Corporate Bond Liquidity During the COVID-19 Crisis*

Mahyar Kargar† Benjamin Lester‡ David Lindsay§ Shuo Liu¶
Pierre-Olivier Weill‖ Diego Zúñiga**

February 5, 2021
First Draft: April 16, 2020

Abstract

We study liquidity conditions in the corporate bond market during the COVID-19 pandemic. We document that the cost of trading immediately via risky-principal trades increased dramatically at the height of the sell-off, forcing customers to shift towards slower, agency trades. Exploiting eligibility requirements, we show that the Federal Reserve's corporate credit facilities had a positive effect on market liquidity. A structural estimation reveals that customers' willingness to pay for immediacy increased by about 200 bps per dollar of transaction at the height of the crisis, and quickly subsided to pre-crisis levels. Dealers' marginal cost increased even more, and did not fully subside.

KEYWORDS: Corporate bonds, liquidity, intermediation, SMCCF, COVID-19
JEL CLASSIFICATION: G12, G14, G21.

---

*We thank Darren Aiello, Roc Armenter, Mitchell Berlin, John Cochrane, Nathan Foley-Fisher, Itay Goldstein, Valentin Haddad, Ricardo Lagos, Borghan Narajabad, Avanidhar Subrahmanyam, Stéphane Verani, James Vickery, Jonathan Vogel, Bin Wei, Haoxiang Zhu, two anonymous referees, and, seminar participants at the Journal of Finance and the Fama-Miller Center virtual conference on the financial consequences of the COVID-19 pandemic, Atlanta Fed/GSU virtual conference on financial stability and the coronavirus pandemic, Search and Matching in Macro and Finance virtual workshop on liquidity in fixed income markets, University of Illinois, UCLA Anderson, and UCLA economics for comments and suggestions. Yiling Pan provided expert research assistance. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. The first version of this paper was circulated on April 16th, 2020.

†University of Illinois at Urbana-Champaign, kargar@illinois.edu, website
‡Federal Reserve Bank of Philadelphia, benjamin.lester@phil.frb.org, website
§University of California, Los Angeles, lindsayd@ucla.edu
¶Tsinghua University, School of Economics and Management, liushuo3@sem.tsinghua.edu.cn, website
‖University of California, Los Angeles, NBER, and CEPR, poweill@econ.ucla.edu, website
**University of California, Los Angeles, zung@g.ucla.edu

1 Introduction

The COVID-19 pandemic induced an unprecedented shock to the global economy. As the implications of this shock began to crystalize in mid-March, 2020, financial markets plummeted and reports of illiquidity began to surface. One particularly important market that was "under significant stress" (Bernanke and Yellen, 2020) was the $10 trillion corporate bond market, which a March 18 report from Bank of America deemed "basically broken" (Idzelis, 2020). In response, the Federal Reserve introduced several facilities designed to bolster liquidity and reduce the costs and risks of intermediating corporate debt, including the Primary Dealer Credit Facility (PDCF) and the Primary and Secondary Market Corporate Credit Facilities (PMCCF and SMCCF, respectively). The latter two facilities represented a particularly bold intervention, in that they allowed the Fed, for the first time, to make outright purchases of investment-grade corporate bonds issued by US companies, along with exchange-traded funds (ETFs) that invested in similar assets.

The purpose of this paper is to study trading conditions in the US corporate bond market in response to the large economic shock induced by COVID-19, as well as the unprecedented interventions that followed. Given the exogenous nature of the pandemic, set against the backdrop of a well-capitalized financial sector, this episode offers a unique opportunity to identify the nature of shocks that precipitate illiquidity in financial markets, the consequences for transaction costs and consumer surplus, and the efficacy of various policy responses designed to restore liquidity in times of distress.

A central feature of our analysis is the distinction between two types of transactions offered by dealers: "risky-principal" trades, in which a dealer offers a customer-seller immediacy by purchasing the asset directly and storing it on his balance sheet until finding a customer-buyer; and "agency" trades, in which the customer-seller retains the asset while waiting for a dealer to find a customer-buyer to take the other side of the trade. This distinction—which has been studied recently using pre-pandemic data—is crucial for a number of our key insights.[1] We highlight four.

[1] For recent work that studies the distinction between risky-principal and agency trades, see Schultz (2017), Bao, O'Hara, and Zhou (2018) , Choi and Huh (2018), Bessembinder, Jacobsen, Maxwell, and Venkataraman (2018), and

First, distinguishing between risky-principal and agency trades provides a more complete assessment of market liquidity by accounting for both the cost of trading and the time it takes to trade. Indeed, we show that focusing on transaction costs alone understates the deterioration in liquidity after the COVID-19 shock. Second, studying the cost and quantity of these distinct types of trades—in concert with a structural model—allows us to disentangle two widely cited (but not mutually exclusive) sources of illiquidity: a large, unexpected increase in customers' demand for immediacy, sometimes called a "dash for cash"; and a decrease in dealers' willingness to supply immediacy by absorbing assets onto their balance sheets, either because of rising costs or binding regulatory constraints. A key finding is that matching the data requires large shocks to *both* demand and supply at the onset of the crisis. Third, studying the evolution of demand and supply factors against the timeline of the Fed's interventions offers new insights into the efficacy of various policies. In particular, we show that the surge in customers' demand for immediacy receded almost immediately, and fully, after the announcement of the Fed's interventions, whereas the negative shock to dealers' supply of immediacy responded more gradually, and only partially. Finally, our analysis of these differentiated transaction services highlights the important—but, perhaps, overlooked—point that the consumer surplus of market participants is not well-approximated by average transaction costs. In particular, we provide a quantitative estimate of the change in consumer surplus throughout the crisis, and demonstrate the importance of accounting for changes in customers' preferences, along with changes in the mix of risky-principal and agency trades.

After providing some background information in Section 2, we begin our analysis in Section 3 by documenting trading conditions in the corporate bond market in response to the panic of mid-March and the Fed's interventions that followed. Using data from the Trade Reporting Compliance Engine (TRACE), we first construct time series to measure the costs of risky-principal and agency trades in the corporate bond market. We find that the cost of risky-principal trades increased significantly during the COVID-induced panic, reaching a peak of more than 250 basis points (bps), while the cost of agency trades increased much more modestly. As the premium paid for

Goldstein and Hotchkiss (2020). To the best of our knowledge, we are the first (and only) paper to employ this distinction to study liquidity conditions during the COVID-19 pandemic.

risky-principal trades increased, we show that customers substituted towards agency trades: the fraction of total volume executed as agency trades increased by as much as 15% at the height of the sell-off, and remained elevated even months after the initial panic subsided. Hence, the average trade was not only more expensive, but also more likely to be slower or of "lower quality."

As trading shifted from risky-principal to agency transactions, we show that the dealer sector as a whole absorbed *no* inventory, on net, during the most tumultuous period of trading. Therefore, when the demand for transaction services surged, it was customers themselves who ultimately stepped up to provide additional liquidity. In fact, it was only after the announcement of the Federal Reserve's interventions that dealers began to absorb inventory onto their balance sheets, and trading conditions started to improve. Indeed, after the announcement of the Fed's credit facilities, the quantity of corporate debt held by dealers more than doubled relative to pre-COVID levels. At the same time, the cost of risky-principal trades decreased significantly, but remained approximately twice the levels observed before the pandemic.

While these observations establish the coincidence of key interventions and improvements in market liquidity, they do not establish a causal relationship. To further explore the effects of interventions on market liquidity, we exploit restrictions on the types of bonds that could be purchased through the Fed's corporate credit facilities. In particular, using a difference-in-differences approach, we use restrictions on bond ratings and time-to-maturity to identify the change in trading costs induced by the announcement of the SMCCF. We find that, immediately after the announcement of the SMCCF, the cost of trading bonds that were eligible for purchase by the Fed decreased substantially relative to the cost of trading ineligible bonds. Later, when the program was expanded in both size and scope, we show that the trading costs of all bonds fell.

Hence, our findings suggest that the Fed's interventions had significant effects on transaction costs and trading activity in the corporate bond market. However, the observations described above also lead to several important questions. Why did the Fed's interventions improve trading conditions so quickly, but not fully? Did the announcement and implementation of these policies (at least partially) restore liquidity by easing investors' concerns and halting the "dash for cash"?

Or should the efficacy of these interventions be attributed to easing dealers' balance sheet concerns, thus increasing their willingness to "lean against the wind" (Weill, 2007)? Given the deterioration of both the cost *and quality* of intermediation services during the COVID-19 crisis, what was the effect on the surplus of customers in the US corporate bond market during this period?

To confront these questions, and interpret our empirical findings more generally, in Section 4 we construct a parsimonious equilibrium model of a market for vertically differentiated transaction services: low-quality, meant to capture agency trades; and high-quality, meant to capture risky-principal trades. We assume that customers prefer high-quality transaction services, but they are more costly for dealers to produce. Within this framework, we characterize the differential impact of two types of shocks—to customers' relative demand for high-quality risky-principal trades and to dealers' cost of supplying these transaction services—on equilibrium prices and allocations.

Then, using our estimates of relative prices and quantities in concert with our theoretical framework, we estimate key parameters of the model, which allows us to identify shocks to customers' demand for immediacy, at the height of the crisis and during the interventions that followed. We confirm that a large, sudden increase in the demand for immediacy was a crucial source of illiquidity early in the crisis. In fact, we estimate that customers' willingness to pay for each inframarginal unit of risky-principal trade (rather than an agency trade) increased by approximately 200 bps at the height of the crisis. However, we also find that this shock alone cannot explain what we observe in the data: to rationalize the observation that customers ultimately substituted towards agency trades, we show that there must have also been an even larger shock to dealers' marginal cost of supplying immediacy. Hence, understanding the market turmoil of March 2020 requires studying both the origins of the "dash for cash" *and* the factors that dissuaded dealers from absorbing selling pressure onto their balance sheets.

Studying the behavior of shocks to customers' demand for immediacy, and dealers' cost of supplying it, against the timeline of policy announcements and implementation also reveals new insights regarding the channels through which the Fed's interventions operated. In particular, we find that the demand shock receded quickly, and fully, soon after the *announcement* of the PDCF and SMCCF—that is, the announcement alone seems to have effectively reversed the initial "dash for cash." The supply shock, however, appears to have lingered even after the Fed began purchasing bonds. While there are multiple explanations for this, we document one plausible candidate: the total volume of customer-dealer transactions remained elevated through June which—when combined with binding balance sheets constraints—could explain why the relative cost (fraction) of risky-principal trades remained elevated (depressed) months after markets appear to have calmed.

Finally, we leverage our theoretical framework to map easily quantifiable objects, such as trading costs and the fraction of risky-principal trades, into an object that is harder to measure but important for policymakers—namely, the consumer surplus of market participants. We document that the surplus per unit of transaction experienced a significant decline at the height of the market turmoil, a fall of about 30 bps, and remained approximately 10 bps below pre-crisis levels even at the end of June. Our results highlight the fact that consumer surplus does not track average transaction costs, as the latter abstracts from the effects of changes in preferences and changes in the composition of risky-principal and agency trades. In particular, compared with average transaction costs, the fall in consumer surplus was more modest at the height of the crisis, but more pronounced after the initial shock subsided.

1.1 Related literature

Given the size of the COVID-19 shock, and the historic nature of the Fed's response, it is not surprising that a number of recent papers have emerged to study financial markets since the onset of the pandemic. Our paper belongs to the literature focused on the corporate bond market, which we discuss in more detail below, but shares much in common with studies of other markets, including the market for Treasuries and other government debt (Duffie, 2020; He, Nagel, and Song, 2020; Fleming and Ruela, 2020; Schrimpf, Shin, and Sushko, 2020), as well as the market for asset-backed securities (Foley-Fisher, Gorton, and Verani, 2020; Chen, Liu, Sarkar, and Song, 2020).

In the corporate bond market, Falato, Goldstein, and Hortaçsu (2020) study the effect of the pandemic on outflows from bond mutual funds, and the role that the Fed's corporate credit facilities played in reversing these outflows. Ma, Xiao, and Zeng (2020) also explore outflows in fixed-income mutual funds, including those that invest in corporate bonds and Treasuries. They derive a pecking order theory of liquidation, which explains why selling pressure was strongest in the most liquid sectors of these markets. Haddad, Moreira, and Muir (2020) focus primarily on the behavior of credit spreads during the crisis, and attempt to identify the mechanism through which the Fed's interventions improved market conditions.[2] Though different along many dimensions, these three papers all argue that a large, sudden increase in customers' demand for immediacy played a crucial role in the deterioration of market liquidity in March, 2020. We, too, identify such a shock, but find that matching the data also requires a significant shock to the dealers' cost of supplying immediacy.

Our paper is most closely related to contemporaneous work by O'Hara and Zhou (2020) and Boyarchenko, Kovner, and Shachar (2020), who also investigate liquidity conditions in the corporate bond market during the COVID-19 crisis, and the effects of the Fed's interventions.[3] Despite some overlap, the three papers differ (and complement one another) in several important ways. For example, using the regulatory version of TRACE—which contains dealer identities—O'Hara and Zhou (2020) document the heterogeneous response of different dealers to the Fed's interventions. This allows them to control for dealer fixed effects and to disentangle the effects of the PDCF and the SMCCF, among other things. Boyarchenko et al. (2020) also use the regulatory version of TRACE, along with data on the volume of bonds (or shares of ETFs) purchased by the Fed's corporate credit facilities. This allows them to decompose the effects of the Fed's interventions into direct "purchase effects" and indirect "announcement effects."[4]

While our paper makes a number of distinct contributions relative to these contemporaneous studies, we highlight several aspects of our analysis that are particularly important. First, our approach to measuring trading conditions accounts for two channels through which market liquidity

[2]For related work on the behavior of credit risk/spreads throughout the crisis, and the effects of the Fed's interventions, see Nozawa and Qiu (2020) and D'Amico, Kurakula, and Lee (2020).

[3]In more recent work, Gilchrist, Wei, Yue, and Zakrajšek (2020) quantify the effects of the SMCCF on credit spreads and transaction costs using a regression discontinuity approach. Using a different methodology to construct their sample, they find qualitatively similar results regarding the effects of the SMCCF on transaction costs for eligible and ineligible bonds, though their quantitative magnitudes differ from ours. We discuss this further in Section 3.

[4]For more on the purchase effects of the SMCCF, see Flanagan and Purnanandam (2020).

can deteriorate—customers can face higher transaction costs or longer waiting times for executing a trade—and hence provides a multi-dimensional assessment of market conditions during the crisis. Second, in contrast with the papers cited above, we develop a theoretical framework that, when combined with our empirical estimates, allows us to construct quantitative estimates of the shocks that precipitated the COVID-19 crisis.[5] Finally, we use these estimates of shocks to demand, along with bounds on shocks to supply, to study the efficacy of various policy interventions, and the implications for consumer surplus, at the height of the crisis and beyond.

2 Background

The COVID-19 Shock. Despite reports of a potentially lethal virus spreading in China, US equity markets reached all-time highs on February 19, 2020. Just two weeks later, as the scope of the COVID-19 coronavirus and the duration of its effects became apparent, financial markets around the world entered a period of turmoil. For example, between March 5 and March 23, the S&P 500 fell more than 25%. In the corporate bond market, the ICE Bank of America AAA US Corporate Index Option-Adjusted spread increased by about 150 bps over this same period, while the corresponding spread for high-yield (HY) corporate debt increased by more than 500 bps.[6] As the price of equities and debt plummeted, reports of illiquidity in key financial markets emerged. Such reports were especially troubling in the corporate bond market, as many large US firms would almost surely need access to capital in light of the impending shocks to their balance sheets.[7]

Two complementary factors were cited as the root of the panic in the corporate bond market. The first was a surge in the demand for immediacy, or so-called "dash for cash," as investors pulled out of corporate bond funds in droves. For example, Falato et al. (2020) report that, between the

[5] Along this dimension, our paper is related to Goldberg and Nozawa (2020), who use a structural VAR approach to identify demand and supply shocks in the corporate bond market during (and after) the 2007-2009 financial crisis. However, our identification strategy is different from theirs, as is our focus: their primary concern is the asset pricing implications of liquidity supply shocks.

[6] See Ebsim, Faria-e Castro, and Kozlowski (2020) for a more comprehensive analysis of credit spreads during this time period.

[7] Indeed, Darmouni and Siani (2020) document that corporate bond issuance reached historic levels in the Spring of 2020, after the Fed's interventions, despite a relatively healthy banking sector.

months of February and March, the average corporate bond fund experienced cumulative outflows of approximately 9% of net asset value—by far the largest outflows in the last decade. At the same time, market participants reported that dealers were either unable or unwilling to supply customers with immediacy by absorbing corporate debt onto their balance sheet. In a *Wall Street Journal* article titled "The Day Coronavirus Nearly Broke the Financial Markets," Baer (2020) writes about the experience of Vikram Rao, the head bond trader of Capital Group, after calling senior executives for an explanation on why broker-dealers wouldn't trade:

> [T]hey had the same refrain: There was no room to buy bonds and other assets and still remain in compliance with tougher guidelines imposed by regulators after the previous financial crisis [...] One senior bank executive leveled with him: "We can't bid on anything that adds to the balance sheet right now."

Federal Reserve Interventions. In response to signs of illiquidity in several key financial markets, the Federal Reserve introduced a number of new facilities designed to bolster liquidity and reduce trading costs. On the evening of March 17, the Federal Reserve revived the aforementioned PDCF, offering collateralized overnight and term lending to primary dealers. By allowing dealers to borrow against a variety of assets on their balance sheets, including investment-grade corporate debt, this facility intended to reduce the costs associated with holding inventory and intermediating transactions between customers.[8]

On March 23, the Federal Reserve proposed even more direct interventions in the corporate bond market through the PMCCF and SMCCF. These facilities were designed to make outright purchases of corporate bonds issued by investment-grade US companies with remaining maturity of five years or less. The facilities were also allowed to purchase shares in US-listed exchange-traded funds (ETFs) that invested in US investment-grade corporate bonds. On April 9, these

[8] In addition to the facilities that we highlight in our analysis here, it is also noteworthy that the Federal Reserve temporarily relaxed the supplementary leverage ratio (SLR) rule—first on April 1 and again on May 15, 2020—to ease balance sheet constraints and increase banks' ability to lend to households and businesses. By excluding US Treasury securities and reserves from the calculation of the SLR rule for holding companies, the rule change was primarily intended to increase liquidity in the Treasury market. However, to the extent that it relaxed dealers' balance sheet constraints, the effects could clearly extend to the corporate bond market as well, as we discuss later in the text. To read more about the rule change, see the April 1, 2020 and the May 15, 2020 press releases.

corporate credit facilities were expanded in size and extended to allow for purchases of ETFs that invested in high-yield corporate bonds.[9] Interestingly, though many of the effects of these corporate credit facilities were observed immediately after they were announced (and expanded), the Federal Reserve did not actually begin purchasing bonds until May 12. We provide a more detailed description of this timeline, and of the Federal Reserve's corporate facilities, in Appendix B.

