The Option Value of Municipal Liquidity*

Andrew Haughwout, Benjamin Hyman, Or Shachar

Federal Reserve Bank of New York

October, 2022

Abstract

We estimate the option value of the Municipal Liquidity Facility (MLF), introduced for the first time to address municipal turmoil during COVID-19. Using a regression discontinuity that exploits lending eligibility population cutoffs, we find that while the liquidity option improved overall municipal bond market functioning, *lower-rated* bonds with MLF access traded at higher prices and were issued more frequently. This suggests a potential credit-risk sharing channel on top of the Fed's role as liquidity-provider of last resort. By contrast, local governments recalled a large share of employees in response to liquidity interventions, but recalls were only sustained for higher-rated municipalities.

*The views expressed here are the authors' and are not necessarily the views of the Federal Reserve Bank of New York or the Federal Reserve System. We thank Ben Lahey, Rebecca Landau, Nicholas Ritter, and Christopher Simard for outstanding research assistance. We owe special thanks to John Bagely, Nick Frost, Kent Hiteshew, Matt Lieber, and Shawn Nee for their continued assistance with the data and expert guidance on the institutional details concerning the muni market. We also thank Kirill Borusyak, Anthony DeFusco, Gilles Duranton, Dan Garrett, Daniel Green, Paul Goldsmith-Pinkham, Bev Hirtle, Bob Inman, Ben Keys, Anna Kovner, Stephan Luck, Byron Lutz, Melissa Moye, and Wilbert van der Klaauw, for valuable feedback on our early-stage results, and various seminar participants from the Virtual Municipal Finance Workshop, UC Irvine, NYU Furman, UKY Martin, the AEI, the UEA and NTA Annual Meetings, the Federal Reserve System Regional Meetings, the Brookings Municipal Finance Conference, and the NBER Summer Institute for their comments. Lastly, we are grateful to Gray et al. (2020) for sharing their codebase with our team. All errors are our own. Contact: andrew.haughwout@ny.frb.org; ben.hyman@ny.frb .org; or.shachar@ny.frb.org.

1 Introduction

The State and Local sector is the dominant government service provider in the US, employing nearly 20 million workers. State and local government activity directly contributed 10.9 percent of GDP in 2019, well above the federal sector's direct contribution of 6.6 percent (BEA, 2020). A long-standing debate concerns the optimal level of debt with respect to growth rates, yet less is known about the extent to which municipalities are liquidity constrained and thus what the value of lifting such constraints would be, especially in periods of fiscal distress. Particularly important to municipal functioning but often overlooked is the short-term municipal bond market, which allows localities to maintain smooth spending paths when revenues and expenditures are temporally misaligned. In this paper, we estimate the option value of municipal liquidity by exploiting the Federal sector's primary market intervention in municipal debt markets in response to the COVID-19 shock.

The COVID-19 pandemic left state and local governments to manage both health and economic crises, necessitating borrowing from an almost frozen municipal bond market. The Municipal Liquidity Facility (MLF) was then introduced by the Federal Reserve in April 2020 to provide stop gap cash flow financing to states and municipalities by purchasing up to $500 billion of short-term notes issued by these entities. That is, the facility ensured that municipalities would have had a willing buyer of their debt at a predetermined interest rate. Eligible issuers were initially limited to US states, cities with populations greater than 1 million, and counties with population greater than 2 million. By targeting the largest issuers, where the bulk of trading occurs, the thinking was that this push would help the market stabilize, while additional issuers could be covered by the parent state's ability to "downstream" to ineligible issuers. On April 27, 2020, MLF lending eligibility cutoffs were *expanded*, granting direct Federal Reserve support to city issuers with populations greater than 250,000, and county issuers with populations greater than 500,000. The new population thresholds significantly expanded the number of entities that could borrow directly from the MLF – 87 cities and 140 counties as well as their political subdivisions and instrumentalities.

The April 27, 2020, change in MLF eligibility allows us to undertake a regression discontinuity (RD) design around these new population cutoffs. The underlying identification relies on the idea that outcomes for entities that fall just around the population cutoff have very similar ex-ante characteristics, and that outcomes would have evolved smoothly absent the cutoff. The April 27, 2020, reform provides sufficient mass close to the cutoffs to estimate the causal effect of the option to access primary market liquidity on improvements in secondary market functioning, and retention of public sector employees. This empirical strategy allows us to make three novel contributions to our understanding of the Fed's role as a buyer of last resort: (1) the RD provides a unique setting to estimate credible causal effects;

(2) the intervention's "primary-only" approach, which contrasted interventions in the corporate bond market, allow us to test the extent to which a combination of primary and secondary interventions are required for market stabilization; (3) RD heterogeneity across the credit ratings distribution provides a novel mechanism test for understanding the liquidity versus credit effects of the intervention.

First, combining the cross-sectional variation from the RD with the timing of the policy announcement helps us overcome the empirical challenge of separately identifying the effects of the initial MLF announcement from other interventions, such as the Payroll Protection Program (PPP) and the Main Street Lending program, that were simultaneously implemented around the same period.

Second, unlike the other Federal market interventions in the U.S. Treasury and corporate bond markets, the MLF offered solely *primary* market liquidity. This allows us to study the role of the Fed as a lender separately from its actions as a market maker of last resort, and to offer a quantitative estimate for the effectiveness of primary market interventions in stabilizing both primary and secondary markets. Importantly, though the facility terms target primary issuance of notes with maturities of up to 36 months, we find substantial impacts on long-term secondary market yields. To our knowledge, this feedback channel between short-term primary issuance and long-term secondary market liquidity in the municipal sector has not been previously documented. One example of how this could happen is short-term notes could be directly applied to long-term general obligation (GO) debt servicing. Short-term optionality could enter the operational budget affecting long term pass-through if localities have flexible constitutions or statutes which permit such fungibility. More generally, to the extent that cash is fungible in municipal budgets, this channel could be at work.

We start with analyzing yields in the secondary market and issuance in the primary market during the COVID-19 crisis. We find that, for the most part, municipal bond markets returned to normal functioning following the totality of Federal interventions that were introduced between mid-March and the end of April 2020. Nonetheless, low-rated investment grade (IG) bonds remained relatively distressed. Using our RD, however, we find that low-rated IG (A and BBB) city and county issuers' bonds traded at higher prices immediately following facility access. Yields of low-rated IG issuers that were narrowly eligible for emergency lending declined by roughly 75 basis points (bps) relative to observationally equivalent issuers that narrowly missed eligibility—a sizable magnitude that closely resembles the overall market spread between BBB yields and higher-rated issuers. These differences in investor perceptions also translated to a large differential increase in primary issuance among low-rated IG issuers. Although there were only two issuers issuing new bonds *directly* through the MLF, the *option* to borrow from the facility at back stop rates improved issuers' ability to issue at favorable terms on private markets. These results imply that absent the MLF, issuers would have likely been constrained in their ability to issue new debt throughout the crisis.

Next, extending the contribution of the paper beyond the evaluation of MLF, we uncover the link between market liquidity and the real economy. We study whether municipalities' smoother funding paths induced local budget officers to retain more public sector employees or increase hours in light of mass furloughs and separations. Smoother municipal funding paths have at least two mechanisms through which the provision of such liquidity could be related to employment decisions. First, short-term borrowing may serve to preserve the cash flow required to maintain employment when receipt of revenues is delayed or otherwise mismatched with the timing of payroll. They may also support positive externalities emanating from high fiscal multipliers (Chodorow-Reich, 2019, and Yi, 2020). Second, and less obvious, is a long-term debt refunding channel. By facilitating the issuance of debt at lower cost, the availability of liquidity may allow the issuer to lower its debt service costs through refunding (Ang et al., 2017), thereby freeing up additional resources for service provision and reducing the duration of an economic downturn (Auerbach et al., 2020). This second channel is especially relevant when interest rates are falling, as they were during the early stages of the COVID-19 pandemic.

We focus on the MLF's impact on May and June 2020 public sector employment outcomes, and also estimate a full path of longer-run effects. While our results show a tight link between the MLF and improvements in market functioning, the time lag between the MLF announcements and their potential real outcomes limits our ability to fully attribute the effect to MLF. Specifically, the CARES Act announced on March 27 also provided access to aid for cities and counties with populations over 500,000, whereas those under this cutoff had to rely on their parent states for downstreamed aid. This raises the possibility that estimates capture a combination of CARES aid *and* MLF interventions through one of the two cutoffs (i.e. 500,000), though both policies reflect similar liquidity interventions and thus likely share the same theoretical sign on employment and underlying economic interest. Nevertheless, we are able to show that city and county governments retained more service-providing government employees in response to the *totality* of both the MLF option and direct CARES aid. In our preferred specification, eligible issuers retained about 422 to 517 more local service-providing employees relative to observationally equivalent ineligible local governments in the two months following the cutoff announcement. On a baseline year-on-year decline of about 1,674 employees, this reflects a substantial recovery effect size of roughly 25% to 30%. Almost all of the retained service-providing employees were concentrated in *Education and Health Services*, the majority of which is comprised of educational institution employees.

Consistent with priority hiring for education sector employees, year-on-year RD estimates are positive during lockdown months, revert to zero during summer months when school labor demand is low, with positive effects sustaining beyond the summer for *high-rated* governments (the vast majority of issuers). The result that less fiscally constrained governments appear more elastic in hiring, especially while shutdowns were largely still in place, is consistent with the hypothesis that state and local governments may have over-weighted the worst possible outcomes based on past experience, furloughing education sector employees even though realized revenue shortfalls were far lower than originally anticipated (this view was also articulated in Sheiner, 2021).

In the remainder of the paper, we probe why MLF access was perceived as more valuable lower down on the ratings distribution rather than as neutral across the distribution, and discuss implications of our results for optimal policy. We use two further strategies to gauge the mechanism behind this result. We first estimate the probability that one of the three major nationally recognized statistical rating organizations (NRSROs) downgraded any of the issuer's bonds, and find modest evidence that downgrade probabilities differentially increased when the issuer was revealed as ineligible for facility access. Using a complementary asset price decomposition approach following Boyarchenko et al. (2022) and Schwert (2017), we also quantify the credit-risk channel over a broader set of issuers, which implies a non-trivial role for credit risk in determining yields during peak distress. These results combined suggest the presence of a potential credit-risk sharing channel beyond the Federal Reserve's role as a buyer of last resort – a willingness to credit-risk share even if not widely exercised in practice.

Given our new estimate of credit-risk sharing, a parameter which enters a broader welfare calculus regarding the efficiency of macroeconomic stabilization policies, one natural question is when does municipal credit-risk sharing represent a socially efficient policy improvement? From a financial frictions viewpoint, several influential papers have suggested that sub-optimal risk sharing could arise if institutional investors are overly concentrated in locally exempt bonds—a home market bias that distorts the efficient spread of risk (Poterba, 1989, Pirinsky and Wang, 2011, Schwert, 2017, Babina et al., 2021). Other market structure constraints may also result in inefficient muni market pricing (Garrett, 2020). When such frictions are also heterogeneously correlated with the mean income of underlying geographies, inequality weights must also be taken into consideration. Future work will need to contextualize the magnitude of potential externalities and multipliers from government functioning, including the social value of public services, to help answer this question. A final potential implication of our empirical results is that a "primary-only" intervention may have sufficiently persistent effects that a smaller sized intervention may achieve similar results in other markets.

