2004)) reflecting investor preference for assets that enable the quickest possible disengagement from the market when necessary.

This study contributes to the extant literature in several important and significant ways. We are the first to examine whether the resiliency dimension of liquidity is priced in bond yields. Second, we are also the first to test whether the aforementioned three dimensions of liquidity – trading costs, depth, and resiliency – are priced in *conjunction*, as opposed to being priced separately. Third, an important methodological contribution we make is to use the principles underlying the empirical measure of resiliency developed (for limit-order-book markets) by Kempf et al. (2015) to define and estimate a new measure for the resiliency of over-the-counter dealer markets (like corporate bond markets). Fourth, we analyze the relative pricing relevance of both bond-specific and market-wide dimensions of liquidity. Fifth, we show that most of the Longstaff (2004) "flight-to-liquidity" premium is a liquidity premium directly related to the conventional measures of liquidity – spreads, depth, and resiliency. However, we also do find that the non-default spread in some periods, particularly periods of crisis, impounds a tiny additional "flight-to-extreme-liquidity" premium that, in the spirit of the quote of former Federal Reserve Bank Chairman Alan Greenspan cited at the start of Longstaff (2004), reflects a strong investor preference for assets that enable quickest possible disengagement from the market when necessary. Finally, consistent with Longstaff (2004), we do not find significant evidence of a "quality spread" arising from government guaranteed bonds being perceived inferior to direct government obligations.

Chapter 3: Director Networks and Firm Value[25]

1. Introduction

The board of directors of a corporation is responsible for making decisions on major corporate issues and establishing policies related to management such as setting CEO compensation and firing and hiring the CEO. The director network, defined as the connections, both current and former, between a firm's board of directors and board members at other firms, may allow well-connected boards to perform these crucial tasks more effectively. Connected directors may not only have better access to information about value-increasing management practices (Mizruchi, 1990; Mol, 2001), but also have more influence over fellow directors and management to ensure these practices are implemented (DeMarzo et al., 2003). Moreover, better connected directors may have better access to suppliers, customers or politicians through their network which can lead to strategic economic benefits for the firm. Conversely, a well-connected board could also have negative effects on firm value. For instance, more connected directors could be more distracted (Fich and White, 2003; Loderer and Peyer, 2002; Fich and Shivdasani, 2006) or they may spread value-destroying management practices or misleading information (Bizjak et al., 2009; Snyder et al., 2009; Armstrong and Larcker, 2009). To determine whether the benefits of connected boards outweigh the costs, we use an exogenous shock to board connectedness to examine if director networks are valuable.

To determine the effect of director networks on firm value, we use the sudden deaths of well-connected directors as a shock to the networks of directors who sit on the

[25] This chapter is based on collaborative work with Tor-Erik Bakke, Scott C. Linn, and Hamed Mahmudi.

same board as the deceased director (interlocked directors). The death of the well-connected director severs the network tie between the interlocked director and the deceased director's network. This represents a negative shock to the director network of the interlocked director's firm (director-interlocked firms). By looking at the announcement returns of the director-interlocked firms and using a difference-in-differences methodology, we find that this negative shock to director networks reduces firm value suggesting that director networks are value-enhancing.

Existing work finds evidence of a positive association between director networks and firm value (Larcker et al., 2013). However, due to the pervasive endogeneity of director choice and firm value, convincingly establishing causality has eluded researchers. For instance, better connected directors may choose to sit on the boards of better performing firms, or an omitted variable such as investment opportunities may be correlated with both director connectedness and firm value. Moreover, well-connected directors are usually more experienced and talented making identifying the effect of directors' connectedness on firm value difficult. Our experimental setup helps overcome this endogeneity problem as we identify exogenous shocks to board connectedness. The sudden death of interlocked directors is unlikely to be correlated with value-relevant omitted variables which could contaminate inference. Importantly, by studying how the sudden death of well-connected directors affects the value of interlocked firms (and not the deceased director's firm itself), we are able to separate out the effect of board connectedness on firm value from the effect of other value-relevant director attributes. Finally, the randomness of the sudden death breaks the endogenous matching between directors and firms.

Our sample consists of Canadian public firms from 2000 to 2012. We focus on professional networks. That is, directors are connected if they currently or previously served on the same board. The advantage of focusing on professional connections is that they are observable, objective and not subject to sample selection concerns. For instance, unlike many educational ties where directors may have simply co-existed in the same environment, directors that served on the same board have had repeated face-to-face interactions and a working relationship. The disadvantage is that we miss other types of social connections that could also facilitate the flow of information and affect the centrality of a director in the network.

Next, we compute commonly used network centrality measures for each director in the network. To identify significant exogenous shocks to the director network we focus on the sudden deaths of the most connected directors. Our sample consists of the sudden death of seven well-connected directors which results in 128 directors at 159 interlocked firms that experienced negative shocks to their director networks. These director-interlocked firms lose access to the deceased director's network and are therefore considered the treatment group.

The shock to the connectedness of director-interlocked firms is economically and statistically significant. We find that the eigenvector centrality of treated firms falls by about 1% relative to control firms.[26] Moreover, it is likely that the change in centrality measures understates the magnitude of the shock as it implicitly assumes that readjusting the network is frictionless. In reality adjusting one's network, to compensate for the loss of the well-connected directors network, entails significant frictions in the form of search

[26] Changes in eigenvector centrality best capture the loss of a well-connected director to an interlocked director's network which is the focus of our identification strategy.

costs. Further bolstering the argument that these shocks are significant, Falato et al. (2014) provide evidence that replacing a lost director is time consuming and costly.

To test whether the exogenous elimination of network ties affects firm value, we compare the announcement returns of treated firms to a matched sample of control firms that were unaffected by the director network shock. We show that the shock to director networks, caused by an unexpected death of a director, results in negative cumulative abnormal returns (CARs) for director-interlocked firms relative to the sample of control sample have similar pre-shock firm characteristics. Specifically, in univariate results, we find that relative to control firms, treated firms have around 0.6 percent lower abnormal returns in response to the sudden death of well-connected interlocked directors. When controlling for other various factors in a multivariate seemingly unrelated regression (SUR) framework, this difference is about 0.3 percent but remains highly statistically significant. This indicates that the loss of network connections led to a statistically and economically significant decline in firm value of director-interlocked firms.

We next investigate whether our results are driven by an increase in the busyness of interlocked directors. This is important as our results could be confounded by the fact that the sudden death of the well-connected director has two effects on interlocked directors: (i) a negative shock to the director's network and (ii) an increase in the director's busyness (Falato et al., 2014). To separate these two effects, we redefine treated firms to only include firms who lost a past connection due to the sudden death of one of the well-connected directors. These firms have at least one director who previously served on the same board as the deceased director (a past connection), but do not currently share a director with the deceased director's firm(s). As the loss of a past connection does not increase the busyness of the interlocked director, but does affect the connectedness of the director, this strategy enables us to better isolate the first effect (i). We find that our results continue to hold using only past connections suggesting that our results are not simply an artifact of an increase in director busyness but rather due to a reduction in the connectedness of the firm's directors.

We contribute to several strands of literature. First, we contribute to the broad literature on the value of connections. Faccio and Parsley (2009) show that political connections are valuable; Hochberg et al. (2007) find that more connected venture capital firms perform better; Cohen et al. (2008) and Cohen et al. (2010) show that education connections are valuable to mutual fund managers and equity research analysts respectively; Faleye et al. (2012) show that better-connected CEOs innovate more. We add to this literature by showing that firms benefit from having better-connected boards.

Second, our study fits in the literature that studies director networks. This literature uncovers positive and negative aspects to having a well-connected board. On the one hand, Engelberg et al. (2013) show that CEO pay is increasing in the number and importance of her own connections. Similarly, Barnea and Guedj (2009) and Renneboog and Zhao (2011) show that firms with better connected directors pay their CEOs more, but these firms also grant pay packages with lower pay-performance sensitivity. In addition, Barnea and Guedj (2009) show that well connected directors are more likely to get more directorships and provide softer monitoring.

On the other hand, Horton et al. (2012) show that the positive link between connectedness and director compensation is not due to the connected directors using their power to extract economic rents. Instead, they find evidence that firms compensate directors for their network connections. Moreover, Fogel et al. (2014) show that powerful independent directors are associated with fewer value-destroying M&A bids, more high-powered CEO compensation, more accountability for poor performance, and less earnings management. Helmers et al. (2015) find that better-connected boards spend more on R&D and obtain more patents. Shelley, and Tice (2015) demonstrate that firms with well-connected boards are less likely to both misstate their annual financial statements and adopt practices that reduce financial reporting quality. We contribute by showing that overall well-connected boards are value-enhancing.

Third, we add to the literature that studies the link between board connectedness and firm value. Several studies have found positive associations between the connectedness of a firm's board of directors and its operating performance (Hochberg et al., 2007; Horton et al., 2012; Crespí-Caldera and Pascual-Fuster, 2015). Larcker et al. (2013) show that firms with more connected boards have significantly higher risk-adjusted returns than firms with less connected boards. Stern (2015) demonstrates, using a learning model, that better connected board chairmen (but not directors in general) are associated with more value creation for their firms. In contrast to these papers, we provide causal evidence that having better connected directors increases firm value.

Fogel et al. (2014) provide evidence that the sudden death of powerful directors negatively affects the value of the powerful director's firm. However, unlike in our study, they are unable to distinguish whether the decline in value was due to the loss of the deceased director's talent or due to the loss of the deceased director's connections. As connected directors are likely to be talented, this may confound inference. We get around this challenge by looking at the effect of the sudden death on director-interlocked firms only. Thus, our study contributes by better isolating the effect of director networks on firm value.

The remainder of the chapter proceeds as follows. In Section 2, we discuss the link between the connectedness of a firm's board of directors and firm value. Section 3 describes the data used for our empirical analysis, and the network centrality measures. In Section 4, we discuss the benefits of our identification strategy and the research design. Section 5 contains our empirical results. Finally, we conclude in Section 6.

2. Director Networks and Firm Value

In this section we discuss the link between the connectedness of a firm's board of directors and firm value. We start by discussing the benefits of director networks, and then switch to the potential costs. The potential benefits of having well-connected directors come in three forms. First, directors can use their boardroom networks to gain access to valuable information from other directors. This information could be related to industry trends, market conditions, and regulatory changes or could be information on value-enhancing business practices (e.g., technological innovations, effective corporate governance mechanisms etc.). Thus, well-connected directors are able to make better decisions as they have access to a larger pool of information. (Mizruchi, 1990; Mol, 2001).

Second, well-connected directors may have better access to strategic economic benefits through their networks. For instance, closely connected firms could benefit from collusion and other anti-competitive behavior (Pennings, 1980).[27] Another potential

[27] It is important to note that, although collusion can have a positive effect on firm value, if it leads to the violation business law, the regulatory, litigation, and reputation costs can negatively affect firm value.

strategic benefit is that connected firms may enjoy political favors or superior supplier or customer relationships which could not be possible without access to a large professional network of directors.

Finally, better connected directors may be more influential and therefore better able to prevail in discussions with the rest of the board and management. As demonstrated in a theory paper by DeMarzo et al. (2003), an individual's influence on group opinions depends not only on accuracy, but also on how well-connected the individual is. Thus, a director who is well-connected within the network of directors is more likely to have the power to sway other directors in the board room towards his views. Both the well-connected director's access to superior information and increased power to persuade the board should lead to better firm decisions and enhanced shareholder value.

There are also potential costs to having a well-connected board. Bizjak et al. (2009), Snyder et al. (2009) and Armstrong and Larcker (2009) show that director connections facilitate the propagation of value-destroying governance practices. Moreover, well-connected directors with multiple directorships may be busy and therefore unable to allocate sufficient time and attention to monitoring and advising on all the boards on which they serve. This in turn could negatively affect firm value (Core et al., 1999; Fich and White, 2003; Loderer and Peyer, 2002; Fich and Shivdasani, 2006).

Several papers provide evidence suggesting that the benefits of director networks exceed the costs and that director networks overall increase firm value. Larcker et al. (2013) show that firms with large director networks are associated with superior risk-adjusted returns and greater increases in future profitability than firms with less connected boards. However, endogeneity remains a significant concern. This must be properly addressed before advising firms to go out and hire more connected directors.

In this setting endogeneity concerns are numerous and multi-faceted. One concern is reverse causality. For instance, more connected directors may choose to work for better firms (Masulis and Mobbs, 2012). Moreover, connected directors may also use their networks to correctly anticipate which firms are likely to perform well. Thus, causality may flow from firm value to more connected boards, and not vice versa. Another concern is omitted variables. Any unobservable variable that affects firm value and is correlated with board connectedness can contaminate inference. For example, connected directors may choose to work for firms with better governance or good investment opportunities, both of which are likely to affect firm value. Although, it is possible to find proxies for both governance and investment opportunities, these proxies are imperfect and any measurement error could significantly bias the estimated coefficients.

Another important latent variable is director ability. Fogel et al. (2014) provide evidence that the sudden death of powerful directors negatively affects the value of the powerful director's firm. However, powerful directors are also likely to be talented. By omitting director talent, Fogel et al. (2014) are not able to determine whether it is the loss of the deceased director's connections or talent that causes the decline in firm value. In the identification section (Section 4) we discuss how we tackle these endogeneity concerns and provide persuasive causal evidence that more connected boards increase firm value, but first we present our data and how we measure director connectedness.

3. Data and Network Centrality Measures

3.1 Data

Our sample consists of Canadian public firms in the Clarkson Centre for Business Ethics and Board Effectiveness dataset from 2000 to 2012. We use annual firm-level accounting data from Worldscope and return data from Datastream. We drop all observations with missing or negative total assets. We calculate Tobin's Q as the sum of market capitalization and the book value of debt, scaled by total assets. Leverage is calculated as total debt over total assets (and is treated as missing if less than zero), and ROA is calculated as net income over total assets. Finally, cash and capital expenditures are scaled by total assets. Firm size is *ln(total assets)*. All continuous variables are winsorized at the 1st and 99th percentiles.

3.2. Director network centrality measures

We construct director networks measures for our sample firms from 2000 to 2012 using data from the Clarkson Centre for Business Ethics and Board Effectiveness. Two directors are linked if they (i) currently sit on the same board or (ii) previously sat on the same board.[28] The network is undirected and unweighted. Undirected networks assume that influence and information flow both ways between connected directors. In unweighted networks each link between directors has equal importance (i.e., the intensity of each link is the same).