3 Trading Conditions During the Pandemic

In this section, we describe how market conditions evolved from the sanguine conditions of mid-February through the freefall of mid-March to the post-intervention recovery of April and May. As a first step, we construct time series for several variables of interest: the cost of risky-principal trades, the cost of agency trades, and the fraction of each type of transaction services. We document that, at the height of the selling pressure, the cost of risky-principal trades surged and the fraction of such trades dropped significantly. Conditions improved immediately after the Fed's announcement of the corporate credit facilities, with dealers providing liquidity directly, via risky-principal trades, at significantly lower prices. To test the causal relationship between the Fed's interventions and market liquidity, we exploit the eligibility requirements for bond purchases by the SMCCF. We find that, after the initial announcement, trading costs for eligible bonds fell substantially more than trading costs for ineligible bonds. Later, after the program was expanded in both size and scope, we document more significant declines in trading costs for all bonds.

3.1 Data

We combine the standard TRACE data set (for 2020Q1) with the End-of-Day version (for 2020Q2). We first filter the report data following the standard procedure laid out in Dick-Nielsen (2014). We merge the resulting data set with the TRACE master file, which contains bond grade informa-

[9]The April 9 update also allowed the SMCCF to make direct purchases of bonds that had been downgraded from investment-grade to high-yield status (so-called "fallen angels") after March 22. The facility also allowed purchasing of high-yield ETFs.

tion, and with the Mergent Fixed Income Securities Database (FISD) to obtain bond fundamental characteristics. Following the bulk of the academic literature, we exclude variable-coupon, convertible, exchangeable, and puttable bonds, as well as asset-backed securities, and private placed instruments. We also exclude newly-issued and foreign securities.

For most of our analysis, we use the (filtered) data covering the period from January 2 to June 30, 2020, which contains 7.2 million trades and 30,748 unique bonds. Approximately 60% of the transactions are identified as customer-dealer and 40% as interdealer trades. The average trade size is $218,104 across all transactions, with average total daily volumes for customer-dealer and interdealer trades of $7.25 billion and $3.13 billion respectively. It is worth noting that, in both the standard and End-of-Day versions of TRACE, the trade size for investment-grade and high-yield bonds is top-coded at $5 million and $1 million, respectively.[10]

In all of our plots below, we include vertical dashed lines to highlight several key dates: February 19, when stock markets reached their all-time peaks; March 5, which marks the beginning of the extended fall in equity prices and rise in corporate credit spreads; March 18, the first day of trading after the announcement of the PDCF; March 23, the day that the PMCCF and SMCCF were announced; April 9, the day that the size and scope of the corporate credit facilities were expanded; May 12, the date that the SMCCF started buying bond ETFs; June 16, the day that the SMCCF began purchasing individual bonds; and June 29, the date the PMCCF began operating.

3.2 The cost of trading, fast and slow

To capture the average transaction cost for risky-principal trades, we use the measure of bid-ask spreads proposed by Choi and Huh (2018), CH hereafter. To construct this measure, we first calculate, for each customer trade, the spread

$$2Q \times \frac{\text{traded price} - \text{reference price}}{\text{reference price}},$$

[10]Table A1 in Appendix A presents additional summary statistics for our sample.

where Q is equal to $+1$ (-1) when a customer buys from (sells to) a dealer, and the reference price is taken to be the volume-weighted average price of interdealer trades larger than $100,000 in the same bond-day. Importantly, we restrict our sample so that it only includes trades in which the dealer who buys the bond from a customer holds it for more than 15 minutes. In doing so, we leave out those trades where the dealer had pre-arranged for another party (either a customer or another dealer) to buy the bond immediately.[11] The measure of risky-principal trading costs is aggregated at the bond-day level by taking the volume-weighted average of trade level spreads, and then at the daily level by taking the average in each day across all bonds, weighted by bonds' daily total volume of customer trades where the CH measure is available.

To capture the average transaction cost of agency trades, we calculate a modified version of the Imputed Roundtrip Cost measure described in Feldhütter (2012). To construct this modified imputed roundtrip cost (or "MIRC"), we first identify imputed roundtrip trades (IRT) by matching a customer-sell trade with a customer-buy trade of the same size that takes place within 15 minutes of each other.[12] We exclude interdealer trades in constructing IRTs, so that each IRT only includes one customer-buy trade and one customer-sell trade. Then, to compute the MIRC, we calculate

$$\frac{P_{max} - P_{min}}{P_{max}},$$

where P_{max} (P_{min}) is the largest (smallest) price in the IRT. Within each bond and day, we calculate the daily average roundtrip cost as the average of the bond's MIRC on that day, weighted by trade size. Finally, a daily estimate of average roundtrip cost is the average of roundtrip costs on that day across all bonds, weighted by bonds' total daily trading volumes in the matched IRTs.

Figure 1 plots the two time series, along with the difference between the two. The two measures of transaction costs are relatively stable through February 19, with risky-principal trades

[11]Likewise, in calculating reference prices, we follow CH and exclude interdealer trades executed within 15 minutes of a customer-dealer trade.

[12]In other words, as in earlier papers, we assume that customer-buys and customer-sells that occur in rapid succession are likely to be agency trades. Indeed, in an agency trade, dealers search for counterparties on behalf of customers. When counterparties are found, the two customers are matched by dealers, and two customer-to-dealer trades are recorded in a short time window.

approximately twice as expensive as agency trades. Upon realization of the COVID-induced shock, the cost of risky-principal trades rises dramatically, while the cost of agency trades is more muted. In particular, between Thursday, March 5, and Monday, March 9, the cost of risky-principal trades roughly triples, to approximately 100 bps; over these three trading days, the S&P 500 Index declined more than 12%. A week later, during the most tumultuous period of March 16-18, this series continues to rise, reaching a peak of more than 250 bps, before beginning a steady decline after the announcement of the SMCCF on March 23. The MIRC measure of agency trading costs, in contrast, increases from a baseline around 8 bps to approximately 28 bps, before receding slightly after the Fed's intervention.

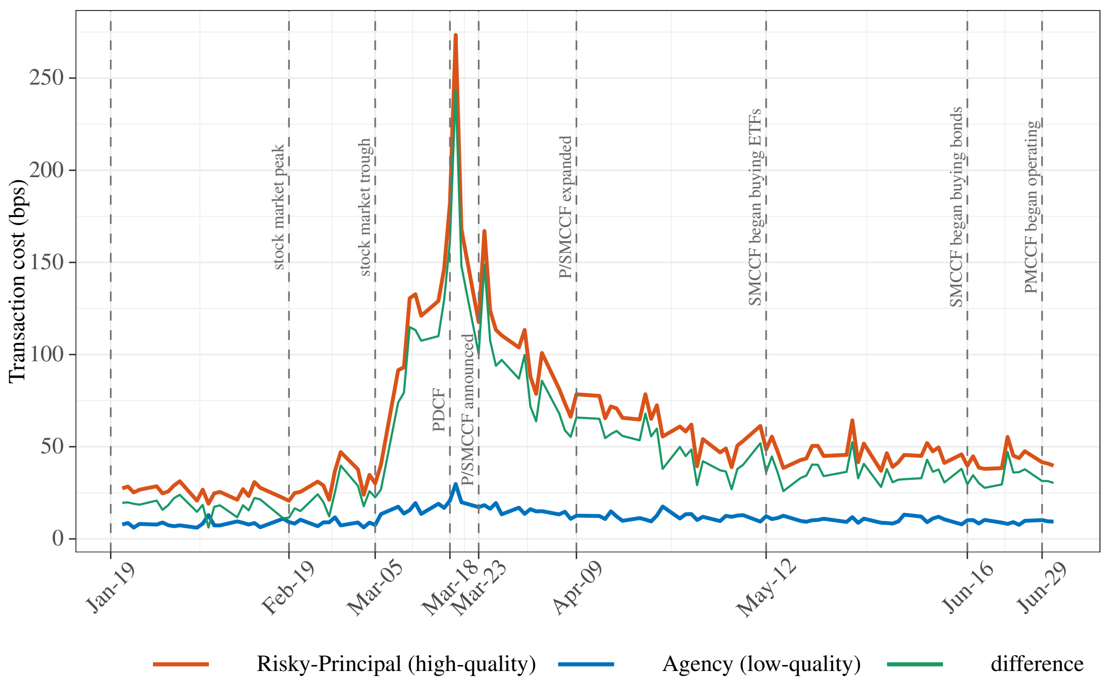

Figure 1. Transaction costs: Risky-principal vs. agency trades. This figure shows the time-series of trading costs for risky-principal trades in red and agency trades in blue, and their difference in green.

One can see that the cost of risky-principal trades, which we interpret as the cost of trading immediately, was considerably more responsive to both the heightened selling pressure induced by the pandemic in mid-March and the Fed's interventions which followed. Moreover, despite considerable improvement in both metrics during the month of April, the cost of risky-principal trades remained elevated through June, which suggests that liquidity conditions remained somewhat strained well after markets appear to have calmed.

Of course, the change in spreads could be driven by a change in the composition of bonds that were traded during this period of distress. For example, perhaps trading volume was unusually high for retail-size trades of illiquid bonds, which typically involve higher transaction costs. Thus, to further clarify the impact of the crisis and ensuing interventions on the cost of risky-principal and agency trades, we turn to formal regressions that allow us to control for bond- and trade-level fixed effects. We consider the following specification

$$y_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}. \tag{1}$$

The dependent variable, y_{ijt}, represents the transaction cost for a type $j \in \{\text{risky-principal, agency}\}$ trade of bond i on day t. The dummy variables Crisis_t and Intervention_t allow us to distinguish between three sub-periods: (i) Pre-crisis, which corresponds to dates before March 5, 2020; (ii) Crisis, which covers the period March 5–23, 2020; and (iii) Intervention, which covers the period after March 23. Hence, the coefficients β_1 and β_2 measure transactions costs relative to the pre-crisis period. Finally, α_i and α_s represent bond and trade size fixed effects, respectively. Bond fixed effects capture bond characteristics that are fixed over time such as industry, par amount, etc.[13] For trade size fixed effects we consider three categories: less than $100,000, between $100,000 and $1 million, and larger than $1 million.

Table 1 presents results for all bonds, as well as the sub-sample of bonds issued by US firms.[14] We include bond and size category fixed effects and cluster standard errors at the bond and day levels in all regressions to account for correlation over time within a bond and across bonds in a given day. Columns (1) and (3) reveal that, during the crisis period of March 5-23, average bond-

[13]We do not have access to the latest credit rating data for all bonds in our sample, just the binary IG/HY classification provided by TRACE. For the sub-sample of bonds where the credit rating is available, we include a credit rating fixed effect in specification (1) to control for potentially time-invariant nature of bond credit ratings. From Table A2 in Appendix C, we see that controlling for bond credit rating leads to very similar results to those in Table 1.

[14]One reason we include the results for the US sub-sample is to demonstrate that the trading cost patterns are similar to the full sample. This is helpful later, in Section 3.5, when we focus on the US sub-sample exclusively.

level trading costs for risky-principal and agency trades increased by approximately 105 bps and 9 bps, respectively, relative to the pre-crisis period. After the Fed's interventions on March 23, trading costs for risky-principal trades fell by approximately 64 bps—more than half the initial spike—while transaction costs for agency trades declined much more modestly. These results are consistent with the aggregate results in Figure 1. Columns (2) and (4) show that the sub-sample of US-issued bonds exhibits roughly the same behavior as the sample of all bonds, though the cost of agency trades for US-issued bonds increased slightly more during the crisis period.

3.3 Substituting agency trades for risky-principal trades

We now establish that, as the premium for risky-principal trades increased, customers responded by substituting towards agency trades. Figure 2 plots the proportion of agency trades by number (left axis) and volume (right axis).[15] During the most tumultuous weeks of trading, between March 5 and March 23, the fraction of agency trades (measured by both number and volume) increased by as much as 15 percentage points, trough to peak, before receding after the March 23 announcement of the corporate credit facilities. Again, this shift toward agency trades has important implications for assessing market liquidity. In particular, if one were simply to measure trading costs across all trades, they would underestimate the erosion in liquidity as the composition of trades shifted from faster, more expensive risky-principal trades to less costly, but slower agency trades.

To study the substitution from risky-principal to agency trades more carefully, we consider a regression with the following specification:

$$\text{Agency}_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}, \tag{2}$$

where Agency_{ijt} is an indicator variable that takes the value one if trade j for bond i on day t is an agency trade and zero otherwise. The variables on the right-hand side of specification (2) are the same as in (1). Under this specification, the coefficients β_1 and β_2 measure the change in the

[15]We discuss how we identify agency trades in depth in Appendix A.

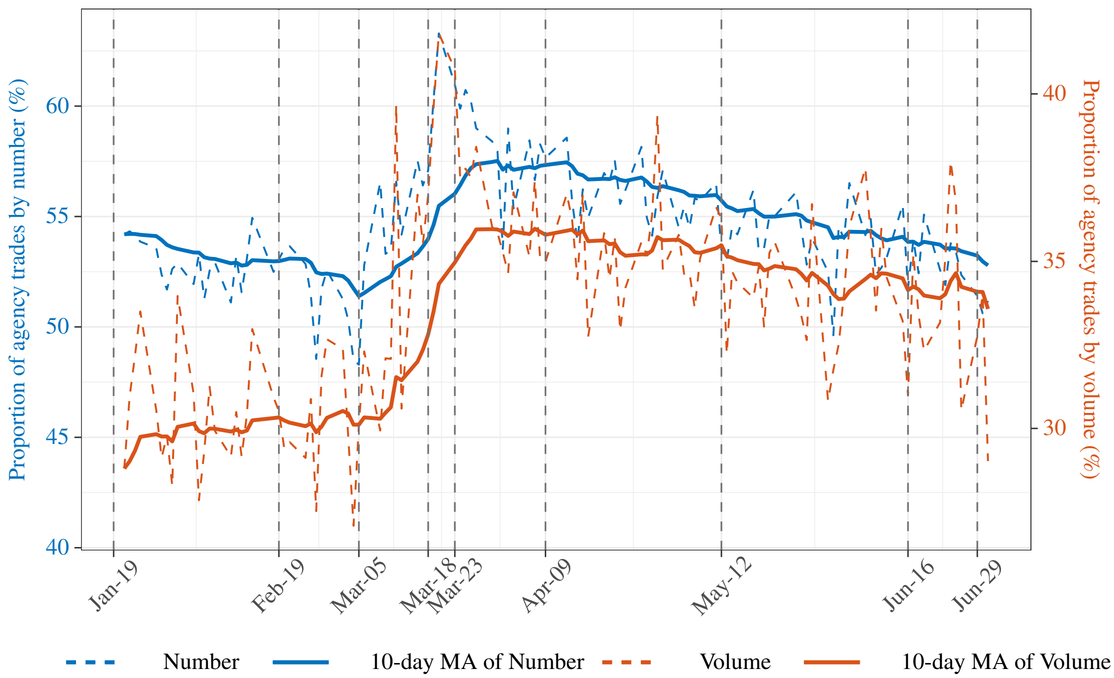

Figure 2. Proportion of agency trades. This figure plots the fraction of agency trades by volume in red (right axis) and by number in blue (left axis).

probability of an agency trade during the crisis and intervention periods, respectively, relative to the pre-crisis period. Table 2 presents results using a linear probability model (OLS), along with logit and probit specifications for robustness.

Column (1) reveals that, during the crisis period of March 5–23, the probability of an agency trade for a given bond, on average, rose by 4.3 percentage points relative to the pre-crisis period. After the Fed interventions on March 23, this probability decreased from the crisis period (by 160 bps) to 2.7 percentage points higher than the pre-crisis period. For the sake of completeness, we report marginal effects calculated at the sample means for logit and probit models in columns (2) and (3); the results are very similar to the linear probability model (OLS) in column (1).[16]

[16]For the interested reader, we also report results from a linear probability model that distinguishes between eligible and ineligible bonds for the SMCCF in Appendix C. We find that the shift towards agency trades was more pronounced among bonds that were eligible for the Fed's purchasing program.

3.4 Dealers' inventory accumulation

As the relative price of risky-principal trades spiked in mid-March, and customers substituted towards agency trades, one might naturally wonder: who was providing liquidity in the corporate bond market? Were dealers "leaning against the wind" and absorbing some of the inventory during the selloff? Or was the shift to agency trades sufficiently large that other customers were ultimately providing liquidity? To answer this question, we construct a measure of the (cumulative) value of bonds that were absorbed over time by the dealer sector. In particular, using the daily Market Sentiment data from FINRA, we subtract the value of bonds that dealers sell to customers from the value of bonds that they buy from customers each day, and then calculate the cumulative sum of the net changes.[17] Figure 3 plots the cumulative net change in inventory held in the dealer sector, both in levels (left axis) and as a fraction of pre-crisis outstanding supply (right axis).

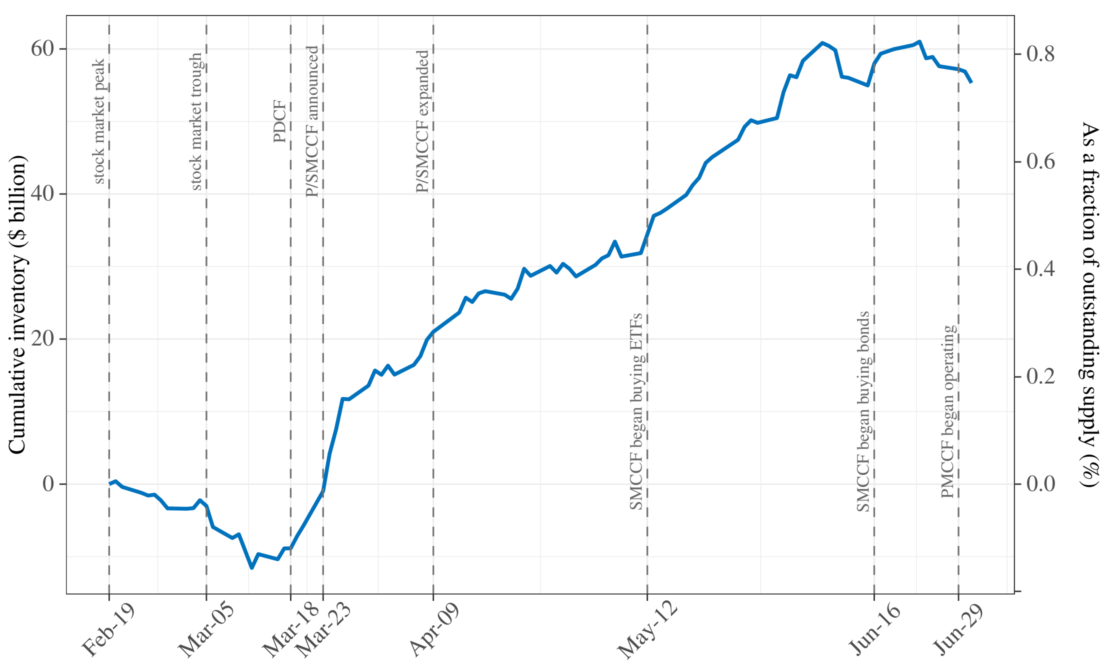

Figure 3. Cumulative inventory change in the dealer sector. This figure plots the cumulative inventory change in the dealer sector in billions of USD (left axis) and as a fraction of total supply in % (right axis). Source: FINRA market sentiment tables.

