

Essays on the Liquidity of Financial Markets

Jackson R. Pfeiffer

April 19, 2021

Dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Industrial Administration (Finance) at the Tepper School of Business at Carnegie Mellon University

Doctoral Committee:

Burton Hollifield (Chair)

Thomas G. Ruchti

Duane J. Seppi

Chester S. Spatt

Chad J. Zutter

Abstract

Liquidity is a central feature in the research of financial markets. Broadly speaking, liquidity describes the ease with which an asset can be traded both quickly and without deviating much from the current market price. In this collection of essays, I investigate three aspects of liquidity – one affecting stocks in the context of a limit order book market, and two in the context of the over-the-counter market governing U.S. corporate bonds. In the first essay, measuring liquidity as the depth of the limit order book at price levels extending from the mid-quote price, I explore the implications of private information which can be drawn from the dynamic levels of liquidity in the order book. In the second essay, measuring illiquidity as the effective bid-ask spread, I create an estimate for the aggregate illiquidity of the broader over-the-counter market for U.S. corporate bonds from the observed illiquidity of bonds that are traded. In the third essay, measuring liquidity as the market’s ability to absorb the large demand to sell bonds resulting from a fire sale, I observe the unintended impacts that Dodd-Frank regulations have on bond market liquidity and the indirect effects they have on the cost of capital faced by corporations when raising debt in the bond market.

Essay 1: Informed Liquidity in Limit Order Book Markets

In this essay, I propose a novel estimator for the presence of asymmetric information among market makers in limit order book markets. Model parameters are structurally estimated using a continuously-updated generalized method of moments regression on historical Nasdaq TotalView-ITCH high-frequency data, which includes every order sent to the exchange. I find statistically significant evidence that market makers are able to anticipate the presence of informed traders in the near future and allocate their provision of liquidity accordingly.

Keywords: market making, asymmetric information, order flow, microstructure

JEL codes: C13, C58, D40, G14

Essay 2: Estimating Aggregate Illiquidity of the Corporate Bond Market

In this essay, I create an estimation procedure for measuring the liquidity available in the U.S. corporate bond market as a whole. The procedure is split into two primary components. First, I use a hidden Markov model framework to estimate the ex-ante probability that each bond issue will trade on a particular day. Next, I estimate the illiquidity of those bonds which trade throughout the day in question, and use a locally weighted regression to estimate what the liquidity of bonds that did not trade would have been had they traded. Both stages of this estimation demonstrate strong external predictive validity and outperform alternative estimation methodologies. I compare my estimate of overall market illiquidity to other macroeconomic factors with known ties to bond market liquidity, extending the results of existing literature performed at the individual bond level to the overall bond market.

Keywords: hidden Markov model, local regression, machine learning

JEL Codes: C11, C13, C14, C38, C53, G14

Essay 3: Liquidity Risk from Dealer Inventory Limits

In this essay, I seek to better understand how the Volcker Rule has affected the provision of liquidity in the market for U.S. corporate bonds by conducting various natural experiments on the observed trading patterns surrounding stress events. I use credit rating downgrades to test the differential impact of liquidity shocks before and after the Volcker Rule went into effect. I find very little evidence that the regulations have deteriorated the market's ability to absorb liquidity shocks, contradicting existing literature. I do, however, find strong evidence that the Volcker rule has increased institutional investors' aversion to holding BBB- rated bonds. Moreover, I show that after the implementation of the Volcker Rule, investment grade bonds on the cusp of junk ratings are priced as if they are already junk bonds due to the increased aversion to holding such bonds. In this way, the cost of capital for the affected firms increased substantially – ostensibly as a result of the Volcker Rule.

Keywords: Dodd-Frank, Volcker Rule, dealer inventories

JEL codes: G14, G23, G24, G28

Contents

1	Informed Liquidity in Limit Order Book Markets	6
1.1	Introduction	6
1.1.1	Contributions to Current Literature	8
1.2	Model of Optimal Liquidity Provision	10
1.2.1	Preliminaries	10
1.2.2	“Agnostic” Market Makers	12
1.2.3	Informed Market Makers	14
1.3	Data	19
1.4	Empirical Model Estimation	21
1.4.1	Distribution of Market Orders	21
1.4.2	Estimation of Model Parameters	23
1.4.3	Optimal Order Size	26
1.4.4	Comparison of Observed Orders to Model-Implied Optimal Orders	27
1.5	Estimator of Relative Liquidity Provisioning	30
1.6	Results	33
1.7	Discussion	39
1.8	Limitations	42
1.9	Future Research	44
1.10	Conclusion	44
2	Estimating Aggregate Illiquidity of the Corporate Bond Market	46
2.1	Introduction	46
2.2	Data	49
2.3	Trade-Based Estimate of Bond Liquidity	51
2.4	Estimating the Probability of Trading	52
2.4.1	Hidden Markov Model Tuning Procedure	53
2.4.2	Performance of Hidden Markov Model Estimator	56

2.5	Matching Estimator	59
2.5.1	LOWESS Regression Tuning Procedure	60
2.5.2	Performance of LOWESS Regression Matching Estimator	62
2.6	Measure of Overall Market Liquidity	66
2.6.1	Relationships between Market Liquidity and Macroeconomic Factors	66
2.7	Limitations	68
2.8	Future Research	70
2.9	Conclusion	70
3	Liquidity Risk from Dealer Inventory Limits	71
3.1	Introduction	71
3.2	Related Literature	72
3.3	Data	73
3.4	Test Design and Results	75
3.5	Discussion	84
3.6	Limitations	87
3.7	Future Research	87
3.8	Conclusion	88
A	References	89
B	Tables	94
C	Figures	121

1 Informed Liquidity in Limit Order Book Markets

1.1 Introduction

I propose a model in which a market maker adds an order to a pre-existing limit order book. A market maker solves for her optimal supply of liquidity by maximizing an expected profit function. Using this model, I can predict the optimal size of an order placed by a market maker who has no private information. I can compare the observed actions of market makers to those which one would expect if they had no private information. The differences can be used to identify when market makers are acting on private information. Moreover, the direction of the divergence - less or more liquidity provision than is optimal - provides insights about the private information market makers possess.

Actions taken by an informed market maker are visible to other market participants. Those participants therefore receive a signal from the informed market maker, and revise their expectations accordingly. In this way, there exist feedback loops, both directly (mechanically altering the empirical price impact function), and indirectly (by shifting the other market participants' expectations about future order flow). Thus, the private information is filtered through the market to the other participants. It is possible to glean a better understanding of this process by investigating the type of information market makers have, how they act on that information, and how the rest of the market responds to those actions.

Central to the notion of market completeness is the ease with which a market participant can place a trade. This, in turn, is in large part a result of the amount of liquidity supplied to the market. I find that market makers can predict the impending presence of informed traders, and adjust their liquidity provisioning to compensate for the presence of such traders. Market completeness, or maximal market completeness, is fundamentally important to the optimal functioning of financial markets. The degradation of completeness reduces the ability of market participants to allocate their resources so as to insure consumption across various states of nature, thereby decreasing the social welfare provided by the markets.

In practice, global capital markets can never be fully complete, spanning every source of

risk. There is, however, some maximal level of completeness that is possible to achieve if all the assets benefit from perfectly efficient trading. As trading frictions of any one asset increase, the ease of trading that asset necessarily diminishes, making the market as a whole less complete. The presence of informed traders in financial markets introduces a significant amount of additional risk to market makers.

In the canonical microstructure framework, there are two types of agents placing market orders: informed traders and uninformed traders. Informed traders have knowledge about a stock. Observing the schedule of liquidity provisioning provided by market makers and taking into consideration their knowledge, these agents place informed trades. By acting upon their information, these traders are signaling what they know to the liquidity providers. Informed traders will place larger orders when they expect to see greater gains based on their information, which in turn sends a larger signal to the market. The presence of this signal motivates the concept of a price impact function, wherein the size and direction of an order affect the price of a stock moving forward, as market makers respond to the signal. Uninformed traders have some exogenous liquidity demand, unrelated to the value of the asset. In this way, the actions of the uninformed traders add noise to the signal of the informed traders, and they are therefore often referred to as noise traders.

Market makers, in turn, attempt to provide liquidity in such a way that their expected losses from trading against an informed trader are compensated by their expected gains from collecting a spread when trading against noise traders. It is this intuition that allows me to establish the optimal spread for a market maker to provide on an asset and to obtain the optimal shape of the order book. The greater the activity of noise traders at any time, the greater the volume of orders being placed in the market, and the greater the volatility. In periods such as this, it is possible for informed traders to hide more of their signal in the noise. It has been shown [Back, 1992] that it is therefore optimal for informed traders to act more aggressively on their knowledge, placing larger orders at these times. Because of the increased activity of informed traders, periods of high volume and volatility are costly to a market maker. If a market maker knows that such a period is beginning to occur, she

would want to prepare herself by reducing her liquidity provisioning until once again, in expectation, her losses from informed traders are balanced by her gains from noise traders.

In response to the presence of informed traders, market makers supply less liquidity to the market than they otherwise would. These actions decrease the short term price efficiency, making it more expensive to execute trades. Insofar as social welfare is derived from the proper functioning of financial markets, market makers' responses to the perceived presence of informed traders therefore carry a social cost borne from the resulting increase in trading frictions. However, by exhibiting such behavior, market makers are able to provide more liquidity to the market when they see that the presence of informed traders is diminished. In this way, while financial markets experience periods of poor functioning as a result of informed market makers, the net welfare effect depends upon the duration of those periods.

If the activity of informed traders is limited to short, intermittent bouts of increased aggressiveness, then the resulting social costs would be limited to similarly short periods. Moreover, for the majority of the time, the provision of liquidity would be enhanced (due to the market makers' knowledge of the presence of informed traders), resulting in better functioning financial markets.

1.1.1 Contributions to Current Literature

My paper adds to the current literature in three ways. First, existing models ignore the existing order book when considering a market maker's optimal behavior. A market maker places orders which maximize her expected profit. The profit function depends on the existing supply of liquidity in the market. Thus, it is important to consider the depth of the existing limit order book when solving for the optimal liquidity provision. Without doing so, one cannot accurately model the behavior of market makers in a limit order book market.

Second, my paper is the first to consider informed market makers outside the context of a dealer market. The bulk of market microstructure theory regarding information asymmetries considers a market in which new information filters into the market purely through the actions of informed traders placing market orders. Such models, however, are most applicable

to a broker-dealer market in which the market makers do not possess any informational advantages beyond the trivial knowledge of their own individual actions, preferences, and constraints. Financial markets operated in large part as dealer markets when the seminal works of Kyle (1985) and Glosten and Milgrom (1985) were introduced. In modern markets characterized by electronic limit order books, however, those traders who possess private information often act as market makers [e.g. Bloomfield, O'Hara, and Saar (2005)], placing the limit orders which make up the spread and provide liquidity to the market. As such, it would be beneficial to take seriously a model which considers the possibility that the placement, revision, or cancellation of a limit order is done by an informed market maker.

While there exists a growing set of research which investigates the dynamics of equilibria for models in which market makers possess private information [e.g., Gould and Vericchia (1985), Biais (1993), Kumar and Seppi (1994), and Calcagno and Lovo (2006)], the literature builds off of the Kyle (1985) framework in which the trading process takes place as a first-price auction. This paper extends the Glosten and Milgrom framework, proposing a model in which market makers place quotes in a limit order market. Moreover, my model allows for inferences to be made about several aspects of private information. Optimal liquidity provisioning is dependent on the future value of the security, the future market order flow, and the existing supply of liquidity. Divergences from a market maker's optimal behavior imply private knowledge about one or more of these channels of information. In this way, my model allows for my measure of asymmetric information to be more flexible in the type of private information market makers can possess than the current literature.

Finally, the measure of asymmetric information I propose is the first I know of to directly identify asymmetric information through the divergences between the observed action of a market maker and the action which would be optimal if said market maker had no private information. Current empirical measures of asymmetric information measure it in one of two ways. Some attempt to identify it directly, through the actions of traders placing market orders [Easley et al. (2002) and Chang, Chang, and Wang (2014)], while others measure asymmetric information indirectly, through observing signs of the adverse selection

which would result from the presence of information asymmetries [Roll (1984) and Yueshen (2016)]. By identifying asymmetric information through the actions of market makers, I am able to investigate both the ways in which market makers act upon their private information as well as the type of private information available to market makers. These aspects of asymmetric information inform both the theoretical market microstructure models and the empirical tests for the presence of asymmetric information. In this way, by developing a better understanding of the nature of market makers' information asymmetries, I am facilitating future work in this area of study.

This essay is organized as follows. Section 1.2 describes the model environment; the setup of the model, the optimal behavior of uninformed market makers, and the behavior of informed market makers. Section 1.3 introduces the data used in testing this estimator, and Section 1.4 describes the empirical methods used in fitting the model to the data. Section 1.5 proposes estimators to measure either the relative provision of liquidity supplied to the market or the amount of price pressure driven by the informed market-makers' actions. The empirical content contained in these estimators is investigated in Section 1.6 and discussed in Section 1.7. Finally, Section 1.8 reviews the limitations of this study, Section 1.9 proposes avenues for future research, and Section 1.10 concludes.

1.2 Model of Optimal Liquidity Provision

1.2.1 Preliminaries

My model builds off of the Glosten (1994) model. I then impose discrete prices and a time priority rule as in Seppi (1997). Consider an environment in which market makers can place limit orders along a finite price schedule. Let $p_1 < p_2 < \dots < p_n$ be the set of prices at which there exist positive quantities of limit sell orders, with corresponding quantities $\{Q_1, Q_2, \dots, Q_n\}$. Similarly, let $p_{-1} > p_{-2} > \dots > p_{-m}$ be the set of prices at which there exist positive quantities of limit buy orders, with corresponding quantities $\{Q_{-1}, Q_{-2}, \dots, Q_{-m}\}$. In this way, p_{-1} represents the price of the best bid, and p_1 represents

the price of the best offer. When a market order is submitted, limit orders are filled first in order of price, favoring price levels beneficial to the market order, then in order of the time at which the market's orders of the same price level were submitted, filling older limit orders first. At any time before a limit order is filled, the agent who submitted the order may either reduce the size of the order (without losing her previously established time preference), or cancel it entirely. For every limit order placed in the market, the agent placing said order faces some fixed costs, γ_0 , and for every order that is successfully executed, the agent faces some variable costs, γ_1 , that depend on the quantity traded. γ_0 can be interpreted as the on-going cost associated with monitoring the order. In this way, when an agent already has an open limit order in the market, she must weigh the expected revenue generated by that limit order against the fixed on-going cost to continue monitoring the order.

The fundamental value of the stock evolves according to the stochastic process

$$dX_t = \mu \cdot dt + \sigma_t \cdot dW_t \quad (1)$$

where μ is the instantaneous drift of the stock price, σ_t^2 is the time t variance of the value, and W_t is a Brownian motion under the physical measure.

Let m_t denote market order quantity at time t where $m < 0$ corresponds to a sell order. In the presence of a market order $m_{t'}$, the expected future value of the stock is given by

$$E[X_{t'} | X_t, m_{t'}] = X_t + \mu(t' - t) + h(m_{t'}) \quad (2)$$

where $h(m) = \alpha m$ is a price impact function. Note the simplifying assumption that the price impact function is linear. Also assume that the distribution of market order quantities is given by the probability mass function $f(m)$ and corresponding cumulative mass function $F(m)$. Further assume that the time between market orders follows an exponential distribution given by the probability density function

$$k(\Delta T_t) = \frac{1}{\psi} e^{-\frac{1}{\psi} \Delta T_t} \quad (3)$$

so that ψ is the expected time between market orders. Note that, because of the memorylessness of the hazard function, at any given time, ψ is the expected time until the next market order arrival¹.

1.2.2 “Agnostic” Market Makers

Assume that the trading structure is set up as a repeated single stage game. First, the fundamental value of the asset is announced, and a continuum of liquidity providers facing perfect competition are exogenously assigned either an existing limit order to monitor or a price level at which they may enter a new limit order. If assigned an existing limit order, liquidity providers are then allowed to either modify or cancel said order. Otherwise, they decide if they want to place an order at their assigned price level, and if so, its size. After a random period of time, a new market order is submitted, a new fundamental value is announced, and liquidity providers are randomly assigned either one of the limit orders which remains after the market order is filled or a price level. Thus, the game repeats itself.

The expected marginal profit on the last unit of a sell limit order at price level p of size q is given by

$$E[(p - \gamma_1 - E[X_{t'}|X_t, m_{t'}])\mathbb{1}_{m_{t'} \geq \bar{Q} + q}] \quad (4)$$

where $\gamma_0 + \gamma_1 q$ is the total order monitoring and processing costs² for an order of size q , and \bar{Q} is the sum of the sizes of all those sell limit orders with either price preference or time preference to the order in question. (4) can be combined with (2) and the distribution of m

¹Easley et. al. (2008) show that the order arrival intensity of market orders is highly persistent. For this reason, it may be beneficial to allow for the order intensity to evolve according to an auto-regressive process in future work.

²In this way, γ_0 is the fixed cost of placing an order and γ_1 is the variable cost. These costs implicitly include the cost of monitoring the additional order, the additional risk the order places on the market maker’s portfolio, and the cost that results from tightening the market maker’s inventory constraints.

to be rewritten as

$$\begin{aligned}
& E[(p - \gamma_1 - E[X_{t'}|X_t, m_{t'}])\mathbb{1}_{m_{t'} \geq \bar{Q} + q}] \\
&= E[(p - \gamma_1 - E[X_t + \mu(t' - t) + \alpha m_{t'}])\mathbb{1}_{m_{t'} \geq \bar{Q} + q}] \\
&= E[(p - \gamma_1 - X_t - \mu E[t' - t] - \alpha E[m_{t'}])\mathbb{1}_{m_{t'} \geq \bar{Q} + q}] \\
&= (p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'}|m_{t'} \geq \bar{Q} + q])\mathbb{P}(m_{t'} \geq \bar{Q} + q)
\end{aligned} \tag{5}$$

where ψ is the expected time until the next market order is placed. Setting (5) equal to zero, I have

$$0 = p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'}|m_{t'} \geq \bar{Q} + q] \tag{6}$$

which I can use to find the optimal order size \hat{q} at price level p for a risk-neutral, profit-maximizing market maker. The market maker will place the order so long as the following condition is satisfied:

$$\sum_{n=1}^{\hat{q}} (p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'}|m_{t'} \geq \bar{Q} + n]) \geq \gamma_0. \tag{7}$$

As a market maker monitors an existing order, if said order does not execute quickly, then she re-evaluates the position. Monitoring an open limit order is costly – both in the attention that it requires and in the opportunity cost of a different limit order that may have a higher probability of executing. Therefore, when monitoring an existing limit order, the market maker will cancel her order if the following condition is not satisfied:

$$\mathbb{P}(m_{t'} \geq \bar{Q}) \sum_{n=1}^{\hat{q}} (p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'}|m_{t'} \geq \bar{Q} + n]) \geq \gamma_0, \tag{8}$$

where $\mathbb{P}(m_{t'} \geq \bar{Q})$ is an added penalty such that the market maker is more likely to cancel an order that is less likely to be filled. With a cancellation correlating to a size of 0 shares, the optimal order size is restricted to $\hat{q} \in [0, \infty]$.

Similarly, on the other side of the book, I have the conditions

$$0 = X_t - \gamma_1 - p + \mu\psi + \alpha E[m_{t'} | m_{t'} \leq -\bar{Q} - \hat{q}], \quad (9)$$

$$\sum_{n=1}^{\hat{q}} (X_t - \gamma_1 - p + \mu\psi + \alpha E[m_{t'} | m_{t'} \leq -\bar{Q} - n]) \geq \gamma_0, \quad (10)$$

and

$$\mathbb{P}(m_{t'} \leq -\bar{Q}) \sum_{n=1}^{\hat{q}} (X_t - \gamma_1 - p + \mu\psi + \alpha E[m_{t'} | m_{t'} \leq -\bar{Q} - n]) \geq \gamma_0. \quad (11)$$

1.2.3 Informed Market Makers

Here, I examine three classes of informed market makers: those that have private information about the fundamental value of the asset, those that have information about the presence of informed traders, and those that have information about the future market order flow but no additional information about the presence of informed traders in that order flow. Market makers with private information about the fundamental value of the asset are shown to behave in a similar fashion as informed traders; they act upon their information to maximize expected profits, and in so doing, release that information to the market. The second and third class of informed market makers do not have information concerning the value of the asset. Rather, the second class have information about the presence (or lack thereof) of informed traders, and the third only has information about the level of future market order flow. The latter two types of informed market maker act upon their knowledge in such a way that the efficient provision of liquidity is maximized; by knowing about the presence of informed traders – or the conditions during which informed traders are most active – they are able to avoid providing liquidity to the market when the probability of trading against informed traders is high and provide more liquidity to the market when the probability is low.

1.2.3.1 Information Regarding the Fundamental Value of the Asset Recall from Equation (4) that the expected marginal profit on the last unit of a sell limit order at price

p and size q is given by

$$E[(p - \gamma_1 - E[X_{t'}|X_t, m_{t'}])\mathbf{1}_{m_{t'} \geq \bar{Q} + q}],$$

where \bar{Q} is the total depth of the limit order book ahead of the order. Suppose that a market maker has private information about the fundamental value of the asset, \tilde{X} . Thus, $E[X_{t'}|X_t, m_{t'}] = \tilde{X}, \forall m_{t'}$, so the expected marginal profit on the last unit of a sell limit order becomes

$$E[(p - \gamma_1 - \tilde{X})\mathbf{1}_{m_{t'} \geq \bar{Q} + q}] = (p - \gamma_1 - \tilde{X})\mathbf{P}(m_{t'} \geq \bar{Q} + q).$$

Without loss of generality, assume that $\tilde{X} < p - \gamma_1$. In such a situation, the informed market maker would earn guaranteed profits for every unit of her limit order that is filled, and would therefore be willing to provide an arbitrarily large amount of liquidity at that price level. Moreover, the market maker would be willing to provide arbitrarily large amounts of liquidity to the sell side of the order book for any price p such that $p > \tilde{X} + \gamma_1$. Note that the market maker would also be willing to provide arbitrarily large amounts of liquidity to the buy side of the order book for any price p such that $p < \tilde{X} - \gamma_1$.

Thus, given the current limit order book B_t , where

$$B_t \equiv (\{Q_{-m}, \dots, Q_{-2}, Q_{-1}, Q_1, Q_2, \dots, Q_n\}, \{p_{-m} < \dots < p_{-2} < p_{-1} < p_1 < p_2 < \dots < p_n\}),$$

the informed market maker's expected marginal profits at the next market order $m_{t'}$'s arrival for a sell limit order are

$$\pi_A(p_a, q|B_t) = (p_a - \gamma_1 - \tilde{X})\mathbf{P}\left(m_{t'} \geq q + \sum_{i=1}^a Q_i\right),$$

and her expected profits for a buy limit order are

$$\pi_B(p_{-b}, q|B_t) = (\tilde{X} - \gamma_1 - p_{-b})\mathbb{P}\left(m_{t'} \leq -q - \sum_{i=1}^b Q_{-i}\right).$$

The informed market maker's optimal strategy is similar to that of the informed trader in the Kyle (1985) model; such a strategy maximizes expected profits (and in so doing greatly increases trade volume) while simultaneously minimizing the information revealed to the other market participants through her actions. While the informed trader attempts to hide her actions in the market order flow from noise traders, the informed market maker attempts to hide her actions in the limit order flow from information agnostic market makers. An important difference in the optimal strategy of an informed market maker and that of an informed trader is that the market maker's strategy must additionally be optimized over the conditional expectation of future market order flow. Whereas the informed trader can directly control the volume she submits to the market, the informed market maker can only directly control the maximum volumes that she will absorb; the market maker only has indirect control over the expected volume she will absorb.

Solving for the informed market maker's optimal strategy is outside the scope of this essay. That being said, inferences can be drawn about how the presence of such market makers and their actions would present themselves in the limit order book. In order to maximize expected profits, these market makers must maximize the amount of liquidity that they absorb at profitable price levels. In this way, they will provide more liquidity to the order book than would otherwise be optimal, but only at certain price levels. In the event that the fundamental value of the asset, which is known to the informed market maker, is below the best-bid price³, then the market maker will add increased levels of liquidity to the ask side of the order book. Conversely, if the fundamental value of the asset is above the

³Prima facie analysis might suggest that an actor with knowledge that the fundamental value of the asset is below the best-bid price will sell the asset and fill the bid in order to capture guaranteed profits. It is conceivable, however, that the informed actor will instead attempt to absorb market buy order flow in an effort to capture greater profits in expectation by maintaining their informational advantage for a longer time period.

best-ask price, then the informed market maker will provide greater liquidity to the bid side of the book.

1.2.3.2 Information Regarding the Presence of Informed Traders Market makers and informed traders are in a perpetual cat and mouse game. Transacting with a trader who is acting on private information leads to losses for the market maker. The premiums that market makers charge must therefore compensate in expectation for the losses stemming from informed traders. Moreover, when market makers believe that they are more likely to be trading against an informed trader, they will charge greater premiums. As market makers charge greater premiums on liquidity, however, they disincentivize trading activity from uninformed noise traders. Thus, profit maximizing market makers attempt to discriminate informed market order flow from the market order flow generated by noise traders. In response, informed traders attempt to hide their order flow in the order flow from noise traders, executing their informed trades in periods of increased trading activity.

If a market maker has no information about the presence of informed traders, then her response to periods of increased market order flow would be to provide less liquidity than normal. Since informed traders hide their trades in periods of increased order flow, the risk of transacting with an informed trader is greater during these periods. Thus, the uninformed market maker must charge higher premiums in periods when market order flow is elevated, forgoing the profits that she would otherwise earn during these periods were she not at risk of trading against an informed trader. Conversely, if a market maker has information about the presence (or lack thereof) of informed traders, then she is better able to discriminate between the market order flow from noise traders and the order flow from informed traders and can better respond to periods of increased market order flow. For example, if a market maker has information that no informed traders are present during a period of increased order flow, then her profit maximizing action will be to provide more liquidity than would otherwise be optimal.

In the context of my model, the price impact, $\tilde{\alpha}$, of an informed trader's order is greater

than the average price impact, α , of a market order, which is itself greater than the price impact, $\tilde{\alpha}$, of a noise trader's order. Since my model is designed as a repeated single stage game, changes in the aggressiveness of market order flow can be represented by changes to the distribution of market order quantities, $f(m)$. A period of increased market order flow can be represented by the distribution $f_H(m)$, and a period of decreased market order flow, by $f_L(m)$, such that $f_L(m)$ is second-order stochastically dominant over $f(m)$, which is itself second-order stochastically dominant over $f_H(m)$.

Recall from Equation (6) that the optimal size \hat{q} of a limit sell order satisfies

$$0 = p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'} | m_{t'} \geq \bar{Q} + \hat{q}] = p - \gamma_1 - X_t - \mu\psi - \alpha \sum_{m=\bar{Q}+\hat{q}}^{\infty} f(m) \cdot m$$

Holding all else constant, in order to satisfy the optimality condition, it must be the case that

$$\frac{1}{\alpha} \propto \sum_{m=\bar{Q}+\hat{q}}^{\infty} f(m) \cdot m.$$

Let $\hat{q}_A(\alpha, f(\cdot) | p, \gamma_1, X_t, \mu, \psi, \bar{Q})$ be a function that maps the tuple $(\alpha, f(\cdot))$ to the size of a limit sell order that satisfies the aforementioned optimality condition, given the other parameters. From the stochastic dominance of $f_L(m)$, $f(m)$, and $f_H(m)$, it is known that for a given q ,

$$\sum_{m=\bar{Q}+q}^{\infty} f_H(m) \cdot m > \sum_{m=\bar{Q}+q}^{\infty} f(m) \cdot m > \sum_{m=\bar{Q}+q}^{\infty} f_L(m) \cdot m.$$

It is also assumed that $\tilde{\alpha} > \alpha > \tilde{\alpha}$. Therefore, I have the following relationships:

$$\begin{aligned} \hat{q}_A(\alpha, f_H(\cdot)) &< \hat{q}_A(\alpha, f(\cdot)) < \hat{q}_A(\alpha, f_L(\cdot)) \\ \hat{q}_A(\tilde{\alpha}, f(\cdot)) &< \hat{q}_A(\alpha, f(\cdot)) < \hat{q}_A(\tilde{\alpha}, f(\cdot)). \end{aligned}$$

Following the same logic, these relationships also hold for the optimal size of a limit buy order. Note that, if a market maker only has information about the presence of informed traders, her response will be symmetric across both sides of the order book. Thus, if a

market maker has information that an informed trader is present in the market, then she will provide less liquidity to both sides of the order book than she otherwise would have given the current market order flow. Alternatively, if a market maker has information that informed traders are not present, then she will provide more liquidity to both sides than she otherwise would.

1.2.3.3 Information Regarding Future Market Order Flow As discussed in Section 1.2.3.2, informed traders are most active during periods of increased market order flow when it is easier for them to hide their informed order flow among the order flow of noise traders. If a market maker has no information about the presence of informed traders, then her response to periods of increased market order flow would be to provide less liquidity than normal. If a market maker without information about the presence of informed traders has information about future market order flow, then her optimal reaction to this information would be to proactively adjust her provision of liquidity ahead of the changing order flow. Thus, if a market maker has information that the market order flow will increase in the near future, increasing the risk of informed traders being present, then she will provide less liquidity than she otherwise would. Alternatively, if a market maker has information that the market order flow will decrease in the near future, then she will provide more liquidity than she otherwise would. Note that, similar to market makers with information about the presence of informed traders, a market maker with information about future market order flow will act symmetrically on both sides of the order book.

1.3 Data

This essay uses Nasdaq's Historical TotalView-ITCH data, which includes messages indicating all orders and trade transactions, time-stamped to the nanosecond. From this, I am able to reconstruct Nasdaq's Level 1 TotalView data stream, which Nasdaq describes as "...the standard Nasdaq data feed for serious traders - [displaying] the full order book depth for Nas-

daq market participants.”⁴ Thus, I am able to match every order to the order book present at the time of the order’s arrival. Moreover, the data contain unique order ID numbers allowing me to track the precise depth of limit orders with either price preference or time preference, \overline{Q} , for orders which are either revised or canceled. Market orders that execute against multiple limit orders are reported as separate orders – one market order for each limit order against which they execute – with identical trade times (recorded in nanoseconds since midnight). I am therefore able to identify these orders by matching market orders in the same direction that execute at the same time, and can reconstruct the full market orders by aggregating their constituent pieces.

I use the order book data for ten highly liquid stocks, [American Airlines (ticker AAL), Apple (AAPL), Bank of America (BAC), Caterpillar (CAT), Cisco Systems (CSCO), Facebook (FB), General Electric (GE), Google (GOOG), Goldman Sachs (GS), and Microsoft (MSFT)], for the 20 trading days in January of 2015. I split this sample into an in-sample training data set consisting of the 15 trading days between January 2, 2015 and January 23, 2015 and an out-of-sample testing data set containing January 26, 2015 through January 30, 2015.

The distributions of market order sizes over the in-sample and out-of-sample periods are presented in Figure 1, and the distributions of limit order sizes over the same periods are presented in Figure 2. The distributions of limit order placements/revisions across price levels away from the bid-ask spread, for both the in-sample and out of sample periods, are shown in Figure 3. Similarly, the distributions of the location in the limit order book at which orders are canceled, for both the in-sample period and the out-of-sample period, are shown in Figure 4. The most activity with respect to limit orders occurs in the five price levels on either side of the order book that are nearest to the bid-ask spread. This is indicative of market makers, who by their nature only act near the spread, limiting the bulk of their activity to price levels within a few ticks of the spread. Therefore, I restrict my attention to limit orders which are placed (or canceled) within five active price levels of the current

⁴“Nasdaq TotalView-ITCH,” <http://www.nasdaqtrader.com/Trader.aspx?id=Totalview2>, (April 12, 2016)

spread.

The behavior observed in a subset of limit order placements and cancellations is consistent with the anecdotal trend of “flickering orders,” in which liquidity providers try to learn about other market participants’ strategies by placing a limit order only to cancel it almost immediately. Placing and canceling limit orders in this fashion neither affects market liquidity in any meaningful way, nor does it resemble the behavior exhibited by limit orders that are submitted with the intent to provide liquidity to the market. I therefore exclude these flickering orders from my sample, identifying them as those limit orders that are canceled within a few seconds of being placed and before a new market order arrives.

A summary of the market activity observed for each of the stocks during the in-sample and out-of-sample periods, after restricting the sample to the five ticks nearest the spread and excluding flickering orders, is presented in Table 1. For each stock and for both time periods, the table reports the number of market orders observed, the number of limit order placements/revisions observed, the number of limit order cancellations observed, the total number of order messages included in the study, as well as the number of flickering orders that were observed and subsequently excluded.

1.4 Empirical Model Estimation

1.4.1 Distribution of Market Orders

In order to match the empirical distribution of market orders, I first allow for the possibility that the probability of a market order being a sell or a buy order is not equal to 0.5. It is immediately apparent from the histogram of market orders (Figure 1) that a significant mass of the distribution exists over orders of sizes which are less than a full lot (between 1 and 99 shares). It is also readily apparent that, stylistically, given $|m| \geq 100$, the order sizes are even lots almost surely. Thus, I impose the following distributional assumption on the size of market orders.

$$f(m) = \begin{cases} (1 - \xi)\lambda_1(1 - \lambda_2)^{m/100}\lambda_2 & \text{if } m \in 100n, n \in \mathbb{N} \\ (1 - \xi)(1 - \lambda_1) \cdot \frac{1}{99} & \text{if } m \in \{1, \dots, 99\} \\ \xi(1 - \phi_1) \cdot \frac{1}{99} & \text{if } m \in \{-1, \dots, -99\} \\ \xi\phi_1(1 - \phi_2)^{-m/100}\phi_2 & \text{if } m \in -100n, n \in \mathbb{N} \end{cases} \quad (12)$$

Note that this distributional assumption implicitly restricts the domain of m to $\pm\{1, \dots, 99\} \cup \pm 100n, n \in \mathbb{N}$. The above pmf implies that the cmf is given by

$$F(m) = \begin{cases} \xi + (1 - \xi)(1 - \lambda_1) + (1 - \xi)\lambda_1(1 - (1 - \lambda_2)^{m/100}) & \text{if } m \geq 100 \\ \xi + (1 - \xi)(1 - \lambda_1) \cdot \frac{1}{99} \cdot m & \text{if } m \in \{1, \dots, 99\} \\ \xi\phi_1 + \xi(1 - \phi_1) \cdot \frac{1}{99}(m + 100) & \text{if } m \in \{-1, \dots, -99\} \\ \xi\phi_1(1 - \phi_2)^{-m/100} & \text{if } m \leq -100 \end{cases} \quad (13)$$

While the parameterization of market orders smaller than a full lot as uniformly distributed clearly doesn't match the empirical distribution precisely (with the empirical distribution showing a spike at orders of 50 shares), the cumulative mass over this region plays a much larger part in determining the optimal size of a limit order than the shape of the distribution over the region. This is because the vast majority of orders are located in the order book in such a way that the sum of the size of limit orders with either price or time preference totals at least 100 shares. Caring only about orders large enough to execute against their own limit order, market makers thus do not care about the exact distribution of such small market orders. Therefore, this essay need not concern itself with the precise distribution over said region, and is safe in its assumption of a uniform distribution.