As is common in the literature (Renneboog and Zhao, 2011; Larker et al., 2013; Berkman et al., 2015; Crespi-Cladera and Pascual-Fuster, 2015) we restrict attention to

[28] In regard to past connections, we use director start dates, and end dates for each position that each director holds to establish if directors previously sat on the same board. This approach allows the network to extend back beyond 2000. One shortcoming is that we miss past connections if at least one of the directors ended a position before 2000.

the director's professional network (i.e., shared directorates). The advantage of focusing on professional connections is that we can observe the entire network. Moreover, no judgement is involved in determining the ties. Finally, directors that served on the same board have had repeated face-to-face interactions and a working relationship. In contrast, educational ties (considered by Fogel et al. (2014)) could range from situations in which directors worked closely together to situations in which directors may have simply co-existed in the same environment. A downside of focusing only on professional connections is that that we miss other types of social connections (i.e., friends, acquaintances, family etc.) that could also facilitate the flow of information and affect the centrality of a director in the network. Unfortunately, data on social ties is not widely available.

Using the start and end dates for each director's position, we are able to construct a separate adjacency matrix for each year from 2000 through 2012. Intuitively, the adjacency matrix represents the network structure in each sample year. More specifically, the adjacency matrix $\boldsymbol{A}$ is a symmetric matrix in which each row and corresponding column refer to an individual director. Director i is then defined as connected to director j ($\boldsymbol{A}[i, j] = \boldsymbol{A}[j, i] = 1$) if the two directors sit on the same firm's board in the same year, or have ever sat on the same board in the same year at some point in the past. If a director leaves the sample completely, and does not return, then all of her connections are severed. This could happen for various reasons from retirement to illness to a career change, as well as death.

Using the adjacency matrices constructed based on our network of directors and UCINET software (see Borgatti et al. (2002)), we calculate four network centrality measures for each director each year: degree, eigenvector, closeness, and betweenness. Degree centrality measure is the number of connections a given director has within the network. Mathematically, the degree centrality for director i is simply the sum of column i (or row i) in the adjacency matrix.

Eigenvector centrality is closely related to degree centrality. Intuitively, eigenvector centrality weights each connection by how important it is. Specifically, eigenvector centrality is an iteratively calculated weighted average of the importance of a director's direct contacts, with weights determined by the importance of their direct connections, and so on. Assuming E_i is the eigenvector centrality measure for director i, and $\boldsymbol{E}$ is a vector containing $[E_1, E_2, \ldots, E_i, \ldots, E_N]$, then the aforementioned iterative calculations will converge to the condition $\boldsymbol{AE} = \lambda\boldsymbol{E}$, where λ is the eigenvalue associated with $\boldsymbol{E}$.[29] The resulting E_i values are then normalized using a Euclidian normalization in order for the sum of the squares of the resulting centrality measures will equal 1 for any given network. This allows for comparison of eigenvector centrality measures between different networks.

A director's closeness centrality captures how close the director is to every other director in the network. Closeness centrality is calculated as the reciprocal of the sum of the shortest distances between the director and every other director in the network. One complication is that in large and complex network, such as the one we study, some directors in isolated subnetworks may have undefined distances to others (i.e., there are some parts of the network they cannot access). To account for this, we follow Fogel et al. (2014), and define director i's closeness, C_i as

[29] Since the adjacency matrix $\boldsymbol{A}$ may have multiple eigenvalues, we apply the Perron-Frobenius theorem to ensure that all $E_i \geq 0$, and use the eigenvector $\boldsymbol{E}$ with the largest λ.

$$C_i = \frac{n_i - 1}{\sum_{i \neq j \in N} g_{ij}} \times \frac{n_i}{N}, \tag{14}$$

where n_i is the size of the subnetwork which contains director i, g_{ij} is the geodesic distance from director i to director j, and N is the size of the entire network. This correction calculates the closeness of a director within a sub-network, and then weights that closeness measure by the relative size of the sub-network, which will correct for a director being highly connected within a very small sub-network (i.e., one firm with a board that has no connections to any other directors at other firms).

A director's betweenness centrality is the number of the shortest-paths between all directors in the network that go through the director. To better understand this measure, consider a spoke-and-hub network. The center hub will lie on every shortest path between the other directors (high betweenness), but a spoke will not lie on any of the shortest paths (low betweenness).

Directors who score highly on any of these four network centrality measures are likely to have more power and influence as well as better access to information. That being said, different centrality measures are important for different reasons. For example, the number of immediate connections a director has – degree centrality – as well as the importance of those connections – eigenvector centrality – may increase the director's power and influence in the board room (Renneboog and Zhao, 2011) and enable directors to better convince or persuade other directors or management. Closeness and betweenness centrality may be more apt to capture a director's ease of accessing valuable information. For example, if a director has a high betweenness centrality, then she is more likely to broker conversations with other directors, gaining insight to potentially valuable information. Similarly, if a director has a high closeness centrality, then his position to access information is advantageous relative to other directors in the network. It can also be argued that betweenness centrality better captures the power of the director as high betweenness implies that the director is on more of the shortest paths within the network and therefore more influential (Lee et al., 2010). Given the subtle differences between the measures, we report and use all four measures in our analysis.

4. Identification Strategy and Research Design

In this section we describe how we identify the effect of director networks on firm value. We focus on the negative shocks to director network stemming from the sudden and unexpected deaths of well-connected directors of Canadian firms from 2000 to 2012. To find the sudden deaths of well-connected directors we first, identify all directors who left the sample between 2000 and 2012. Second, we prioritize the directors with the highest network centrality measures. Specifically, we search for sudden death among the 2100 most connected directors for each year in our sample. Third, to ascertain which of the well-connected directors left the sample due to a sudden death, we hand collect information about the passing of numerous directors from Factiva, obituaries, news media, and press releases. We eliminate all directors who left the sample for a reason other than death (i.e. career change, retirement, etc.). To ensure that the sudden death is unanticipated and exogenous, we also exclude director deaths in which the director retired prior to his passing, or had a prolonged illness which caused them to leave a firm in the year of their death.

Ultimately, we classify seven deaths as sudden and unexpected. Specifically we identify sudden deaths of well-connected directors in 2001, 2002, 2003, 2005, 2006, and two in 2011. To illustrate what we consider a sudden death, consider two examples. One director, Donald Fullerton, died May 29, 2011. His obituary claimed it to be a "sudden but peaceful passing." Another director, John Beddome, died on May 10, 2005 "after a brief and courageous struggle with cancer." We deem each of these director deaths to be sufficiently unexpected so that any the impact of their deaths is not already impounded in market prices. Even if deaths were partially anticipated it is likely that much uncertainty is still resolved on and around the announcement date. Moreover, the suddenness of the deaths implies that the firm did not have readily available replacement. Consistent with this conjecture, Falato et al. (2014) provide evidence that about half of firms, that lost a director due to death, do not fill the director vacancy one or two year after the death. They show that firms fill director vacancies even slower after a sudden death.

For each year in which we identify sudden director death, we create "shocked" adjacency matrices for each year. These shocked matrices are identical to the pre-shock matrices, except for the column and row corresponding to the deceased director, in which we change each element to zero.[30] In other words, the post-shock network structure is identical to the pre-shock network structure except that the well-connected director is removed from the network. To assess the magnitude of the shock to director networks induced by well-connected director deaths, we aggregate the estimated centrality measures at the firm level by averaging the network centrality measures of firms' current directors each year. This is done for both the pre-shock and shocked director networks.

[30] In 2011, there are two chronological shocks. The first shocked adjacency matrix is treated the same as the other shocked matrices, but for the second shocked matrix, we use the first shocked matrix as the pre-shock matrix, and eliminate connections of the second deceased director.

Next, we find the percentage change in firm-level average director centrality by dividing the difference in the shocked and pre-shocked centrality by the pre-shock centrality value. This gives us a relative measure of how much a given death affected the network centrality of each firm's board of directors.

How large are our network shocks? To determine this we compare how the shock affected firms with direct connections to the deceased directors firms (treated firms) to firms without this direct connection (control firms). We find that the shock to director networks is economically and statistically significant. Eigenvector centrality is significantly shocked, dropping 0.91% more for treated firms relative to control firms. The other network centrality measures also experience statistically significant drops, but the magnitudes are smaller (degree, closeness, and betweenness centrality are differentially shocked by -0.26%, -0.02%, and -0.15%, respectively). This is not unexpected as eigenvector centrality is the centrality measure that is best suited to capture the loss of an important individual connection (as is the case in our setting).

We also regress the percentage changes in average firm network centrality on a treatment dummy. The regression coefficient on the treatment dummy captures the DID estimate of the effect of the network shocks on the network centrality measures. We also include a number of control variables as well as industry fixed effects in the regressions. The controls include board size (number of firm directors), size, market-to-book, and profitability. The standard errors are panel-corrected standard errors (PCSE). The results of running these regressions are in Table 20 Panel A. The most notable result is that all four of the centrality measures were negatively shocked by the deaths of these directors (all except one is also statistically significant).

It is important to keep in mind that the estimated shocks to the network centrality measures likely understate the true impact of the network shock. This is because the calculation of the post-shock adjacency matrix assumes that directors can adjust their networks immediately and without any costs (e.g., search costs). For instance, in the case of betweenness and closeness centrality the post-shock network recalculates all the shortest paths. In reality the readjustment of the network is unlikely to be frictionless, but both time-consuming and costly. Taken in conjunction, this evidence suggests that the deaths of the seven well-connected directors did in fact have a negative impact on the network centrality of connected firms.

To identify the impact of the board connectedness on firm value, we conduct an event study around each sudden death. Treated firms are defined as any firm that had a director interlock with the deceased director's firm. These firms lost access to the well-connected director's network and are therefore subject to a network shock. Given that the deaths were sudden and unexpected, announcement returns should capture the value implications of the network shock on the firm value of director interlocked firms. Moreover, the unexpected nature of the shock also ensures that we have exogenous variation in director networks allowing the identification of a causal effect. We compare the abnormal returns of treated firms to a baseline of similar control firms, that were unaffected by the network shock (i.e., do not have an interlock with the deceased director's firm). To the extent that the market anticipated how firms would react to loss of network connections, this difference-in-differences test can be interpreted as the causal effect of director connectedness on firm value.

It is important to recognize that we exclude the deceased director's firm from our analysis (i.e., these firms are not part of our treatment group). The deceased director firms could see drops in value for two reasons, due to (i) the loss of the deceased director's network connections and (ii) the loss of the deceased director's talents, experience and knowledge. By focusing our analysis only on director-interlocked firms we are able to, unlike Fogel et al. (2014), to isolate the effect of the shock to board connectedness on firm value.

To avoid violations of the parallel trends assumption, it is useful to test if observable firm characteristics of treated and control firm similar in the pre-shock period. Descriptive statistics of the pre-shock firm characteristics are displayed in Table 17 Panel A. We see that on average treated firms are much larger, have more board members, and are much better connected in the director network. The samples also differ in terms of cash holdings, Tobin's Q, and profitability (ROA).

Given the significant differences in pre-shock characteristics of the treated and untreated firms, we employ a matching procedure to obtain more similar treatment and control samples. Specifically, from the subsample of control firms in the same 1-digit SIC code, we limit the possible matches for each treated firm to its 7 nearest neighbors in pre-shock average director degree centrality, and then match each treated firm with the 3 nearest neighbors in pre-shock firm size. Matching is done with replacement. This results in a control sample of 477 firm-years, or 3 matched firms for each of the 159 treated firm-years. Descriptive statistics for the treated and matched control sample are displayed in Table 18. We use three different methods to test for differences in the distribution of the two samples. The difference in means is tested using both a pooled difference-in-means t-test and a paired difference t-test, while the difference in medians is tested using a two-sided (rank-sum) Wilcoxon Z-test.

As can be seen from Table 18 Panel A, the pre-shock samples are similar in both firm size and Tobin's Q. Because we are using Canadian firms, we are unable to utilize the full Fama and French (1993) 3-factor model to calculate abnormal returns, but only a market model, therefore a treatment and control sample matched on both size and book-to-market is important to alleviate concerns about systematic bias in our measurement of abnormal returns. We also see that the treatment and control firms are similar on most other dimensions, including board size, cash holdings, leverage, capital expenditures and return on assets. We do find statistically significant differences in both the means and medians in network centrality; however, the economic significance of the difference is fairly minuscule. For example, while the mean control firm has directors with an average of 29.14 connections, the mean treated firm has director with an average of 31.92 connections, a relatively small difference. Overall, the matched samples are similar on observables, which makes it less likely that a differential trend during the event windows is biasing our results.

5. Empirical Results

5.1. Main Results

We start our presentation of the empirical results, by examining the cumulative abnormal returns (CARs) around the announcement of the sudden director deaths. To implement our tests we first calculate abnormal returns for all firms in our sample using the market model. We use returns on the S&P/TSX Composite Index, including dividends, obtained from Datastream as the market return in the model. Betas are estimated using data from 230 trading days prior to the death of each well-connected director, as we exclude the 30 days prior to the event date from the estimation window to mitigate contamination. We focus on event windows (0), (0,+1), (-1,+1) as well as (-2,+2) to allow for potential leakage of information prior to the announcement. Day zero is the announcement date of director deaths. Leakage is a possibility in the cases in which the director is admitted to a hospital and passes away relatively quickly, however, leakage is unlikely if the director's death is due a stroke, heart attack or accident.

Next, we calculate cumulative abnormal returns (CARs) for each event window for the treated and control firms separately. As the CARs are clustered over seven different event periods, cross-correlation may bias our standard errors downward and lead to over rejection of the null hypothesis (Kothari and Werner, (2006)). To adjust for the cross correlation, we calculate the t-statistics for the difference of the mean CARs from zero using the technique in Kolari and Pynnönen (2010). As can be seen in Table 19, treated firms have abnormal returns that are negative and statistically significant on the day the death is announced, however, in other event windows the abnormal returns are not statistically significant.

However, we are more interested in how the treated and control firms reacted differentially to the death of the well-connected director. Therefore, we compare the differences of the CARs of the treated and control firms in the event windows. Treated firms have event-day abnormal returns of -0.36%, compared to 0.23% for control firms, resulting in a difference-in-differences (DID) estimate of -0.59%. This effect is statistically significant at the 1% level. When we expand the event window, we find the

DID estimate is -0.44% in the (0, +1) event window, and -0.45% in the (-1, +1) event window, both of which are statistically significant using a paired-difference t-test. These results strongly suggest with the negative shock to director connectedness results in a decrease in firm value.

Although, the treatment and the matched control group are similar (see section 4), it is possible that omitted variables could be driving our univariate findings. Thus, we further test the effect of these network shocks in a multivariate setting. This allows the inclusion of control variables and industry fixed effects. We use a seemingly unrelated regression (SUR) framework that allows coefficient estimates to vary for each shock. Residuals are assumed to be correlated within each shock, but uncorrelated between different shocks. In the regressions we control for board size (number of firm directors), size, market-to-book, and profitability. Another potential concern is that the residuals are correlated within each of the shocks leading to biased estimates of standard errors. To adjust for this standard errors are panel-corrected standard errors (PCSE).

