Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

The Relationship between Market Depth and Liquidity Fragility in the Treasury Market

Andrew Meldrum, Oleg Sokolinskiy

2025-014

Please cite this paper as:
Meldrum, Andrew, and Oleg Sokolinskiy (2025). "The Relationship between Market Depth and Liquidity Fragility in the Treasury Market," Finance and Economics Discussion Series 2025-014. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2025.014.

The Relationship between Market Depth and Liquidity Fragility in the Treasury Market*

Andrew Meldrum† Oleg Sokolinskiy‡

February 14, 2025

Abstract

Analysis of market liquidity often focuses on measures of the current cost of trading. However, investors and policy-makers also care about what would happen to liquidity in the event of an adverse shock. If liquidity were to deteriorate rapidly at times when investors were seeking to rebalance portfolios, this could amplify the effects of shocks to the financial system even if liquidity is high most of the time. We examine the potential for such fragility of liquidity in the Treasury market. We show that a reduction in the availability of resting orders to trade ("market depth") increases liquidity fragility, likely because lower depth increases the dependence of low trading costs on prompt replenishment of resting orders. Our results apply to all major benchmark Treasury securities individually, which enables us to establish analogous conclusions for market-wide liquidity fragility.

Keywords: liquidity, fragility, Treasury market, price impact, volatility, market depth, hidden Markov model

JEL codes: G01, G10, G12, C51, C58

*This article represents the views of the authors, and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or other members of its staff. We are grateful for helpful comments from and discussions with Matteo Aquilina, Dobrislav Dobrev, Michael Fleming, Yesol Huh, and Edith Liu, as well as participants at a Federal Reserve System seminar.

†Board of Governors of the Federal Reserve System. Email: andrew.c.meldrum@frb.gov

‡Board of Governors of the Federal Reserve System. Email: oleg.v.sokolinskiy@frb.gov

1 Introduction

This paper is concerned with fragility of Treasury market liquidity—that is, the possibility that while liquidity is generally good it could deteriorate in the face of shocks. Understanding the drivers of Treasury market liquidity is important to market participants and policymakers alike. In normal circumstances, the market is highly liquid, meaning that investors can trade easily and cheaply in large volumes. However, investors are not only concerned with the current liquidity of the market; what also matters is how liquid the market would be at times when they need to adjust positions. For example, if liquidity would deteriorate at precisely the time when an investor would need to sell a security, potentially amplifying movements in its price, that security may carry a liquidity premium even if the market is liquid most of the time. Investors therefore need to understand the fragility of liquidity. Potentially fragile liquidity in response to shocks may also present a financial stability concern if it could cause the financial system to amplify the effects of shocks.

In this paper, we show that while a reduction in the volume of quotes posted to electronic trading platforms (commonly known as "market depth") does not necessarily mean that trading costs increase in normal circumstances, it does mean that the fragility of liquidity—the probability of a sudden increase in trading costs—increases. In recent years market depth has declined. This trend does not necessarily mean that the market is less liquid in normal times; in fact, the market can support higher volume of trading at low execution cost with a lower level of depth than in the past as a consequence of two things. First, market participants now employ more sophisticated execution algorithms, splitting large orders into multiple smaller orders over time to reduce their price impact. Second, high-speed liquidity providers can more quickly replenish quotes on the order book in response to incoming orders. In this environment, the continued ability and willingness of market makers to replenish orders at the best prices, rather than the level of quotes posted to the order book at any one time, has become the key to the cost of trading remaining low.

The microstructure of the Treasury market is key to understanding what drives its liquidity because the microstructure defines how market participants interact and trade terms are established. We study on-the-run Treasuries traded electronically via inter-dealer brokers (IDB)—a sector capturing a large proportion of traded volumes in the most liquid, benchmark securities. Specifically, we consider trades and quotes from the IDB with the largest market share, the BrokerTec Alternative Trading System (ATS). BrokerTec matches orders following a central limit order book (CLOB) protocol: Market participants can provide liquidity by submitting buy and sell limit orders with associated prices and quantities that cannot be immediately matched (e.g., a limit buy order with an associated price below the best ask) and are posted to the CLOB. Market participants can also consume liquidity by submitting "marketable" limit orders that can be immediately matched (e.g., a limit buy order with an associated price at or above the best ask), thereby reducing the amount of liquidity remaining on the book. Market depth at the best prices is the sum of posted volumes over all orders at the best bid and ask prices in the CLOB. The greater the market depth, the larger the transaction size that can be executed without immediate mechanical impact on the price.

Market makers provide liquidity by posting both buy and sell orders at various levels of the CLOB, seeking to profit from a bid-ask spread. To reduce their exposure to informed traders, modern market makers prefer to post modest order sizes and then have the option to replenish the order book. In turn, market participants avoid submitting large orders that would fully exhaust the limit orders available at the best prices, with the expectation that market makers will respond by replenishing the order book, allowing execution of further small marketable orders at preferable prices.

An increase in uncertainty about future prices causes market makers' risk of trading against informed traders to increase. Market makers will therefore respond by decreasing the amount of orders they post to the order book and potentially by being more cautious about replenishing orders. Such a reduction in depth does not necessarily mean that the cost of trading increases, providing the speed with which they replenish quotes remains sufficient to meet incoming order flow without large price moves. However, we show that it does increase the risk of trading costs increasing, presumably because there is a greater chance that quote replenishment will be inadequate to meet incoming order flow.

A key challenge when modeling liquidity is its latent nature. While liquidity may appear observable in a CLOB-based market (because market participants observe how much is available to instantaneously trade at various price levels), the willingness of market makers to replenish quotes at the best prices is unobserved. We use the natural framework of a Hidden Markov model (HMM) to capture this latent nature of liquidity. In our HMM, we capture the fragility of liquidity as the endogenous probability of a transition to a bad liquidity state. A worse liquidity state is defined by price impact being more sensitive to certain conditioning factors, signifying lower willingness of market makers to replenish quotes given the same degree of adverse changes to those factors. Most importantly, we allow price impact to be inversely related to interest rate volatility, as explored by Bouchaud et al. (2009) among others. As mentioned above, the intuition is that at times of high volatility, the risks to market makers associated with adverse price movements are high, so market makers will tend to provide less liquidity and charge more for the liquidity they provide. We extend the set of variables driving price impact to encompass uncertainty about volatility, since this adds to tail risk, also increasing the cost of market making. We also consider volatility persistence, as this factor may increase inventory management costs for traditional dealers. Finally, we allow the transition probabilities between liquidity states to depend on market depth, which is fundamental for explaining liquidity fragility.

Our paper extends the existing literature in a number of directions. First, we consider our contributions to the financial stability literature. While the level of liquidity metrics is commonly seen as a financial stability indicator (see Board of Governors of the Federal Reserve System (2020) and O'Hara (2004) among others), we take a more forward-looking approach and treat the probability of a sudden liquidity deterioration as the more relevant indicator, since it is arguably more relevant to how much the financial system amplifies shocks. Our approach is related to Cespa and Vives (2017), who study liquidity fragility within a theoretical framework that allows for endogenous dynamics leading to liquidity fragility. Cespa and Foucault (2014) and Raman et al. (2020) explore the fragility that arises from illiquidity spillovers between assets. In contrast, we seek for a predictor of liquidity fragility that is observable and not a spillover from another market. Our finding that low market depth raises the probability of a liquidity deterioration has an intuitive interpretation because the lower depth is, the more rapidly the order book must be replenished to support a given flow of marketable orders without moving prices substantially.

Next we consider our contributions to the econometric modeling of liquidity, specifically, in how our HMM addresses the latent nature of liquidity and the non-linear relationship between price impact and factors driving the cost of liquidity provision. First, our approach is related to that of Flood et al. (2016), who link liquidity states identified using an HMM to observed economic variables. However, our framework integrates explanatory variables directly into the HMM, both in observation and state transition equations. Second, Duffie et al. (2023) explore the non-linear relationship between liquidity and volatility via quantile regression analysis. Our HMM framework allows for a non-linear relationship between price impact and volatility through different conditional price impact distributions in endogenously determined latent liquidity states, whereas a quantile regression may be seen as requiring states to correspond to the quantiles of liquidity measures. Thus, an HMM allows for endogenous state dynamics with the introduction of covariates into the state transition specification, which proves particularly beneficial for studying liquidity fragility. Third, our study is related to Hautsch and Huang (2012a), who show a link between price impact and depth in that the price impact of a limit order depends on how deep in the central limit order book (CLOB) it rests; the results of Nguyen et al. (2020) and Aronovich et al. (2021) similarly point to a rich dependence structure between liquidity, volume, and depth. Fourth, our study complements the results of Nguyen et al. (2020), who develop a model for the interaction of liquidity and volatility at high-frequencies. They find that market depth and trading volume significantly affect intraday volatility. We show that at a trading-day time scale lower market depth leads to future states characterized by higher price impacts of trade flows and, thus, all else equal, greater volatility. Finally, Fraenkle et al. (2011) allow the slope and intercept in price impact regressions to depend on explanatory variables such as trading volume and volatility, which highlights how other liquidity measures can directly affect price impact, but they do not consider market depth.

The remainder of this paper proceeds as follows. In Section 2, we motivate our econometric framework with exploratory analysis of the non-linear association between price impact and volatility, as well as the relationship between price impact and depth.[1] In Section 3, we introduce our liquidity modeling framework and apply it to the 10-year benchmark Treasury security. In Section 4, we first reaffirm our results for other benchmark Treasury securities through HMMs with security-specific latent liquidity states. Then, we extend our results to market-wide liquidity fragility in the framework of an HMM with latent liquidity states that are shared across Treasury securities. In Section 5, we offer some concluding remarks.

2 Motivating an HMM

In this section, we extend the literature relating liquidity to volatility and establish the rationale for modeling the relationship using a nonlinear HMM framework. In Section 2.1, we show that the effect of a given change in volatility on price impact increases with the level of volatility. This finding justifies a nonlinear relationship between price impact and volatility, which can be accommodated within the HMM framework, where the effect of volatility on price impact can vary between the latent liquidity states. In Section 2.2, we provide intuition for why market depth can serve as an indicator of the fragility of liquidity, which motivates the inclusion of market depth as a variable that can affect the probability of transitioning between different liquidity states in the HMM.

[1] In Appendix A, we set out our model for estimating price impact.

2.1 Price Impact and Volatility

Liquidity is generally worse at times of high volatility, as noted by Chordia et al. (2005), among others. Intuitively, market intermediaries are less willing to provide liquidity when the risk of large price moves is relatively high. However, the precise form of the relationship has not been explored in detail. We show that each unit of volatility affects liquidity more when volatility is high. Specifically, we consider a smooth threshold regression specification

$$\theta_t = \beta_1 V_t + (\beta_2 - \beta_1) G(V_t, \kappa_1) V_t + (\beta_3 - \beta_2) G(V_t, \kappa_2) V_t + \epsilon_{i,t}, \tag{1}$$

Here, θ_t is price impact on day t. On the right-hand side of Equation (1), V_t is a measure of the swaption-implied volatility; $G(V_t, \kappa_i) := \frac{1}{1+\exp^{-(V_t - \kappa_i)}}$ is the first-order logistic smooth transition function, with κ_1 and κ_2 being threshold levels of volatility. A two-threshold model is statistically preferred to a linear specification and is sufficient for the purpose of establishing the non-linearity of price impact in volatility to motivate the introduction of an HMM in Section 3.