Under the above parameterization, by the properties of conditional expectations of uniform distributions and geometric distributions, for $A > 0$,

$$E[m|m \geq A] = \begin{cases} \frac{(100-A)}{99}(1 - \lambda_1)\frac{99+A}{2} + (1 - \frac{(100-A)}{99})(1 - \lambda_1)\frac{100}{\lambda_2} & \text{if } A < 100 \\ A + \frac{100}{\lambda_2} & \text{if } A \geq 100. \end{cases} \quad (14)$$

Similarly, for $A < 0$,

$$E[m|m \leq A] = \begin{cases} \frac{(100+A)}{99}(1-\phi_1)\frac{A-99}{2} - (1 - \frac{(100+A)}{99})(1-\phi_1)\frac{100}{\phi_2} & \text{if } A > -100 \\ A - \frac{100}{\phi_2} & \text{if } A \leq -100. \end{cases} \quad (15)$$

1.4.2 Estimation of Model Parameters

This essay's structural model contains ten parameters which must be estimated,

$$\theta = \{\mu, \alpha, \xi, \lambda_1, \lambda_2, \phi_1, \phi_2, \psi, \gamma_0, \gamma_1\}.$$

Given the assumption of the characterization of the stochastic process of the stock's underlying value X_t , I have the process

$$\Delta X_t = \mu \Delta T_t + \alpha m_t + \varepsilon_t,$$

which I can use to arrive at the moment condition

$$E[e_1(\theta; y_t)] = E[\Delta X_t - \mu \Delta T_t - \alpha m_t | \text{executed against visible liquidity}] = 0, \quad (16)$$

where ΔX_t denotes the change in mid-quote price caused by the execution of the market order m_t . When a market order executes against hidden liquidity, the remainder of the hidden limit order becomes visible, resulting in a tightening of the bid-ask spread. Since I measure the stock price as the mid-quote price, this implies that $\alpha < 0$. Clearly the price impact function should be monotonically increasing, so I must remove these observations from my moment matching condition.

Since $\xi = \mathbb{P}(m_t < 0)$, I can create the moment condition

$$E[e_2(\theta; y_t)] = E[\mathbb{1}_{m_t < 0} - \xi] = 0 \quad (17)$$

Since I assume that the distribution of market orders follows a geometric distribution for

orders of even lots, I have the moment conditions

$$E[e_3(\theta; y_t)] = E\left[m_t - \frac{100}{\lambda_2} \mid m_t \geq 100\right] = 0 \quad (18)$$

$$E[e_4(\theta; y_t)] = E\left[-m_t + \frac{100}{\phi_2} \mid m_t \leq -100\right] = 0 \quad (19)$$

Since I assume that the distribution of market orders smaller than a full lot follow uniform distributions, I need only to match the distribution of mass between orders of even lots and orders smaller than a full lot. I can do this by matching the mean of the distribution to the sample average, resulting in the moment conditions

$$E[e_5(\theta; y_t)] = E\left[m_t - 50(1 - \lambda_1) - \frac{100\lambda_1}{\lambda_2} \mid m_t > 0\right] = 0 \quad (20)$$

$$E[e_6(\theta; y_t)] = E\left[-m_t + 50(1 - \phi_1) + \frac{100\phi_1}{\phi_2} \mid m_t < 0\right] = 0 \quad (21)$$

In order to find the average time until a new market order is placed, ψ , I can simply fit a hazard function to the time intervals between market orders such that the pdf is $\frac{1}{\psi}e^{-\frac{1}{\psi}\Delta T_t}$, where ΔT_t is the time elapsed since the previous market order when m_t is entered, arising in the moment condition

$$E[e_7(\theta; y_t)] = E[\Delta T_t - \psi] = 0 \quad (22)$$

I am now able to derive moment conditions for γ_1 by using the break-even conditions (6) and (9).

$$E[e_8(\theta; y_t)] = E\left[p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'} \mid m_{t'} \geq \bar{Q} + q] \mid \text{limit sell order}\right] = 0 \quad (23)$$

$$E[e_9(\theta; y_t)] = E\left[X_t - \gamma_1 - p + \mu\psi + \alpha E[m_{t'} \mid m_{t'} \leq -\bar{Q} - \hat{q}] \mid \text{limit buy order}\right] = 0 \quad (24)$$

The inequality constraints governing whether a limit order should be placed (Equations (7) and (10)) and whether an existing limit order should be canceled (Equations (8) and

(11)) lead to a set of weak inequality moment conditions for γ_0 .

$$E[e_{10}(\theta; y_t)] = E \left[\sum_{n=1}^{\hat{q}} \left(p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'} | m_{t'} \geq \bar{Q} + n] \right) - \gamma_0 \mid \text{limit sell} \right] \geq 0 \quad (25)$$

$$E[e_{11}(\theta; y_t)] = E \left[\sum_{n=1}^{\hat{q}} \left(X_t - \gamma_1 - p + \mu\psi + \alpha E[m_{t'} | m_{t'} \leq -\bar{Q} - n] \right) - \gamma_0 \mid \text{limit buy} \right] \geq 0 \quad (26)$$

$$E[e_{12}(\theta; y_t)] = E \left[\gamma_0 - \mathbf{P} \left(m_{t'} \geq \bar{Q} \right) \sum_{n=1}^{\hat{q}} \left(p - \gamma_1 - X_t - \mu\psi - \alpha E[m_{t'} | m_{t'} \geq \bar{Q} + n] \right) \mid \text{cancel sell} \right] \geq 0 \quad (27)$$

$$E[e_{13}(\theta; y_t)] = E \left[\gamma_0 - \mathbf{P} \left(m_{t'} \leq -\bar{Q} \right) \sum_{n=1}^{\hat{q}} \left(X_t - \gamma_1 - p + \mu\psi + \alpha E[m_{t'} | m_{t'} \leq -\bar{Q} - n] \right) \mid \text{cancel buy} \right] \geq 0 \quad (28)$$

Inequality constraints pose a problem for conventional moment-based estimation techniques. However, Moon and Schorfheide (2008) show that models for which inequality moment conditions provide overidentifying information can be modified in such a way that they can be included in a generalized method of moment estimator and that the inclusion of these conditions can reduce the mean-squared estimation error. I rewrite the moment conditions e_{10} through e_{13} such that $E[e_{10}(\theta; y_t) - \zeta] = 0$, $E[e_{11}(\theta; y_t) - \zeta] = 0$, $E[e_{12}(\theta; y_t) - \eta] = 0$, and $E[e_{13}(\theta; y_t) - \eta] = 0$ for $\zeta, \eta \geq 0$. In this way, the slackness parameters, ζ and η , can be interpreted as the average opportunity costs of placing a new limit order and maintaining an existing limit order, respectively.

Using these thirteen moment conditions, setting e_1 through e_7 equal to zero for all observations of market orders, setting e_8 through e_{11} equal to zero for all limit order placements, and setting e_{12} and e_{13} equal to zero for all limit order cancellations, I am able to estimate the twelve parameters $\theta = \{\mu, \alpha, \xi, \lambda_1, \lambda_2, \phi_1, \phi_2, \psi, \gamma_0, \gamma_1, \zeta, \eta\}$. I perform a Continuous-Updating Generalized Method of Moments (CUGMM) regression,

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmin}} \left[\frac{1}{T} \sum g(\theta; y_t) \right]^\top \widehat{W}(\hat{\theta}) \left[\frac{1}{T} \sum g(\theta; y_t) \right] \quad (29)$$

where $g'(\theta; y_t) = [e_1(\theta; y_t), \dots, e_{14}(\theta; y_t)]$ and $\widehat{W}(\hat{\theta})$ is a continuously updated weighting

matrix such that

$$\widehat{W}(\widehat{\theta}) = \left(\frac{1}{T} \sum g(\widehat{\theta}; y_t) g(\widehat{\theta}; y_t)^\top \right)^{-1} \quad (30)$$

is estimated simultaneously with $\widehat{\theta}$. Hansen, Heaton, and Yaron (1996) find, in a comparison of Monte-Carlo simulations for various GMM estimation procedures, that CUGMM tends to outperform two-step feasible GMM and iterated GMM procedures.

The results of my estimation for each stock are reported in Table 2. Interestingly, for six of the ten stocks in my sample, the fixed cost γ_0 is negative. This suggests that, on average, the on-going monitoring costs associated with these stocks are surpassed by the on-going benefit of holding such an order. Such a result could be indicative of traders acting through the limit order book, placing orders that reflect their heterogeneous beliefs about the valuation of the asset. The result could also be indicative of market makers acting upon a non-linear price impact function with respect to market orders.

1.4.3 Optimal Order Size

Recall from Equations (6) and (9) that the optimal order size, \widehat{q} , is that which sets $E[m_{t'} | m_{t'} \geq \overline{Q} + \widehat{q}]$ or $E[m_{t'} | m_{t'} \leq -\overline{Q} - \widehat{q}]$ equal to $\frac{p - \gamma_1 - X_t - \mu\psi}{\alpha}$ or $\frac{p + \gamma_1 - X_t - \mu\psi}{\alpha}$, respectively. Therefore, given the current price of the equity, X_t , the model parameters $\{\gamma_1, \mu, \psi, \alpha\}$, and the amount of liquidity in the book with either price or time priority, \overline{Q} , the optimal order size for price level p is wholly dependent upon the conditional distribution of market orders. With the distributional assumptions I impose on the parameterization of market orders, if $\overline{Q} \geq 100$, there exists a closed form solution to the optimal size of a limit order. Specifically, if $\overline{Q} \geq 100$, the optimal order size for limit sell orders is given by

$$\widehat{q}^A = \frac{p - \gamma_1 - X_t - \mu\psi}{\alpha} - \overline{Q} - \frac{100}{\lambda_2}, \quad (31)$$

and the optimal order size for limit buy orders is given by

$$\widehat{q}^B = \frac{100}{\phi_2} - \overline{Q} - \frac{p + \gamma_1 - X_t - \mu\psi}{\alpha}. \quad (32)$$

Given the parameterization of this model, however, for $\bar{Q} < 100$, the above closed form solution is not guaranteed to hold. The solution has the possibility of becoming quadratic, with no guarantee of real roots. Thus, the estimation procedure requires a search function for the zero of the derivative of the expected profit function.

1.4.4 Comparison of Observed Orders to Model-Implied Optimal Orders

I am able to evaluate the ability for my model to capture the behavior of model makers by comparing the sizes of observed limit orders to the optimal order size projected by my model. Reported in Figure 5 are the in-sample and out-of-sample distributions across price levels of the count of limit orders for the ten stocks which are optimally provisioned, relatively thin, and relatively thick according to my model. Over both the in-sample and out-of-sample periods, more than a third of the orders (38% and 37%, respectively) are equal in size to that which is predicted by the model. Over the in-sample period, 23% of the orders provided more liquidity than is predicted by the model, and 39% of the orders provide less liquidity. Over the out-of-sample period, 24% of the orders were larger than predicted, and 40% were smaller than predicted. Not only does this demonstrate that my model is able to explain the provision of liquidity well, the out-of-sample performance shows that my model has strong external predictive validity.

This result provides evidence that contradicts Sandås's (2001) conclusions about market makers' behavior. Sandås (2001) rejects a similar model, concluding that the aggregate depth of the empirical order book was too thin relative to his model of the aggregate order book. The effectiveness of this model to capture the behavior of market makers placing individual orders within the context of the rest of the order book, however, suggests that Sandås's findings are in part driven by the latency with which agents are able to observe the market. Sandås's trading mechanism allows liquidity providers to continue to add or modify limit orders until they no longer wish to make any changes, and only then does a market order arrive. While Sandås finds that the aggregate liquidity provided by the limit order book is sub-optimal, his trading mechanism assumes that all market makers are able

to view and respond to changes in the market before the next market order arrives. My model relaxes this assumption and in so doing, makes it possible to examine the behavior of individual market makers. When market makers do respond to changes in the market, I find that the majority of their actions provide as much if not more liquidity than would be optimal according to my model.

The optimal order size for any given price level can be used to calculate the maximal expected profit available for limit orders placed at said price levels. Other strategic order placement strategies notwithstanding, one would expect market makers to place their orders at the price level which offers them the highest possible expected profit. With this in mind, I can compare the price levels at which I observe orders to the predicted price levels in order to investigate the strategic considerations of market makers with respect to future limit order flow. Placing a limit order at a price level away from the spread leaves open the possibility for another market maker to place an order which undercuts the order in question. Moreover, placing a limit order away from the spread increases the probability that, should a new limit order be placed, it undercuts the former. For example, if an order is placed four ticks away from the spread, the likelihood of an order five or more ticks away being fulfilled diminishes because of the new liquidity in the book, and the price impact of a market order large enough to fulfill an order at those price levels increases. Therefore, the price levels one to three ticks away from the spread become relatively more profitable in expectation, and this increased desirability makes it more likely that a new limit order will be placed such that it has price priority over the former order.

My procedure for examining the choice of price levels is as follows. For each limit order placement, I take the limit order book on the same side as the order immediately preceding the order submission and find all active price levels within ten ticks of the bid ask spread and the existing depth at each price. For orders that are submitted at a price not currently represented in the limit order book, I add said price level to my record of the order book with a quantity of zero. For each price level in the order book, I predict the optimal order size given the price and existing depth of the order book in front of the hypothetical order.

Next, for each of these predicted orders, I find the expected profit of the order, conditional on being executed, then discount that by the probability of being executed in order to find the unconditional expected profit for the order. Finally, I compare the price level of the hypothetical order that maximizes expected profit to the price level of the observed order, and record whether the observed order was placed closer to the bid-ask spread than the profit maximizing order, at the same price level, or further away from the spread.

Reported in Figure 6 is the distribution across price levels of the count of limit orders which are placed at the optimal price, those placed relatively too near the bid-ask spread, and those placed relatively too far away from the bid-ask spread. The vast majority of limit orders – 81% of orders placed during the in-sample period and 83% of orders placed during the out-of sample period – are placed at price levels too near the bid-ask spread. Interestingly, for both the in-sample and out-of-sample periods, about 91% of those orders that are placed at the model-projected optimal price are observed on the buy side of the book. Moreover, for those orders that are placed at a price level too far away from the bid-ask spread, 91% and 92% of the orders are observed on the buy side of the book over the in-sample and out-of-sample periods, respectively.

This demonstrates a potential inconsistency between the behavior of market makers and my model of optimal liquidity provision. Given the model's predictions of optimal limit orders, it appears that market makers are placing their orders too aggressively. Since profit maximizing orders are determined using the unconditional probability that they will be executed against by the next market order, this finding cannot be a result of market makers positioning their orders in the book such that the next market order is likely to fill the limit order given the current size of the order book. Rather, I argue that this observation provides evidence that market makers strategically position their orders based on their expectations of future limit order flow. The aggressiveness with which limit orders are placed with respect to their position in the order book suggests that future limit order flow is a major consideration for market makers when choosing how to provide liquidity to the market.

1.5 Estimator of Relative Liquidity Provisioning

Focusing my attention on those orders which diverge in size from the optimal \hat{q} given by the model with no private information, I can attempt to extract the private information of the market makers. While looking at the limit orders' sizes alone cannot reveal the specific channel through which the information is influencing the market maker's decision, it can provide implications of the decision, as well as the relative amount of information. I can thus define a measure of relative liquidity provision at time T (RLP_T) in as general a form as possible.

$$RLP_T = \sum_{\tau \in [0, T]} \omega(\tau, T) \mathcal{I}(q_\tau, Q_\tau, \hat{q}_\tau) \quad (33)$$

where $\omega(\cdot)$ is a weighting function that measures the residual impact of the order on the level of asymmetric information at time T , and $\mathcal{I}(\cdot)$ measures the informational content (available to the econometrician) of the order, given the order q_τ , the composition of the order book immediately prior to the order Q_τ , and the optimal \hat{q} given by the model with no private information.

Any sensible $\omega(\tau, T)$ must be strictly monotonically decreasing in $T - \tau$ so that older orders play less of a roll in determining the level of asymmetric information at time T . One natural weighting function is

$$\omega(\tau, T) = e^{\rho(\tau - T)} \quad (34)$$

where ρ is a parameter that controls the speed of decay of the weighting over time. Note that with this weighting function, since $e^{\rho(\tau - (T+t))} = e^{\rho(\tau - T)} e^{-\rho t}$, I can split the calculation of the estimator into intervals.

$$RLP_{T+t} = e^{-\rho t} \cdot RLP_T + \sum_{\tau \in [T, T+t]} e^{\rho(\tau - (T+t))} \mathcal{I}(q_\tau, Q_\tau, \hat{q}_\tau) \quad (35)$$

Structurally, it would be beneficial for (abusing notation a bit) $|\mathcal{I}(q_\tau, Q_\tau, \hat{q}_\tau)|$ to be strictly monotonically increasing in $|\hat{q}_\tau - q_\tau|$, $1/\bar{Q}_\tau$, and the probability of the order executing against the next market order. Since the probability of a limit order being filled is inversely

proportional to the aggregate depth of the order book out to the price level at which the order is placed, I can define

$$\tilde{Q}_\tau = \begin{cases} \sum_{n=1}^j Q_{n,\tau} & \text{if the order is a sell limit order} \\ \sum_{n=1}^j Q_{-n,\tau} & \text{if the order is a buy limit order} \end{cases} \quad (36)$$

where $Q_{n,\tau}$ is the aggregate depth of the order book at price level n immediately after the order at time τ is placed. The probability of the order executing against the next market order on its side of the book is then

$$P_\tau = \begin{cases} \mathbb{P}(m \geq \tilde{Q}_\tau) = 1 - F(\tilde{Q}_\tau) & \text{if the order is a sell limit order} \\ \mathbb{P}(m \leq \tilde{Q}_\tau) = F(-\tilde{Q}_\tau) & \text{if the order is a buy limit order} \end{cases} \quad (37)$$

As discussed in Section 1.2.3, there are three types of informed market makers, of which two behave symmetrically on either side of the book and one acts asymmetrically. Those informed market makers that act symmetrically will either increase their provision of liquidity to both sides of the order book or decrease their provision of liquidity to both sides of the book, depending on the information upon which they are acting. Conversely, those informed market makers who act asymmetrically have information about the fundamental value of the asset, and will increase their provision of liquidity on one side of the order book while simultaneously decreasing their provision of liquidity on the other side of the book. Thus, I construct two different information functions – one that measures the symmetric actions taken by informed market makers without knowledge about the fundamental value of the asset, and one that measures the asymmetric directional actions taken by market makers with information about the fundamental value of the asset.

The symmetric information function $\mathcal{I}^S(q_\tau, Q_\tau, \hat{q}_\tau)$ can be given by

$$\begin{aligned} \mathcal{I}^S(q_\tau, Q_\tau, \hat{q}_\tau) &= \begin{cases} (q_\tau - \hat{q}_\tau^A)(1 - F(\tilde{Q}_\tau)) & \text{if the order is a sell limit order} \\ (q_\tau - \hat{q}_\tau^B)F(-\tilde{Q}_\tau) & \text{if the order is a buy limit order} \end{cases} \quad (38) \\ &= [(q_\tau - \hat{q}_\tau^B)F(-\tilde{Q}_\tau)]^{\mathbb{1}_B} [(q_\tau - \hat{q}_\tau^A)(1 - F(\tilde{Q}_\tau))]^{1-\mathbb{1}_B}, \end{aligned}$$

and the directional information function $\mathcal{I}^D(q_\tau, Q_\tau, \hat{q}_\tau)$ is given by

$$\begin{aligned} \mathcal{I}^D(q_\tau, Q_\tau, \hat{q}_\tau) &= \begin{cases} (\hat{q}_\tau^A - q_\tau)(1 - F(\tilde{Q}_\tau)) & \text{if the order is a sell limit order} \\ (q_\tau - \hat{q}_\tau^B)F(-\tilde{Q}_\tau) & \text{if the order is a buy limit order} \end{cases} \quad (39) \\ &= [(q_\tau - \hat{q}_\tau^B)F(-\tilde{Q}_\tau)]^{\mathbb{1}_B} [(\hat{q}_\tau^A - q_\tau)(1 - F(\tilde{Q}_\tau))]^{1-\mathbb{1}_B}, \end{aligned}$$

where $\mathbb{1}_B$ is again an indicator function for the order being a bid, \hat{q}_t^A is the optimal sell limit order size given by Equation (31), and \hat{q}_t^B is the optimal buy limit order size given by Equation (32).

The construction of the symmetric information function $\mathcal{I}^S(q_\tau, Q_\tau, \hat{q}_\tau)$ allows it to be positive whenever an order provides more liquidity than the average optimal depth and negative when the order provides less liquidity. The construction of the directional information function $\mathcal{I}^D(q_\tau, Q_\tau, \hat{q}_\tau)$ allows it to be positive when the ask side of the book provides less liquidity than would otherwise be optimal and the bid side of the book provides more liquidity than would otherwise be optimal. In this way, given the behavior of market makers with information about the fundamental value of the asset discussed in Section 1.2.3.1, the directional information function is positive when market makers have information that the asset's fundamental value is above the current mid-quote price and is negative when market makers have information that the value is below the mid-quote price.

By plugging Equation (34) and Equation (38) into Equation (33), I then arrive at a

measure of symmetric actions taken by informed market makers

$$RLP_T^S = \sum_{\tau \in [0, T]} e^{\rho(\tau-T)} [(q_\tau - \hat{q}_\tau^B) F(-\tilde{Q}_\tau)]^{\mathbb{1}_B} [(q_\tau - \hat{q}_\tau^A) (1 - F(\tilde{Q}_\tau))]^{1-\mathbb{1}_B}. \quad (40)$$

Similarly, by plugging Equation (34) and Equation (39) into Equation (33), I arrive at a measure of directional actions taken by informed market makers

$$RLP_T^D = \sum_{\tau \in [0, T]} e^{\rho(\tau-T)} [(q_\tau - \hat{q}_\tau^B) F(-\tilde{Q}_\tau)]^{\mathbb{1}_B} [(\hat{q}_\tau^A - q_\tau) (1 - F(\tilde{Q}_\tau))]^{1-\mathbb{1}_B}. \quad (41)$$

Here, ρ determines the “stickiness” of any one observation’s impact on the estimate. Intuitively, the impact shouldn’t last too long. For this reason, $\hat{\rho}$ is calibrated to $\frac{1}{5*60*10^9}$ so that $e^{\hat{\rho}(\tau-T)} = e^{\text{No. of 5 min intervals since order}}$.

A representative day of the directional and symmetric estimators for American Airlines (AAL) is presented in Figures 7 and 8, respectively.

1.6 Results

My estimator being a measure of liquidity supply relative to that which is optimal in my model of liquidity provisioning, I can compare the value of the estimator to various aspects of the price process and market order flow in order to examine to what extent deviations from the model-implied optimal level of liquidity can identify the actions of informed market makers.

Information Regarding Future Market Order Flow First, I investigate whether there is evidence that market makers are acting on information about market order flow in the near future. Since my model of optimal liquidity provisioning is static, the current level of my symmetric estimator of relative liquidity provisioning is not necessarily indicative of informed market maker activity. Rather, as market makers respond to changing market conditions, the estimator is expected to reflect those responses. There is much empirical

evidence of the presence of volatility clustering in financial markets. Volatility generally follows an autoregressive process, wherein volatility is likely to remain high if it is currently high and remain low if it is currently low. It is possible, therefore, that the market makers' liquidity supply is simply a response to current market volatility and is not informative about future volatility beyond that which is due to the autoregressive nature of volatility. Thus, I will use the change in the value of my estimator as an indicator of potential informed market maker activity. While the level of the estimator may be reflective of the current market order activity, the change in the estimator's value is indicative of informed market makers' actions.

As discussed in Section 1.2.3.3, if market makers have information about future market order flow but do not have information about the presence of informed market makers, then they will decrease their provision of liquidity ahead of an increase in market order flow and will likewise increase their provision of liquidity ahead of a decrease in market order flow. Several factors regarding the distribution of market order flow affect the optimal provision of liquidity - namely, the standard deviation of the size of market orders, the average order size, the maximum order size, and the total volume of market order flow over the course of the interval. Therefore, I regress the change in these factors on the change of the symmetric measure of relative liquidity provisioning. To test the estimator's predictive power over the factors in the near future, I regress the change in these factors over one minute periods, comparing the change in the value of my estimator in the preceding interval to the change in market order flow.

$$(R1) : \Delta\sigma_t^m = \beta \cdot \Delta RLP_{t-1}^S + \varepsilon_t$$

$$(R2) : \Delta\overline{|m|}_t = \beta \cdot \Delta RLP_{t-1}^S + \varepsilon_t$$

$$(R3) : \Delta\max\{|m|\}_t = \beta \cdot \Delta RLP_{t-1}^S + \varepsilon_t$$

$$(R4) : \Delta\text{sum}\{|m|\}_t = \beta \cdot \Delta RLP_{t-1}^S + \varepsilon_t$$

If market makers have information about future market order flow, then I expect that the

coefficients for these regressions would be negative.

The results of these regressions are reported in Table 3. There are only two factors for a singular stock for which coefficients are negative with any level of significance. Over the in-sample period, the coefficients related to the change in standard deviation of market orders and maximum market order size for Goldman Sachs are negatively related to the change in RLP^S . For all other stocks, both in-sample and out-of-sample, the only coefficients that hold at any level of statistical significance are positive. Thus, there is strong evidence that market makers do not change their provision of liquidity in response to information about future market order flow. Such a result makes sense; traders tend to increase their activity when liquidity is relatively high and decrease their market order activity when the provision of liquidity is relatively low. Therefore, even if market makers had information about future market order flow, changing their provision of liquidity to reflect this information would cause market participants to adjust their market order flow in response, nullifying whatever information the market makers had.

Information Regarding the Presence of Informed Traders Next, I investigate the existence of market makers with information regarding the presence of informed traders. As discussed in Section 1.2.3.2, market makers with information about the presence of informed traders would present themselves by responding differently to changing levels of market activity than would otherwise be expected. Namely, if there were market makers with information that informed traders were active in the market, then they would provide less liquidity to both sides of the book than they otherwise would given the current market order flow, and if there were market makers with information that informed traders were not active, then they would provide more liquidity to both sides of the order book than they otherwise would. The presence of informed traders in the market cannot be measured directly and must therefore be estimated via a proxy. Since informed traders place trades with higher price impacts than noise traders, the amount by which prices change relative to the volume

of market order flow observed and the depth of the limit order book⁵ can act as an estimate of abnormal price impact, and therefore, a proxy for the presence of informed traders.

For each stock in my sample, I first partition each day into 1-minute intervals and find the level of market order flow throughout each period, the relative provision of liquidity throughout each period, as well as the change in price over each period. Then, I determine the size of price change that is unexplained by limit order book depth and market order flow by regressing the magnitude of price change observed over each period on the market order flow throughout each period as well as the estimate of relative liquidity provisioning at the end of each period (R5). The residuals of this regression represent abnormal price impacts; a positive residual indicates that prices changed more than they normally would given the market order flow and depth of the order book.

$$(R5) : |\Delta X_t| = \beta_0 + \beta_1 \cdot RLP_t^S + \beta_2 \cdot \sigma_t^m + \beta_3 \cdot \overline{|m|}_t \\ + \beta_4 \max\{|m|\}_t + \beta_5 \text{sum}\{|m|\}_t + \nu_t$$

The results of this regression are reported in Table 4. There are several instances for which the factors relating to the distribution of market order flow do not correlate with the change in price in the expected direction. For example, there are four stocks for which either average or maximum market order size is negatively correlated with the magnitude of price changes over an interval. Insofar as the market order flow from noise traders has little impact on price, it is plausible that an increase in market order flow being correlated with a decrease in price impact is indicative of increased levels of noise trader activity.

Since noise traders have a smaller price impact than informed traders, if market makers had information about the presence (or lack thereof) of informed traders, then periods of greater liquidity provisioning, relative to current market order flow, would be associated with smaller price movements, relative to the level of liquidity provisioning and level of market

⁵The price impact of market orders is mechanically tied to the depth of the order book; less depth in the order book means that market orders will consume a greater portion of liquidity and will therefore have a greater price impact. Thus, the relative depth of the order book must be accounted for when identifying periods of abnormal price impact.

order flow. Thus, in order to test for the presence of market makers that have information about the presence of informed traders, I perform a regression of the unexplained price impacts on the change of the symmetric measure of relative liquidity provisioning over the previous interval.

$$(R6) : \hat{v}_t = \beta \cdot \Delta RLP_{t-1}^S + \varepsilon_t$$

If market makers have information about the presence of informed traders, then I expect the coefficient of this regression to be negative.

The results of regression (R6) are reported in Table 5. There is only one stock for which the coefficient is negative and statistically significant during the in-sample period. However, over the out-of-sample period, five of the ten stocks have negative coefficients with statistical significance. Thus, there is evidence that some market makers possess information about the presence of informed traders and decrease their provision of liquidity ahead of periods during which those informed traders are active.

Information Regarding the Fundamental Value of the Asset Next, I investigate the hypothesis that market makers have information about the fundamental value of the asset. As discussed in Section 1.2.3.1, if market makers have information about the value of the asset, then they will decrease their provision of liquidity on the side of the order book in the same direction as the fundamental value of the asset and increase their provision of liquidity on the opposite side of the book. Thus, I regress the change in price over an interval on the change in my directional estimator of relative liquidity provisioning in the preceding interval.

$$(R7) : \Delta X_t = \beta \cdot \Delta RLP_{t-1}^D + \varepsilon_t$$

I also perform this regression after controlling for the autoregressive nature of asset returns.

$$(R8) : \Delta X_t = \beta_1 \cdot \Delta X_{t-1} + \beta_2 \cdot \Delta RLP_{t-1}^D + \varepsilon_t$$

Given the construction of my directional estimator of relative liquidity provisioning, if market makers have information about the fundamental value of the asset, then I expect the value of the coefficient acting on RLP_{t-1}^D to be positive.

The result of these regressions are reported in Table 6. Without controlling for the autoregressive nature of price returns, there are four stocks over the in-sample period and two stocks over the out-of-sample period for which there is a positive coefficient with statistical significance. After controlling for the autoregressive nature of price returns, five of the ten stocks have statistically significant positive coefficients over the in-sample period, and two stocks have significant positive coefficients over the out-of-sample period. Therefore, there is evidence that some market makers have information about the fundamental value of the asset and act upon that information.

Adverse Selection Finally, in order to test the influence of the measure of liquidity provisioning on the the price impact of market orders, I first perform a standard regression of the change in price on the size of market orders, controlling for the instantaneous price drift.

$$(R9) : \Delta X_t = \mu \cdot \Delta T_t + \alpha \cdot m_t + \varepsilon_t$$

The results of this are reported in Table 7. For each market order, I can then find the empirical price impact and compare it to the projected price impact. Regressing the residual price impact on the value of the symmetric estimator at the time the order was placed, I can then observe the effects of the market makers' liquidity provisioning decisions on the price execution of market orders.

$$(R10) : (\Delta X_t - \hat{\mu} \cdot \Delta T_t) / m_t - \hat{\alpha} = \beta \cdot RLP_t^S + \varepsilon_t$$

The results of this regression are also reported in Table 7. More costly execution of market orders is indicative of the presence of adverse selection. By the construction of the symmetric estimator, if it is able to act as a proxy for adverse selection, I would expect a smaller

estimator value to be correlated with a larger residual. Indeed, in the regression I find that the coefficient for the estimator is negative for all ten stocks over both the in-sample and out-of-sample periods and is significant at the 1% level for nine of those stocks over both time periods. Thus, it is evident that my estimator serves as a good proxy for the costs to investors borne from informed market makers.

1.7 Discussion

There are several different possible explanations for why a liquidity provider would deviate from the model-implied optimal quantity. As previously mentioned, liquidity providers may be informed about the fundamental value of the asset, future market order flow, or the presence of informed traders in the market. A market maker with private information about the fundamental value of the asset would place limit orders in such a way that takes advantage of this information - providing less liquidity in the direction of the anticipated price movement and more liquidity on the opposite side of the book. A market maker who is informed about future market order flow would be able to position her orders in anticipation of this order flow, providing less liquidity when the order flow increases than she would otherwise provide. Finally, if a market maker has information about the presence of informed traders in the market, then she would provide less liquidity when those traders are active.

It is also possible that a market maker will deviate from the expected optimal liquidity provision not because of private information but because of issues stemming from the structure of my model. Namely, it is possible that market makers deviate from the model's prediction either because they are responding to changing state variables or because they have market power. It should be noted that the estimation procedure used herein implicitly assumes a static model in which the dispersion of market order flow remains constant in average over the course of the month. Insofar as market makers are reacting to changes in current market conditions, it may therefore be the case that the deviations in liquidity provision we observe are a response to varying market conditions. Alternatively, it is possible that liquidity providers wield market power. My model assumes that market makers are

facing perfect competition. If they are not facing perfect competition, then market makers who have limit orders placed at several different price levels would systematically provide less liquidity at price levels close to the spread, positioning themselves in such a way that they extract more profits from the market order flow.

If divergences in the supply of liquidity are driven only by changing market conditions, then the actions of market makers would not reflect any anticipation of future market order flow and asset price returns beyond that of the autoregressive tendencies of the conditions. In order to determine whether market makers possess and act upon private information, I compare various predictions implied by the type of information they might possess.