Table 20 reports the cross-sectional regressions with CARs on the left hand side and an indicator variable that equals one if the firm is a member of our treated group on the right hand side. The results for the announcement day suggest that the treated firms experience abnormal returns that are smaller than those for the control firms by 0.3% which is statistically significant at the 5% level. In the other windows, the results are economically similar, but statistically insignificant. If the markets efficiently process the implications of the deaths for director networks, then this result is perhaps not surprising.

At the mean, a -0.30% decrease in firm value is approximately equivalent to a $19.7 million loss in market capitalization per treated firm.[31]

We also find that these results are qualitatively robust to instead using a paired difference specification in which each pair is a treated firm and a matched control firm. To this end, we regress the difference between treatment and control firm abnormal returns on the differences between treatment and control firm characteristics. The paired difference regression has the advantage that statistical power is improved in matched-pair regressions due to the additional information (i.e., which treated firm that is matched to which control firm) that is disregarded in pooled regressions. The downside of this specification is that control variables must take the form of matched-differences, which eliminates the possibility of industry fixed effects. This concern is mitigated by the fact that the control firms are matched to industry peers.

The results using the paired difference specifications are in Table 20 Panel B. We again find that treated firms, those with direct connections to the deceased director, have event-day announcement returns that are lower than control firms (0.23%). Expanding the event window to include the day following the death (i.e., (0,+1)), we see that treated firms' stocks had returns that were 0.31% lower compared to their control firms. Collectively these results are consistent with the negative shocks to director centrality reducing firms value differentially in our treated firms relative to our control firms.

5.2 Busyness

The value effect evident in the previous regressions is possibly due to two economic mechanisms. Firm value either dropped due to the exogenous severing of

[31] This is calculated as -0.30% × $6,564,604,348 where the latter number is the average market capitalization of treated firms.

director network connections, or because board members from treated firms must now work more for the firm of the deceased, and thus neglect their other firms. This is the busyness effect hypothesized by Falato et al. (2014). They find, using the sudden death of interlocked directors, a negative abnormal market reaction for interlocked firms which they attribute to the increased busyness of the interlocked directors.[32] This confounding effect of busyness is a threat to the internal validity of our results as the firm value of the director-interlocked firms (treated firms) could decrease either due to a negative shock to director networks or because its board is more distracted.

We tackle this challenge by showing that our results continue to hold in a sample where director busyness is unaffected. To avoid the contamination of the increased busyness effect, we focus on a subsample of firms which do not share an interlocked director with the firm of the deceased, but which do have a past professional connection to the deceased director. In other words, we analyze the returns of firms which have directors that previously sat on boards with the deceased, but did not at the time of his death. This allows us to isolate the effect of a change in network centrality without any confounding change in busyness. This is because the director with only a past connection to the deceased will experience a loss of connectedness, but will not be incurring an increased workload due to the death.[33] Thus, in this subsample the shock only affects the

[32] Falato et al. (2014) define interlock as when two directors not only sit at on the same board but also on the same committee.

[33] It is possible for a firm to have both a current and past connection to the deceased director. This would occur if one of a firm's directors is currently interlocked with the deceased director's firm and another director previously sat on a board with the deceased director. However, we verify that this does not occur for any of the treated firms in this study. Thus, treated firms have either a current or past connections to the deceased director, but not both.

director's network and not his busyness allowing us to better identify the effect of the negative network shock on firm value.

Panel B of Table 18 displays the subsample of treated firms that only have past connection with the deceased directors, and thus no busyness-effect contamination, and their matched control firms. We can see that treated and control groups have comparable means and medians for most observable pre-shock firm characteristics. Pre-shock network centrality measures are statistically different between treated and control firms, but are mostly economically similar. This is comforting as it suggests that the parallel trends assumption is more likely to hold in this setting.

Table 19 Panel B reports the univariate results. First, treated firms exhibit economically and statistically significant drops in firm value during the event windows. On the announcement date firm value drops 0.55%. If we expand the window to also include the day after the announcement we find an even larger 0.91% drop in firm value. This result is highly robust to different definitions of the event window. Second, we also find that treated firms experience significantly larger deceases in value relative to control firms. We find a -0.69% event-day DID estimate indicating a significant decrease in firm value. When we expand the event window, we find DID CAR estimates of -1.09%, -1.07%, and -0.90% for the (0, +1), (-1, +1), and (-2, +2) event windows, respectively, all statistically significant at the 1% level.

These results continue to hold in a multivariate setting. We use the same specifications as in Table 20, but only retain treated firms (and their matched control firms) that have a past connection with the deceased directors. In Panel A of Table 21 we regress abnormal returns on a treatment dummy, controls variables and industry fixed effects using the SUR methodology. The coefficient on the Treatment dummy is negative and statistically significant in all event windows (except having a p-value of 0.103 in the (-2,+2) window). This indicates that firms that experience a negative network shock see decreases in firm value relative to control firms, despite having no shock to busyness. For the announcement date this differential decline in value is 0.41%. This effect becomes more pronounced in the (0, +1) event window, decreasing by 0.69%. Both of these effects are significant at the 5% significance level. The (-1, +1) event window shows a 0.63% differential in abnormal returns, with a p-value of 0.057.

We repeat the pair difference regressions from Table 20 (i.e., each pair is a treated firm and its matched control firms). The results are in Panel B of Table 21. Here the intercept (the DID estimate) remains fairly similar to the pooled specification above, however the statistical significance increases. We find a -0.37% differential change in firm value on the event-day, a -0.68% change in the (0, +1) event window, a -0.64% change in the (-1, +1) event window, and a -0.61% change in the (-2, +2) event window – all statistically significant at either the 5 or 1% level. Using the average market capitalization of this treated subsample, the economic magnitude of these abnormal returns ranges from $25 to 46 million per firm – value which is being lost due to the negative network shock. It is also important to note that we also see a highly significant change in degree centrality for this subsample. This suggests that the network shock had a material adverse effect on board connectedness, and further suggest that the observed decline in firm value is indeed due to the change in director centrality.

In sum, using a subsample of treated firms which were only connected to the deceased director in the past allows us to tease out the effect of a loss in network connectedness on firm value without confounding change in busyness. The results suggest that network shocks lead to reductions in firm value independently of any busyness effect. This is not to suggest that busyness is not important or that busyness does not have adverse effects on firm value as found in Falato et al. (2014) suggest. Even though we show that network shocks reduce firm value, it is likely that both factors matter in practice.

We perform some additional tests to further separate the busyness and network channels. Following Falato et al. (2014), we postulate that if the deceased directors who sat on smaller committees, then the sudden deaths should have a greater impact the busyness of the interlocked directors. In contrast, if the deceased director sat on a larger committee the shock to busyness of the interlocked director is smaller. Thus, if busyness is driving our results we expect to find that our results are stronger when the deceased director sat on a smaller committee. To accomplish this, using a triple difference methodology where third difference is whether the deceased directors sat on small or large committees.

To implement these tests, we first find that the median committee size of the seven deceased directors is 7 directors. We then create a dummy variable (*Big Committee*) that equals 1 if the observation is related to the death of a director whose average committee size was greater than the median of 7, and 0 otherwise. Next, we use the paired-difference specification for the regression so that the dummy variable can directly be interpreted as the differential impact of the busyness effect. In contrast with the previous test, we also limit the sample to firms that were currently interlocked. This is done as past connections are unaffected by busyness.

If the busyness effect is prevalent in our sample, the *Big Committee* dummy variable will be significantly positive in these regressions with CARs as the dependent variable. A positive coefficient implies that the effect of the network shocks on firm value is attenuated when the deceased director sits on larger committees. In Table 22 we run these triple difference tests using SUR regressions, and the same control variables as previously, but adding industry fixed-effects since the constant is no longer necessary for interpretation.

Interestingly, we do not find that the *Big Committee* dummy variable is significantly positive. In fact, the dummy variable is negative in all specifications, and significant at the 5% level when looking at cumulative abnormal returns in the (-2, +2) event window. This is inconsistent with the busyness channel. In Panel B of Table 22 we run similar triple difference regressions except that we have the matched-difference in changes in centrality measures on the left hand side. We find that in terms of degree, closeness, and betweenness centrality, the sudden deaths of the directors on big-committees shocked the treated firms significantly more than the deaths of the directors on small committees. This is consistent with larger shocks to director networks leading to larger value-effects for *Big Committee* firms. Eigenvector centrality, however was shocked differentially more, suggesting that while the deceased directors on big committees were more connected and central, they may not have been as important. In sum, the results of the SUR regressions in Table 22 provide addition evidence that the value effect which we observe is not due to interlocked director busyness, but instead due to the severing of ties in director networks.

6. Concluding Remarks

We use exogenous variation provided by the sudden death of well-connected directors to isolate the impact of board connectedness on firm value. To this end we study the abnormal returns of interlocked firms, whose interlocked director suffers a negative shock to his network of board connections, relative to control firms who are unaffected by the shock. We find that the negative network shock leads to about a 0.6% decrease in firm value.

Our approach sidesteps many of the identification challenges faced by other papers. Given that the director deaths we study are unexpected and sudden, the variation in director networks we study is unlikely to be correlated with important omitted variables that affect firm value. Moreover, the sudden deaths break up the endogenous matching in the director labor market, whereby highly connected directors choose better performing firms, making reverse causality less of a concern in our setting. By focusing our analysis on the interlocked firms (and not the deceased director's firm), we are able to isolate the impact of director networks from potential confounding variables such as director talent and experience. Finally, we find, by studying past connections, we find that our results are not an artifact of the increase in busyness of interlocked directors following the sudden deaths.

Our findings are important as it is difficult to draw causal inference between board characteristics, such as director networks, and shareholder value. Moreover, the recent regulatory interest in this area makes our findings topical and highly relevant as it suggests firm performance can be improved by having a better connected board. We acknowledge that our test does not allow us to disentangle the specific channel through which director networks affect value. For instance, the loss of connection could lead to a loss of access to information. Or it could be due to a decline in the power and influence of the director. Thus, in future research it would be interesting to ascertain why director networks are valuable.

References

Acharya, V.V., Pedersen, L. H., 2005. Asset pricing with liquidity risk. Journal of Financial Economics 77, 375-410.

Amihud, Y., 2002. Illiquidity and Stock Returns: Cross-section and Time-series Effects. Journal of Financial Markets 5, 31-56.

Amihud, Y., Mendelson, H., 1980. Dealership market: Market-making with inventory. Journal of Financial Economics 8, 31-53.

Anand, A., Hua, J., and McCormick, T., 2016. Make-Take Structure and Market Quality: Evidence from the US Options Markets. Management Science 62, 3271-3290.

Angel, J. J., Harris, L. E., Spatt, C. S., 2011. Equity trading in the 21st century. The Quarterly Journal of Finance 1, 1-53.

Armstrong, C., Larcker, D., 2009. Discussion of "The impact of the options backdating scandal on shareholders" and "Taxes and the backdating of stock option exercise dates". Journal of Accounting and Economics 47, 50-58.

Bao, J., Pan, J., Wang, J., 2011. The Illiquidity of Corporate Bonds. Journal of Finance 66, 911-946.

Barnea, A., Guedj, I., 2009. Director networks. Working Paper.

Battalio, R., Corwin, S. A., and Jennings, R., 2016. Can Brokers Have it All? On the Relation between Make- Take Fees and Limit Order Execution Quality. The Journal of Finance 71, 2193-2238.

Bessembinder, H., Kahle, K., Maxwell, W., Xu, D., 2009. Measuring Abnormal Bond Performance. Review of Financial Studies 22, 4219-4258.

Boehmer, E., and Kelley, E. K., 2009. Institutional investors and the informational efficiency of prices. Review of Financial Studies 22, 3563-3594.

Bizjak, J., Lemmon, M., Whitby, R., 2009. Option backdating and board interlocks. Review of Financial Studies 22, 4821-4847.

Black, J. R., Stock, D., Yadav, P. K., 2016. The pricing of different dimensions of liquidity: Evidence from government guaranteed bonds. Journal of Banking & Finance 71, 119-132.

Borgatti, S., Everett, M., Freeman, L., 2002. UCINET 5 for Windows: Software for social network analysis. Analytic Technologies, Harvard, MA.

Brennan, M. J., Subrahmanyam, A., 1996. Market microstructure and asset pricing: On the Compensation for Illiquidity in Stock Returns. Journal of Financial Economics, 41, 441-464.

Brogaard, J., Hendershott, T., and Riordan, R., 2014. High-frequency trading and price discovery. Review of Financial Studies 27, 2267-2306.

Brolley, M., and Malinova, K., 2013. Informed trading and maker-taker fees in a low-latency limit order market. Working Paper. Available at SSRN 2178102.

Cardella, L., Hao, J., Kalcheva, I., 2015. Make and take fees in the US equity market. Working Paper. Available at SSRN 2149302.

Chordia, T., Roll, R., Subrahmanyam, A., 2000 (a). Co-Movements in Bid-Ask Spreads and Market Depth. Financial Analysts Journal 56, 23-27.

Chordia, T., Roll, R., Subrahmanyam, A., 2000 (b). Commonality in Liquidity. Journal of Financial Economics 56, 3-28.

Chordia, T., Roll, R., and Subrahmanyam, A., 2008. Liquidity and market efficiency. Journal of Financial Economics 87, 249-268.

Chordia, T., Roll, R., and Subrahmanyam, A., 2011. Recent trends in trading activity and market quality. Journal of Financial Economics 101, 243-263.

Chung, D., and Hrazdil, K., 2010. Liquidity and market efficiency: A large sample study. Journal of Banking & Finance 34, 2346-2357.

Cohen, L., Frazzini, A., Malloy, C., 2008. The small world of investing: The use of social networks in bank decision-making. Journal of Political Economy 116, 951-979.

Cohen, L., Frazzini, A., Malloy, C., 2010. Sell- side school ties. The Journal of Finance 65, 1409-1437.

Colliard, J. E., and Foucault, T., 2012. Trading fees and efficiency in limit order markets. Review of Financial Studies 25, 3389-3421.

Core, J., Holthausen, R., Larcker, D., 1999. Corporate governance, chief executive officer compensation, and firm performance. Journal of financial economics 51, 371-406.

Crespí-Cladera, R., Pascual-Fuster, B., 2015. Optimal board independence with non-strictly independent directors, Working Paper.

DeMarzo, P., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and uni-dimensional opinions. The Quarterly Journal of Economics 118, 909-968.

Dick-Nielsen, J., 2009. Liquidity Biases in TRACE. Journal of Fixed Income 19, 43-55.