We estimate Equation (1) using our estimate of price impact for the 10-year Treasury Note. Specifically, it is the incremental price move associated with purchasing $500 million of the security using marketable orders spread over the trading day split across 1-minute periods in proportion to the total volumes in each of those periods. This estimate of price impact is conventional in that it comes from a regression of price changes on trade and order flows, estimated separately for each day using intraday data. However, we introduce the novelty of allowing for non-linearity in trade and order flows. Appendix A provides further details on how we estimate price impact.[2] As the measure of volatility corresponding to the 10-year Treasury Note, we use the annualized basis point implied volatility obtained from end-of-day observations of the 1-month-ahead option contract on the 10-year swap,

[2] Appendix A also shows that price impacts associated with alternative execution strategies are strongly positively correlated over time, suggesting that we would obtain qualitatively robust inferences about the overall dynamics of market liquidity by focusing on our baseline execution strategy.

obtained from TP ICAP. The sample period runs from April 1, 2014 to December 31, 2023. We omit days with shortened trading hours, leaving a total of 2,374 days. Table 1 reports β_i $(i = 1, 2, 3)$ estimates for the optimal volatility thresholds. The estimated thresholds are at 94 and 147 basis points. The slope coefficients are all positive and statistically significant, and are significantly greater for higher levels of volatility than for lower levels. We conclude that this result suggests that any model relating volatility to price impact should allow for a nonlinear relationship in which the sensitivity is greater as volatility increases.

Table 1: Threshold regression of price impact on volatility

The table reports estimates of the threshold regression model of Equation (1). Numbers in brackets indicate standard errors. *** indicates two-sided p-values of less than 0.001.

	Model 1
β_1	0.097***
	(0.001)
β_2	0.137***
	(0.001)
β_3	0.196***
	(0.004)
κ_1	93.882***
	(0.744)
κ_2	146.806***
	(0.749)
R^2	0.897

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

2.2 Price Impact and Market Depth

Figure 1 shows daily time series of price impact and market depth. We compute depth as the time-weighted mean of the total posted quantities at the best bid and ask prices between 7 a.m. and 4.45 p.m. (with weights determined by the lengths of time between consecutive market events – trades or cancellations). The upper two panels both show price impact but using different vertical axis scales, while the lower panel shows market depth. The sample correlation between price impact and market depth is fairly negative, at -0.75, which is intuitive because falling depth and rising price impact are both associated with a deterioration in liquidity conditions. The period following the onset of the COVID-19 pandemic in March 2020 stands out as seeing by far the highest values of price impact, as well as the largest decline in depth. There were also notable spikes in price impact coinciding with other days of local troughs in depth identified by Aronovich et al. (2021), as indicated by the vertical dashed lines: the Treasury market "flash crash" of October 15, 2014, where prices moved sharply before reversing quickly; the U.K. "Brexit" referendum on June 24, 2016; the sharp spike in the VIX index of equity market volatility on February 6, 2018; the onset of concerns among financial market investors about the global growth outlook on August 2, 2019; a sharp worsening in the market strains associated with the COVID-19 pandemic on March 9, 2021; and the flash event of February 25, 2021. In addition, we pick out the Monday following the closure of Silicon Valley Bank, that is, March 13, 2023. There were also less sharp increases in price impact and declines in depth around early 2015 and late 2016, which were likely associated with spillovers from stresses in European sovereign debt markets and the notable increase in yields following U.S. elections, respectively. And depth has declined and price impact has risen notably since late 2021, which market commentary has linked to high interest rate volatility associated with uncertainty about the economic outlook.

Although peaks in price impact tend to coincide with troughs in depth, price impact and depth appear to have substantially different time-series dynamics. For example, the sample autocorrelation functions in Figure 2 show that depth is substantially more persistent than price impact. We also observe substantial differences between the dynamics of market depth and price impact following the onset of episodes of market stress. Figure 3 shows depth and price impact from 10 days before to 15 days after the seven events highlighted in Figure 1. To facilitate comparison across episodes, we normalize all variables to have the value of unity at the beginning of the event window. Depth recovered relative quickly following the Brexit referendum on June 24, 2016, and the VIX spike on February 6, 2018, although it still took 2 or 3 weeks before depth had retraced the large majority of the initial decline. In comparison, the recoveries following the flash crash of October 15, 2014, the August 2, 2019 global growth concerns, the February 25, 2021 flash event, and onset of the March 2023 banking-sector stresses were slower. And the recovery following the most severe episode of market dysfunction, the onset of the

Figure 1: Price Impact and Market Depth

The charts show times series of price impact (upper and center panels) and market depth (lower panel) for the 10-year Treasury Note. The upper and center panels plot the same estimates but using different y-axis scales. The vertical dashed lines correspond to: (a) October 15, 2014; (b) June 24, 2016; (c) February 6, 2018; (d) August 2, 2019; (e) March 9, 2020; (f) February 25, 2021; and (g) March 10, 2023.

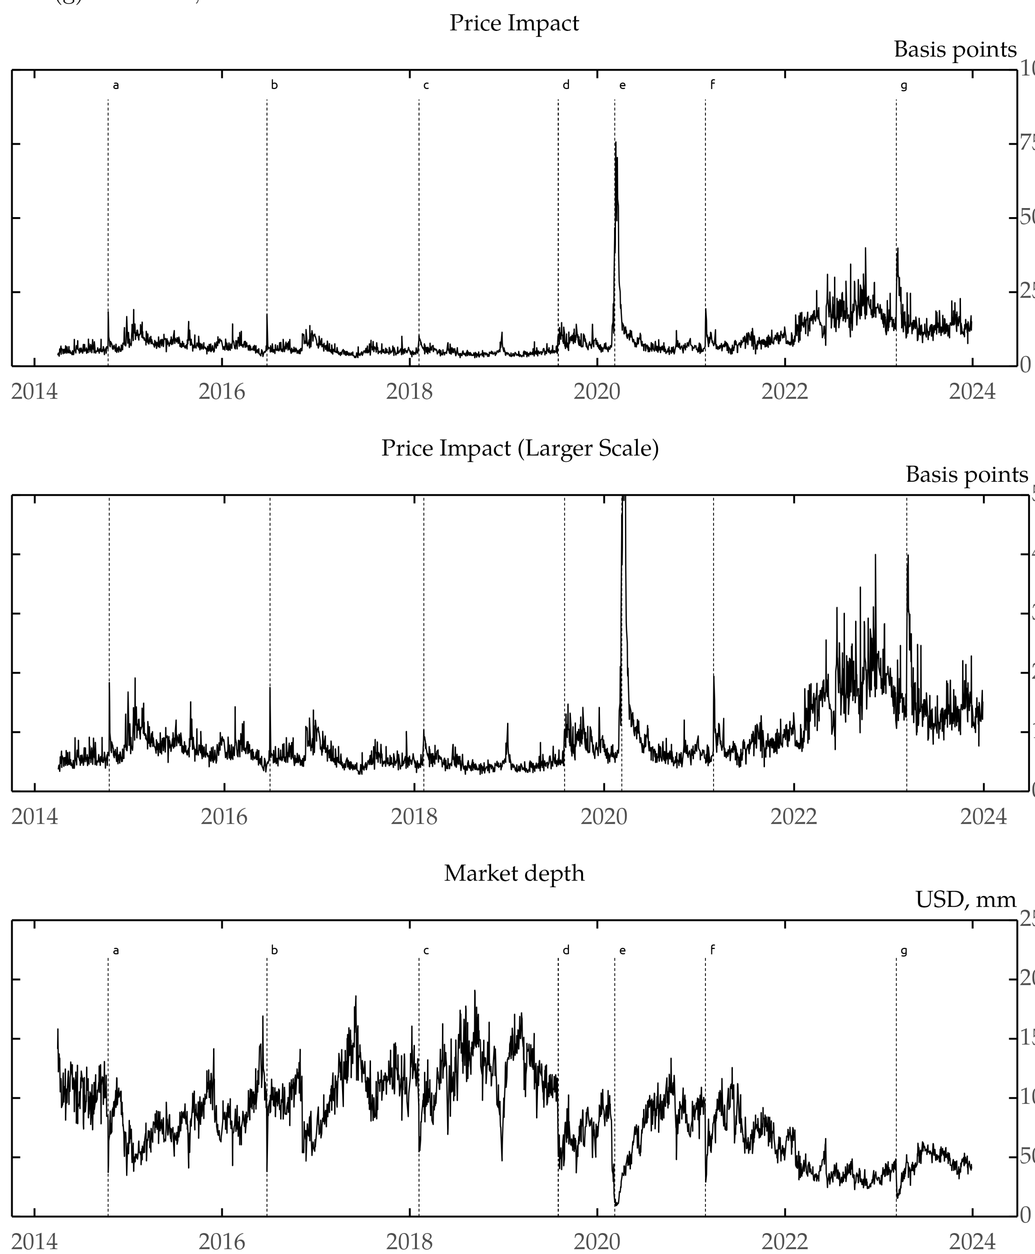

Sources: Repo Inter Dealer Broker Community; authors' calculations.

COVID-19 pandemic in March 2020 took substantially longer still. Price impact rose immediately at the onset of all episodes. In some episodes, price impact fell back relatively rapidly compared with the slower recovery of depth; this observation is consistent with Aronovich et al. (2021), who show that bid-ask spreads tend to recover substantially faster than market depth following these episodes. A plausible explanation for this is that market participants rapidly adjusted to lower depth by trading in smaller sizes to allow quotes on the order book to be replenished. In other cases—most notably March 2020 and March 2023—the peak of price impact did not come until several days after the trough in depth, suggesting that the low level of depth may increase the risk of a spike in price impact.

Figure 2: Autocorrelation Functions of Price Impact and Market Depth

The bars show the sample autocorrelations of market depth and our baseline estimate of price impact from March 1, 2014 to March 31, 2023 at various lags.

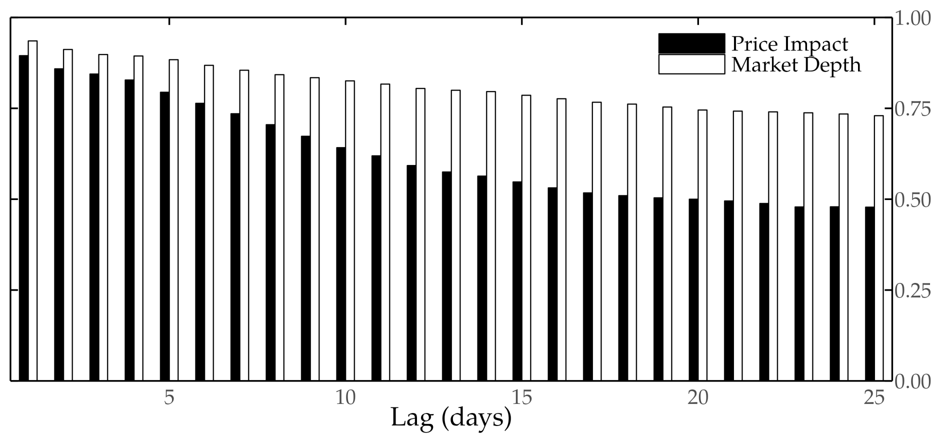

Sources: Repo Inter Dealer Broker Community; authors' calculations.