[17]The Market Sentiment data is available through FINRA TRACE Market Aggregate Information. We use this data, as opposed to the standard or End-of-Day TRACE data, because it is not top-coded and hence allows for a more accurate assessment of the inflow and outflow of bonds in the dealer sector.

Several aspects of Figure 3 are striking. First, during the most tumultuous period of trading, the dealer sector absorbed, on net, *no* additional inventory despite the considerable selling pressure from customers. In fact, dealers actually *reduced* inventory holdings and became net sellers. Hence, during this period, it was indeed other customers that were supplying liquidity to the market. Second, dealers' reluctance to absorb inventory appears to have changed substantially around the dates corresponding to the Fed's announcement of the Primary Dealer Credit Facility (March 18) and the Primary and Secondary Market Corporate Credit Facilities (March 23). Lastly, dealers continued to accumulate inventory through April and May. Indeed, from March 18, the data indicates that dealers absorbed more than $50 billion in corporate debt, or roughly doubled their inventory holdings relative to pre-pandemic levels.[18]

3.5 The effects of the Fed's intervention

The results above suggest that the Fed's interventions—in particular, the March 23 announcement of the SMCCF—had a significant effect on dealers' willingness to absorb inventory onto their balance sheets, and hence on market liquidity. In this section, we exploit the eligibility requirements specified in the SMCCF to test this hypothesis more formally.

According to the original term sheet, a bond is eligible to be purchased through the SMCCF if it has an investment-grade rating on March 23, 2020; if it has a time-to-maturity of five years or less; and if its issuer is domiciled in the US.[19] However, the Fed has a considerable degree of discretion to determine whether a foreign issuer is domiciled in the US. Indeed, in the Fed's SMCCF transaction-level disclosures, we found many cases in which the holding firm of the

[18]From Table L.130 of the Flow of Funds, at the end of 2019Q4, security brokers and dealers held $54 billion in corporate and foreign bonds on the asset side of their balance sheets.

[19]The original March 23 term sheet can be found here. Initially, there was an additional eligibility criterion for the SMCCF on March 23: eligible issuers excluded firms that were expected to receive direct financial assistance from the then-pending CARES Act. This criterion (and others) were later added to the SMCCF term sheet on April 9. See Appendix B for more details.

security is a non-US entity.[20] Given this lack of clarity, we chose to focus on US firms exclusively, and classify a bond as eligible based on credit rating and time-to-maturity alone.[21]

To start, we repeat the regression specified in (1) with two modifications. First, we separate the sample of bonds into those that were eligible for purchase through the SMCCF and those that were not. Second, we separate the intervention period into two sub-periods. The first sub-period, which we call the "SMCCF," covers from March 23-April 8, 2020. During this period, it appeared that only investment-grade bonds would be eligible for purchase. The second sub-period, which we call the "SMCCF expansion," starts on April 9, when the Fed announced that it was increasing the size of the program and expanding the set of eligible bonds to include high-yield debt.

Table 3 reports the results. Column (2) reveals that the initial decline in trading costs was largely driven by bonds that were eligible for the SMCCF: the price of risky-principal trades for ineligible bonds declined much more modestly immediately after the March 23 announcement, relative to the crisis period, while the price of agency trades for ineligible bonds actually increased during this time period. After the program was expanded on April 9, in both scope and size, the price of risky-principal trades for all bonds declined significantly.

To further explore the causal effect of the SMCCF on bond market liquidity during the crisis, we consider a difference-in-differences regression over a sub-sample of our data from March 6 to April 9, 2020. These dates are chosen to exclude the pre-crisis period, when spreads were very low, and the post-expansion period, when the set of bonds available for purchase through the SMCCF was widened to include high-yield bonds. In particular, we use the specification

$$y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}, \quad (3)$$

where, as before, y_{ijt} represents our measures of transaction costs; Eligible_t takes the value of 1 if the bond in trade j has an investment-grade rating and time-to-maturity of five years or less on

[20] SMCCF transaction-level disclosures are available here. We provide additional details of this issue, including examples, in Appendix A.

[21] Recall from Table 1 that transaction costs for US firms behaved very similarly to all bonds in our sample.

March 23, 2020; SMCCF_t takes the value of 1 if the trade occurs between March 23 and April 9, 2020; and α_s controls for size fixed effects.[22]

Unlike specification (1), we do not include bond fixed effects in the baseline specification (3), but instead control for industry fixed effects (α_k) and bond-specific characteristics such as bond age, amount outstanding, and time-to-maturity ($X_{i,t}$). However, for robustness, we also include results allowing for bond fixed effects, as well as credit rating fixed effects. To ensure that treatment and control groups do not overlap, we remove all trades in bonds that were downgraded from IG to HY. Finally, we drop all foreign bonds and focus only on bonds issued by US firms.

Table 4 contains our results. As is standard in difference-in-differences regressions, β_1 is the primary coefficient of interest. The first key takeaway is that the SMCCF had a significant effect on the cost of risky-principal trades for eligible bonds relative to ineligible bonds. The quantitative magnitude of this effect is approximately 50 bps, and is robust to a variety of alternative specifications. For example, in column (2) we include a credit rating fixed effect, which absorbs some of the effects of eligibility related to the ratings restriction, leaving (roughly speaking) the effects of eligibility based on time-to-maturity. In columns (3) and (4), we allow for bond-specific fixed effects, which increases the explanatory power of the regressions (i.e., R^2) but does not significantly change the estimates of β_1.

The second noteworthy result is that, for risky-principal trades, β_2 is not statistically different from zero under any of our specifications. Hence, it appears that the *announcement* of the initial SMCCF did not have significant spillover effects on the cost of risky-principal trades for ineligible bonds. However, this does not rule out the potential for spillover effects from the actual *purchase* of eligible bonds, which began on May 12, 2020. In particular, by purchasing bonds and relaxing dealers' balance sheet constraints, the SMCCF could potentially increase dealers' willingness to

[22]One potential complication in distinguishing between eligible and ineligible bonds based on maturity is that the criteria for eligibility are determined at the Fed's time of purchase. Therefore, for example, a bond that would be characterized as ineligible when the SMCCF was announced on March 23, 2020, might, in fact, be purchased by the Fed in November 2020 (since the program remained active until December 31, 2020). To make sure that this complication does not affect our main results, in Tables A9–A12 in Appendix C.4, we recreate Tables 1–4, leaving out all trades involving bonds with maturity 5-6 years on March 23, 2020. Since these bonds represent a small fraction of our transactions, it turns out that our results are largely unaffected.

purchase any bond. If this is true, then some of the post-expansion decline in the costs of risky-principle trades for ineligible bonds (reported in Table 3) could be attributed to spillover effects from the Fed's bond purchases.

Columns (5)–(8) indicate that the announcement of the SMCCF on March 23 also reduced the cost of agency trades for eligible bonds.[23] One possible explanation is that, by establishing itself as a buyer of last resort, the Federal Reserve reduced the risk to private investors from purchasing eligible corporate bonds. According to this logic, it is possible that the announcement of the SMCCF made it easier for dealers to locate customer-buyers, hence reducing the spreads they charged on agency trades for eligible bonds. Note that this mechanism could also explain why the cost of agency trades for ineligible bonds went up in the immediate aftermath of the SMCCF announcement: if budget-constrained customers substituted from ineligible to eligible bonds, it would become more difficult for dealers to locate customer-buyers for ineligible bonds, driving spreads up.

In Appendix C, we provide several additional robustness checks for the results discussed above. In particular, in Tables A4 and A5, we show that the impact of the SMCCF on the trading cost of eligible bonds is even more pronounced if we limit our sample to those bonds that are just above and below the eligibility thresholds for and credit rating, respectively. In addition, in Tables A6–A8, we show that small and large trades are responsible for the entire liquidity improvement documented in Table 4: small trades (with par volume of $100,000 or less) become much more liquid after the SMCCF announcements, while large trades (with volume larger than $1 million) also exhibit a significant decline in trading costs. Odd-lot trades (with volume between $100,000 and $1 million), however, are essentially unaffected by the Fed's intervention.

[23]Note that, looking at the overall effect $(\beta_1 + \beta_2 + \beta_3)$, column (6) indicates that, after controlling for credit rating, the cost of agency trades for eligible bonds decreased after SMCCF announcement.

4 A structural analysis

The empirical analysis above highlights that the US corporate bond market experienced a significant decline in liquidity at the onset of the COVID-19 crisis, which was partially reversed by the Fed's interventions. Though informative, the facts we document leave several key questions unanswered. For one, what was the nature of the shocks that led to a lack of liquidity? Why did the policies that were implemented appear to restore liquidity relatively quickly, but only partially? And how did these shocks and the ensuing interventions affect the surplus of the customers in this market?

To confront these questions, we now construct a parsimonious equilibrium model of the market for immediacy and use it to conduct a structural analysis of our empirical observations. Our analysis reveals that, at the onset of the crisis, the market was hit by large shocks to both customer demand (the "dash for cash") and dealer supply (rising "balance sheet costs"). After the announcement of the Fed's key policy interventions, the demand shock subsided relatively quickly, and fully, while the supply shock recovered more gradually, and only partially. Relative to the pre-crisis period, we estimate that the loss in surplus per unit of transaction reached approximately 30 bps at the height of the crisis, and only partially recovered afterwards.

4.1 A theoretical framework

There are two types of agents: a measure N of customers and a measure one of dealers, all of whom are price takers. Each customer seeks to trade one share of an asset, and we do not distinguish between purchases and sales; this simplification allows us to study the determinants of transaction costs, though it is worth noting that our model is silent on the determinants of the asset's price. Since there are N customers with unit demand, the aggregate demand for transactions is exogenous and equal to N. However, while the total number of transactions is exogenous, the composition is not. Namely, we assume that customers demand vertically differentiated transaction services supplied by dealers at a convex cost: low-quality (l) transaction services, interpreted as agency trades, and high-quality (h) transaction services, interpreted as risky-principal trades.

Customers have quasi-linear utility for transaction services and for cash. Specifically, the problem of a customer is to choose how much low- and high-quality transaction services to demand from dealers at each time t in order to solve

$$\begin{aligned} &\max_{x_{lt},x_{ht}} \; u(x_{lt}, x_{ht}) + \theta_t x_{ht} - p_{lt} x_{lt} - p_{ht} x_{ht} \\ &\text{sub. to} \quad x_{lt} + x_{ht} = 1. \end{aligned} \tag{4}$$

We assume that $u(x_{lt}, x_{ht})$ is increasing, concave, twice continuously differentiable, and satisfies $u_h(x_{lt}, x_{ht}) - u_l(x_{lt}, x_{ht}) \geq 0$, where the h and l subscripts denote first partial derivatives with respect to x_{ht} and x_{lt}, respectively. This condition simply means that the customer values high-quality transaction services more than low-quality transaction services. As will become clear shortly, the term $\theta_t x_{ht}$ generates shocks to customers' relative demand for risky-principal trades.

Assuming interior solutions, the first-order optimality conditions are

$$p_{jt} = u_j(x_{lt}, x_{ht}) + \mathbf{1}_{\{j=h\}}\theta_t, \qquad j \in \{l, h\}. \tag{5}$$

On the other side of the market, dealers choose their supply of transaction services, X_{lt} and X_{ht}, in order to maximize profits,

$$p_{lt} X_{lt} + p_{ht} X_{ht} - C(X_{lt}, X_{ht}),$$

where $C(X_{lt}, X_{ht})$ is some continuous, convex, and twice continuously differentiable cost function. This leads to the first-order optimality conditions

$$p_{jt} = C_j(X_{lt}, X_{ht}), \qquad j \in \{l, h\}. \tag{6}$$

Finally, the market clearing conditions for transaction services are simply

$$X_{jt} = N_t x_{jt}, \qquad j \in \{l, h\}. \tag{7}$$

An equilibrium is thus described by a sequence $\{x^\star_{lt}, x^\star_{ht}, X^\star_{lt}, X^\star_{ht}, p^\star_{lt}, p^\star_{ht}\}$ such that, at each time t, given the level of aggregate transaction demand N_t, equations (4)–(7) are satisfied.

4.2 Comparative statics

Combining the assumption of fixed-size demand, (4), with the customers' first-order optimality conditions in (5), we can express customers' *demand for immediacy* as a single equation in two unknowns, $(p_{ht} - p_{lt})$ and x_{ht}:

$$p_{ht} - p_{lt} = u_h(1 - x_{ht}, x_{ht}) - u_l(1 - x_{ht}, x_{ht}) + \theta_t. \tag{8}$$

This equation defines the inverse demand for immediacy: the relationship between the price premium for risky-principal trades and customers' marginal utility for upgrading from slow, agency trades to fast, risky-principal trades. As anticipated above, θ_t is a demand shock that generates a parallel shift of this inverse demand curve.

Exploiting the market clearing conditions in (7), in conjunction with (4) and (6), similar steps reveal an equation that captures the dealers' *willingness to supply immediacy*:

$$p_{ht} - p_{lt} = C_h\left(N_t(1 - x_{ht}), N_t x_{ht}\right) - C_l\left(N_t(1 - x_{ht}), N_t x_{ht}\right).$$

Hence, the equilibrium characterization reduces to a price premium, $p_{ht} - p_{lt}$, and a fraction of risky-principal trades, x_{ht}, that lies at the intersection of these demand and supply schedules.

This simple representation offers a parsimonious, transparent framework to analyze the effects of various types of shocks. First, shocks to consumers' relative preference for immediate, risky-principal trades, as captured by θ_t, shift the demand for immediacy but not the supply. As is evident from Figure 4a, a positive innovation to θ_t induces an increase in the relative price of risky-principal trades, along with an increase in the equilibrium fraction of such trades.

Alternatively, under natural conditions,[24] a surge in customer-to-dealer trading volume, as captured by an increase in N_t, has no effect on each customer's demand for immediacy, but causes an upward shift in the supply curve, as providing risky-principal transaction services becomes more costly as the total volume of transaction services grows. Of course, any shock to the cost function $C(\cdot, \cdot)$—perhaps due to an increase in intermediaries' risk aversion or cost of funding—would engender a similar shift in the supply curve. As is evident from Figure 4b, an upward shift in dealers' supply curve causes an increase in the price premium paid for risky-principal trades, but a decrease in the equilibrium fraction of such trades.

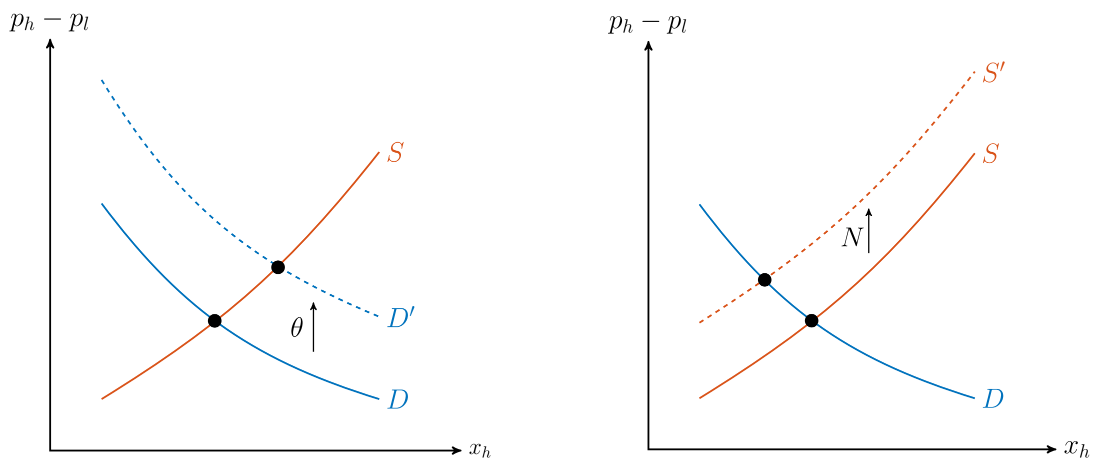

(a) Shock to customers' preferences **(b)** Shock to dealers' cost

Figure 4. Demand and supply shocks.

4.3 Estimating the model

In the data, we observed, simultaneously, an increase in the price premium $p_{ht} - p_{lt}$ (Figure 1) and a decrease in the fraction of risky-principal trade x_h (Figure 2). According to the model, this is indicative of a supply shock. As noted above, a supply shock could have been generated by an increase in the total volume of transaction services which, when combined with binding balance

[24]For the interested reader, we spell out these conditions in Appendix D.

sheet constraints, would make it more costly for dealers to supply risky-principal trades. Indeed, Figure 5 illustrates that customer-to-dealer volume was about 50 percent higher during the crisis, relative to the (average in the) same months in 2016-2019.

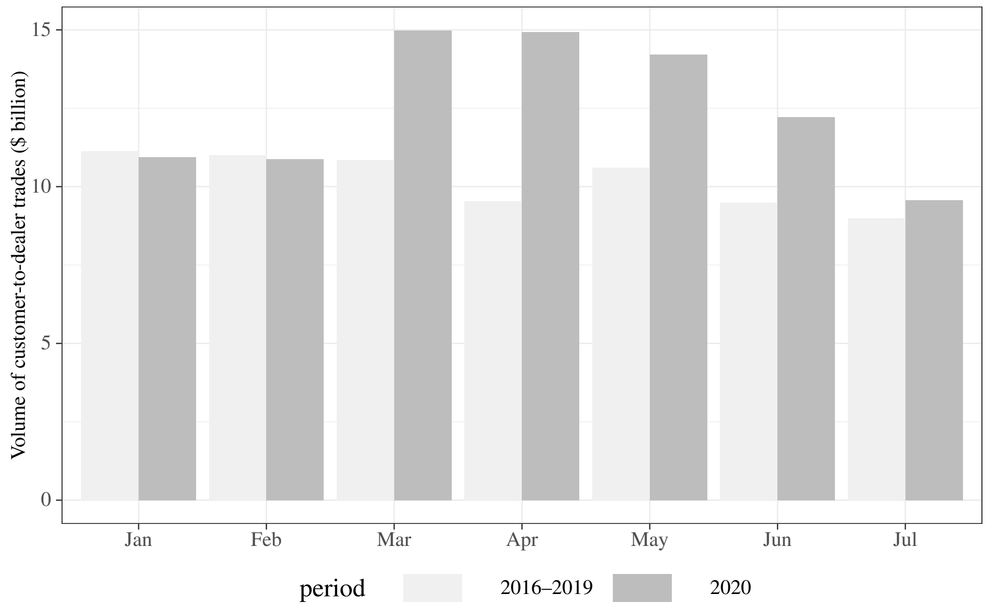

Figure 5. Customer-to-dealer volume. This figure plots the monthly volume of customer-to-dealer trades for 2020 and the 2016–2019 average in billions of USD.

However, the data does not rule out shocks to relative demand, as it is entirely possible that both the inverse supply and demand curves shifted in the same direction. To separately identify demand from supply shocks, we proceed in two steps. In this section, after imposing a specific functional form on customers' preferences, we estimate the parameter that determines the shape of the inverse demand curve by exploiting shifts to the supply curve that have occured during periods *outside of the crisis*, i.e., in "normal" times. Then, in Section 4.4, we use our estimated inverse demand curve to decompose movements in the price premium and the fraction of risky-principal trades *during the crisis* into movements that occur along the demand curve—caused by innovations to N_t or shocks to the cost function—and movements caused by innovations to θ_t, which shift the demand curve.