Dick-Nielsen, J., Feldhütter, P., Lando, D., 2012. Corporate Bond Liquidity Before and After the Onset of the Subprime Crisis. Journal of Financial Economics 103, 471-492.

Dong, J., Kempf, A., Yadav, P.K., 2010. Resiliency, the Neglected Dimension of Market Liquidity: Empirical Evidence from the New York Stock Exchange. Working Paper.

Elton, E., Gruber, M., Agrawal, D., Mann, C., 2001. Explaining the Rate Spread on Corporate Bonds. Journal of Finance 56, 247-277.

Engelberg, J., Gao, P., Parsons, C., 2013. The price of a CEO's rolodex. Review of Financial Studies 26, 79-114.

Ericcson, J., Renault, O., 2006. Liquidity and Credit Risk. Journal of Finance 61, 2219-2250.

Falato, A., Kadyrzhanova, D., Lel, U., 2014. Distracted directors: Does board busyness hurt shareholder value? Journal of Financial Economics 113, 404-426.

Faleye, O., Kovacs, T., Venkateswaran, A., 2014. Do better-connected CEOs innovate more? Journal of Financial and Quantitative Analysis 49, 1201-1225.

Fama, E. F., 1965. The behavior of stock-market prices. The Journal of Business 38, 34-105.

Fama, E. F., and French, K. R., 1988. Permanent and temporary components of stock prices. Journal of political Economy 96, 246-273.

Fama, E., French, K., 1993. Common risk factors in the returns on stocks and bonds. Journal of Financial Economics 33, 3-56.

Federal Register, Part VII, FDIC, 12 CFR Part 370, Temporary Liquidity Guarantee Program; Final Rule, November 26, 2008.

Fich, E., Shivdasani, A., 2006. Are busy boards effective monitors? The Journal of Finance 61, 689-724.

Fich, E., White, L., 2003. CEO compensation and turnover: The effects of mutually interlocked boards. Wake Forest L. Rev. 38, 935.

Fogel, K., Ma, L., Morck, R., 2014. Powerful independent directors, Working Paper.

Fotak, V., Raman, V., Yadav, P. K., 2014. Fails-to-deliver, short selling, and market quality. Journal of Financial Economics 114, 493-516.

Foucault, T., Kadan, O., Kandel, E., 2005. Limit order book as a market for liquidity. Review of Financial Studies 18, 1171-1217.

Friewald, N., Jankowitsch, R., Subrahmanyam, M., 2012. Illiquidity or credit deterioration: A Study of Liquidity in the US Corporate Bond Market during Financial Crises. Journal of Financial Economics 105, 18-36.

Garbade, K., 1982. Securities Markets. New York: McGraw-Hill.

Glosten, L.R., Harris, L.E., 1988. Estimating the components of the bid/ask spread. Journal of Financial Economics 21, 123-142.

Glosten L.R., Milgrom P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71-100.

Hasbrouck, J., 1991. The summary informativeness of stock trades: An econometric analysis. Review of Financial Studies 4, 571-595.

Hasbrouck, J., 1993. Assessing the quality of a security market: A new approach to transaction-cost measurement. Review of Financial Studies 6, 191-212.

Hansch, O., Naik, N.Y., Viswanathan, S., 1998. Do inventories matter in dealership markets? Evidence from the London Stock Exchange. The Journal of Finance 53, 1623-1656.

Harris, L., 1990, Liquidity, trading rules and electronic trading systems, New York University Salomon Center Monograph Series in Finance and Economics 1990-4.

Harris, L., 2003. Trading & Exchanges: Market Microstructure for Practitioners. Oxford University Press.

He, Milbradt, 2014. Endogenous liquidity and defaultable bonds. Econometrica 82, 1443-1508.

He, Z., Xiong, W., 2012. Rollover Risk and Credit Risk. The Journal of Finance 67, 391-430.

Helmers, C., Patnam, M., Rau, P., 2015. Do board interlocks increase innovation? Evidence from a corporate governance reform in India? Evidence from a Corporate Governance Reform in India (April 19, 2015).

Ho, T., Stoll, H. R., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47-73.

Ho, T.S., Stoll, H. R., 1983. The dynamics of dealer markets under competition. The Journal of Finance 38, 1053-1074.

Hochberg, Y., Ljungqvist, A, Lu, Y., 2007. Whom you know matters: Venture capital networks and investment performance. The Journal of Finance 62, 251-301.

Holden, C.W., Jacobsen, S.E., Subrahmanyam, A., 2014. The empirical analysis of liquidity. Kelley School of Business Research Paper.

Hong, G., Warga, A., 2000. An Empirical Study of Bond Market Transactions. Financial Analysts Journal 56, 32-46.

Horton, J., Millo, Y., Serafeim, G., 2012. Resources or power? Implications of social networks on compensation and firm performance. Journal of Business Finance & Accounting 39, 399-426.

Kempf, A., Mayston, D., Gehde-Trapp, M., Yadav, P., 2015. Resiliency, a Dynamic View of Liquidity. Working Paper.

Kyle, A, 1985. Continuous Auctions and Insider Trading. Econometrica 6, 1315-1335.

Larcker, D., So, E., Wang, C., 2013. Boardroom centrality and firm performance. Journal of Accounting and Economics 55, 225-250.

Lee, C., and Ready, M. J., 1991. Inferring trade direction from intraday data. The Journal of Finance 46, 733-746.

Lee, S.H.M., Cotte, J., Noseworthy. T.J., 2010. The role of network centrality in the flow of consumer influence. Journal of Consumer Psychology 20, 66-77.

Lin, H., Wang, J., Wu, C., 2011.Liquidity risk and expected corporate bond returns. Journal of Financial Economics 99, 628-650.

Loderer, C., Peyer, U., 2002. Board overlap, seat accumulation and share prices. European Financial Management 8, 165-192.

Longstaff, F., 2004. The Flight-to-Liquidity Premium in U.S. Treasury Bond Prices. Journal of Business 77, 511-526.

Longstaff, F., Mithal, S., Neis, E., 2005. Corporate Yield Spreads: Default Risk or Liquidity? New Evidence from the Credit Default Swap Market. Journal of Finance 50, 2213-2253.

Lutat, M., 2010. The effect of maker-taker pricing on market liquidity in electronic trading systems–empirical evidence from European equity trading. Working Paper. Available at SSRN 1752843.

Malinova, K., Park, A., 2015. Subsidizing liquidity: The impact of make/take fees on market quality. The Journal of Finance 70, 509-536.

Masulis, R.W., Mobbs, S., 2014. Independent director incentives: Where do talented directors spend their limited time and energy?. Journal of Financial Economics, 111, 406-429.

Menkveld, A. J., 2013. High frequency trading and the new market makers. Journal of Financial Markets 16, 712-740.

Mizruchi, M., 1990. Cohesion, structural equivalence, and similarity of behavior: An approach to the study of corporate political power. Sociological Theory 8, 16-32.

Mol, M., 2001. Creating wealth through working with others: Inter-organizational relationships. The Academy of Management Executive 15, 150-152.

Naik, N.Y, Yadav, P.K., 2003. Do dealer firms manage inventory on a stock-by-stock or a portfolio basis?. Journal of Financial Economics 69, 325-353.

NASDAQ Stock Market. (2014, Dec. 19) Form SR-NASDAQ-2014-128. Retrieved from https://www.sec.gov/rules/sro/nasdaq/nasdaqarchive/nasdaqarchive2014.shtml.

Obizhaeva, A., Wang, J., 2012. Optimal trading strategy and supply/demand dynamics. Journal of Financial Markets 16, 1-32.

Omer, T., Shelley, M., Tice, F., 2014. Do well-connected directors affect firm value? Journal of Applied Finance 24, 17-32.

Pástor, L., Stambaugh, R. F., 2003. Liquidity Risk and Expected Stock Returns. Journal of Political Economy 111, 642-685.

Pedersen, M., 2009. Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches. Review of Financial Studies 22, 435-480.

Pennings, J., 1980. Interlocking directorates. Jossey-Bass Inc Pub.

Renneboog, L., Zhao, Y., 2011. Us knows us in the UK: On director networks and CEO compensation. Journal of Corporate Finance 17, 1132-1157.

Roberts, M. R. Whited, T.M., 2013. Endogeneity in Empirical Corporate Finance. Handbook of the Economics of Finance.

SEC. (2016, June 10) Recommendation for an Access Fee Pilot. Retrieved from https://www.sec.gov/spotlight/emsac/emsac-regulation-nms-recommendation-61016.pdf.

Skjeltorp, J. A., Sojli, E., and Tham, W. W. (2012, June). Identifying cross-sided liquidity externalities. In Asian Finance Association 2013 Conference.

Snyder, P., Priem, R., Levitas, E., 2009, August. The diffusion of illegal innovations among management elites. In Academy of Management Proceedings 1, 1-6.

Stern, L., 2015. A learning-based approach to evaluating boards of directors. Working Paper.

Stoll, H.R., 1989. Inferring the components of the bid- ask spread: theory and empirical tests. The Journal of Finance 44, 115-134.

Tetlock, P. C., 2008. Liquidity and prediction market efficiency. Working paper. Available at SSRN: http://dx.doi.org/10.2139/ssrn.929916.

Appendix A: Tables and Figures

Table 1: Propensity Score Matching Probit

This table displays the regression coefficients and (mean) marginal effects of the probit model used for propensity score matching. Stock data is from TAQ and CRSP. Stocks are excluded from this the regression if they are missing data on mispricing, time inside the NBBO, adverse selection, volume, price, bid-ask spread, or industry (SIC code). Stocks were also excluded if the pre-shock window contained less than 80 trading days, was classified as a financial vehicle (NAICS 525990), the stock was not listed on either the NASDAQ or NYSE or had a price difference with the treated stock of more than 10%. Variables are averaged over the 4-month pre-experiment period. Treated Dummy is a binary variable equal to 1 if the stock was affected by the NASDAQ Access Fee Experiment, and 0 otherwise. Other variables are described in Section 3. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	Treated Dummy	**Marginal Effects** (× 100,000)
Constant	-4.1344***	-1.6492
	(0.000)	
Nasdaq Volume	2.2E-6***	8.74E-7
	(0.000)	
Bid-Ask Spread ($)	0.0637	0.0254
	(0.939)	
Price	-0.0074	-0.0029
	(0.207)	
MAPE	1.8555	0.7402
	(0.582)	
Psuedo R^2	0.549	
Treated Firms	14	
Untreated Firms	7559	

Table 2: Matched Sample Pre-shock Comparison

This table contains comparisons of the means and medians of the treated and matched-control pre-shock samples on a stock-day level. Potential match stocks are dropped from the sample if they are missing data on mispricing, time inside the NBBO, adverse selection, volume, price, bid-ask spread, or industry (SIC code). Potential matches were also dropped if the pre-shock window contained less than 80 trading days, or was classified as a financial vehicle (NAICS 525990). Each treated stock is then matched with five untreated stocks, with replacement, based on price and propensity score – the fitted value of the probit displayed in Table 1.

Variable	Treated			Control			Diff
	Obs.	Mean	Med.	Obs.	Mean	Med.	T
Best Ask is on Nasdaq	1,176	0.4890	0.4377	5,876	0.4492	0.3922	5.9062
Best Bid is on Nasdaq	1,176	0.4889	0.4395	5,876	0.4524	0.3969	5.5359
Best Bid and Ask are on Nasdaq	1,176	0.2396	0.1321	5,876	0.1890	0.0619	7.0173
Best Bid or Ask are on Nasdaq	1,176	0.7332	0.7416	5,876	0.7106	0.7162	3.7465
Adverse Selection Cost (1 min)	1,176	0.0017	0.0003	5,876	0.0006	0.0002	3.7609
Adverse Selection Cost (15 min)	1,176	0.0017	0.0003	5,876	0.0005	0.0002	4.0706
Adverse Selection Cost (30 min)	1,176	0.0012	0.0003	5,876	0.0005	0.0002	2.5642
Adverse Selection Cost (60 min)	1,176	0.0012	0.0003	5,876	0.0005	0.0002	2.7516
Mean Absolute Pricing Error (MAPE)	1,176	0.0219	0.0003	5,876	0.0048	0.0003	10.4240
Variance of Mispricing (σ_{ϕ}^{2})	1,176	0.0452	0.0000	5,876	0.0028	0.0000	9.6881
Nasdaq Volume Share	1,176	0.3659	0.3411	5,876	0.3583	0.3411	3.0087
Log(Nasdaq Volume)	1,176	14.1028	14.2802	5,876	13.7706	13.9284	28.3482
Log(Volume)	1,176	15.0780	15.0943	5,876	14.5286	14.6080	18.4594
Price	1,176	24.8315	25.3363	5,876	24.4064	23.1726	0.6863
Bid-ask Spread (S)	1,176	0.0228	0.0119	5,876	0.0133	0.0107	8.8361
Bid-ask Spread (%)	1,176	0.0014	0.0009	5,876	0.0013	0.0005	2.5120

Table 3: Pricing Efficiency Effect

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to make-take fee level on market efficiency. The dependent variables are the mean absolute pricing error (MAPE) and variance of pricing error innovations on a stock-day level. The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, a dummy variable equaling 1 for observations during the experiment, and an interaction of the two dummy variables as well as control variables described in Section 3. Standard errors for these panel regressions are clustered by stock and date. Two-tailed *p*-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	MAPE	**MAPE**	$\boldsymbol{\sigma^2_{\phi}}$	$\boldsymbol{\sigma^2_{\phi}}$
Treated Dummy	0.0180**	0.0176**	0.0613**	0.0606**
	(0.024)	(0.025)	(0.016)	(0.016)
Experiment Dummy	0.0002	-0.0007**	0.0003	-0.0007**
	(0.252)	(0.040)	(0.329)	(0.040)
Treated x Experiment	**0.0017*** **	**0.0017*** **	**0.0046** **	**0.0049** **
	(0.000)	**(0.000)**	**(0.032)**	**(0.029)**
Price		-0.0000		-0.0000
		(0.972)		(0.920)
Bid-Ask Spread (%)		0.1018		-0.8604
		(0.853)		(0.587)
Log(Volume)		0.0009		0.0009
		(0.124)		(0.150)
Constant	0.0054**	-0.0074	0.0038	-0.0068
	(0.040)	(0.166)	(0.104)	(0.285)
Obs.	20,952	20,919	20,952	20,919
R^2	0.0418	0.0438	0.0613	0.0617