Contribution and Related Literature To our knowledge, our paper is the first to establish a *causal* link between MLF optionality, market functioning, and issuers' health. Contemporaneous papers, Bi and Marsh (2020), Bordo and Duca (2021), Li and Lu (2020), and Fritsch et al. (2021) that study the impact of MLF on the functioning of the municipal bond markets use time series or event study approaches, which can be biased by other interventions that were implemented during the same time period (and in the case of MLF, the same day).

Our paper is also distinct from other contemporaneous papers that have examined interventions in other markets, such as the corporate bond market (e.g., O'Hara and Zhou, 2021, Boyarchenko et al., 2022, Haddad et al., 2021), that also show a positive impact of interventions on market functioning. While participant heterogeneity has been shown to play an important role in the design of market interventions, the municipal sector's diverse landscape (see e.g., Li and Schürhoff, 2019, and Cestau et al., 2019) would have likely generated considerable challenges in designing a secondary market intervention as was done in corporate credit markets. However, studying the transmission mechanism of primary market support in the municipal bond market when a direct secondary market intervention is unavailable allows us to better isolate the Fed's impact as a buyer of last resort. Our findings show that while interventions in the primary market of the municipal bond market improved the functioning of the secondary market, both the path and cross-sectional impact have been distinct from the experience of corporate bond markets. This is especially important to study further given the stickiness in issuer-underwriter relationships as shown in Chen et al. (2022), and the interaction between municipal bond mutual fund flows and issuance size and frequency as shown in Adelino et al. (2021).

Second, our employment and credit downgrade results are consistent with Adelino et al. (2017), who find that credit ratings impact borrowing constraints, and subsequently have real implications for employment and expenditure outcomes; as well as Cornaggia et al. (2018) who find a significant link between credit ratings, yields, and borrowing costs. Novel to our setting, and to the COVID-19 pandemic, we find that recovery in municipal employment (the first real-economy margin likely to respond to a municipal bond market intervention) was driven by the education sector.

Third, we complement the municipal bond literature studying yields components (e.g., Ang et al., 2014, Wang et al., 2008, Schwert, 2017). In particular, our findings are consistent with Wang et al. (2008) and Schwert (2017), who attribute a large proportion of the observed yield to default risk in municipal bond pricing. Our estimates of MLF access imply a large investor perceived spread, but only among the lowest rated issuers, which we interpret as reflecting credit downgrade and default risk.

Our study also builds on methodological contributions from a series of municipal bond market papers, including Harris and Piwowar (2006), Green et al. (2007b), and Green et al. (2010), who study secondary market transaction costs; and Green et al. (2007a) and Schultz (2012) who focus on the pricing and price dispersion of newly issued bonds. We also draw from work by Novy-Marx and Rauh (2012) quantifying the sovereign default channel. Lastly, related to evaluating Federal Reserve facilities, our methodological approach most closely links to work by Moore (2017) who analyzed the Term Auction Facility by comparing marginal winners and losers, and Luck and Zimmermann (2020) who find strong employment effects from quantitative easing (QE) policies via bank lending channels.

Finally, Auerbach et al. (2020) examine federal and state fiscal conditions in the context of the pandemic, both generating their own estimates and evaluating other papers. They find generally, that the pandemic has induced smaller tax losses than may have been initially expected, for example, when imputing expected losses from the Great Recession. One specific paper examining the effect of CARES Act aid on employment during COVID-19, Green and Loualiche (2020), finds that sales-tax reliant localities with stricter balanced budget requirements (BBRs) (proxied by lower rainy day funds) likely suffered greater job losses due to an inability to deficit spend across budget cycles, and it was in those cases that fiscal aid was more effective. This finding, in part, prompts us to include state fixed effects in our preferred specification to absorb variation across states emanating from the stringency of their underlying BBRs, yet our main results do not rely on these fixed effects.

The rest of the paper proceeds as follows. Section 2 provides institutional background on municipal debt markets and the unique market distress exhibited during the COVID-19 pandemic. In Section 3, we discuss the main bond market and real economy data sets used, matching procedures and sample restrictions. In Section 4 we first describe our methodological approach and and then present the main results on bond-level outcomes, public sector employment, and mechanism tests. Section 5 complements the RD analysis by studying yields' components. We provide sensitivity and robustness tests to our results in Section 6, and conclude with a discussion of what is learned from the main results in light of a perceived low take-up puzzle regarding the number of issuers that made use of the Municipal Liquidity Facility in Section 7.

2 Municipal Distress and Interventions during COVID-19

In this section we discuss municipal bond market conditions in the run-up to the pandemic, how recent dynamics may have diverged from prior trends and contributed to peak disruption, and a detailed timeline of Federal sector interventions that supported municipal debt markets during this period.

2.1 The Municipal Bond Market

The $3.8 trillion municipal bond market contains more than 50,000 issuers and 1 million individual bonds, making it approximately half the size of the corporate bond market with 10 times as many issuers.

Roughly 90% of this market is exempt from federal income tax,[1] and more than 80% is rated investment grade. Consequently, municipal bond default rates have historically been low (Appleson et al., 2012). As of May 2020, 26% of outstanding debt was issued directly by state, city, county, and other local governments, 41% by utilities, service, and transit issuers, 21% by school districts, and 8% by public hospitals. Unlike treasury and corporate bond markets, 70% of municipal debt is held by retail investors (a third of which is in mutual funds and ETFs) seeking tax advantages associated with municipal bond returns.[2] Unlike the corporate sector, municipal debt is also commonly issued in deals containing many different tenors as independent bonds, facilitating more predictable budget smoothing.

Government issuers in this market are often required to balance their operating budgets, and can usually only borrow long term in order to finance infrastructure investments. In 2015, 48 of 50 US states had some form of balanced budget requirement (BBR), while 39 of 50 had strong constitutional or statutory requirements (Brookings Tax Policy Center, 2015), limiting deficit spending across fiscal cycles.[3] GO bonds, which constitute approximately 30% of the long-term municipal market, are not secured by a specific revenue source but are instead backed by the "full faith and credit of the taxing authority" and typically finance capital projects like bridges and schools. The large remainder of the long term market (60%) is dominated by Revenue bonds issued by public enterprises and secured by defined revenue sources (such as transit revenue, airport feels, bridge tolls, etc.).[4]

Less well known but central to this paper, state and local governments also frequently leverage the $440 billion short-term municipal note market to bridge cash flow gaps within fiscal years. As localities depend on revenues that are only received at specific intervals, budget officers seek to smooth spending in anticipation of such receipts, and can do so by issuing tax anticipation notes (TANs), revenue anticipation notes (RANs), and bond anticipation notes BANs. Other notes include tax and revenue anticipation notes (TRANs), tender option bonds (TOBs), and Variable Rate Demand Notes (VRDNs), the latter of which comprises roughly half of the short-term market. Short-term notes are typically secured by the revenues expected to be received later in the fiscal year, and are paid off when said revenues arrive. A classic example is the proceeds from final settlements of state income tax returns due April 15 (the federal tax filing deadline), which can include taxable unearned income and capital gains (typically not withheld). Other examples include quarterly property tax receipts, or expected surges in airport fees during the holiday season.

[1]Some bonds are "double-exempt" for local residents, applicable to tax liabilities at both Federal and State levels (e.g. CA state bonds for CA residents), whereas others are "triple-exempt" at the local level as well (e.g. New York City bonds).

[2]Outstanding debt by sector calculated from Bloomberg, as of 5/19/2020. Retail share calculated from the Board of Governors, "Financial Accounts of the United States, Z.1, as of 1Q2020." Retail holdings calculated from MSRB.

[3]Some states further prohibit or limit GO bond issues, including Arizona, Colorado, Idaho, Indiana, Iowa, Kansas, Kentucky, Nebraska, North Dakota, South Dakota, and Wyoming.

[4]MSRB, Bloomberg, calculated as of 5/21/2020.

While many of these features of the municipal market have remained relatively constant for several decades, one notable change in recent years has been the increasing share of muni holdings in mutual funds, which nearly doubled to $800 billion since 2010 (from 10% to 20% of the market) and by some measures now reflect one third of overall muni holdings (WSJ, Oct 2019). Importantly, municipal holdings by mutual funds have also become more concentrated, with over 70% of muni mutual funds held by the 10 largest fund families (investment banks).[5] This trend played a key role in the municipal liquidity crisis when the COVID-19 shock hit financial markets.

2.2 Turmoil in Municipal Debt Markets

Figure 1 shows the performance of secondary market municipal bond yields during this period, and highlights the liquidity crisis that occurred in early- to mid-March. The left panel shows mean weekly yields for the universe of secondary-market traded AAA muni bonds—the safest in the market—as a ratio to comparable US treasuries. In the period leading up to the mid-March turmoil, municipal bonds were trading at yields below US treasuries (ratios below 1) due to their special tax-exempt status. The market then deteriorated with yields rising steeply for all bond tenors but most sharply for short-run debt, reflecting deteriorating demand conditions as investors anticipated severe dislocation in the near term.[6] The precipitous spike in yields has been partially attributed to selling by municipal bonds mutual funds. After a record surge of muni inflows into mutual funds that netted $90 billion over the 12 months ending in February 2020, inflows reversed sharply in March 2020, with municipal bond mutual funds experiencing outflows of $43 billion in March (Cipriani et al., 2020a) as investors rushed for liquidity. The ensuing sell-off may have flooded the market due to seller concentration (Li et al., 2021), bringing muni demand and new debt issuance to a near standstill by mid-March.

After the first Federal intervention in the muni market was announced on March 23 (indicated by the left-most dotted line in Figure 1) yields started to decrease. Yields finally reverted back close to their pre-COVID levels in June 2020 after a series of further Federal interventions. The right panel of Figure 1, however, reveals significant heterogeneity in this resumption to normalcy—the lowest rated IG bonds (BBB) continued to exhibit distress despite recovery in lower risk credit bins. One potential explanation for these heterogeneous recovery rates is that investors may have struggled to assess the change in credit default risk for low-rated issuers weathering the first wave of the pandemic, consistent with notoriously low price discovery in municipal markets due to their low trading frequency.[7] As Green

[5] Calculations from "Financial Accounts of the United States, Z.1, as of 1Q2020".

[6] It is common practice in the municipal bond market that newly issued bonds target an at-issuance coupon rate of 5% by modifying initial offering prices. Because coupon interest is fixed at issuance and comprises the numerator of the current yield metric shown in Figure 1 (with prices reflected in the denominator), large spikes in yields thus tend to reflect declining prices as demand shifts downward.

[7] One additional example of opaqueness in municipal debt markets relates to undisclosed private muni loans, extensively

et al. (2010) note, illiquid assets may "rise faster than they fall". This differential recovery motivates our main heterogeneous specifications in which we examine effects separately by pre-crisis creditworthiness levels.

The primary market for municipal bonds was also affected. As the pandemic set in, governments endured a clear primary issuance shortfall relative to historical trend, which approached a market freezing point at its nadir around March 20th (see Figure A.11.). This is not surprising given that the ability of state and local governments to issue new municipal bonds on primary markets is critically linked to secondary market performance. If demand is low for bonds already traded in the market, then the additional supply of new bonds, like those needed to bridge government cash flow shortfalls, may be both constrained, and potentially issued at higher cost than during an expansion—even in the presence of a monetary policy environment with lower prevailing interest rates. Local governments supplying new bonds in this environment may face increased intermediation fees from market makers and underwriters seeking to arrange initial offerings, and may have to resort to compensating investors with high interest rates for retaining illiquid positions or additional credit risk.