If a market maker possesses private information about the fundamental value of the asset, she would place limit orders in such a way that takes advantage of this information - providing less liquidity in the direction of the anticipated price movement and more liquidity on the opposite side of the book. The construction of my directional estimator of relative liquidity provisioning leverages the above reasoning, and is positive when the actions of market makers reflect the hypothesized behavior of a market maker who anticipates positive price returns in the near future. When the provision of liquidity mirrors the behavior of a market maker who anticipates negative price returns in the near future, my directional estimator is negative. I use my directional estimator to test whether market makers have information about the fundamental value of the stocks in my sample, and I find evidence that some market makers exhibit behavior that indicates they indeed possess such information. Regressions (R7) and (R8) demonstrate that the change in liquidity provisioning has a strong predictive power over future price returns for several stocks in my sample. For these stocks, therefore, there is evidence that market makers possess and act on private information regarding the fundamental value of the stock throughout my sample.

If a market maker knows that an informed trader is present in the market, she would want to reduce her liquidity provisioning on either side of the order book ahead of the informed trader's orders. The results of regression (R6) demonstrate that, for one stock over the in-sample period and five stocks over the out-of-sample period, my symmetric measure

of relative liquidity provisioning is strongly negatively correlated with the abnormal price impact of market orders submitted to NASDAQ in the near future. Therefore, it appears that there exist market makers who have information about the presence of informed traders in the market and adjust their provision of liquidity accordingly.

A market maker without information about the presence of informed traders, but with information about future market order flow, would want to adjust her provision of liquidity accordingly. As discussed in Section 1.2.3.3, a market maker with such information would want to decrease her provision of liquidity ahead of periods of increased market order flow and increase her provision of liquidity ahead of periods of decreased market order flow. I find no evidence that market makers behave in this way. The results of regressions (R1) through (R4) show that the change of liquidity provisioning has a strong positive correlation with the change in market order flow in the near future. This result is intuitive. When liquidity provisioning is low and the order book is thinner than usual, market orders are more expensive. Thus, when the book is thin, one would expect traders to wait for more liquidity before placing their orders.

In considering the broader usefulness of this estimator, it is beneficial to investigate whether it can speak to the welfare effects of the actions taken by informed market makers. Perhaps the most direct evidence of the instantaneous health of the market for an asset is its price efficiency. Thus, the welfare effects stemming from the the liquidity provisioning decisions of informed market makers can be measured by the price impact of market orders. Regressions (R9) and (R10) test the impact that the symmetric measure of relative liquidity provisioning has on the price impact functions of each stock. As expected, when market makers decrease their providing of liquidity, the price impact of a market order increases. An increase of the price impact of market orders is indicative of decreased price efficiency, making it more costly for traders placing market orders. Clearly, therefore, the actions taken by market makers when they sense the presence of informed traders temporarily damage the liquidity of financial markets, diminishing the social welfare they provide.

However, the periods during which market makers supply less liquidity than normal are

short in duration. Therefore, while the behavior of informed market makers – reducing their supply of liquidity in anticipation of increased informed trader activity – makes order execution more costly, these effects are transient. This aligns with the empirical findings of Yueshen (2016) that, over the last decade in the U.S. equity markets, the average price impact of market orders has decreased, while the dispersion of price impact has increased. By having knowledge about the presence of informed traders, market makers are therefore able to provide more liquidity on average, improving the price execution of the average order.

1.8 Limitations

The model of optimal limit order placements proposed herein is estimated as a static model. That is to say, the estimation of this model assumes that the parameters describing the evolution of stock values and the flow of market orders remain static over the sample period. In reality, market makers must constantly re-evaluate their expectations of these features and update their strategies accordingly. Indeed, the results of regressions (R3)-(R6), presented in Table 4, suggest that market makers adjust their provision of liquidity based on recent market order flow. The estimator of relative liquidity provisioning therefore captures both the effects of changing market order flow and the effects of market makers acting on private information. Regressions (R7)-(R10), the results of which are presented in Table 5, suggest that after controlling for recent market order flow, market makers act on private information about future market order flow that extends beyond recent market order flow. While this result provides evidence that market makers have and act on private information in this manner, the estimator of relative liquidity provisioning is not purely a measure of private information.

This essay's model of optimal limit order placements is also based on the assumption that market makers are only placing and monitoring one limit order at any given time. In reality, each market maker is responsible for a portfolio of limit orders at different price levels on both sides of the order book. This limitation stems from the information available to me. NASDAQ does not release identifying information about who places orders sent to

the exchange. It is impossible, therefore, to match limit orders with one another by market participant, which is necessary if one were to construct a model in which market makers hold a portfolio of several outstanding limit orders at the same time. In this way, my model is unable to capture more complex strategies that may be employed by market makers.

Additionally, my model of optimal limit order placements implicitly assumes that market makers are able to view and respond to the market with zero latency. The latency with which market makers are able to interact with the exchange is an important consideration in constructing their strategies. As this latency increases for a market maker, the risk that new limit orders will enter into the market, undercutting her position, increases. Thus, when placing a new limit order, the market maker must optimize her order not based on the current volume of limit orders that have price and time preference but rather her expectation of this volume at the time of the next market order. The risk that a limit order will be undercut non-linearly increases with the distance of the order from the bid-ask spread. By excluding this risk from my model, its effects on optimal liquidity provisioning must instead be captured by the model's estimations of the costs associated with placing and monitoring a limit order. However, given the nature of this risk and the treatment of costs in the model, capturing this risk through these costs biases the model's estimates of optimal limit order size. For limit orders placed very close to the bid-ask spread, the model underestimates the optimal limit order size, and for orders placed far away from the bid-ask spread, the model overestimates the optimal size.

Finally, the sample of this model is restricted to ten stocks over the course of one month. This limitation arises from two separate considerations that affect the chosen sample in different ways. The restriction of stocks included in the sample is due to data availability; extending the sample to a broader array of stocks is cost prohibitive. While the cost of data is also a consideration with respect to the time window of the sample, the restriction of the sample window to one month is also due to computational limitations. The estimations performed for each stock in this study are highly memory intensive, and are dependent upon the number of orders submitted to the exchange. For several stocks in the current sample, the

memory required to estimate the model parameters exceeds 100GB. To expand the sample window beyond one month, the use of a computing cluster would be required.

1.9 Future Research

This research, and specifically the model of optimal liquidity provisioning, can be extended in two important ways. First, model parameters can be estimated on a rolling basis rather than as a static model. While allowing the estimation to be conducted on a rolling basis would decrease the transparency of the model, it would help to isolate the private information held by the market makers. Second, the model's implicit assumption that market makers can view and respond to the market with zero latency can be relaxed to allow for their expectations of future limit order flow to enter the model explicitly. Extending the model in this way would allow for a more nuanced investigation into how market makers respond to considerations of future limit order flow in addition to future market order flow. Such an extension would also allow for a better understanding of the importance of latency to market makers' optimal behavior.

1.10 Conclusion

The model proposed herein of information agnostic market makers' optimal order placement does a good job capturing the bulk of orders sent to the Nasdaq exchange. This basic result contradicts Sandås's (2001) conclusions that a similar model, which solves for the optimal aggregate order book depth using expected profit maximizing conditions, fails to explain the empirically observed shape of the order book. When solving for the optimal order size of an individual limit order placed within the context of an existing order book, my model can correctly predict the order size of over 37% of limit orders sent to the exchange, both during the in-sample and out-of-sample periods.

Insofar as the model accurately reflects the behavior of market makers who are acting without private information, it is reasonable to assume that a divergence from this model indicates that the market maker is acting on some private information. Using the differences

between the model-implied order sizes, given the existing supply of liquidity in the order book when the orders were placed, and the observed order sizes, I am therefore able to create estimators both for the presence of asymmetric information and for the implications it has on future market order flows. The estimators proposed herein contain predictive power over forthcoming asset price returns and the presence of informed traders in the market. This result suggests that the deviations from the model predictions are indeed indicative of market makers acting on private information.

Market makers' liquidity provisioning is not only responsive to evolving order flow; it correctly anticipates the market order flow activity of informed traders in the immediate future. Liquidity providers are able to get out of the way of the abnormal market order flow that otherwise would be detrimental to their profitability. Moreover, the periods of significant illiquidity in the order book are short lived. In this way, by punishing these informed traders with a relatively thin order book, market makers are able to afford to provide extra liquidity to everyone else in the market.

2 Estimating Aggregate Illiquidity of the Corporate Bond Market

2.1 Introduction

There is a rich body of work focusing on the estimation of liquidity in opaque, over-the-counter markets such as the U.S. corporate bond market. The methods by which liquidity is estimated can be broadly separated into two distinct groups: those which directly estimate liquidity via observed trading activity [Roll (1984), Amihud (2002), Feldhütter (2012), Dick-Nielsen et al. (2012), Chen et al. (2007)], and those which indirectly estimate liquidity through the use of proxies. Proxy-based approaches have been used to estimate liquidity both at the individual bond level [Mahanti et al. (2008), Longstaff et al. (2005)] and at the aggregate level [Duffie and Singleton (1997), Collin-Dufresne et al. (2001), Grinblatt (2003), Campbell and Taksler (2003)]. Currently, however, the majority of this branch of market microstructure research only attempts to directly estimate the liquidity of an individual bond issue, stopping short of using observed trading activity to estimate the broader liquidity of corporate bond markets as a whole. To the best of my knowledge, Bao et al. (2011) represents the only attempt to create an aggregate liquidity measure that is based on observed trading activity. In this essay, I construct a measure of the illiquidity of the overall U.S. corporate bond market that is estimated using observed trading activity.

The difficulty of using a liquidity measure that is directly estimated from trading activity stems from the fact that the majority of outstanding corporate bonds do not trade on any given day. If an estimate of market liquidity is constructed via a simple aggregation of the observed liquidity of those bonds that do trade, then the estimate will be biased in favor of the type of bonds that are traded on that day, and will therefore not be representative of the liquidity available to the broader market. Bao et al. (2011) construct an estimate of overall bond market liquidity by taking the median monthly Roll liquidity estimate of the 1,035 most actively traded investment grade bonds in their sample. Bao's estimation procedure

restricts the sample to bonds that trade on at least 75% of the business days during which they are outstanding. As will be demonstrated in the next section, such a restriction is not representative of the trading patterns for the majority of U.S. corporate bonds. Thus, the estimate of aggregate liquidity is heavily biased towards a small subset of the most liquid bonds in the market. For this reason, while useful for the investigation of the connections between liquidity and credit spreads, the aggregate measure is inappropriate for estimating the liquidity of the bond market as a whole.

I propose a measure of overall market liquidity that combats this bias by first estimating the latent liquidity of bonds that are not traded, then aggregating the liquidity across the entire market. In this context, I define the latent liquidity of a bond that is not traded to be the liquidity that would have been observed had the bond traded. The latent liquidity of un-traded bonds is estimated using the observed liquidity of similar bonds that are traded. Thus, the majority of this paper focuses on the method by which that measure is extended to bonds that are not traded.

I can decompose the liquidity of a bond into two components: that which is driven by the overall liquidity of the market, and that which is particular to the individual bond issue. A contention of this paper is that the component of a particular issue's liquidity which is exogenous to the market is largely driven by the ex-ante probability that it will trade. The willingness of a dealer to buy a bond is, in large part, driven by her expected costs either of searching for the opposite side to the trade or of keeping the bond in her inventory until such time that another party wishes to buy the bond in question. Alternatively, a dealer's willingness to sell a bond which she holds in her inventory is determined by weighing the opportunity costs associated with keeping the bond in her inventory against the premium she expects the other party will be willing to pay due to the search costs associated with finding another dealer who is also holding the bond in their inventory. The spread that dealers charge and the volume they are willing to absorb must therefore be related to the probability that the bond will trade. A bond issue with a high probability of trading will be easier to sell and have lower search costs than one with a low probability.

The liquidity of a bond issue is directly linked to dealers' willingness and ability to trade the asset. In turn, the dealers' ability is shaped by their expectations about the ongoing demand for liquidity in the respective bond issue. Having constrained inventories, the dealers must consider the potential impact a trade would have on their ability to absorb future liquidity. This inventory consideration creates a feedback effect wherein bond dealers prepare to provide liquidity for the issues which they expect to trade, thereby reducing the trade costs for those issues, making it more likely that the issues will be traded. Conversely, if dealers think that there will not be much ongoing demand for liquidity in an issue, they will not be as willing or able to provide liquidity, increasing the trade costs, making it less likely that the issue will be traded. In this way, the dealers' expectations are not only self-fulfilling, they are a factor driving the liquidity with which bonds trade.

While each bond has its own (path dependent) liquidity patterns, as the saying goes, a rising tide lifts all boats. In this case, rising levels of liquidity in the market lift the liquidity of all bond issues. When considering the impact that market liquidity has on individual bonds, however, it is important to note that the asymmetric distribution of trading activity across bonds leads to a non-linear impact on different bonds. To extend the metaphor, while a rising tide lifts all boats, in this case it lifts some boats more than others. It is useful, therefore, to observe the composition of market activity in some stylized sense.

At any given time, a very small minority of bond issues are responsible for the majority of all trading activity and volume. A bond issue in this minority is traded almost daily, so the probability of the issue trading is consistently quite high and the fluctuations in liquidity must be driven primarily by outside market forces. On the opposite end of the spectrum, those bonds which are dormant primarily face the challenge of dealers' willingness to absorb the demand for liquidity and search for an opposite side to the trade. For this reason, one would expect that a dormant bond's liquidity is primarily determined by the ex-ante probability that it will trade.

The decomposition of liquidity into an issue's individual liquidity (which is driven by its probability of trading) and market driven liquidity (which affects issues in a non-linear way)

leads to a procedure which is broken into two primary parts. First, I create an estimate for the (path dependent) probability of trading given the bond issue’s characteristics and trade history. I estimate this propensity to trade by implementing a hidden Markov model, which leverages Bayesian updates given an issue’s trade history to better predict the underlying factors which inform the probability of trading. I then use this probability, along with other bond characteristics, to help inform my estimation of what the liquidity would have been for an issue which didn’t trade had it actually traded. Since market liquidity affects bonds asymmetrically, I must estimate the effect that the market would have on a particular issue by interpolating between a group of similar bonds. I perform this interpolation using a locally weighted multivariate polynomial regression (LOWESS), which estimates the latent liquidity for each observation without trades by performing a weighted least squares regression on a localized subset of data, selecting the localized subset via a nearest neighbor meta-model.

The organization of this essay is as follows. Section 2.2 introduces the data used in the construction of the liquidity measure and briefly examines the observed distribution of trading activity. Section 2.3 describes the trade-based estimate of individual bond liquidity. Section 2.4 discusses the hidden Markov model estimate of the probability a bond will trade; the structure of the model and the procedure by which it is tuned are described, and the performance of the model is compared to alternatives. Section 2.5 explains the LOWESS regression model used to extend liquidity estimates to non-traded bonds, details the tuning procedure for the model, and compares its performance to alternative models. Section 2.6 aggregates the bond-level liquidity estimates into a measure of overall market liquidity and explores the relationships between market liquidity and various macroeconomic factors with ties to liquidity. Section 2.7 discusses the limitations of my estimator, Section 2.8 discusses potential avenues for future research, and Section 2.9 concludes.

2.2 Data

Using FINRA’s Trade Reporting and Compliance Engine (TRACE) enhanced historical data from January 1, 2006 through March 31, 2018, I select all bond issues that I am able to match

with their respective issue dates and maturity dates⁶. I am left with 19,993 issues spanning 2,173 different corporations. For each of these bonds, I gather all of the trades that are reported to TRACE. All FINRA-registered broker-dealers are obligated to report any transaction of a TRACE-eligible security⁷ in which they engage. Thus, for each bond, my data set includes the transaction data for all of the trades in which broker-dealers participated. This transaction data includes the date and time of the trade, the volume traded, the price at which the trade took place, and a unique identifier for the trade.

Occasionally, there are errors in the transaction data that are reported to FINRA. For these transactions, corrected transaction data are reported to FINRA, and these transactions are assigned a new unique identifier while simultaneously retaining the original identifier as a separate variable. I am therefore able to identify those trades for which inaccurate transaction data was reported, remove all such errant transactions, and replace them with their corrected counterparts.

The distribution of trading activity across bond issues is highly skewed. On average, trades are observed on 24.2% of the days during which bonds in this sample are outstanding. Of the 19,993 bonds in the sample, however, 11,475 have trades observed on fewer than 5% of the days over which they are outstanding, while 983 bonds have trades observed on at least 90% of the days over which they are outstanding. The distribution of trading frequency is displayed in Figure 9.

The distribution of daily trading volume of individual bonds (shown in Figure 10) is even more skewed than the distribution of trading frequency, following an approximately log-normal distribution with a median of \$450,000. 10% of daily volumes are greater than \$10 million, and the maximum daily trade volume observed for an individual bond is slightly

⁶In order to appropriately train and use my estimator of the probability of trading, it is necessary to know the issue date and maturity date of the bonds being estimated. Bonds must only enter into the estimator while they are outstanding. If, during the training of the model, a bond enters the estimator before its issue date or remains in the estimator after its maturity date, then the projected probability of trading will be biased downward for all bonds and the error of the model will be skewed upward. Moreover, after training the model, if a bond enters the estimator before its issue date, then the projected probability of trading will be biased downward for the lifetime of the bond.

⁷All corporate bonds are TRACE-eligible. However, the trade history of non-investment-grade rated bonds are only published in the enhanced historical TRACE data set.

over \$485 million. When restricting the sample to those bonds that trade on at least 90% of the business days over which they are outstanding, I find the median daily volume to be \$1.99 million. Thus, not only do a small portion of the bonds trade much more frequently than the rest of the bonds in the sample, but these actively traded bonds also tend to have much larger daily volumes than the remainder of bonds in the sample.

2.3 Trade-Based Estimate of Bond Liquidity

There are several measures that can be used to directly estimate the liquidity of a bond given observed trading activity [Roll (1984), Amihud (2002), Feldhütter (2012), Dick-Nielsen et al. (2012)]. For the purpose of this paper, I aim to measure the cost of liquidity by estimating the spread demanded by broker-dealers trading a particular bond. There are two such measures proposed in the literature, Roll (1984) and Feldhütter (2012). Roll (1984) proposes to estimate the spread of a security by assuming that consecutive trades occur on opposite sides of the spread. With this assumption in place, we know that there is a negative auto-correlation in the series of price returns and can therefore estimate the spread as $2\sqrt{-cov(\Delta p_t, \Delta p_{t-1})}$. Feldhütter (2012) estimates the spread by leveraging the microstructure observation that bond trades are often matched by the dealer, with both sides of the transaction identified, before a transaction takes place. Matching trades of equal volume that occur within a short time-window, we are able to construct what Feldhütter denotes an “imputed round-trip trade.” The difference between the price at which a dealer buys a bond from a customer and the price at which she sells a bond to a client is the “round-trip cost” of the series of trades⁸. Feldhütter extends this approach to structurally estimate a search model, but the methodology can be adopted to directly estimate the spread.

I create my trade-based estimate of individual bond liquidity by adopting Feldhütter’s notion of imputed round-trip trades and their accompanying round-trip costs. For bond

⁸Note that the matching algorithm used to find imputed round-trip trades allows for the possibility that a dealer sells the bond to second dealer, who in turn sells the bond to her own client. In this way the imputed round-trip cost estimates the total spread that is imposed upon investors. Since dealers often trade bonds among themselves before eventually selling the bond to a client, other measures of spread have the possibility of underestimating the effective spread affecting investors.

trades that occur on a given day, I isolate trades of equal volume. For each series of trades of a given volume that occur within 15 minutes of one another, I classify the series of trades as an imputed round-trip trade. The difference between the maximum and minimum price observed in the imputed round-trip trade is the round-trip cost. Imputed round-trip trades with zero round-trip cost are denoted “immediate matches” a la Green et al (2007), and are removed from the estimator. For the bond and day in question, my estimate of the individual bond liquidity is the log-volume weighted average of the round-trip costs of all the imputed round-trip trades that occurred on that day.

2.4 Estimating the Probability of Trading

The challenge of creating a forward-looking estimate of the propensity for a bond issue to trade is as follows. Having only the characteristics and realized trade history of an issue, I wish to establish the probability that the bond will trade on the next trading day. For this exercise, I use a hidden Markov model⁹, which assumes that bonds belong to one of three (unobserved) type states and establishes a probability distribution of these types. The distribution over states updates at the end of each trading day given the realized activity that day. In this way, the model’s probability distribution undergoes Bayesian updates each day evolving towards a path-dependent estimate of the current type to which the bond belongs. Each day, given the bond’s characteristics, I also calculate a conditional probability of trading for a bond belonging to each type. These type-dependent probabilities of trading are then combined with the probability distribution over the type space to arrive at an overall probability of trading.

⁹Hidden Markov models are a form of dynamic Bayesian network models in which the data generating process is driven by an underlying state. The states are un-observable to the econometrician, and are assumed to follow an action-dependent Markov process. In this sense, while the states are hidden, the observed actions feed into Bayesian updates which inform the probability over possible states. This updating procedure, when included in an optimization, not only determines the most likely final state but also the most likely path of states which would result in the observed realizations. As more observations are included, the econometrician’s estimate of the underlying state thus becomes more statistically powerful, leading to more accurate predictions of the dependent variable.

2.4.1 Hidden Markov Model Tuning Procedure

Suppose that there is a hidden Markov model describing the probability that a bond issue will trade on a given day. There are I issues which we observe over T time periods. An issue can belong to one of three unobserved states: actively traded, commonly traded, and dormant. The probability of an issue being in each of the states is described by a 1x3 vector, π , with an initial distribution $\pi_0 = [\alpha_1, \alpha_2, (1 - \alpha_1 - \alpha_2)]$ for $\alpha_1, \alpha_2 \in (0, 1)$ such that $\alpha_1 + \alpha_2 < 1$. $\pi(1)$ represents the probability that the bond issue is the active type, and $\pi(3)$ is the probability that the bond issue is of the dormant type.

The issues can transition between the three states, and do so following action-dependent rules of motion. These rules can be expressed through transition matrices wherein transition probabilities vary depending on whether the issue traded in the period or not, with the matrix Tr_1 representing the transition probabilities given a trade occurs and the matrix Tr_0 the probabilities given no trade.

$$Tr_1 = \begin{bmatrix} 1 & 0 & 0 \\ 1 - \gamma_2 & \gamma_2 & 0 \\ 0 & 1 - \gamma_3 & \gamma_3 \end{bmatrix} Tr_0 = \begin{bmatrix} \lambda_1 & 1 - \lambda_1 & 0 \\ 0 & \lambda_2 & 1 - \lambda_2 \\ 0 & 0 & 1 \end{bmatrix},$$

where $\gamma_2, \gamma_3, \lambda_1, \lambda_2 \in (0, 1)$. With this construction, $Tr_a(i, j)$ gives the probability that a bond issue of type i transitions to type j after action a is observed. Thus, if I observe a bond trading on date t , $\pi_t = \pi_{t-1} * Tr_1$ gives the updated probability distribution for bond type after date t . In this way, the transition matrices act as simple Bayesian updates to our estimate of the bond's probability distribution over the latent states. Iterating over a bond's trade history therefore provides an efficient path-dependent estimate of the current distribution over states.

Further assume that the bonds have state-dependent functions mapping various bond characteristics to the probability of trading. Letting $y_{i,t}$ represent the binary variable indi-

cating whether bond i is traded on date t , I have

$$y_{i,t} = \begin{cases} 1 & \text{if } \beta_0^s + \beta_1 X_{i,t} + \varepsilon_{s,i,t} > 0, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

where β_0^s is a state-dependent parameter, β_1 is a vector of parameters that applies equally across states, $X_{i,t}$ is a vector of bond characteristics, and $\varepsilon_{s,i,t}$ is an error term following the type-I extreme value distribution. Thus, the logistic function $F_s(X_{i,t})$ provides the conditional probability that a trade will occur given state s .

Combining these conditional probabilities of a trade occurring with the probability distribution over latent states yields the unconditional probability that a trade will occur for bond,

$$\mathbb{P}(y_{i,t} = 1) = \sum_S \pi_{i,t-1}(s) \cdot F_s(X_{i,t}) = \pi_{i,t-1} F(X_{i,t}), \text{ for } F(X_{i,t}) = \begin{bmatrix} F_1(X_{i,t}) \\ F_2(X_{i,t}) \\ F_3(X_{i,t}) \end{bmatrix}.$$

By leveraging the Markov property of the distribution over latent states, I have the probability distribution over latent states for bond i after date t ,

$$\pi_{i,t} = \pi_0 \prod_{n=1}^t Tr_1^{y_{i,n}} Tr_0^{(1-y_{i,n})}.$$

Substituting this into the unconditional probability that a trade will occur for bond i on date t , I have

$$\mathbb{P}(y_{i,t} = 1) = \pi_0 \prod_{n=1}^{t-1} \left(Tr_1^{y_{i,n}} Tr_0^{(1-y_{i,n})} \right) F(X_{i,t}).$$

Applying the unconditional probability that a trade will occur across all bonds $i \in I$ and dates $t \in T$, with some abuse of notation I am therefore able to find a maximum likelihood estimate of the parameters $\hat{\theta} = [\hat{\alpha}_1, \hat{\alpha}_2, \hat{\gamma}_2, \hat{\gamma}_3, \hat{\lambda}_1, \hat{\lambda}_2, \hat{\beta}_0^1, \hat{\beta}_0^2, \hat{\beta}_0^3, \hat{\beta}_1]$ describing the initial state distribution, both the transition probability matrices, and the state-dependent

trade probabilities by solving

$$\begin{aligned}
\hat{\theta} &= \underset{\theta \in \Theta}{\operatorname{argmax}} \prod_I \prod_T \left[\left(\pi_0 \prod_{n=1}^{t-1} \left(Tr_1^{y_{i,n}} Tr_0^{(1-y_{i,n})} \right) F(X_{i,t}) \right)^{y_{i,t}} \right. \\
&\quad \left. \times \left(1 - \pi_0 \prod_{n=1}^{t-1} \left(Tr_1^{y_{i,n}} Tr_0^{(1-y_{i,n})} \right) F(X_{i,t}) \right)^{(1-y_{i,t})} \right] \\
&= \underset{\theta \in \Theta}{\operatorname{argmax}} \left[\sum_I \sum_T y_{i,t} \ln \left(\pi_0 \prod_{n=1}^{t-1} \left(Tr_1^{y_{i,n}} Tr_0^{(1-y_{i,n})} \right) F(X_{i,t}) \right) \right. \\
&\quad \left. + \sum_I \sum_T (1 - y_{i,t}) \ln \left(1 - \pi_0 \prod_{n=1}^{t-1} \left(Tr_1^{y_{i,n}} Tr_0^{(1-y_{i,n})} \right) F(X_{i,t}) \right) \right].
\end{aligned}$$

In order to evaluate the out-of-sample performance of the Hidden Markov Model estimator for the probability of trading, I split the data into a training set (X^{Train}, Y^{Train}) and a test set (X^{Test}, Y^{Test}) such that

$$t \in T^{Train} \quad \forall (\vec{x}_{it}, y_{it}) \in (X^{Train}, Y^{Train}),$$

where T^{Train} comprises the first $\sim 78\%$ of the dates in the data set. Specifically, $t \leq$ October 13, 2015 $\forall t \in T^{Train}$. I train the model by performing the above described maximum likelihood estimation on the training data set, which is comprised of a matrix of bond characteristics, X^{Train} , consisting of indicator variables for the day of the week (with indicators for Tuesday through Friday) and a vector of indicator variables for whether or not bonds are traded for each day, Y^{Train} . The parameters that are found by performing this maximum likelihood estimation are presented in Table 8. Using the parameters given by that optimization, I am able to use the model to create an estimate of the likelihood of a trade occurring for each day throughout the life of each bond in the complete data set. I am then able to test the in-sample and out-of-sample performance of the estimator. In the following section, I assess the performance of the estimator and compare it to alternative models.

2.4.2 Performance of Hidden Markov Model Estimator

The distribution of probability estimates given by the hidden Markov model is presented in Figure 11. The distribution is strongly bimodal, with a shape resembling that of the distribution of trading frequency (Figure 9). I decompose the probability estimates into the distributions of ex-ante trading probabilities for those observations for which a trade occurs and those for which no trade occurs. These distributions are presented in Figures 12 and 13.

In order to assess the accuracy of my hidden Markov model estimator of the probability that a bond will trade, I benchmark the performance of my model against alternative estimators across various metrics of classification model performance. For each of the alternative models, I train the model on the same training data set (X^{Train}, Y^{Train}) as is used to train the hidden Markov model. After training these models, I use them to generate alternative estimates of the probability that each bond will trade throughout the entire data set. I am then able to compare these probability estimates to the observed trade data. Thus, I can evaluate the in-sample and out-of-sample performance of my model and of the alternative models.

2.4.2.1 Alternative Models

Standard Logistic Regression Since estimating the probability that a bond will trade is a binary classification problem, the natural choice for an alternative model is a logistic regression. Letting $y_{i,t}$ represent the binary variable indicating whether bond i is traded on date t , we have

$$y_{i,t} = \begin{cases} 1 & \text{if } \beta_0 + \beta_1 X_{i,t} + \varepsilon_{i,t} > 0, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

where $\theta = [\beta_0, \vec{\beta}_1]$ is a vector of parameters, $X_{i,t}$ is a vector of bond characteristics, and $\varepsilon_{s,i,t}$ is an error term following the type-I extreme value distribution. Thus, the logistic function $F_\theta(X_{i,t})$ provides the probability that a trade will occur. The logistic regression

model is trained on (X^{Train}, Y^{Train}) by solving for the maximum likelihood estimator

$$\begin{aligned}\hat{\theta} &= \underset{\theta}{\operatorname{argmin}} \prod_I \prod_T F_{\theta}(X_{i,t})^{y_{i,t}} (1 - F_{\theta}(X_{i,t}))^{(1-y_{i,t})} \\ &= \underset{\theta}{\operatorname{argmax}} \sum_I \sum_T [y_{i,t} \ln (F_{\theta}(X_{i,t})) + (1 - y_{i,t}) \ln (F_{\theta}(X_{i,t}))].\end{aligned}$$

The estimated parameters $(\beta_0, \vec{\beta}_1)$ of the logistic regression model are presented in Table 9.

Weighted Logistic Regression As discussed in the description of the data in Section 2.2, the trading activity is quite imbalanced - with a significant portion of the bonds trading very infrequently. This stylized fact leads to the existence of many more observations $y_{i,t}$ in which no trade is observed than there are observations in which a trade is observed. This imbalance will bias the estimates given by the standard logistic regression downward by effectively assigning greater weight to the observations in which no trade is observed. Weighted logistic regressions help to control for the bias introduced by imbalanced data sets by applying class-weights that tune the relative importance of observations belonging to the underrepresented class. Since this model is being applied to a binary classification problem, the class-weights are fully parameterized by a single class-weight hyper-parameter $\omega > 0$. The weighted logistic regression model extends the standard logistic regression model, and is trained on the same training data by solving for the maximum likelihood estimator

$$\hat{\theta} = \underset{\theta}{\operatorname{argmax}} \sum_I \sum_T [\omega y_{i,t} \ln (F_{\theta}(X_{i,t})) + (1 - y_{i,t}) \ln (F_{\theta}(X_{i,t}))],$$

with the hyper-parameter ω chosen such that the impacts that the classes have on the likelihood estimate are equalized. The estimated parameters $(\beta_0, \vec{\beta}_1)$ and hyper-parameter ω that optimize the weighted logistic model are presented in Table 9.

Unconditional Probability To provide a baseline estimator to the bench-marking of my hidden Markov model against alternative estimators, I also include an unconditional trade probability model $\widehat{y_{i,t}} = \alpha$, where $\alpha \in (0, 1)$. This model is trained on Y^{Train} by

solving for the maximum likelihood estimator

$$\begin{aligned}\hat{\alpha} &= \underset{\alpha}{\operatorname{argmax}} \prod_I \prod_T \alpha^{y_{i,t}} (1 - \alpha)^{(1 - y_{i,t})} \\ &= \underset{\alpha}{\operatorname{argmax}} \sum_I \sum_T [\alpha y_{i,t} + (1 - \alpha)(1 - y_{i,t})].\end{aligned}$$

The maximum likelihood estimator α that is found for the unconditional probability model when fit to the training data is given in Table 9.

2.4.2.2 Bench-marking of Hidden Markov Model against Alternative Models

Receiver Operator Characteristic (ROC) curves illustrate the ability of binary classification models to distinguish between the two classes by graphing the true positive rate against the false positive rate for over the range of various threshold values. The closer a model's ROC curve approaches to the upper left corner of the graph, the greater its ability to distinguish between the two classes. A random classifier model with no ability to distinguish the two classes from one another possesses a ROC curve that bifurcates the ROC space with a diagonal line from the bottom left to the top right of the graph. The area under the ROC curve (ROC AUC) is a metric that summarizes the ability for the classification model to discriminate between the two classes, and can be conceptualized as the probability that the model will assign a higher score to a randomly chosen positive observation than it will a randomly chosen negative observation.

The ROC curves illustrating the in-sample and out-of-sample performance of the hidden Markov model and its alternatives are presented in Figures 14 and 15, respectively. The hidden Markov model clearly outperforms the alternative models, both in-sample (over the training data set) and out-of-sample (over the test data set). The ROC AUC of the hidden Markov model and the alternative models, both in-sample and out-of-sample, is presented in Table 10. The ROC AUC of the hidden Markov model over the training data set is 0.9473, and the AUC over the test data set is 0.9519, demonstrating a high level of external predictive validity. The next-best alternative model - the weighted logistic regression model - achieves AUCs of 0.8268 and 0.8272 over the training and test data sets, respectively.

The weighted logistic regression shows a marked improvement over the standard logistic regression, which achieves AUCs of 0.5759 and 0.5758 over the training and test data sets, meaning that it differentiates days in which a trade occurs from days in which there are no trades only slightly better than random chance. As would be expected, the unconditional probability model has no ability to differentiate between the two classes; the model has an AUC of 0.5 over both the training and test data sets.

While the ROC AUC summarizes a classification model’s ability to differentiate between the two classes, it does not take into account the confidence of the model’s predictions. That is to say, the ROC AUC does not measure a model’s ability to assign higher probability to samples that are more likely to be positive. To measure this aspect of classification model accuracy, it is useful to consider the log loss of the model, which is given by $-\frac{1}{N} \sum_{i=1}^N (y_i \ln(\hat{y}_i) + (1 - y_i) \ln(1 - \hat{y}_i))$. Given this construction, a lower log loss is indicative of a model with better ability to predict the classes with higher probability. The log loss of the hidden Markov model and of the alternative models, both in-sample and out-of-sample, is presented in Table 10. Once again, the hidden Markov model outperforms the alternative models over both the training and test data sets. The log loss of the hidden Markov model is 0.2490 over the training data and 0.2325 over the test data. The next-best alternative model is again the weighted logistic regression model, which has a log loss of 0.4897 over the training data and 0.4827 over the test data.