Table 4: Treatment and Reversal Effects

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to make-take fee level on market efficiency. The dependent variables are the mean absolute pricing error (MAPE) and variance of pricing error innovations on a stock-day level. The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, a dummy variable equaling 1 for observations during the experiment, and an interaction of the two dummy variables as well as control variables described in Section 3. The sample in Panel A contains observations from Oct. 2014 – May 2015 (omitting observations after the NASDAQ experiment) to test the effect of instituting the pilot. Panel B contains observations from Feb. 2015 – Sept. 2015 (omitting observations before the NASDAQ experiment) to test for the reversal effect after the pilot ceases. Standard errors for these panel regressions are clustered by stock and date. Two-tailed *p*-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Oct. 2014 – May 2015 (Before and During Pilot)

	MAPE	**MAPE**	$\boldsymbol{\sigma_\phi^2}$	$\boldsymbol{\sigma_\phi^2}$
Treated Dummy	0.0171**	0.0150**	0.0424**	0.0397**
	(0.024)	(0.029)	(0.016)	(0.017)
Experiment Dummy	0.0008**	0.0017*	0.0013**	0.0024*
	(0.017)	(0.059)	(0.011)	(0.078)
Treated x Experiment	**0.0026*** **	**0.0027** **	**0.0235*** **	**0.0237** **
	(0.000)	**(0.012)**	**(0.000)**	**(0.013)**
Price		0.0001		0.0001
		(0.495)		(0.593)
Bid-Ask Spread (%)		1.7726*		2.2134
		(0.093)		(0.202)
Log(Volume)		0.0062*		0.0080
		(0.070)		(0.110)
Constant	0.0048**	-0.0905*	0.0028*	-0.1212*
	(0.041)	(0.065)	(0.099)	(0.096)
Obs.	13,901	13,901	13,901	13,901
R^2	0.0430	0.0566	0.0699	0.0741

Panel B: Feb. 2015 – Sept. 2015 (During and After Pilot)

	MAPE	**MAPE**	$\boldsymbol{\sigma_\phi^2}$	$\boldsymbol{\sigma_\phi^2}$
Treated Dummy	0.0189**	0.0180**	0.0798	0.0775
	(0.025)	(0.026)	(0.159)	(0.161)
Experiment Dummy	-0.0004	-0.0032*	-0.0007*	-0.0064**
	(0.115)	(0.076)	(0.063)	(0.050)
Treated x Experiment	**0.0008** **	**0.0012*** **	**-0.0139**	**-0.0126**
	(0.042)	**(0.004)**	**(0.181)**	**(0.153)**
Price		-0.0000		-0.0001
		(0.862)		(0.874)
Bid-Ask Spread (%)		-0.2004		-1.3039
		(0.746)		(0.516)
Log(Volume)		0.0013		0.0025*
		(0.112)		(0.064)
Constant	0.0060**	-0.0094	0.0048	-0.0235*
	(0.039)	(0.129)	(0.106)	(0.057)
Obs.	13,900	13,867	13,900	13,867
R^2	0.0417	0.0457	0.0665	0.0686

Table 5: Percentage of NASDAQ quotes inside the NBBO

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to make-take fee level on the percentage of NASDAQ quotes inside the NBBO. The samples for all regressions include observations from Oct. 2014 – Sept. 2015. The dependent variables are the amount of quotes on the NASDAQ which were the best bid, ask, either, or both. The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, a dummy variable equaling 1 for observations during the experiment, and an interaction of the two dummy variables. Panel A contains no controls, while Panel B contains controls described in Section 3. Standard errors for these panel regressions are clustered by stock and date. Two-tailed p-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Without Controls

	Best Bid	**Best Ask**	**Best Either**	**Best Both**
Treated Dummy	0.0142	0.0126	-0.0039	0.0291
	(0.775)	(0.798)	(0.935)	(0.573)
Experiment Dummy	-0.0078	-0.0070	-0.0130	-0.0010
	(0.219)	(0.271)	(0.175)	(0.793)
Treated x Experiment	**-0.1988*** **	**-0.1988*** **	**-0.2716*** **	**-0.1243*** **
	(0.000)	**(0.000)**	**(0.000)**	**(0.000)**
Constant	0.4197***	0.4226***	0.6791***	0.1621***
	(0.000)	(0.000)	(0.000)	(0.000)
Obs.	20,966	20,966	20,966	20,966
R^2	0.0490	0.0470	0.0953	0.0149

Panel B: With Controls

	Best Bid	**Best Ask**	**Best Either**	**Best Both**
Treated Dummy	0.0138	0.0125	-0.0048	0.0296
	(0.780)	(0.799)	(0.920)	(0.564)
Experiment Dummy	-0.0101	-0.0090	-0.0183*	0.0002
	(0.168)	(0.217)	(0.065)	(0.972)
Treated x Experiment	**-0.1987*** **	**-0.1987*** **	**-0.2712*** **	**-0.1244*** **
	(0.000)	**(0.000)**	**(0.000)**	**(0.000)**
Price	0.0005	0.0006	0.0013	-0.0001
	(0.556)	(0.509)	(0.159)	(0.933)
Log(Volume)	0.0014	0.0010	0.0037	-0.0014
	(0.769)	(0.841)	(0.463)	(0.784)
Constant	0.3875***	0.3944***	0.5987***	0.1842**
	(0.000)	(0.000)	(0.000)	(0.016)
Obs.	20,933	20,933	20,933	20,933
R^2	0.0522	0.0508	0.1110	0.0152

Table 6: Informed Trading Effect

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to make-take fee level on informed trading. Panel A contains no controls, while Panel B contains controls described in Section 3. The samples for all regressions include observations from Oct. 2014 – Sept. 2015. The dependent variables are the 1-, 15-, 30-, and 60-minute adverse selection costs (average losses of market makers due to informed trading). The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, a dummy variable equaling 1 for observations during the experiment, and an interaction of the two dummy variables. Standard errors for these panel regressions are clustered by stock and date. Two-tailed p-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Without Controls

	ASC_1 (x1,000)	**ASC_{15} (x1,000)**	**ASC_{30} (x1,000)**	**ASC_{60} (x1,000)**
Treated Dummy	0.5148**	0.4622	0.2636	0.2910*
	(0.029)	(0.352)	(0.135)	(0.100)
Experiment Dummy	-0.1302	-0.0510	-0.0816	-0.0457
	(0.126)	(0.434)	(0.381)	(0.644)
Treated x Experiment	**-0.1426***	**-0.7452****	**-0.8556*****	**-0.3626**
	(0.067)	**(0.027)**	**(0.008)**	**(0.267)**
Constant	0.6799***	0.6013***	0.5866***	0.5389***
	(0.000)	(0.000)	(0.000)	(0.000)
Obs.	20,930	20,928	20,924	20,926
R^2	0.0018	0.0029	0.0024	0.0010

Panel B: With Controls

	ASC_1 (x1,000)	**ASC_{15} (x1,000)**	**ASC_{30} (x1,000)**	**ASC_{60} (x1,000)**
Treated Dummy	0.5262**	0.4750	0.2871	0.3047*
	(0.027)	(0.340)	(0.102)	(0.085)
Experiment Dummy	-0.0707	0.0085	0.0051	0.0115
	(0.279)	(0.870)	(0.931)	(0.858)
Treated x Experiment	**-0.1385***	**-0.7413****	**-0.8504*****	**-0.3590**
	(0.068)	**(0.027)**	**(0.009)**	**(0.272)**
Price	-0.0149***	-0.0132**	-0.0133***	-0.0109***
	(0.002)	(0.014)	(0.004)	(0.000)
Log(Volume)	-0.0472	-0.0490	-0.0767*	-0.0487**
	(0.320)	(0.250)	(0.074)	(0.028)
Constant	0.0017**	0.0016**	0.0020***	0.0015***
	(0.015)	(0.013)	(0.002)	(0.000)
Obs.	20,930	20,928	20,924	20,926
R^2	0.0052	0.0053	0.0057	0.0030

Table 7: Volume Effect

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to make-take fee level on exchange volumes. The samples for all regressions include observations from Oct. 2014 – Sept. 2015. The dependent variables are the NASDAQ market share, log of NASDAQ volume, and log of total volume on other exchanges (not NASDAQ). The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, a dummy variable equaling 1 for observations during the experiment, and an interaction of the two dummy variables. Panel A contains no controls, while Panel B contains controls described in Section 3. Standard errors for these panel regressions are clustered by stock and date. Two-tailed p-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Without Controls

	Nasdaq Volume Share	**Log(Nasdaq Volume)**	**Log(Volume)**
Treated Dummy	0.1354**	0.5705***	0.1802
	(0.043)	(0.001)	(0.126)
Experiment Dummy	-0.0083	0.9491***	1.0209***
	(0.368)	(0.000)	(0.000)
Treated x Experiment	**-0.0335****	**-0.1145***	**0.0582****
	(0.021)	**(0.074)**	**(0.012)**
Constant	0.4227***	12.7601***	13.0925***
	(0.000)	(0.000)	(0.000)
Obs.	20,933	20,902	20,207
R^2	0.0284	0.0663	0.0784

Panel B: With Controls

	Nasdaq Volume Share	**Log(Nasdaq Volume)**	**Log(Volume)**
Treated Dummy	0.1362**	0.5776***	0.1798
	(0.043)	(0.003)	(0.178)
Experiment Dummy	-0.0098	0.9381***	1.0153***
	(0.285)	(0.000)	(0.000)
Treated x Experiment	**-0.0333****	**-0.1136***	**0.0586*****
	(0.021)	**(0.090)**	**(0.004)**
Price	0.0016***	0.0132***	0.0060
	(0.005)	(0.009)	(0.201)
Constant	0.3841***	12.4470***	12.9523***
	(0.000)	(0.000)	(0.000)
Obs.	20,933	20,902	20,207
R^2	0.0403	0.0842	0.0826

Table 8: Falsification Tests

This table displays results for the multivariate falsification (placebo) analysis examining whether market efficiency changed prior to the NASDAQ pilot for the treated stocks. The dependent variables are the mean absolute pricing error (MAPE) and variance of pricing error innovations on a stock-day level. The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, three dummy variables equaling 1 for observations during November, December, and January, and an interaction of the dummy variables as well as control variables described in Section 3. Standard errors for these panel regressions are clustered by stock and date. Two-tailed *p*-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	MAPE	**MAPE**	$\boldsymbol{\sigma_{\phi}^2}$	$\boldsymbol{\sigma_{\phi}^2}$
Treated Dummy	0.0134	0.0106	0.0246	0.0205
	(0.254)	(0.340)	(0.184)	(0.222)
November	0.0007	0.0016**	0.0012	0.0024*
	(0.364)	(0.025)	(0.121)	(0.085)
December	0.0012**	0.0013**	0.0005**	0.0004
	(0.021)	(0.045)	(0.011)	(0.622)
January	-0.0007	-0.0008	-0.0011*	-0.0014*
	(0.360)	(0.351)	(0.071)	(0.098)
November x Treated	**0.0039***	**0.0034**	**0.0272***	**0.0264***
	(0.070)	**(0.170)**	**(0.098)**	**(0.093)**
December x Treated	**0.0058**	**0.0064**	**0.0287**	**0.0298**
	(0.132)	**(0.181)**	**(0.153)**	**(0.140)**
January x Treated	**0.0057**	**0.0067**	**0.0175**	**0.0193***
	(0.115)	**(0.107)**	**(0.117)**	**(0.084)**
Price		0.0001		0.0002
		(0.324)		(0.319)
Bid-Ask Spread (%)		2.2569**		3.8089*
		(0.042)		(0.068)
Log(Volume)		0.0074**		0.0101**
		(0.024)		(0.045)
Constant	0.0045**	-0.1108**	0.0027*	-0.1571**
	(0.048)	(0.021)	(0.087)	(0.042)
Obs.	7,052	7,052	7,052	7,052
R^2	0.0438	0.0617	0.0637	0.0721

Table 9: Robust Matching Tests

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to make-take fee level on market efficiency using alternate control groups, matched on one of the following: percentage of NASDAQ quotes inside the NBBO, the bid-ask spread (in dollars), the mean absolute pricing error (MAPE), and the percentage of volume traded on the NASDAQ. The dependent variables are MAPE and variance of pricing error innovations on a stock-day level. The dependent variables are regressed on a dummy variable equaling 1 for treated stocks, a dummy variable equaling 1 for observations during the experiment, and an interaction of the two dummy variables as well as control variables described in Section 3. Standard errors for these panel regressions are clustered by stock and date. Two-tailed *p*-values are in parenthesis below the corresponding coefficients. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Match Criteria	MAPE		Nasdaq Volume Share		Bid-Ask Spread ($)		% Quotes at NBBO	
Dependent Variable	**MAPE**	**MAPE**	**MAPE**	**MAPE**	**MAPE**	**MAPE**	**MAPE**	**MAPE**
Treated Dummy	0.0067	-0.0206	0.0209	0.0207	0.0006	0.0067	0.0006	0.0067
	(0.654)	(0.258)	(0.150)	(0.118)	(0.974)	(0.677)	(0.368)	(0.677)
Experiment Dummy	0.0015	0.0020**	-0.0001	-0.0001	0.0020***	0.0013	0.0020	0.0013**
	(0.107)	(0.024)	(0.725)	(0.696)	(0.000)	(0.201)	(0.317)	(0.020)
Treated x Experiment	**0.0018***	**0.0024***	**0.0035*****	**0.0036*****	**0.0014****	**0.0016****	**0.0014****	**0.0016****
	(0.057)	**(0.063)**	**(0.000)**	**(0.000)**	**(0.042)**	**(0.023)**	**(0.042)**	**(0.023)**
Price		0.0000		0.0000		-0.0004		-0.0004
		(0.923)		(0.864)		(0.263)		(0.263)
Bid-Ask Spread (%)		0.3134		0.2816***		1.9106		1.9106
		(0.168)		(0.000)		(0.449)		(0.449)
Log(Volume)		0.0080***		0.0003		-0.0036		-0.0036
		(0.002)		(0.831)		(0.628)		(0.628)
Constant	0.0153***	- 0.0794***	0.0010**	-0.0032	0.0214**	0.0760	0.0214**	0.0760
	(0.001)	(0.002)	(0.011)	(0.794)	(0.016)	(0.494)	(0.016)	(0.494)
Obs.	13,812	13,793	11,732	11,732	13,944	13,944	13,944	13,944
R^2	0.0042	0.0853	0.1014	0.1178	0.0005	0.0192	0.0005	0.0192