After a series of Federal interventions, the issuance trend then mean reverts and appears to compensate for missing issuance. Especially apparent is a significant issuance spike in mid-July, when local and state income tax receipts and non-withheld income such as capital gains were realized (deferred in 2020 due to the IRS federal tax deadline extension from April 15 to July 15, enacted as a matter of fiscal policy). Cumulatively, though at least the end of June, 2020 year-to-date primary issuance had narrowly outpaced its five-year historical levels, though as we discuss below, this uptick may not necessarily reflect the socially efficient level of primary issuance given the unusual supply needs of municipal issuers during the pandemic.

2.3 Pass-Through of Market Distress to Budgets and the Real Economy

In a normal expansionary year, such as 2019, state and local governments incur expenses such as payroll, debt service, etc., but only receive expected revenue at distinct intervals, as shown in the timeline below. One prominent example of such a revenue source is the proceeds from final settlements of state income tax returns, typically due April 15 (the federal tax filing deadline), which can include taxable unearned income (typically not withheld).[8] When governments are short of revenue yet have guaranteed sources of future receipts, they can issue short-term anticipation notes on municipal bond markets to help smooth out the temporary shortfall. Take for example, the state of New York, whose tax base includes large amounts of capital gains realized in the prior calendar year. Taxes on these amounts are typically

documented in Ivanov et al. (2021).

[8]Other examples include quarterly property tax receipts, or expected surges in airport fees during the holiday season.

received on April 15. New York may issue a 4-month TAN in anticipation of this revenue on Jan 1, 2019, maturing April 30th, 2019 (one month after its fiscal cycle closes). This would generate bond proceeds to smooth out cash flow, and would be secured by revenue incoming on April 15, enabling the government to start spending part of the expected settlement amount in the interim.

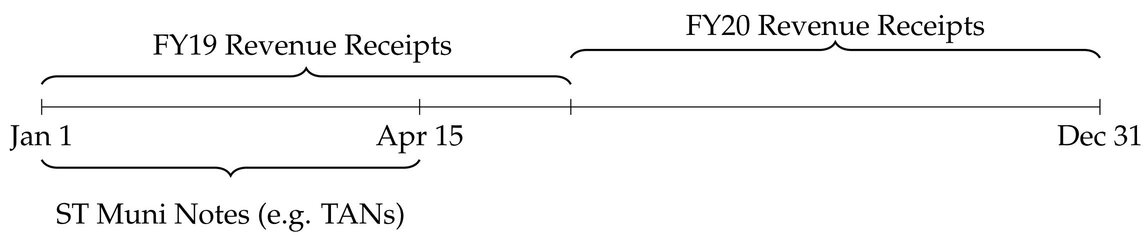

However, when an economy experiences a major income shock accompanied by a bond market liquidity crisis, as was the case with the COVID-19 pandemic and the muni market, there are three distinct ways in which this enters fiscal budgeting, shown in red in the timeline that follows.

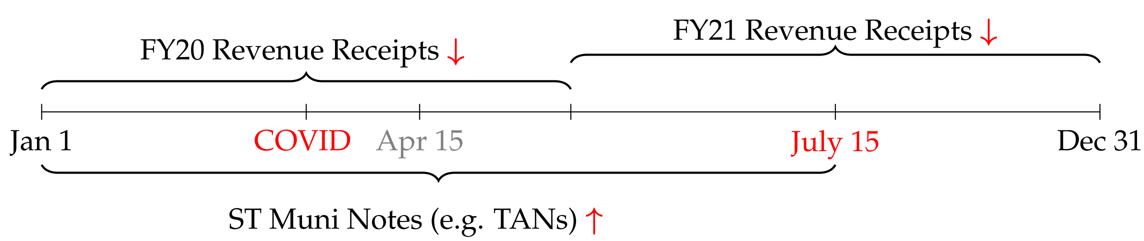

First, income receipts decline in both the current and upcoming fiscal cycles. While tax assessments may be revised downwards, 2021 budgets were for the most part already set when the COVID pandemic took hold. This led to a decline in the revenue that would normally cover these planned expenditures. Second, as a matter of fiscal policy, the IRS extended the 2020 federal tax deadline from April to July 15th, increasing the quantity of short-term notes governments desired to issue on the market to plug the additional cash flow gap. Third, while the temporary misalignment of expenditures and revenue needs (even if unexpectedly large in a specific year) are normally able to be remedied through the bond market, the simultaneous large financial sell off of fixed income assets effectively drove investor demand for munis toward zero, leaving governments unable to borrow. This was the challenge faced by the US economy in mid-March of 2020; local governments could not borrow on private municipal markets despite their increasing need to do so because of demand conditions, potentially constraining payroll obligations.

2.4 Federal Interventions during COVID-19

In light of both *real* municipality liquidity challenges and a municipal *market* liquidity crisis, several Federal Sector agencies undertook interventions to support municipal bond markets. In this section we describe the key interventions that relate to this paper.

March 23, 2020: MMLF Accepts Some Types of Municipal Debt as Collateral: Initially announced on March 18, 2020, the Federal Reserve System expanded the Money Market Mutual Fund Liquidity Facility (MMLF)—which makes loans to financial institutions secured by high-quality collateral—to accept variable rate muni demand notes (VRDNs) as pledgable collateral.[9] Henceforth, we refer to the time period prior to this first announcement as the "pre-period" or placebo period when studying effects of subsequent facility announcements before any were yet established. However, it bears noting that two days earlier, the Internal Revenue Service (IRS) and US Treasury formally extended the Federal tax filing deadline from April 15 to July 15 to provide household tax payment relief on 2019 liability payments due in April.[10] Further, on this same day, several additional new Federal Reserve facilities were announced or expanded.[11] For these reasons, any time series analysis across these various announcements have the potential to confound effects from a bundle of interventions.

April 9, 2020: MLF Announced. As part of a broader $2.3 trillion support package that included support for the Payroll Protection Program (PPP), Main Street Lending program, expansions to corporate credit facilities (PMCCF and SMCCF) as well as the Term Asset-Backed Securities Loan Facility (TALF), the MLF was announced on April 9 "to help state and local governments better manage the cash flow pressures they are facing as a result of the increase in state and local government expenditures related to the COVID-19 pandemic and the delay and decrease of certain tax and other revenue."[12] The MLF was established with a $35b US Treasury initial equity investment (appropriated through the March 27 CARES act) as backstop for a maximum of $500b in direct Federal Reserve lending of short-term notes (TANs, RANs, TRANs, BANs) to municipal issuers via a Special Purpose Vehicle (SPV).

There was considerable concern in the design phase of the facility that facing up to 50,000 issuers would render the operational arm of the facility less effective, and risk a slower speed to intervention.[13] Consequently, eligible issuers were initially limited to US states, cities with populations greater than 1 million, and counties with population greater than 2 million. By targeting the largest issuers, where the bulk of trading occurs, the thinking was that this push would help the market stabilize, and additional issuers could be covered by the parent state's ability to "downstream" to ineligible issuers. That is, "downstreaming", whereby for example, a state or higher-level government issues

[9]The MMLF was established "to prevent outflows from prime and muni money market funds from turning into an industry-wide run, as happened in September 2008", and was supported by $10 billion of credit protection from the U.S. Treasury's Exchange Stabilization Fund (see (Cipriani et al., 2020b)). While the MMLF had accepted some forms of short-term municipal debt on March 20, the types of pledgble collateral were greatly expanded on March 23, so we start the analysis here.

[10]See March 21, 2020 IRS Extension Announcement for details.

[11]See Federal Reserve System Board of Governors Press Release, March 23, 2020 "Federal Reserve announces extensive new measures to support the economy" for a description.

[12]See Federal Reserve Bank of New York MLF Page, April 9, 2020 and April 9, 2020, MLF Term Sheet for further details.

[13]See Examination of the Municipal Liquidity Facility, 09/17/20, Congressional Oversight Committee

on behalf of a cash-flow managing instrumentality or political subdivision is permitted, but the extent of downstreaming is regulated by state legislative approval. Further, the MLF established maximum borrowing caps for each issuer at 20% of each issuer's 2017 "Own-Source General and Utility Revenue" (OSGUR), which implicitly limits the extent of downstreaming.[14] In our setting, the presence of downstreaming is expected to attenuate any effects we find through our regression discontinuity analysis, as ineligible issuers are more likely to be beneficiaries of downstreaming when subdivisions do not already have direct MLF access. We thus interpret all RD estimates as a lower bound on the direct MLF access effect in this paper.

Tenors were originally limited to 2-year maturities, and priced at a penalty rate to a private market index by issuer credit rating.[15] The lending eligibility cutoffs established were also initially distinct from cutoffs used to determine direct CARES Act aid through the Coronavirus Relief Fund (CRF), through which cities and counties above 500,000 in population were eligible for more direct aid—an important consideration we return to when interpreting effects on public sector employment.[16]

April 27, 2020: MLF Population Cutoffs Expanded. Central to our empirical strategy, on this day eligibility was expanded to cities with populations greater than 250,000, counties with populations greater than 500,000, and multi-state entities were included as eligible issuers.[17] Tenors were also extended to allow up to 3-year maturity notes to provide maneuvering room for localities in which Balanced Budget Requirements (which usually preclude deficit borrowing across fiscal cycles) are statutorily amendable or flexible. Terms regarding minimum credit ratings were also established—"an Eligible Issuer that is not a Multi-State Entity must have been rated at least BBB-/Baa3 as of April 8, 2020, by two or more major nationally recognized statistical rating organizations ("NRSROs"). An Eligible Issuer that is not a Multi-State Entity and that was rated at least BBB-/Baa3 as of April 8, 2020, but is subsequently downgraded, must be rated at least BB-/Ba3 by two or more major NRSROs at the time the Facility makes a purchase." (MLF Term Sheet, April 27, 2020). Finally, this announcement added an explicit sunset date of December 31, 2020.

Extensions: Several notable extensions were announced thereafter, including the publication of the

[14]OSGUR caps use 2017 as a base as it reflects the latest year prior to 2020 in which the U.S. Census Bureau's Annual Survey of State and Local Government Finances administers a full government census (conducted every 5 years, with sampled surveys in the interim). For a list of OSGUR caps per issuer, see Federal Reserve Municipal Facility Limits.

[15]Details regarding the penalty pricing grid are shown in Appendix A.3.2.

[16]CRF aid was first allocated to states based on population, with a floor for small states giving them greater aid per person. Within states, any jurisdiction with population greater than 500,000 received *direct* US Treasury access to the population share of the state allocation * 45%, whereas the remaining 55% of was controlled by the state. By contrast, localities with populations *under* 500,000 had to rely on downstreaming from their underlying states (who in this case controlled 100% of the local allocation), leaving aid to smaller localities potentially more politicized, less certain, or slower to materialize.

[17]The MLF term sheet defines multi-state entity as an entity created by a compact between two or more States, approved by the United States Congress, acting pursuant to its power under the Compact Clause of the United States Constitution MLF Term Sheet, April 27, 2020. These include for example, authorities like the Port Authority of New York and New Jersey.