These observations raise questions of how market depth and price impact are related, and whether it matters that depth remains low after a stress event even though price impact has returned to more normal levels. Given market participants' unwillingness to trade through multiple levels of the book, as noted in Hautsch and Huang (2012a), a decrease in top-of-the-book depth causes optimal execution to prescribe splitting the parent order into smaller child orders. Hypothetically, market

Figure 3: Price Impact and Market Depth following Episodes of Known Liquidity Strains

The charts show market depth and price impact from 10 days before until 15 days after various episodes of known liquidity strains. All series are normalized to have value one 10 days before the event.

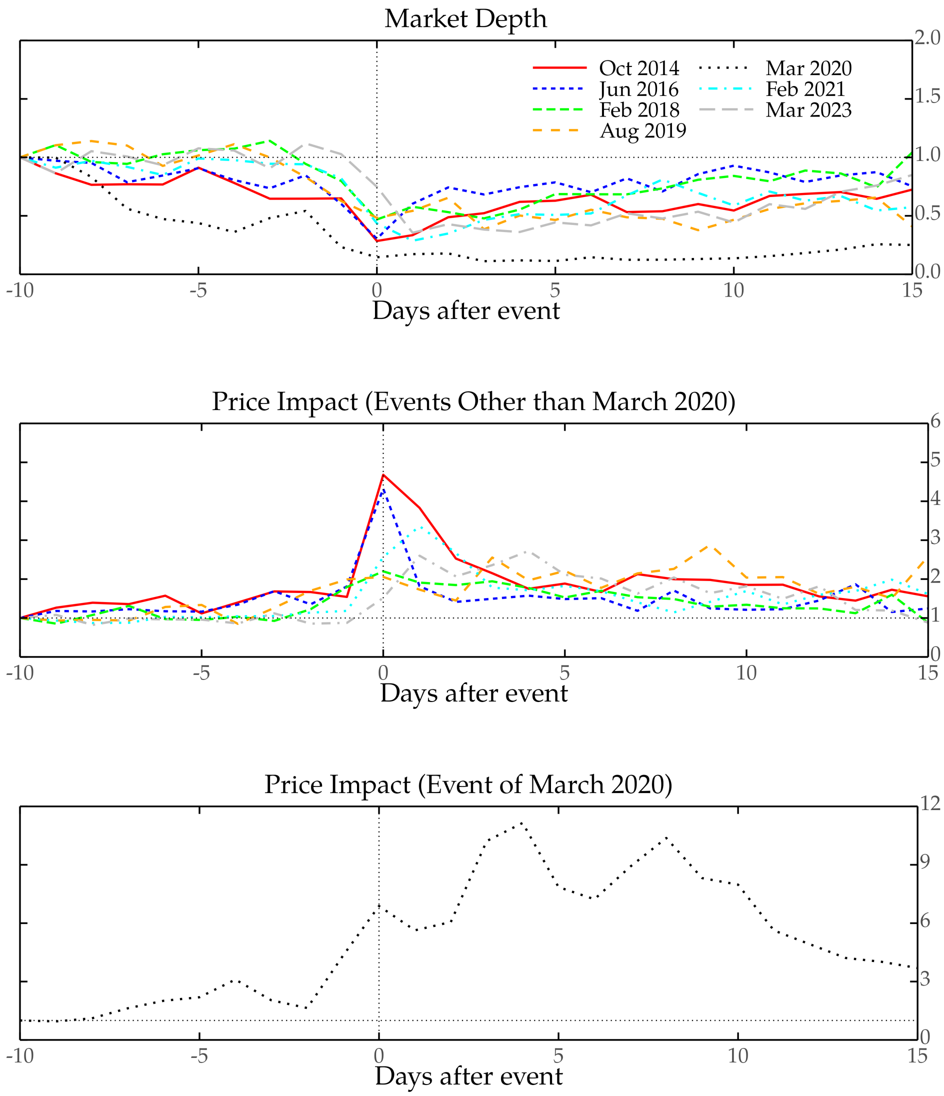

Sources: Repo Inter Dealer Broker Community; authors' calculations.

makers may replenish the order book sufficiently rapidly to result in either no or moderate increase in the price impact of execution; the point that prompt order-book replenishment matters for the cost of trading has been made previously by Dobrev and Meldrum (2020) and Aronovich et al. (2021). The immediate reduction in depth following a market stress event is often associated with a pull-back in high-speed liquidity provision, but this liquidity provision returns faster than the willingness of market participants to quote in large size; hence, depth remains low but price impact falls back.

One tempting interpretation of these results is that depth is not a meaningful indicator of liquidity conditions, because high-speed liquidity replenishment means that trading costs can be kept low even though most of the willingness of market participants' to provide liquidity remains hidden. However, a conjecture broadly in line with Raman et al. (2014) and Board of Governors of the Federal Reserve System (2020) is that greater dependence on sophisticated quoting and execution algorithms in such circumstances may nevertheless make liquidity more *fragile*, meaning that low depth makes it more likely that price impact could increase again if there is a further pull-back in high-speed liquidity provision—perhaps because high-frequency trading firms may be inclined to scale back their activity if some other stress event hits or because elevated trading volumes will overwhelm the speed of quote replenishment. Indeed, as mentioned above, some of the spikes in price impact following the events discussed above did not come until after sometime following the initial trough in depth. Our HMM is natural framework for studying more systematically whether the low level of depth is associated with liquidity becoming more fragile, meaning an elevated risk of price impact rising.

3 Hidden Markov Model for Latent Liquidity States

Having provided intuition for the appropriateness of the Hidden Markov Model (HMM) framework, we apply it to modeling the liquidity of the 10-year Treasury Note. In Section 3.1, we set out the observation equation component of the general

HMM framework. In Section 3.2, we consider an HMM with constant probabilities of transitioning between the latent liquidity states. This model demonstrates how the sensitivity of price impact to volatility—as well as some other potential explanatory variables—varies over different liquidity states. Our results show that the sensitivities of price impact to volatility, the surprise component of uncertainty about volatility, as well as the surprise component of volatility persistence, all vary over time with the liquidity state. In Section, 3.3, we extend the HMM to allow the transition probabilities to vary over time as a function of market depth. Our results provide clear evidence that the fragility of liquidity decreases with market depth, that is, greater market depth reduces the probabilities of transitions to worse liquidity states.

3.1 Hidden Markov Model: Observation Equation

Our HMM explains price impact, θ, via a state-dependent linear function of covariates X_t:

$$\theta_t = \beta_i X_t + \epsilon_{i,t}, \tag{2}$$

where $\epsilon_{i,t} \sim \mathcal{N}(0, \sigma_i^2)$ for three latent liquidity states $i = 1, 2, 3$.[3]

We estimate various versions of the HMM with different vectors of covariates X_t. However, we always include a measure of interest rate volatility. This choice is motivated by the analysis in Section 2.1, which demonstrates that price impact is positively related to interest rate volatility and that price impact may become more sensitive to volatility as volatility increases. We label the states such that the slope coefficients with respect to interest rate volatility decrease with the state index in the measurement equation, Equation (2); thus, state 1 is a "low-liquidity" state in which price impact is most sensitive to volatility, state 2 is a "medium-liquidity" state, and state 3 is a "high-liquidity" state.

The characterization of hidden liquidity states based on factor sensitivities corresponds well to interpreting latent liquidity in terms of the willingness of market

[3] In the models we consider, Bayesian and Akaike information criteria favor models with three states to models with two states.

makers to replenish the order book. Specifically, price impact is more sensitive to volatility exactly when market makers reduce their replenishment speed more per unit increase of volatility. The price impact of trade flow is by its very definition a measure of the market participants' willingness to replenish the order book – market price changes in response to trades when the queues at the best prices are not replenished or resting orders are cancelled.[4] Thus, price impact is both a measure of execution costs and the market makers' willingness to replenish the order book.

3.2 Hidden Markov Model: Constant Transition Probabilities Matrix

We start our investigation with an HMM characterized by constant probabilities of transitioning between the states:

$$p_{i,j,t} = p_{i,j}, \tag{3}$$

where $p_{i,j,t}$ denotes the trading day-t probability that $s_{t+1} = j$ conditional on $s_t = i$, and is a fixed model parameter. We estimate this model (and others considered in this paper) by maximum likelihood, using the method of Visser and Speekenbrink (2010).

Table 3 reports the resulting parameter estimates. We first consider the observation equation parameter estimates in the four considered specifications, labeled Models I through IV. The observation equation of Model I includes only the volatility of rates as an explanatory variable. Here, implied volatility from a tenor-matched (that is, 10-year) swaption contract with 1-month expiration is our measure of interest rate volatility. The slope coefficients, β_i^V on volatility are all positive and highly significant, meaning that higher volatility is associated with greater price impact in all states. Moreover, the differences between volatility coefficients among states are also large and significant, in line with the hypothesis that price impact is more sensitive to volatility when liquidity is low.

[4]Most trades do not take more liquidity than is available at the best price.

Table 2: HMM with Fixed Transition Probabilities Matrix: Parameter Estimates

The table reports parameter estimates for the HMMs described by Equation (2) and fixed transition probability matrix that explains the 10-year Treasury Note liquidity. The sample period is from April 1, 2014 to December 31, 2023. β^{V} coefficients corresponds to the effect of swaption implied volatility, β^{VV} – to the effect of the surprise component of the volatility of swaption implied volatility, and β^{CS} – to the effect of surprise component of volatility persistence. Sub-indices correspond to states, $i = 1, 2, 3$, ranked from lowest to highest liquidity. Panel $\mathcal{B}$ contains estimates of the transition equation parameters: $p_{i,j}$ is the probability of transitioning from latent liquidity state i to state j. Numbers in brackets indicate standard errors based on the Hessian. A $\cdot$, *, or ** indicates two-sided p-values of less than 0.1, 0.05, and 0.01, respectively.