Table 10: Chapter 2 Variable Descriptions

Panel A: Descriptive Statistics

	Units	Obs.	Mean	Std. Dev.	Min.	Med.	Max.
Non-Default Spread	Pct. Yield	38,106	0.2072	0.2280	-0.2177	0.1492	1.2779
Bid-Ask Spread	bps of Price	19,386	15.97	21.20	0.00	8.30	118.71
Amihud	%Δ per $1M	40,752	10.839	34.436	0.0000	0.373	250.31
Resiliency		14,925	0.4713	0.2681	0.0210	0.4457	0.9978
Market Bid-Ask Spread	bps of Price	45,712	6.9E-04	8.3E-04	0.0000	4.2E-04	7.5E-03
Market Amihud	%Δ per $1M	46,710	0.0003	0.0007	0.0000	0.0001	0.0078
Market Resiliency		45,602	0.1135	0.1809	-0.5313	0.0662	1.1116
Coupon	Pct. of Par	47,145	2.0223	0.8045	0.2305	2.1250	3.2500
VIX		47,135	24.848	8.177	13.450	22.660	68.510
Volume	Dollars	47,145	1.4E+07	6.4E+07	5.0000	2.7E+06	8.4E+09

(continued on next page)

Panel B: Correlation Matrix *(Bond-specific liquidity variables have been orthogonalized to each other after logarithms are calculated)*

	Non-Default Spread	*ln(Bid-Ask Spread)*	*ln (Amihud)*	*ln (Resil)*	*ln(Market Bid-Ask Spread)*	*Ln (Market Amihud)*	*Ln (Market Resil)*	*Coupon*	*VIX*	*Volume*
Non-Default Spread	1.000	0.292	-0.040	0.140	0.156	0.196	-0.122	0.086	0.771	0.169
ln(Bid-Ask Spread)	0.292	1.000	-0.253	-0.072	0.088	0.108	-0.065	-0.021	0.225	-0.078
ln (Amihud)	-0.040	-0.253	1.000	-0.089	-0.020	-0.003	0.021	0.230	-0.087	-0.152
ln (Resil)	0.140	-0.072	-0.089	1.000	0.033	0.035	-0.014	0.121	0.122	0.103
ln(Market Bid-Ask Spread)	0.156	0.088	-0.020	0.033	1.000	0.103	-0.007	0.005	0.143	0.027
ln(Market Amihud)	0.196	0.108	-0.003	0.035	0.103	1.000	-0.021	0.035	0.170	0.047
ln(Market Resil)	-0.122	-0.065	0.021	-0.014	-0.007	-0.021	1.000	0.022	-0.061	-0.002
Coupon	0.086	-0.021	0.230	0.121	0.005	0.035	0.022	1.000	0.046	-0.064
VIX	0.771	0.225	-0.087	0.122	0.143	0.170	-0.061	0.046	1.000	0.151
Volume	0.169	-0.078	-0.152	0.103	0.027	0.047	-0.002	-0.064	0.151	1.000

Panel C: Observations from Issuer Credit Ratings

Rating	*AAA*	*AA+*	*AA*	*AA-*	*A+*	*A*	*A-*	*BBB+*	*BBB*	*BBB-*	*BB+*	*WORSE*
Bonds	8	20	5	7	35	83	43	10	4	6	3	0
Obs.	193	6,670	261	1,549	7,199	20,925	2,889	1,711	219	740	320	0

Table 11: Pricing of Liquidity Dimensions

This table displays results for the multivariate analysis of the pricing of the three dimensions of liquidity in the non-default spread (NDS). Each of the three proxies for the liquidity dimensions – bid-ask spread, Amihud measure, and resiliency – have been orthoganalized to the other two. Following Elton et al. (2001), the coupon rate controls for state taxes within the non-default spread. Bond-, Day-, and Firm-fixed effects are used as controls in Models (2), (3), and (4), respectively.

The sample for this unbalanced panel regression consists of bonds guaranteed by the FDIC under the DGP. Standard errors are clustered by bond and date. Two-tailed *p*-values are in parenthesis. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Variable	(1) NDS	(2) NDS	(3) NDS	(4) NDS
Ln(Bid-Ask Spread)	0.0590***	0.0597***	0.0064***	0.0597***
orthogonalized	(0.000)	(0.000)	(0.000)	(0.000)
Ln(Amihud)	0.0026	0.0089***	0.0044***	0.0067**
orthogonalized	(0.312)	(0.000)	(0.000)	(0.011)
–Ln(Resiliency)	0.0460***	0.0503***	0.0037**	0.0514***
orthogonalized	(0.000)	(0.000)	(0.034)	(0.000)
Coupon	0.0412***	*Subsumed by Fixed Effects*	0.0051	0.0189
	(0.004)		(0.591)	(0.175)
Constant	0.1467***	*Subsumed by Fixed Effects*		
	(0.000)			
Adj. R^2	0.130	0.183	0.874	0.157
Bonds	65	65	65	65
Days	958	958	958	958
Obs.	10,122	10,122	10,122	10,122
Fixed Effects	None	Bond	Day	Firm

Table 12: Pricing of Market-wide Liquidity Dimensions

This table displays results for the multivariate analysis of the pricing of market-wide dimensions of liquidity in the non-default spread (NDS). Each of the three proxies for the bond-specific liquidity dimensions – bid-ask spread, Amihud measure, and resiliency – have been orthoganalized to the other two. Following Elton et al. (2001), the coupon rate controls for state taxes within the non-default spread. Model 1 presents results without any fixed effects in the regression model specification. Models 2, 3, and 4 present results with bond fixed effects, which is what is directly relevant for the research question being investigated. Models 1 and 2 cover the entire sample period, while Model 3 includes only the financial crisis period and Model 4 includes only the post-financial-crisis period.

The sample for this unbalanced panel regression consists of bonds guaranteed by the FDIC under the DGP. Standard errors are clustered by bond and date. Two-tailed *p*-values are in parenthesis. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	**(2)**	**(3)**	**(4)**
Variable	**NDS**	**NDS**	**NDS**	**NDS**
Ln(Bid-Ask Spread)	0.0485***	0.0501***	0.0326***	0.0184***
Orthogonalized	(0.000)	(0.000)	(0.000)	(0.000)
Ln(Amihud)	0.0030	0.0081***	0.0003	0.0067***
Orthogonalized	(0.192)	(0.000)	(0.928)	(0.000)
–Ln(Resiliency)	0.0403***	0.0450***	0.0411***	0.0096***
Orthogonalized	(0.000)	(0.000)	(0.000)	(0.000)
Ln(Market Spread)	0.0051***	0.0045***	0.0083***	0.0012***
	(0.000)	(0.000)	(0.000)	(0.001)
Ln(Market Amihud)	0.0296***	0.0263***	0.0128	0.0103***
	(0.000)	(0.000)	(0.370)	(0.000)
–Ln(Market Resiliency)	0.1604***	0.1550***	0.0631	0.0512***
	(0.000)	(0.000)	(0.702)	(0.001)
Coupon	0.0329***	*Subsumed by Fixed Effects*		
	(0.000)			
Constant	0.5028***	*Subsumed by Fixed Effects*		
	(0.006)			
Adj. R^2	0.164	0.212	0.173	0.138
Bonds	65	65	64	64
Days	887	887	254	633
Obs.	9,419	9,419	3,842	5,577
Sample	Full	Full	Crisis	Post-Crisis
Fixed Effects	None	Bond	Bond	Bond

Table 13: Pricing of Market-wide Liquidity Dimensions

This table displays results for testing whether the non-default spread (NDS) impounds a residual non-default yield spread. In order to directly interpret the constant as a residual non-default yield spread, the liquidity variables have been monotonically transformed so that the constant will evaluate the remaining non-default spread when liquidity variables are taken at values corresponding with perfect liquidity. Each of the three proxies for the bond-specific liquidity dimensions – bid-ask spread, Amihud measure, and resiliency – have been orthoganalized to the other two. Following Elton et al. (2001), the coupon rate controls for state taxes within the non-default spread. Following Dick-Nielsen et al. (2012) and others, we control for the differential in liquidity pricing during crisis periods by splitting the sample into crisis (2008-09) and post-crisis (2010-12) subsamples. Models 1 and 4 cover the entire sample period, while Models 2 and 5 include only the financial crisis period and Models 3 and 6 include only the post-financial-crisis period. The sample for this unbalanced panel regression consists of bonds guaranteed by the FDIC under the DGP. Standard errors are clustered by bond and date. Two-tailed *p*-values are in parenthesis. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Variable	**NDS**	**NDS**	**NDS**	**NDS**	**NDS**	**NDS**
Constant	0.0097	0.0838*	0.0831***	-0.0064	0.0871	0.0768***
	(0.759)	(0.071)	(0.000)	(0.840)	(0.130)	(0.000)
ln(100×Bid-Ask Spread+1)	0.4482***	0.2242***	0.1509***	0.3549***	0.1684**	0.1434***
orthoganolized	(0.000)	(0.004)	(0.000)	(0.000)	(0.013)	(0.000)
ln(100×Amihud+1)	0.0017	-0.0041	0.0072***	0.0031	-0.0004	0.0071***
orthoganolized	(0.554)	(0.224)	(0.000)	(0.212)	(0.902)	(0.000)
ln(100×(1-Resiliency)+1)	0.0243***	0.0267***	0.0066***	0.0213***	0.0234***	0.0063***
orthoganolized	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ln(100×Market Spread+1)				62.112***	64.655***	17.842***
				(0.000)	(0.000)	(0.000)
ln(100×Market Amihud+1)				0.4852***	0.2474	0.1161***
				(0.002)	(0.271)	(0.001)
ln(100×(1-Market Resiliency)+1)				-0.0016	-0.0095	-0.0013
				(0.685)	(0.215)	(0.392)
Coupon	0.0462***	0.0925***	-0.0048	0.0314***	0.0706***	-0.0050
	(0.001)	(0.000)	(0.205)	(0.009)	(0.000)	(0.201)
Adj. R2	0.089	0.067	0.053	0.161	0.128	0.073
Bonds	65	64	64	65	6	64
Days	958	275	683	904	256	648
Obs.	10,144	4,108	6,036	9,601	3,882	5,719
Sample	Full	Crisis	Post-Crisis	Full	Crisis	Post-Crisis
Fixed Effects	None	None	None	None	None	None

Table 14: Analysis of the Residual Non-Default Yield Spread

This table displays results for regression specifications analyzing the residual non-default yield spread within the non-default spread (NDS). In order to directly interpret the constant as a residual non-default yield spread, the liquidity variables have been monotonically transformed so that the constant will evaluate the remaining non-default spread when liquidity variables are taken at values corresponding with perfect liquidity. The VIX level has been demeaned so that the intercept can be interpreted as the residual non-default yield spread when the VIX is taken at the mean of the regression sample. Each of the three proxies for the bond-specific liquidity dimensions – bid-ask spread, Amihud measure, and resiliency – have been orthoganalized to the other two. Following Elton et al. (2001), the coupon rate controls for state taxes within the non-default spread. Following Dick-Nielsen et al. (2012) and others, we control for the differential in liquidity pricing during crisis periods by splitting the sample into crisis (2008-09) and post-crisis (2010-12) subsamples. Model 1 covers the entire sample period, while Model 2 includes only the financial crisis period and Model 3 includes only the post-financial-crisis period. The sample for this unbalanced panel regression consists of bonds guaranteed by the FDIC under the DGP. Standard errors are clustered by bond and date. Two-tailed *p*-values are in parenthesis. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	**(2)**	**(3)**
Variable	**NDS**	**NDS**	**NDS**
Constant	0.1188***	0.2855***	0.0801***
	(0.000)	(0.000)	(0.000)
VIX	0.0175***	0.0222***	0.0034***
demeaned	(0.000)	(0.000)	(0.000)
ln(100×Bid-Ask Spread+1)	0.1473***	-0.0234	0.1368***
orthoganolized	(0.000)	(0.386)	(0.000)
ln(100×Amihud+1)	0.0049***	0.0017	0.0072***
orthoganolized	(0.002)	(0.500)	(0.000)
ln(100×(1-Resiliency)+1)	0.0090***	0.0018	0.0061***
orthoganolized	(0.000)	(0.553)	(0.000)
ln(100×Market Spread+1)	32.034***	23.927***	17.245***
	(0.000)	(0.000)	(0.000)
ln(100×Market Amihud+1)	0.1950**	0.1162	0.0896***
	(0.015)	(0.181)	(0.004)
ln(100×(1-Market Resiliency)+1)	-0.0016	-0.0017	-0.0014
	(0.381)	(0.537)	(0.303)
Coupon	0.0144**	0.0342**	-0.0055
	(0.046)	(0.031)	(0.157)
Adj. R2	0.621	0.773	0.120
Bonds	65	64	64
Days	904	256	648
Obs.	9,601	3,882	5,719
Sample	Full	Crisis	Post-Crisis
Fixed Effects	None	None	None

Table 15: Changes Specification Regressions

This table displays results for robustness tests intended to show that the pricing of liquidity dimensions remains statistically significant when controlling for the possible non-stationarity in the non-default spreads (NDS). Since the NDS is close to a non-stationary variable, we include the lagged level of the NDS. Following Dick-Nielsen et al. (2012) and others, we control for the differential in liquidity pricing during crisis periods by splitting the sample into crisis (2008-09) and post-crisis (2010-12) subsamples in Models 4 and 5, respectively. The sample for this unbalanced panel regression consists of bonds guaranteed by the FDIC under the DGP. Standard errors are clustered by bond and date. Two-tailed *p*-values are in parenthesis. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Variable	**ΔNDS**	**ΔNDS**	**ΔNDS**	**ΔNDS**	**ΔNDS**
Lagged NDS	-0.1845***	-0.1943***	-0.2074***	-0.0775***	-0.7726***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
ln(Bid-Ask Spread)	0.0162***	0.0179***	0.0166***	-0.0002	0.0167***
orthoganolized	(0.000)	(0.000)	(0.000)	(0.910)	(0.000)
ln(Amihud)	0.0014	0.0033***	0.0031***	-0.0008	0.0058***
orthoganolized	(0.143)	(0.004)	(0.004)	(0.366)	(0.000)
-ln(Resiliency)	0.0093***	0.0101***	0.0086***	-0.0007	0.0076***
orthoganolized	(0.000)	(0.000)	(0.000)	(0.697)	(0.000)
ln(Market Spread)			0.0013***	0.0006	0.0010***
			(0.000)	(0.221)	(0.001)
ln(Market Amihud)			0.0051***	-0.0014	0.0080***
			(0.005)	(0.539)	(0.000)
-ln(Market Resiliency)			0.0428**	0.0215	0.0413***
			(0.014)	(0.451)	(0.004)
Coupon	0.0103***	*Subsumed by Fixed Effects*			
	(0.007)				
Constant	0.0112	*Subsumed by Fixed Effects*			
	(0.273)				
Adj. R2	0.130	0.141	0.150	0.051	0.565
Bonds	64	64	64	63	64
Days	957	957	886	253	633
Obs.	10,110	10,110	9,408	3,833	5,575
Sample	Full	Full	Full	Crisis	Post-Crisis
Fixed Effects	None	Bond	Bond	Bond	Bond