MLF pricing grid (May 11) and revision to the pricing grid (August 11). On June 3, the MLF was expanded to allow state governors to specially designate up to 2 ineligible cities or counties as eligible, as well as two revenue-bond issuers (RBIs) per state.[18] The cap of governor-designated city and county issuers varies between 0 and 2 per state depending on the number of already-eligible cities and counties, as some small states had fewer eligible issuers according to the April 27 cutoffs.[19] Like downstreaming, to the extent that this Governor privilege is utilized when issuers do no already have direct access to the MLF, our RD estimates should be interpreted as lower bounds on the true effect. Because these subsequent extensions were minor relative to earlier announcements, we label all activity beyond April 27 as the "post-period".

3 Data Sources and Underlying Variation

3.1 Linking Trades and Primary Issuance to Census Populations

We construct a dataset that allows us to link high-frequency municipal bond trades and new primary issuance to MLF-eligible and ineligible localities. We begin with the universe of unique Bloomberg issuers which are classified by type (e.g. city, county, state, school district, etc.) at the 6-digit base CUSIP-level. These contain both issuers with active debt, as well as small issuers who do not have (or never had) outstanding debt, and allow us to first flag all city and county issuers. We then acquire and clean the universe of secondary market municipal bond trades from the Municipal Securities Rulemaking Board (MSRB) via their Electronic Municipal Market Access (EMMA) service. Since 1998, MSRB has required registered dealers to report all municipal bond transactions, which include information about their CUSIP, date and time of trade, price and yield, maturity, coupon, par volume traded, total amount of the bond, and a flag for whether the dealer bought from a customer, sold to a customer, or whether the transaction was an inter-dealer trade. Importantly, the detailed security description has full information on the issuer's name and geography.

For our RD strategy, the main analysis period is contained in calendar year 2020, however, to provide an additional year of pre-trend we keep all trades occurring after January 1, 2019, and end our sample frame on November 20, 2020. After restricting our analysis to city and county issuers, this results in 2,849,015 trades of 194,068 bonds across 7,936 unique issuers in our MSRB-Bloomberg matched data.[20] To establish issuer lending eligibility cutoffs, the MLF makes use of two specific Census Bureau files: a

[18]While RBI issuer access to the MLF may be reflected in aggregate estimates of market yields, our main RD design is restricted to cities and counties and thus remains insulated from this designation—we therefore do not focus on it.

[19]See MLF Term Sheet, June 3, 2020 for further details.

[20]Cities: 1,930,235 trades, 141,976 bonds, 6,015 issuers; Counties: 918,780 trades, 52,092 bonds, 1,921 issuers.

2018 population list for cities and towns, and a 2019 population list for counties.[21] These files contain detailed place (government) names and populations, but they do not include CUSIP identifiers. We thus extensively clean MSRB issuer names to match these Census place name lists, and are able to match 6,015 / 6,842 of MSRB city issuers and 1,921 / 1,947 of MSRB county issuers to populations, respectively.

Figure 2 shows the distribution of matched city and county issuers and trades prior to all Federal announcements, focusing on municipalities with populations above 100,000 and cities less than 1.5 million for exposition. First, we note that the initial MLF eligibility cutoff established on 4/9/20, while able to target the most traded bonds, ultimately resulted in very few eligible issuers (blue bars). Were this cutoff to have remained in place, 24 issuers with 7,561 outstanding bonds would have been MLF-eligible. The revision of cutoffs to their 4/27/20 levels (red bars) expanded eligibility to a total of 203 issuers with 37,690 outstanding bonds (prior to the pandemic).[22] For these reasons, we focus all subsequent analysis on the 4/27/20 cutoffs, as any variation across the initial cutoffs would be smooth across the new thresholds. While a formal Frandsen (2017) test (analogous to a McCrary, 2008, test for a discrete running variable) is provided in Section 6, the smooth mass around the new cutoffs combined with institutional details behind the selection of 250,000 and 500,000 as cutoffs (including their apparent use of round number heuristics) also provides preliminary evidence that cutoffs were not chosen to target specific types of issuers or anomalies in the population distribution.

To study the effects on new primary issuance, we turn to Mergent which contains the universe of municipal issuance. In addition to the 194,068 trading bonds in our MSRB post-2019 sample, we are able to add an additional 9,002 bonds from Mergent that are newly issued but do not trade.[23] Mergent also provides key bond characteristics, including offering amount, the source of funds and use of proceeds, coupon type (fixed, variable, or zero), the tax status of the coupon payments, callability and first call date, insurance status and the identity of the insurer, and pre-refunding status and timing. We discuss these variables, cleaning procedures, and tax adjustments in detail in Appendix A.3.

To analyze effects by creditworthiness, we construct time-varying issuer credit ratings from ratings data on their underlying bonds. We first acquire monthly ratings updates at the bond level from the three major NRSROs: Standard and Poor's (S&P), Moody's and Fitch.[24] At the CUSIP level, we ascribe

[21] U.S. Census Bureau, "Annual Estimates of the Resident Population: April 1, 2010 to July 1, 2018" for cities, as of April 6, 2020; and U.S. Census Bureau, "Population, Population Change, and Estimated Components of Population Change: April 1, 2010 to July 1, 2019" for counties, as of April 6, 2020. As these files cover populations greater than 50,000, we further use U.S. Census Bureau 2010-2019 Populations, All Places Files to fill in populations when less than 50,000. This name cleaning process is involved, and described in detail in Appendix A.3.3.

[22] Of all eligible MLF issuers, only Madison, WI, did not have trading bonds after Jan 1 2019 in our data. We otherwise have fully coverage of city and county eligible issuers.

[23] Among city and county issuers, non-traded bonds are generally constrained to localities with populations less than 25,000, and therefore newly issued bonds are largely irrelevant to our analysis of yields. We otherwise, are unable to match 25% of Mergent non-trading bonds to Census populations due to mismatches in Mergent issuer names and Census issuer names.

[24] These are assembled via Bloomberg.

all trades that occur within the month following each ratings observation their long-term, short-term and muni bond rating (when available). We then concord all long-term, short-term, and muni-ratings—each system specific to their NRSRO—to an aggregated S&P Ratings bin (AAA, AA, A, BBB, BB, CCC, C, D) using a custom concordance file.[25] Finally, we construct issuer plurality ratings within issuer-month, equal to the most common aggregated bond rating across NRSROs and CUSIPs that month. This allows us to subset on pre-pandemic fixed issuer ratings when implementing heterogeneous specifications. In Figure 3 below, we show the resulting distribution of plurality ratings (at the trade-level) when fixing the rating to January (or February) 2020, and when letting the plurality rating vary by month.

Figure 3 shows that the ratings distribution is relative stable over time, as well as heavily skewed toward high-credit issuers. In panel (b) we show the distribution for an arbitrary bandwidth of 100,000 around city and county cutoffs, highlighting the distribution stability of this sub-sample (further discussed in Section 3.2). To the extent that there are monthly credit downgrades that occur differentially after the pandemic, or for MLF-eligible versus non-eligible issuers, we use disaggregated ratings and the bond level. This allows us to ask whether the probability that a bond was downgraded relative to the prior month changed overall, and heterogeneously by baseline plurality ratings (for example within January 2020 BBB's). Generally, we are able to match 96.3% of trading bonds to a monthly plurality rating, where 3.7% of the unmatched reflect the inability to find a plurality when ratings are split.

3.2 Underlying Variation in Yields

We now turn to highlighting the underlying variation in our matched sample. In our main RD design specification, we choose the running variable's window of analysis using an optimal bandwidth selection procedure (following Calonico et al., 2014) which takes choice and publication bias out of the econometrician's hands. In this section, we first show some of the key variation underlying our design using an intuitive symmetric population bandwidth of 100,000 below (ineligible) and 100,000 above (eligible) MLF population cutoffs. In Figure 4, we show unconditional weekly mean yields (i.e. without covariate adjustments) separately by issuers narrowly eligible for MLF (blue dashed series), and similar-sized issuers that narrowly miss MLF eligibility (red solid line).

Panels (a) and (b) show city and county yields exhibited relatively parallel pre-trends that rose together during the liquidity crisis in the run up to the first Federal intervention on March 23rd (left black dashed line).[26] Cities and counties continued to track each other throughout the initial MLF announcement (middle gray dashed line) until the announcement of the April 27 cutoffs—the main

[25]See Appendix A.3.4 for further details on these concordances and the construction of plurality ratings.

[26]This is a necessary condition for any regression discontinuity design as the bandwidth approaches 0 in the limit; here we already begin to see observational equivalence (in the dependent variable) with an arbitrary symmetric cutoff of 100,000.

focus of this design.[27] Panel (a) shows investor-perceived measures of MLF access (yields) did not differentially materialize in cities after the third announcement's cutoffs were mandated. While there appears to be mild visual evidence that narrowly ineligible county bonds began trading at lower prices (higher yields) than eligible counties after the third announcement, this difference in not statistically significant in formal regressions. Both plots show full resumption to pre-crisis levels in the long run.

However, turning to panels (c) and (d), the main results of the paper become apparent from simply inspecting the path of these means across the ratings distribution. In these panels and subsequent analyses, we pool the two highest and two lowest IG issuer plurality ratings from January 2020 (fixing issuer ratings at their pre-crisis levels) for statistical power—a choice motivated by the sparse BBB mass shown in panel (b) of Figure 3. While high-rated bonds share the same pattern as the overall market, among lower-rated bonds, a strong wedge begins to emerge exactly when the new cutoffs are made binding, immediately after the third announcement. Importantly, the timing of the yield wedge, exactly in the week the revision was announced, also minimizes concerns that chosen population cutoffs may have been used for other COVID-related policies, which would have appeared differentially in the pre-trend before April 27.

Analyzing the levels of these means, low-rated MLF-*ineligible* bonds (red series) remained significantly distressed beyond the last Federal intervention, in line with the aggregate BBB municipal market trend initially presented in Figure 1. By contrast, low-rated MLF-*eligible* bonds appear to recover fully to pre-crisis levels, mimicking the behavior of high-rated bonds when supported by the MLF.[28] In some weeks, this wedge is as large as 100 bps—roughly half the pre-crisis spread level. The question of how to interpret these differential effects that are only present among low-rated bonds, forms the basis of our formal analysis, and is discussed in Section 4.

The sample used in Figure 4 and in the estimation section applies the following three restrictions: (1) we two-tail winsorize all yields at the 1% level; (2) we drop all yields with likely classical measurement error (prices less than 50, and greater than 150, following Green et al., 2010);[29] (3) we drop any secondary market trades occurring within 90 days of primary issuance following Schwert (2017).[30] After imposing these restrictions, the pooled low-ratings sample (the smallest in our analysis and thus the one

[27]To ensure bin means do not overlap announcements, we mean-collapse all yields that fall within inter-announcement intervals to a single observation centered at each interval's midpoint (inclusive of the Monday of the each announcement).

[28]This pattern also emerges strongly when looking separately at A and BBB categories, as we show in additional results in Figure A.7. We also examine whether selection on composition can explain these results in Figure A.10, which shows that while there is some evidence that revenue bonds are traded disproportionately following the announcement, this pattern is not sustained throughout the series. Furthermore, these patterns are not driven by refinancing: among the bonds trading in the 2020 A and BBB sample within 100,000 of the cutoff, only 4 were redeemed after the April 27 announcement, all of which were advanced refunding deals among MLF-*ineligible* issuers.