	Model I	Model II	Model III	Model IV
Panel $\mathcal{A}$.	**Observation equation**			
β_1^V	0.191**	0.163**	0.185**	0.163**
	(0.007)	(0.003)	(0.005)	(0.003)
β_1^{VV}		1.92**		1.829**
		(0.123)		(0.131)
β_1^{CS}			-0.715**	-0.096
			(0.092)	(0.069)
σ_1	9.317**	5.575**	7.804**	5.607**
	(0.547)	(0.259)	(0.453)	(0.264)
β_2^V	0.12**	0.115**	0.12**	0.117**
	(0.001)	(0.001)	(0.001)	(0.001)
β_2^{VV}		0.633**		0.697**
		(0.067)		(0.057)
β_2^{CS}			0.011	0.057**
			(0.017)	(0.018)
σ_2	1.886**	1.631**	1.883**	1.639**
	(0.061)	(0.05)	(0.059)	(0.054)
β_3^V	0.084**	0.081**	0.084**	0.083**
	(0.001)	(0.001)	(0.001)	(0.001)
β_3^{VV}		0.194**		0.285**
		(0.048)		(0.044)
β_3^{CS}			0.102**	0.106**
			(0.01)	(0.01)
σ_3	1.029**	0.918**	1.003**	0.924**
	(0.027)	(0.029)	(0.027)	(0.034)
Panel $\mathcal{B}$.	**Transition equation**			
$p_{1,1}$	0.841**	0.926**	0.844**	0.923**
	(0.037)	(0.02)	(0.036)	(0.021)
$p_{1,2}$	0.155**	0.071**	0.152**	0.073**
	(0.036)	(0.02)	(0.035)	(0.021)
$p_{1,3}$	0.004	0.004	0.004	0.004
	(0.007)	(0.004)	(0.007)	(0.004)
$p_{2,1}$	0.026**	0.018**	0.028**	0.018**
	(0.006)	(0.005)	(0.007)	(0.006)
$p_{2,2}$	0.935**	0.936**	0.933**	0.93**
	(0.01)	(0.01)	(0.011)	(0.011)
$p_{2,3}$	0.039**	0.045**	0.039**	0.052**
	(0.008)	(0.008)	(0.008)	(0.009)
$p_{3,1}$	0.001	0.003	0.001	0.004
	(0.002)	(0.002)	(0.001)	(0.002)
$p_{3,2}$	0.029**	0.043**	0.028**	0.039**
	(0.006)	(0.008)	(0.006)	(0.008)
$p_{3,3}$	0.97**	0.955**	0.971**	0.957**
	(0.006)	(0.008)	(0.006)	(0.008)

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

Looking beyond volatility as a potential explanatory variable for price impact, we hypothesize that market makers' risk considerations beyond volatility – heavy-tails of the rate distribution – affect supply of liquidity. Volatility of volatility induces heavy-tails of the rate distribution. Since volatility tends to be volatile when its level is high, we isolate the surprise component from this relationship to measure the incremental impact of heavy tails. Specifically, we estimate the surprise component of the volatility of volatility in two steps. First, we obtain the filtered dynamic conditional volatilities from a AR(1)-GARCH(1,1) model with Student-t innovations applied to swaption volatilities. Second, we regress this time series on swaption volatilities to obtain the residuals; these residuals are an estimate of the surprise component of the volatility of volatility. In Model II, we extend the set of explanatory variables to include the surprise component of the volatility of swaption-implied volatility. The effects of the surprise component of the volatility of volatility, β_i^{VV}, are highly significant and positive in all liquidity states. The surprise component of volatility of volatility has larger effects in lower liquidity states: $\beta_i^{VV} > \beta_j^{VV}$ for $i < j$. Since, heavy tails and volatility are both risks to market makers, we can extend the definition of low liquidity states to states where liquidity is sensitive to rate risks, in general.

Shocks to volatility persistence affect liquidity supply through the balance sheet channel – the risk of maintaining an inventory of securities over multiple trading days. This channel is relevant for traditional dealers, but not principal trading firms (PTFs) – traditional dealers often carry considerable positions overnight, while PTFs tend to avoid such exposures. Model III includes volatility and the surprise component of volatility persistence as explanatory variables. Again, the rationale for retaining only the surprise component is to allow for the typical relationship between volatility and its persistence. We estimate the surprise component of volatility persistence in two steps. First, we calculate the difference between the 3-month and 1-month swaption volatilities – the *volatility calendar spread*. Second, we obtain residuals from a regression of the volatility calendar spread on the 1-month swaption volatility – these residuals are our surprise component of volatility persistence estimates.

We hypothesize that the surprise component of volatility persistence increases price impact through the balance sheet channel. However, there is also the confounding effect that we are likely to encounter during periods of extreme but brief market stress—when 1-month volatility spikes due to localized episodes of market dysfunction, 3-month volatility will not change in accordance with the typical relationship—the surprise component of volatility persistence would then be negative and large in magnitude.[5] Under such circumstances, a negative surprise component of volatility persistence reflects the severity of a localized market dysfunction. Such dysfunctions are either transitory by nature, as in the case of flash events, or are unlikely to last due to the critical role of the Treasury market and the resultant response of regulators.

In summary, the coefficient estimates β^{CS} in Model III validate our intuition: β^{CS} is positive in the high and moderate liquidity states, $i = 2, 3$, and negative in the lowest liquidity state, $i = 1$. In moderate-to-good liquidity states, when market dysfunction is near impossible, the balance sheet channel dominates. In contrast, the worst liquidity state coincides with some periods of potential market dysfunction, and the above described confounding effect dominates.

Finally, Model IV contains the full set of considered covariates: volatility, the surprise component of the volatility of volatility, and the surprise component of volatility persistence. All observation equation parameter estimates are significant with the exception of β_1^{CS}.

We conclude this section with the description of the estimated time-invariant transition probability matrix. The observations are robust across the four model specifications. Irrespective of the current state, the probability of remaining in the same state in the next period is high (above 0.92 in Model IV), suggesting notable persistence of liquidity states. The probability of remaining in the same state in the next period is greater for the higher liquidity states. Overall, transitioning to a higher liquidity state has a higher probability than transitioning to a lower

[5] A flash event is a prime example of a localized market dysfunction.

liquidity state. Transitioning between adjacent liquidity states is more probable than migrating directly between worst and best liquidity states.

3.3 Hidden Markov Model with Endogenous Transition Probabilities Matrix

We conjecture that low depth presents a vulnerability of market liquidity to a future shock to the willingness of market participants to replenish the order book. Extending the model to allow for endogenous transition probabilities permits us to explore the role of market depth in transitioning between latent liquidity states. In this case, we model the transition probabilities using the multinomial logistic specification of Visser and Speekenbrink (2010) and Zucchini and MacDonald (2009). Specifically, let s_d denote the liquidity state on day d. Then, the time-varying probability of transitioning between states is

$$p_{i,j,t} = \frac{\exp\left(\delta_{i,j} + \gamma_{i,j} z_t\right)}{1 + \sum_{k=2,3} \exp\left(\delta_{i,k} + \gamma_{i,k} z_t\right)}, \tag{4}$$

where $p_{i,j,t}$, with $j \neq 1$, denotes the trading day-t probability that $s_{t+1} = j$ conditional on $s_t = i$, and z_t is average market depth on day t. To ensure that probabilities add up (over j) to one, the probabilities of transitions to "lowest-liquidity" state 1 must be parameterized as

$$p_{i,1,t} = \frac{1}{1 + \sum_{k=2,3} \exp\left(\delta_{i,k} + \gamma_{i,k} z_t\right)}. \tag{5}$$

As a consequence of this normalization, there are no parameter estimates for $\delta_{i,1}$ and $\gamma_{i,1}$.

Table 3 reports parameter estimates for the HMM with an endogenous transition probabilities matrix. Considering first the measurement equation of the models reported in Panel $\mathcal{A}$, the main results for the measurement equations of the model, which we discussed in Section 3.2, are robust to allowing for this form of time-variation in the transition probabilities. We therefore focus on the transition equations of the models, reported in Panel $\mathcal{B}$.

The estimates of γ parameters, capturing the effect of market depth, are statistically significant at the 1 percent level, with the exception of $\gamma_{1,3}$ in Models I and II and $\gamma_{3,2}$ in all models. For all specifications, depth effects in the low and moderate liquidity states, $\gamma_{i,j}$ $(i = 1, 2,\ j = 2, 3)$, are positive: greater market depth increases the probability of a transition to a better liquidity state and reduces the probability of a transition to the low-liquidity state. $\gamma_{1,3}$ are estimated with greater uncertainty because there are few transitions from the low to the high-liquidity state, bypassing the moderate liquidity state. In the high-liquidity state 3, $\gamma_{3,3}$ are positive for all specifications indicating that higher market depth increases the probability of remaining in the high-liquidity state. In Models II and IV, $\gamma_{3,2}$ are negative suggesting that higher market depth reduces the probability of a transition to a moderate liquidity state. In Models I and III, $\gamma_{3,2}$ are positive, albeit insignificant. However, positive $\gamma_{3,2}$ can also be seen as higher market depth causing a further reduction in the probability of transitioning to the low liquidity state. Comparing Models I and III against Models II and IV, we see that accounting for the surprise component of the volatility of swaption-implied volatility in the observation equation moderates the magnitudes the effects of depth on transition probabilities. Comparison of the effects of depth on the transition probabilities suggests their greater strength when liquidity is lower, since $\gamma_{i,j} > \gamma_{i+1,j}$ for $i = 1, 2, 3$ and $j = 2, 3$.

To further illustrate the effect of depth on the transition probabilities, Figure 4 shows the transition probabilities conditional on various levels of depth, as implied by Model IV. As depth decreases, the probability of transitioning to a worse liquidity state increases. For example, when depth is \$10 million, the probability of transitioning from moderate liquidity, state 2, to low liquidity, state 1, is about nine in ten; and, in contrast, when depth is \$120 million, this probability decreases to essentially zero.

Figure 5 shows the estimated smoothed probabilities of being in a given latent liquidity state at each point in time, as implied by Model IV, and computed using an algorithm from Visser and Speekenbrink (2010). The lower panel plots price

Table 3: HMM with Endogenous Transition Probabilities Matrix: Parameter Estimates

The table reports parameter estimates for the HMMs in Equations 2 and 4 using data on price impact, volatility, and market depth for the period from April 1, 2014 to December 31, 2023. β^V coefficients corresponds to the effect of swaption implied volatility, β^{VV} – to the effect of the surprise component of the volatility of swaption implied volatility, and β^{CS} – to the effect of surprise component of volatility persistence. Sub-indices correspond to states, $i = 1, 2, 3$, ranked from lowest to highest liquidity. Panel $\mathcal{B}$ contains estimates of the transition equation parameters, following the notation of Equation (4). Numbers in brackets indicate standard errors based on the Hessian. A $\cdot$, *, or ** indicates two-sided p-values of less than 0.1, 0.05, and 0.01, respectively.