We plot this basic intuition in Figure 6. Note that this strategy is, by construction, largely independent of the shape of the supply curve and the nature of the shocks that shift it. However, by separately identifying shocks to the demand for immediacy from shocks to supply, we can study how the two responded differently to policy, and the quantitative implications for consumer surplus throughout the pandemic, which we do in Sections 4.4 and 4.5, respectively.

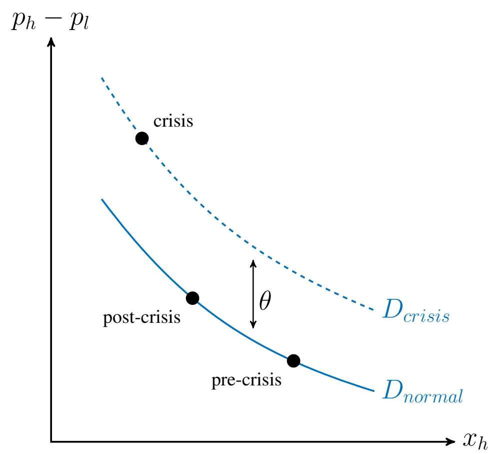

Figure 6. Identifying relative preference shocks, θ.

Parametric specification. Since x_h and x_l represent the market shares of high- and low-quality transactions in a vertically differentiated market, respectively, a natural choice for the demand curve is a logit specification.[25] In particular, we assume that, for each dollar of transaction service, the utility function of a consumer is given by

$$\theta_t x_{ht} - \sigma \left[x_{lt} \log(x_{lt}) + x_{ht} \log(x_{ht}) \right]. \tag{9}$$

It is well known (see, e.g., page 77 of Anderson et al., 1992) that this specification is equivalent to assuming that, for each dollar of transaction services, a consumer chooses between agency and risky-principal trades with net utilities $1 - p_{lt} + \varepsilon_{lt}$ and $1 + \theta_t - p_{ht} + \varepsilon_{ht}$, respectively, where ε_{lt} and ε_{ht} are independently and identically distributed (IID) over time and across consumers according to Gumbel distribution with location parameter zero and scale parameter σ.[26]

[25]Classic references for this specification include McFadden (1973), Anderson, De Palma, and Thisse (1992), and Berry (1994).

[26]In Appendix E, we provide a more detailed derivation of the logit demand function in a discrete choice framework.

Given this parametric assumption, the inverse demand for high- versus low-quality transaction services takes a log-linear form:

$$p_{ht} - p_{lt} = -\sigma \log\left(x_{ht}/x_{lt}\right) + \theta_t. \tag{10}$$

As one would expect, a larger price premium, $p_{ht} - p_{lt}$, results in a lower demand for risky-principal trades, x_{ht}. In addition, one sees that the shape of the demand curves depend on just one semi-elasticity parameter, σ.

As is well known, a simple OLS regression of the price premium $p_{ht} - p_{lt}$ on log quantities, $\log(x_{ht}/x_{lt})$, would yield a biased estimate of σ, since relative quantities are, in general, correlated with the relative demand shock θ_t. Hence, we use an instrumental variable (IV) approach to estimate the parameter of interest, σ.

To do so, consider an arbitrary instrument Z_t for the log relative quantities $\log(x_{ht}/x_{lt})$. Using (10), one easily sees that

$$\beta_{IV} = -\frac{\mathrm{Cov}\left(Z_t, p_{ht} - p_{lt}\right)}{\mathrm{Cov}\left(Z_t, \log(x_{ht}/x_{lt})\right)} = \sigma - \frac{\mathrm{Cov}\left(Z_t, \theta_t\right)}{\mathrm{Cov}\left(Z_t, \log(x_{ht}/x_{lt})\right)}.$$

Hence, as is well known, β_{IV} is an unbiased estimator of σ when the instrument Z_t is uncorrelated with the demand shock θ_t.

A binary IV approach. Consider observations about prices and relative quantities in two periods: a pre-crisis period, such as January 2020; and a post-crisis period, such as June 2020. Suppose the instrument Z_t takes the value zero in the pre-crisis period and 1 in the post-crisis period. Since

$$\mathrm{Cov}(Z_t, \theta_t) = \Pr(Z_t = 1)(1 - \Pr(Z_t = 1))\left(\mathbb{E}\left[\theta_t \mid Z_t = 1\right] - \mathbb{E}\left[\theta_t \mid Z_t = 0\right]\right),$$

it follows that an IV estimate based on Z_t is consistent if $\mathbb{E}\left[\theta_t \mid Z_t = 1\right] = \mathbb{E}\left[\theta_t \mid Z_t = 0\right]$. In other words, as long as the relative demand shock in the post-crisis period has returned to its pre-crisis average, then it is uncorrelated with the binary instrument, Z_t. Of course we also need the binary instrument to be relevant, i.e., correlated with the relative quantities. However, this is verified empirically because, as shown in Figure 2, relative quantities in the post-crisis period are lower than in the pre-crisis period. We interpret this observation as follows: as shown in Figure 5, trading volume in June remained elevated relative to pre-pandemic levels, which, in our model, shifts the marginal cost of providing transaction services, creating a supply shock. In equilibrium, the relative quantities demanded by consumers are reduced along a fixed demand curve.

The binary IV approach leads to the following candidate estimate

$$\hat{\sigma} = -\frac{\frac{1}{T_1}\sum_{t:Z_t=1}(p_{ht}-p_{lt}) - \frac{1}{T_0}\sum_{t:Z_t=0}(p_{ht}-p_{lt})}{\frac{1}{T_1}\sum_{t:Z_t=1}\log(x_{ht}/x_{lt}) - \frac{1}{T_0}\sum_{t:Z_t=0}\log(x_{ht}/x_{lt})},$$

where T_0 and T_1 are the lengths of pre- and post-crisis periods in days, respectively. For the estimation, we set the pre-crisis period to run between January 15, 2020 and February 14, 2020, and the post-crisis period to run between June 1, 2020 and June 30, 2020. We obtain an estimate of $\hat{\sigma} = 100.09$ with a standard deviation of 15.40.[27]

A high-frequency IV approach. An alternative approach to estimating σ is to use high frequency variation in trading conditions that affect prices and quantities but are unlikely to be attributable to aggregate shocks to customers' preferences for risky-principal trades. According to our theory, changes in N_t—which could equivalently represent the total volume of trade or the total number of trades (since trade size is fixed)—change the dealers' cost of supplying transaction services, but do not change individual consumers' relative demand for risky-principal trades.[28]

Hence, we consider two series for our instrument N_t: the total volume of customer-dealer trades and the total number of such trades in each day. Importantly, we exclude the crisis period March 1,

[27] We provide more details about the binary IV estimation method in Appendix F.

[28] Intuitively, one could imagine that dealers cannot fully adjust the size of their balance sheets or their trading infrastructure in the short term (our time period is a day). Thus, in the short run, an increase in the number of market participants or trading volume could put upward pressure on dealers' cost of providing immediacy, changing the relative prices or risky-principal and agency trades without shifting relative demand.

2020 to April 15, 2020, as shocks to relative demand during this period are likely significant and correlated with changes in both measures of N_t. We also seasonally-adjust and detrend (by adding month dummies) $\log(N_t)$ for both measures, so that shocks represent the residual deviation. After constructing these series, the formal exclusion restriction for the IV estimate to be consistent is $\text{Cov}(\log(N_t), \theta_t) = 0$. Relative to the binary IV, this approach has one advantage: it does not assume that the relative demand shock θ_t has returned to normal in June.

Table 5 presents the estimates of σ that emerge from our high frequency IV. The estimates range from 70 to 73 depending on the instrument, falling near the lower bound of the confidence interval of the binary-IV estimates.[29]

4.4 Demand Shocks, Supply Shocks, and Policy Implications

Given an estimate of the semi-elasticity parameter σ, along with the time series for the price premium ($p_{ht} - p_{lt}$) and the ratio of agency trades to risky-principal trades (x_{lt}/x_{ht}), we can infer the sequence of shocks to customers' relative demand for risky-principal trades (θ_t). Using our estimate of σ from the binary IV described above, $\hat{\sigma} = 100.09$, Figure 7 plots the time series relative to a pre-crisis benchmark, $\theta_t - \theta_0$, where θ_0 is the inferred value of θ on January 2, 2020.[30] We highlight three key takeaways.

First, θ_t experiences a dramatic increase during the most tumultuous weeks of March—our estimates suggest that customers' willingness to pay for each inframarginal unit of risky-principal trade (rather than an agency trade) increased by approximately 200 bps at the height of the crisis—but then recedes quickly after the announcements of the Fed's interventions. To give a sense of the effects of these shocks, note that we can decompose the change in the price premium if we assume

[29] As a robustness check, in Appendix F, we combine both our instruments and estimate the overidentified system using two-stage least square (2SLS). We also provide first-stage results for each IV regression verifying that each one is a valid instrument. For our 2SLS where we use both number and volume as instruments, we provide test statistics for the Sargan-Hansen test of overidentifying restrictions.

[30] As the estimate of $\hat{\sigma}$ from the high-frequency IV estimation lies within two standard deviations of the estimate from the binary IV, our results are largely independent of which estimate we use, as is clear from the error bands in the plots below.

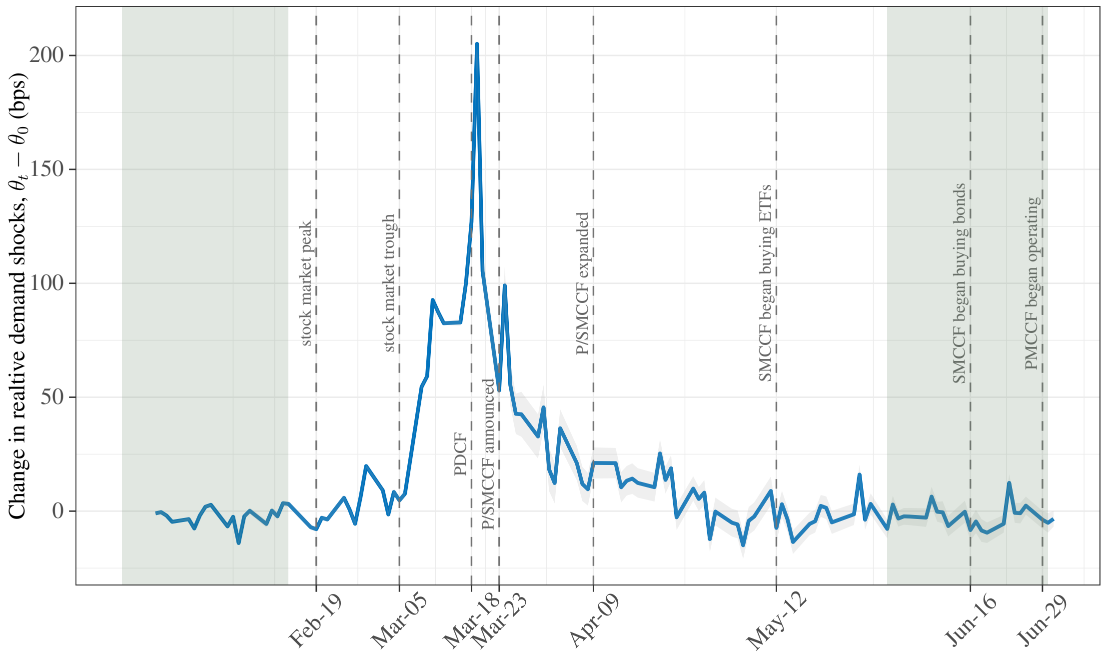

Figure 7. The change in the estimated relative demand shocks for risky-principal trades. This figure plots the time-series of the change in the estimated relative demand shocks for risky-principal trades relative to a pre-crisis benchmark on January 2, 2020, $\theta_t - \theta_0$, implied from Equation (10). The shaded areas correspond to the early and late periods used to estimate the semi-elasticity parameter σ.

that supply is perfectly inelastic:

$$p_{ht} - p_{lt} - (p_{h0} - p_{l0}) = \theta_t - \theta_0 + \sigma \left[\log \left(x_{lt}/x_{ht} \right) - \log \left(x_{l0}/x_{h0} \right) \right].$$

In general, this decomposition depends on the relative elasticity of demand and supply. However, we believe that the case of a perfectly inelastic supply is a natural benchmark, for two reasons. First, it delivers an upper bound for the contribution of demand shocks to the price premium. Second, the assumption of an inelastic supply accords with reports of binding balance-sheet constraint at the height of the crisis, and is seemingly confirmed by our evidence on dealers' inventory accumulation from Figure 3. According to this decomposition, a large portion of the spike in the cost of risky-principal trades (relative to agency trades) at the height of the crisis can be explained by relative demand shocks. For example, between March 5 and April 9, the average change in the price premium was 75 bps and the portion explained by the demand shifter was 58 bps, or approximately 75%. Hence, our results are consistent with other studies that highlight the "dash for cash" as an important driver of the turmoil in the corporate bond market (such as Falato et al., 2020; Ma et al., 2020; Haddad et al., 2020).

However, an increase in the demand for immediacy alone would generate an *increase* in x_h, as in Figure 4a, which is opposite of what we observe in the data. Therefore, the second key takeaway is that the onset of the pandemic must have also induced an even larger negative shock to dealers' willingness to use their balance sheet space to accommodate the surge in selling pressure. Indeed, within the context of Figure 6, our results suggest that the relative supply of risky-principal trades would have to experience a significant shift to the left in order to induce a drop in x_h. Thus, in addition to a surge in customers' demand for immediacy, our analysis reveals an equally important shock to dealers' costs of supplying immediacy; simply put, matching the data on prices *and quantities* requires an increase in the expected cost of dealers adding inventory to their balance sheets, as documented in the Treasury market by He et al. (2020).[31]

Finally, studying the behavior of the series $\{\theta_t, p_{ht} - p_{lt}, x_{ht}\}$ against the timeline of the Fed's interventions reveals important, new insights into the channels through which various policies impacted market liquidity. To start, the time path of relative preference shocks suggests that the *announcement* of the Fed's interventions were enough to halt and reverse the "dash for cash" that began in the second week of March. Indeed, by mid-April—a full month before the Fed's bond purchases actually began—we find that customers' preferences had returned to pre-crisis levels.[32] Hence, the expectation of price support from the Fed appears to have quickly and effectively ended the rush among customers to liquidate corporate debt immediately. Moreover, Haddad et al. (2020) and Ma et al. (2020) document that the initial dash for cash was most concentrated among more liquid investment-grade bonds. Therefore, the decline in θ_t that we find after the announcement of the Fed's interventions helps to explain why bid-ask spreads for risky-principal trades decline

[31]Studying the market for mortgage-backed securities (MBS), Chen et al. (2020) also find that the combined liquidity constraints of customers are dealers are responsible for the severe price dislocations observed during the COVID-19 pandemic.

[32]Notice that this result is not a trivial consequence of our underlying identification assumptions, but rather a confirmation of it: while our estimation assumes that relative preferences return to normal by early June, the data suggests that relative preferences, in fact, returned to normal much earlier.

more significantly for eligible (investment-grade) bonds, relative to ineligible (high-yield) bonds, during the period between March 23 and April 8, as documented in Table 4.

As θ_t receded and the demand curve shifted down, we also observe that x_h does not decrease, but rather increases slightly. Hence, it must be that the announced interventions also reduced dealers' perceived cost of supplying risky-principal trades, shifting their supply curve to the right and inducing them to absorb inventory onto their balance sheet, as documented in Figure 3. Lastly, the fact that the price premium and the fraction of agency trades remained elevated, even months after the initial shock appears to have passed, suggests that the supply curve did not return to its original location, i.e., that balance sheet costs remained higher than pre-crisis levels despite calmer markets and the Fed's interventions. These costs could derive from persistently high trading volume (as documented in Figure 5), from expectations of future price declines or volatility, or from losses incurred on other parts of the dealers' balance sheets.

4.5 Customer Surplus

In this section, we study how the shocks induced by the COVID-19 pandemic—and the interventions that followed—affected the consumer surplus of customers, or "customers' surplus."[33]

Theory. Given equilibrium prices and allocations, we can define the consumer surplus per dollar unit of transaction as

$$s_t = \theta x^\star_{ht} + u(x^\star_{lt}, x^\star_{ht}) - p^\star_{lt} x^\star_{lt} - p^\star_{ht} x^\star_{ht}.$$

[33]Note that we are intentionally not making any statements about optimal policy interventions or design. Doing so would, at the very least, require a richer model of the relationship between corporate bond market liquidity and real investment decisions, along with micro-foundations for customers' preferences (i.e., households), possible linkages between corporate bond market liquidity and other funding markets, and so on. While certainly interesting, these extensions are beyond the scope of the current paper.

Substituting (4) and (8), one can see that

$$
\begin{aligned}
s_t &= u(1-x^\star_{ht}, x^\star_{ht}) - x^\star_{ht}\left[u_h(1-x^\star_{ht}, x^\star_{ht}) - u_l(1-x^\star_{ht}, x^\star_{ht})\right] - p^\star_{lt} \\
&= \int_0^{x^\star_{ht}} \left[u_h(1-y,y) - u_l(1-y,y)\right] dy - x^\star_{ht}\left[u_h(1-x^\star_{ht}, x^\star_{ht}) - u_l(1-x^\star_{ht}, x^\star_{ht})\right] - p^\star_{lt} \\
&= -\int_0^{x^\star_{ht}} \left[u_{hh}(1-y,y) - 2u_{lh}(1-y,y) + u_{ll}(1-y,y)\right] y\, dy - p^\star_{lt}, \qquad (11)
\end{aligned}
$$

where the final equality follows from integration by parts. The term $-\left[u_{hh} - 2u_{lh} + u_{ll}\right]$ in the integral represents the slope of the inverse demand curve. Hence, the integral measures the area between the price premium and the inverse demand curve and so captures the surplus from upgrading from low-quality to high-quality transaction services. The second term in the consumer surplus is the cost of purchasing the baseline low-quality services.

Notice that the demand shock θ does not appear in the consumer surplus equation (11). The reason is that, for any fixed $x^\star_h$, the consumer surplus is the same regardless of the location of demand. Indeed, a parallel shift in the inverse demand curve increases the willingness to pay for all units by the same amount, θ. Therefore, the consumer surplus does not change, since it depends on the difference between the willingness pay for each infra-marginal unit and the willingness to pay for the marginal unit $x^\star_h$ (the price).

With the logit specification, $-\left(u_{hh} - 2u_{lh} + u_{ll}\right) = 1/[y(1-y)]$ and so we obtain a simple, closed-form expression for the change in consumer surplus between time zero and time t:

$$
s_t - s_0 = -\sigma \log\left(x_{lt}/x_{l0}\right) - \left(p_{lt} - p_{l0}\right). \qquad (12)
$$

Estimate of consumer surplus. Using Equation (12), Figure 8 plots the change in consumer surplus, per unit of transaction, over time. The figure reveals, not surprisingly, that there was a sharp, significant decline in customer surplus at the height of the market turmoil in mid-March, 2020. However, the figure also reveals that this decline was persistent: customer surplus per unit of transaction remained approximately 10 bps below pre-crisis levels even at the end June.