[29]Prices at par, for the most part, vary closely around 100 at time of initial issuance.

[30]The results in Figure 4 are nearly identical without these restrictions.

warranting the most power concerns) is comprised of 759 issuers, 7,248 bonds, and 94,849 trades in the post-period; with 4.75% of bonds falling with 100,000 population above the cutoff, and 3.53% within 100,000 below the cutoff.[31]

3.3 Public Sector Employment Data

In order to estimate effects on real-economy outcomes we leverage an additional source, the U.S. Census Bureau Quarterly Census of Employment and Wages (QCEW) that encompass local public sector employment by county-month and sector at eligible and non-eligible localities.

We use the May 26, 2021 revised release of government employment data from QCEW, which includes a monthly stock of all workers that were employed on the 12th day of the month through the end of 2020Q4. The QCEW is based on ES-202 unemployment insurance filings at the state level, required for filing payroll taxes for federal UI-covered workers, so our sample requires a minimum level of part-time work for an employee to be included. This data covers roughly 90% of U.S. employees, excluding self-employed workers, some agricultural workers, and informal workers, and roughly 96% of state and local government workers (BLS-QCEW).[32] The QCEW is unique in its ability to provide a census of employees at a relatively granular sector-by-geography resolution. The most disaggregated geography in this data is at the county level, however, public sector employees are also broken out by level of government: we use *Local Government Employees* (Ownership Title = 3) employed in *Total, All Industries* (NAICS = 10) as our preferred measure, which excludes state and federal employees but aggregates city and county employees at the county level. For roughly 67% of the data, sector-specific components of government employment are readily available, and allow us to separately estimate effects for goods-producing (NAICS = 101) and service-providing (NAICS = 102) public employees, the latter of which is largely comprised of *Education and Health Services* employees (NAICS = 1025).[33]

4 RD Empirical Framework

In this section we use an RD design to estimate the MLF access effects. We later complement this analysis with a spread decomposition approach to support our proposed mechanism hypothesis and generalize

[31] In Appendix A.3.5, we further enumerate the list of 43 A and BBB city and county issuers that fall within 100,000 in population of the cutoff, along with the number of bonds issued for each.

[32] Contractors are also covered by the QCEW (BLS), and reflect about 10% of local public sector employment (the sample studied in this paper), and include goods-producing employees. See Hyman (2018) for an explanation of sample coverage using similar data.

[33] We are able to confirm for this subset, that goods and services sector sub-components add up to aggregate employment totals, and thus consider these reliable. Cells that do not add up are withheld "so as to protect the identifiable information of respondents", and coded to missing in QCEW data as a small cell disclosure protection (`https://www.bls.gov/cew/questions-and-answers.htm` question 9).

to the broader market (Section 5) .

4.1 Regression Discontinuity with MLF-Eligible Populations

We estimate the option value of MLF access using a pooled RD strategy, analyzing effects in two distinct periods: the pre-period (prior to March 23, 2020) and the post-period (after April 27, 2020), corresponding to pre-crisis or "placebo" and post-cutoff revision or "treatment" groups.[34]

$$\begin{aligned} Y_{n(bi)t} &= \alpha + \beta_t * \mathbb{1}(pop \geq cutoff)_i + \gamma_t * (pop - cutoff)_i \\ &+ \delta_t * \mathbb{1}(pop \geq cutoff)_i (pop - cutoff)_i + \mathbf{X}_{bit} + \varepsilon_{n(bi)t} \end{aligned} \tag{1}$$

Here, we track outcome $Y \in$ {Yields, 1(Primary Issuance), 1(Credit Downgrade), Public Sector Employment}. When the outcome is yields, we measure trade n of bond b from issuer i in period t. $1(pop \geq cutoff)$ is an indicator variable that takes a value of 1 if the running variable (population relative to cutoff) is greater than or equal to zero. We use "relative population" as the running variable to be able to stack cities and counties with different cutoffs together in one specification. The γ and δ terms estimate separate polynomial slopes on each side of the cutoff, and capture the parametric relationship between city/county size and the outcome variable. α estimates the cutoff-intercept of the γ term's slope from the left-hand side, and has the convenient interpretation of representing the control group (MLF-ineligible) mean close to the cutoff. β_t is the parameter of interest, and estimates the difference in intercepts between the left polynomial and right polynomial at the cutoff, or the jump in the regression function at the cutoff.[35]

We follow Calonico et al. (2014) and Gray et al. (2020), in using a data-driven optimal bandwidth selection procedure to select asymmetric bandwidths on either side of the cutoff for estimating each polynomial, which subsequently helps determine the parameter of interest.[36] The procedure chooses bandwidth boundaries that minimize the integrated mean-squared error (IMSE) of the regression, which is linearly separable in bias and variance terms that are traded off in the bandwidth selection (smaller bandwidths are less biased but higher variance; larger bandwidths are more precise but also more biased). We use a non-parametric first-order polynomial (hence the linear terms) with an asymmetric triangular kernel that weights observations closer to the cutoff with values closer to 1, and diminishing weights moving away from the cutoff (and weight=0 beyond the bandwidth boundary).[37]

[34] In Figure A.2, we also provide estimates from a dynamic bi-weekly RD specification.

[35] For continuous running variables, this builds on a well known statement that if the average potential outcomes are continuous functions of the running variable at the cutoff, the difference between the limits of the observed outcomes of two groups (as the score converges to the cutoff) is in fact the average treatment effect Hahn et al. (2001).

[36] This method takes one choice parameter—bandwidth—out of the econometrician's hands, thus minimizing susceptibility to publication bias and "p-hacking".

[37] We provide the implied kernel weights from the IMSE-optimal bandwidth selection procedure in Figure A.8. The order of

In practice, the choice of kernel has little sway over regression results, while the choice of bandwidth can be massively influential (Cattaneo et al., 2019a). While a consistent estimate of β_t does not require covariates to be identified (and in fact, should be robust to estimation without covariates), in small samples, the precision of estimates can be improved by their inclusion. We thus include the **X** vector to account for mostly time-invariant bond and issuer characteristics in subsequent regressions, however all results are robust to their exclusion. Lastly, we conservatively cluster standard errors on the running variable (relative population) so that inference is drawn from between-issuer differences rather than within-issuer—consistent with recommendations for a discrete running variable (Frandsen, 2017).

Importantly, we also estimate Equation 1 separately for the four pre-crisis plurality credit ratings types of issuers (AAA, AA, A and BBB), which we interpret as a reduced form test of risk-sharing. As discussed earlier, in our main estimation, we pool the two highest and two lowest IG issuer plurality ratings from January 2020 for statistical power. Our hypothesis is that if the Federal sector directly shares credit-downgrade or credit-default risk, one would expect lower rated eligible and non-eligible issuers to have stronger price spreads due to comprising a higher share of bonds on the default margin. If instead there are neutral effects, this direct access most likely reflects an "added MLF liquidity effect" beyond what is already happening in the aggregate.

4.1.1 Identifying Assumptions and Balance Test

The key identifying assumption of the RD is that potential outcomes are smooth across population relative to MLF eligibility cutoffs, absent the cutoff announcement itself.[38] Any selection on the running variable could violate this by breaking the associated "no manipulation" or "no sorting" requirements, which ensure that the RD is locally as good as random. As shown, this is unlikely to be problematic in our setting. However, similar to randomized controlled trials, this identifying assumption generally has a natural falsification test (should pre-treatment data be available) that helps support the validity of the assumption: estimating the RD on baseline data prior to the experiment itself.

In our context, we have collected many months of a pre-period placebo outcome data, prior to March 23, 2020, in which we can formally test this assumption on our trade-level dataset. Table 1 presents Equation 1 estimates for the pooled city and county sample in the period from Jan 1, 2020 to March 23, 2020, where each row reflects a separate regression in which a new optimal bandwidth is selected depending on the outcome variable, thus producing different observation counts per variable.

Each regression is estimated unconditionally, without adding covariates on the right-hand side. "Discontinuity" and "Control Mean" columns correspond to β and α, respectively, from Equation

polynomial is chosen to best fit the underlying data.

[38] For a discrete running variable, the "local continuity" assumption is softened to "local randomization".

1.[39] Among the dependent variables listed, there are only a few cases in which a significant break is detected in the placebo sample. When they are detected, the control mean suggests that they are not economically meaningful or large in magnitude, and thus simply a result of randomization (sampling) error. Furthermore, estimates on the important yield variables, which include both a month-to-month and year-on-year trend variable, are particularly small relative to their control means.[40] While randomization errors are more likely to occur within smaller samples, such as estimates within lower-rated bins, the combination of balance in the placebo period and relatively parallel pre-trends in dependent variables shown earlier, suggests that the local randomization assumptions hold in this setting.

4.1.2 Differenced RD Modification for Employment Effects

We also desire to estimate the effects of MLF eligibility on city and county public employment outcomes in our RD framework, as local employment should be one of the first margins to respond to interventions if there are real-economy effects of the policies.[41] Recall that MLF eligibility could directly support the retention of employees by allowing the matching of cash in- and outflows. In addition, to the extent that MLF eligibility enhances an issuer's ability to refund outstanding long term debt, this could reduce debt service costs, freeing up cash flow that could be used to provide services and maintain payroll. We face, however, a challenge in addressing this very important issue. The data only allow us to observe local public employment *combined* for city and county employees at the county-month level. That is, while the QCEW provides exceptional coverage of employment and breakdowns by sector, it does not decompose total local public employment into its city and county sub-components—consequently, any effort to estimate RD effects for counties could be made difficult by the possibility that an untreated county's public employment to the left of the county population cutoff is partially covered by an MLF-eligible *city*; symmetrically, an untreated city could be covered by an MLF-eligible county. Both of these cases would attenuate relative employment effects toward 0.

For these reasons, we opt for a strategy comparing combined city and county eligible employees within *counties* on the right of the cutoff, with "neither-eligible" counties to the left of the cutoff (i.e. counties in which neither the county nor any cities it contains are eligible). We propose a novel adaption in the spirit of difference-in-discontinuity designs to difference out any unintended selection concerns resulting from this measurement strategy in which we first pre-sort cities and counties into four cases: (a) the county itself is the only MLF-eligible issuer within the entities spanned by the county; (b) only 1 city

[39] A full set of summary statistics can be found in Table A.4.

[40] For more details on the how the remaining variables are constructed, see Appendix A.3.

[41] While of interest, we consider additional real economic effects such as local job multipliers outside the scope of this paper, which is predominantly concerned with stablization effects.

is MLF-eligible within the county, while the county itself is ineligible; (c) both the county and one of its underlying cities are MLF-eligible; (d) neither the county nor any of its contained cities are MLF-eligible. We show representative maps for cases (a), (b) and (c) in Appendix Figure A.1.

In our preferred strategy to estimate employment effects, we stack all case (a) "county-only" types (spanning 59/119 MLF-eligible counties) with case (c) "double-eligible" types (spanning 60/119 eligible counties) on the right of the cutoff of the *county* running variable, and compare these with case (d) "neither-eligible" types to the left of the cutoff along the same running variable.[42] The modified RD identification assumption is that potential outcomes are smooth around each running variable's *stacked* sample. To the extent that these pre-sorted samples introduce further selection concerns, and to address the potential problem that any RD randomization error could magnify level differences emanating from differently sized geographies around the cutoff in these modified samples, our preferred specifications consider the first-difference change and percent change in employment (year-on-year) rather than its level. Year-on-year changes in employment in a given month should difference out any aggregation errors arising from this sorting procedure, and as county employment levels have a long right tail, focusing on changes also helps reduce noise in these estimates.