	Model I	Model II	Model III	Model IV
Panel $\mathcal{A}$. **Observation equation**				
β_1^V	0.190**	0.165**	0.185**	0.167**
	(0.007)	(0.004)	(0.006)	(0.004)
β_1^{VV}		1.675**		1.512**
		(0.108)		(0.132)
β_1^{CS}			-0.675**	-0.162*
			(0.092)	(0.073)
σ_1	8.949**	5.712**	7.825**	5.776**
	(0.525)	(0.27)	(0.468)	(0.299)
β_2^V	0.119**	0.115**	0.12**	0.116**
	(0.001)	(0.001)	(0.001)	(0.001)
β_2^{VV}		0.427**		0.463**
		(0.056)		(0.065)
β_2^{CS}			0.026	0.044*
			(0.017)	(0.018)
σ_2	1.729**	1.509**	1.762**	1.555**
	(0.061)	(0.053)	(0.064)	(0.063)
β_3^V	0.083**	0.081**	0.084**	0.083**
	(0.001)	(0.001)	(0.001)	(0.001)
β_3^{VV}		0.157**		0.236**
		(0.044)		(0.045)
β_3^{CS}			0.095**	0.096**
			(0.01)	(0.01)
σ_3	0.974**	0.917**	0.958**	0.910**
	(0.033)	(0.028)	(0.029)	(0.028)
Panel $\mathcal{B}$. **Transition equation**				
$\delta_{1,2}$	-9.893**	-8.286**	-10.000**	-8.988**
	(2.501)	(1.955)	(2.656)	(2.701)
$\delta_{1,3}$	-107.188	-200.496	-22.938**	-284.616**
	(144.634)	(132.53)	(7.075)	(91.647)
$\gamma_{1,2}$	0.292**	0.200**	0.300**	0.230**
	(0.079)	(0.056)	(0.088)	(0.082)
$\gamma_{1,3}$	1.892	3.338	0.528**	4.696**
	(2.321)	(2.142)	(0.148)	(1.482)
$\delta_{2,2}$	-3.999**	-3.023**	-4.277**	-3.117**
	(1.137)	(0.935)	(1.276)	(1.002)
$\delta_{2,3}$	-8.317**	-7.562**	-8.889**	-7.467**
	(1.379)	(1.52)	(1.515)	(1.474)
$\gamma_{2,2}$	0.135**	0.100**	0.144**	0.104**
	(0.027)	(0.018)	(0.031)	(0.02)
$\gamma_{2,3}$	0.16**	0.122**	0.174**	0.126**
	(0.029)	(0.023)	(0.033)	(0.025)
$\delta_{3,2}$	-2.564	3.586*	-5.417	3.366.
	(3.357)	(1.588)	(3.474)	(1.968)
$\delta_{3,3}$	-5.428.	-0.042	-7.287*	0.294
	(2.957)	(0.429)	(3.128)	(1.123)
$\gamma_{3,2}$	0.058	-0.032	0.094.	-0.029
	(0.046)	(0.02)	(0.05)	(0.024)
$\gamma_{3,3}$	0.118**	0.039**	0.144**	0.036**
	(0.041)	(0.005)	(0.046)	(0.012)

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

Figure 4: HMM Transition Matrix

The figure shows the transition probability matrices from the HMM conditional on various levels of market depth, as implied by Model IV. The levels of market depth are $10 million, $30 million, $60 million, and $120 million. In each case, the matrix shows the probabilities of transitioning from the states on the vertical axis to the states on the horizontal axis. Darker shading highlights higher probabilities.

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

impact. The estimated state probabilities make intuitive sense, with higher probabilities of residing in a low-liquidity state generally being higher at times of spikes in price impact. Until around late 2019 or early 2020, the probability of being in the high-liquidity state 3 was generally very high, with just a handful of episodes of being in the medium-liquidity state. These episodes generally coincided with moderate increases in price impact. Since 2020, however, the medium-liquidity state has shifted to being the predominant state, with the high-liquidity state becoming relatively rare and the low-liquidity state becoming much more common. The two most notable periods of low-liquidity in the sample appear to be associated with the COVID-related market turmoil in early 2020 and the low liquidity as economic uncertainty rose toward the end of the sample.

Figure 5: HMM-Implied Probabilities of Being in Each State

The upper panel shows the estimated probabilities of being in each state, as implied by Model IV. The blue, grey, and red areas indicate the probabilities of being in the high-, medium, and low-liquidity states, respectively. The lower panel shows a time series of price impact. The vertical dashed lines in the lower panel correspond to: (a) October 15, 2014; (b) June 24, 2016; (c) February 6, 2018; (d) August 2, 2019; (e) March 9, 2020; (f) February 25, 2021; and (g) March 13, 2023.

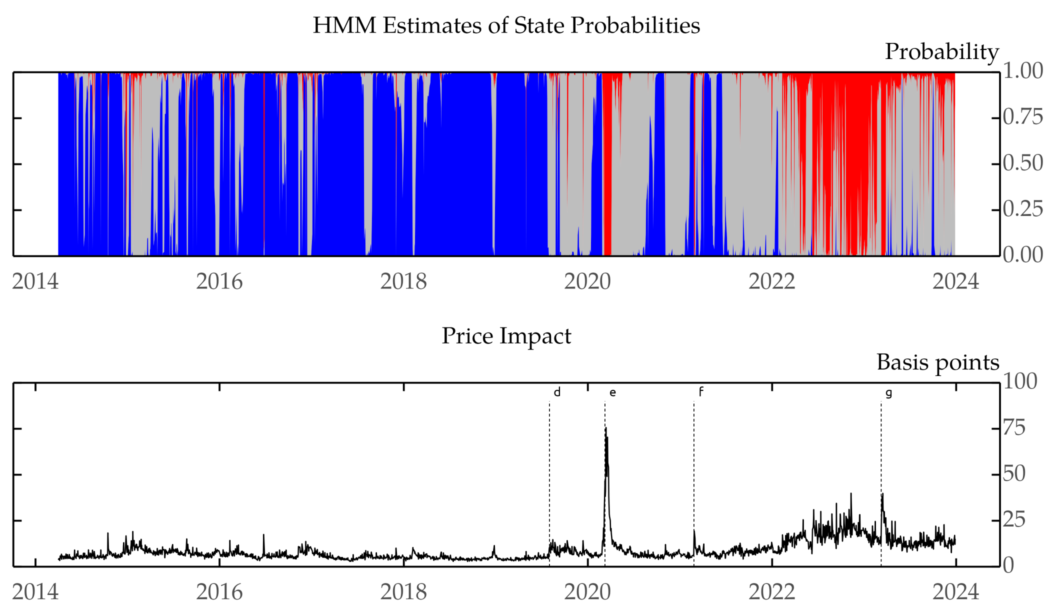

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

The bottom panel of Figure 1 suggests that post the event of August 2, 2019, market depth appears to have been persistently lower than before. In Appendix B, we establish the robustness of our results to this potential structural break in liquidity conditions.

4 Market-Wide Liquidity Fragility

In Section 4.2 we first confirm our results for all major benchmark Treasury securities by fitting HMMs with security-specific latent liquidity states. The validity of our analysis across multiple securities suggests the feasibility of extending our approach to modeling market-wide liquidity fragility. Consequently, in Section 4.1, we introduce an HMM with common latent liquidity states to assess market-wide liquidity fragility.

4.1 General Validity of Results Across Benchmark Securities

In this section, we estimate HMMs separately for each considered benchmark Treasury security. This approach allows latent liquidity states to be security-specific rather than market-wide. Then, we consider market-wide states in Section 4.2.

Table 4 reports HMM estimates for the 2-, 5-, 10-, and 30-year on-the-run Treasury securities.[6] Reported parameter estimates support the universal applicability of our Section 3 results to on-the-run Treasury securities. First, there is strong evidence of state-dependent sensitivities of price impact to considered uncertainty-related factors. There is a common pattern of how sensitivities vary with liquidity states, which was described in detail for the case of the 10-year benchmark security in Section 3.1. Second, market depth contributes to explaining transitions between liquidity states – this evidence establishes depth as a universal driver of liquidity fragility for all considered benchmark securities. Once again, the pattern of how the

[6] Since a bond's interest rate sensitivity depends on maturity and price impact depends on the amount of risk transfer, we re-scale price impacts for 2-, 5-, and 30-year Treasuries in order for them to correspond to roughly the same amount of risk transfer as in the case of the 10-year Treasury Note trade of them same notional amount. We omit the 20-year Treasury Bond due to the shorter sample of available data since the reintroduction of the security in 2020.

effect of depth changes with liquidity state is generally shared by all securities.

4.2 Market-Wide Liquidity Fragility

In this section, we explore whether our results extend to market-wide liquidity fragility. To this end, we turn the HMM in Equations (2) and (4) into a multivariate model with common liquidity states for the 2-, 5-, 10-, and 30-year on-the-run Treasury securities. Reinterpreting Θ_t as a vector consisting of price impacts of each security on day t, the observation equation becomes:

$$\Theta_t = B_i \mathbf{V}_t + \epsilon_{i,t}, \tag{6}$$

where i indexes the latent liquidity state; B_i is a diagonal matrix

$$B_i = \mathbf{diag}\left(\beta_{i,2y}, \beta_{i,5y}, \beta_{i,10y}, \beta_{i,30y}\right)$$

with the diagonal built from security-specific state i sensitivity parameters; $\mathbf{V}_t$ is a vector of 1-month swaption volatilities of tenors corresponding to securities' original maturities; and $\epsilon_{i,t} \sim \mathcal{N}(0, \Sigma_i)$ is an observation error vector for latent liquidity state i.[7] Each Σ_i is diagonal, thereby forcing the common states to encapsulate all cross-sectional dependence.

Since liquidity states are now common for all securities, we consider a larger 4-state model, with the implications discussed below. In order to identify drivers of liquidity fragility, we make the transition probability matrix endogenous – the time-varying probability of transitioning between states is

$$p_{i,j,t} = \frac{\exp\left(\delta_{i,j} + \gamma_{i,j}\mathbf{PC}_t^1\right)}{1 + \sum_{k=2,3,4} \exp\left(\delta_{i,k} + \gamma_{i,k}\mathbf{PC}_t^1\right)}, \tag{7}$$

where $p_{i,j,t}$, with $j \neq 1$, denotes the trading day-t probability that $s_{t+1} = j$ conditional on $s_t = i$. In contrast to the transition equation for the univariate HMM in

[7] Due to a large parameter space, we restrict attention to the model with swaption volatilities as covariates in the observation equation.

Table 4: HMM with Endogenous Transition Probabilities Matrix: Parameter Estimates

The table reports parameter estimates for the HMMs in Equations 2 and 4 using data on price impact, volatility, and market depth for the period from April 1, 2014 to December 31, 2023. β^V coefficients corresponds to the effect of swaption implied volatility, β^{VV} – to the effect of the surprise component of the volatility of swaption implied volatility, and β^{CS} – to the effect of surprise component of volatility persistence. Sub-indices correspond to states, $i = 1, 2, 3$, ranked from lowest to highest liquidity. Panel $\mathcal{B}$ contains estimates of the transition equation parameters, following the notation of Equation (4). Numbers in brackets indicate standard errors based on the Hessian. A $\cdot$, *, or ** indicates two-sided p-values of less than 0.1, 0.05, and 0.01, respectively.