2.5 Matching Estimator

Having estimates for each bond issue’s ex-ante probability of trading, I turn my attention to the the estimation of bond liquidity. For those bonds which do not trade, the goal is to estimate what the liquidity would have been by looking at the realized liquidity of those issues which are most similar. The demand for a robust matching algorithm leads naturally to the implementation of a nearest neighbors search. As this is an interpolation problem, it is especially useful to consider nearest neighbor search algorithms which include multiple observations in order to add consistency to the estimates. Such an estimator therefore acts

as a kind of kernel density estimator, which provides strong non-parametric estimates with low asymptotic error.

For this matching estimator, I use a locally weighted multivariate polynomial regression (LOWESS), as developed by Cleveland (1978) and extended by Cleveland, Devlin, and Grosse (1988). The LOWESS regression is a non-parametric regression model that synthesizes several weighted least squares regressions - each applied to localized subsets of the data - in a nearest neighbor meta-model. In its most general form, a LOWESS regression is carried out in the following manner. For each point x_0 , the local regression output \hat{y}_0 is given by $x_0\theta$, where $\theta_0 = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^m w_i (y_i - x_i\theta)^2$ for the m observations nearest x_0 , with weights w_i given by some weighting function that satisfies the conditions specified in Cleveland (1978).

I apply a LOWESS regression to each day separately, using the observed liquidity of bonds that trade on a given day to estimate the liquidity of bonds that do not trade on that day. The parameters of each LOWESS regression are separately found for each point at the time of estimation and cannot be tuned directly. I therefore calibrate the matching estimator by tuning hyper-parameters that indirectly affect the selection and weighting of the neighboring observations used for each estimation. The specific procedure by which these hyper-parameters are tuned is described in the following section.

2.5.1 LOWESS Regression Tuning Procedure

For the purpose of this discussion, let y_{it} represent the liquidity of bond i on date t , and let x_{iht} represent characteristic h of bond i on date t . Let X be the set of bond characteristics, such that each row $\vec{x}_{it} \in X \quad \forall i \in I, t \in T$ is a vector of characteristics for bond i on date t . These characteristics include (a) the ex-ante probability that the bond will trade, as estimated by the hidden Markov model, (b) the number of days since the previous day the bond traded, (c) the trading volume of the bond issue over the previous week, (d) the observed liquidity of the bond issue over the previous week, (e) an indicator of whether other bond issues belonging to the same corporation were traded, (f) the number of trades of other

bond issues belonging to the same corporation, (g) the cumulative trading volume of other bond issues belonging to the same corporation, (h) the cumulative trading volume of other bond issues belonging to the same corporation over the previous week, (i) the number of days since any other bond issue belonging to the same corporation traded, (j) the average observed liquidity of other bond issues belonging to the same corporation, and (k) the average observed liquidity of other bond issues belonging to the same corporation over the previous week.

Restricting the data to observations for which the trade-based liquidity measure y_{it} is observable, I am left with a set of data X^{Trade} and corresponding Y^{Trade} with which I can calibrate and test my matching estimator. The remaining data $X^{NoTrade}$ is the set of bond characteristics for which no liquidity data is observable (and for which this estimator aims to project the latent liquidity \tilde{y}_{it} that would have been observed had the bonds traded). To maintain consistency with the procedure used to train and test the Hidden Markov Model from the previous section, I further split X^{Trade} and Y^{Trade} into a training set (X^{Train}, Y^{Train}) and a test set (X^{Test}, Y^{Test}) such that

$$t \in T^{Train} \quad \forall (\vec{x}_{it}, y_{it}) \in (X^{Train}, Y^{Train}),$$

where T^{Train} is the same training window as used in the tuning of the Hidden Markov Model.

In order to maintain tractability with my estimation procedure and avoid the costs associated with high-dimensional nearest neighbor searches, I first perform dimension reduction on the bond characteristics data. To accomplish this, I perform principal component analysis on the training set of bond characteristics X^{Train} . I project the matrix X^{Train} onto a two-dimensional orthogonal space such that the resulting matrix $C^{Train} = X^{Train}W_2$ preserves the maximal amount of variance from the original data¹⁰. This process has the dual effect of preserving the tractability of a nearest neighbor search over the set of bond characteris-

¹⁰Note here that for the bond characteristic data in question, I find that C^{Train} preserves over 90% of the variance from X^{Train} . Thus, C^{Train} provides an acceptably small amount of information loss from X^{Train} , with the span of the orthogonal bases approximating a dimension reduction subspace for the row space of X^{Train} .

tics while simultaneously increasing the signal-to-noise ratio of the resulting set of principal components.

Let n_t be the number of observations on date t . Thus, $N^{Train} = \sum_{T^{Train}} n_t$ is the total number of observations in the training set. The matching estimator is calibrated by finding the hyper-parameters $(\tilde{q}, \tilde{\alpha}_1, \tilde{\alpha}_2)$ such that

$$(\tilde{q}, \tilde{\alpha}_1, \tilde{\alpha}_2) = \underset{(q, \alpha_1, \alpha_2)}{\operatorname{argmin}} \sqrt{\frac{1}{N^{Train}} \sum_{T^{Train}} \sum_I (y_{it} - \hat{y}_{it})^2},$$

where $\hat{y}_{it} = c_{it}\theta_{it}$ such that

$$\theta_{it} = \underset{\theta}{\operatorname{argmin}} \sum_{j=1}^{m_t} w_j \left(y_{jt} - c_{jt} \begin{bmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{bmatrix} \theta \right)^2,$$

with weights w_j given by the tri-cube function applied to the euclidean distance between $c_{it} \begin{bmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{bmatrix}$ and $c_{jt} \begin{bmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{bmatrix}$ for the $m_t = n_t q$ observations at date t nearest c_{it} . In this way, the hyper-parameters (q, α_1, α_2) control the behavior of the LOWESS regression. $q \in (0, 1]$ determines the number of neighboring observations to include in each local regression as a portion of the total number of observations for each day. $\alpha_1, \alpha_2 > 0$ scale the principal component vectors, thereby determining the relative importance of each principal component in the calculation of distance between two observations. The value of these hyper-parameters selected for the LOWESS regression model are reported in Table 11.

2.5.2 Performance of LOWESS Regression Matching Estimator

In order to assess the accuracy of my LOWESS regression matching estimator of the latent liquidity of bonds that are not observed to trade, I benchmark the model against alternative matching estimators across various metrics of regression model performance. For each of the alternative models, I train the model on the same training data set (X^{Train}, Y^{Train}) as is used to train the LOWESS regression model. After training these models, I use them to

generate alternative estimates of the liquidity of bonds based on the observed liquidity of other bonds on the day in question throughout the entire data set. I am then able to compare these liquidity estimates to the observed trade data. Thus, I can evaluate the in-sample and out-of-sample performance of my model and of the alternative models.

2.5.2.1 Alternative Models

Ordinary Least Squares (OLS) Regression When choosing an interpolation model, an OLS regression model is a natural consideration and is therefore taken as an alternative model. The formulation of this approach is similar to that of the LOWESS regression model, with the exception being that for each day an OLS regression is performed on the entirety of the observations on that day. For each day, the liquidity estimates of bonds for which liquidity is observable are regressed against their respective principal components. Note that with this formulation, there are no hyper-parameters that can be tuned. Therefore, the daily OLS regression model requires no pre-training.

K-Nearest Neighbor Search Given the non-linear relationship that exists between bond characteristics, it is useful to consider a k-nearest neighbor search as an alternative model. Rather than performing a local regression on the neighboring observations (as is done in the LOWESS regression model) for each day, a simple average is taken of the liquidity observations neighboring the observation in question. In this way, the k-nearest neighbor model is a less complex estimator as compared to the LOWESS regression model. That being said, the computational complexity of the k-nearest neighbor search meta-model in the LOWESS regression dwarfs that of the weighted local regressions. Thus, by employing a k-nearest neighbor model rather than a LOWESS regression, a significant amount of accuracy is sacrificed to achieve only a minor improvement in computational speed.

The k-nearest neighbor model is calibrated by solving for the hyper-parameter \hat{k} such

that

$$\hat{k} = \underset{k}{\operatorname{argmin}} \sqrt{\frac{1}{N^{Train}} \sum_{t \in Train} \sum_I (y_{it} - \sum_{j=1}^k y_{jt})^2},$$

where the k nearest observations to c_{it} at date t are determined by the euclidean distance between c_{it} and c_{jt} . In this way, the hyper-parameter $k \in \mathbb{N}^+$ controls the behavior of the k -nearest neighbor model, denoting the number of neighboring observations to include in each estimation. The value of this hyper-parameter chosen for the k -nearest neighbor model is reported in Table 11.

Deep Neural Network Deep neural networks are useful for identifying complex non-linear relationships without the need to impose prior knowledge on the structure of these relationships. Taking a vector \vec{x}_{it} of regularized bond characteristics as an input, each characteristic enters the neural network through its own input node, independent from the other input nodes. Each input node is mapped to the next layer of nodes via an array of scalar weights such that it is connected to every node in the next layer. The next layer is the first “hidden” layer of the network; neither its inputs nor its outputs are directly visible to the econometrician. For each node in this hidden layer, the input to the node is the sum of all the nodes from the previous layer, scaled by their respective weights corresponding to this node. The output of each node in this hidden layer is determined by a Rectified Linear Unit (ReLU)¹¹ activation function $f(x) = \max\{0, x\}$. The outputs from each node of the first hidden layer are mapped to each node of the next hidden layer via an array of scalar weights. Each hidden layer performs in the same way as the the first hidden layer, with nodes taking the sum of their inputs and passing that through the activation function before sending their output to each of the nodes in the next layer, scaled by an array of weights. The outputs from the final hidden layer are mapped via scalar weights to one output node. The output node takes the sum of the scaled inputs from each of the nodes in the final hidden layer and

¹¹While the ReLU activation function is not differentiable at 0, it has been found that gradient descent learning algorithms still perform well in practice [Goodfellow, et al. (2016)]. The computational simplicity, representational sparsity, and local linearity of the ReLU activation function make it a popular choice for training deep neural networks [Glorot, et al. (2011)].

passes that sum as the network’s prediction of the bond’s liquidity measure.

The deep neural network is trained as follows. First, the model is initialized with small random weights for each of the connections between the nodes in one layer and the nodes in the next layer. Next, the matrix of bond characteristics are passed through the network, and the network’s outputs are compared to the true liquidity estimates, finding the mean squared error of the predictions. Using stochastic gradient descent, with gradients found via back-propagation, the weights of the model are updated such that the mean squared error of the model is minimized. The hyper-parameters (d, w) , where d denotes the number of hidden layers in the neural network (the network’s depth) and w denotes the number of nodes in each layer (the network’s width), are optimized using k-fold cross validation to minimize the average out-of-sample mean squared error. The value of these hyper-parameters selected for the deep neural network model are reported in Table 11.

2.5.2.2 Bench-marking of LOWESS Regression against its Alternative Models

The coefficient of determination, R^2 , measures the proportion of total variance in the dependent variable that is explained by the model. When evaluated on the in-sample training data, R^2 is a representation of the explanatory power of the model. When evaluated on the out-of-sample test data, however, R^2 provides a representation of the predictive power of the model. The in-sample and out-of-sample R^2 statistics for the LOWESS regression and the alternative models are presented in Table 12. Also presented in Table 12 are the in-sample and out-of-sample root-mean-square error (RMSE) statistics for the LOWESS regression and its alternatives. While R^2 measures the explanatory and predictive power of the models, the RMSE serves as a measure of the accuracy of the model predictions. The LOWESS regression model outperforms the alternative models across all of the measures of performance by a wide margin, with an in-sample R^2 of 0.3202, out-of-sample R^2 of 0.3447, and in-sample and out-of-sample RMSEs of 1.0886 and 1.0256, respectively. The next-best performing model is the deep neural network, which itself outperforms the remaining models across all of the measures of performance. The in-sample and out-of-sample R^2 s of the deep

neural network are 0.1692 and 0.1295, respectively, and the in-sample and out-of-sample RMSEs are 1.2294 and 1.1844, respectively.

2.6 Measure of Overall Market Liquidity

Having applied the LOWESS regression matching estimator to interpolate liquidity estimates for bonds that do not trade, I then aggregate the individual liquidity measures into an estimate of the overall bond market liquidity. To perform this aggregation, for each day with a sufficient number of traded bonds¹², I take the median liquidity estimate of all the bonds in my sample that are outstanding on the date in question. This aggregation method captures the overall market liquidity while remaining robust to individual bonds with abnormal liquidity measures that may not be representative of the overall liquidity available in the market. The evolution of my measure of total market illiquidity throughout the sample is displayed in Figure 16. Using this measure of aggregate market liquidity, I can examine the relationship between overall bond market liquidity and other macroeconomic factors related to liquidity.

2.6.1 Relationships between Market Liquidity and Macroeconomic Factors

First, I investigate the relationship between the latent market liquidity of the U.S. corporate bond market and the Chicago Board Options Exchange(CBOE) Volatility Index (VIX). The VIX is a measure of the market's expectation of stock market volatility over the next 30 days, calculated as the square root of the variance strike of a hypothetical variance swap replicated by a portfolio of out-of-the-money 30-day vanilla S&P 500 options. In periods of

¹²For the LOWESS regression to be feasibly applied to a given day t , it must be the case that the following inequality holds: $qn_t \geq 5 \Rightarrow n_t \geq \frac{5}{q}$, where q is the hyper-parameter that determines the number of nearest neighbors to include in the model's local regressions as a portion of the number of bonds for which liquidity data is observed. Given the training of the LOWESS regression model, the optimal q is found to be equal to 0.0074. Thus, the LOWESS regression can only be implemented on days for which the trade-based liquidity measure is observed for at least 676 bonds. Since my sample consists of all FINRA TRACE historical bond transactions that are observed over the life of the bonds, the concurrent number of outstanding bonds in the sample grows over time. The first date in my sample for which there are a sufficient number of traded bonds is January 5, 2008. Therefore, my estimate of overall bond market activity spans the period from January 5, 2008 through March 31, 2018.

high stock market volatility, it is common for investors to de-risk by shifting out of stocks and into bonds. Given the elevated demand for bonds during this flight to safety, we would expect the cost of liquidity in the bond market to become elevated. Since my measure of latent liquidity is constructed from an estimate of the spread demanded from broker-dealers, we would therefore expect my measure of latent bond market liquidity to have a positive correlation with the VIX.

Taking my measure of latent bond market liquidity and comparing it to the daily close value of the VIX, I find that the correlation between the two measures $\rho = 0.6828$. Testing the null hypothesis $H_0 : \rho \leq 0$ against the alternative hypothesis $H_a : \rho > 0$, I find that I am able to reject the null hypothesis at the 1% confidence level. I also perform a linear regression of the VIX on my measure of latent market liquidity, the results of which are reported in Table 13, and find a significant positive relationship. This strong relationship between latent market liquidity and stock market volatility reinforces the finding in existing literature [Chordia et al. (2005), Campbell and Taksler (2003)] that stock and bond market volatility and liquidity are correlated. Chordia et al.(2005) investigate the relationship between the stock market and government bonds, and Campbell and Taksler (2003) explore the effect that firm-level equity volatility has on corporate bonds. I extend these results, establishing a link between stock market volatility and U.S. corporate bond market liquidity at the aggregate level.

Next, I investigate the relationship between the latent market liquidity and credit spreads. There is a well established “credit spread puzzle” wherein a significant portion of the spreads on corporate bonds cannot be fully explained by their default risk and other credit risk determinants [Collin-Dufresne et al. (2001), Huang and Huang (2012)]. Much research has been conducted in an attempt to explain the variation in credit spreads. Recent work suggests that a bulk of the remaining spread can be explained by illiquidity, especially when examining the credit spread of individual bonds [Longstaff et al. (2005), Edwards et al. (2007), Bao et al. (2011)]. This essay makes no effort to address the credit spread puzzle at the individual security level. Rather, I examine the relationship between aggregate market

illiquidity and average credit spreads.

To serve as a proxy for the portion of overall credit spreads unexplained by default risk¹³, I take the Intercontinental Exchange Bank of America Merrill Lynch US Corporate Master Index Option-Adjusted Spread. This index tracks the market capitalization weighted average of the option-adjusted spreads (OASs) of USD denominated investment grade corporate debt that is publicly issued in the U.S. domestic market. Since my measure of latent liquidity is constructed from an estimate of the spread demanded from broker-dealers, we would again expect my measure of latent bond market liquidity to have a positive correlation with the OAS index. Taking my measure of latent bond market liquidity and comparing it to the daily close value of the OAS index, I find that the correlation between the two measures $\rho = 0.7452$. Testing the null hypothesis $H_0 : \rho \leq 0$ against the alternative hypothesis $H_a : \rho > 0$, I find that I am able to reject the null hypothesis at the 1% confidence level. I also perform a linear regression of the OAS index on my measure of latent market liquidity, the results of which are reported in Table 13, and find a significant positive relationship between the measure of latent market liquidity and credit spreads. Based on this linear regression, I further find that my measure of latent market liquidity can explain about 55% of the variation in credit spreads. This result reinforces the Longstaff et al. (2005) finding that “nondefault components” account for approximately 50% of the spreads of investment grade corporate bonds.

2.7 Limitations

I am unable to directly observe dealers’ willingness to provide liquidity. For this reason, there are important aspects to liquidity that cannot be captured by this estimator. A dealer may, for example, be willing to purchase a bond that is not expected to trade and is difficult to resell in order to strengthen her network ties with the client who sells her the bond.

¹³The difference between the credit spread of high-yield corporate bonds and the credit spread of investment grade corporate bonds is commonly taken as a proxy for the additively separable portion of credit spreads that is explained by default risk. A corollary to the assumption underlying this proxy for the portion of spreads explained by default risk is that the credit spread of investment grade corporate bonds serves as a proxy for the portion of credit spreads unexplained by default risk.

Such a consideration on the dealer's part is dependent upon the client attempting to sell said bond, and cannot be captured by this estimator. The accuracy of this estimator on a given day therefore relies on the sample of clients trading with the dealers that day being representative of the population of bond investors. Moreover, this estimator may not capture dealers' willingness and ability to absorb large volumes. While the estimator is built upon a volume-weighted average of the effective bid-ask spread, the robustness of the measure to large trading volumes is dependent upon the observation of large trading volumes. If the average trading volume on a given day is smaller than usual, then the estimator will likely be biased, reporting more liquidity than would be available if larger volumes were traded.

Given the assumption that similar bonds have similar levels of liquidity, the accuracy of this estimator is dependent upon the observation of bonds that are similar in nature to each of the non-traded bonds. If no similar bond is traded, as would happen if the only bonds that trade on a given day belong to the most actively traded highly liquid subset of bonds, then this estimator will be biased toward the most liquid bonds on that day. Additionally, while this estimator is trained on a sample of bonds with a high proportion of seldom-traded bonds, it is trained only on those bonds for which trades are observed. Therefore, the accuracy of this estimator is un-tested on bonds that never trade. This limitation is exacerbated by the possibility that these excluded bonds are dissimilar to the bonds upon which the estimator is trained in ways other than the propensity to trade.

The above-mentioned limitations of this estimator affect the intermediate step of estimating the latent liquidity of bonds that do not trade on a given day, and primarily bias the estimates in the direction of having more available liquidity. That being said, the final aggregation into a measure of overall liquidity is performed via taking the median value of active bond-level liquidity estimates. While this aggregation technique is robust to outliers, be they due to idiosyncratic liquidity shocks or caused by biased or otherwise inaccurate estimates of latent liquidity, it does not capture higher level notions such as the skewness and kurtosis of the distribution of liquidity across the U.S. corporate bond market.

2.8 Future Research

There are several potential avenues for future research. First, the time-series behavior of the overall market liquidity can be used to investigate the impact that various market interventions have had on the liquidity of the bond market as a whole. Second, the relationships between overall bond market liquidity and other macroeconomic factors are explored herein only at a cursory level; I investigate the extent to which overall liquidity is correlated with CBOE's VIX and the ICE BofA US Corporate Master Index OAS. These relationships can be studied in greater detail, and the relationships between market liquidity and additional macroeconomic factors can be explored in order to better understand the ways in which these factors are connected to overall bond market liquidity. Third, as discussed in the previous section, the aggregation technique used in forming the measure of overall liquidity does not capture the skewness and kurtosis of the distribution of liquidity across the U.S. corporate bond market. It may prove beneficial to explore the ways in which these other distributional measures evolve over time and how they are related to other macroeconomic factors.

2.9 Conclusion

In this essay, I construct a measure of aggregate illiquidity in the U.S. corporate bond market. This illiquidity measure employs a two-step approach to estimate the latent illiquidity of bonds that do not trade from the observed illiquidity of similar bonds that do trade. For each step, I find that my approach has a high level of out-of-sample predictive validity and that my approach outperforms competing alternative models. I explore the relationships between my measure of overall market illiquidity and both stock market volatility and the portion of overall bond market credit spreads unexplained by default risk. For each of these factors, I find a strong relationship exists between it and overall bond market illiquidity. These results reinforce the findings of existing literature and extend their findings to the overall bond market illiquidity.

3 Liquidity Risk from Dealer Inventory Limits

3.1 Introduction

Title VI of the Dodd-Frank Wall Street Reform and Consumer Protection Act – commonly referred to as the “Volcker rule” – established regulations intended to curb proprietary trading on behalf of commercial banks. The method by which the regulation achieves this, however, may have unintentionally lead to a decrease in the liquidity available for affected securities. Namely, the Volcker rule attempts to address proprietary trading by placing strict inventory limits on each trading desk, which simultaneously restricts those dealers’ ability to absorb shocks to the demand for liquidity.

An unintended consequence of the Volcker Rule and Basel III regulations has been the significant reduction of dealer inventories for the majority of bond issues. As reported by Barclays Capital, “uncertainty about the implications of the limitations on proprietary trading included in the Volcker rule [led] dealers to reduce inventories” [Meli (2011)]. This dissolution of inventories leads to potentially worse market health. Duffie (2012) warns of the potential adverse effects of this regulation, including the degradation of the capacity for market makers to provide their services.

If a dealer is unable to immediately find a counter-party to take on the opposite side of a trade, then she would increase the spread she is offering on her quote to compensate for the risk of not being able to find anyone. The extent to which dealers can exhibit this behavior, however, is limited by the maximum size of their inventory. With the passage of Dodd-Frank and the Volcker rule, dealers have been given strict limits on the size of inventories they can accumulate. Therefore, the dealers are unable to absorb shocks in liquidity demand. With low inventories, dealers no longer behave as normal market makers – they instead act more like intermediaries, facilitating transactions. If a significant portion of a bond’s holders are attempting to liquidate their positions at the same time, there is a possibility that market makers, unable to hold large inventories, would be unable to provide adequate liquidity to the market.

Many institutional investors, especially pension funds, restrict themselves in their Investment Policy Statements¹⁴ to hold only so-called “investment grade” bonds. This common self-imposed restriction forces them to liquidate their position in a bond issue if a Nationally Recognized Statistical Rating Organization (NRSRO) downgrades the credit rating of the issue from an investment grade rating to a “junk” rating (i.e. any rating BB+ or lower). NRSROs being barred from leaking information about an upcoming downgrade, these downgrading events therefore represent plausibly exogenous shocks to trade flow demands of these often illiquid instruments. When such a downgrade occurs, institutional investors who face the aforementioned restrictions are forced to liquidate their positions, spawning a fire sale.

This provides a potential test bed for a natural experiment to test the impact of lower dealer inventories on the price impact of such liquidity crunches. TRACE data provide trade prices for all bond transactions. Implementing a Diff-in-Diff framework around these regulations, I can test the impact of credit rating changes on price and liquidity in high and low inventory environments. Special interest can be given to the impact of liquidity crunches by focusing on instances in which bonds are downgraded from investment grade to junk.

The remainder of this essay is organized as follows. Section 3.2 discusses the current literature concerning the impact of the Volcker rule on market liquidity. Section 3.3 describes the data used in this essay. The specification of the tests performed and the results of those tests are given in Section 3.4. These results are discussed in Section 3.5. Sections 3.6 and 3.7 assess the limitations of this essay and potential avenues for future research, and Section 3.8 concludes.

3.2 Related Literature

This essay is related to a class of working papers investigating the impacts of regulatory interventions following the 2008 financial crisis. Contradicting anecdotal reports, Trebbi and Xiao (2015) find no evidence of the deterioration of market liquidity resulting from post-crisis regulatory activity. Trebbi and Xiao (2015), however, only examines the normal

¹⁴The self-imposed restrictions of the ten largest U.S. public pension funds are presented in Table 14.

level of liquidity available in the market, and does not capture the effects of recent regulatory reform on periods of abnormal demand for liquidity as exist during fire sales. Dick-Nielsen and Rossi (2016) leverage index exclusions as events which spawn liquidity shocks, and show that the cost of immediacy has increased substantially following the financial crisis. Dick-Nielsen and Rossi (2016) provide evidence in support of Duffie (2012)'s prediction, but their data span only through 2013 and can therefore not make a statement about the enforcement of the Volcker rule specifically.

My work is most closely related to Bao, O'Hara, and Zhou (2016), which not only investigates the effects of the implementation of the Volcker Rule in times of stress, but does so using bond downgrades as events that cause liquidity shocks. Their paper finds that there has been an increase in illiquidity related to stress events. Moreover, they disentangle the impacts stemming from the Volcker rule and those from Basel III requirements, and show that the observed degradation is a result of the Volcker rule and not Basel III. These results, however, are driven by comparing the liquidity of bonds which were downgraded from investment grade to junk ratings with bonds that held junk ratings throughout. This implementation stands to overestimate the impact of the Volcker rule on stress events by failing to account for bond rating changes that do not move bonds from investment grade to junk. Such rating changes would still impact the balance of supply and demand as the market incorporates the new rating information into the price of the bond, and therefore will affect the liquidity observed during this period. In this way, Bao et al. (2016)'s measurement of the effects includes both those which result from forced fire sales and those which result from the market pricing in new information.

3.3 Data

From the Mergent Fixed Income Securities Database (FISD), I obtain the credit rating history of U.S. Corporate bonds for which there was at least one credit rating change between January 1, 2010 and March 31, 2016. This database contains the complete ratings history from each of the three largest NRSROs: Moody's, Fitch, and Standard & Poor's. This

sample consists of 13,338 rating changes. Of these rating changes, 2,398 occurred before the Volcker rule was announced, 3,699 occurred after the Volcker rule was announced but before it was enforced, and 7,291 occurred after the Volcker rule went into effect. A complete summary of the rating changes included in this sample are reported in Table 15.

For each rating change, I identify rating upgrades and downgrades as those rating changes that cross the threshold between investment grade and junk. Upgrades take a rating from BB+ or worse to BBB- or better, and downgrades take a rating from BBB- or better to BB+ or worse. For each change, I also find the amount by which the rating changed and the direction of the change, with a difference of a full rating (from BBB to BB or from BBB to A) being equal to difference of 1 and a downward rating change being represented as a negative change.

For bonds that are covered by more than one NRSRO, for each rating change, I also find the most recent rating from the other rating agencies. Using the ratings from other NRSROs, I identify what I denote “surprise upgrades” and “surprise downgrades.” For the purposes of this essay, a surprise downgrade is a rating change from investment grade to junk for a bond that is either a) not rated by any other NRSROs, or b) rated BBB- or better by all other NRSROs that publish ratings for the bond. Similarly, a surprise upgrade is a rating change from junk to investment grade for a bond that is either a) not rated by any other NRSROs, or b) rated BB+ or worse by all other NRSROs that cover the bond. I also identify what I denote “definitive upgrades” and “definitive downgrades.” Here, a definitive downgrade is a rating change from investment grade to junk for a bond that is either a) not rated by any other NRSROs, or b) rated as BB+ or worse by at least one other NRSRO. A definitive upgrade is defined similarly.

Additionally, for each rating change, I identify what I denote “near upgrades” and “near downgrades.” I define near downgrades to be rating changes that take a rating from BBB or better to BBB- or worse, and define near upgrades to be rating changes that take a rating from BBB- or worse to BBB or better. I also identify what I denote “large upgrades” and “large downgrades.” Here, a large upgrade is a rating change from junk to a rating of BBB

or better, and a large downgrade is a rating change from a rating of BBB or better to junk.

Using FINRA’s Trade Reporting and Compliance Engine (TRACE) enhanced historical data from November 1, 2009 through May 31, 2016, for each credit rating change, I gather all trades that are reported to TRACE for the bond whose rating changed in the four months surrounding the change (i.e. the two months preceding and the two months following the change). I then find the volume-weighted average price for the pre-change and post-change periods. I find the volume-weighted average price over four different time windows: two months, one month, two weeks, and one week. Using these volume-weighted average prices, I find the change in price around each rating change for each of the four time windows.

3.4 Test Design and Results

In order to determine the price impact of credit rating downgrades, I first perform a regression on the following linear relationship:

$$(R1) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Up_{i,t} + \beta_2 \cdot Down_{i,t} + \varepsilon_{i,t},$$

where $\Delta P_{i,t}$ is the change in volume-weighted average price of bond i from before the credit rating change to after the rating change, $\Delta R_{i,t}$ represents the size and direction of the rating change for bond i on date t (in which a negative value corresponds to a downward change), $Up_{i,t}$ is a dummy variable indicating that the rating change upgraded the bond issue from junk grade to investment grade, $Down_{i,t}$ is a dummy variable indicating that the rating change downgraded the bond issue from investment grade to junk grade, and $\varepsilon_{i,t}$ is an error term following a normal distribution. It is possible that effects of credit rating downgrades are short-lived. Therefore, I perform this regression taking the pre-rating and post-rating volume-weighted average price using a two-month window, a one-month window, a two-week window, and a one-week window. Given that credit ratings are correlated with credit spread, with a worse credit rating indicating a higher probability of default and thus a higher credit spread, I expect β_0 to be positive. Credit upgrades from junk to investment grade open up

the possibility for institutional investors to hold the bond, and represent both a decrease in default risk as well as an increase in the marketability of the bond. Thus, I expect β_1 to be positive. Credit downgrades from investment grade to junk, on the other hand, represent both an increase in default risk and a decrease in the marketability of the bond, as institutional investors are no longer able to hold the bond in question. Downgrades are not only expected to cause a sell-off of the bond by institutional investors, but also have the potential of causing a fire sale as these investors liquidate their positions. Therefore, I expect β_2 to be negative. Additionally, given the fact that fire sales are a possible occurrence given a credit downgrade but are not expected to occur after an upgrade, I expect the magnitude of β_2 to be greater than that of β_1 .

The results of this regression are reported in Table 16. For each of the four time windows, the predictions regarding the direction of β_0 , β_1 , and β_2 are all observed, and are all significant at the 1% level. In order to test whether the magnitude of β_2 is greater than that of β_1 , I perform Welch's unequal variances t-test on the following hypotheses:

$$H_0 : |\beta_2| \leq |\beta_1|, \quad H_a : |\beta_2| > |\beta_1|.$$

The results of this test are reported in Table 18. I find that, over all four time windows, I am able to reject the null hypothesis at the 1% significance level. Thus, all four predictions about the price impact of credit rating changes, upgrades, and downgrades are confirmed at the 1% significance level.

In order to determine whether the Volcker Rule has had deleterious effects on the ability of the bond market to absorb liquidity shocks, I perform a regression on the following linear relationship:

$$\begin{aligned} (R2) : \Delta P_{i,t} = & \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Up_{i,t} + \beta_2 \cdot Down_{i,t} \\ & + \beta_3 \cdot Up_{i,t} \times During_t + \beta_4 \cdot Down_{i,t} \times During_t \\ & + \beta_5 \cdot Up_{i,t} \times After_t + \beta_6 \cdot Down_{i,t} \times After_t + \varepsilon_{i,t}, \end{aligned}$$

where $During_t$ is a dummy variable indicating that the rating change occurred after the Volcker Rule regulations were announced but before they were enforced¹⁵, $After_t$ is a dummy variable indicating that the change occurred after the regulations were enforced, and the other factors are defined as in regression (R1). Thus, I employ a Diff-in-Diff approach to measure the effects that the regulations have had on the price impact of credit rating downgrades, with β_3 and β_4 representing the treatment effect of the announcement of the Volcker Rule, and with β_5 and β_6 representing the treatment effect of the enforcement of the regulations. If the Volcker Rule has had a deleterious effect on the ability of the bond market to absorb liquidity shocks, then I expect β_6 to be negative. Moreover, if dealers' reductions in inventory after the announcement of the Volcker Rule had a deleterious effect on the market's ability to absorb liquidity shocks, then I expect β_4 to be negative.

The results of this regression are reported in Table 17. While β_6 is negative for the one month, two week, and one week windows, and β_4 is negative for the one month and two week windows, neither of the two predictions about the effects of the Volcker rule hold with any level of significance over any time window. Thus, using all credit downgrades as an indicator, there is not evidence that the implementation of the Volcker rule has harmed the market's ability to absorb liquidity shocks.