In all employment regressions, we depart from our IMSE-optimal bandwidth selection procedure and instead use a fixed bandwidth due to the much smaller number of observations—an N between 700 and 900 (whereas yields RDs are estimated from tens and sometimes hundreds of thousands of observations). We choose a fixed bandwidth of 400,000 below the cutoff, and 600,000 above the cutoff, allowing us to take advantage of a full window size that spans localities of 1 million in population (between 100,000 and 1.1 million). As before, we employ a triangular kernel that weights observations near the cutoff close to 1, and observations at the bandwidth boundary equal to 0. As we did for yields, we two-tail winsorize employment values at the 0.1% level prior to taking first differences. For consistency with yields results, we include state fixed effects in some regressions (and month fixed effects when the outcome is not year-on-year first differenced), however all employment effects are robust to being estimated unconditionally without covariates.

4.2 RD Effects on Secondary Market Yields

Table 2 presents our main results on the investor-perceived option value of access to the MLF, estimated from secondary market municipal bond yields (in basis points).

[42] In Table A.3, we also separately compare case (b) cities (spanning 21/86 eligible cities) against case (d) types along the *city* running variable (which requires constructing city employment by calculating the employment weight of each county spanned by the city). This sample has the appeal that the city MLF eligibility cutoff (250,000) is distinct from the CARES eligibility cutoff (500,000), however power calculations in the appendix show we are unable to detect large effect sizes. Furthermore, estimating effects on the city running variable also requires strong assumptions to assign county employment weights to multiple *untreated* cities within a given county.

Panel (a) pools all trades post-April 27 into one treatment sample, whereas panel (b) pools all pre-March 23 observations into one placebo sample (starting from January 1, 2020). Each row representatives a different sub-sample of issuers and separate regression, as discussed above. As before, the "discontinuity" and "control mean" columns correspond to β and α from Equation 1 respectively. In this regression table, we implement the sampling restrictions for yields discussed in Section 3, and include calendar month fixed effects to absorb time series variation in overall levels, and state fixed effects to reduce variability from differences in exemption rules and budget balance requirements (which can limit downstreaming) across states.[43]

Consistent with the event study plots, we fail to detect a statistically significant differential effect between eligible and ineligible issuers in the post-period overall, nor by city or county. The highest rated issuers appear to also be unaffected by the announcement of MLF eligibility.[44] Moving down the ratings gradient however, we estimate a strong statistically significant yield spread for low-rated issuers. The point estimate suggests that narrowly MLF-eligible issuer bonds trade, on average, 75 bps lower (i.e. at a higher price) than observationally equivalent issuers that narrowly miss the cutoff. This wedge represents about 24% to 28% of baseline yields (depending on whether one benchmarks to the pre-crisis (263.29) or post-crisis (307.69) period respectively), a magnitude that closely corresponds to the overall municipal market wedge between BBB yields and higher-rated issuers as shown initially in Figure 1. In fact, when estimated separately by BBB and A, effects among BBB-rated issuers are even stronger than this average (see additional results in Figure A.7).

It is worth highlighting that including state fixed effects in our main specification also implicitly tax-adjusts yield estimates—since the above table is restricted to 2020, state fixed effects would account for any differential state tax rates and exemption rules on muni income accrued to the marginal investor residing in each state, who is presumed to take advantage of doubly-exempt bonds (i.e. muni interest income that is exempt from federal *and* state tax liabilities). In Section 6, we report the sensitivity of our RD results to directly adjusting tax-exempt bonds using NBER-TAXSIM effective rates for a representative high-wealth household investor in each state, as well as excluding state and month fixed effects, all of which show our results are qualitatively robust to excluding all covariates and do not rely on any specific tax adjustment method.[45]

Encouragingly, for all subgroups, estimates in the pre-period are close to zero. This suggests that

[43] State fixed effects are estimated using all data within the optimal bandwidth, spanning close to 759 issuers (the total number in the combined A and BBB group.)

[44] In regressions by ratings bin, we pool cities and counties together for statistical power.

[45] NBER-TAXSIM provides these rates net of local cross-deductibility rules, including mortgage interest deductions. The latest rates provided are for 2018, which we assume are the effective rates salient in 2020. See 2018 "Wages Federal Rate" for the effective rate applied to federal-only exempt bonds, and "Wages Total Rate" for the sum of federal and state effective tax rates applied to each state's underlying bonds (`https://users.nber.org/~taxsim/state-rates/maxrate.html`).

balance is achieved even among these smaller sample sizes using only a parsimonious number of standard controls (state and month fixed effects). To explore the precision and sensitivity of our preferred estimates transparently, Figure 5 shows RD scatter plots associated with high- and low-rated issuers.

Here, we begin to see a pattern of which issuer types investors appear most responsive to across the relative population distribution (the running variable). Among high-rated bonds, yields are neutral across city (and county) size distributions. This neutrality is preserved for low-rated issuers that are MLF-eligible (flat slope), yet for the ineligible issuers, an upward-sloping relationship appears in which larger locations have higher yields, but only up to the cutoff. These estimates remain nearly identical when controlling for bond type (GO versus revenue bond), tenor length, remaining bond duration, as well as trade week, day of week, maturity size, and amount outstanding.[46] To the extent that there are concerns that the RD is picking up underlying volatility beyond these controls, we note that RD polynomial terms on either side of the cutoff account for any remaining differential volatility by relative population. This is further confirmed by placebo RD plots in Section 6 which do not show any strong pre-existing pattern in volatility.

In Table 3, we further show sensitivity of our main result (column 2) to changing both the bandwidth selection method, the kernel used, and the order of polynomial. In all specifications, we find a large negative yield effect for low-rated bonds, with placebo effects which are not statistically distinguishable from zero. Adding a second-order polynomial does little to our preferred specification, consistent with a linear fit being sufficient, as suggested visually in Figure 5. To address whether our results are influenced by trades further away from the cutoff (which would load more weight on the polynomial choice), in Figure 6 we conduct permutation tests in which we reestimate the RD at wider fixed (symmetric) bandwidths away from the cutoff. The results from these tests are remarkably stable as one moves further away from the cutoff, and are similar in magnitude to our main coefficient (75.43) no matter what kernel is used. Together, these additional tests suggest that observations further away from the cutoff are less influential. Even using only 20,000 observations (almost half of the 37,977 in our main sample)—corresponding to a symmetric bandwidth of 40,000 in population around the cutoff—we find substantial impacts on low-rated IG yields.[47]

We interpret these robust results as suggestive of a setting in which pre-existing uncertainty in the pricing environment for low-rated municipal issuers (those with more bonds on the margin of default) may have been amplified by the pandemic. This credit risk seems to have been priced differentially by investors, depending on whether issuers had the ability to borrow from the Federal sector. In the

[46]See sensitivity table shown in Section 6.

[47]In additional falsification tests in which we randomly vary the cutoff away from the "true" cutoff, among all possible permutations across the support, only the estimate at the true cutoff is statistically different from zero.

following section, we then ask whether such perceptions also impacted new debt issuance.

4.3 RD Effects on Primary Issuance

We start our examination of whether investor perceptions, which created a wedge in secondary market yields, also induced additional new primary issuance of munis differentially for MLF-eligible issuers by showing graphically in Figure 7 the cumulative stock of bonds that were newly issued (defined here as having a first offering during calendar year 2020), by each week of the 2020 calendar year.[48]

Panel (a) shows that issuers of all ratings (overall mean) within the symmetric bandwidth above and below the pooled cutoff shared roughly equivalent trends in new issuance, with a slightly higher level among lower population issuers due to the greater density of cities and counties of smaller populations versus larger populations. While overall issuance does appear to converge later on in September (where the two series intersect), this is not sustained as a strong pattern in the long run. By contrast, turning to low-rated issuers whose bonds we confirmed were trading at higher prices after the MLF cutoffs were revised in late April, we see a different pattern. First, the baseline level of new issuance in panel (b) is close to zero from January to July of 2020 in both samples due to the high inherent distress levels of these issuers, with the exception of one issuer above the cutoff that issued a series of differently-maturing bonds prior to the pandemic. In July, we observe bonds narrowly above the cutoff beginning to issue sooner, followed by a step function pattern of new issuance throughout 2020. The deferred timing of this issuance (in July) relative to the immediately binding effects on yields, may be linked to the beginning of fiscal budget cycles, or coincide with the deferred IRS tax filing deadline of July 15 which may have revealed a clearer picture of municipal revenue positions that put new issuance into motion.[49]

Importantly, only 5 low-rated city and county governments (within the 100k bandwidth) issued new bonds after the MLF cutoffs were announced, 4 issuers (on a denominator of 23) above the cutoff and 1 issuer (on a denominator of 20 issuers) below the cutoff (see Appendix A.3.5 for denominators). With the caveat that this does not provide bountiful variation to pick out a small signal from the noise, we now show the formal RD estimation associated with the probability that a bond was *ever* newly issued after (and before) the MLF cutoffs were announced, as well as the intensive margin amount of issuance. To do this we collapse each CUSIP (bond or note) into a single post- and pre-MLF period observation, and assign each bond a value of 1 if it was newly issued during that period, and 0 otherwise. Here, we depart from optimal bandwidths due to the limited size of the collapsed dataset, and show results without

[48]We choose a cumulative approach due to the lumpy nature and sparsity of primary issuance. Later, we also formally test the probability that a bond is newly issued separately in the pre- and post-MLF periods.

[49]To rule out seasonal mean reversion in primary issuance as a key driver of these results, we also report the distribution of 2019 new issuance as a benchmark expansion year in Figure A.9, which does not provide any evidence that seasonal mean reversion is a key factor.

covariates for transparency. Each regression uses a fixed bandwidth of 400,000 below and 600,000 above the cutoff, allowing us to take advantage of a full window size that spans localities of 1 million in population (between 100,000 and 1.1 million), and a triangular kernel to weight observations close to the cutoff with weight=1, and closer to the bandwidth boundary with weight=0.

Table 4 shows that the probability a bond is issued increases by 8 percentage points (on a baseline share of 10 percent), which is largely driven by an increase of 22 percentage points among low-rated bonds. Due to the sparse nature of primary issuance, we do not have ample variation to disentangle these effects completely from differential issuance in the placebo period—indeed, in panel (b) larger municipalities are 7 percentage points more likely to issue in the pre-period relative to marginally smaller ones. Thus, rather than interpreting these point estimates directly, we take these along with the larger coefficients on total issuance for lower rated bonds as a qualitative evidence that having MLF optionality encouraged primary issuance, and investors likely priced this ability to issue in secondary markets, especially among low rated bonds, as shown previously.

4.4 Differenced RD Effects on Public Sector Employment

We next turn to effects on local public sector employment using the differenced RD modification strategy previously discussed. In Table 5, we stack all case (a) "county-only" and case (c) "double-eligible" types (spanning 60/119 eligible counties) on the right of the cutoff of the *county* running variable, and compare these with case (d) "neither-eligible" types to the left of the cutoff along the same running variable. Again, each row represents a separate RD regression. As discussed, unlike our yields analysis, our preferred outcome variable in this context is differenced year-on-year such that post-period results reflect pooled annual changes in county local government employment for May and June (between 2019 and 2020), while the placebo shows the same effects for differenced January and February pooled (columns (3) and (4)). As described previously, we estimate effects separately for goods- and service-providing public sector employees, although we can only do so for the 67% of the sample in which these sector breakdowns are reliably available, which results in different observation counts.