	2-y Note	5-y Note	10-y Note	30-y Bond
Panel $\mathcal{A}$. Observation equation				
β_1^V	0.209**	0.162**	0.167**	0.471**
	(0.003)	(0.002)	(0.004)	(0.004)
β_1^{VV}	1.079**	1.038**	1.512**	3.335**
	(0.138)	(0.092)	(0.132)	(0.121)
β_1^{CS}	-0.066	0.033	-0.162*	0.009
	(0.054)	(0.041)	(0.073)	(0.022)
σ_1	8.66**	5.37**	5.776**	7.378**
	(0.296)	(0.177)	(0.299)	(0.203)
β_2^V	0.13**	0.105**	0.116**	0.35**
	(0.002)	(0.001)	(0.001)	(0.002)
β_2^{VV}	0.468**	0.827**	0.463**	1.308**
	(0.078)	(0.046)	(0.065)	(0.164)
β_2^{CS}	0.033.	0.024	0.044*	0.273**
	(0.017)	(0.016)	(0.018)	(0.039)
σ_2	1.408**	1.227**	1.555**	3.471**
	(0.049)	(0.042)	(0.063)	(0.11)
β_3^V	0.079**	0.072**	0.083**	0.254**
	(0.001)	(0)	(0.001)	(0.002)
β_3^{VV}	0.144**	0.302**	0.236**	0.114
	(0.043)	(0.026)	(0.045)	(0.183)
β_3^{CS}	0.048**	0.027**	0.096**	0.133*
	(0.009)	(0.009)	(0.01)	(0.052)
σ_3	1.027**	0.805**	0.91**	2.389**
	(0.026)	(0.022)	(0.028)	(0.099)
Panel $\mathcal{B}$. Transition equation				
$\delta_{1,2}$	-7.803**	-7.19**	-8.988**	-6.077*
	(2.108)	(1.081)	(2.701)	(2.723)
$\delta_{1,3}$	-91.645	-17.984**	-284.616**	-21.571
	(242.266)	(4.745)	(91.647)	(83.307)
$\gamma_{1,2}$	0.153**	0.129**	0.23**	0.46
	(0.051)	(0.028)	(0.082)	(0.386)
$\gamma_{1,3}$	0.868	0.279**	4.696**	0.935
	(2.114)	(0.07)	(1.482)	(1.465)
$\delta_{2,2}$	-3.616**	-2.44**	-3.117**	-4.201.
	(1.289)	(0.924)	(1.002)	(2.181)
$\delta_{2,3}$	-6.323**	-5.933**	-7.467**	-11.508**
	(1.339)	(1.32)	(1.474)	(3.031)
$\gamma_{2,2}$	0.078**	0.075**	0.104**	0.919**
	(0.021)	(0.015)	(0.02)	(0.292)
$\gamma_{2,3}$	0.074**	0.077**	0.126**	1.226**
	(0.021)	(0.017)	(0.025)	(0.354)
$\delta_{3,2}$	1.285	-0.607	3.366.	1.873
	(1.114)	(3.965)	(1.968)	(1.54)
$\delta_{3,3}$	2.741**	0.738	0.294	0.41
	(0.79)	(2.131)	(1.123)	(0.348)
$\gamma_{3,2}$	-0.003	0.005	-0.029	0.103
	(0.004)	(0.031)	(0.024)	(0.23)
$\gamma_{3,3}$	0.004*	0.024	0.036**	0.498**
	(0.002)	(0.017)	(0.012)	(0.189)

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

Equation (4), now the first principal component of market depths of all considered securities, $\mathbf{PC}_t^1$, explains the latent liquidity state dynamics. Similar to Equation (5), to ensure that probabilities add up (over j) to one, the probabilities of transitions to "lowest-liquidity" state 1 must be parameterized as

$$p_{i,1,t} = \frac{1}{1 + \sum_{k=2,3,4} \exp\left(\delta_{i,k} + \gamma_{i,k}\mathbf{PC}_t^1\right)}. \tag{8}$$

Once again. as a consequence of this normalization, there are no parameter estimates for $\delta_{i,1}$ and $\gamma_{i,1}$.

Table 5 reports the model parameter estimates, with Panels $\mathcal{A}$ and $\mathcal{B}$ dedicated to the observation and transition equations, respectively. We label the states from lowest to highest liquidity. The distinction between States 2 and 3 is the notably worse liquidity of the 2- and 5-year Treasury Notes in state 2 relative to state 3.

Since all significant $\gamma_{i,j}$ coefficient estimates are positive for transitions to states with no less liquidity ($j \geq i$), we obtain evidence that greater overall market depth reduces market-wide liquidity fragility. Interpretation of a positive, albeit insignificant, $\gamma_{3,2}$ is more nuanced, as it not only increases the probability of a transition to state 2 but also diminishes the probability of a transition to an even worse state 1, through increasing the denominator in Equation (8). Moreover, the magnitudes of $\gamma_{3,3}$ and $\gamma_{3,4}$ are notably greater than that of $\gamma_{3,2}$. Reaffirming the results for individual securities, we once again find the effect of depth to be stronger and more significant in the three lower liquidity states. In the worst liquidity state, the effect of depth is significant for the probability of escaping to the better liquidity state 2, as such transitions to an adjacent state are sufficiently common. In contrast, direct transitions from state 1 to either state 3 or 4 are relatively rare, making the standard errors for $\gamma_{1,j}$ ($j = 3, 4$) large. However, we can also interpret the combination of positive $\gamma_{1,j}$ ($j = 3, 4$) as greater depth decreasing the probability of remaining in the worst liquidity state 1. Insignificant estimates of $\{\gamma_{4,j}\}$ ($j = 2, 3, 4$) suggest that depth plays a lesser role when liquidity is highest—market makers replenish the central limit order book promptly and, thus, depth plays a diminished role.

That said, $\gamma_{4,4} > \gamma_{4,j}$ $(j = 2, 3)$ and has the greatest significance among all depth sensitivities in the highest liquidity state 4. In summary, our results on the pivotal role of depth carry over to market-wide liquidity fragility.

5 Conclusion

In this paper, we have drawn attention to liquidity fragility—the likelihood of a deterioration in liquidity—in the benchmark Treasury market. We showed that market depth significantly affects liquidity fragility in all maturity sectors of the Treasury market – times of low market depth are associated with an increased probability of low liquidity states in the future. This could reflect a more severe degree of liquidity fragility associated with greater reliance on high-speed quote replenishment when depth is relatively low. The universal validity of our results for benchmark securities enables us to analyze market-wide liquidity fragility through a model with common liquidity states for all benchmark Treasury securities. Our results remained valid when considering the fragility of liquidity for all benchmark Treasury securities, as it is significantly affected by a measure of aggregate market depth.

We contributed to the literature on the econometric modeling of latent liquidity—our Hidden Markov Model (HMM) relates price impact to features of volatility, conditional on latent liquidity states. Unlike in previous studies using such models to analyze liquidity, we allowed observable variables to affect transition probabilities. As a result, our HMM captured the non-linear dependence of price impact on the key uncertainty-related economic factors affecting market makers. In addition, our analysis benefitted from a novel framework for quantifying price impact that allows trade flow imbalance and non-marketable order flow imbalance to have different and nonlinear effects on prices.

We also considered uncertainty-related factors that affect market makers. There may be other systemic factors that explain additional variation in liquidity, like balance sheet constraints explored in Duffie et al. (2023), among others. We leave

Table 5: HMM with Endogenous Transition Probabilities Matrix: Parameter Estimates

The table reports parameter estimates for a multivariate HMM, Equations (6 - 8), with four common liquidity states for the 2-, 5-, 10-, and 30-year Treasury securities, using data on price impact, volatility, and market depth for the period from April 1, 2014 to December 31, 2023. Panel $\mathcal{A}$ contains parameter estimates for the observation equation. Sub-indices correspond to states, $i = 1, 2, 3, 4$, ranked from lowest to highest liquidity. Panel $\mathcal{B}$ contains estimates of the transition equation parameters. The notation corresponds to Equation (7); the parameters are arranged so that rows correspond to departure states, while columns correspond to arrival states. Numbers in brackets indicate standard errors based on the Hessian. A $\cdot$, *, or ** indicates two-sided p-values of less than 0.1, 0.05, and 0.01, respectively.

	State 1	State 2	State 3	State 4
Panel $\mathcal{A}$. **Observation equation**				
$\beta_{i,2-y}$	0.313**	0.187**	0.110**	0.076**
	(0.012)	(0.002)	(0.002)	(0.001)
$\beta_{i,5-y}$	0.238**	0.149**	0.099**	0.069**
	(0.008)	(0.002)	(0.001)	(0.000)
$\beta_{i,10-y}$	0.235**	0.130**	0.109**	0.080**
	(0.012)	(0.001)	(0.001)	(0.000)
$\beta_{i,30-y}$	0.610**	0.419**	0.408**	0.319**
	(0.025)	(0.005)	(0.004)	(0.002)
Panel $\mathcal{B}$. **Transition equation**				
To state:		State 2	State 3	State 4
	From state:			
$\delta_{1,\cdot}$	State 1	32.335**	177.251	69.348
		(10.096)	(170.449)	(196.023)
$\gamma_{1,\cdot}$		10.923**	85.766	34.718
		(3.384)	(87.117)	(114.256)
$\delta_{2,\cdot}$	State 2	9.86**	12.796**	8.801*
		(2.25)	(2.467)	(3.568)
$\gamma_{2,\cdot}$		3.083**	5.593**	4.984**
		(0.898)	(1.026)	(1.63)
$\delta_{3,\cdot}$	State 3	11.529*	15.968**	14.503**
		(4.915)	(4.94)	(4.942)
$\gamma_{3,\cdot}$		4.841*	7.078**	7.727**
		(2.274)	(2.298)	(2.302)
$\delta_{4,\cdot}$	State 4	8.725	10.703	12.131
		(7.544)	(7.556)	(7.553)
$\gamma_{4,\cdot}$		1.075	1.495	2.033
		(1.578)	(1.573)	(1.568)

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

the integration of such systemic factors into a richer HMM for future research. A related question is the extent to which non-linearity in the relationship between price impact and uncertainty may be due to balance sheet constraints and other systemic factors. Finally, another strand of future research could verify whether the statistical regularities that we uncovered in the Treasury market hold in other CLOB-driven markers and, if not, then what specific aspects of the Treasury market microstructure drive our conclusions.

References

T. Adrian, M. J. Fleming, and E. Vogt. An Index of Treasury Market Liquidity: Evidence from 30 Years of Limit Order Book Data. *Federal Reserve Bank of New York Staff Report Number 827*, 2023.

A. Aronovich, D. Dobrev, and A. Meldrum. The Treasury Market Flash Event of February 25, 2021. *FEDS Notes, May 14, 2021*, 2021.

J. Białkowski, S. Darolles, and G. Le Fol. Improving VWAP Strategies: A Dynamic Volume Approach. *Journal of Banking and Finance*, 32(9):1709–1722, 2008.

Board of Governors of the Federal Reserve System. Financial Stability Report, November, 2020.

J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How Markets Slowly Digest Changes in Supply and Demand. In T. Hens and K. R. Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*, pages 57–160. Elsevier, 2009.

J. Brogaard, T. Hendershott, and R. Riordan. Price Discovery without Trading: Evidence from Limit Orders. *Journal of Finance*, 74(4):1621–1658, 2019.

G. Cespa and T. Foucault. Illiquidity Contagion and Liquidity Crashes. *The Review of Financial Studies*, 27(6):1615–1660, 2014.

G. Cespa and X. Vives. High Frequency Trading and Fragility. *Available at SSRN*, 2017.

L. Chollete, R. Næs, and J. A. Skjeltorp. *What Captures Liquidity Risk? A Comparison of Trade and Order-Based Liquidity Factors.* Norges Bank Working Paper Number 3, 2007.

T. Chordia, R. Roll, and A. Subrahmanyam. Co-movements in Bid-Ask Spreads and Market Depth. *Financial Analysts Journal*, 56(5):23–27, 2000.

T. Chordia, R. Roll, and A. Subrahmanyam. Market Liquidity and Trading Activity. *Journal of Finance*, 56(2):501–530, 2001.