“Surprise” Downgrades There is significant overlap in ratings coverage between Moody's, Fitch, and Standard & Poor's, and thus significant overlap in the bonds that are downgraded. For a bond that has already been downgraded by another rating agency, the additional informational value of a second or third downgrade is limited. It is possible, therefore, that by performing regressions (R1) and (R2) using all credit downgrades as an indicator, I am diluting the effects of the Volcker rule, and thereby causing the negative results

¹⁵Starting as a proposal in a February 22, 2010 Wall Street Journal op-ed endorsed by several past Secretaries of the Treasury, Dodd-Frank passed the Senate vote on May 20, 2010. Initial Volcker rule regulations were proposed on October 11, 2011, and final rules were announced on July 21, 2012. After issues implementing the new regulations, the ultimate form of the regulations began being enforced on January 14, 2014. Note that, as reported by Barclays Capital, dealers had already started to reduce their inventories to align with the new regulatory environment before it started being enforced. As such, it is important to measure the impacts of the various treatment stages.

that are observed. In order to control for the attenuating effects that pre-existing ratings from other credit rating agencies would have on the price impact of credit rating downgrades, I re-perform the analysis above using “surprise” credit rating downgrades. Thus, I perform regressions on the following linear relationships:

$$(R3) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t}^s + \beta_1 \cdot Up_{i,t}^s + \beta_2 \cdot Down_{i,t}^s + \varepsilon_{i,t}$$

$$(R4) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t}^s + \beta_1 \cdot Up_{i,t}^s + \beta_2 \cdot Down_{i,t}^s \\ + \beta_3 \cdot Up_{i,t}^s \times During_t + \beta_4 \cdot Down_{i,t}^s \times During_t \\ + \beta_5 \cdot Up_{i,t}^s \times After_t + \beta_6 \cdot Down_{i,t}^s \times After_t + \varepsilon_{i,t},$$

where $\Delta R_{i,t}^s$ is the difference in ratings between the new rating and the most recent rating from any agency covering bond i , $Up_{i,t}^s$ is a dummy variable indicating that the rating change is the first rating among any active ratings that takes the bond issue from junk grade to investment grade (i.e. “surprise upgrades”), and $Down_{i,t}^s$ is a dummy variable indicating that the rating change represents the first rating among any active ratings that takes bond i from investment grade to junk grade (i.e. “surprise downgrades”). The other factors are defined as in regressions (R1) and (R2). Again, I perform these regressions taking volume-weighted average prices using the four different time windows. For regression (R3), I expect the same relationships to hold as in regression (R1). Namely, I expect that $\beta_0 > 0$, $\beta_1 > 0$, $\beta_2 < 0$, and $|\beta_2| > |\beta_1|$. For regression (R4), assuming “surprise downgrades” indeed carry more weight than other downgrades, if the announcement and enforcement of the Volcker rule harmed market liquidity in stress-related events, then I expect β_4 and β_6 to be negative.

The results of regression (R3) are reported in Table 19. For all time windows except two months, the predictions regarding the direction of β_0 , β_1 , and β_2 are all observed, and are all significant at the 1% level. For the two month window, the predictions regarding the directions of β_0 , β_1 , and β_2 are all observed, with β_0 being significant at the 1% level, β_1 being significant at the 10% level, and β_2 not having any significance. I re-perform the unequal variances t-test to test the prediction that $|\beta_2| > |\beta_1|$, the results of which are

reported in Table 21. For the two week and one week time windows, I am able to reject the null hypothesis at the 1% level, but I am unable to reject the null hypothesis for the two month and one month time windows.

The results of regression (R4) are reported in Table 20. While β_4 and β_6 are negative over all four time windows, the prediction regarding β_4 only holds at the 1% significance level for the one month time window and at the 5% significance level for the two week time window, and the prediction regarding β_6 only holds at the 10% significance level for the one week time window. Thus, using “surprise downgrades” as an indicator, there is only weak evidence that the Volcker rule has degraded the market’s ability to absorb liquidity shocks.

“Definitive” Downgrades While “surprise downgrades” may provide more information about credit risk than other downgrades, they do not necessarily force institutional investors to liquidate their holdings. Among the ten largest U.S. public pension funds, in the case of split-ratings (i.e. when NRSROs assign the same bond different credit ratings) only one fund restricts itself to investing in bonds for which all ratings are investment grade; three of the ten funds restrict themselves to investing in bonds for which at least one rating is investment grade. A complete listing of the restrictions of the ten largest U.S. public pension funds is reported in Table 14. Given the differing restrictions prescribed by various funds’ investment policies, it is possible that “surprise downgrades” do not carry the risk of sparking fire sales. For this reason, I re-perform the analysis using “definitive” credit rating downgrades. Here, “definitive upgrades” are the second or third rating change that takes the rating from junk to investment grade, and “definitive downgrades” are the second or third rating change that takes the rating from investment grade to junk status. If no other rating agencies cover the bond in question, then any upgrade or downgrade is taken to be “definitive.”

Thus, I perform regressions on the following linear relationships:

$$(R5) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Up_{i,t}^d + \beta_2 \cdot Down_{i,t}^d + \varepsilon_{i,t}$$

$$\begin{aligned}
(R6) : \Delta P_{i,t} = & \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Up_{i,t}^d + \beta_2 \cdot Down_{i,t}^d \\
& + \beta_3 \cdot Up_{i,t}^d \times During_t + \beta_4 \cdot Down_{i,t}^d \times During_t \\
& + \beta_5 \cdot Up_{i,t}^d \times After_t + \beta_6 \cdot Down_{i,t}^d \times After_t + \varepsilon_{i,t},
\end{aligned}$$

where $Up_{i,t}^d$ is a dummy variable indicating that the rating change is a “definitive upgrade,” $Down_{i,t}^d$ is a dummy variable indicating that the rating change is a “definitive downgrade,” and the other factors are defined as in regressions (R1) and (R2). Again, I perform these regressions taking volume-weighted average prices using the four different time windows. For regression (R5), I expect the same relationships to hold as in regression (R1); I expect that $\beta_0 > 0$, $\beta_1 > 0$, $\beta_2 < 0$, and $|\beta_2| > |\beta_1|$. For regression (R6), assuming “definitive downgrades” indeed carry more weight than “surprise downgrades,” and if the announcement and enforcement of the Volcker rule harmed market liquidity in stress-related events, then I expect β_4 and β_6 to be negative.

The results of regression (R5) are reported in Table 22. For all time windows, the predictions regarding the direction of β_0 , β_1 , and β_2 are all observed. For β_0 , the prediction is significant at the 1% level over all time periods. For β_1 , the prediction is only significant at the 5% level over the one week time window, and significant at the 10% level over the two month and two week time windows. For β_2 , the prediction is significant at the 1% level over the two month, one month, and two week time levels. Again, I re-perform the unequal variances t-test to test the prediction that $|\beta_2| > |\beta_1|$, the results of which are reported in Table 24. I find that, over the two month, one month, and two week time windows, I am able to reject the null hypothesis at the 1% significance level.

The results of regression (R6) are reported in Table 23. β_4 is positive for all four time windows, and β_6 is only negative for the one month and two week time windows. Moreover, the prediction for β_6 does not hold at any significance level over the two time windows for which $\beta_6 < 0$. While β_4 is positive over all four time windows, and β_6 is positive over the two month and one week time windows, the prediction regarding β_4 only holds at the 10% significance level for the two month time window and at the 5% significance level for one two week time window, and the prediction regarding β_6 only holds at the 10% significance level

for the one week time window. Thus, using “definitive downgrades” as an indicator, there is no evidence that the Volcker rule has degraded the market’s ability to absorb liquidity shocks.

“Near” Downgrades The analysis heretofore offers only limited support of the hypothesis that the Volcker rule has degraded the bond market’s ability to absorb liquidity shocks. One possible explanation for this lack of observed impact is that institutional investors have become more risk averse in their bond holdings. It is reasonable to suspect that institutional investors are sophisticated. Realizing the implications of the inventory reduction as a result of the Volcker rule, those investors who face the credit rating constraints would choose to hold less debt that has a risk of losing its investment grade rating – removing bonds from their portfolio before they are downgraded to junk ratings in order to avoid potential fire sales. Such behavior would reduce the volume of debt facing forced liquidation, therefore mitigating the impact of a downgrade.

In order to test this possibility, I re-perform the analysis using “near downgrades.” Here, “near downgrades” are rating changes that take a rating from BBB or better to BBB- or worse, and “near upgrades” are rating changes that take a rating from BBB- or worse to BBB or better. I perform regressions on the following linear relationships:

$$(R7) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Up_{i,t}^{BBB} + \beta_2 \cdot Down_{i,t}^{BBB-} + \varepsilon_{i,t}$$

$$(R8) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Up_{i,t}^{BBB} + \beta_2 \cdot Down_{i,t}^{BBB-} \\ + \beta_3 \cdot Up_{i,t}^{BBB} \times During_t + \beta_4 \cdot Down_{i,t}^{BBB-} \times During_t \\ + \beta_5 \cdot Up_{i,t}^{BBB} \times After_t + \beta_6 \cdot Down_{i,t}^{BBB-} \times After_t + \varepsilon_{i,t},$$

where $Up_{i,t}^{BBB}$ is a dummy variable indicating that the rating change is a “near upgrade,” $Down_{i,t}^{BBB-}$ is a dummy variable indicating that the rating change is a “near downgrade,” and the other factors are defined as in regressions (R1) and (R2). For regression (R7), I expect the same directional relationships to hold as in regression (R1); I expect that $\beta_0 > 0$, $\beta_1 > 0$,

and $\beta_2 < 0$. However, given that the entire sample period is included in this regression, and the hypothesized aversion to holding BBB- rated bonds is only a result of the Volcker rule, I do not expect the magnitude of β_2 to be greater than that of β_1 . For regression (R8), if the announcement and enforcement of the Volcker rule has caused investors to become more risk averse to holding BBB- rated bonds, then I expect β_4 and β_6 to be negative.

The results of regression (R7) are reported in Table 25. For all time windows, the predictions regarding the direction of β_0 , β_1 , and β_2 are all observed. For β_0 , the prediction is significant at the 1% level over all time periods. For β_1 , the prediction is only significant at the 5% level over the two month time window, and significant at the 10% level over the two week time window. For β_2 , the prediction is only significant at the 10% level over the two month, one month, and two week time levels. While I do not maintain the expectation that the magnitude of β_2 to be greater than that of β_1 , I re-perform the unequal variances t-test to test the prediction that $|\beta_2| > |\beta_1|$, the results of which are reported in Table 27. I find that I am only able to reject the null hypothesis for the one month time window at the 5% significance level.

The results of regression (R8) are reported in Table 26. β_4 is negative for the two month, one month, and two week time windows, and the prediction for β_4 holds at the 10% significance level for the two month and one month time windows. β_6 is negative over all four time windows, and the prediction holds at the 1% significance level for the two month and one month time windows, at the 5% significance level for the two week time window, and at the 10% level for the one week time window. Thus, there is strong evidence that the enforcement of the Volcker rule has caused investors to become more averse to holding BBB-rated bonds.

The hypothesis that the Volcker rule has increased institutional investors' aversion to holding BBB- rated bonds is based on the supposition that this aversion is due to said investors seeking to avoid fire sales. After a bond is rated BBB-, these investors will slowly sell off their holdings of the bond, causing a slow decrease in the price of the bond. Therefore, I expect the price impact of "near downgrades" to increase as the time window increases. To

test this, I perform three separate unequal variance t-tests on the following sets of hypotheses:

$$\begin{aligned}
 H_0 : |\beta_6^{2M}| &\leq |\beta_6^{1M}|, & H_a : |\beta_6^{2M}| &> |\beta_6^{1M}| \\
 H_0 : |\beta_6^{1M}| &\leq |\beta_6^{2W}|, & H_a : |\beta_6^{1M}| &> |\beta_6^{2W}| \\
 H_0 : |\beta_6^{2W}| &\leq |\beta_6^{1W}|, & H_a : |\beta_6^{2W}| &> |\beta_6^{1W}|,
 \end{aligned}$$

where β_6^{2M} is the impact of the Volcker rule on the price impact of “near downgrades” over a two month window, β_6^{1M} is the impact of the Volcker rule over a one month window, β_6^{2W} is the impact over a two week window, and β_6^{1W} is the impact over a one week window.

The results of these tests are reported in Table 28. I am able to reject all three null hypotheses at the 1% confidence level. This provides strong evidence that, after a bond is rated BBB-, institutional investors slowly sell off their holdings of said bond. Rather than immediately affecting the price of the bond, the price impact of such a rating change slowly increases over time. Thus, investors are able to liquidate their holdings of such bonds without sparking a fire sale.

“Large” Downgrades Given the strong evidence that institutional investors have become averse to holding BBB- rated bonds, in order to isolate potential forced liquidation events, I re-perform the previous analysis using “large downgrades.” Here, “large downgrades” are rating changes that take a rating from BBB or better to BB+ or worse, and “large upgrades” are rating changes that take a rating from BB+ or worse to BBB or better. I also restrict the sample to rating changes for bonds for which only one NRSRO issued a rating. In this way, I remove the downgrades of bonds for which investors have already had the opportunity to divest their holdings when the bond was rated BBB-, and I remove the possibility that investors are not forced to liquidate their holdings given the downgrade. Thus, the downgrades identified in this analysis have the greatest likelihood of causing a fire sale.

Over the sample period, there are no large upgrades for bonds covered by only one rating agency, and there are only nine large downgrades for bonds covered by one agency. Moreover,

only three such downgrades occur after the enforcement of the Volcker rule¹⁶. Nonetheless, I perform regressions on the following linear relationship:

$$(R9) : \Delta P_{i,t} = \beta_0 \cdot \Delta R_{i,t} + \beta_1 \cdot Down_{i,t}^x + \beta_2 \cdot Down_{i,t}^x \times After_t + \varepsilon_{i,t},$$

where $Down_{i,t}^x$ is a dummy variable indicating that the rating change is a “large downgrade,” and the other factors are defined as in regressions (R1) and (R2). Again, I perform this regression taking volume-weighted average prices using the four different time windows. Following previous logic regarding the impact of credit ratings and downgrades on bond prices, I expect β_0 to be positive and β_1 to be negative. If the enforcement of the Volcker rule harmed market liquidity in stress-related events, then I expect β_2 to be negative.

The results of regression (R9) are reported in Table 29. β_0 is positive for all time windows, is significant at the 1% level for the one week window, and is significant at the 5% level for the one month and two week windows. β_1 is negative and significant at the 1% level for the two month and one month windows, but is positive for the two week and one week windows. β_2 is negative for the two week and one week time windows. β_2 is significant at the 1% level for the two week time window, but it is not significant at any confidence level for the one week time window. Therefore, using “large downgrades” as an indicator for potential forced liquidation events, there is little evidence that the Volcker rule has degraded the market’s ability to absorb liquidity shocks.

3.5 Discussion

As discussed in the previous section, there is little to no evidence that the Volcker rule has deteriorated the bond market’s ability to absorb liquidity shocks. This result contradicts Bao et al. (2016)’s finding that there has been an increase in illiquidity related to stress events as a result of the Volcker rule. It is plausible that the test design employed by Bao et al. (2016) overestimated the impact of credit downgrades by failing to control for the

¹⁶It should be noted that the limited number of observations for which the “large downgrade” flag applies may lead to unreliable regression results.

impact of non-downgrade credit rating changes. I do, however, find that the Volcker rule has increased institutional investors' aversion to holding bonds with a BBB- rating. I argue that this aversion is a result of investors seeking to avoid the possibility of holding bonds for which credit downgrades might spark a fire sale. Thus, as a result of the Volcker rule, these investors are acting as if the Volcker rule has harmed the market's ability to absorb liquidity shocks.

An increased aversion to holding BBB- rated bonds has the potential to impact the real economy through distorting the issuance of new bonds at a rating near junk. As illustrated in Figure 17, the general relationship between a bond's credit rating at the time of issuance and its yield at origination remains consistent throughout the implementation of the Volcker rule. Before the Volcker rule was enforced, there is a jump in yields between BBB- and BB+ rated bonds, reflecting a jump in yields between investment grade and junk bonds. After the Volcker rule went into effect, however, there is a jump in yields between BBB and BBB- with no such jump between BBB- and BB+ rated bonds. This jump in yields is indicative of institutional investors treating bonds issued at a BBB- rating as if they were junk after the Volcker rule was enforced.

I further investigate this relationship by performing several regressions of yield at origination on credit rating, the results of which are reported in Tables 30 and 31. Table 30 reports the results of regressing the yield at time of issuance on bond ratings, using various combinations of controls. Table 31 reports the results of similar regressions, restricting the sample to bond ratings ranging from BBB to BB. I find a sharp rise in the yield of bonds issued with BBB- ratings relative to nearby ratings following the enforcement of the Volcker rule. This rise is significant at the 1% level in all but one of the regressions performed – for that regression, the rise in yield is significant at the 10% level.

This result is consistent with the distribution of bonds that are issued at each respective credit rating. In the years immediately preceding the passing of Dodd-Frank, bonds with ratings of BBB and BBB- comprised approximately 9% and 8.5% of new issues, respectively. These proportions sharply declined after Dodd-Frank passed, falling to about 5.5% and 5%

respectively. Declining again to about 3.3% and 2.3% after the regulations of the Volcker rule were proposed, new BBB and BBB- issues maintained this low level of activity until the final regulations went into effect in January of 2014, at which time they fell to 1% and 0.75% of total issues.

The steady decline of new, near-junk rated, bonds is indicative of firms electing not to issue new debt at the now higher yield. As yields increase, firms whose debt is rated near junk face a higher cost of capital if raising a new bond issue, and are therefore less likely to issue debt unless they plan to fund a riskier endeavor. In response, investors must demand even higher yields for holding this debt. This tension represents a cycle of adverse selection which translates the financial market regulations into distortions of the decisions made in the real economy.

Therefore, while I find little to no evidence that the Volcker rule has deteriorated the market's ability to absorb liquidity shocks, I argue that the regulation has had the unintended consequence of distorting the issuance of new bonds at a rating near junk, negatively impacting the real economy. I find significant evidence that the Volcker rule has increased institutional investors' aversion to holding U.S. corporate debt with a BBB- credit rating. Such an aversion, I argue, is due to the investors seeking to avoid the potential of a credit downgrade forcing the liquidation of their position and thus sparking a fire sale. Insofar as this increased aversion is caused by fears of illiquidity related to stress events as a result of the Volcker rule, then I argue that the distortions caused by the Volcker rule would be eliminated if an exception to the inventory limits established by the regulation were carved out such that said inventory limits did not apply to bonds that were downgraded from investment grade to junk for some limited period of time. Such an exception would allay investor fears that credit downgrades would lead to liquidity shocks. This would remove their increased aversion to holding BBB- rated debt, and would therefore remove the adverse selection associated with issuing new debt with a BBB- credit rating.

In this way, establishing an exception to inventory limits for downgraded bonds would remove the distortions caused by the Volcker rule. Moreover, given the limited scope of

such an exception, it would not affect the primary purpose of the Volcker rule – to restrict bond dealers from engaging in proprietary trading. While dealers would be able to absorb liquidity shocks caused by downgrades, they would not feasibly be able to construct private positions to capitalize upon the exception.

3.6 Limitations

While I find little to no evidence that the Volcker rule has had deleterious effects on the bond market’s ability to absorb liquidity shocks, this result only applies to potential liquidity shocks caused by credit rating downgrades of U.S. corporate bonds. Given the strong evidence that the Volcker rule has increased institutional investors’ aversion to holding bonds with a BBB- credit rating, and the fact that the majority of downgrades affect bonds with BBB- ratings, the likelihood that a credit downgrade will cause a liquidity shock is limited. Moreover, there are too few large downgrades (downgrades that affect bonds with ratings of BBB or better) for bonds covered by only one NRSRO to provide a reliable estimation of the impact of the Volcker rule on the illiquidity surrounding those downgrades that are most likely to cause liquidity shocks. It is still plausible, therefore, that the Volcker rule has caused an increase in illiquidity related to stress events, but that this illiquidity is not observable through the identification of stress events using credit rating downgrades.

3.7 Future Research

As previously discussed, given institutional investors’ response to the Volcker rule, credit rating downgrades may not be able to identify liquidity shocks of sufficient size to be affected by the regulatory inventory limits faced by bond dealers. Using different events to identify liquidity shocks may provide greater insight into how the Volcker rule has affected the market’s ability to absorb liquidity shocks.

In this essay, I identify a plausible pathway by which the Volcker rule has distorted the issuance of new bonds at a rating near junk, negatively impacting the real economy. I do not, however, measure the complete impact of those distortions. As previously stated, these

distortions are expected to cause bonds issued with a BBB- credit rating to become more risky. Future work can therefore measure the impact of the Volcker rule on this debt by comparing the default rate of bonds issued with a BBB- rating after the regulations went into effect with the default rate of similar bonds issued before the Volcker rule.

3.8 Conclusion

In this essay, by identifying stress events using credit rating downgrades, I find little to no evidence that the Volcker rule has harmed the bond market's ability to absorb liquidity shocks. Rather, I find strong evidence that the Volcker rule has increased institutional investors' aversion to holding BBB- rated bonds. I show that this increased aversion to BBB- debt translates into a higher yield at issuance for bonds with such a rating, and that the proportion of debt issued at this rating has decreased substantially since the enforcement of the Volcker rule. I argue that this demonstrates a distortion of the decisions made regarding the issuance of new debt, which harms the real economy. Finally, I recommend a change to the Volcker rule that I argue would eliminate these distortions by removing institutional investors' increased aversion to holding BBB- rated bonds.

A References

1. Amihud, Y., 2002, "Illiquidity and stock returns: cross-section and time-series effects," *Journal of Financial Markets*, 5.1, 31-56.
2. Back, K., 1992, "Insider Trading in Continuous Time," *Review of Financial Studies*, 5.3: 387-409.
3. Bao, Jack, Maureen O'Hara, and Alex Zhou, 2016, "The Volcker Rule and Market-Making in Times of Stress," *Finance and Economics Discussion Series*, 2016-102. Washington: Board of Governors of the Federal Reserve System, <https://doi.org/10.17016/FEDS.2016.102>.
4. Bao, J., Pan, J., Wang, J., 2011, "The illiquidity of corporate bonds," *Journal of Finance*, 66.3, 911-946.
5. Beltran, H., Grammig, J., and Menkveld, A., 2016, "Understanding Limit Order Depth: Conditioning on Trade Informativeness," Available at RePEc: <http://repec.org/esLATM04/up.23536.1081939804.pdf>
6. Biais, B., Hillion, P., and Spatt, C., 1995, "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50.5: 1655-89.
7. Bloomfield, R., O'Hara, M., and Saar, G., 2005, "The 'make or take' decision in an electronic market: Evidence on the evolution of liquidity," *Journal of Financial Economics*, 75.1: 165-199.
8. Campbell, J. and Taksler, G., 2003, "Equity Volatility and Corporate Bond Yields," *Journal of Finance*, 58.6: 2321-2350.
9. Chen, H., Cui, R., He, Z., and Milbradt, K., 2018, "Quantifying Liquidity and Default Risks of Corporate Bonds over the Business Cycle," *Review of Financial Studies*, 31.3: 852-897.

10. Chen, L., Lesmond, D., and Wei, J., 2007, "Corporate Yield Spreads and Bond Liquidity," *Journal of Finance*, 62.1: 119-149.
11. Chordia, T., Sarkar, A., and Subrahmanyam, A., 2005, "An Empirical Analysis of Stock and Bond Market Liquidity," *Review of Financial Studies*, 18.1: 85-129.
12. Chung, K., Van Ness, B., and Van Ness, R., 1999, "Limit Orders and the Bid-Ask Spread," *Journal of Financial Economics*, 53.2: 255-287.
13. Cleveland, W., 1979, "Robust Locally Weighted Regression and Smoothing Scatterplots," *Journal of the American Statistical Association*, 74.368: 829-836.
14. Cleveland, W., Devlin, S., and Grosse, E., 1988, "Regression by local fitting: Methods, properties, and computational algorithms," *Journal of Econometrics*, 37.1: 87-114.
15. Cohen, K., Maier, S., Schwartz, R., and Whitcomb, D., 1981, "Transaction Costs, Order Placement Strategy, and Existence of the Bid-Ask Spread," *Journal of Political Economy*, 89.2: 287-305.
16. Collin-Dufresne, P., Goldstein, R., and Martin, S., 2001, "The determinants of Credit Spread Changes," *Journal of Finance*, 56.6: 2177-2207.
17. Dick-Nielsen, J., P. Feldhütter, and D. Lando, 2012, "Corporate bond liquidity before and after the onset of the subprime crisis," *Journal of Financial Economics*, 103.3: 471-492.
18. Dick-Nielsen J. and M. Rossi, 2016, "The Cost of Immediacy for Corporate Bonds," Working Paper, Copenhagen Business School and Texas A&M University.
19. Duffie, D., 2012, "Market-making Under the Proposed Volcker Rule," Working Paper, Stanford University.
20. Duffie, D. and Singleton, K., 1997, "An Econometric Model of the Term Structure of Interest-Rate Swap Yields," *Journal of Finance*, 52.4: 1287-1321.

21. Easley, D., Engle, R., O'Hara, M., and Wu, L., 2008, "Time-Varying Arrival Rates of Informed and Uninformed Trades," *Journal of Financial Econometrics*, 6.2: 171-207.
22. Feldhütter, P., 2012, "The same bond at different prices: identifying search frictions and selling pressures," *Review of Financial Studies*, 25.4, 1155-1206.
23. Foucault, T., 1999, "Order Flow Composition and Trading Costs in a Dynamic Limit Order Market," *Journal of Financial Markets*, 2.2: 99-134.
24. Glorot, X., Bordes, A., and Bengio, Y., 2011, "Deep Sparse Rectifier Neural Networks. Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics," *Proceedings of Machine Learning Research* 15: 315-323.
25. Glosten, L., 1994, "Is the Electronic Limit Order Book Inevitable?" *Journal of Finance*, 49.4: 1127-1161.
26. Glosten, L., and Milgrom, P., 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14.1: 71-100
27. Goldstein, M. and Kavajecz, K., 2000, "Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the NYSE," *Journal of Financial Economics*, 56.1: 125-149.
28. Goodfellow, I., Bengio, Y., and Courville, A., 2016, *Deep Learning (Adaptive Computation and Machine Learning series)*, The MIT Press.
29. Green, R., Hollifield, B., and Schürhoff, N., 2007, "Financial Intermediation and the Costs of Trading in an Opaque Market," *Review of Financial Studies*, 20.2: 275-314.
30. Grinblatt, M., 2003, "An Analytic Solution for Interest Rate Swap Spreads," *International Review of Finance*, 2.3: 113-149.

31. Hansen, P., Heaton, J., and Yaron, A., 1996, "Finite-Sample Properties of Some Alternative GMM Estimators," *Journal of Business & Economic Statistics*, 14.3: 262-280.
32. Huang, J. and Huang, M., 2012, "How much of the Corporate-Treasury Yield Spread is Due to Credit Risk?," *Review of Asset Pricing Studies*, 2.2: 153-202.
33. Hollifield, B., Miller, R., and Sandås, P., 2004, "Empirical Analysis of Limit Order Markets," *Review of Economic Studies*, 71.4: 1027-1063.
34. Kyle, A., 1985, "Continuous Auctions and Insider Trading" *Econometrica*, 53.6: 1315-1336.
35. Kumar, P. and Seppi, D., 1994, "Information and Index Arbitrage," *The Journal of Business*, 67.4: 481-509.
36. Madhavan, A., Richardson, M., and Roomans, M., 1997, "Why Do Security Prices Change? A Transaction-Level Analysis of NYSE Stocks," *Review of Financial Studies*, 10.4: 1035-1064.
37. Meli, J., 2011 "Credit Outlook," *Barclays Capital Global Outlook*.
38. Menkveld, A., 2013, "High Frequency Trading and the New Market Makers." *Journal of Financial Markets* 16.4: 712–740.
39. O'Hara, M., 2015, "High Frequency Market Microstructure." *Journal of Financial Economics*, 116.2: 257–270.
40. Parlour, C., 1998, "Price Dynamics in Limit Order Markets," *Review of Financial Studies*, 11.4: 789-816.
41. Parlour, C. and Seppi, D., 2003, "Liquidity-Based Competition for Order Flow," *Review of Financial Studies*, 16.2: 301-343.
42. Ranaldo, A., 2004, "Order aggressiveness in limit order book markets," *Journal of Financial Markets*, 7.1: 53-74.

43. Roll, R., 1984, "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market," *Journal of Finance*, 39.4: 1127-1139.
44. Sandås, P., 2001, "Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market," *Review of Financial Studies* 14.3: 705-734.
45. Seppi, D., 1997, "Liquidity Provision with Limit Orders and a Strategic Specialist," *Review of Financial Studies*, 10.1: 103-150.
46. Trebbi, F. and K. Xiao, 2015, "Regulation and Market Liquidity," Working Paper, University of British Columbia.
47. Wen, Y., 2012, "Information Asymmetry and Order Aggressiveness in Limit Order Markets," Available at SSRN: <http://ssrn.com/abstract=2136833>
48. Yueshen, B., 2016, "Uncertain Market Making," *Finance Down Under 2016 Building on the Best from the Cellars of Finance*. Available at SSRN: <http://ssrn.com/abstract=2667837>

B Tables

Table 1: Summary Data Statistics

Ticker	Market Orders	In-Sample		Total Included	(Excluded) Flickering Orders
		Limit Order Placement/Revision	Cancellation		
AAL	186,680	932,327	1,065,648	2,184,655	1,565,899
AAPL	684,584	5,164,812	5,765,407	11,614,803	7,790,732
BAC	106,181	1,009,475	1,292,938	2,408,594	1,964,235
CAT	115,550	757,949	1,088,388	1,961,887	1,754,239
CSCO	136,260	1,515,061	2,000,254	3,651,575	2,105,962
FB	345,725	2,338,654	3,117,182	5,801,561	4,930,529
GE	73,753	1,024,688	1,410,278	2,508,719	2,027,056
GOOG	132,733	790,110	921,185	1,844,028	1,486,236
GS	87,027	354,092	598,623	1,039,742	1,254,857
MSFT	235,526	2,194,676	2,912,029	5,342,231	4,313,775

Ticker	Market Orders	Out-of-Sample		Total Included	(Excluded) Flickering Orders
		Limit Order Placement/Revision	Cancellation		
AAL	78,102	470,566	518,332	1,067,000	669,452
AAPL	315,161	2,087,883	2,154,991	4,558,035	2,556,940
BAC	30,226	358,987	481,305	870,518	641,750
CAT	52,989	324,415	345,636	723,040	538,176
CSCO	52,296	587,730	760,319	1,400,345	737,685
FB	146,145	823,473	1,141,308	2,110,926	1,488,826
GE	28,534	354,773	455,091	838,398	668,333
GOOG	52,050	293,429	292,042	637,521	501,612
GS	23,748	97,118	172,491	293,357	396,330
MSFT	138,760	968,904	1,005,233	2,112,897	1,644,862

This table reports the number of each type of order observed for the sample of stocks included in the training sample and outside the training sample, as well as the number of excluded “flickering orders.” The counts of limit order placements/revisions and limit order cancellations are restricted to those orders within five ticks from the bid-ask spread that are not identified as “flickering orders.”

Table 2: Estimated Model Parameters

Ticker	Price Process			Market Order Distribution					Market Maker Costs			
	ψ	μ	α	ξ	λ_1	λ_2	ϕ_1	ϕ_2	γ_0	γ_1	ζ	η
AAL	1.4793	$-5.6713 \cdot 10^{-5}$	$2.1490 \cdot 10^{-5}$	0.5616	0.5140	0.3702	0.8099	0.6484	-0.1895	0.2052	0	0
AAPL	0.2848	$-2.8041 \cdot 10^{-5}$	$1.7223 \cdot 10^{-5}$	0.4902	0.8056	0.5491	0.7821	0.6379	-0.1069	0.1185	0	0.0004
BAC	1.6302	$-2.3653 \cdot 10^{-5}$	$1.0093 \cdot 10^{-6}$	0.5205	0.9275	0.1630	0.9223	0.1433	-0.2169	0.0786	0	0.0041
CAT	1.9961	$1.5707 \cdot 10^{-5}$	$1.3807 \cdot 10^{-4}$	0.4768	0.6665	1.0166	0.6356	0.7449	0.0102	0.0296	0	0
CSCO	1.2726	$2.3356 \cdot 10^{-6}$	$1.9669 \cdot 10^{-6}$	0.4801	0.7707	0.3045	0.8880	0.2665	-0.0057	0.1004	0	0.0069
FB	0.5428	$3.8032 \cdot 10^{-6}$	$6.4512 \cdot 10^{-6}$	0.4971	0.8203	0.6448	0.8112	0.6281	-0.0379	0.1373	0.3173	0.0583
GE	1.8938	$5.7583 \cdot 10^{-6}$	$3.4073 \cdot 10^{-6}$	0.4346	0.9503	0.3332	0.9105	0.3024	0.1680	0.0637	0.1161	0.1722
GOOG	2.2770	$-2.8178 \cdot 10^{-4}$	$1.0793 \cdot 10^{-3}$	0.5164	0.1495	0.9085	0.1071	0.4721	-0.9223	1.5247	0	0
GS	3.5921	$3.9197 \cdot 10^{-6}$	$5.3170 \cdot 10^{-4}$	0.4585	0.5008	0.9996	0.5232	0.7331	0.5908	0.2779	0.0021	0.4604
MSFT	0.6796	$-4.5100 \cdot 10^{-6}$	$7.3331 \cdot 10^{-6}$	0.4804	0.8788	0.4847	0.8682	0.4235	0.3629	0.0924	0	0.3818

This table reports the estimates of model parameters found using the Continuous-Updating Generalized Method of Moments (CUGMM) regression described in Section 1.4.2 on the training sample (January 2, 2015 through January 23, 2015). ψ represents the average time, in seconds, between market orders. μ represents the instantaneous price drift of the stock, and α describes the average price impact of a market order. Given the parameterization of the distribution of the size of market orders, we have the following parameters. ξ is the unconditional probability that a market order is a sell order. λ_1 is the probability that the size of a market buy order is greater than or equal to a full lot (100 shares). $100 \cdot \lambda_2$ is the average size of a market buy order, conditional on being greater than or equal to a full lot. ϕ_1 is the probability that the size of a market sell order is greater than or equal to a full lot. $100 \cdot \phi_2$ is the average size of a market sell order, conditional on being greater than or equal to a full lot. Finally γ_0 and γ_1 are the fixed and variable costs, respectively, of either placing a new limit order or continuing to maintain an existing order. γ_0 can be interpreted as the on-going cost associated with monitoring the order. ζ and η can be interpreted as the average opportunity costs associated with placing a new limit order and continuing to maintain an existing order, respectively.