We begin our interpretation by analyzing employment effects in levels in columns (1) and (2). While we fail to detect statistically significant MLF effects in levels, the large standard error underscores that employment was highly volatile in the pre-crisis (placebo) period. Importantly, the control mean in these columns (marginally-ineligible counties) indicates our sample counties initially had roughly 18,835 local government employees on average—heavily skewed to the right by large counties.[50] The small effect sizes contrasted with large baseline means also suggest that any effects on levels are likely

[50]Baseline means vary across models because the "control group" is the intercept of the RD polynomial from the left-hand side of the cutoff, which may vary when a control adjusts the optimal bandwidth over which the polynomial is projected.

estimated with noise. Given this, in conjunction with our main selection concern that subsetting counties to isolate precise employment effects may generate unintended imbalances in baseline characteristics, we use year-on-year differences in our preferred specification in column (3). Here, the unconditional control group shows counties experienced an average annual decline of roughly 1,717 local public sector employees in May and June of 2020 (relative to the same months in 2019).

While effects on overall employment are not significant at conventional levels, the point estimates for service-providing public employment sectors in columns (3) and (4) suggest MLF counties with at least one MLF-eligible entity increased local government service employment by 422 to 517 employees relative to the control group. Interpreting estimates with respect to the control mean in column (3), this amounts to an average employment decline of 1,157 to 1,252 employees relative to a baseline decline of 1,674 employees for ineligible counties—an effect size of 25.2% to 30.9%, indicating a striking degree of employment smoothing in response to emergency liquidity.[51]

Figure 8 show RD scatter plots associated with our preferred estimates of employment changes (column (4) of Table 5), both overall (Panel a) and for service-providing (Panel b) sectors.[52] These plots first reveal a largely linear and negative employment-county size gradient for local public employees—the larger the locality, the larger the year-on-year level decline in employment. Interestingly, the same plots in the placebo period produce the opposite slope—the largest counties also had the largest growth in public employment in the run up to the pandemic, albeit with a milder gradient (see placebo effects in Section 6). While noisy due to the sparseness of the QCEW data, the plots reveal a modest upward level shift in *overall* employment when MLF liquidity is available, and a more pronounced statistically significant shift for service-providing sectors in panel (b), which is mostly comprised of employees at educational institutions.

4.4.1 Interpreting Magnitude of Employment Effects

To this point, we have interpreted the large May and June 2020 public sector employment effects after the April 27 MLF eligibility cutoffs (250,000 for cities, 500,000 for counties) as fully attributable to the MLF emergency lending option. However, as noted previously in Section 2.4, the CARES Act announced on March 27 also provided for *more direct access* to aid for cities and counties with populations over 500,000, whereas those under this cutoff had to rely on their overlying states for downstreamed aid, which could be lagged and politically less certain. This raises the possibility that earlier estimates capture a combination of both direct aid and lending eligibility through one of the two cutoffs (i.e. 500,000),

[51] In percentages, column (6) shows a 1.69 percentage point gain in employees on a baseline loss of 7.96%; a similar effect size of 21.2%.

[52] RD plots associated with unconditional effects from column (3) of Table 5 are nearly identical.

though both policies reflect similar liquidity interventions and thus likely share the same theoretical sign on employment and underlying economic interest. To begin to asses whether there may be separate contributions from MLF optionality and direct CARES Act aid, in Figure 9 we provide dynamic monthly year-on-year estimates similar to column (4) from Table 5, to learn from the differential timing of CARES and MLF cutoff announcements.

Plotting RD estimates dynamically presents an important observation on April 12—employees retained several weeks after the CARES Coronavirus Relief Fund (CRF) provisions were enacted, but also several weeks *before* MLF population cutoffs were revised. In the data, all *direct* CRF aid was fully dispersed (and assumed to be salient to and anticipated by budget officers) between April 15 and May 6. We thus interpret the departure from zero that we detect on April 12 as the employment effect of anticipated (direct) CARES aid for larger localities, whereas May and June employment effects reflect the totality of any current and past realized aid as well as the MLF option. The similar point estimates between April and May provide initial suggestive evidence that CARES aid may have been material in affecting hiring and retention decisions, however, unobserved downstreamed aid may also have become fully dispersed by May or June, equalizing aid on both sides of the shared 500,000 population cutoff, leaving the MLF as the sole source of employment differences in May and beyond. There is also evidence that some localities imposed deadlines for applying to downstreamed aid.[53] Lastly, the dynamics in Figure 9 also reveal an interesting pattern related to the school-year calendar, a feature we return to when interpreting these effects. It also bears mentioning that monthly estimates in placebo months (January to March) are also reassuringly concentrated around zero prior to the onset of the pandemic in the United States.

To highlight the difficulty in decomposing total employment effects into those emanating from CARES versus MLF more formally in Figure 10, we compare variation in MLF lending caps, calculated as 20% of each government's 2017 "own-source general and utility revenue" (OSGUR) as per MLF regulations,[54] and the amount of aid either apportioned directly to each local government through the CARES Act based on the CRF population formula (right of the cutoff), or received as indirect downstreamed aid to county governments, hand-assembled (incompletely) by the National Association of Counties (NACO Aid). We unfortunately are unable to locate any data on downstreamed aid to cities, which are therefore omitted from the plot.[55]

[53]For example, New Jersey had a Novemeber 2020 deadline for downstreamed aid applications (New Jersey Local Government Emergency Fund).

[54]See April 9, 2020, MLF Term Sheet for further details on MLF lending cap calculation.

[55]According to the CARES Act CRF Allocation Formula, no population can be double-covered by both its underlying city and county for aid, however, each entity above the cutoff can independently claim direct aid on behalf of its constituents. We use direct CARES aid data received by either cities or counties from USASpending.gov (the US Treasury's official tracker of fiscal expenditures). Nevertheless, this analysis requires total county aid, which is comprised of direct county aid in addition to

Of note, the two different scales of the y-axes in Figure 10 underscore that counties at the 500,000 cutoff were eligible for both an outsized amount of *potential* lending, and a much smaller but still material amount of more direct aid, making comparisons difficult. Downstreamed data is largely incomplete on the right, and there is nothing precluding downstreamed aid to fully close the gap to the left of the cutoff were full data available. In fact, it can be seen by extending the slope to the left of the cutoff, that some states followed the same formula for smaller localities even though this was not legally binding.[56], [57]

In Table A.3, we also separately compare case (b) cities (spanning 21/86 eligible cities) against case (d) types along the *city* running variable (which requires constructing city employment by calculating the employment weight of each county spanned by the city). This sample has the added appeal that the city MLF eligibility cutoff (250,000) is distinct from the CARES eligibility cutoff (500,000), but power calculations in the Appendix show we are unable to detect large effect sizes. Furthermore, estimating effects on the city running variable also requires strong assumptions to assign county employment weights to multiple *untreated* cities within a given county. We do, however, unambiguously find effects on yields when looking at either cutoff independently, confirming that indeed the effects we find on financial markets appear to be driven by the MLF (see Figure A.3).

We conclude from this exploration that while there are unambiguously large employment effects from the liquidity package introduced by the CARES Act through its lending (MLF) and direct aid components, some degree of the sizable effects may be due to the coupling of these two policies. While the exact extent to which CARES versus MLF (and their interaction) contributed is left to future work, municipal liquidity interventions appear to have generated substantial recovery in labor markets.

4.4.2 Employment Effects by Credit Ratings and Calendar Month

One interesting pattern that arose in Figure 9 was that the employment response to the totality of direct CARES aid and MLF lending eligibility was only positive in non-summer months. We further examine this result, adding heterogeneity by January 2020 plurality ratings in Figure 11.

When examining effects separately by ratings, we first confirm the general pattern of only detecting

direct aid paid to a county's underlying cities. Toward this end, we split the direct city aid across counties based on the county CRF cap were it to have accessed the aid directly. Because aid received is calculated with the "latest" population vintage, we estimate the county-level allocation with minor error.

[56]Within states, any jurisdiction with population greater than 500,000 received *direct* US Treasury access to the population share of the state allocation * 45%, whereas the remaining 55% of was controlled by the state. By contrast, localities with populations *under* 500,000 had to rely on downstreaming from their underlying states (who in this case controlled 100% of the local allocation), leaving aid to smaller localities potentially more politicized, less certain, or slower to materialize.

[57]While the kinked nature of the CARES population formula appears at first to provide a potential test in which a regression kink design (RKD) may help discern between these two policies, the underlying kink is in fact a flat schedule in the variation of interest—aid per person. That is, when both the outcome variable and treatment variable are highly correlated with the running variable (as it is here), there is in fact no kink in the denominator since the aid allocation formula is based on population itself, and therefore no meaningful variation in aid to leverage.

positive effects from April to June, and October to December. This school year pattern is broadly consistent with the notion that emergency liquidity—whether aid or a lending option—supported educational institution retention, which is particularly striking given that the majority of schools were still under lockdown from April to June. Interestingly, Figure 11 also reveals that these employment returns were only sustained in the long run when the underlying government was a highly-rated issuer (which as we showed earlier, reflects the majority of municipal bond issuers).

If indeed it is the case that the least fiscally constrained issuers were most responsive to the additional liquidity, both during the months of April to June when government shutdowns were still widely in place, and through the close of the 2020 calendar year, this is consistent with the hypothesis that state and local governments may have over-weighted the worst possible outcomes based on past crisis experience, furloughing education sector employees even though realized revenue shortfalls were far lower than originally anticipated (this view was also articulated in Sheiner, 2021).

5 Pricing Credit Risk or Liquidity?

So far we showed that MLF access impacted investor-perceived yields, but only at the low end of the credit ratings distribution. This prompts the question, what exactly are investors pricing? To investigate the mechanisms underlying these yields, we perform two additional exercises. First, in Appendix A.4.1, we estimate the effect of MLF access on the probability an NRSRO downgrades a bond (a lagged but very salient type of observable credit risk), and find modest evidence that MLF-eligible issuers had lower downgrade probabilities than observationally equivalent ineligible issuers. While this is consistent with yields predominantly pricing credit risk, in what follows we implement a more formal decomposition test to separate yields changes arising from illiquidity versus credit risk.