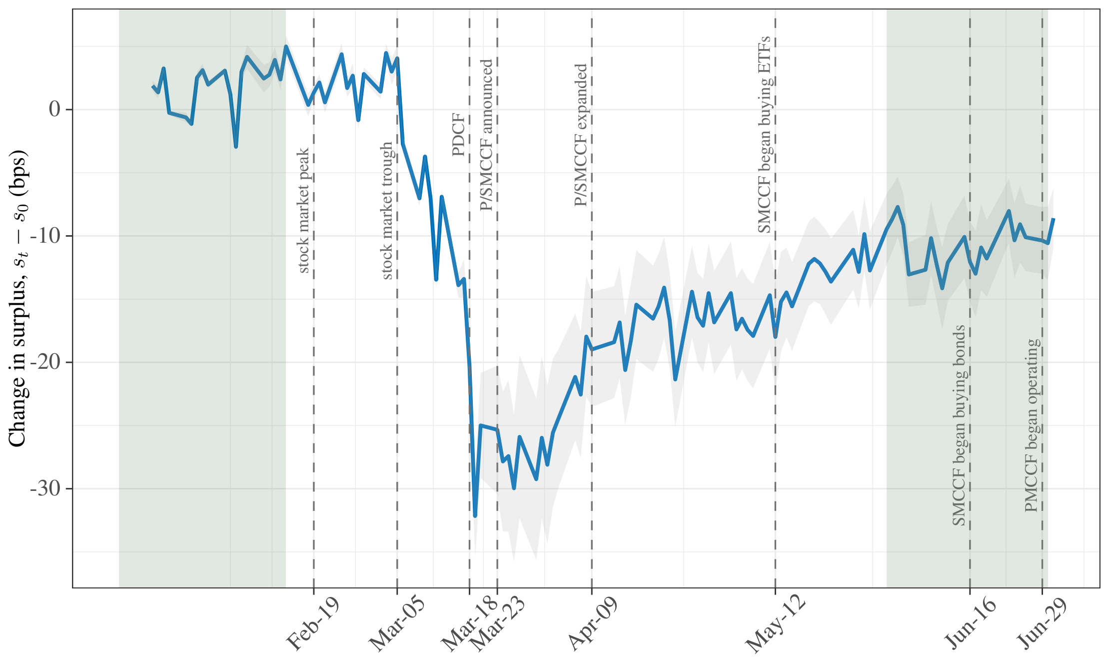

Figure 8. The change in consumer surplus in the logit demand specification. This figure shows the change in surplus relative to a pre-crisis benchmark on January 2, 2020, $s_t - s_0$, according to Equation (12). The shaded areas correspond to the early and late periods used to estimate the semi-elasticity parameter σ.

More generally, the figure conveys the important insight that customer surplus is not well-approximated by average transaction costs, which is perhaps the most commonly used metric for market liquidity and often used as a proxy for the well-being of customers. Intuitively, studying average transaction costs alone abstract from changes in preferences, as well as changes in the composition of risky-principal and agency trades. As a result, compared with average transaction costs, the fall in consumer surplus was more modest at the height of the crisis but more pronounced after the initial shock subsided.

Decomposition. Indeed, one can actually decompose the change in surplus into these three components—the change in average transaction costs, relative preferences for high- vs. low-quality transactions, and the composition. Formally, an application of the Envelope Theorem allows us to write the change in surplus as:

$$ds_t = d\theta_t x_{ht} - dp_{ht} x_{ht} - dp_{lt} x_{lt}.$$

Differentiating average transaction costs, $a_t \equiv p_{lt}x_{lt} + p_{ht}x_{ht}$, and substituting yields

$$\begin{aligned} ds_t &= -da_t + d\theta_t x_{ht} + (p_{ht} - p_{lt})\, dx_{ht} \\ &= -da_t + \left[dp_{ht} - dp_{lt} + \sigma \frac{dx_{ht}}{x_{ht}(1 - x_{ht})} \right] x_{ht} + (p_{ht} - p_{lt})\, dx_{ht}, \end{aligned} \tag{13}$$

where the second equality follows from the logit specification. The three terms represent the changes in surplus induced by the change in average transaction costs, the change in relative preference shocks, and the change in composition, respectively.

Figure 9 plots this decomposition, and illustrates that the initial fall in consumer surplus is more modest than the increase in transaction costs would suggest because customers value risky-principal trades significantly more during this period (even though the share of such trades falls). The figure also reveals that consumer surplus remains depressed in the latter part of our sample in part because the fraction of risky-principal trades remains below pre-crisis levels.

More generally, the decomposition reveals several important insights. For one, separately identifying shocks to customers' demand for immediacy and dealers' supply is crucial for understanding the relationship between consumer surplus and prices; for example, had the observed spike in the price premium early in the crisis been driven primarily by supply shocks, and not demand shocks, then consumer surplus *would* have tracked average transaction costs more closely. Moreover, to the extent that some interventions affected demand shocks more than supply shocks, identifying the sources of illiquidity is crucial for understanding which policies can help restore liquidity, and by how much.

5 Conclusion

It often takes a bad shock to discover whether or not a market is liquid, and to expose any sources of illiquidity. Unfortunately, many shocks to financial markets originate *inside* financial intermediaries, and hence the aggregate shock *is* a liquidity shock. In this sense, the COVID-19

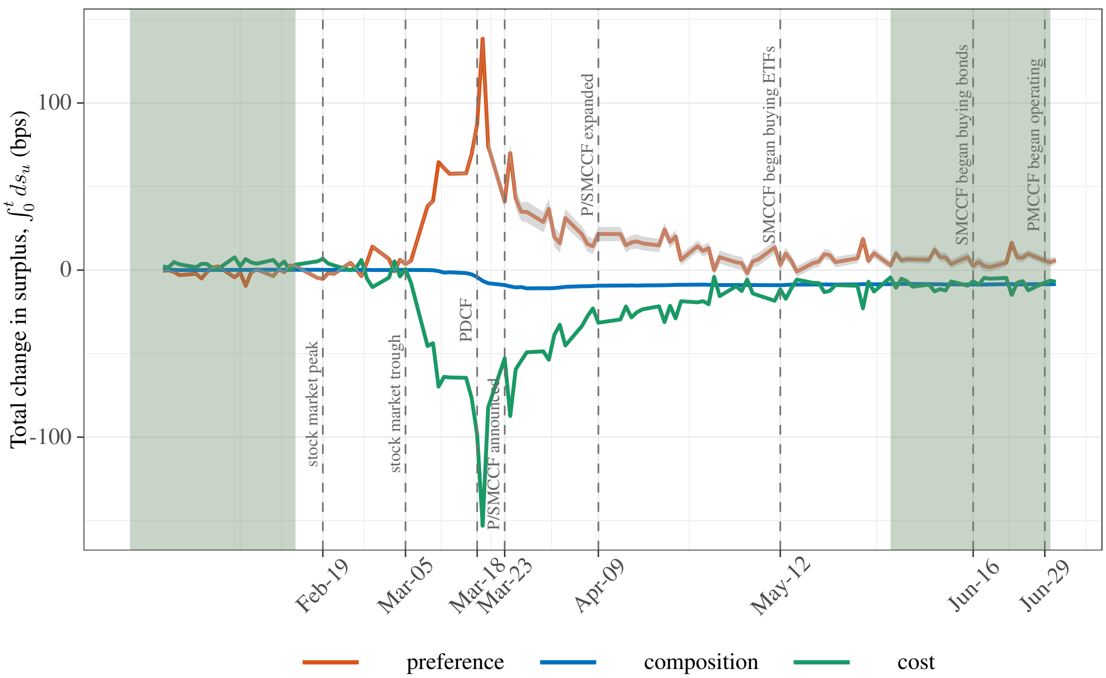

Figure 9. Decomposition of the total change in consumer surplus in the logit demand specification. This figure shows the decomposition of the total change in surplus, from integrating Equation (13), into three components induced by: (1) relative preference shocks (solid red line), (2) change in the average transaction costs (solid green line), and (3) substitution from risky-principal to agency trades (solid blue line). The shaded areas correspond to the early and late periods used to estimate the semi-elasticity parameter σ.

pandemic—a truly exogenous, large shock that did not originate in the banking sector—offers a unique opportunity to study market conditions, the shocks that precipitate episodes of illiquidity, and the implications for transaction costs, policy, and consumer surplus.

In this paper, we study trading conditions in the US corporate bond market from many angles as the COVID-19 pandemic unraveled. However, a key insight is that distinguishing between risky-principal and agency trades offers not only a more complete assessment of market conditions, but also a unique window into the sources of illiquidity, the efficacy of policy interventions, and the consequences for the surplus of customers in this large, important market. In particular, we find that the initial panic was caused by shocks to both customers' demand for immediacy *and* dealers' willingness to supply it. The former shock receded quickly, and almost fully, after the mere announcement of the Fed's intention to enter the market and purchase bonds. The latter shock, however, lingered months after markets appeared to calm, indicating that elevated trading volume in conjunction with balance sheet constraints remain a risk in times of crisis.

While this is an important first step, much work remains to be done. Perhaps most importantly, further examination of dealers' balance sheets and changes (or heterogeneity) in regulatory requirements could allow us to pinpoint the precise source of dealers' unwillingness to "lean against the wind" during times of crisis. Identifying and understanding these constraints would allow us to design better policies to balance the crucial trade-off between risk-taking and liquidity provision often at the heart of liquidity provision in financial markets. We leave this work for the future.

References

Anderson, Simon P., Andre De Palma, and Jacques-Francois Thisse, 1992, Discrete choice theory of product differentiation (MIT press).

Baer, Justin, 2020, The Day Coronavirus Nearly Broke the Financial Markets, The Wall Street Journal , May 20.

Bao, Jack, Maureen O'Hara, and Alex Zhou, 2018, The Volcker rule and market making in times of stress, Journal of Financial Economics 130, 95–113.

Bernanke, Ben, and Janet Yellen, 2020, The Federal Reserve must reduce long-term damage from coronavirus, Financial Times , March 18.

Berry, Steven T, 1994, Estimating discrete-choice models of product differentiation, The RAND Journal of Economics 242–262.

Bessembinder, Hendrik, Stacey Jacobsen, William Maxwell, and Kumar Venkataraman, 2018, Capital commitment and illiquidity in corporate bonds, The Journal of Finance 73, 1615–1661.

Boyarchenko, Nina, Anna Kovner, and Or Shachar, 2020, It's what you say and what you buy: A holistic evaluation of the corporate credit facilities, Working paper, FRB New York.

Chen, Jiakai, Haoyang Liu, Asani Sarkar, and Zhaogang Song, 2020, Cash-forward arbitrage and dealer capital in MBS makers: COVID-19 and beyond, Working paper, University of Hawaii, FRB, and Johns Hopkins University.

Choi, Jaewon, and Yesol Huh, 2018, Customer liquidity provision: Implications for corporate bond transaction costs, Working paper, UIUC and Federal Reserve Board.

D'Amico, Stefania, Vamsidhar Kurakula, and Stephen Lee, 2020, Impacts of the fed corporate credit facilities through the lenses of ETFs and CDX, Working paper, FRB Chicago.

Darmouni, Olivier, and Kerry Y. Siani, 2020, Crowding out bank loans: Liquidity-driven bond issuance, Working paper, Columbia Business School.

Dick-Nielsen, Jens, 2014, How to clean enhanced TRACE data, Working paper, CBS.

Duffie, Darrell, 2020, Still the world's safe haven? Redesigning the US treasury market after the COVID-19 crisis, Hutchins Center Working paper #62.

Ebsim, Mahdi, Miguel Faria-e Castro, and Julian Kozlowski, 2020, Corporate bond spreads and the pandemic, Federal Reserve Bank of Saint Louis, On the Economy Blog, Posted April 9, 2020.

Falato, Antonio, Itay Goldstein, and Ali Hortaçsu, 2020, Financial fragility in the COVID-19 crisis: The case of investment funds in corporate bond markets, Working paper, FRB, Wharton, and University of Chicago.

Feldhütter, Peter, 2012, The same bond at different prices: Identifying search frictions and selling pressures, The Review of Financial Studies 25, 1155–1206.

Flanagan, Thomas, and Amiyatosh Purnanandam, 2020, Corporate bond purchases after covid-19: Who did the fed buy and how did the markets respond?, Working paper, University of Michigan.

Fleming, Michael J., and Francisco Ruela, 2020, Treasury market liquidity during the COVID-19 crisis, Working paper, FRB New York.

Foley-Fisher, Nathan, Gary Gorton, and Stéphane Verani, 2020, The dynamics of adverse selection in privately-produced safe debt markets, Working paper, FRB and Yale School of Management.

Gilchrist, Simon, Bin Wei, Vivian Z. Yue, and Egon Zakrajšek, 2020, The Fed Takes on Corporate Credit Risk: An Analysis of the Efficacy of the SMCCF, Working paper, NYU, FRB, Emory, and BIS.

Goldberg, Jonathan, and Yoshio Nozawa, 2020, Liquidity supply in the corporate bond market, The Journal of Finance , Forthcoming.

Goldstein, Michael A., and Edith S. Hotchkiss, 2020, Providing liquidity in an illiquid market: Dealer behavior in us corporate bonds, Journal of Financial Economics 135, 16–40.

Haddad, Valentin, Alan Moreira, and Tyler Muir, 2020, When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed's response, The Review of Financial Studies , Forthcoming.

He, Zhiguo, Stefan Nagel, and Zhaogang Song, 2020, Treasury inconvenience yields during the covid-19 crisis, Working paper, University of Chicago and Johns Hopkins University.

Idzelis, Christine, 2020, The Corporate Bond Market Is 'Basically Broken,' Bank of America Says, Institutional Investor , March 19.

Ma, Yiming, Kairong Xiao, and Yao Zeng, 2020, Mutual fund liquidity transformation and reverse flight to liquidity, Working paper, Columbia and Wharton.

McFadden, Daniel, 1973, Conditional logit analysis of qualitative choice behavior, in P. Zarembka, ed., Frontiers in Econometrics, chapter 4, 105–142 (Academic Press: New York).

Nozawa, Yoshio, and Yancheng Qiu, 2020, Corporate bond market reactions to quantitative easing during the COVID-19 pandemic, Working paper, HKUST.

O'Hara, Maureen, and Xing (Alex) Zhou, 2020, Anatomy of a liquidity crisis: Corporate bonds in the Covid-19 crisis, Journal of Financial Economics , Forthcoming.

Schrimpf, Andreas, Hyun Song Shin, and Vladyslav Sushko, 2020, Leverage and margin spirals in fixed income markets during the covid-19 crisis, Working paper, BIS.

Schultz, Paul, 2017, Inventory management by corporate bond dealers, Working paper, University of Notre Dame.

Weill, Pierre-Olivier, 2007, Leaning against the wind, Review of Economic Studies 74, 1329–1354.

Table 1. Trading costs during the COVID-19 crisis. This table presents regression results for the following specification: $y_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}$. The dependent variables are our measures of transactions costs for risky-principal and agency trades. Crisis_t and Intervention_t are dummies which take the value of 1 if day t falls into the Crisis and Intervention sub-periods defined above. There are three trade size categories: less than $100,000, between $100,000 and $1 million, and larger than $1 million. The sample starts on January 3 and ends on June 30, 2020. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:			
	Risky-principal		Agency	
	All	US Only	All	US Only
	(1)	(2)	(3)	(4)
Crisis	105.19***	104.76***	8.70***	9.99***
	(13.08)	(13.78)	(1.71)	(2.06)
Intervention	41.54***	40.34***	8.04***	8.26***
	(4.06)	(4.32)	(0.76)	(1.06)
Bond FE	Yes	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes	Yes
Observations	741,579	603,913	249,392	160,539
Adjusted R^2	0.18	0.19	0.28	0.28

Note: *p<0.1; **p<0.05; ***p<0.01

Table 2. Probability of an agency trade for all bonds. This table presents regression results for the following specification from: $\text{Agency}_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}$. The dependent variable, Agency_{ijt}, is an indicator variable that takes the value 1 if trade j for bond i on day t is an agency trade and 0 otherwise. Columns (1), (2), and (3) report result for the linear probability (OLS), logit, and probit models, respectively. We report marginals effects calculated at the sample means for logit and probit models in columns (2) and (3). Crisis_t and Intervention_t are dummies which take the value of 1 if day t falls into Crisis and Intervention sub-periods defined above. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. In logit and probit specifications, the pseudo-R^2 is defined as $1 - L_1/L_0$, where L_0 is the log likelihood for the constant-only model and L_1 is the log likelihood for the full model with constant and predictors. The sample starts on January 3 and ends on June 30, 2020. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:		
	Probability of agency trade		
	OLS	Logit	Probit
	(1)	(2)	(3)
Crisis	0.043***	0.043***	0.042***
	(0.010)	(0.010)	(0.010)
Intervention	0.027***	0.027***	0.027***
	(0.003)	(0.003)	(0.003)
Bond FE	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes
Observations	7,095,617	7,095,617	7,095,617
Adjusted R^2	0.104		
Pseudo R^2		0.079	0.079
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 3. Trading costs across eligible and ineligible bonds during the initial and expanded interventions. This table presents regression results for the following specification: $y_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{SMCCF Expansion}_t + \varepsilon_{ijt}$. The dependent variables are measures of transaction costs for risky-principal and agency trades. Crisis_t is a dummy which takes the value of 1 if day t falls into the Crisis sub-periods defined above. SMCCF_t and SMCCF Expansion_t are dummies that take the value of 1 if the trading day t is between March 23 and April 9, and after April 9, 2020, respectively. The SMCCF eligibility criteria were expanded to include fallen angels on April 9, 2020. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. A bond is considered eligible if it has an investment-grade rating and time-to-maturity of five years or less on March 23, 2020. The sample begins on January 3 and ends on June 30, 2020. Only US firms are included in the regressions. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:					
	Risky-principal			Agency		
	All	Eligible	Ineligible	All	Eligible	Ineligible
	(1)	(2)	(3)	(4)	(5)	(6)
Crisis	107.97***	108.88***	106.67***	10.55***	15.94***	7.39***
	(14.77)	(15.14)	(16.60)	(2.30)	(3.35)	(1.91)
SMCCF	83.17***	61.94***	96.63***	13.43***	11.97***	14.43***
	(7.99)	(8.37)	(9.82)	(0.96)	(1.22)	(1.36)
SMCCF Expansion	26.55***	14.15***	33.05***	6.09***	4.25***	7.31***
	(2.69)	(2.60)	(3.63)	(0.88)	(0.92)	(1.21)
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	602,430	219,624	382,806	159,653	56,264	103,389
Adjusted R^2	0.19	0.18	0.19	0.29	0.21	0.29
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 4. The Effects of Fed Intervention: difference-in-differences. This table presents regression results for the following difference-in-differences specification from equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transaction costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of five years or less on March 23, 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. The sample begins on March 6 and ends on April 9, 2020. Only US firms are included and bonds that change credit grade are excluded. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	−57.70***	−41.72***	−47.24***	−41.45***	−10.25***	−12.85***	−9.59**	−9.85***
	(11.80)	(12.27)	(10.21)	(10.34)	(2.99)	(3.11)	(3.44)	(3.47)
SMCCF	−1.89	−21.75	−14.30	−20.03	6.33***	8.10***	4.56**	4.72**
	(14.58)	(14.64)	(14.65)	(14.43)	(2.00)	(2.11)	(1.97)	(2.02)
Eligible	2.86	−14.81			0.37	9.93***		
	(14.24)	(11.36)			(3.15)	(3.69)		
log(Amt outstanding)	−30.33***	−31.88***			−3.62***	−1.87***		
	(7.25)	(9.19)			(0.64)	(0.65)		
log(Time-to-maturity)	15.40***	16.77***			4.00***	5.53***		
	(4.96)	(4.99)			(0.85)	(1.26)		
log(Age)	27.61***	28.84***			4.93***	5.24***		
	(7.54)	(6.40)			(1.10)	(1.14)		
Trade size category FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	158,647	146,143	158,649	146,143	47,628	45,324	47,630	45,324
Adjusted R^2	0.04	0.05	0.20	0.20	0.08	0.10	0.25	0.26

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5. Estimating the logit demand parameter σ using high-frequency IV approach. This table presents the IV estimates of the logit demand parameter σ. In column (1), the seasonally-adjusted log volume of trades is used as an instrument. In column (3), the seasonally-adjusted log number of trades is used as an instrument. Standard error are the maximum of robust and the usual standard errors. The pre-crisis runs from January 3, 2020 until February 29, 2020. The post-crisis data begins on April 15, 2020 and runs until July 31, 2020. We exclude holidays, weekends and half trading days. Our estimates and standard errors are transformed using the delta-method where appropriate.