Table 3: Distribution of Current Market Order Flow Regressed on Recent ΔRLP^S

In-Sample Results				
Ticker	Volatility	Ave. Order Size	Max. Order Size	Trade Volume
AAL	$8.4711 \cdot 10^{-4***}$	$2.1534 \cdot 10^{-4***}$	$1.1915 \cdot 10^{-2***}$	$1.0500 \cdot 10^{-3***}$
AAPL	$7.4758 \cdot 10^{-4***}$	$1.7061 \cdot 10^{-4***}$	$1.4732 \cdot 10^{-2***}$	$1.0668 \cdot 10^{-3***}$
BAC	$8.4530 \cdot 10^{-3***}$	$2.1652 \cdot 10^{-3**}$	$5.9497 \cdot 10^{-2***}$	$1.3973 \cdot 10^{-3***}$
CAT	$8.7445 \cdot 10^{-4}$	$4.1592 \cdot 10^{-4}$	$6.7682 \cdot 10^{-4}$	$1.6605 \cdot 10^{-3***}$
CSCO	$4.3763 \cdot 10^{-4*}$	$1.0128 \cdot 10^{-4}$	$5.2339 \cdot 10^{-3***}$	$3.9563 \cdot 10^{-4***}$
FB	$4.5392 \cdot 10^{-4***}$	$9.2799 \cdot 10^{-5***}$	$6.4228 \cdot 10^{-3***}$	$6.5742 \cdot 10^{-4***}$
GE	$9.0049 \cdot 10^{-3***}$	$2.2969 \cdot 10^{-3**}$	$6.2415 \cdot 10^{-2***}$	$1.5568 \cdot 10^{-3***}$
GOOG	$7.9925 \cdot 10^{-4***}$	$2.2333 \cdot 10^{-4*}$	$6.2818 \cdot 10^{-3***}$	$3.7201 \cdot 10^{-3***}$
GS	$-1.3856 \cdot 10^{-3**}$	$-5.7422 \cdot 10^{-4}$	$-5.2336 \cdot 10^{-3**}$	$4.6875 \cdot 10^{-3***}$
MSFT	$2.9698 \cdot 10^{-3***}$	$8.8947 \cdot 10^{-4***}$	$3.4825 \cdot 10^{-2***}$	$3.6170 \cdot 10^{-3***}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Out-of-Sample Results				
Ticker	Volatility	Ave. Order Size	Max. Order Size	Trade Volume
AAL	$4.6350 \cdot 10^{-4**}$	$1.5346 \cdot 10^{-4***}$	$5.2002 \cdot 10^{-3***}$	$2.1125 \cdot 10^{-4***}$
AAPL	$1.0089 \cdot 10^{-3**}$	$2.0207 \cdot 10^{-4***}$	$1.9633 \cdot 10^{-2***}$	$1.9125 \cdot 10^{-4}$
BAC	$1.1426 \cdot 10^{-2***}$	$3.3587 \cdot 10^{-3**}$	$6.4570 \cdot 10^{-2***}$	$1.4440 \cdot 10^{-3***}$
CAT	$2.2417 \cdot 10^{-5}$	$1.2508 \cdot 10^{-5}$	$-7.5158 \cdot 10^{-4}$	$2.0500 \cdot 10^{-3***}$
CSCO	$2.5790 \cdot 10^{-3***}$	$5.6359 \cdot 10^{-4*}$	$2.9411 \cdot 10^{-2***}$	$7.0958 \cdot 10^{-4***}$
FB	$1.1035 \cdot 10^{-4}$	$2.8110 \cdot 10^{-5}$	$1.3321 \cdot 10^{-3}$	$2.8065 \cdot 10^{-4***}$
GE	$8.9435 \cdot 10^{-3***}$	$2.5957 \cdot 10^{-3*}$	$5.8001 \cdot 10^{-2***}$	$1.9859 \cdot 10^{-3***}$
GOOG	$8.0878 \cdot 10^{-5}$	$2.1358 \cdot 10^{-5}$	$1.2827 \cdot 10^{-3}$	$4.0558 \cdot 10^{-04**}$
GS	$1.9023 \cdot 10^{-3**}$	$1.2401 \cdot 10^{-3**}$	$9.2468 \cdot 10^{-3***}$	$3.3417 \cdot 10^{-3***}$
MSFT	$1.4264 \cdot 10^{-3***}$	$2.6448 \cdot 10^{-4*}$	$2.4008 \cdot 10^{-2***}$	$3.0772 \cdot 10^{-7}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the in-sample and out-of-sample results of regressions (R1) through (R4). These regressions explore the predictive power that changes to the symmetric measure of relative liquidity provisioning, RLP^S , have on changes to various factors that describe the current distribution of market orders. Four factors describing the distribution of market orders over one minute intervals are investigated: the distributional volatility of market order flow (R1), the average size of market orders (R2), the maximum order size (R3), and the total volume of market orders (R4). For each factor's regression, a negative coefficient implies that market makers have knowledge about future market order flow and are provisioning liquidity accordingly.

Table 4: Price Impact Regressed on Current RLP^S and Market Order Flow

Ticker	In-Sample Results					
	Intercept	RLP^S	Volatility	Ave. Order Size	Max. Order Size	Trade Volume
AAL	$5.4405 \cdot 10^{-4***}$	$-8.3851 \cdot 10^{-9***}$	$-1.6523 \cdot 10^{-7}$	$-4.7985 \cdot 10^{-8}$	$3.4993 \cdot 10^{-8***}$	$1.7331 \cdot 10^{-6***}$
AAPL	$3.6017 \cdot 10^{-4***}$	$-4.5900 \cdot 10^{-10***}$	$-5.4088 \cdot 10^{-8}$	$-4.9188 \cdot 10^{-7***}$	$1.4963 \cdot 10^{-8***}$	$5.2930 \cdot 10^{-7***}$
BAC	$2.1343 \cdot 10^{-4***}$	$-2.8346 \cdot 10^{-9***}$	$8.8751 \cdot 10^{-9}$	$6.0112 \cdot 10^{-8***}$	$-1.5833 \cdot 10^{-9}$	$2.2571 \cdot 10^{-6***}$
CAT	$3.0178 \cdot 10^{-4***}$	$-4.5798 \cdot 10^{-8***}$	$-1.3816 \cdot 10^{-6***}$	$1.0255 \cdot 10^{-6***}$	$1.8605 \cdot 10^{-7***}$	$1.2210 \cdot 10^{-6***}$
CSCO	$1.9893 \cdot 10^{-4***}$	$-1.3112 \cdot 10^{-9***}$	$-6.6188 \cdot 10^{-8***}$	$1.0542 \cdot 10^{-7***}$	$3.1870 \cdot 10^{-9}$	$2.2090 \cdot 10^{-6***}$
FB	$2.7769 \cdot 10^{-4***}$	$-1.3554 \cdot 10^{-9***}$	$5.8677 \cdot 10^{-8}$	$-1.7672 \cdot 10^{-7}$	$-2.7302 \cdot 10^{-9}$	$9.4310 \cdot 10^{-7***}$
GE	$1.5792 \cdot 10^{-4***}$	$-8.2677 \cdot 10^{-9***}$	$1.1087 \cdot 10^{-8}$	$8.5654 \cdot 10^{-8***}$	$-9.2962 \cdot 10^{-9**}$	$2.0518 \cdot 10^{-6***}$
GOOG	$2.4226 \cdot 10^{-4***}$	$-2.3620 \cdot 10^{-8***}$	$-6.5172 \cdot 10^{-7***}$	$3.6400 \cdot 10^{-7}$	$1.6783 \cdot 10^{-7***}$	$1.3037 \cdot 10^{-6***}$
GS	$3.6333 \cdot 10^{-4***}$	$-1.1263 \cdot 10^{-7***}$	$-7.4659 \cdot 10^{-7***}$	$3.9491 \cdot 10^{-7*}$	$1.3429 \cdot 10^{-7***}$	$1.5024 \cdot 10^{-6***}$
MSFT	$2.0482 \cdot 10^{-4***}$	$-2.2526 \cdot 10^{-9***}$	$-5.9639 \cdot 10^{-7***}$	$7.9095 \cdot 10^{-7***}$	$5.7893 \cdot 10^{-8***}$	$9.4286 \cdot 10^{-7***}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Ticker	Out-of-Sample Results					
	Intercept	RLP^S	Volatility	Ave. Order Size	Max. Order Size	Trade Volume
AAL	$1.9202 \cdot 10^{-4*}$	$2.5516 \cdot 10^{-9***}$	$-3.1496 \cdot 10^{-7}$	$-6.3956 \cdot 10^{-7}$	$4.9125 \cdot 10^{-8}$	$-3.1861 \cdot 10^{-7}$
AAPL	$9.3357 \cdot 10^{-5}$	$1.9734 \cdot 10^{-9***}$	$8.7033 \cdot 10^{-8}$	$-4.4362 \cdot 10^{-7}$	$6.0012 \cdot 10^{-10}$	$-3.5510 \cdot 10^{-8}$
BAC	$1.3369 \cdot 10^{-5}$	$-1.9925 \cdot 10^{-9}$	$-1.2792 \cdot 10^{-7***}$	$9.3238 \cdot 10^{-8**}$	$2.3771 \cdot 10^{-8**}$	$-6.4219 \cdot 10^{-7}$
CAT	$6.3433 \cdot 10^{-5}$	$-3.7999 \cdot 10^{-8***}$	$-1.6442 \cdot 10^{-6***}$	$1.7323 \cdot 10^{-6***}$	$2.2593 \cdot 10^{-7***}$	$-1.5597 \cdot 10^{-6***}$
CSCO	$1.4911 \cdot 10^{-4***}$	$3.8368 \cdot 10^{-10}$	$5.1161 \cdot 10^{-8}$	$-2.2769 \cdot 10^{-7***}$	$-4.2516 \cdot 10^{-9}$	$-4.0320 \cdot 10^{-7**}$
FB	$4.2145 \cdot 10^{-5}$	$1.1506 \cdot 10^{-10}$	$6.6110 \cdot 10^{-7**}$	$-7.2938 \cdot 10^{-7}$	$-6.0814 \cdot 10^{-8***}$	$-4.3409 \cdot 10^{-8}$
GE	$2.7858 \cdot 10^{-5}$	$-2.8902 \cdot 10^{-9*}$	$-1.2080 \cdot 10^{-7***}$	$7.0703 \cdot 10^{-8*}$	$2.6771 \cdot 10^{-8***}$	$-8.1788 \cdot 10^{-7***}$
GOOG	$2.2220 \cdot 10^{-5}$	$1.3626 \cdot 10^{-8***}$	$-5.6352 \cdot 10^{-6***}$	$4.5299 \cdot 10^{-6***}$	$6.9692 \cdot 10^{-7***}$	$-1.9022 \cdot 10^{-7}$
GS	$7.9081 \cdot 10^{-5}$	$1.3522 \cdot 10^{-8**}$	$-2.8851 \cdot 10^{-6***}$	$1.5563 \cdot 10^{-6**}$	$5.0228 \cdot 10^{-7***}$	$-8.2255 \cdot 10^{-7*}$
MSFT	$2.0128 \cdot 10^{-4***}$	$-2.2985 \cdot 10^{-9**}$	$-1.0861 \cdot 10^{-7}$	$-9.3074 \cdot 10^{-8}$	$1.0610 \cdot 10^{-8}$	$-3.5835 \cdot 10^{-7***}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the in-sample and out-of-sample results of regression (R5), which shows the relationship between the change in price over an interval and the relative provision of liquidity over that period as well as the distribution of market order flow over that period. The change in price over an interval that is unexplained by the market order flow and liquidity provision over that interval will be used in regression (R6) to test whether market makers have information about the presence of informed traders.

Table 5: Unexplained Price Impact Regressed on Change of RLP^S Over Previous Period

Ticker	In-Sample	Out-of-Sample
AAL	$-4.9120 \cdot 10^{-10}$	$-5.3749 \cdot 10^{-09***}$
AAPL	$-6.4825 \cdot 10^{-11}$	$-1.7225 \cdot 10^{-09***}$
BAC	$1.7015 \cdot 10^{-09***}$	$-1.4942 \cdot 10^{-09}$
CAT	$-1.1578 \cdot 10^{-09}$	$7.2310 \cdot 10^{-09***}$
CSCO	$8.3471 \cdot 10^{-11}$	$-1.6246 \cdot 10^{-09***}$
FB	$-1.7924 \cdot 10^{-10}$	$1.1237 \cdot 10^{-09***}$
GE	$7.5276 \cdot 10^{-10}$	$1.7530 \cdot 10^{-09*}$
GOOG	$-1.8877 \cdot 10^{-09*}$	$-2.7350 \cdot 10^{-09**}$
GS	$1.6629 \cdot 10^{-08***}$	$-2.2550 \cdot 10^{-08***}$
MSFT	$6.7075 \cdot 10^{-10**}$	$1.2028 \cdot 10^{-09***}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the results of regression (R6). The residuals from regression (R5) represent the change in price over an interval that is unexplained by the market order flow and liquidity provision over that interval. These residuals are regressed on the change in the provision of liquidity in the preceding interval. Here, a negative coefficient corresponds to market makers decreasing their provision of liquidity ahead of periods during which market orders have an unusually large impact on prices.

Table 6: Price Change Regressed on Recent ΔRLP^D

Ticker	In-Sample		
	Regression (R7)	Regression (R8)	
	ΔRLP_{t-1}^D	ΔX_{t-1}	ΔRLP_{t-1}^D
AAL	$4.2566 \cdot 10^{-09**}$	$-1.4000 \cdot 10^{-02}$	$4.2588 \cdot 10^{-09**}$
AAPL	$-3.8163 \cdot 10^{-09***}$	$-2.7340 \cdot 10^{-03}$	$-3.7939 \cdot 10^{-09***}$
BAC	$1.2879 \cdot 10^{-09*}$	$-1.3278 \cdot 10^{-02}$	$1.2949 \cdot 10^{-09*}$
CAT	$1.4708 \cdot 10^{-08**}$	$-1.2650 \cdot 10^{-02}$	$1.4694 \cdot 10^{-08**}$
CSCO	$-8.3015 \cdot 10^{-10}$	$-3.1464 \cdot 10^{-03}$	$-8.3508 \cdot 10^{-10}$
FB	$-2.5590 \cdot 10^{-10}$	$-7.4748 \cdot 10^{-02***}$	$-6.7407 \cdot 10^{-11}$
GE	$-1.1423 \cdot 10^{-09}$	$-2.0359 \cdot 10^{-02}$	$-1.1039 \cdot 10^{-09}$
GOOG	$2.2776 \cdot 10^{-09}$	$1.4970 \cdot 10^{-02}$	$2.2939 \cdot 10^{-09}$
GS	$1.7864 \cdot 10^{-08}$	$1.8775 \cdot 10^{-02}$	$1.8237 \cdot 10^{-08*}$
MSFT	$3.3149 \cdot 10^{-09*}$	$-2.1119 \cdot 10^{-02}$	$3.3941 \cdot 10^{-09*}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Ticker	Out-of-Sample		
	Regression (R7)	Regression (R8)	
	ΔRLP_{t-1}^D	ΔX_{t-1}	ΔRLP_{t-1}^D
AAL	$-5.2015 \cdot 10^{-09**}$	$-7.7113 \cdot 10^{-02***}$	$-4.5798 \cdot 10^{-09**}$
AAPL	$-1.2518 \cdot 10^{-09}$	$-3.4017 \cdot 10^{-02}$	$-1.0434 \cdot 10^{-09}$
BAC	$6.9468 \cdot 10^{-10}$	$-9.4427 \cdot 10^{-02***}$	$7.7180 \cdot 10^{-10}$
CAT	$6.9142 \cdot 10^{-08***}$	$-1.0041 \cdot 10^{-02}$	$6.9330 \cdot 10^{-08***}$
CSCO	$-2.5751 \cdot 10^{-09}$	$-2.5623 \cdot 10^{-02}$	$-2.4892 \cdot 10^{-09}$
FB	$4.3323 \cdot 10^{-09***}$	$8.3533 \cdot 10^{-02***}$	$4.7430 \cdot 10^{-09***}$
GE	$-1.6075 \cdot 10^{-09}$	$-1.0701 \cdot 10^{-01***}$	$-1.5358 \cdot 10^{-09}$
GOOG	$-2.8976 \cdot 10^{-09*}$	$1.0154 \cdot 10^{-01***}$	$-3.4427 \cdot 10^{-09**}$
GS	$1.5475 \cdot 10^{-08}$	$-1.3231 \cdot 10^{-01***}$	$1.5528 \cdot 10^{-08}$
MSFT	$-2.9104 \cdot 10^{-09}$	$2.0643 \cdot 10^{-03}$	$-2.9199 \cdot 10^{-09}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the results of regressions (R7) and (R8) over the in-sample and out-of-sample periods. Regression (R7) shows the relationship between the change in price over an interval and the change in the directional liquidity estimator RLP^D over the previous interval. Regression (R8) controls for the autoregressive nature of asset price returns, showing the relationship between the change in price over an interval and the change in price as well as the change in RLP^D over the previous interval. For both regressions, a positive coefficient acting on ΔRLP^D indicates that changes to the balance of the limit order book are predictive of future asset price returns.

Table 7: Price Impact of Market Order Regressed on Current RLP^S Value

Ticker	In-Sample		
	Regression (R9)		Regression (R10)
	Instantaneous Drift	Price Impact	RLP^S
AAL	$-8.7103 \cdot 10^{-5***}$	$1.8990 \cdot 10^{-5***}$	$-2.5504 \cdot 10^{-9***}$
AAPL	$4.7955 \cdot 10^{-5}$	$1.5960 \cdot 10^{-5***}$	$-4.3431 \cdot 10^{-10***}$
BAC	$-3.2154 \cdot 10^{-5***}$	$2.5596 \cdot 10^{-6***}$	$-3.5702 \cdot 10^{-11}$
CAT	$-1.3218 \cdot 10^{-4***}$	$1.0622 \cdot 10^{-4***}$	$-2.0600 \cdot 10^{-8***}$
CSCO	$-4.4888 \cdot 10^{-6}$	$3.1722 \cdot 10^{-6***}$	$-1.1904 \cdot 10^{-10***}$
FB	$-1.3963 \cdot 10^{-5}$	$1.2314 \cdot 10^{-5***}$	$-5.6309 \cdot 10^{-10***}$
GE	$-2.6110 \cdot 10^{-5***}$	$6.2396 \cdot 10^{-6***}$	$-5.3493 \cdot 10^{-10***}$
GOOG	$-6.4754 \cdot 10^{-4***}$	$3.5245 \cdot 10^{-4***}$	$-1.7054 \cdot 10^{-7***}$
GS	$-4.4317 \cdot 10^{-4***}$	$2.9331 \cdot 10^{-4***}$	$-1.6164 \cdot 10^{-7***}$
MSFT	$-3.1419 \cdot 10^{-6}$	$1.1234 \cdot 10^{-5***}$	$-5.6527 \cdot 10^{-10***}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Ticker	Out-of-Sample		
	Regression (R9)		Regression (R10)
	Instantaneous Drift	Price Impact	RLP^S
AAL	$-1.7541 \cdot 10^{-4***}$	$2.4159 \cdot 10^{-5***}$	$-7.9252 \cdot 10^{-10***}$
AAPL	$3.4153 \cdot 10^{-4***}$	$6.6096 \cdot 10^{-6***}$	$-5.1977 \cdot 10^{-10***}$
BAC	$-4.6142 \cdot 10^{-5***}$	$3.2974 \cdot 10^{-6***}$	$-2.0296 \cdot 10^{-11}$
CAT	$-3.9524 \cdot 10^{-5}$	$6.5479 \cdot 10^{-5***}$	$-4.5932 \cdot 10^{-9***}$
CSCO	$9.1064 \cdot 10^{-6}$	$5.1714 \cdot 10^{-6***}$	$-2.4917 \cdot 10^{-10***}$
FB	$9.3878 \cdot 10^{-6}$	$1.0151 \cdot 10^{-5***}$	$-9.0984 \cdot 10^{-11***}$
GE	$-4.4760 \cdot 10^{-5***}$	$5.4312 \cdot 10^{-6***}$	$-6.9861 \cdot 10^{-10***}$
GOOG	$-3.9607 \cdot 10^{-4**}$	$3.8963 \cdot 10^{-4***}$	$-8.1497 \cdot 10^{-8***}$
GS	$-3.5862 \cdot 10^{-4***}$	$3.5122 \cdot 10^{-4***}$	$-5.0625 \cdot 10^{-7***}$
MSFT	$-1.8191 \cdot 10^{-4***}$	$5.6876 \cdot 10^{-6***}$	$-4.2519 \cdot 10^{-11***}$

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the results of regressions (R9) and (R10) performed over the in-sample and out-of-sample periods. Regression (R9) finds the average price impact of a market order, controlling for the instantaneous price drift. Regression (R10) uses the residual price impact from regression (R9) to find the effect that the current value of the estimator (RLP_t^S) has on the unexplained price impact. For regression (R10), a negative coefficient multiplying the current value of the estimator implies that the price impact of an arriving market order is lower than normal when market makers are placing relatively larger limit orders.

Table 8: Estimated Parameters of the Hidden Markov Model

Parameters		Estimate
Initial State	α_1	0.0797
Distribution	α_2	0.3404
Transition Matrix	γ_2	0.9154
Trade	γ_3	0.8724
Transition Matrix	λ_1	0.8250
No Trade	λ_2	0.9374
State	β_0^1	3.5026
Dependent	β_0^2	-0.0164
Factors	β_0^3	-6.1739
	$\beta_1(1)$	0.6078
Other	$\beta_1(2)$	-0.0829
Factors	$\beta_1(3)$	0.1827
	$\beta_1(4)$	-0.3651

This table summarizes the parameters that optimize the hidden Markov model over the in-sample training data. Using these parameters, it is possible to find the initial state distribution that is estimated for the sample. The initial state distribution is estimated as

$$\pi_0 = [\alpha_1, \alpha_2, (1 - \alpha_1 - \alpha_2)] = [0.0797, 0.3404, 0.5799].$$

Thus, over half of the bonds in the study are expected to begin their life cycle in the dormant state for which few trades are expected, and only $\sim 8\%$ of bonds are expected to begin their life cycle as actively traded bonds. This initial state distribution is similar to the distribution of trade frequencies across bond issues (Figure 9). Similarly, it is possible to find the transition matrices estimated by the model. The transition matrix given a trade occurs is estimated as

$$Tr_1 = \begin{bmatrix} 1 & 0 & 0 \\ 1 - \gamma_2 & \gamma_2 & 0 \\ 0 & 1 - \gamma_3 & \gamma_3 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0.0846 & 0.9154 & 0 \\ 0 & 0.1276 & 0.8724 \end{bmatrix},$$

and the transition matrix given no trade occurs is estimated as

$$Tr_0 = \begin{bmatrix} \lambda_1 & 1 - \lambda_1 & 0 \\ 0 & \lambda_2 & 1 - \lambda_2 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 0.8250 & 0.1750 & 0 \\ 0 & 0.9374 & 0.0626 \\ 0 & 0 & 1 \end{bmatrix}.$$

Table 9: Estimated Parameters of Alternative Trading Probability Models

Model	Parameters
Logistic Regression	β_0 : -1.1214
	$\beta_1(1)$: 0.0119
	$\beta_1(2)$: 0.0356
	$\beta_1(3)$: 0.0600
	$\beta_1(4)$: 0.0995
	$\beta_1(5)$: -0.2621
Weighted Logistic Regression	ω : 3.0707
	β_0 : -0.6045
	$\beta_1(1)$: 3.1277
	$\beta_1(2)$: 0.2707
	$\beta_1(3)$: -0.8349
	$\beta_1(4)$: -0.4617
Unconditional Probability	α : 0.2457

This table summarizes the parameters (and in the case of the weighted logistic regression model, the hyper-parameter ω) that optimize the performance of the alternative trading probability estimators over the in-sample training data set. For the logistic regression and the weighted logistic regression models, $\beta_1(1)$ is the parameter corresponding to the indicator variable of a trade occurring on the previous business day. From the value of $\beta_1(1)$ for the weighted logistic regression model, it can be seen that for a given bond, trading activity being observed on the previous business day is strongly correlated with a trade occurring on the day in question. However, for the standard logistic regression model, the value of $\beta_1(1)$ is much smaller, showing that the presence of a trade on the previous business day has significantly less weight on the estimated probability of trading. This marked decrease in impact on trade probability is due to the imbalanced nature of the data; without the class-weight ω , the proportion of observations without a trade observed overwhelms the impact of a trade occurring on the previous day.

Table 10: In-Sample and Out-of-Sample Performance of Trading Probability Estimators

Model	In-Sample		Out-of-Sample	
	ROC AUC	Log Loss	ROC AUC	Log Loss
Hidden Markov Model	0.9473	0.2490	0.9519	0.2325
Logistic Regression	0.5759	0.5565	0.5758	0.5377
Weighted Logistic Regression	0.8268	0.4897	0.8272	0.4827
Unconditional Probability	0.5	0.5575	0.5	0.5385

This table summarizes the discriminatory power and predictive accuracy of the hidden Markov model and the alternative trade probability estimators. The ROC AUC measures the discriminatory power of the classification model, with an AUC of 1 corresponding to a perfect classifier and an AUC of 0.5 corresponding to a classifier that differentiates between classes no better than random chance. The log loss measures the accuracy of the classification model, with a lower log loss corresponding to a more accurate probability estimate. Both in-sample, over the training data set, and out-of-sample, over the test data set, the hidden Markov model outperforms the other estimators by a large margin, with a high level of discriminatory power (as demonstrated by a high ROC AUC) and a high level of predictive accuracy (as demonstrated by a low log loss). As expected, the weighted logistic regression significantly improves upon the standard logistic regression with respect to discriminatory power, as measured by ROC AUC. However, with respect to the accuracy of trade probability predictions, the weighted logistic regression’s improvement over the standard logistic regression is much more modest. The unconditional probability model assigns the same trade probability to every observation. Therefore, as expected, the unconditional probability model is unable to differentiate between classes.

Table 11: Estimated Matching Model Hyper-Parameters

Model	Hyper-Parameters
LOWESS Regression	q : 0.0074 α_1 : 0.9839 α_2 : 0.9942
OLS Regression	No hyper-parameters
K-Nearest Neighbor Search	k : 90
Deep Neural Network	d : 4 w : 64

This table summarizes the hyper-parameters that optimize the performance of the matching estimators over the in-sample training data set. For the LOWESS regression model, the hyper-parameter q determines the number of neighboring observations that are used in each local regression on a given day t as a proportion of the number of bonds that trade on that day n_t , where the number of neighbors to include is given by qn_t . α_1 and α_2 are hyper-parameter weights on the principal components that determine the impact each component has on the distance measurement used to find the nearest observations. The OLS regression model does not employ any hyper-parameters and therefore requires no pre-training. The behavior of the k-nearest neighbor search model is fully determined by the hyper-parameter k , which gives the number of neighboring observations to use in its estimation. The structure of the deep neural network model is described by the hyper-parameters d and w . In this model, d describes the depth of the model, or the number of hidden layers within the model. The hyper-parameter w describes the width of the model, or the number of nodes contained on each hidden layer of the model.

Table 12: In-Sample and Out-of-Sample Performance of Matching Estimators

Model	In-Sample		Out-of-Sample	
	RMSE	R^2	RMSE	R^2
LOWESS Regression	1.0886	0.3202	1.0256	0.3447
OLS Regression	1.3158	0.0068	1.2583	0.0136
K-Nearest Neighbor Search	1.3090	0.0207	1.2508	0.0253
Deep Neural Network	1.2294	0.1692	1.1844	0.1295

This table summarizes the explanatory and predictive power of the LOWESS Regression and the alternative matching estimators, as well as the accuracy of their predictions. The in-sample R^2 describes the power of the estimators to explain the variation of the target variable in the training sample, and the out-of-sample R^2 measures the power of the models to predict the variation of the target variable in the testing sample. The in-sample and out-of-sample RMSEs measure the accuracy of the models' predictions over the training and test samples, respectively. The LOWESS regression outperforms the other estimators by a large margin, followed by the deep neural network estimator. The ordinary least squares (OLS) regression and the k-nearest neighbor (KNN) search perform very similarly to one another in terms of RMSE, with the KNN search algorithm slightly outperforming the OLS regression with respect to R^2 .

Table 13: Relationship Between Aggregate Illiquidity and Other Macroeconomic Factors

Coefficient	Regression	
	VIX	OAS
Intercept	4.7248***	0.0539*
Agg. Illiq.	37.6468***	4.6276***
R^2	0.4662	0.5553

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the results of regressing each macroeconomic factor on my measure of aggregate bond market illiquidity. Here, “Intercept” reports the intercept of the regression, and “Agg. Illiq.” reports the coefficient multiplying my measure of aggregate illiquidity.

Table 14: Credit Rating Restrictions of the 10 Largest U.S. Pension Funds

Pension Fund	Credit Rating Restriction
California Public Employees' Retirement System (CalPERS)	At least one rating must be investment grade.
California State Teachers' Retirement System (CalSTRS)	At least two ratings* must be investment grade.
New York State Common Retirement Fund (NYSCRF)	At least two ratings* must be investment grade.
New York City Retirement System (NYCRS)	At least two ratings* must be investment grade.
State Board of Administration of Florida (Florida SBA)	At least one rating must be investment grade.
Teacher Retirement System of Texas (TRST)	No specific restriction; mostly externally managed.
New York State Teachers' Retirement System (NYSTRS)	At least two ratings* must be investment grade.
State of Wisconsin Investment Board (SWIB)	All ratings must be investment grade.
North Carolina Retirement System (NCRS)	At least one rating must be investment grade.
Washington State Investment Board (WSIB)	At least two ratings* must be investment grade.

*If only one rating exists, it must be investment grade.

This table reports the credit rating restrictions for U.S. Corporate Bonds held by the ten largest U.S. public pension funds, as measured by total assets. Of the ten pension funds, only one fund (SWIB) requires that all ratings must be investment grade. Five of the ten funds require that at least two ratings are investment grade (unless the bond is only rated by one NRSRO, in which case it must be investment grade). Three of the ten funds only require that at least one rating is investment grade. One of the ten pension funds (TRST), has no explicit restrictions on credit ratings.

Table 15: Summary Statistics of Rating Changes Included in Sample

Type of Rating Change	# Included in Time Period			Total Included
	1/1/2010 to 7/20/2012	7/21/2012 to 1/13/2014	1/14/2014 to 3/31/2016	
All Changes	2398	3699	7291	13388
All Upgrades	126	167	150	443
All Downgrades	41	167	423	631
Fitch Changes	287	574	1045	1906
Fitch Upgrades	25	11	38	74
Fitch Downgrades	16	48	88	152
Moody's Changes	1692	1666	3039	6397
Moody's Upgrades	84	35	48	167
Moody's Downgrades	11	74	192	277
S&P Changes	419	1459	3207	5085
S&P Upgrades	17	121	64	202
S&P Downgrades	14	45	143	202
Surprise Upgrades	67	68	65	200
Surprise Downgrades	30	101	226	357
Definitive Upgrades	76	139	88	303
Definitive Downgrades	13	72	207	292
Near Upgrades	113	273	507	893
Near Downgrades	107	186	697	990
Large Upgrades	2	20	15	37
Large Downgrades	23	61	90	174

This table describes the distribution of rating changes that are included in the sample. The period from 1/1/2010 to 7/20/2012 includes changes that occurred before the Volcker rules were announced, the period from 7/21/2012 to 1/13/2014 includes changes that occurred after the new regulations were announced but before they went into effect, and the period from 1/14/2014 to 5/31/2016 includes changes that occurred after the Volcker rule regulations were enforced. Here, a “surprise” upgrade or downgrade refers to the first rating upgrade or downgrade that occurred among the three credit rating agencies. A “definitive” upgrade or downgrade refers to the second or third rating upgrade or downgrade that occurred among the rating agencies. “Near” upgrades and downgrades take ratings across the threshold between BBB and BBB-, whereas “large” upgrades or downgrades represent upgrades or downgrades that neither end nor begin at a BBB- rating.

Table 16: Price Impact of Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R (\beta_0)$	0.2975***	0.2999***	0.2718***	0.2801***
$Up (\beta_1)$	0.5843***	0.5191***	0.5423***	0.5787***
$Down (\beta_2)$	-0.7043***	-0.7253***	-0.7400***	-0.6983***
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Table 17: Effects of Volcker Rule on Price Impact of Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R (\beta_0)$	0.2981***	0.3091***	0.2774***	0.2831***
$Up (\beta_1)$	1.2638***	0.9392***	1.0833***	1.1088***
$Up \times During (\beta_3)$	-0.9122*	-0.6185	-0.7728**	-0.7044**
$Up \times After (\beta_5)$	-0.9922*	-0.5682	-0.7351*	-0.7679**
$Down (\beta_2)$	-0.9635	-0.0241	-0.3487	-0.5817
$Down \times During (\beta_4)$	0.4836	-0.8377	-0.4187	0.0436
$Down \times After (\beta_6)$	0.1971	-0.6949	-0.4042	-0.1824
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

These tables report the results of regressions (R1) and (R2), regressing the change in volume-weighted average price surrounding credit rating changes on the size and direction of the rating changes, as well as flags for upgrades from junk to investment grade and flags for downgrades from investment grade to junk status. Table 17 further applies a Diff-in-Diff framework to study the effects of the announcement and enforcement of the Volcker rule. Regressions are performed on volume-weighted average prices calculated over the time windows of two months, one month, two weeks, and one week. Here, ΔR represents the size and direction of a credit rating change, Up represents a rating upgrade from junk to investment grade, $Down$ represents a rating downgrade from investment grade to junk, $During$ represents the period after the Volcker rule was announced but before it was enforced, and $After$ represents the period after the Volcker rule was enforced.

Table 18: Test of Relative Impact of Upgrades and Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$ \beta_1 $	0.5843	0.5191	0.5423	0.5787
(SE)	(0.2157)	(0.1742)	(0.1538)	(0.1393)
n	443	406	371	325
$ \beta_2 $	0.7043	0.7253	0.7400	0.6983
(SE)	(0.1966)	(0.1543)	(0.1319)	(0.1139)
n	631	612	597	566
t	4.2486	7.9019	7.8419	4.7632
d.f.	923	830	739	621
$\mathbb{P}(H_0)$	$1.1848 \cdot 10^{-5}$	$4.3299 \cdot 10^{-15}$	$7.7716 \cdot 10^{-15}$	$1.1869 \cdot 10^{-6}$

This table reports the results of performing Welch’s unequal variances t-test on the following hypotheses regarding the relative size of the coefficients regression (R1) reported in Table 16:

$$H_0 : |\beta_2| \leq |\beta_1|, \quad H_a : |\beta_2| > |\beta_1|.$$

t reports the t-statistic from performing the above-described test. d.f. reports the pooled degrees of freedom for the test, as found by rounding the result of the Welch-Satterwaite equation down to the nearest integer.