T. Chordia, R. Roll, and A. Subrahmanyam. Determinants of Daily Fluctuations in Liquidity and Trading Activity. *Cuadernos de Economía*, 40(121):728–751, 2003.

T. Chordia, A. Sarkar, and A. Subrahmanyam. An Empirical Analysis of Stock and Bond Market Liquidity. *The Review of Financial Studies*, 18(1):85–129, 2005.

R. Cont and J.-P. Bouchaud. Herd Behavior and Aggregate Fluctuations in Financial Markets. *Macroeconomic Dynamics*, 4(2):170–196, 2000.

R. Cont, A. Kukanov, and A. Stoikov. The Price Impact of Order Book Events. *Journal of Financial Econometrics*, 12(1):47–88, 2014.

R. Cont, M. Cucuringu, and C. Zhang. Price Impact of Order Flow Imbalance: Multi-level, Cross-sectional and Forecasting. *Working paper*, 2021.

D. Dobrev and A. Meldrum. What Do Quoted Spreads Tell Us about Machine Trading at Times of Market Stress? Evidence from Treasury and FX Markets during the COVID-19-related Market Turmoil in March 2020. *FEDS Notes, September 25, 2020*, 2020.

D. Duffie, M.J. Fleming, F.M. Keane, C. Nelson, O. Shachar, and P. Van Tassel. Dealer Capacity and US Treasury Market Functionality. *FRB of New York Staff Report*, (1070), 2023.

J. D. Farmer, L. Gillemot, F. Lillo, S. Mike, and A. Sen. What Really Causes Large Price Changes? *Quantitative Finance*, 4(4):383–397, 2004.

J. D. Farmer, A. Gerig, F. Lillo, and S. Mike. Market Efficiency and the Long-memory of Supply and Demand: Is Price Impact Variable and Permanent or Fixed and Temporary? *Quantitative Finance*, 6(02):107–112, 2006.

M. D. Flood, J. Liechty, and T. Piontek. Systemwide Commonalities in Market Liquidity. *Office of Financial Research Working Paper Number 15-11*, 2016.

J. Fraenkle, S. Z. Rachev, and C. Scherrer. Market Impact Measurement of a VWAP Trading Algorithm. *Journal of Risk Management in Financial Institutions*, 4(3): 254–274, 2011.

C. Frei and N. Westray. Optimal Execution of a VWAP Order: A Stochastic Control Approach. *Mathematical Finance*, 25(3):612–639, 2015.

N. Hautsch and R. Huang. Limit Order Flow, Market Impact and Optimal Oder Sizes: Evidence from NASDAQ TotalView-ITCH Data. In F. Abergel, J.-P. Bouchaud, T. Foucault, C.-A. Lehalle, and M. Rosenbaum, editors, *Market Microstructure: Confronting Many Viewpoints*. John Wiley and Sons, Ltd., 2012a.

N. Hautsch and R. Huang. The Market Impact of a Limit Order. *Journal of Economic Dynamics and Control*, 36(4):501–522, 2012b.

R. A. Korajczyk and R. Sadka. Pricing the Commonality across Alternative Measures of Liquidity. *Journal of Financial Economics*, 87(1):45–72, 2008.

A. S. Kyle. Continuous Auctions and Insider Trading. *Econometrica*, pages 1315–1335, 1985.

F. Lillo and J. D. Farmer. The Long Memory of the Efficient Market. *Studies in Nonlinear Dynamics and Econometrics*, 8(3), 2004.

E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, and R. N. Mantegna. Market Impact and Trading Profile of Large Trading Orders in Stock Markets. *Working Paper*, 2009.

R. Næs, J. A. Skjeltorp, and B. A. Ødegaard. Liquidity at the Oslo Stock Exchange. *Working paper*, 2008.

G. Nguyen, R. Engle, M. J. Fleming, and E. Ghysels. Liquidity and Volatility in the US Treasury Market. *Journal of Econometrics*, 217(2):207–229, 2020.

Maureen O'Hara. Liquidity and Financial Market Stability. *National Bank of Belgium Working Paper*, (55), 2004.

M. C. Pham, H. M. Anderson, H. N. Duong, and P. Lajbcygier. The Effects of Trade Size and Market Depth on Immediate Price Impact in a Limit Order Book Market. *Journal of Economic Dynamics and Control*, 120:103992, 2020.

V. Raman, M. A. Robe, and P. K. Yadav. Electronic Market Makers, Trader Anonymity and Market Fragility. *Working paper*, 2014.

V. Raman, M.A. Robe, and P.K Yadav. Man vs. Machine: Liquidity Provision and Market Fragility. *Available at SSRN*, 2020.

I. Visser and M. Speekenbrink. depmixs4: an R Package for Hidden Markov Models. *Journal of Statistical Software*, 36:1–21, 2010.

K. Xu, M. D. Gould, and S. D. Howison. Multi-level Order-flow Imbalance in a Limit Order Book. *Market Microstructure and Liquidity*, 4:1950011, 2018.

W. Zucchini and I. L. MacDonald. *Hidden Markov Models for Time Series: An Introduction using R.* Chapman and Hall/CRC, 2009.

Appendix A Estimating Price Impact

We rely on price impact as our measure of price impact – price impact is the effect on market price of trading a given economically significant order size over a set horizon. Given the small commissions, it is the effect of trading on price that constitutes the bulk of trading costs, and price impact is designed to directly assess this effect. The less liquid the market, the more trades and orders will move the price. Price impact is not observed directly but must be estimated using an econometric model. In this section, we explain our novel approach to estimating price impact.

Our approach to estimating price impact is carefully designed to account for several important dimensions of liquidity on inter-dealer broker platforms that use the central limit order book (CLOB) protocol. Trading on such platforms accounts for a large share of dealer-to-dealer transactions in on-the-run Treasury securities.

A CLOB consists of bid and ask queues of "limit orders" (that is, quotes) to respectively buy or sell specified quantities of securities, with the best bid price being below the best ask price. An important aspect of CLOB-based trading is that only a small fraction of desired transaction volume is typically immediately available at the best price. Therefore, market participants wishing to execute a large position change, or "parent order," will typically split the parent order into smaller "child orders" in accordance with some execution algorithm. The first choice when designing an execution algorithm is whether (i) to consume liquidity on the CLOB by submitting marketable orders that cross the spread and are immediately executable, or (ii) to execute trades by *providing* liquidity by posting additional limit orders with prices such that they cannot be immediately matched, in the anticipation that the orders will be matched to marketable orders placed by other market participants later. Executing a trade using "non-marketable" limit orders may be cheaper than using marketable orders because it avoids paying the bid-ask spread but comes at the risk that the order will not be matched quickly or that the price may move in an unfavorable direction such that the limit order will not be matched at all (meaning that the order may need to be canceled and a new order placed at a less favorable price). In practice, a sophisticated execution strategy may use a mixture of marketable and non-marketable orders (see, for example, Moro et al. (2009)). The second dimension of an execution strategy that matters for the price impact of a large parent order is how the market participant splits it into smaller child orders over time, in order to lower the overall price impact.

Our measure of price impact is based on a model relating returns to net trade and non-marketable order flows. The model takes the form:

$$\begin{aligned} r_{t-1,t} = {} & \beta_m sign\left(TF_{t-1,t}\right)\left|TF_{t-1,t}\right|^{\gamma_m} \\ & + \beta_n sign\left(NMOF_{t-1,t}\right)\left|NMOF_{t-1,t}\right|^{\gamma_n} + \varepsilon_t. \end{aligned} \tag{9}$$

where $t = 0, \ldots, T$ are the boundaries of one-minute intraday periods, which together form a partition of the studied trading day; $TF_{t-1,t}$ is the trade flow imbalance, measured as the difference between buyer- and seller-initiated trade volumes between $t-1$ and t; $NMOF_{t-1,t}$ is the non-marketable order flow imbalance, measured as order flow imbalance as calculated in Cont et al. (2014) between $t-1$ and t, less the effect of trades, which are accounted for in trade flow;[8] and $r_{t-1,t} = \frac{P_t - P_{t-1}}{P_{t-1}}$ is the simple return based on changes in mid-quoted prices between $t-1$ and t (P_{t-1} and P_t, respectively). The inclusion of both trade flow and non-marketable order flow as explanatory variables, with different impacts on prices, is motivated by the results of Brogaard et al. (2019), who show that in equity markets marketable orders individually have greater price impact than limit orders, but that limit orders contribute more to price discovery overall. The feature that the effects of trade flow and non-marketable order flow are nonlinear is motivated by the review in Bouchaud et al. (2009), and, in particular, the asymmetric liquidity hypothesis of Lillo and Farmer (2004) and Farmer et al. (2006). These studies reconcile the fact that the direction of order flow imbalances displays long memory with market efficiency: an order has a smaller price impact when following another order in the same direction.

We estimate the model in equation (9) for the 2-, 5-, 10-, and 30-year on-the-run nominal Treasury securities using nonlinear least squares separately for each day between April 1, 2014 and December 31, 2023. We omit days with shortened trading hours, leaving a total of 2,374 days. We consider only the most active trading hours, between 7 a.m. and 4.45 p.m., giving a maximum of 585 one-minute observations per trading day. For model estimation purposes we retain only intraday periods with at least one trade. We also constrain the values of γ_m to be below 1.8, which avoids extreme estimates of this parameter on a small proportion (less than 1 percent) of days.

In a framework that allows trade and non-marketable order imbalances to have different and nonlinear effects on prices, there is not a single measure of price impact

[8] Specifically, $NMOF_{t-1,t} = OFI_{t-1,t} - TF_{t-1,t}$, where $OFI_{t-1,t} = \sum_{n=N(t-1)+1}^{N(t)} \left[q_n^b I\left(P_n^b \geq P_{n-1}^b\right) - q_{n-1}^b I\left(P_n^b \leq P_{n-1}^b\right) - q_n^s I\left(P_n^s \leq P_{n-1}^s\right) + q_{n-1}^s I\left(P_n^s \geq P_{n-1}^s\right) \right]$; $N(t)$ is the index of the last order book event—a limit order or cancellation—of the period between $t-1$ and t; q_n^b and q_n^s are the sizes of the queue at the best bid price, P_n^b, and ask price, P_n^s, respectively, at event n; and $I(.)$ is an indicator function taking the value 1 when the argument is true and 0 otherwise.

for the execution of a large parent order, since it depends on the time horizon for execution, how the parent order is split over time into child orders, and to what extent it is executed aggressively using marketable orders or passively using non-marketable orders. In our analysis, we consider a benchmark execution strategy under which a parent order is executed using marketable orders spread across the course of a single trading day using a "volume-weighted average price" (VWAP) strategy—that is, the parent order is spread in proportion to aggregate trade volume in each 1-minute time interval. This benchmark seems reasonable for two reasons. First, Fraenkle et al. (2011) demonstrate the optimality of VWAP execution, while Frei and Westray (2015) argue that market participants consider it to be a fair execution benchmark, and Moro et al. (2009) provide evidence from equity markets that parent orders appear to be split across time in line with the overall market activity.[9] Using this baseline strategy, we reaffirm previous results that price impact is positively related to interest rate volatility (see, for example, Cont and Bouchaud (2000), Farmer et al. (2004) and Chordia et al. (2005)). And second, Pham et al. (2020) demonstrate that spreading large parent orders over time yields considerable savings in terms of price impact.