	Dependent variable:	
	$(p_h - p_l)$	
	IV (vol)	IV (num)
	(1)	(2)
$\log(x_h/x_l)$	70.17^{**}	73.31^{**}
	(33.63)	(31.02)
Post-crisis	8.25	7.75
	(5.75)	(5.14)
Constant	75.55^{***}	78.15^{***}
	(27.94)	(25.79)
Observations	113	113
Adjusted R^2	0.41	0.39
Note:	*p<0.1; **p<0.05; ***p<0.01	

Appendix

A Data and Definitions

A.1 Data description

We use data from the Trade Reporting Compliance Engine (TRACE), made available by the Financial Industry Regulation Authority (FINRA). The raw TRACE data provides detailed information on all secondary market transactions self-reported by FINRA member dealers. These include bond's CUSIP, trade execution time and date, transaction price ($100 = par), the volume traded (in dollars of par), a buy/sell indicator, and flags for dealer-to-customer and inter-dealer trades. To construct our sample, we combine two versions of TRACE: the standard version (2020Q1), and the End-Of-Day version (2020Q2).

We first filter the report data following the procedure laid out in Dick-Nielsen (2014). We merge the resulting data set with the TRACE master file, which contains bond grade information, and with the Mergent Fixed Income Securities Database (FISD) to obtain bond fundamental characteristics. Following the bulk of the academic literature, we exclude bonds with optional characteristics, such as variable coupon, convertible, exchangeable, and puttable, as well as, asset-backed securities, and private placed instruments. Table A1 provides summary statistics for our sample.

In our empirical specifications, we exclude newly-issued securities (with age less than 90 days), as on-the-run bonds tend to trade differently than off-the-run securities. Since our sample only contains about 130 days, the age and time-to-maturity of a particular bond will vary little over time. Thus, we do not include the standard cross-sectional controls related to the bond's age or time-to-maturity. Furthermore, since we exclude newly-issued bonds, over time, the age (maturity) of any bond will increase (decrease) by one day each day. Thus, the average age (maturity) of our bonds will increase (decrease) monotonically over time, meaning these controls will also correlate with the time trends we are documenting.

We also distinguish between bonds that are eligible for the SMCCF and ineligible bonds. In

Appendix B, we present a detailed description of eligibility criteria for the SMCCF. We define a bond as eligible if it has investment-grade rating and time-to-maturity of five years or less on March 23, 2020, when the SMCCF was first announced. The eligibility criteria also state that the firm must be a US-domiciled corporation. Specifically, the Fed restricts its purchases to bonds where

> The issuer is a business that is created or organized in the United States or under the laws of the United States with significant operations in and a majority of its employees based in the United States.

This criterion leaves the Fed with a considerable degree of discretion. For instance, if a foreign-domiciled corporation uses a US subsidiary to issue dollar-dominated debt, our firm-level data identify the firm as non-US. We would then classify its bonds as foreign, making them ineligible for the SMCCF. However, under the Fed's definition of a US issuer, the bonds may be eligible for purchase. Using the Fed's SMCCF transaction-level disclosures, we find that in many cases, the holding firm of the security is a non-US entity.[34] One such example is British American Tobacco (BAT), a firm listed on the London Stock Exchange and domiciled in the UK. Our firm-level data correctly identifies this firm as foreign; however, its bonds were purchased by the Fed.[35] These bonds were issued by a US wholly-owned subsidiary of BAT, BAT Capital Corporation. Since this subsidiary is guaranteed and wholly-owned by BAT, it is very challenging to correctly classify these bonds as US-domiciled. We, therefore, do not use US vs. non-US as an SMCCF eligibility criterion in our regressions discussed below and focus only on US firms.

Moreover, we do not have access to the latest credit rating data for all bonds in our sample. For the sub-sample of bonds where the credit rating is available, we include a credit rating fixed effect to control for potentially time-invariant nature of bond credit ratings.

A.2 Dates highlighted in the figures

We choose the following dates to highlight in the figures with vertical, dashed lines:

[34] SMCCF transaction-specific disclosures are provided by the Federal Reserve, available here.

[35] On July 10, 2020, the Fed reported that BAT's bonds were purchased as part of the SMCCF (CUSIP 05526DAZ8).

January 19: beginning of the series, chosen to start the sample period one month before the stock market peak.

February 19 stock market peak.

March 5: beginning of extended fall in equity prices and rise in corporate credit spreads.

March 18: first day of trading after announcement of Primary Dealer Credit Facility (announced evening of March 17).

March 23: announcement of Primary and Secondary Market Corporate Credit Facilities.

April 9: expansion of PMCCF and SMCCF (in both size and scope).

May 12: the SMCCF began purchasing eligible ETFs.

June 16: the SMCCF began purchasing individual corporate bonds.

June 29: the PMCCF began operating.

A.3 Identifying agency trades

We define agency trades as two trades in a given bond with the same trade size that take place within 15 minutes of each other. For each bond, we divide its trading sample into three groups: customer-sell-to-dealer (C2D), dealer-sell-to-customer (D2C), and interdealer (D2D) trades. Our identification of agency trades includes the following steps:

1. We match each trade X in group C2D with a trade Y in group D2C that has the same trade size and happens within 15 minutes of X. If there are several trades in D2C satisfying these conditions, we choose the trade that takes place closest in time to X. The identified pair of agency trades is then (X, Y). After this step, we denote the collection of unmatched trades in C2D as u-C2D and that in D2C as u-D2C.

2. We match each trade in u-C2D with a trade in group D2D by the same algorithm. We then obtain a collection of unmatched trades in D2D, denoted by u-D2D.

3. We match each trade in u-D2D with one in u-D2C following the same algorithm.

4. We repeat steps 1–3 using all remaining unmatched trades in the three groups while relaxing the matching criteria. In each agency trade pair, we require the second trade to happen within 15 minutes of the first trade, but it can have a smaller trade size than the first one. By doing so, we consider the situation in which dealers split the trade volumes when they behave as matchmakers.

5. Finally, within all the remaining unmatched trades after steps 1–4, we identify trades with field `remuneration == "C"` in TRACE (commission is included in the price) as agency trades, because, by FINRA's definition, broker-dealers receive commissions only when they intermediate agency trades.

B Corporate Credit Facilities

In March 23, 2020, the Federal Reserve Bank of New York established the Primary Market Corporate Credit Facility (PMCCF) and the secondary Market Corporate Credit Facility (SMCCF). The purpose of the PMCCF was to sustain funding for corporate debt while the SMCCF was meant to support liquidity in the corporate bond market. These corporate credit facilities were funded by a 75 billion dollars investment, to be leveraged up to 750 billion. The SMCCF started its purchases of ETFs on May 12 and of corporate bonds on June 16. The PMCCF started its operations on June 29. On December 31, 2020, the corporate credit facilities stopped their purchases.[36]

B.1 Bond eligibility criteria for the SMCCF

The Federal Reserve established eligibility criteria for the purchases of corporate bonds. We provide excerpts from the Fed's own communications that detail these conditions.[37]

> **Eligible individual corporate bonds:** The Facility may purchase individual corporate bonds that, at the time of purchase by the Facility: (i) were issued by an eligible issuer;

[36] For more details, see the Frequently Asked Questions for PMCCF and SMCCF from the New York Fed, available here.

[37] Source: Secondary Market Corporate Credit Facility Term Sheet, last updated on July 28, 2020.

(ii) have a remaining maturity of 5 years or less; and (iii) were sold to the Facility by an eligible seller.

Eligible issuers for individual corporate bonds: To qualify as an eligible issuer of an eligible individual corporate bond, the issuer must satisfy the following conditions:

1. The issuer is a business that is created or organized in the United States or under the laws of the United States with significant operations in and a majority of its employees based in the United States.
2. The issuer was rated at least BBB−/Baa3 as of March 22, 2020, by a major nationally recognized statistical rating organization ("NRSRO"). If rated by multiple major NRSROs, the issuer must be rated at least BBB−/Baa3 by two or more NRSROs as of March 22, 2020.
   (a) An issuer that was rated at least BBB−/Baa3 as of March 22, 2020, but was subsequently downgraded, must be rated at least BB−/Ba3 as of the date on which the Facility makes a purchase. If rated by multiple major NRSROs, such an issuer must be rated at least BB−/Ba3 by two or more NRSROs at the time the Facility makes a purchase.
   (b) In every case, issuer ratings are subject to review by the Federal Reserve.
3. The issuer is not an insured depository institution, depository institution holding company, or subsidiary of a depository institution holding company, as such terms are defined in the Dodd-Frank Act.
4. The issuer has not received specific support pursuant to the CARES Act or any subsequent federal legislation.
5. The issuer must satisfy the conflicts of interest requirements of section 4019 of the CARES Act.

C Additional Empirical Results

C.1 Transaction costs: impact of credit rating

We do not have access to the latest credit rating data for all bonds in our sample, just the binary IG/HY classification provided by TRACE. For the sub-sample of bonds where the credit rating is available, we include a credit rating fixed effect in specification (1) and run the following regressions

$$y_{ijt} = \alpha_i + \alpha_s + \alpha_r + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt},$$

where α_r represents credit rating fixed effects to control for potentially time-invariant nature of bond credit ratings. In Table A2, we repeat the results in Table 1 for the sub-sample of bonds for which we have credit rating data. We see that the results are very similar to the ones from Table 1.

C.2 The fraction of agency trades

In Table A3, we repeat the OLS regression in column (1) of Table 2 but focusing only on bonds issued by US firms. In columns (2) and (3) we repeat the regression in column (1) restricting the sample to eligible and ineligible bonds, respectively. As before, a bond is considered eligible if it has an IG credit rating and remaining time-to-maturity of five years or less.

Results in column (1), for US bonds, are very similar to what shown in column (1) of Table A3 for all bonds. From columns (2) and (3), we observe that the shift towards agency trades was much more pronounced among bonds that were eligible for the Fed's purchasing program. The probability of an agency trade for a given eligible bond, on average, rose by approximately seven percentage points relative to the pre-crisis period. After the Fed interventions on March 23, this probability decreased from the crisis period (by 200 bps) to five percentage points higher than the pre-crisis period. For ineligible bonds, in contrast, the probability of an agency trade rose by only 1.9 percentage points relative to the pre-crisis period and remained relatively unchanged after the Fed intervention.

C.3 Impact of Fed announcements

In this subsection, we present several robustness checks for the difference-in-differences (DID) results in Section 3.5.

C.3.1 Bonds close to the eligibility threshold for rating and maturity

First, in Table A4, we repeat the regressions in Table 4 but focusing only on bonds just above and below the SMCCF eligibility threshold for time-to-maturity (TTM): bonds with four to six years left to maturity.

Next, in Table A5, we repeat the regressions in Table A4 but adding the extra restriction that the bonds should be close to the IG-HY threshold. In particular, we only include bonds that in addition to having TTM of four, five and six years, are also rated at the bottom tier of investment-grade (BBB+/Baa1, BBB/Baa2, and BBB−/Baa3) or the top tier of high-yield (BB+/Ba1, BB/Ba2, and BB−/Ba3).

C.3.2 Trade costs for different trade size bins

Here we run the regressions in (4) but with the trades of a particular size category in a different regression. In Tables A6-A8, we show that small and large trades are responsible for the entire liquidity improvement documented in Table 4: small trades (with par volume of $100,000 or less) become much more liquid after the Fed's CCF announcements followed by large trades (with volume larger than $1 million). Liquidity of odd-lot trades (with volume between $100,000 and $1 million) seem to be unaffected by the Fed's intervention. Curiously we fail to find an affect for Odd-lot trades. There is some empirical evidence, e.g., Feldhütter (2012), suggesting that trades with different sizes are affected differently by market turmoil.

C.4 Excluding bonds with 5–6 years left to maturity

One potential complication in distinguishing between eligible and ineligible bonds based on maturity is that the criteria for eligibility are determined at the Fed's time of purchase. Therefore, for example, a bond that would be characterized as ineligible when the SMCCF was announced on March 23, 2020, might, in fact, be purchased by the Fed in November 2020 (since the program remained active until December 31, 2020). To make sure that this complication does not affect our main results, in Tables A9–A12 below, we recreate Tables 1–4 from the main text, leaving out all trades involving bonds with maturity 5-6 years on March 23, 2020. Since these bonds represent a small fraction of our transactions, it turns out that our results are largely unaffected.

D Natural Conditions for the Dealers' Cost Function

The following proposition summarizes the conditions under which a surge in selling pressure, as captured by an increase in N_t, has no effect on each customer's relative demand for immediacy, but causes an upward shift in the supply curve, as providing risky-principal transaction services becomes more costly as the total volume of transaction services grows.

Proposition 1 *Let* $(x_l^\star, x_h^\star, X_l^\star, X_h^\star, p_l^\star, p_h^\star)$ *be an equilibrium for a given* N*. If*

$$\frac{\partial}{\partial N}\left[C_h\left(Nx_l^\star, Nx_h^\star\right)\right] \geq \frac{\partial}{\partial N}\left[C_l\left(Nx_l^\star, Nx_h^\star\right)\right] \geq 0,$$

then, in response to a marginal increase in N*:*

- *The cost of all transaction services go up: both* $p_h^\star$ *and* $p_l^\star$ *increase;*
- *The cost of high-quality transaction services go up by more:* $p_h^\star - p_l^\star$ *increases;*
- *Customers substitute towards low-quality transaction services:* $x_h^\star$ *decreases.*

Proof. Assuming interior solutions, the first-order condition of the customers and dealers, along with the market clearing condition, yields the following system of equations

$$p_h - p_l = u_h(1-x_h, x_h) - u_l(1-x_h, x_h) \tag{D.1}$$

$$p_l = C_l(N(1-x_h), Nx_h) \tag{D.2}$$

$$p_h = C_h(N(1-x_h), Nx_h).$$

Combining the three equations lead an implicit function for x_h:

$$u_h(1-x_h, x_h) - u_l(1-x_h, x_h) = C_h(N(1-x_h), Nx_h) - C_l(N(1-x_h), Nx_h). \tag{D.3}$$

Since the functions $x \mapsto u(1-x, x)$ is concave, and the function $x \mapsto C(N(1-x), Nx)$ is convex, it follows that the left-hand side of the equation is decreasing in x, while the right-hand side is increasing in x. The condition stated in the Proposition implies that, locally, the right-hand side is increasing in N. Therefore, the solution x_h to this equation is, locally, decreasing in N. It then follows from Equation (D.1) that $p_h - p_l$ is increasing as well.

The only result that remained to be shown is that p_l is, locally, increasing in N. To do so, we totally differentiate equation (D.2) with respect to N:

$$\frac{dp_l}{dN} = C_{ll} \times \left(1 - x_h - N\frac{dx_h}{dN}\right) + C_{lh} \times \left(x_h + N\frac{dx_h}{dN}\right),$$

where we use double subscript for second derivatives, and we omit the arguments of C_{ll} and C_{lh} to simplify notations. Let

$$\varepsilon \equiv \frac{N}{x_h}\frac{dx_h}{dN},$$

denote the elasticity of high-quality transaction services, x_h, with respect to total transaction demand, N. Plugging back, we obtain:

$$\frac{dp_l}{dN} = C_{ll} \times (1 - x_h(1+\varepsilon)) + C_{lh} \times x_h (1+\varepsilon).$$

Next, applying the Implicit Function Theorem to Equation (D.3), we obtain the following explicit expression for the elasticity ε:

$$\varepsilon = \frac{N}{x_h} \frac{(1-x_h)(C_{lh} - C_{ll}) + x_h (C_{hh} - C_{lh})}{\partial^2 u - N (C_{hh} - 2C_{lh} + C_{ll})}.$$

where $\partial^2 u \equiv u_{ll} - 2u_{lh} + u_{hh} \leq 0$ by concavity. Plugging back into the equation for dp_l/dN, we obtain after some algebra that:

$$\frac{dp_l}{dN} \geq 0 \Leftrightarrow N \left(C_{ll}C_{hh} - C_{lh}^2\right) \geq \partial^2 u \left((1-x_h)C_{ll} + x_h C_{lh}\right).$$

The left-hand side, $N (C_{ll}C_{hh} - C_{lh}^2)$ is positive because C is convex. As for the right-hand side, recall our maintained assumption that, holding (x_l, x_h) fixed, $C_l(Nx_l, Nx_h)$ is increasing in N. Taking derivatives, and replacing x_l by $1 - x_h$, this means that $(1-x_h)C_{ll} + x_h C_{lh} \geq 0$. Keeping in mind that $\partial^2 u \leq 0$, we obtain that the right-hand side is negative, concluding the proof.

E Microfoundations for the Logit Demand

In this section, we derive the utility function in Equation (9) for the logit demand specification from a discrete choice model. We consider a model with a measure-one continuum of ex-ante identical customers who each wish to fulfill a transaction. Each customer faces two choices, either an agency (low-quality) or a risky-principal (high-quality) transaction. The price of these transactions are p_l and p_h, respectively.[38] A customer's utility of choosing each kind of transaction is either

$$1 + \theta - p_h + \varepsilon_h$$

[38] In this section, subscript t is omitted for simplicity.

if he chooses a risky-principal trade or

$$1 - p_l + \varepsilon_l$$

if he chooses an agency trade. The term $\theta \geq 0$ captures the idea that risky-principal trades are more valuable. The terms ε_h and ε_l capture idiosyncratic preference shocks and we assume that they are both independently and identically distributed across the population and drawn from a Gumbel distribution with location parameter zero and scale parameter σ with CDF $G(\varepsilon) = \exp(-\exp(-\varepsilon/\sigma))$. Each customer, after observing his own ε_h and ε_l, chooses the type of trade he prefers.