Table 19: Price Impact of Surprise Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R^s (\beta_0)$	0.3419***	0.3239***	0.2807***	0.2237***
$Up^s (\beta_1)$	0.5975*	0.6910***	0.6904***	0.7347***
$Down^s (\beta_2)$	-0.3488	-0.5690***	-0.8330***	-1.1962***
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Table 20: Effects of Volcker Rule on Price Impact of Surprise Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R^s (\beta_0)$	0.3439***	0.3393***	0.2924***	0.2326***
$Up^s (\beta_1)$	0.7725	0.9486**	1.2351***	1.3151***
$Up^s \times During (\beta_3)$	-0.1234	-0.1801	-0.6003	-0.6478
$Up^s \times After (\beta_5)$	-0.4128	-0.5680	-0.9122	-0.9295*
$Down^s (\beta_2)$	-0.0092	0.5960	0.1086	-0.2707
$Down^s \times During (\beta_4)$	-0.9872	-2.0419***	-1.3417**	-0.8270
$Down^s \times After (\beta_6)$	-0.0887	-0.8873	-0.8533	-1.0458*
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

These tables report the results of regressions (R3) and (R4), regressing the change in volume-weighted average price surrounding credit rating changes on the size and direction of the rating changes, as compared to the existing credit ratings from other rating agencies, as well as flags for “surprise” upgrades from junk to investment grade and flags for “surprise” downgrades from investment grade to junk status. Table 20 further applies a Diff-in-Diff framework to study the effects of the announcement and enforcement of the Volcker rule. Regressions are performed on volume-weighted average prices calculated over the time windows of two months, one month, two weeks, and one week. Here, ΔR^s represents the size and direction of a credit rating change, relative to the most recent credit rating from any of the NRSROs, Up^s represents a “surprise” rating upgrade from junk to investment grade, $Down^s$ represents a “surprise” rating downgrade from investment grade to junk, *During* represents the period after the Volcker rule was announced but before it was enforced, and *After* represents the period after the Volcker rule was enforced.

Table 21: Test of Relative Impact of Surprise Upgrades and Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$ \beta_1 $	0.5975	0.6910	0.6904	0.7347
(SE)	(0.3197)	(0.2658)	(0.2417)	(0.2226)
n	200	173	149	126
$ \beta_2 $	0.3488	0.5690	0.8330	1.1962
(SE)	(0.2649)	(0.2086)	(0.1779)	(0.1538)
n	357	343	334	315
t	-5.1398	-2.6353	3.0718	9.7184
d.f.	380	310	249	197
$\mathbb{P}(H_0)$	1	0.9956	0.0012	0

This table reports the results of performing Welch's unequal variances t-test on the following hypotheses regarding the relative size of the coefficients of regression (R3) reported in Table 19:

$$H_0 : |\beta_2| \leq |\beta_1|, \quad H_a : |\beta_2| > |\beta_1|.$$

t reports the t-statistic from performing the above-described test. d.f. reports the pooled degrees of freedom for the test, as found by rounding the result of the Welch-Satterwaite equation down to the nearest integer.

Table 22: Price Impact of Definitive Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R (\beta_0)$	0.3392***	0.3520***	0.3397***	0.3647***
$Up^d (\beta_1)$	0.4296*	0.3137	0.3576*	0.3975**
$Down^d (\beta_2)$	-1.1371***	-1.0022***	-0.6314***	-0.1756
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Table 23: Effects of Volcker Rule on Price Impact of Definitive Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R (\beta_0)$	0.3440***	0.3549***	0.3419***	0.3679***
$Up^d (\beta_1)$	1.4828***	0.8502**	0.8855**	0.8595***
$Up^d \times During (\beta_3)$	-1.4361**	-0.8109	-0.7833*	-0.6426
$Up^d \times After (\beta_5)$	-1.3667*	-0.6067	-0.6297	-0.6375
$Down (\beta_2)$	-2.1482*	-0.6775	-0.6176	-1.7365**
$Down^d \times During (\beta_4)$	2.2564*	0.2761	0.5695	1.9839**
$Down^d \times After (\beta_6)$	0.6494	-0.5433	-0.2019	1.4987*
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

These tables report the results of regressions (R5) and (R6), regressing the change in volume-weighted average price surrounding credit rating changes on the size and direction of the rating changes, as well as flags for “definitive” upgrades from junk to investment grade and flags for “definitive” downgrades from investment grade to junk status. Table 23 further applies a Diff-in-Diff framework to study the effects of the announcement and enforcement of the Volcker rule. Regressions are performed on volume-weighted average prices calculated over the time windows of two months, one month, two weeks, and one week. Here, ΔR represents the size and direction of a credit rating change, Up^d represents a “definitive” rating upgrade from junk to investment grade, $Down^d$ represents a “definitive” rating downgrade from investment grade to junk, *During* represents the period after the Volcker rule was announced but before it was enforced, and *After* represents the period after the Volcker rule was enforced.

Table 24: Test of Relative Impact of Definitive Upgrades and Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$ \beta_1 $	0.4296	0.3137	0.3576	0.3975
(SE)	(0.2597)	(0.2112)	(0.1875)	(0.1712)
n	303	274	248	214
$ \beta_2 $	1.1371	1.0022	0.6314	0.1756
(SE)	(0.2691)	(0.2110)	(0.1805)	(0.1569)
n	292	283	277	262
t	16.7730	17.6798	7.2961	-5.9318
d.f.	591	554	514	446
$\mathbb{P}(H_0)$	0	0	$5.6477 \cdot 10^{-13}$	1

This table reports the results of performing Welch’s unequal variances t-test on the following hypotheses regarding the relative size of the coefficients of regression (R5) reported in Table 22:

$$H_0 : |\beta_2| \leq |\beta_1|, \quad H_a : |\beta_2| > |\beta_1|.$$

t reports the t-statistic from performing the above-described test. d.f. reports the pooled degrees of freedom for the test, as found by rounding the result of the Welch-Satterwaite equation down to the nearest integer.

Table 25: Price Impact of Credit Rating Near Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R (\beta_0)$	0.3533***	0.3664***	0.3451***	0.3614***
$Up^{BBB} (\beta_1)$	0.3157**	0.1875	0.1907*	0.1295
$Down^{BBB-} (\beta_2)$	-0.2573*	-0.2239*	-0.2035*	-0.1371
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

Table 26: Effects of Volcker Rule on Price Impact of Credit Rating Near Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$\Delta R (\beta_0)$	0.3695***	0.3780***	0.3530***	0.3697***
$Up^{BBB} (\beta_1)$	0.9744**	0.4121	0.3589	0.3190
$Up^{BBB} \times During (\beta_3)$	-1.0589**	-0.4541	-0.3648	-0.3664
$Up^{BBB} \times After (\beta_5)$	-0.6009	-0.1690	-0.1239	-0.1632
$Down^{BBB-} (\beta_2)$	0.9701**	0.7166**	0.3699	0.2389
$Down^{BBB-} \times During (\beta_4)$	-0.9538*	-0.7371*	-0.2997	0.0432
$Down^{BBB-} \times After (\beta_6)$	-1.4691***	-1.1425***	-0.7438**	-0.5549*
N	13388	12282	10943	9556

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

These tables report the results of regressions (R7) and (R8), regressing the change in volume-weighted average price surrounding credit rating changes on the size and direction of the rating changes, as well as flags for upgrades from ratings of BBB- or worse to ratings of BBB or better and flags for downgrades from BBB or better to BBB- or worse. Table 26 further applies a Diff-in-Diff framework to study the effects of the announcement and enforcement of the Volcker rule. Regressions are performed on volume-weighted average prices calculated over the time windows of two months, one month, two weeks, and one week. Here, ΔR represents the size and direction of a credit rating change, Up^{BBB} represents a rating upgrade from BBB- or worse to BBB or better, $Down^{BBB-}$ represents a rating downgrade from BBB or better to BBB- or worse, *During* represents the period after the Volcker rule was announced but before it was enforced, and *After* represents the period after the Volcker rule was enforced.

Table 27: Test of Relative Impact of Near Upgrades and Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$ \beta_1 $	0.3157	0.1875	0.1907	0.1295
(SE)	(0.1518)	(0.1231)	(0.1120)	(0.1036)
n	893	812	700	588
$ \beta_2 $	0.2573	0.2239	0.2035	0.1371
(SE)	(0.1492)	(0.1210)	(0.1073)	(0.0944)
n	990	901	818	754
t	-3.2599	2.1523	0.7488	0.4365
d.f.	1857	1689	1469	1232
$\mathbb{P}(H_0)$	0.9994	0.0158	0.2271	0.3313

This table reports the results of performing Welch’s unequal variances t-test on the following hypotheses regarding the relative size of the coefficients of regression (R7) reported in Table 25:

$$H_0 : |\beta_2| \leq |\beta_1|, \quad H_a : |\beta_2| > |\beta_1|.$$

t reports the t-statistic from performing the above-described test. d.f. reports the pooled degrees of freedom for the test, as found by rounding the result of the Welch-Satterwaite equation down to the nearest integer.

Table 28: Test of Relative Impact of Near Downgrade Time Spans

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
$ \beta_6 $	1.4691	1.1425	0.7438	0.5549
(SE)	(0.4683)	(0.3712)	(0.3253)	(0.2919)
n	697	617	556	514

Factor	H_a		
	$ \beta_6^{2M} > \beta_6^{1M} $	$ \beta_6^{1M} > \beta_6^{2W} $	$ \beta_6^{2W} > \beta_6^{1W} $
t	9.1509	11.5756	5.5625
d.f.	1311	1169	1067
$\mathbb{P}(H_0)$	0	0	$1.6801 \cdot 10^{-8}$

This table reports the results of performing Welch’s unequal variances t-test on the following hypotheses regarding the relative size of the coefficients reported in Table 26:

$$\begin{aligned}
 H_0 : |\beta_6^{2M}| &\leq |\beta_6^{1M}|, & H_a : |\beta_6^{2M}| &> |\beta_6^{1M}| \\
 H_0 : |\beta_6^{1M}| &\leq |\beta_6^{2W}|, & H_a : |\beta_6^{1M}| &> |\beta_6^{2W}| \\
 H_0 : |\beta_6^{2W}| &\leq |\beta_6^{1W}|, & H_a : |\beta_6^{2W}| &> |\beta_6^{1W}|,
 \end{aligned}$$

where β_6^{2M} is the impact of the Volcker rule on the price impact of “near downgrades” over a two month window, β_6^{1M} is the impact of the Volcker rule over a one month window, β_6^{2W} is the impact over a two week window, and β_6^{1W} is the impact over a one week window. t reports the t-statistic from performing the above-described test. d.f. reports the pooled degrees of freedom for the test, as found by rounding the result of the Welch-Satterwaite equation down to the nearest integer.

Table 29: Effects of Volcker Rule on Price Impact of Large Credit Rating Downgrades

Factor	Time Window			
	Two Months	One Month	Two Weeks	One Week
ΔR (β_0)	0.1292	0.3157**	0.3468**	0.4192***
$Down^x$ (β_1)	-6.3711***	-7.2874***	0.6569	0.1956
$Down^x \times After$ (β_2)	4.2652	7.0650***	-0.6314***	-0.4576
N	1456	1158	847	616

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the results of regression (R9), regressing the change in volume-weighted average price surrounding credit rating changes on the size and direction of the rating changes, as well as flags for “large” downgrades from a rating of BBB or better to junk status, as well as a Diff-in-Diff flag to study the effects of the enforcement of the Volcker rule. Regressions are performed on volume-weighted average prices calculated over the time windows of two months, one month, two weeks, and one week. Here, ΔR represents the size and direction of a credit rating change, $Down^x$ represents a “large” rating downgrade, and $After$ represents the period after the Volcker rule was enforced.

To allow for the greatest probability that rating downgrades sparked a fire sale, the sample is restricted to bonds that are rated by only one NRSRO. Over the sample period, there are no large upgrades for bonds covered by only one rating agency, and there are only nine large downgrades for bonds covered by one agency. Moreover, only three such downgrades occur after the enforcement of the Volcker rule. It should be noted that the limited number of observations for which the “large downgrade” flag applies may lead to unreliable regression results.

Table 30: Response of Yield Curve to Volcker Rule Taking Effect

Factor	Regression					
	(1)	(2)	(3)	(4)	(5)	(6)
AAA Yield	1.7852***	1.9609***	1.8172***	1.9966***	1.1380***	1.2817***
Rating Yield Slope	0.8357***	0.8769***	0.7867***	0.8276***	0.3165***	0.3722***
Junk Rated	1.4290***	1.4725***	1.1922***	1.1953***	2.0284***	2.0616***
Junk Rated×Slope	–	–	0.4809***	0.5069***	1.1715***	1.1630***
BBB Premium	0.4764***	0.3665***	0.5925***	0.4814***	0.7318***	0.6390***
BBB- Premium	0.6467***	0.5092***	0.7774***	0.6386***	1.2009***	1.0744***
Term	–	–	–	–	0.1512***	0.1450***
% Time Through Sample	-1.6513***	-2.2408***	-1.6554***	-2.2494***	-1.5994***	-2.0310***
AAA×Volcker	0.5613***	-2.2781***	0.5599***	-2.3050***	0.6930***	-1.3690***
Rating×Volcker	–	-0.1395***	–	-0.1212***	–	-0.1656***
Junk Rated×Volcker	–	-0.2464*	–	-0.0506	–	-0.1325
Junk×Slope×Volcker	–	–	–	-0.2520***	–	-0.1891***
BBB×Volcker	-0.1576	-0.0293	-0.1541	-0.0665	-0.0890	0.0164
BBB-×Volcker	0.6084***	0.8915***	0.6119***	0.8490***	0.1716*	0.3814***
Term×Volcker	–	–	–	–	–	0.0193***
% Time×Volcker	–	4.0472***	–	4.0573***	–	2.8821***
Adjusted R^2	0.665	0.669	0.657	0.670	0.838	0.848

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively. N = 26570

This table reports the results of regressing the yield at time of issuance on bond ratings. “AAA Yield” gives the average yield of AAA rated bonds. “Rating Yield Slope” gives the additional yield for each full step in credit rating (for example, from AAA to AA, or from A to BBB). “Junk Rated” represents the average additional yield for bonds with a junk rating. “Junk Rated×Slope” gives the change in the yield slope for bonds with junk ratings. “BBB” and “BBB-” give the additional yields for the respective ratings above that which is given by the slope of the yield curve. “Term” gives the additional yield attributable to an additional year of the bond’s term to maturity. “% Time Through Sample” reports the change in yield attributable to the time at which the bond was issued, in order to account for the drift in overall yields throughout the period. Finally, the “Volcker” coefficients report the changes to the first five coefficients which are observed after the Volcker rule went into effect.

The “BBB-×Volcker” coefficient shows a sharp rise in the yield at origination of bonds with BBB- ratings following the enforcement of the Volcker rule.

Table 31: Yield Curve as Volcker Rule was Implemented

Factor	Regression Dates	
	1/1/2010 to 3/31/2016	7/21/2012 to 3/31/2016
BBB- Yield	4.5538***	3.0295***
Junk Premium	1.2072***	1.1451***
Term	0.0791***	0.0778***
BBB Premium	-0.4687***	-0.2872***
BB Premium	0.3704***	0.3401**
% of Time Through Sample	-2.6713***	–
BBB- \times Volcker	1.3331***	0.6723***
Junk \times Volcker	-0.8071***	-0.7397***
Term \times Volcker	0.0168***	0.0170**
BBB \times Volcker	-0.4587***	-0.7183***
BB \times Volcker	-0.1327	-0.0747
N	2128	992
Adjusted R^2	0.516	0.536

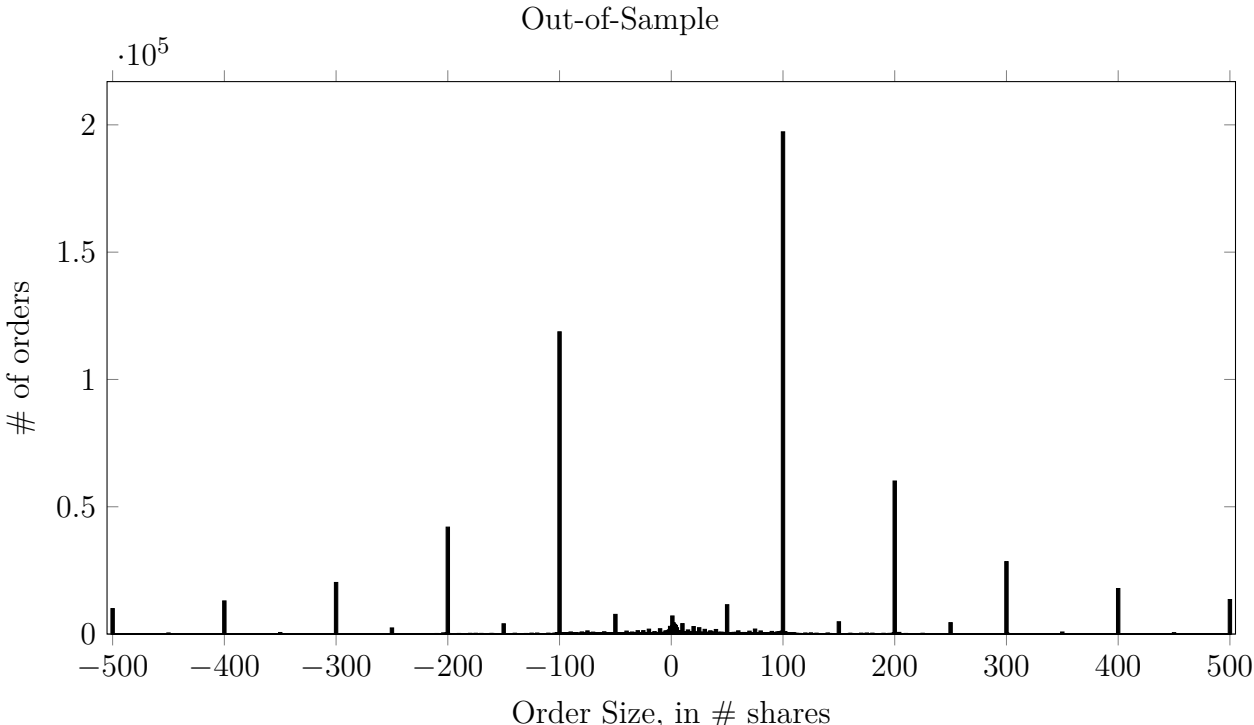

Note: *, **, and *** indicate significance at 0.10, 0.05, and 0.01, respectively.

This table reports the results of regressing the yield at time of issuance on bond ratings ranging from BBB to BB. “BBB- Yield” gives the yield of BBB- rated debt. “Junk Premium” gives the additional yield for a bond with either BB+ or BB ratings. “Term” represents the additional yield for a bond added by each additional year to maturity. “BBB Premium” gives the additional yield of BBB rated debt (a negative premium represents a discount). “BB Premium” represents the additional yield of BB rated debt. In this way, the average yield of a BB bond with a 7 year maturity is given by “BBB- Yield” + “Junk Premium” + “BB Premium” + 7 \times “Term.” For the sample which spans January 1, 2010 to March 31, 2016, “% Time Through Sample” reports the change in yield attributable to the time at which the bond was issued in order to account for the drift in overall yields throughout the period. Finally, the “Volcker” coefficients report the changes to the first five coefficients which are observed after the Volcker rule went into effect.

The “BBB- \times Volcker” coefficient shows a sharp rise in the yield at origination of bonds with BBB- ratings following the enforcement of the Volcker rule.

C Figures

Figure 1: Distribution of Market Order Sizes


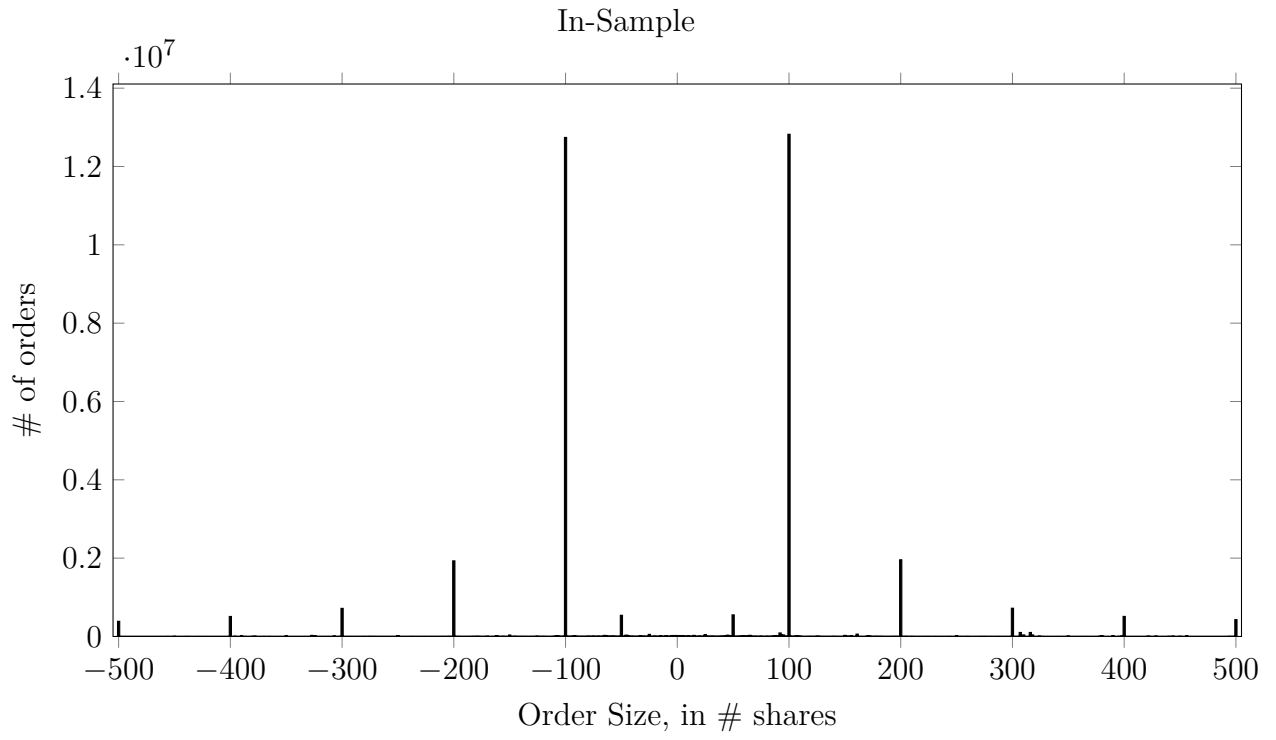


The figure above shows a histogram of the size of market orders for the ten stocks in the sample placed on NASDAQ between January 2, 2015 and January 23, 2015 (in-sample), as well as a histogram of the size of market orders for the ten stocks in the sample placed on NASDAQ between January 26, 2015 and January 30, 2015 (out-of-sample). Over half of the observed orders were of sizes greater than or equal to 100 shares. Of these, a large majority are orders of 100 shares, and the rest stylistically follow a geometric distribution with mass points at multiples of 100 shares. Among orders of sizes less than 100 shares, there are spikes at multiples of 50 shares, but the majority of the mass tends towards very small orders. When aggregated over the ten stocks in the sample, market orders are biased toward the buy side, but this is largely due to AAPL; at the individual security level, eight of the ten stocks are biased toward the sell side. The distribution of market orders placed in the out-of-sample period closely resembles that of orders placed in the in-sample period.

The mass of small orders may be a result of Reg NMS and the protected quotes established by the regulation. Rule 611 requires that exchanges covered by the rule route part (or all) of an incoming market order to another covered exchange if the best quote on the other exchange is better than the best quote on the original exchange's platform. Therefore, the small orders observed on NASDAQ could possibly be a partial order being routed to NASDAQ from another exchange, the remainder of an order which was placed on NASDAQ but partially routed to other exchanges, or even simply an intentionally small order.


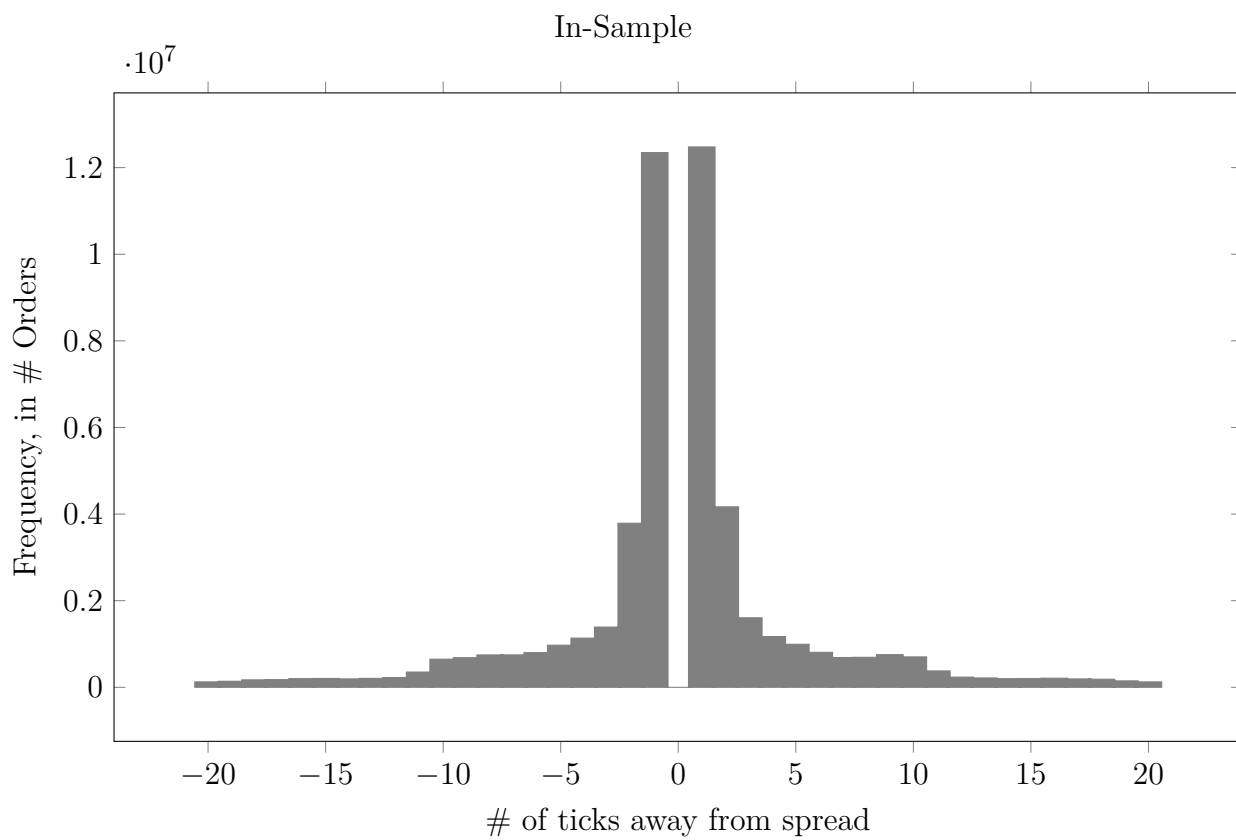
Reg NMS only protects the best quotes on protected markets when the order is placed. For example, suppose that the best quote on exchange A is beaten by the best two quotes on exchange B. If exchange A receives a market order, it is only obligated to re-route the portion of the order which will be filled by the best quote on exchange B; it is not required to give the order the best execution possible. For this reason, it is conceivable that some traders are taking advantage of Reg NMS by placing several small orders back-to-back rather than placing one full order. In this way, the trader would effectively be guaranteed to receive the best execution possible among all protected exchanges.

Figure 2: Distribution of Limit Order Sizes



The figure above shows a histogram of the size of limit orders for the ten stocks in the sample placed near the bid-ask spread on NASDAQ between January 2, 2015 and January 23, 2015 (in-sample), as well as a histogram of the size of limit orders for the ten stocks in the sample placed on NASDAQ between January 26, 2015 and January 30, 2015 (out-of-sample). Similar to market orders, the vast majority of limit orders are orders of 100 shares, with the rest stylistically following a geometric distribution with mass points at multiples of 100 shares. Again, the out-of-sample distribution closely resembles that of the in-sample period.



Figure 3: Distribution of Limit Order Placements/Revisions across Prices



The figure above is a histogram showing the distribution of locations in the limit order book at which limit orders are placed/revise, both for the in-sample period from January 2, 2015 through January 23, 2015, as well as the out-of-sample period from January 26, 2015 through January 30, 2015. The in-sample and out-of-sample distributions of limit order placements/revisions are very similar to one another. The overwhelming majority of limit orders on either side of the limit order book are placed within five ticks of the bid-ask spread. Limit order activity remained elevated between six and ten ticks away from the spread. Beyond ten ticks from the spread, the velocity of limit order activity rapidly decays as the distance from the bid-ask spread increases.

The limit order activity between six and ten ticks away from the spread is consistent with these orders being placed to serve as liquidity back-stops. It is beneficial to a market maker to supply liquidity some distance away from the bid-ask spread in order to both capture additional profit from (infrequent) block orders and reduce the likelihood that a price jump will occur due to a sudden influx of market order flow. In this way, these orders placed at prices farther away from the bid-ask spread act as a liquidity back-stop, and are placed with strategies that differ from the ordinary provision of liquidity

Figure 4: Distribution of Limit Order Cancellations across Prices



The figure above is a histogram showing how far away from the bid-ask spread those limit orders that are canceled are at the time of cancellation, both for the in-sample period from January 2, 2015 through January 23, 2015, as well as the out-of-sample period from January 26, 2015 through January 30, 2015. The in-sample and out-of-sample distributions of limit order cancellations are very similar to one another.

The majority of all limit orders were canceled before being filled. Since order cancellations are commonplace, their distribution across locations in the limit order book closely follows that of limit orders (Figure 3).

This behavior is consistent with the anecdotal trend of “flickering orders,” in which liquidity providers try to learn about other market participants’ strategies by placing a limit order only to cancel it immediately afterward. Limit orders placed and canceled in such a way do not provide any meaningful liquidity to the market. Moreover, such orders are placed and canceled following a completely different strategy than that which governs the behavior of ordinary liquidity provisioning. The placement and cancellation of limit orders that are identified as flickering orders are therefore removed from the sample.

The limit order activity between six and ten ticks away from the spread is consistent with these orders being placed to serve as liquidity back-stops. These orders, which were never intended to provide liquidity close to the bid-ask spread, are placed with strategies that differ from the ordinary provision of liquidity, and are therefore canceled for different reasons than those that govern the cancellation of ordinary limit orders. For this reason, I restrict my sample to limit order activity that takes place within five ticks of the spread.

After excluding flickering orders from the sample, and restricting the sample to limit order activity within five ticks of the spread, the majority of the remaining limit orders are canceled before being filled. The prevalence of cancellations in the restricted sample is consistent with the possibility that market makers are responding to changing market conditions, canceling those orders which are either too far away from the spread to be useful or too aggressive to be profitable in expectation. In this way, the remaining limit order cancellations provide insight into the strategy governing market makers’ provision of liquidity.

Figure 5: Distribution of Divergent Limit Orders

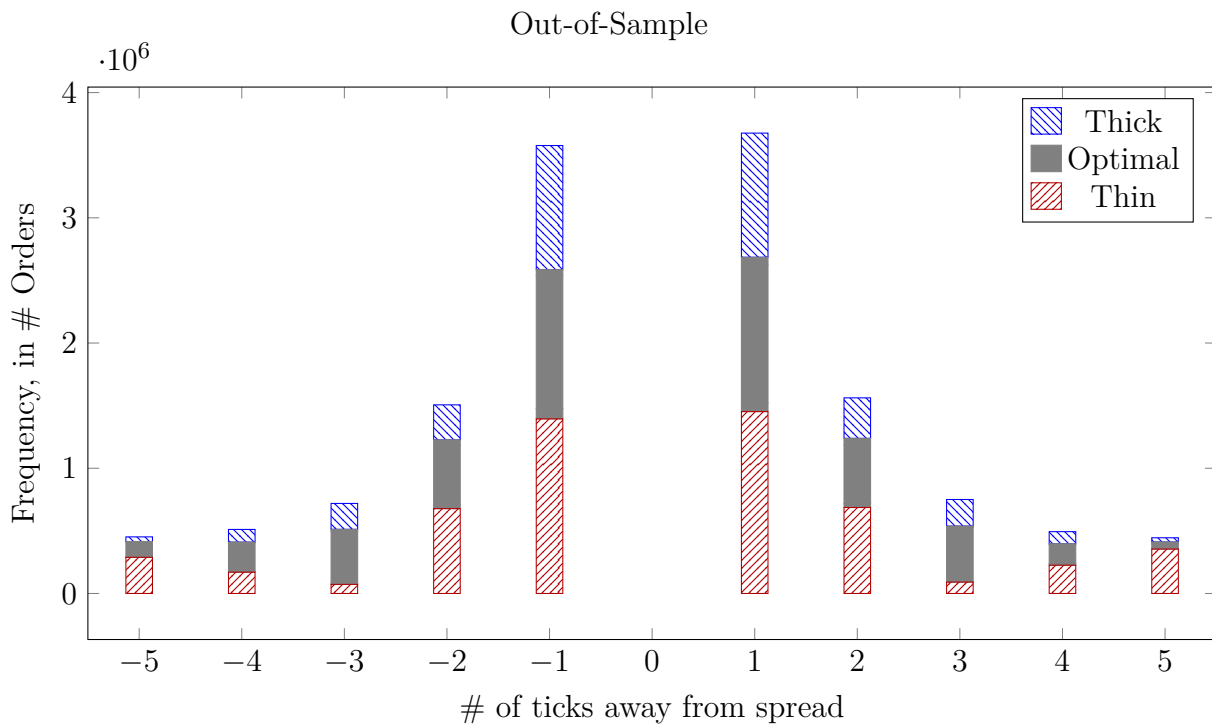
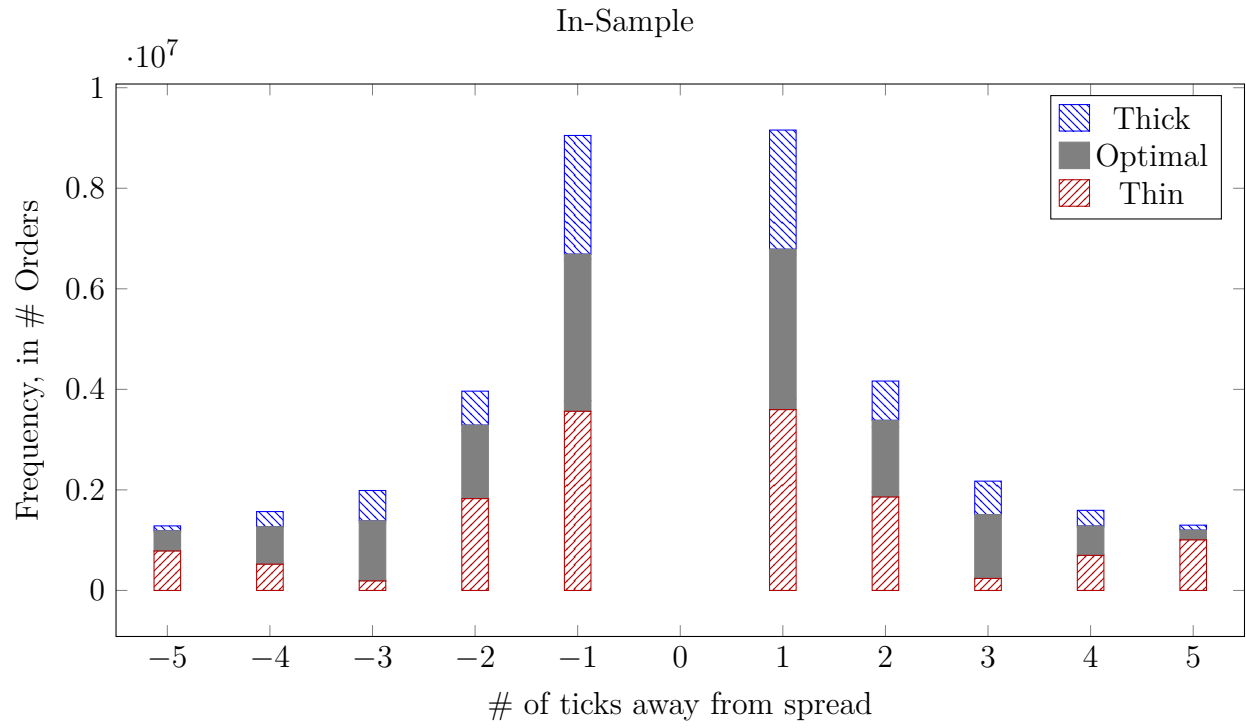


Figure 5 displays, for the ten stocks in both the in-sample period and the out-of-sample period, the relative frequency of limit orders which provide less liquidity than that which would be optimal on average, those which provide more liquidity, and those orders which are precisely the same size as the (static) model's prediction. The frequency of optimal limit orders, relatively large (thick) limit orders, and relatively small (thin) limit orders in the out-of-sample period are very similar to the frequencies in the in-sample period. Over the in-sample period, 38% of the orders are the same size of my model's estimated optimal size, 39% of orders provide less liquidity than my model predicts as optimal, and 23% of orders provide more liquidity. Over the out-of-sample period, 37% of orders are the same size as predicted, 40% provide less liquidity than predicted, and 24% provide more than predicted.