More specifically, our baseline strategy considers the price impact of a purchase of \$500 million—equivalent to about 1.5 percent of median daily trade volumes over the period—executed using marketable orders spread across one-minute periods through the day, in proportion to the volume traded within each minute. In this baseline case, we have $x_{t-1,t} = \xi \frac{V_{t-1,t}}{V_d}$ and $y_{t-1,t} = 0$, where $V_{t-1,t}$ and V_d are the existing trading volumes between times $t-1$ and t, and in aggregate for day d, respectively.[10] This splitting in child orders corresponds to an ideal VWAP algorithm, named so because it targets the volume-weighted average price over its execution horizon. While a market participant would not in practice be able to forecast trading volumes perfectly, as discussed above, we can nevertheless think

[9]While it is not straightforward to achieve VWAP execution because it requires an accurate forecast of the share of each intraday period in aggregate trading volume, Białkowski et al. (2008) suggest such a model.

[10]To simplify calculations, we do not account for the contribution of the marginal trades to volumes, which is a reasonable approximation due to their relatively small size.

of the resulting measure of price impact as being representative of the impact of aggregate trading, as well as corresponding to a recognized execution benchmark. Of course the choice of the $500 million parent order over the period of one trading day is somewhat arbitrary, in that we picked a range of sizes that are not outsized relative to aggregate daily volume. A feature of a nonlinear model is that the scaling of price impact with trade size may be different on each day. Thus, there may be a concern that the results we discuss below are not robust to varying the trade size. However, in preliminary analysis (unreported here but available upon request) shows that the scaling becomes close to linear as the execution horizon lengthens to cover the active hours of the entire trading day. This result means that our main conclusions are likely to be robust to varying the size of the parent order.

We conclude this section with an discussion of how price impact improves on other common measures of liquidity. A significant portion of the literature, represented Chordia et al. (2001) to name one of the most seminal papers in the area, associates liquidity with measures of bid-ask spreads and market depth, defined as the total amount of quotes at the best prices resting on the order book.[11] Bid-ask spread has limited variability in the benchmark Treasury market due to the relatively large tick size: bid-ask spread is often a single tick. Moreover, Dobrev and Meldrum (2020) notices that market participants may not execute during periods of temporarily wide bid-ask spreads, making common measures of bid-ask spread biased. An even greater disadvantage of the bid-ask spread is that it captures only an instantaneous effect of trading a small order and neglects any informational consequences of trading a large order. In principle, market depth at the top of the book captures the volume of marketable orders that can be executed instantaneously without moving prices, so intuitively should be related to price impact. However, market depth has only limited value as a proxy for price impact because the large majority of willingness to post quotes is unobserved. When market participants wish to execute a large parent order, rather than exhausting all available quotes at

[11] Aggregate measures of liquidity may also be constructed, as in Duffie et al. (2023) who proxy liquidity with the first principal component of six measures associated with liquidity.

the best prices, they will typically split it into a number of smaller child orders, each of which are matched with a portion of the posted quotes, which are then quickly replenished (see, for example, Moro et al. (2009)). And because market participants adapt to changing market environments, a deterioration in depth may not translate perfectly and immediately into higher price impact. Specifically, in response to lower depth, market participants may reduce the size of their child orders to avoid trading through multiple levels of the book.[12] In an analogous manner, if quotes at the best prices are temporarily exhausted, market participants tend to avoid executing at relatively unattractive prices and instead wait for quotes at better prices to be replenished, as shown by Dobrev and Meldrum (2020). Thus temporarily widening bid-ask spreads have less effect on price impact through the adaptation of market participants. While price impact and depth are negatively correlated, price impact tends to recover faster following an increase in volatility, suggesting that trading patterns quickly adapt to a lower level of depth. Given our study of the relationship between price impact and market depth, our paper is related to those that examine the signals from multiple measures of liquidity. While one strand of the related literature has shown that various liquidity measures contain much common information (e.g. Chordia et al. (2000), Chordia et al. (2003)), Næs et al. (2008), and Korajczyk and Sadka (2008)), several other studies have focused more on the different signals from different liquidity measures, as we do. In particular, Chollete et al. (2007) argue that there can be major disagreements between liquidity measures, especially during stress events. In summary, depth and bid-ask spread do not have the direct link to price impact, as opposed to our usage of price impact.

Furthermore, we permit trade flow imbalances and non-marketable order flow imbalances to have different effects on prices and to affect prices non-linearly, unlike in the majority of previous studies. The regression-based approach was first established by Kyle (1985), although our specific implementation is close in spirit to Adrian et al. (2023), in that we estimate price impact separately on a day-by-day

[12] Hautsch and Huang (2012a) and Hautsch and Huang (2012b), among others, show that orders potentially trading through multiple levels of the book tend to be shunned by market participants.

basis based on higher-frequency observations. Inclusion of non-marketable order flow imbalance is important because of previous evidence that order flow imbalance is superior to trade flow imbalance in explaining price changes (see, for example, Cont et al. (2014)).[13]

Appendix B Robustness to Starting the Sample in August 2019

One potential concern with analysis in Section 3 is that there may have been a structural break in liquidity conditions in recent years. Since the event of August 2, 2019, market depth does appear to have been persistently lower than before that event, as shown in the lower panel of Figure 1. That could reflect a secular shift toward a market structure that relies more on fast quote replenishment, rather than a large standing volume of quotes. If were the case, a model that allowed for switching back to states that applied in the early part of the sample may be inefficient. To provide some assurance that such a structural break is not driving the results for recent years, we also consider our preferred HMM, Model IV, estimated using a sample starting after the event of August 2, 2019. Figure B.1 shows the estimated probabilities of being in each state and Table B.1 reports the parameter estimates.

The results are mostly similar to the ones obtained with the model estimated for the full sample. Most of the time, there is a high probability of being in a relatively high-liquidity state in which the sensitivity of price impact to volatility is relatively low. During the market stress of March 2020 and for much of the period since 2022, however, there was a relatively high probability of being in a relatively low-liquidity state with higher sensitivity of price impact to volatility. The positive coefficients on depth in the transition equation imply that a reduction in depth increases the probability of transitioning to the low-liquidity state. Finally, we obtain further evidence of the greater impact of depth on transition probabilities when liquidity is

[13]The results of Cont et al. (2014) were subsequently extended by Xu et al. (2018) and Cont et al. (2021) to incorporate order flow imbalance at multiple levels of the CLOB.

low relative to when liquidity is high, since $\gamma_{1,i} > \gamma_{2,i} > \gamma_{3,i}$ for $i = 2, 3$. Thus, we conclude that the main insights from our full-sample model are robust to estimating an HMM only using more recent data.

Figure B.1: HMM-Implied Probabilities of Being in Each State: Sample Starting in August 2019

The upper panel shows the estimated probabilities of being in each state, as implied by Model IV estimated using data starting in August 2019. The red, grey and blue areas indicate the probabilities of being in the low-, medium- and high-liquidity states, respectively. The lower panel shows a time series of price impact. The vertical dashed lines in the lower panel correspond to: (d) August 2, 2019; (e) March 9, 2020; (f) February 25, 2021; and (g) March 13, 2023.

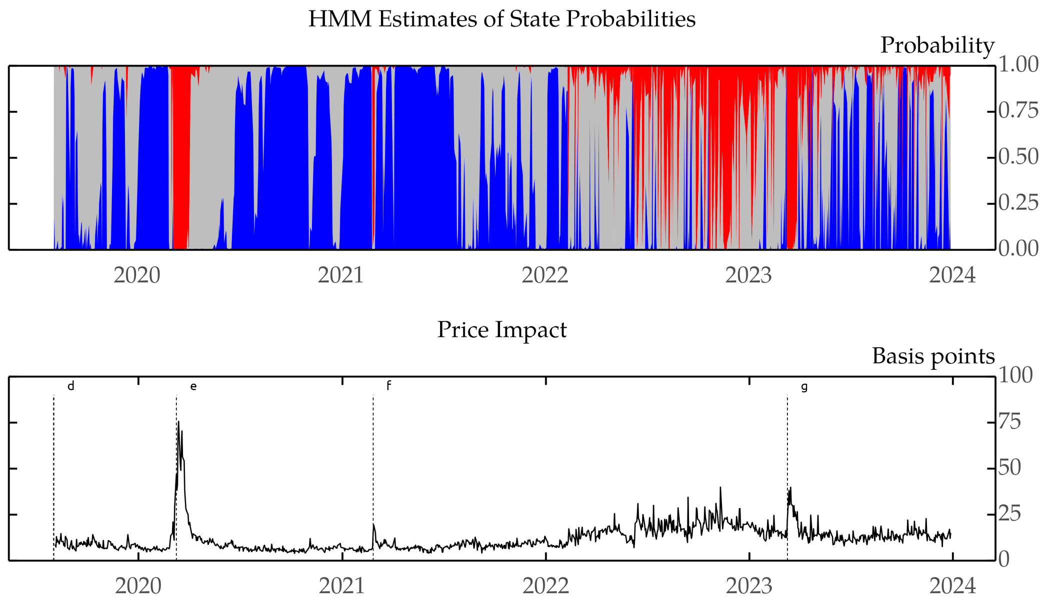

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.

Table B.1: HMM Parameter Estimates: Sample Starting in August 2019

The table reports parameter estimates for the HMM in Equations 2 and 4 using data on price impact, volatility, and market depth for the period from August 3, 2019 to December 31, 2023. Numbers in brackets indicate standard errors based on the Hessian. A $\cdot$, *, or ** indicates two-sided p-values of less than 0.1, 0.05, and 0.01, respectively.

Panel $\mathcal{A}$. Observation equation		Panel $\mathcal{B}$. Transition equation	
β_1^V	0.176**	$\delta_{1,2}$	-11.728*
	(0.006)		(5.178)
$beta_1^{VV}$	1.348**	$\delta_{1,3}$	-148.014
	(0.184)		(193.629)
β_1^{CS}	-0.262*	$\gamma_{1,2}$	0.396*
	(0.111)		(0.176)
σ_1	7.245**	$\gamma_{1,3}$	3.474
	(0.471)		(4.372)
β_2^V	0.131**	$\delta_{2,2}$	-6.316**
	(0.001)		(2.423)
β_2^{VV}	0.464**	$\delta_{2,3}$	-8.549**
	(0.074)		(2.464)
β_2^{CS}	0.119**	$\gamma_{2,2}$	0.203**
	(0.023)		(0.066)
σ_2	1.661**	$\gamma_{2,3}$	0.213**
	(0.084)		(0.067)
β_3^V	0.099**	$\delta_{3,2}$	-4.543.
	(0.001)		(2.438)
β_3^{VV}	0.223**	$\delta_{3,3}$	-7.125**
	(0.083)		(2.302)
β_3^{CS}	0.111**	$\gamma_{3,2}$	0.104*
	(0.013)		(0.047)
σ_3	1.067**	$\gamma_{3,3}$	0.164**
	(0.051)		(0.045)

Sources: Repo Inter Dealer Broker Community; TP ICAP, Swaptions and Interest Rate Caps and Floors Data; authors' calculations.