Thus, the fraction of consumers that will choose risky-principal trades is

$$x_h = \Pr\left(1 + \theta - p_h + \varepsilon_h \geq 1 - p_l + \varepsilon_l\right) = \Pr\left(\theta - p_h + p_l \geq \varepsilon_l - \varepsilon_h\right).$$

We can show that the CDF of $\varepsilon_l - \varepsilon_h$ is $\Pr(\varepsilon_l - \varepsilon_h < z) = 1/\left(1 + \exp(-z/\sigma)\right)$.

$$\begin{aligned}\Pr(\varepsilon_l - \varepsilon_h < z) &= \int_{-\infty}^{\infty} \exp\left(-e^{-\varepsilon_h/\sigma} e^{-z/\sigma}\right) \exp\left(-e^{-\varepsilon_h/\sigma}\right) e^{-\varepsilon_h/\sigma}\left(1/\sigma\right) d\varepsilon_h, \\ &= \frac{1}{1 + e^{-z/\sigma}} \int_{-\infty}^{\infty} \exp\left(-e^{-(\varepsilon_h/\sigma)\left(1 + e^{-z/\sigma}\right)}\right)\left(1 + e^{-z/\sigma}\right) e^{-\varepsilon_h/\sigma}(1/\sigma) d\varepsilon_h.\end{aligned}$$

But the last integral is equal to $\exp(-\exp(-\varepsilon_h/\sigma)(1 + \exp(-z/\sigma)))$ which evaluates to 1.

Taking $z = \theta - p_h + p_l$ and since $x_l = 1 - x_h$, we can write

$$x_h/x_l = \exp\left(\frac{\theta - p_h + p_l}{\sigma}\right).$$

The inverse demand specification in Equation (10) immediately follows:

$$p_{ht} - p_{lt} = \theta_t - \sigma \log\left(x_{ht} - x_{lt}\right).$$

The aggregate surplus over all customers is given by

$$s = x_h \left(1 + \theta - p_h + \mathbb{E}\left[\varepsilon_h \mid \varepsilon_h - \varepsilon_l \geq \theta - p_h + p_l\right]\right) + x_l \left(1 - p_l + \mathbb{E}\left[\varepsilon_l \mid \varepsilon_h - \varepsilon_l \leq \theta - p_h + p_l\right]\right).$$

The term $x_h \mathbb{E}\left[\varepsilon_h \mid \varepsilon_h - \varepsilon_l \geq \theta - p_h + p_l\right]$ can be written as

$$\int_{-\infty}^{\infty} \varepsilon_h G'(\varepsilon_h) G(\varepsilon_h + z) d\varepsilon_h.$$

Taking the same steps as before and recalling that $x_h = 1/\left(1 + \exp(-z/\sigma)\right)$, this can be rewritten as

$$x_h \int_{-\infty}^{\infty} \varepsilon_h e^{-e^{-\frac{\varepsilon_h}{\sigma}} x_h^{-1}} e^{-\frac{\varepsilon_h}{\sigma}} x_h^{-1} (1/\sigma) d\varepsilon_h.$$

Here, it is useful to temporarily introduce a new variable α defined by $e^{\alpha/\sigma} = x_h^{-1}$ and substitute x_h^{-1} above to obtain

$$x_h \int_{-\infty}^{\infty} \varepsilon_h e^{-e^{-\frac{\varepsilon_h - \alpha}{\sigma}}} e^{-\frac{\varepsilon_h - \alpha}{\sigma}} (1/\sigma) d\varepsilon_h.$$

Notice that this integral is now the expected value of a Gumbel random variable with location α and scale σ. That expected value is $\alpha + \sigma\gamma$ where γ is the Euler-Mascheroni constant. Thus,

$$x_h \mathbb{E}\left[\varepsilon_h \mid \varepsilon_h - \varepsilon_l \geq \theta - p_h + p_l\right] = x_h(\alpha + \sigma\gamma).$$

Using our definition of α, we can find that $\alpha = -\sigma \log(x_h)$ and therefore

$$x_h \mathbb{E}\left[\varepsilon_h \mid \varepsilon_h - \varepsilon_l \geq \theta - p_h + p_l\right] = x_h \sigma \left(\gamma - \log(x_h)\right).$$

A symmetry argument readily shows that

$$x_l \mathbb{E}[\varepsilon_l \mid \varepsilon_h - \varepsilon_l \leq \theta - p_h + p_l]) = x_l \sigma \left(\gamma - \log(x_l)\right).$$

Finally, substituting this result in the original aggregate utility expression, we obtain an expression for the logit surplus that is consistent with Equation (9) (up to an additive constant):

$$
\begin{aligned}
s &= x_h \left(1 + \theta - p_h - \sigma \log(x_h)\right) + x_l \left(1 - p_l - \sigma \log(x_l)\right) + \gamma \\
&= \theta x_h - \sigma \left[x_l \log(x_l) + x_h \log(x_h)\right] - p_h x_h - p_l x_l + (1 + \gamma) .
\end{aligned}
$$

F Estimation Details

F.1 Binary IV

In this section, we provide more details about the binary IV estimation method for σ discussed in Section 4.3. Let's consider a linear supply and demand system for transaction services. The demand equation is:

$$
\log(x_{ht}/x_{lt}) = 1/\sigma \left[\theta_\tau + \varepsilon_t - (p_{ht} - p_{lt})\right] ,
$$

where, with some abuse of notation, θ_τ is the average relative demand shifter in period $\tau \in \{A, B\}$, and ε_t is a mean zero shock. The supply equation is assumed to have the following form

$$
\log(x_{ht}/x_{lt}) = b_0 + b_1 \left(p_{ht} - p_{lt}\right) + b_2 \log(N_t) + \eta_t,
$$

where η_t is a mean zero supply shock, and N_t is the customer-to-dealer volume. Consistent with our model, N_t enters in the supply equation because it increases the marginal cost of providing transaction services. It does not enter the demand equation because the total utility for transaction services is linearly homogeneous: the utility of a given bundle of transaction service is the same for each dollar of transaction. We assume, moreover, that volume is smaller in the earlier period

(A) than in the later period (B).[39] Formally, this can be written $\log(N_t) = \log(\overline{N}_\tau) + \xi_t$, where ξ_t is a mean zero shock, $\tau \in \{A, B\}$ and $\overline{N}_A > \overline{N}_B$.

Solving the system of simultaneous equations gives

$$p_{ht} - p_{lt} = \frac{\theta_\tau + \varepsilon_t - \sigma\left(b_0 + b_2 \log(\overline{N}_\tau) + b_2\xi_t + \eta_t\right)}{1 + b_1\sigma},$$

and

$$\log(x_{ht}/x_{lt}) = \frac{b_0 + b_2 \log(\overline{N}_\tau) + b_2\xi_t + \eta_t}{1 + b_1\sigma} + \frac{b_1\left(\theta_\tau + \varepsilon_t\right)}{1 + b_1\sigma}$$

One sees that:

$$\mathbb{E}\left[p_{ht} - p_{lt} \mid t \in \tau\right] = \frac{\theta_\tau - \sigma\left(b_0 + b_2 \log(\overline{N}_\tau)\right)}{1 + b_1\sigma},$$

while

$$\mathbb{E}\left[\log(x_{ht}/x_{lt}) \mid t \in \tau\right] = \frac{b_1\theta_\tau + \left(b_0 + b_2 \log(\overline{N}_\theta)\right)}{1 + b_1\sigma}$$

It thus follows that, if $\theta_A = \theta_B$:

$$\sigma = -\frac{\mathbb{E}\left[p_{ht} - p_{lt} \mid t \in B\right] - \mathbb{E}\left[p_{ht} - p_{lt} \mid t \in A\right]}{\mathbb{E}\left[\log(x_{ht}/x_{lt}) \mid t \in B\right] - \mathbb{E}\left[\log(x_{ht}/x_{lt}) \mid t \in A\right]}.$$

Thus, an estimator of σ is obtained by replacing the conditional expectations by sample averages. To be more precise, let

$$Y_t = \begin{pmatrix} p_{ht} - p_{lt} & \log(x_{ht}/x_{lt}) \end{pmatrix}'$$

denote the vector of observations at time t. Assume for now that all the vector of disturbances, $(u_t, v_t, w_t)'$ are IID over time with finite covariance matrices. The estimate of the mean vector over

[39] Using an unpaired two-sample t-test, we test whether that the volume is smaller in the earlier period (A) than in the later one (B).We first check whether the variance of log volume is similar in the two periods using an F-test. The null hypothesis is that the variance of log volume is the same in both period. The value of the test statistic is 0.32, corresponding to a p-value of 0.012. Hence, we fail to reject the equality of variances for log volume (in level variances seem different, which is why we did it in logs). Next, we do a one-sided t-test. The null hypothesis is that volume is smaller in the period A than in period B. The value of the test statistic is -0.97 corresponding to a p-value of 0.83. Again we fail to reject the null that volume is smaller in period A than in period B.

periods A and B are:

$$\hat{Y}_\tau = \frac{1}{T_\tau}\sum_{t\in\tau} Y_t$$

where T_τ denotes the number of observations in period τ. Then, the weak Law of Large Numbers implies that sample means converge in probability to their population counterpart, $\overline{Y}_\tau$, as T_τ goes to infinity. The Central Limit Theorem implies that $\sqrt{T_\tau}\left(\hat{Y}_\tau - \overline{Y}_\tau\right)$ is asymptotically normally distributed with mean zero and covariance matrix S. Moreover, these two random variables are also independent. Since $\left(Y_t - \overline{Y}_\tau\right)$, $\tau \in \{A, B\}$ are IID, an unbiased and consistent estimator of S is

$$\begin{aligned}\hat{S} &= \frac{T_A - 1}{T_A + T_B - 2}\frac{1}{T_A - 1}\sum_{t\in A}\left(Y_t - \hat{Y}_A\right)\left(Y_t - \hat{Y}_A\right)' + \frac{T_B - 1}{T_A + T_B - 2}\frac{1}{T_B - 1}\sum_{t\in B}\left(Y_t - \hat{Y}_B\right)\left(Y_t - \hat{Y}_B\right)' \\ &= \frac{1}{T-2}\left(\sum_{t\in A}\left(Y_t - \hat{Y}_A\right)\left(Y_t - \hat{Y}_A\right)' + \sum_{t\in B}\left(Y_t - \hat{Y}_B\right)\left(Y_t - \hat{Y}_B\right)'\right)\end{aligned}$$

The estimate of σ can be written as

$$\hat{\sigma} = f\left(\hat{Y}_A, \hat{Y}_B\right) = -\frac{\hat{Y}_{A2} - \hat{Y}_{A1}}{\hat{Y}_{B2} - \hat{Y}_{B1}},$$

where the 1 and 2 subscripts denote the first and second coordinates of the Y vector. By standard delta method, we thus have that

$$\begin{aligned}\sqrt{T}\left(\hat{\sigma} - \sigma\right) &= \sqrt{T}\left(f\left(\hat{Y}_A, \hat{Y}_B\right) - f\left(\overline{Y}_A, \overline{Y}_B\right)\right) \\ &= \sqrt{T}\left(\frac{\partial f}{\partial Y_A'}\left(\hat{Y}_A - \overline{Y}_A\right) + \frac{\partial f}{\partial Y_B'}\left(\hat{Y}_B - \overline{Y}_B\right)\right) \\ &= \sqrt{\frac{T}{T_A}}\sqrt{T_A}\frac{\partial f}{\partial Y_A'}\left(\hat{Y}_A - \overline{Y}_A\right) + \sqrt{\frac{T}{T_B}}\sqrt{T_B}\frac{\partial f}{\partial Y_B'}\left(\hat{Y}_B - \overline{Y}_B\right).\end{aligned}$$

By symmetry it is clear that $\sqrt{T_A}\frac{\partial f}{\partial Y_A'}\left(\hat{Y}_A - \overline{Y}_A\right)$ and $\sqrt{T_B}\frac{\partial f}{\partial Y_B'}\left(\hat{Y}_B - \overline{Y}_B\right)$ have the same asymptotically normal distribution, with mean 0 and variance

$$\frac{\partial f}{\partial Y_A'} S \frac{\partial f}{\partial Y_A}, \text{ where } \frac{\partial f}{\partial Y_A'} \equiv \begin{pmatrix}\frac{1}{\overline{Y}_{B2} - \overline{Y}_{A2}} & -\frac{\overline{Y}_{B1} - \overline{Y}_{A1}}{\left(\overline{Y}_{B2} - \overline{Y}_{A2}\right)^2}\end{pmatrix}.$$

It thus follows that the estimate $\hat{\sigma}$ is asymptotically normal with mean σ and standard deviation

$$\left(\frac{1}{T_A} + \frac{1}{T_B}\right)^{1/2} \left(\frac{\partial f}{\partial Y_A'} S \frac{\partial f}{\partial Y_A}\right)^{1/2}$$

A consistent estimate of the standard deviation is found by replacing S by $\hat{S}$, and the population means in $\partial f / \partial Y_A'$ by their sample counterparts.

For the estimation we set the early period A to run between 2020-01-15 and 2020-02-14, and the late period B to run between 2020-06-01 and 2020-06-30. We use the raw series for x_h and x_l, not their moving average, so as not to artificially reduce standard errors. We obtain an estimate of $\hat{\sigma} = 100.09$ with a standard deviation of 15.4

F.2 High-frequency IVs

In this section we provide additional details and robustness results for the high-frequency IV method discussed in Section 4.3. Table A13 presents our estimates for σ from an overidentified IV. All specifications contain a post-crisis dummy variable allowing for a one time shift in the demand curve. This method gives us an estimate of σ that lies between the two individual IV estimates in columns (2) and (3) from Table 5. The weak instrument test statistic is 11.172. Since here we have an overidentified system, we obtain the Sargan's J test statistics. The value of the test statistic is 0.057 and the p-value of 0.81. So we fail to reject the validity of the overidentification restrictions.

In Table A14, we present the first stage of our IV regressions for each instrument. We find that both log number and volume of trades are valid instruments for $\log(x_h/x_l)$. The F-stat for the weak instruments test is 16.731 in the case of volume and 22.509 for the number of trades. In both cases we therefore reject the null that the instruments are weak at the 1% level. Both instruments have a correlation of 0.88.

For robustness, in Table A15, we repeat the IV regressions in Table 5 using the full sample from January to July, with indicators for the Crisis, SMCCF, and SMCCF expansion sub-periods defined in the main text. The IV estimates are now larger than the ones in Table 5, due to the potentially larger bias in the coefficients. As shown in Table A16, the first stage of the IV regression shows the both instrument remains valid in the full sample.

Table A1. Summary statistics. This table provides mean, standard deviation, median, 5th and 95th percentiles of the average daily number of trades and volume by counterparty type, proportion of agency trades, proportion of trades on IG bonds, proportion of trades on the bonds eligible for SMCCF, and daily average trading cost for risky-principal (CH) and agency trades (MIRC) for eligible and ineligible bonds respectively. "num" refers to number of trades, and the "vol" refers to volume of trades in par value. A bond is considered eligible for the SMCCF if it has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. Source: TRACE and FISD. The sample starts on January 3 and ends on June 30, 2020.

	Mean	Std.dev	Q05	Q50	Q95
daily num. all trades	51,908	8,871	39,642	52,237	67,005
daily num. interdealer	20,972	3,746	16,073	20,963	27,282
daily num. customer	30,936	5,278	23,622	31,137	38,949
daily num. customer-bought	16,114	2,859	12,187	16,239	20,630
daily num. customer-sold	14,822	2,952	11,202	14,497	19,752
daily vol. all trades ($b)	10.38	1.90	7.38	10.53	13.25
daily vol. interdealer ($b)	3.13	0.61	2.19	3.17	4.04
daily vol. customer ($b)	7.25	1.39	5.18	7.19	9.48
daily vol. customer-bought ($b)	3.75	0.66	2.83	3.75	4.79
daily vol. customer-sold ($b)	3.50	0.80	2.35	3.45	4.94
prop. agency (num)	0.54	0.03	0.50	0.53	0.59
prop. agency (vol)	0.33	0.04	0.29	0.33	0.38
prop. IG (num)	0.74	0.03	0.70	0.74	0.78
prop. IG (vol)	0.82	0.03	0.79	0.82	0.86
prop. eligible (num)	0.33	0.04	0.30	0.33	0.40
prop. eligible (vol)	0.27	0.04	0.23	0.27	0.34
daily avg. CH (bps)	51.30	35.21	22.38	40.66	123.29
daily avg. CH of Eligible bonds (bps)	28.11	26.34	10.71	20.85	65.86
daily avg. CH of ineligible bonds (bps)	61.32	41.78	27.50	47.67	147.32
daily avg. MIRC (bps)	10.57	3.38	7.31	9.62	17.43
daily avg. MIRC of Eligible bonds (bps)	5.35	2.78	2.91	4.43	11.45
daily avg. MIRC of ineligible bonds (bps)	12.33	4.12	8.19	11.18	20.51

Table A2. Robustness: Trading costs during the COVID-19 crisis adding credit rating FEs. This table presents regression results for the following specification: $y_{ijt} = \alpha_i + \alpha_s + \alpha_r + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}$. The dependent variables are our measures of transactions costs for risky-principal and agency trades. Crisis_t and Intervention_t are dummies which take the value of 1 if day t falls into the Crisis and Intervention sub-periods defined above. α_r represents credit rating fixed effects. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. The sample starts on January 3 and ends on June 30, 2020. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:			
	Risky-principal		Agency	
	All	US Only	All	US Only
	(1)	(2)	(3)	(4)
Crisis	105.78***	103.96***	8.94***	9.40***
	(13.03)	(13.57)	(1.83)	(1.98)
Intervention	37.68***	37.42***	6.10***	5.73***
	(4.35)	(4.46)	(0.72)	(0.81)
Bond FE	Yes	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes	Yes
Credit Rating FE	Yes	Yes	Yes	Yes
Observations	677,728	598,887	197,450	158,730
Adjusted R^2	0.19	0.19	0.29	0.30
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A3. Robustness: Probability of an agency trade for US bonds (OLS only). This table presents regression results for the following specification from: $\text{Agency}_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}$. The dependent variable, Agency_{ijt}, is an indicator variable that takes the value 1 if trade j for bond i on day t is an agency trade and 0 otherwise. Only US firms are included in the regression. Crisis_t and Intervention_t are dummies which take the value of 1 if day t falls into Crisis and Intervention sub-periods defined above. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. A bond is considered eligible if it has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. The sample starts on January 3 and ends on June 30, 2020. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:		
	Probability of agency trade		
	All	Eligible	Ineligible
	(1)	(2)	(3)
Crisis	0.043***	0.068***	0.025***
	(0.010)	(0.014)	(0.008)
Intervention	0.027***	0.044***	0.016***
	(0.003)	(0.004)	(0.004)
Bond FE	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes
Observations	5,770,765	2,337,519	3,433,246
Adjusted R^2	0.104	0.071	0.123
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table A4. DID robustness: only include bonds with 4 to 6 years left to maturity. This table presents regression results for the following DID specification from Equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transactions costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. The sample begins on March 6 and ends on April 9, 2020. Only US firms, bonds with 4, 5, or 6 years left to maturity on the intervention date are included. Bonds that change credit grade are excluded. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	−93.26**	−76.08***	−61.24***	−61.42***	−7.45*	−9.13**	−3.38	−3.79
	(39.33)	(27.87)	(17.83)	(17.89)	(4.36)	(4.22)	(4.14)	(4.09)
SMCCF	13.88	−0.46	−9.05	−9.14	2.82	5.15**	1.46	1.75
	(35.04)	(25.14)	(16.58)	(16.64)	(2.37)	(2.47)	(2.14)	(2.14)
Eligible	54.22	12.27			−0.77	11.87**		
	(50.71)	(33.55)			(4.42)	(5.46)		
log(Amount outstanding)	−3.86	−23.05**			−4.06***	−1.64		
	(16.73)	(10.71)			(1.03)	(1.07)		
log(Time-to-maturity)	−66.38	−102.01			8.79	30.42*		
	(134.08)	(90.03)			(17.08)	(17.69)		
log(Age)	28.46*	31.43**			0.99	2.55*		
	(15.69)	(12.45)			(1.95)	(1.45)		
Trade size category FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	30,743	30,430	30,744	30,430	9,182	9,004	9,183	9,004
Adjusted R^2	0.03	0.07	0.20	0.20	0.11	0.14	0.29	0.30
Note:								*p<0.1; **p<0.05; ***p<0.01