Figure 6: Distribution of Limit Orders by Optimal Price


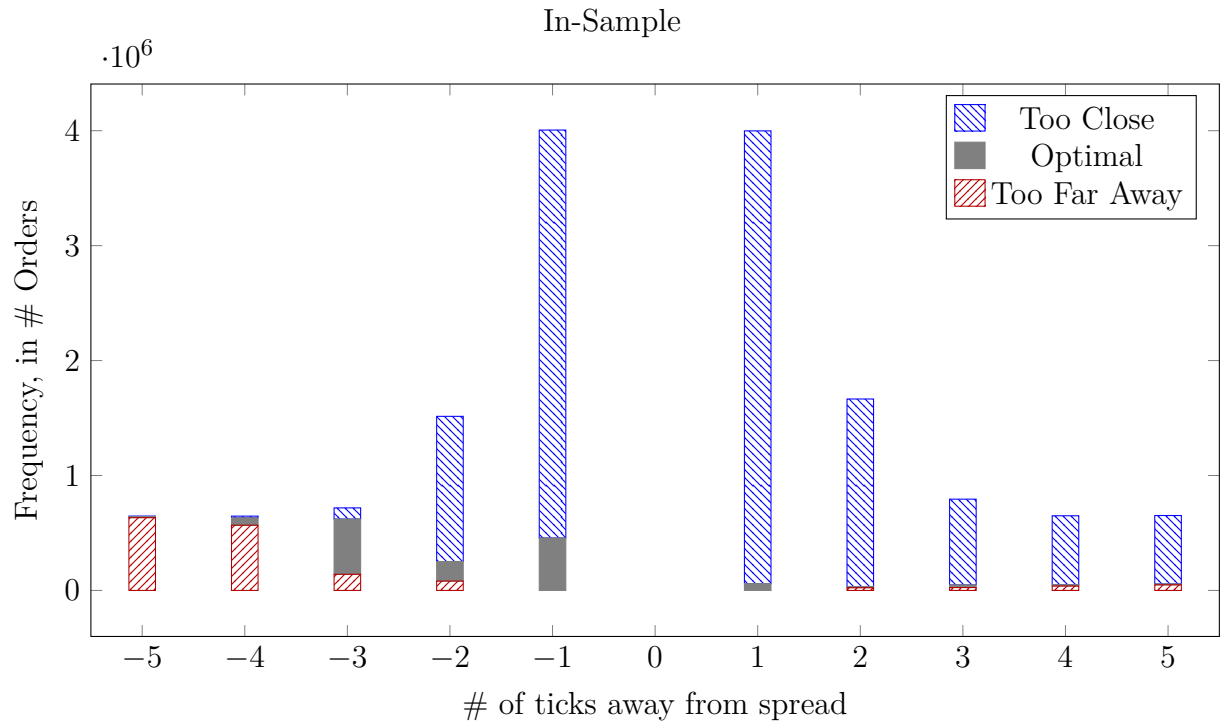




Figure 6 displays, for the ten stocks in both the in-sample period and the out-of-sample period, the relative frequency of limit orders which (according to the model) are placed at a price level that is too close to the bid-ask spread, at the most profitable price level, and a price level that is too near the spread. Over the in-sample period, 81% of the orders are placed at a price level too near the spread, 10% of the orders are placed at a price level too far away from the spread, and only 8% of the orders are placed at the model-projected optimal price level. The distribution of orders in the out-of sample period is very similar to the distribution of the in-sample period, with 8% of the orders being placed at the model-projected optimal price level, 83% of the orders placed too near the bid-ask spread, and 9% of the orders place too far away from the spread.

Figure 7: RLP^D Throughout Day - AAL - 1/30/15



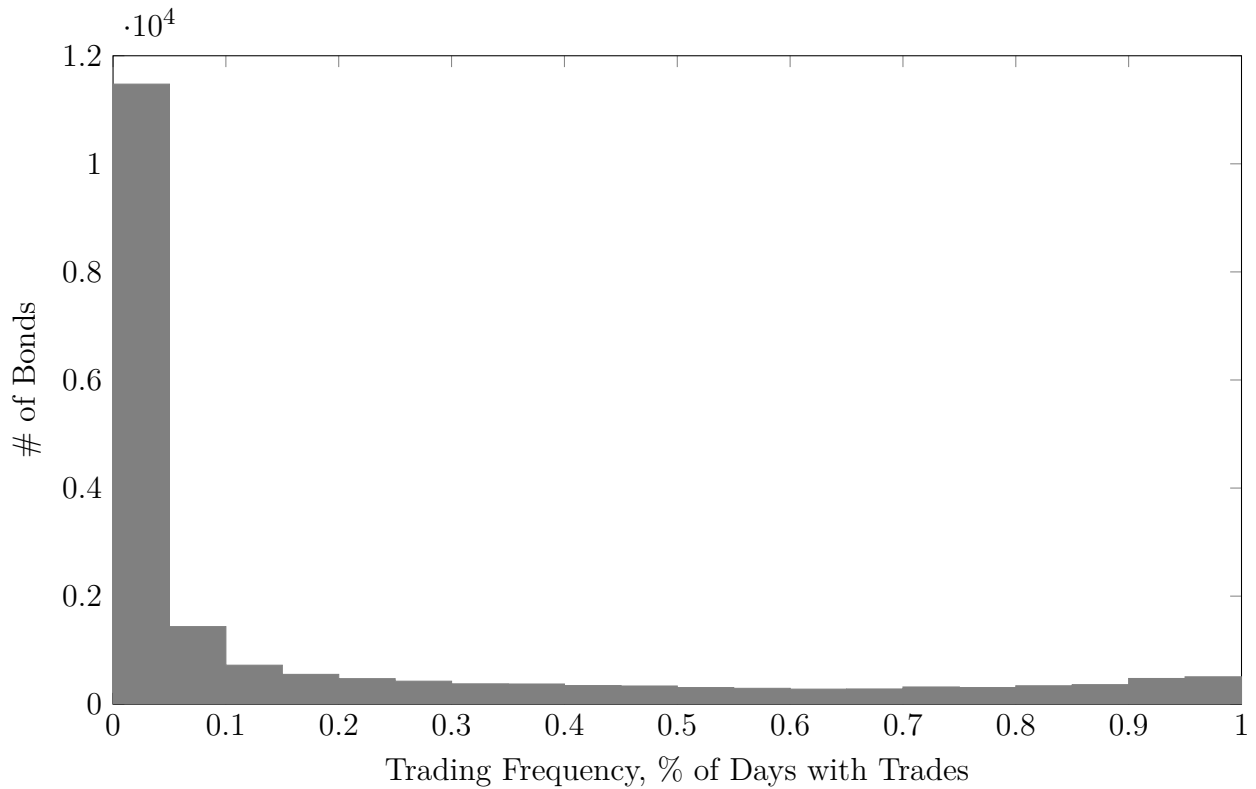
This figure shows the time series progression of the directional (RLP^D) estimator of relative liquidity provisioning throughout January 30, 2015 for American Airlines (AAL). This day and stock are representative of the patterns seen in the evolution of the directional estimator's value. The estimator is characterized by a process centered around zero that is punctuated by several short lived spikes in either direction and occasional prolonged periods of large values in either direction.

Figure 8: RLP^S Throughout Day - AAL - 1/30/15



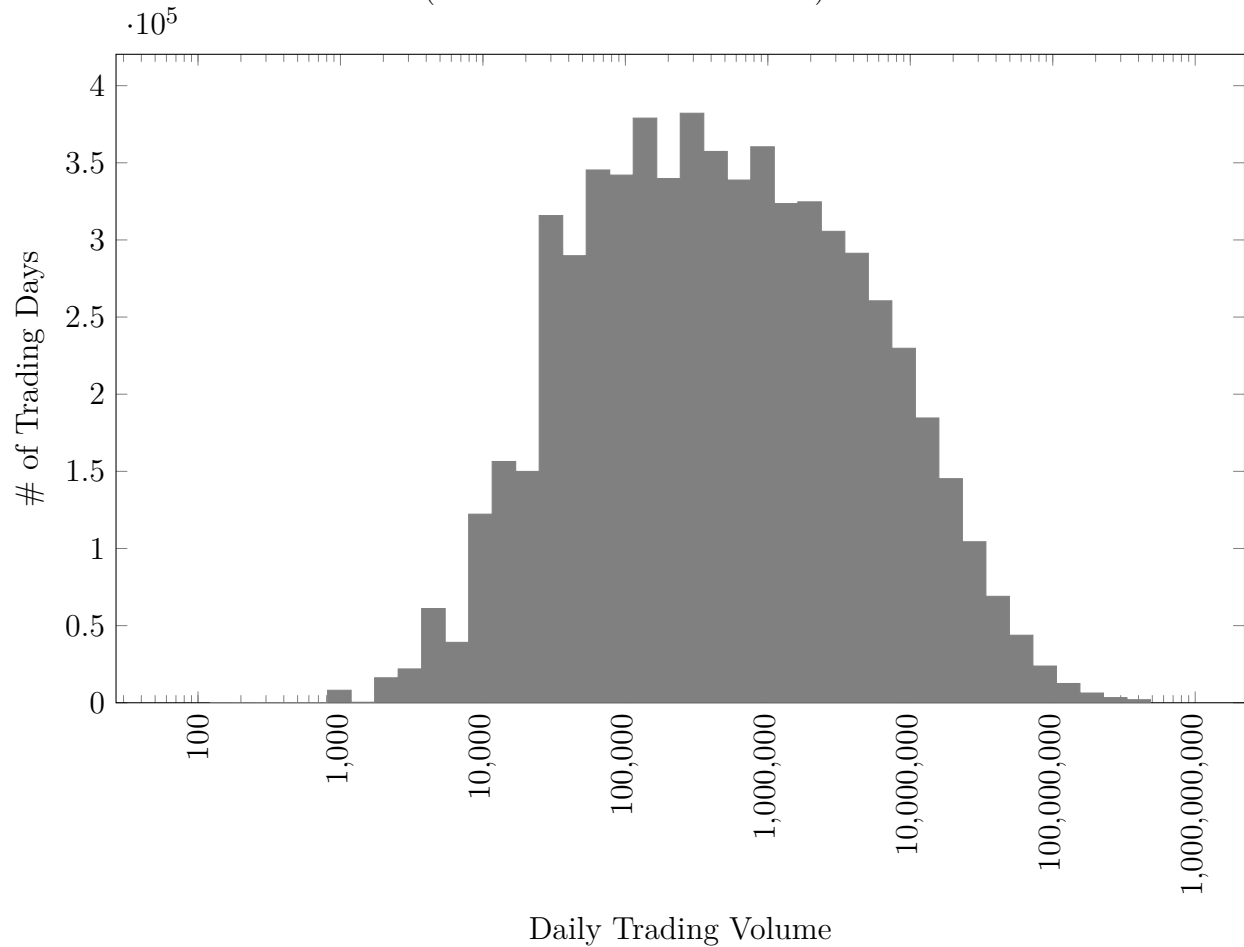
This figure shows the time series progression of the symmetric (RLP^S) estimator of relative liquidity provisioning throughout January 30, 2015 for American Airlines (AAL). This day and stock are representative of the patterns observed in the evolution of the symmetric estimator's value. The estimator appears to be centered above zero, with certain prolonged periods of large negative values.

Figure 9: Distribution of Trading Frequency Across Bonds




The above figure is a histogram depicting the distribution of trading frequency across bonds in my sample between January 1, 2006 and March 31, 2018. The distribution is heavily skewed downward; over half of the bonds in the sample have trades observed on fewer than 5% of the days during which they remain outstanding, with the number of bonds decreasing as the frequency of trading days increases. The distribution is also weakly bimodal, with a slight increase in the number of bonds observed as the frequency of trading days increases above 60%.

Figure 10: Distribution of Daily Trading Volume
(When Trades Are Observed)



The above figure is a histogram depicting the distribution of daily trading volumes (for days when trades are observed) across all dates and bonds in my sample between January 1, 2006 and March 31, 2018. The daily trade volumes follow an approximately log-normal distribution.

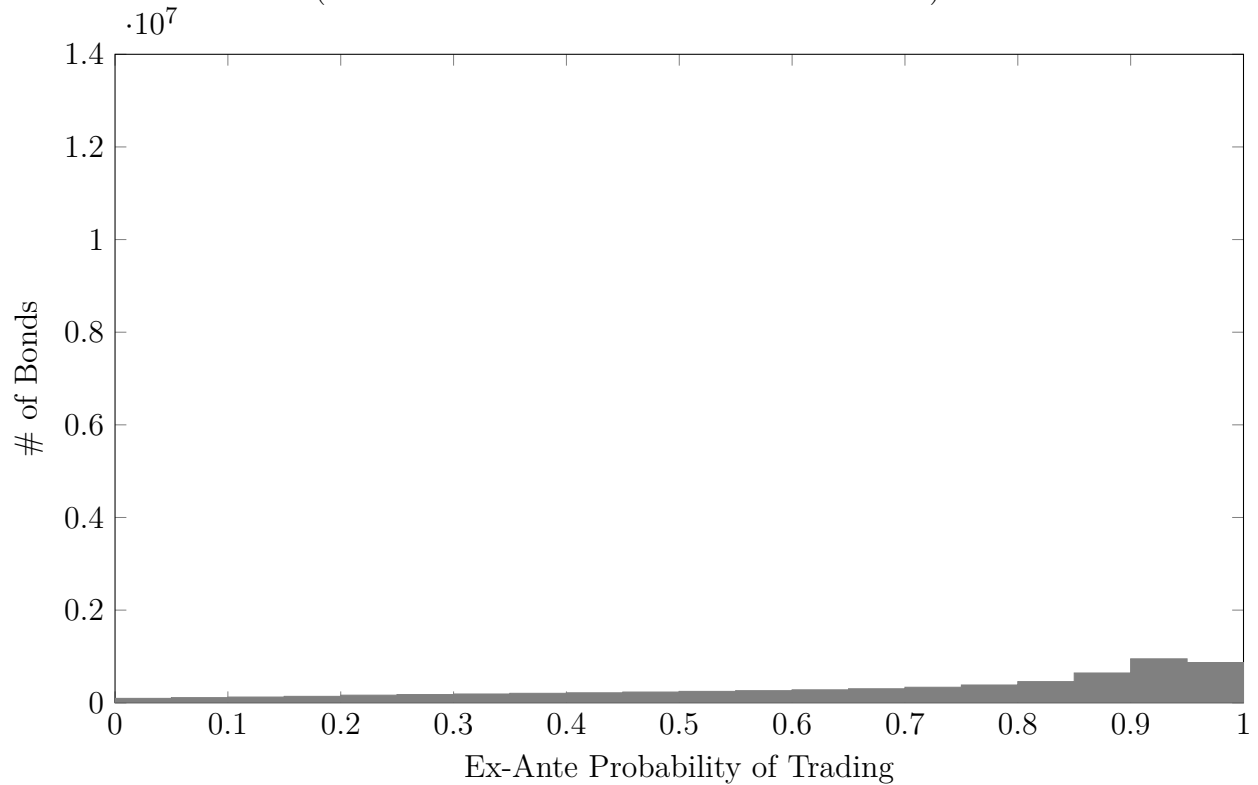
Figure 11: Distribution of Ex-Ante Trade Probabilities



The above figure is a histogram displaying the distribution of the ex-ante probabilities that a bond will trade generated by the hidden Markov model for bonds in my sample between January 1, 2006 and March 31, 2018. This distribution closely resembles the distribution of trading frequency shown in Figure 9.

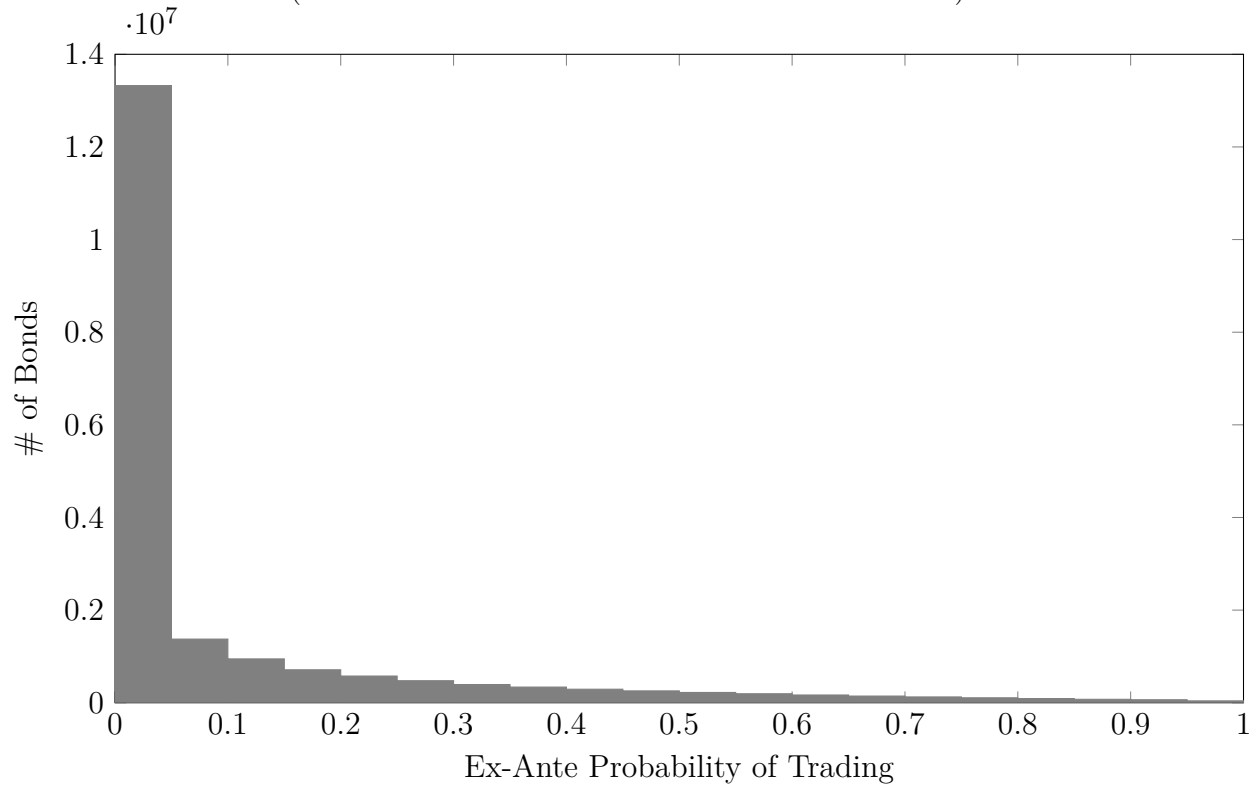
Figure 12: Distribution of Ex-Ante Trade Probabilities

(Observations for which Trades are Observed)



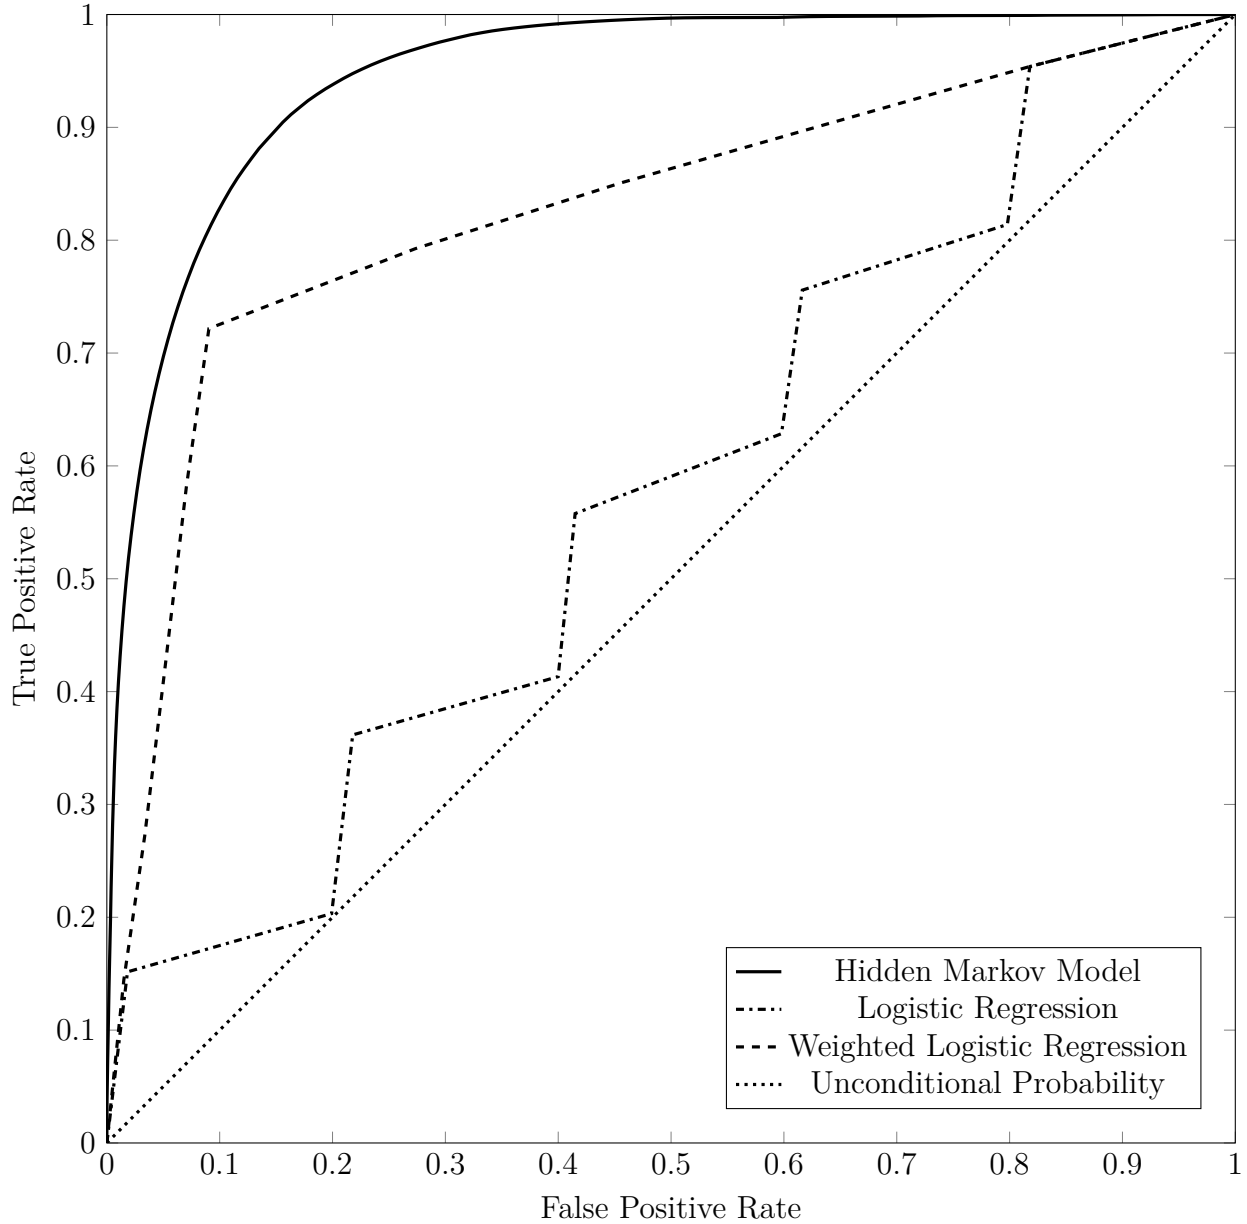
The above figure is a histogram displaying the distribution of the ex-ante probabilities that a bond will trade for observations for which a trade is observed. This decomposition of the distribution of trading probabilities shown in Figure 11 into the portion of the distribution corresponding to observed trades demonstrates the accuracy with which the hidden Markov model predicts trades.

Figure 13: Distribution of Ex-Ante Trade Probabilities
(Observations for which No Trades are Observed)




The above figure is a histogram displaying the distribution of the ex-ante probabilities that a bond will trade for observations for which no trades are observed. This decomposition of the distribution of trading probabilities shown in Figure 11 into the portion of the distribution corresponding to the absence of trades demonstrates the accuracy with which the hidden Markov model predicts that no trade will occur.

Figure 14: In-Sample ROC Curves for HMM and Alternative Models



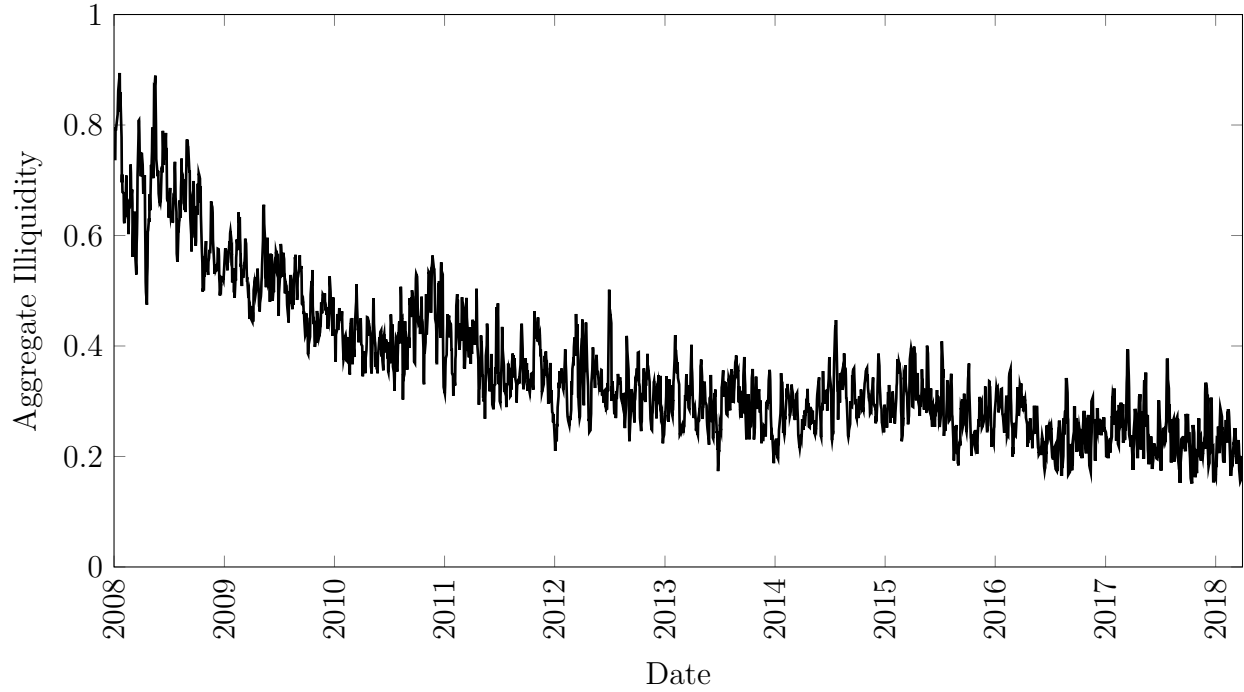
This figure shows the receiver operator characteristic (ROC) curves over the training sample for the hidden Markov model and the alternative trade probability models. The closer a model's ROC curve approaches to the upper left corner of the graph, the greater its ability to distinguish between the two classes. A random classifier model with no ability to distinguish the two classes from one another possesses a ROC curve that bifurcates the ROC space with a diagonal line from the bottom left to the top right of the graph.

Figure 15: Out-of-Sample ROC Curves for HMM and Alternative Models



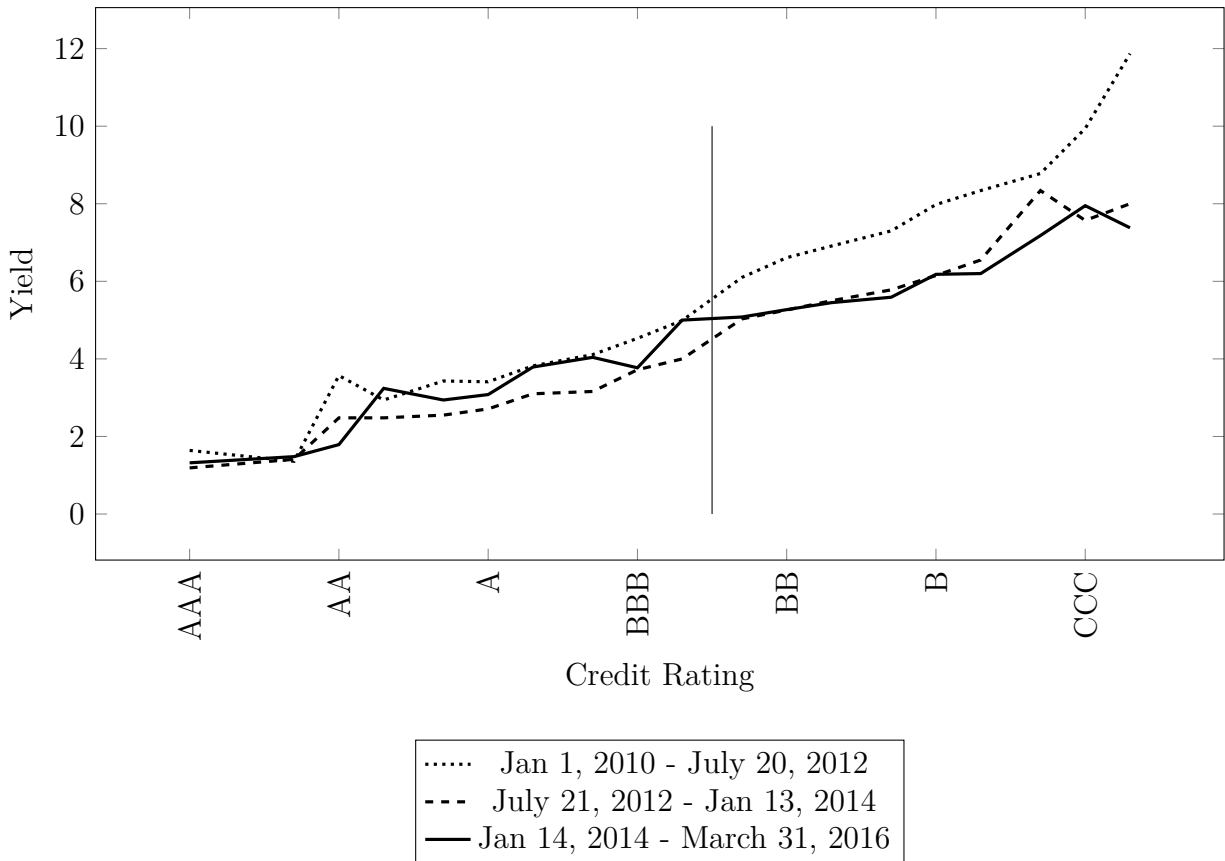
This figure shows the receiver operator characteristic (ROC) curves over the test sample for the hidden Markov model and the alternative trade probability models. The closer a model's ROC curve approaches to the upper left corner of the graph, the greater its ability to distinguish between the two classes. A random classifier model with no ability to distinguish the two classes from one another possesses a ROC curve that bifurcates the ROC space with a diagonal line from the bottom left to the top right of the graph.

Figure 16: Aggregate U.S. Corporate Bond Market Illiquidity



This figure shows the evolution of my measure of aggregate bond market illiquidity between January 5, 2008 and March 31 2018. There is a clear downward trend in the illiquidity of the overall bond market, indicating that the liquidity available to investors has been increasing steadily through the sample period.

Figure 17: Average Yield for Credit Rating at Time of Offer



This figure shows the average yield at the time of origination for each credit rating of U.S. corporate bonds issued between January 1, 2010 and March 31, 2016. Bonds issued between January 1, 2010 and July 20, 2012 were issued before the final Volcker rule regulations were announced. The bonds issued between July 21, 2012 and January 13, 2014 were issued after the final regulations were announced but before the Volcker rule was enforced, and bonds issued after January 14, 2014 were issued after the regulations began being enforced. The black vertical line indicates the threshold between investment grade credit ratings and junk ratings.