Table 16: Vector Autoregressions

This table displays results for one-lag vector autoregression testing of the impact of the lagged three dimensions of liquidity (orthogonalized to each other) on the non-default spreads (NDS) and the dimensions themselves. The contemporaneous VIX level is included in the VAR to control for market volatility. Panel A contains the crisis subsample while panel B contains the post-crisis subsample. The sample for this unbalanced panel regression consists of bonds guaranteed by the FDIC under the DGP. Two-tailed *p*-values are in parenthesis. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Crisis Subsample

	(1)	(2)	(3)	(4)
Variable	**NDS_{id}**	**ln(Bid-Ask Spread)$_{id}$**	**ln(Amihud)$_{id}$**	**ln(Resiliency)$_{id}$**
$NDS_{i,d-1}$	0.8025***	-0.0197	0.5886*	0.2824**
	(0.000)	(0.903)	(0.076)	(0.038)
ln(Bid-Ask Spread)$_{i,d-1}$	0.0049***	0.0946***	0.2440***	-0.0066
orthogonalized	(0.003)	(0.000)	(0.000)	(0.726)
ln(Amihud)$_{i,d-1}$	0.0020***	0.0067	0.3733***	0.0104
orthogonalized	(0.008)	(0.522)	(0.000)	(0.238)
ln(Resiliency)$_{i,d-1}$	0.0053***	0.0147	0.2373***	0.0239
orthogonalized	(0.006)	(0.579)	(0.000)	(0.279)
Constant	-0.0582***	-0.1294	0.9394***	-0.0301
	(0.000)	(0.163)	(0.000)	(0.697)
VIX_{id}	0.0042***	0.0176***	-0.0394***	0.0018
	(0.000)	(0.000)	(0.000)	(0.622)
Adj. R^2	0.926	0.042	0.142	0.012
Bonds	53	53	53	53
Days	274	274	274	274
Obs.	2094	2094	2094	2094
Sample	Crisis	Crisis	Crisis	Crisis
Fixed Effects	None	None	None	None

(continued on next page)

Panel B: Post-Crisis Subsample

Variable	(1) NDS_{id}	(2) $ln(Bid\text{-}Ask\ Spread)_{id}$	(3) $ln(Amihud)_{id}$	(4) $ln(Resiliency)_{id}$
$NDS_{i,d-1}$	0.4963***	1.3543***	-0.5500	0.1757
	(0.000)	(0.000)	(0.313)	(0.343)
$ln(Bid\text{-}Ask\ Spread)_{i,d-1}$	0.0019	0.1010***	0.1799***	0.0190
orthogonalized	(0.105)	(0.000)	(0.000)	(0.132)
$ln(Amihud)_{i,d-1}$	0.0008	0.0098	0.2559***	0.0226***
orthogonalized	(0.187)	(0.416)	(0.000)	(0.001)
$ln(Resiliency)_{i,d-1}$	0.0005	0.0135	0.2200***	0.0644***
orthogonalized	(0.796)	(0.708)	(0.000)	(0.002)
Constant	0.0166***	-0.6024***	0.9338***	-0.0260
	(0.001)	(0.000)	(0.000)	(0.638)
VIX_{id}	0.0025***	0.0053	-0.0080	0.0006
	(0.000)	(0.198)	(0.255)	(0.808)
Adj. R^2	0.357	0.022	0.061	0.006
Bonds	50	50	50	50
Days	570	570	570	570
Obs.	2461	2461	2461	2461
Sample	Post-Crisis	Post-Crisis	Post-Crisis	Post-Crisis
Fixed Effects	None	None	None	None

Table 17: Vector Autoregressions

Panel A presents the pre-shock descriptive statistics for all of the untreated firms in the Canadian sample. Data for board and network variables are collected from the Clarkson Centre. Other firm-level data are collected from Worldscope. Panel B presents descriptive statistics for the treated firms. A firm is considered treated if one of its current directors had a network connection, past or present, with the deceased director. Firms in which the deceased director currently sitting on the board are excluded from both the untreated and treated samples.

Panel A: Untreated

	Obs.	Mean	St. Dev.	Min	Median	Max
Board Size	3,537	7.941	3.809	1	7	25
Degree Centrality	3,537	15.429	10.200	1	12.875	47.583
Eigenvector Centrality	3,537	0.007	0.011	0	0.001	0.056
Closeness Centrality	3,468	0.165	0.090	0	0.197	0.276
Betweenness Centrality	3,537	7,201	7,980	0	4,679	38,003
Ln(Assets)	2,363	13.410	2.359	7.533	13.494	20.357
CapEx/Assets	2,354	0.087	0.098	0	0.055	0.475
Cash/Assets	2,359	0.136	0.183	0	0.061	0.823
Leverage	1,985	0.265	0.203	0.001	0.234	0.963
Tobin's Q	2,299	1.696	1.606	0.109	1.204	10.229
ROA	2,361	-0.033	0.237	-1.557	0.024	0.269

Panel B: Treated

	Obs.	Mean	St. Dev.	Min	Median	Max
Board Size	247	11.798	3.956	3	11	23
Degree Centrality	247	30.160	10.428	7.333	30.538	47.583
Eigenvector Centrality	247	0.021	0.015	0	0.019	0.056
Closeness Centrality	247	0.238	0.025	0.166	0.238	0.276
Betweenness Centrality	247	15,242	8,970	1,361	13,827	38,003
Ln(Assets)	181	15.608	2.065	10.396	15.560	20.357
CapEx/Assets	181	0.061	0.064	0	0.043	0.402
Cash/Assets	181	0.079	0.130	0	0.028	0.823
Leverage	167	0.255	0.176	0.001	0.240	0.816
Tobin's Q	179	1.102	0.771	0.109	1.045	5.018
ROA	181	0.025	0.120	-0.937	0.031	0.252

Table 18: Matched Sample Pre-shock Firm Characteristics

This table contains comparisons of the means and medians of the treated and matched-control pre-shock samples. Firms are dropped from the sample if they are missing data on assets, industry (SIC code), or abnormal returns. Each treated firm is then matched with three untreated firms, with replacement, based on industry, pre-shock firm size, and pre-shock degree centrality. Panel A contains all treated and control firms. Panel B contains only treated firms with no current director interlock and their matched control firms. In the Treated Mean column, the means of the treated and control samples are tested using a pooled sample t-test. In the Treated Median column, the distributions of the treated and control samples are tested using a non-parametric rank-sum test. In the Paired Difference column, the control firm is subtracted from the paired treated firm and the means are tested against zero using a standard t-test. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: All treated and matched-control firms

	Matched Control			Treated			Paired	
Variable	N	Mean	Median	N	Mean	Median	N	Difference
Board Size	477	13.195	13.000	159	13.189	13.000	477	-0.006
Degree	477	29.140	29.500	159	31.915***	32.333***	477	2.775***
Eigenvector	477	0.021	0.020	159	0.023*	0.020**	477	0.003***
Closeness	477	0.232	0.236	159	0.239**	0.239**	477	0.007***
Betweenness	477	12,313	10,845	159	14,818***	13,171***	477	2,505***
Ln(Assets)	477	15.887	15.656	159	15.930	15.830	477	0.043
CapEx/Assets	477	0.066	0.044	159	0.061	0.041	477	-0.005
Cash/Assets	477	0.064	0.023	159	0.068	0.026	477	0.004
Leverage	460	0.265	0.243	149	0.254	0.240	435	-0.012
Tobin's Q	477	1.095	0.946	158	1.046	0.994	474	-0.046
ROA	477	0.033	0.029	159	0.042	0.032	477	0.008**

(Continued on next page)

Panel B: Firms with no current interlock

	Matched Control			Treated			Paired	
Variable	**N**	**Mean**	**Median**	**N**	**Mean**	**Median**	**N**	**Difference**
Board Size	201	12.746	12.000	67	12.642	13.000	201	-0.104
Degree	201	29.463	30.333	67	32.399**	32.333**	201	2.936***
Eigenvector	201	0.020	0.021	67	0.024*	0.020	201	0.003***
Closeness	201	0.242	0.248	67	0.247	0.246	201	0.005***
Betweenness	201	12,574	11,453	67	15,646***	13,171**	201	3,072***
Ln(Assets)	201	16.004	15.872	67	15.947	16.056	201	-0.057
CapEx/Assets	201	0.061	0.041	67	0.056	0.032	201	-0.005
Cash/Assets	201	0.075	0.024	67	0.078	0.026	201	0.003
Leverage	194	0.234	0.215	63	0.250	0.240	184	0.008
Tobin's Q	201	1.173	1.022	67	1.053	0.973	201	-0.121*
ROA	201	0.042	0.033	67	0.048	0.033	201	0.006

Table 19: Matched Sample Post-Shock Changes

This table contains comparisons of the means and medians of the treated and matched-control abnormal returns and changes in network centrality. Firms are dropped from the sample if they are missing data on assets, industry (SIC code), or abnormal returns. Each treated firm is then matched with three untreated firms, with replacement, based on industry, pre-shock firm size, and pre-shock degree centrality. Panel A contains all treated and control firms. Panel B contains only treated firms with no current director interlock and their matched control firms. In the Treated Mean column, the means of the treated and control samples are tested using a pooled sample t-test. In the Treated Median column, the distributions of the treated and control samples are tested using a non-parametric rank-sum test. In the Paired Difference column, the control firm is subtracted from the paired treated firm and the means are tested against zero using a standard t-test. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: All treated and matched-control firms

	Matched Control			Treated			Paired	
Variable	N	Mean	Median	N	Mean	Median	N	Difference
AR (0)	477	0.23%	0.12%	159	-0.36%***	-0.20%***	477	-0.59%***
CAR (-2, +2)	477	0.39%	0.44%	159	0.35%	0.35%	477	-0.04%
CAR (-1, +1)	477	0.40%	0.18%	159	-0.05%	0.09%	477	-0.45%*
CAR (0, +1)	477	0.19%	0.07%	159	-0.26%	-0.18%*	477	-0.44%**
%ΔDegree	477	-0.03%	0.00%	159	-0.29%***	-0.19%***	477	-0.26%***
%ΔEigenvector	477	5.38%	-0.29%	159	5.37%	-1.13%	477	-0.01%
%ΔCloseness	477	-0.04%	-0.01%	159	-0.04%	-0.03%***	477	0.00%
%ΔBetweenness	477	0.03%	-0.01%	159	-0.01%	-0.06%***	477	-0.05%

(continued on next page)

Panel B: Firms with no current interlock

	Matched Control			Treated			Paired	
Variable	N	Mean	Median	N	Mean	Median	N	Difference
AR (0)	201	0.14%	0.11%	67	-0.55%***	-0.08%**	201	-0.69%***
CAR (-2, +2)	201	0.06%	0.21%	67	-0.84%**	0.26%	201	-0.90%***
CAR (-1, +1)	201	0.22%	0.14%	67	-0.85%***	-0.29%**	201	-1.07%***
CAR (0, +1)	201	0.18%	0.21%	67	-0.91%***	-0.38%***	201	-1.09%***
%ΔDegree	201	-0.06%	0.00%	67	-0.16%**	0.00%***	201	-0.10%***
%ΔCloseness	201	-0.06%	-0.01%	67	-0.05%	-0.02%	201	0.02%
%ΔEigenvector	201	7.47%	0.99%	67	12.69%	0.86%	201	5.22%**
%ΔBetweenness	201	0.14%	0.08%	67	0.01%	0.08%	201	-0.13%

Table 20: Multivariate Analysis of Shock – Full Matched Sample

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to the director network. The pooled sample in Panel A contains all of the treated firms and their matched-control firms as separate observations. The dependent variables are percentage changes in the four network centrality measures following the shock to the director network and the cumulative abnormal returns in the event-windows surrounding the directors' deaths. The dependent variables are regressed on a dummy variable equaling 1 for treated firms, as well as control variables and industry fixed effects (using 1 digit SIC codes). The intercept term is subsumed by the fixed-effects. Panel B contains a sample of paired differences of treated and control firms. The control firm is subtracted from the treated firm for the dependent and control variables. No fixed effects are used in Panel B, allowing the intercept to be inferred as the difference-in-differences coefficient. All regressions use unbalanced panel, seemingly unrelated regression methodology, allowing the residuals on the seven event-days (director deaths) to be correlated. The variance-covariance matrix is adjusted using panel-corrected standard errors. Two-tailed p-values are in parenthesis below their corresponding coefficient. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Pooled Specifications

Dependent variable	%Δ Degree	%Δ Eigenvector	%Δ Closeness	%Δ Betweenness	AR (0)	CAR (-2,+2)	CAR (-1,+1)	CAR (0,+1)
Treated Dummy	**-0.282*****	**-0.367**	**-0.016*****	**-0.151*****	**-0.297****	**-0.103**	**-0.108**	**-0.254**
	(0.000)	**(0.329)**	**(0.000)**	**(0.000)**	**(0.032)**	**(0.738)**	**(0.664)**	**(0.215)**
Ln(Assets)	0.011**	0.466***	-0.001	0.015**	-0.096*	-0.366***	-0.294***	-0.197**
	(0.041)	(0.009)	(0.254)	(0.011)	(0.079)	(0.002)	(0.003)	(0.015)
Tobin's Q	0.016*	0.170	0.000	0.008	0.042	-0.230	-0.223	-0.135
	(0.079)	(0.518)	(0.939)	(0.418)	(0.666)	(0.289)	(0.196)	(0.336)
ROA	-0.113	0.174	-0.025**	-0.046	-1.917	-6.971***	-2.458	-2.151
	(0.274)	(0.945)	(0.013)	(0.654)	(0.114)	(0.008)	(0.243)	(0.210)
Board Size	0.008***	-0.038	0.001***	0.001	-0.004	-0.024	0.020	0.019
	(0.000)	(0.529)	(0.000)	(0.651)	(0.853)	(0.604)	(0.602)	(0.551)
R^2	0.432	0.022	0.092	0.092	0.043	0.051	0.032	0.038
Obs.	635	635	635	635	635	635	635	635
Fixed Effects	Industry	Industry	Industry	Industry	Industry	Industry	Industry	Industry

(continued on next page)