Table A5. DID robustness: only include bonds with 4 to 6 years left to maturity and rating close to the IG/HY threshold. This table presents regression results for the following DID specification from Equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transactions costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. The sample begins on March 6 and ends on April 9, 2020. Only US firms, bonds with 4, 5, or 6 years left to maturity that are rated at the bottom tier of IG (BBB+/Baa1, BBB/Baa2, and BBB−/Baa3) or the top tier of HY (BB+/Ba1, BB/Ba2, and BB−/Ba3) are included. Bonds that change credit grade are excluded. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	−94.92**	−86.46**	−73.52**	−73.52**	−2.82	−5.38	−0.16	−0.16
	(45.73)	(43.93)	(30.56)	(30.56)	(3.58)	(3.80)	(3.90)	(3.90)
SMCCF	46.47	37.41	16.30	16.30	2.54	5.14**	2.41	2.41
	(29.20)	(24.24)	(17.65)	(17.65)	(1.89)	(2.04)	(2.34)	(2.34)
Eligible	63.68	46.64			−1.18	7.74		
	(46.38)	(52.37)			(4.19)	(7.49)		
log(Amount outstanding)	−9.03	−18.19			−3.94***	−2.58*		
	(21.30)	(16.18)			(1.47)	(1.45)		
log(Time-to-maturity)	−160.39	−129.63			40.68	45.95		
	(150.14)	(130.02)			(30.85)	(30.34)		
log(Age)	36.44	30.68*			−3.29	−1.56		
	(22.44)	(18.18)			(3.10)	(2.31)		
Trade size category FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	14,124	14,124	14,124	14,124	4,595	4,595	4,595	4,595
Adjusted R^2	0.04	0.05	0.16	0.16	0.12	0.13	0.28	0.28

Note: *p<0.1; **p<0.05; ***p<0.01

Table A6. DID robustness: only include trades with par volume < $100,000, i.e., micro trades. This table presents regression results for US firms for the following DID specification from Equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transactions costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. The sample begins on March 6 and ends on April 9, 2020. Only trades that are less than $100,000 in par volume, i.e., micro trades, are included. Bonds that change credit grade are excluded. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	−77.67***	−54.42***	−58.51***	−45.93***	−16.21***	−19.37***	−14.86***	−15.66***
	(19.69)	(18.51)	(12.66)	(12.56)	(4.45)	(4.53)	(5.24)	(5.27)
SMCCF	3.05	−28.30	−18.24	−30.64	9.45***	11.48***	6.63***	7.26***
	(24.08)	(23.12)	(20.75)	(21.29)	(3.06)	(2.97)	(2.28)	(2.36)
Eligible	0.94	−22.17			6.48	15.67***		
	(23.80)	(18.37)			(4.66)	(4.64)		
log(Amt outstanding)	−38.79***	−32.77**			−4.24***	−2.03**		
	(11.74)	(13.90)			(0.79)	(0.79)		
log(Time-to-maturity)	11.54*	7.93			4.26***	5.98***		
	(6.90)	(7.99)			(1.08)	(1.58)		
log(Age)	39.02***	37.74***			7.44***	8.23***		
	(13.24)	(11.15)			(1.85)	(1.84)		
Industry FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	92,300	82,694	92,301	82,694	28,556	27,182	28,556	27,182
Adjusted R^2	0.05	0.08	0.35	0.37	0.05	0.08	0.26	0.27

Note: *p<0.1; **p<0.05; ***p<0.01

Table A7. DID robustness: only include trades with \$100,000 $\leq$ par volume $<$ \$1 million, i.e., odd-lot trades. This table presents regression results for US firms for the following DID specification from Equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transactions costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. The sample begins on March 6 and ends on April 9, 2020. Only trades that are greater than \$100,000 and less than \$1 million in par volume, i.e., odd-lot trades, are included. Bonds that change credit grade are excluded. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	−11.67	−8.28	−6.37	−5.44	−1.21	−1.66	−1.84	−1.01
	(12.57)	(13.87)	(13.85)	(14.15)	(2.50)	(2.35)	(2.56)	(2.34)
SMCCF	−27.03*	−31.91*	−36.84**	−37.76**	2.42	2.82*	2.58	1.74
	(15.14)	(16.45)	(15.07)	(15.25)	(1.78)	(1.51)	(2.00)	(1.82)
Eligible	−0.33	−4.16			0.23	2.10		
	(12.17)	(13.48)			(1.99)	(1.81)		
log(Amt outstanding)	−18.54***	−25.29***			−2.97***	−2.40**		
	(4.34)	(3.84)			(0.99)	(1.02)		
log(Time-to-maturity)	26.48***	33.45***			3.86***	3.00***		
	(2.85)	(3.26)			(0.60)	(0.54)		
log(Age)	15.49***	17.94***			3.02***	2.43***		
	(3.10)	(3.08)			(0.58)	(0.76)		
Industry FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	36,406	34,457	36,407	34,457	10,089	9,775	10,089	9,775
Adjusted R^2	0.03	0.03	0.07	0.07	0.04	0.03	0.20	0.22

Note: *p<0.1; **p<0.05; ***p<0.01

Table A8. DID robustness: only include trades with par volume $\geq$ \$1 million, i.e., large trades. This table presents regression results for US firms for the following DID specification from Equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transactions costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of 5 years or less on the March 23 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. The sample begins on March 6 and ends on April 9, 2020. Only trades that are greater than or equal to \$1 million in par volume, i.e. large trades, are included. Bonds that change credit grade are excluded. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	−24.22***	−22.17**	−29.02***	−30.56***	−6.45***	−7.96***	−1.81	−2.08
	(9.08)	(8.78)	(8.94)	(8.90)	(2.49)	(2.85)	(3.19)	(3.19)
SMCCF	6.70	4.25	5.12	6.66	4.43	5.76*	−0.83	−0.63
	(9.27)	(10.06)	(9.91)	(9.96)	(2.82)	(3.03)	(3.16)	(3.18)
Eligible	−0.16	−9.19			−19.98***	−7.83***		
	(7.51)	(10.00)			(1.57)	(2.75)		
log(Amt outstanding)	−19.14***	−26.76***			−1.10	0.31		
	(2.67)	(3.17)			(0.98)	(0.94)		
log(Time-to-maturity)	19.55***	22.15***			2.49***	4.99***		
	(2.33)	(2.99)			(0.76)	(0.90)		
log(Age)	14.35***	15.53***			0.42	0.81		
	(2.43)	(2.99)			(1.28)	(1.28)		
Industry FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	29,941	28,992	29,941	28,992	8,983	8,367	8,985	8,367
Adjusted R^2	0.02	0.02	0.02	0.02	0.13	0.18	0.28	0.28

Note: *p<0.1; **p<0.05; ***p<0.01

Table A9. Trading costs during the COVID-19 crisis: excluding bonds with 5–6 years left to maturity. This table presents regression results for the following specification: $y_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}$. The dependent variables are our measures of transactions costs for risky-principal and agency trades. Crisis_t and Intervention_t are dummies which take the value of 1 if day t falls into the Crisis and Intervention sub-periods defined above. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. All trades for bonds with 5–6 years left to maturity are excluded from the regressions. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:			
	Risky-principal		Agency	
	All	US Only	All	US Only
	(1)	(2)	(3)	(4)
Crisis	106.28***	106.23***	9.88***	11.46***
	(13.25)	(13.70)	(1.62)	(1.96)
Intervention	53.36***	55.36***	9.56***	10.59***
	(5.73)	(6.17)	(0.80)	(1.18)
Bond FE	Yes	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes	Yes
Observations	699,232	522,452	225,238	133,248
Adjusted R^2	0.18	0.19	0.26	0.27
Note:				*p<0.1; **p<0.05; ***p<0.01

Table A10. Probability of an agency trade for all bonds: excluding bonds with 5–6 years left to maturity. This table presents regression results for the following specification from: $\text{Agency}_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{Intervention}_t + \varepsilon_{ijt}$. The dependent variable, Agency_{ijt}, is an indicator variable that takes the value 1 if trade j for bond i on day t is an agency trade and 0 otherwise. Columns (1), (2), and (3) report result for the linear probability (OLS), logit, and probit models, respectively. We report marginals effects calculated at the sample means for logit and probit models in columns (2) and (3). Crisis_t and Intervention_t are dummies which take the value of 1 if day t falls into Crisis and Intervention sub-periods defined above. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. In logit and probit specifications, the pseudo-R^2 is defined as $1 - L_1/L_0$, where L_0 is the log likelihood for the constant-only model and L_1 is the log likelihood for the full model with constant and predictors. The sample starts on January 3 and ends on June 5, 2020. All trades for bonds with 5–6 years left to maturity are excluded from the regressions. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:		
	Probability of agency trade		
	OLS	Logit	Probit
	(1)	(2)	(3)
Crisis	0.036***	0.035***	0.035***
	(0.009)	(0.009)	(0.009)
Intervention	0.032***	0.032***	0.031***
	(0.003)	(0.003)	(0.003)
Bond FE	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes
Observations	6,486,640	6,486,640	6,486,640
Adjusted R^2	0.115		
Pseudo R^2		0.087	0.087
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A11. Trading costs across eligible and ineligible bonds during the initial and expanded interventions: excluding bonds with 5–6 years left to maturity. This table presents regression results for the following specification: $y_{ijt} = \alpha_i + \alpha_s + \beta_1 \times \text{Crisis}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{SMCCF Expansion}_t + \varepsilon_{ijt}$. The dependent variables are measures of transaction costs for risky-principal and agency trades. Crisis_t is a dummy which takes the value of 1 if day t falls into the Crisis sub-periods defined above. SMCCF_t and SMCCF Expansion_t are dummies that take the value of 1 if the trading day t is between March 23 and April 9, and after April 9, 2020, respectively. The SMCCF eligibility criteria were expanded to include fallen angels on April 9, 2020. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. A bond is considered eligible if it has an investment-grade rating and time-to-maturity of five years or less on March 23, 2020. The sample begins on January 3 and ends on June 5, 2020, when the SMCCF expanded eligibility criterion to fallen angels. Only US firms are included in the regressions. All trades for bonds with 5–6 years left to maturity are excluded from the regressions. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:					
	Risky-principal			Agency		
Crisis	110.82***	112.46***	109.09***	12.22***	16.52***	9.64***
	(14.81)	(15.61)	(15.78)	(2.24)	(3.53)	(1.85)
SMCCF	95.35***	64.75***	116.42***	14.86***	12.26***	16.84***
	(8.53)	(8.58)	(9.62)	(1.08)	(1.26)	(1.52)
SMCCF Expansion	34.37***	16.95***	45.95***	7.76***	5.05***	10.06***
	(3.26)	(2.85)	(4.84)	(1.07)	(0.97)	(1.65)
Bond FE	Yes	Yes	Yes	Yes	Yes	Yes
Trade size category FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	521,933	200,761	321,172	133,070	50,192	82,878
Adjusted R^2	0.19	0.19	0.19	0.27	0.21	0.27

Note: *p<0.1; **p<0.05; ***p<0.01

Table A12. The Effects of Fed Intervention: difference-in-differences: excluding bonds with 5–6 years left to maturity. This table presents regression results for the following difference-in-differences specification from equation (3): $y_{ijt} = \alpha_s + \alpha_k + \beta_1 \times \text{SMCCF}_t \times \text{Eligible}_t + \beta_2 \times \text{SMCCF}_t + \beta_3 \times \text{Eligible}_t + \gamma \times X_{i,t} + \varepsilon_{ijt}$. The dependent variables are measures of transaction costs for risky-principal and agency trades. SMCCF_t is a dummy that takes the value of 1 if day t falls between March 23 and April 9, and 0 otherwise. Eligible_t takes the value of 1 if the bond has an investment-grade rating and time-to-maturity of five years or less on March 23, 2020. X_{it} controls for log(Amt outstanding), log(Age), and log(Time-to-maturity): logs of bond's amount outstanding, years since bond issuance, and years to maturity, respectively. There are three trade size categories: less than \$100,000, between \$100,000 and \$1 million, and larger than \$1 million. The sample begins on March 6 and ends on April 9, 2020. Only US firms are included and bonds that change credit grade are excluded. Only US firms are included in the regressions. All trades for bonds with 5–6 years left to maturity are excluded from the regressions. Clustered standard errors at the day and bond levels are shown in parentheses.

	Dependent variable:							
	Risky-principal				Agency			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
SMCCF × Eligible	-60.29^{***}	-48.64^{***}	-45.86^{***}	-39.23^{***}	-11.06^{***}	-7.75^{***}	-9.60^{**}	-9.38^{**}
	(11.54)	(11.17)	(10.85)	(11.07)	(2.48)	(2.66)	(3.48)	(3.45)
SMCCF	1.29	−16.85	−11.11	−17.67	6.82^{***}	5.91^{**}	4.62^{**}	4.49^{**}
	(16.07)	(13.39)	(14.34)	(14.16)	(2.43)	(2.56)	(2.14)	(2.09)
Eligible	-35.43^{***}	-34.26^{***}			-3.48^{***}	-2.03^{***}		
	(7.80)	(10.59)			(0.63)	(0.66)		
log(Amt outstanding)	15.22^{***}	20.28^{***}			3.95^{***}	4.26^{***}		
	(4.17)	(4.21)			(0.82)	(1.17)		
log(Time-to-maturity)	24.14^{***}	25.14^{***}			5.34^{***}	6.03^{***}		
	(7.01)	(6.10)			(1.11)	(1.22)		
Trade size category FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industy FE	Yes	Yes	No	No	Yes	Yes	No	No
Bond FE	No	No	Yes	Yes	No	No	Yes	Yes
Credit rating FE	No	Yes	No	Yes	No	Yes	No	Yes
Observations	142,584	130,103	142,585	130,103	43,211	40,964	40,964	43,212
Adjusted R^2	0.04	0.05	0.20	0.20	0.08	0.10	0.25	0.24

Note: *p<0.1; **p<0.05; ***p<0.01

Table A13. Estimating the logit demand parameter σ: overidentified IV. This table presents the IV estimate of the logit demand parameter σ. Both log volume and number of customer-to-dealer trades are used as instruments. The pre-crisis runs from January 3, 2020 until February 29, 2020. The post-crisis data begins on April 15, 2020 and runs until July 31, 2020. We exclude holidays, weekends and half trading days. Standard error are the maximum of robust and the usual standard errors.

	Dependent variable:
	$(p_h - p_l)$ IV (num & vol)
	(1)
$\log(x_h/x_l)$	73.10**
	(31.04)
Post-crisis	7.78
	(5.16)
Constant	77.97***
	(25.81)
Observations	113
Adjusted R^2	0.39
Note:	*p<0.1; **p<0.05; ***p<0.01

Table A14. The first stage of the IV regressions for estimating the logit demand parameter σ. We regress the log ratio of fraction of risky-principal and agency trades on the log of the seasonally-adjusted daily aggregate number and volume of trades in columns (1) and (2), respectively. The pre-crisis runs from January 3, 2020 until February 29, 2020. The post-crisis data begins on April 15, 2020 and runs until July 31, 2020. We exclude holidays, weekends and half trading days. Standard errors are given by the maximum of robust and the usual standard errors.

	Dependent variable:	
	$\log(x_h/x_l)$	
	(1)	(2)
log(Volume of trades)	-0.22^{***}	
	(0.05)	
log(Number of trades)		-0.44^{***}
		(0.09)
Post-crisis	-0.12^{***}	-0.13^{***}
	(0.02)	(0.02)
Constant	0.82^{***}	0.84^{***}
	(0.01)	(0.01)
Observations	113	113
Adjusted R^2	0.45	0.47
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table A15. IV robustness for estimating the logit demand parameter σ: full sample. This table presents OLS and IV estimates of the logit demand parameter σ using the entire 2020 transactions from TRACE from January 3, 2020 until August 12, 2020. In column (2), the seasonally-adjusted log number of trades is used as the instrument. In column (3), the seasonally-adjusted log volume of trades is used as the instrument. The full TRACE 2020 sample is used. Crisis_t is a dummy which takes the value of 1 if day t falls into the Crisis sub-periods defined above. SMCCF_t and SMCCF Expansion_t are dummies that take the value of 1 if the trading day t is between March 23 and April 9, and after April 9, 2020, respectively. We exclude holidays, weekends and half trading days. Standard error are the maximum of robust and the usual standard errors. Our estimates and standard errors are transformed using the delta-method where appropriate.

	Dependent variable:		
	$(p_h - p_l)$		
	IV (vol)	IV (num)	IV (num & vol)
	(1)	(2)	(3)
$\log(x_h/x_l)$	107.73**	129.37***	129.79***
	(43.16)	(42.95)	(43.13)
Crisis	68.75***	64.06***	63.97***
	(17.05)	(13.21)	(13.17)
SMCCF	38.97***	33.05**	32.93**
	(14.15)	(13.56)	(13.59)
SMCCF expansion	1.80	−1.77	−1.84
	(7.89)	(7.44)	(7.47)
Constant	107.34***	125.27***	125.63***
	(35.85)	(35.70)	(35.86)
Observations	143	143	143
Adjusted R^2	0.70	0.68	0.68
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A16. The first stage of the IV regressions for estimating the logit demand parameter σ: full sample. We regress the price difference between high and low cost trades on the log of the seasonally-adjusted number of trades in that day. Standard errors are given by the maximum of robust and the usual standard errors. This uses the full 2020 TRACE sample. The sub-periods are defined in the text.

	Dependent variable:	
	$\log(x_h/x_l)$	
	(1)	(2)
log(Volume of trades)	-0.24^{***}	
	(0.05)	
log(Number of trades)		-0.48^{***}
		(0.09)
Crisis	-0.17^{***}	-0.18^{***}
	(0.03)	(0.03)
SMCCF	-0.18^{***}	-0.19^{***}
	(0.04)	(0.04)
SMCCF expansion	-0.12^{***}	-0.13^{***}
	(0.02)	(0.02)
Constant	0.83^{***}	0.85^{***}
	(0.02)	(0.02)
Observations	143	143
Adjusted R^2	0.48	0.51
Note:	*p<0.1; **p<0.05; ***p<0.01	