Panel B: Paired-Difference Specifications

Dependent variable	%Δ Degree *difference*	%Δ Eigenvector *difference*	%Δ Closeness *difference*	%Δ Betweenness *difference*	AR (0) *difference*	CAR (-2,+2) *difference*	CAR (-1,+1) *difference*	CAR (0,+1) *difference*
Intercept	**-0.259*** **	**-0.913*** **	**-0.015** **	**-0.149*** **	**-0.232** **	**-0.148**	**-0.168**	**-0.311** **
	(0.000)	**(0.000)**	**(0.000)**	**(0.000)**	**(0.012)**	**(0.492)**	**(0.332)**	**(0.041)**
Ln(Assets)	0.008	0.408	0.000	0.010	-0.098	-0.032	-0.136	-0.338*
difference	(0.623)	(0.159)	(0.656)	(0.375)	(0.355)	(0.898)	(0.499)	(0.059)
Tobin's Q	0.014	0.108	0.000	0.000	0.066	-0.114	-0.746***	-0.609***
difference	(0.357)	(0.664)	(0.545)	(0.991)	(0.581)	(0.687)	(0.001)	(0.002)
ROA	-0.222	-0.039	-0.011	-0.097	-1.525	-9.835***	0.652	-0.895
difference	(0.172)	(0.987)	(0.117)	(0.251)	(0.297)	(0.003)	(0.803)	(0.697)
Board Size	0.013***	0.006	0.000***	0.003*	-0.030	-0.001	0.041	0.029
difference	(0.000)	(0.899)	(0.001)	(0.074)	(0.210)	(0.983)	(0.371)	(0.472)
R^2	0.052	0.004	0.013	0.008	0.009	0.026	0.027	0.033
Obs.	474	474	474	474	474	474	474	474
Fixed Effects	None	None	None	None	None	None	None	None

Table 21: Multivariate Analysis of Shock – Non-Interlocked Subsample

This table displays results for the multivariate difference-on-differences analysis on the effect of a shock to the director network. The pooled sample in Panel A contains the treated firms with no current interlock to the deceased director and their matched-control firms as separate observations. The dependent variables are the cumulative abnormal returns in the event-windows surrounding the directors' deaths. The dependent variables are regressed on a dummy variable equaling 1 for treated firms, as well as control variables and industry fixed effects (using 1 digit SIC codes). The intercept term is subsumed by the fixed-effects. Panel B contains a sample of paired differences of treated and control firms. The control firm is subtracted from the treated firm for the dependent and control variables. No fixed effects are used in Panel B, allowing the intercept to be inferred as the difference-in-differences coefficient. All regressions use unbalanced panel, seemingly unrelated regression methodology, allowing the residuals on the seven event-days (director deaths) to be correlated. The variance-covariance matrix is adjusted using panel-corrected standard errors. Two-tailed p-values are in parenthesis below their corresponding coefficient. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Pooled Specifications

Dependent variable	AR(0)	CAR(-2,+2)	CAR(-1,+1)	CAR(0,+1)
Treated Dummy	**-0.409****	**-0.647**	**-0.631***	**-0.694****
	(0.034)	**(0.103)**	**(0.057)**	**(0.013)**
Ln(Assets)	-0.083	-0.294**	-0.261**	-0.129
	(0.263)	(0.050)	(0.036)	(0.236)
Tobin's Q	-0.093	-0.211	-0.331	-0.232
	(0.481)	(0.430)	(0.132)	(0.213)
ROA	1.476	3.471	3.458	0.693
	(0.361)	(0.275)	(0.173)	(0.753)
Board Size	0.010	-0.017	0.028	0.018
	(0.749)	(0.789)	(0.601)	(0.685)
R^2	0.059	0.105	0.088	0.062
Obs.	268	268	268	268
Fixed Effects	Industry	Industry	Industry	Industry

(continued on next page)

Panel B: Paired-Difference Specifications

Dependent variable	AR(0) *difference*	CAR(-2,+2) *difference*	CAR(-1,+1) *difference*	CAR(0,+1) *difference*
Intercept	**-0.366****	**-0.614****	**-0.637*****	**-0.681*****
	(0.011)	**(0.039)**	**(0.010)**	**(0.001)**
Ln(Assets)	-0.371**	-0.111	-0.375	-0.591**
difference	(0.017)	(0.728)	(0.146)	(0.014)
Tobin's Q	-0.074	-0.347	-0.983***	-0.799***
difference	(0.661)	(0.318)	(0.001)	(0.001)
ROA	2.103	5.416	9.187***	1.477
difference	(0.354)	(0.218)	(0.010)	(0.631)
Board Size	-0.014	0.013	0.028	0.014
difference	(0.732)	(0.883)	(0.694)	(0.812)
R^2	0.037	0.010	0.075	0.069
Obs.	201	201	201	201
Fixed Effects	None	None	None	None

Table 22: Triple Difference Analysis of Shock – Interlocked Subsample

This table displays results for the multivariate difference-in-difference-on-differences analysis on the effect of a shock to the director network for the deaths of directors on big committees vis-á-vis directors on small committees. Panel A contains a sample of paired differences of the treated firms with a current interlock to a firm of the deceased director and their matched-control firms as separate observations. The dependent variables are the cumulative abnormal returns in the event-windows surrounding the directors' deaths. The dependent variables are regressed on a dummy variable equaling 1 for observations stemming from the death of a director who had an average committee size greater than 7 (the median), as well as control variables and industry fixed effects (using 1 digit SIC codes). Panel B contains the same sample and specification, but with the percentage changes in network centrality measures as the dependent variables. All regressions use unbalanced panel, seemingly unrelated regression methodology, allowing the residuals on the seven event-days (director deaths) to be correlated. The intercept term is subsumed by the fixed-effects in both panels. The variance-covariance matrix is adjusted using panel-corrected standard errors. Two-tailed p-values are in parenthesis below their corresponding coefficient. ***, **, and * represent statistical significance at the 1%, 5% and 10% levels, respectively.

Panel A: Abnormal Returns

Dependent variable	AR(0) *difference*	CAR(-2,+2) *difference*	CAR(-1,+1) *difference*	CAR(0,+1) *difference*
Big Committee Dummy	**-0.035**	**-1.594****	**-0.438**	**-0.310**
	(0.911)	**(0.013)**	**(0.434)**	**(0.533)**
Ln(Assets)	-0.036	-1.000**	-0.339	-0.430
difference	(0.808)	(0.011)	(0.268)	(0.112)
Tobin's Q	0.108	-1.010**	-1.146***	-0.919***
difference	(0.549)	(0.034)	(0.001)	(0.003)
ROA	-3.356*	-21.028***	-6.161*	-1.943
difference	(0.077)	(0.000)	(0.070)	(0.532)
Board Size	-0.065**	-0.066	-0.064	-0.060
difference	(0.022)	(0.358)	(0.264)	(0.229)
R^2	0.163	0.290	0.243	0.222
Obs.	273	273	273	273
Fixed Effects	Industry	Industry	Industry	Industry

(continued on next page)

Panel B: Network Centrality Measures

Dependent variable	%Δ Degree *difference*	%Δ Eigenvector *difference*	%Δ Closeness *difference*	%Δ Betweenness *difference*
Big Committee Dummy	**-0.209*****	**2.324*****	**-0.002***	**-0.152*****
	(0.000)	**(0.001)**	**(0.076)**	**(0.006)**
Ln(Assets)	0.008	0.321	-0.001	0.007
difference	(0.592)	(0.490)	(0.531)	(0.690)
Tobin's Q	-0.001	0.723*	0.000	-0.023
difference	(0.931)	(0.098)	(0.935)	(0.123)
ROA	-0.238	8.513**	-0.012*	-0.203
difference	(0.140)	(0.028)	(0.091)	(0.172)
Board Size	0.010***	0.082	0.000	0.003
difference	(0.001)	(0.305)	(0.206)	(0.337)
R^2	0.204	0.134	0.071	0.044
Obs.	273	273	273	273
Fixed Effects	Industry	Industry	Industry	Industry

Figure 1: Residual Non-Default Yield Spread by Issuer Credit Rating

This figure displays the average residual non-default yield spread for bonds grouped by issuer credit ratings at the time of observation. The residual non-default yield spreads are calculated by including rating fixed-effects in the regressions modelled in Table 14, thus these residual non-default yield spreads assume the mean level of the VIX in each respective regression: 26.22 (Full Sample), 31.84 (Crisis subsample), and 22.40 (Post-crisis subsample).

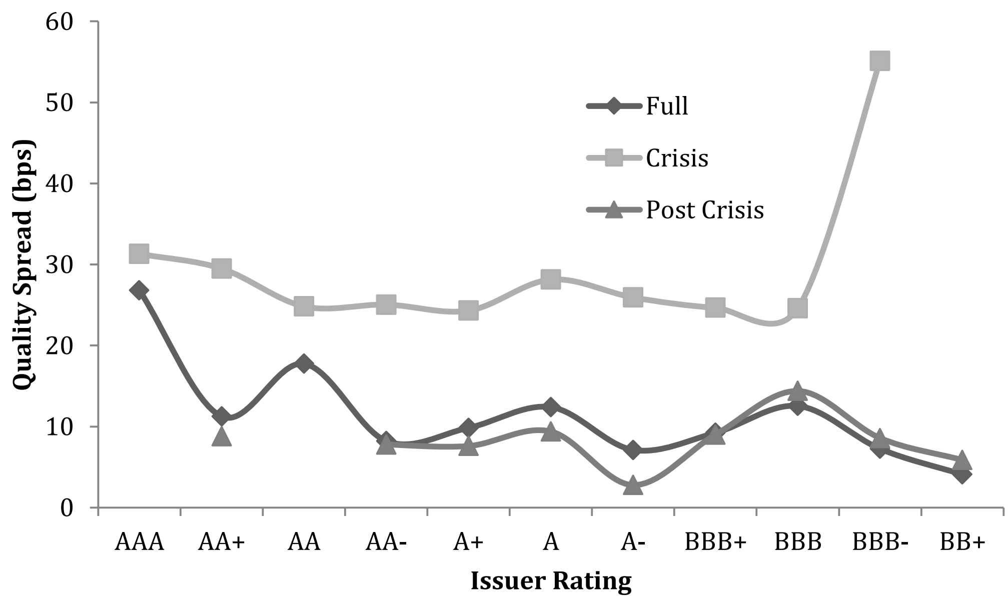

Figure 2: Impulse Response Functions

This figure displays the response of the non-default spread (NDS), ln(Bid-Ask Spread), ln(Amihud measure), and ln(Resiliency) to a one standard deviation shock to each of the variables while controlling for contemporaneous market volatility using the VIX level. The three liquidity variables are orthogonalized to each other. Panel A contains responses during the financial crisis (2008-2009) while Panel B contains responses during the post-crisis period (2010-2012).

(Continued on next page)

Panel A: Crisis Subsample

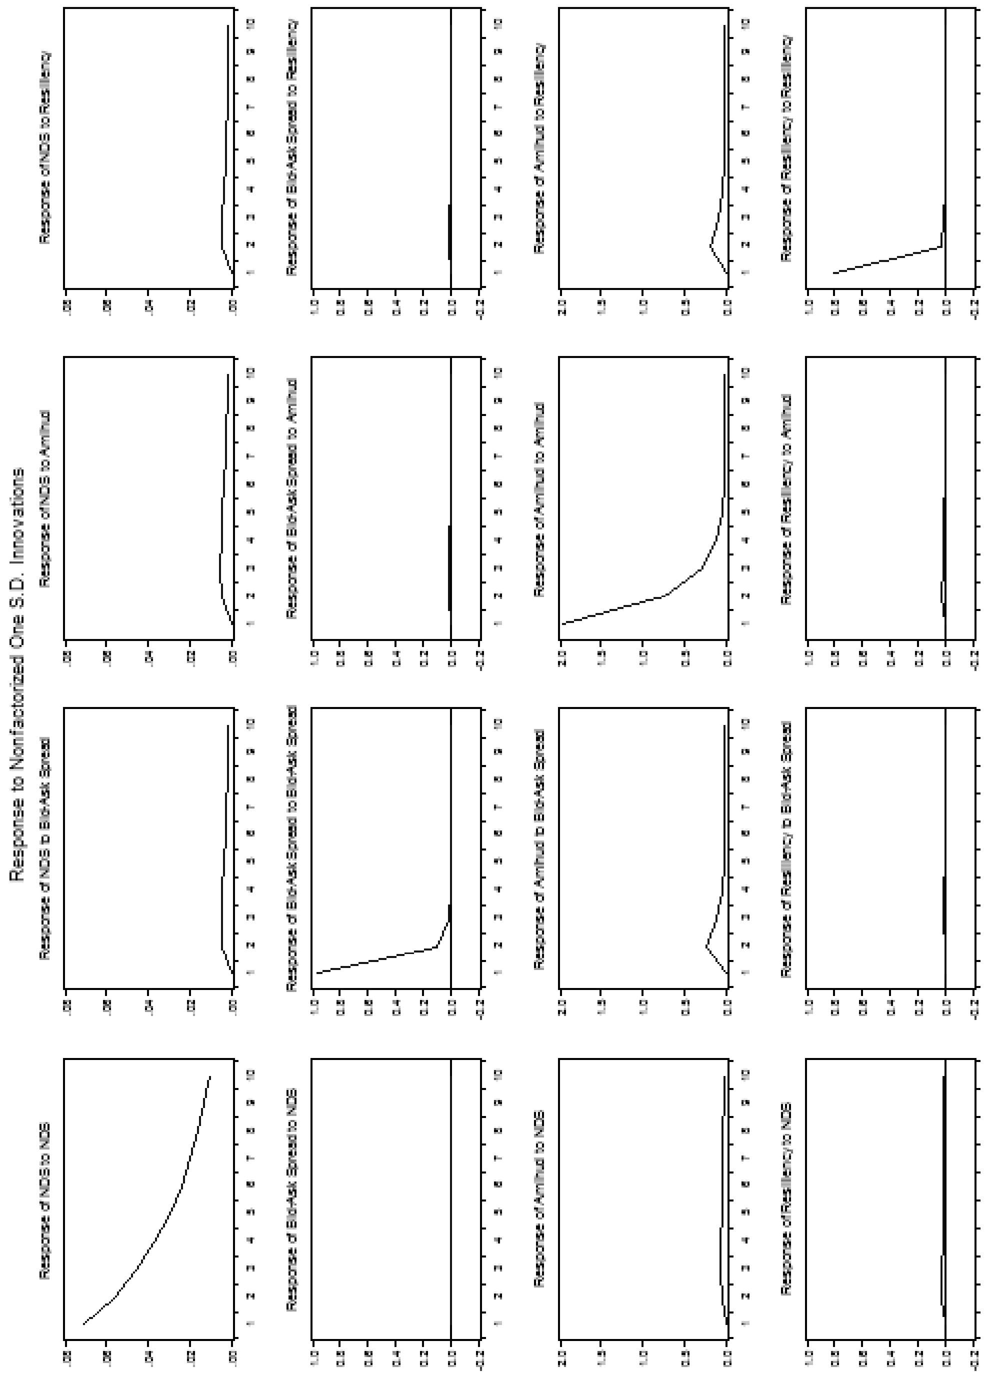

(Continued on next page)

Panel B: Post-Crisis Subsample