The infrequent trading of municipal bonds and a potentially large number of unexecutable quotes pose a challenge when measuring liquidity and credit risk premia at a high frequency for many issuers. We thus follow the transaction-based decomposition in Schwert (2017), applying the methodology at a *weekly* frequency rather than monthly. We first estimate the liquidity component of yields directly, then interpret the predicted residual as compensation for credit default risk. For bond b issued by issuer i in state s with yield y_{bist}, we begin by computing the duration-matched, tax-adjusted spread for each bond, as:

$$\hat{y}_{bist} = \begin{cases} \left(1 - \tau_f^{deduct} - \tau_s^{deduct}\right)^{-1} y_{bist} - y_{bist}^{UST} & \text{if federal and state exempt} \\ \left(1 - \tau_f^{deduct}\right)^{-1} y_{bist} - y_{bist}^{UST} & \text{if federally exempt, but state taxable} \\ y_{bist} - y_{bist}^{UST} & \text{if taxable} \end{cases} \quad (2)$$

where y_{bist}^{UST} is the yield-to-maturity on the bond implied by the risk-free Treasury curve at the time of transaction minus an interpolated maturity-matched swap rate following Gürkaynak et al. (2007); τ_f^{deduct} and τ_s^{deduct} are the annual federal and state tax rates for a representative wealthy household taken from the NBER-TAXSIM model, which also adjusts for local and state deductions and assumes the marginal investor buys bonds within their resident state to take advantage of state tax exemptions.[58]

Next, we directly estimate a weekly liquidity measure for each bond as the first principal component λ_{bt} of a set of bond-week liquidity measures that include: Amihud price impact, Imputed round-trip cost, Roll measure, and price dispersion.[59] Appendix A.4.2 provides details on the calculation of these liquidity measures, and shows liquidity deterioration for both AAA/AA and A/BBB city and county bonds in March 2020, although the liquidity of the lower-rated bonds declined more deeply and took longer to recover to its pre-pandemic level (if at all, depending on the illiquidity measure).[60] To quantify the relative contribution of illiquidity to yields in the build up to the crisis, we estimate the sensitivity of spreads to liquidity β_t^{λ} by regressing the duration-matched, tax-adjusted yield on the liquidity measure, conditional on observed ratings:

$$\hat{y}_{bist} = \alpha_t + \beta_t^{\lambda} \lambda_{bt} + \delta_t Ratings_{bt} + \varepsilon_{bist}, \quad (3)$$

where $Ratings_{bt}$ contains a series of bond ratings dummy variables (AAA, AA, A, BBB, BB and below, NR), meant to capture observable credit risk as represented by NRSRO ratings, which may be correlated with liquidity.[61] The liquidity spread component of yields (reflecting the *illiquidity* of a bond) is then

[58] This is similar to using a state-deducted formula such as $1/(1-\tau_f)(1-\tau s)$, which does not adjust for local deductions (see Schwert, 2017, for an example). We use the NBER-TAXSIM annual "Total Rate, Wages" for bonds that are exempt at both federal and state levels (in Mergent, *tax_code=EXMP* and *state_tax=N or missing*), and the "Federal Rate, Wages" for bonds exempt at the federal-level but not at the state-level (in Mergent, *state_tax=Y*), taken from `https://users.nber.org/~taxsim/state-rates/maxrate.html` (and apply 2018 tax rates, the final year of the data, to 2019 and 2020 bonds).

[59] Appendix A.4.2 includes a specification where we calculate the liquidity factor including the effective bid-ask spread in addition to these four liquidity measures, and the results are qualitatively unchanged. Our choice to not include the effective bid-ask spread in the main specification is related to the sample restrictions that are required for the bid-ask spread calculation, and are broadly discussed in the appendix.

[60] A broader sample of bonds that include *state* bonds also exhibits a similar liquidity deterioration, which supports the generalizability of our city and county findings to the broader municipal bond market.

[61] Unlike in the previous section, the decomposition here requires including time-varying ratings rather than fixed issuer ratings, as time-varying ratings capture the sticky nature of credit rating revisions. Toward this end, we assign each bond the minimum ratings among the three NRSROs, to be consistent with the MLF Term Sheet.

calculated as

$$l_{bt} \quad = \quad \beta_t^{\lambda} \left(\lambda_{bt} - \lambda_{p01} \right) \tag{4}$$

where λ_{p01} is the first percentile of the liquidity measure's first principal component, thus benchmarking to a very liquid municipal bond within a given week t. The default spread—the yield compensation for credit default risk—is then calculated as the remaining portion of the yield spread after adjusting for liquidity

$$\gamma_{bist} \quad \equiv \quad \hat{y}_{bist} - \hat{l}_{bt} \tag{5}$$

which accounts for both observed credit risk (through ratings) and unobserved credit risk. This allows us to generate the following decomposition for city and county bonds by January 2020 plurality ratings group, plotting weekly means of default risk (γ_{bist}) and illiquidity ($\hat{l}_{bt}$) components, in Figure 12.

Panels (a) and (b) of Figure 12 demonstrate that while prior to the crisis, both observable illiquidity measures and residual credit-default risk appeared to contribute meaningfully to yield spreads, the rise in city/county yields for both high and low-rated issuers seems to have been mostly driven by the price of default risk, with a higher contribution of credit risk to long run yields distress for low-rated issuers—evinced by both higher underlying levels and higher contributions to adjusted yields. While our estimates of the default component are upper bounds insofar as observed liquidity measures are imperfect or incomplete, the large credit-default residual suggests that investor perceptions of right-tail risk may have changed as the crisis deepened. As market functioning returned to normal, uncertainty became more resolved and observed credit ratings may also have had more time to adjust. However, consistent with market yields trends, credit default risk remained elevated for A and BBB bonds, whereas the default risk of safer AAA and AA bonds reverted back to their mean more quickly. To understand how greater illiquidity during the crisis did not result in a higher decomposition coefficient on liquidity during this same period, we show in panels (c) and (d), weekly percent changes in both the liquidity first principal component and adjusted yields. These panels clearly show that while illiquidity rose dramatically during this period relative to pre-trend (consistent with Appendix A.4.2), adjusted yields grew even faster on a larger base, resulting in a somewhat stable coefficient on liquidity. The bottom line is that the decomposition fails to attribute the large rise in yields to illiquidity, and instead assigns the lion's share of weight to default risk.

The large magnitude on credit risk is consistent with Schwert (2017), who found that credit-default risk accounted for between 74% and 84% of mean spreads using a sample mostly reflecting expansion years, suggesting that there may be a higher degree of underlying credit risk in the municipal bond market than is commonly appreciated. Here we show that this was especially true for low-rated issuers in the run up to peak distress during the pandemic. The high measured price of credit risk in the municipal bond market can be due to high systematic risk, or a consequence of tax-treatment segmentation and the inability of local retail investors at the margin to diversify idiosyncratic municipal risk. To more precisely quantify our own estimates and compare with Schwert (2017), we report summary statistics at the bond-week level in Table 6 for the transaction-based bond spread decomposition.

In our preferred (unweighted) mean specification, we find that in the pre-crisis period, default risk accounted for about 73.2% and 79.4% of yields for high- and low-rated bonds respectively. During the crisis, these rose to 89.1% and 88.5% respectively, and remained elevated post-intervention, especially among low-rated bonds (86.3% versus 82.4%). We interpret the persistently higher default risk share among lower rated bonds as a potential explanatory factor in the slower recovery of lower rated bonds in the overall market, and as suggestive evidence for why MLF seems to have an outsized impact on lower rated bonds (whose credit risk was particularly elevated during this period).

6 Robustness and Sensitivity

In Appendix A.2., we show a battery of robustness and sensitivity tests. For the main RD results, these include (but are not limited to) dynamic specifications over event time which also compare RD with difference-in-differences (Figure A.2), robustness of yields results to a potentially confounding CARES Act cutoff (Figure A.3), sensitivity to the inclusion of controls (Table A.1), pre-MLF placebo RD plots (Figure A.4 and Figure A.5), formal manipulation tests (Figure A.6), and several checks for composition changes in trading over time. In Appendix A.4., we also show additional results on credit downgrades, as well as full details underlying our yields decomposition.

7 Conclusion

The municipal bond market serves a key role in assisting municipalities to maintain smooth spending paths, both in normal times when revenues and expenditures are misaligned, and in times of stress, when a recession or natural disaster hits. The onset of the COVID-19 pandemic, which simultaneously shocked incomes and resulted in a liquidity crisis in the municipal bond market, and the subsequent introduction of Federal sector facilities, present an opportunity to estimate the option value of municipal liquidity and its impact on bond market activity and public sector hiring decisions.

We show that while overall secondary market yields and primary issuance for the most part returned to normal market functioning as a result of the totality of Federal interventions that were introduced between mid-March and the end of April 2020, low-rated IG bonds (A and BBB) remained relatively distressed. Using an RD design that exploits lending eligibility cutoffs introduced by the Federal sector's MLF on April 27, 2020, however, we find that low-rated IG government issuers' bonds traded at higher prices with facility access. Yields among low-rated IG issuers that were narrowly eligible for emergency lending immediately declined roughly 75 bps relative to yields of issuers that narrowly missed eligibility—a magnitude that closely corresponds to the overall municipal market spread between BBB yields and higher-rated issuers. We find modest evidence that eligible issuers experienced lower credit downgrade probabilities overall, and stronger evidence that MLF optionality translated to differential primary issuance on private markets. Using a complementary asset price decomposition approach, we quantify the credit-risk channel over a broader set of issuers, which implies a non-trivial role for credit risk in determining yield spreads during crises.

In contrast to investors, we show that local governments responded to the totality of MLF optionality and direct CARES aid by retaining 25% to 30% more service-providing public sector employees (reflecting educational institutions) in light of mass furloughs and separations that were prevalent during this period. These employment smoothing improvements are detected predominantly among high-rated issuers (the majority of governments) and were strong even when schools remained shut down. This is consistent with a viewpoint that state and local governments may have over-weighted the worst possible outcomes (based on past experience) in furloughing education sector employees given that realized revenue shortfalls were far lower than originally anticipated (Sheiner, 2021). When the additional liquidity through Federal interventions became available as potential buffer for the locality's revenue position, this may have counteracted some of the uncertainty that induced the initial furloughs. Taken together, our results imply that municipal debt market and employment outcomes would likely have been worse absent the MLF facility's operation and CARES aid provisions. The combination of a "primary-only" MLF intervention and low direct take-up, also suggest that the "option" alone may have sent a sufficiently strong signal that complementary secondary market interventions, such as those implemented in the corporate bond market, may not always be warranted.

Interestingly, the results also suggest the presence of a potential credit-risk sharing channel on top of the Fed's role as buyer of last resort. This brings up a number of questions that future research will need to address. Is emergency liquidity only valuable at the low end of the ratings distribution because in response to adverse shocks, extra liquidity provides implicit insurance against credit risk? If so valuable, why did only two issuers take up direct facility lending?[62] We provide two candidate explanations here: the first relates to the combination of the MLF's restricted eligibility and the design of the penalty pricing grid (shown in Appendix A.3.2). While part of the facility's rationale was that issuers simply could not find markets to match their target issuance due to a general liquidity freeze, a more nuanced view is that eligible issuers could in fact find matches, especially after the market began to normalize, but only used the MLF when its prices were more favorable than those offered through a competitive bidding process. It would, for example, be in an eligible issuer's interest to access the MLF if the issuer's rating was substantially higher than the bond being issued, as the facility penalty prices are at the *issuer* rather than *issuance* level. Indeed, on some occasions eligible issuers appear to have conducted just such a comparison. In one of those cases, the state of New Jersey decided to issue publicly after considering MLF rates,[63] while in another, MTA chose MLF.[64] A second explanation relates to stigma in borrowing, which may have deterred issuers as suggested by Armantier et al. (2015) and Moore (2017) who studied the Fed's Discount Window and Term Auction Facility, respectively.

A second question is if the facility provided short-term assistance to new primary issuance, how did this support pass through to long-term secondary market yields? We begin the discussion by noting that sufficient liquidity is necessary for municipalities to meet spending obligations like payroll, debt service, and other fixed payments, as well as unforeseen circumstances.[65] Municipality access to short-term borrowing can allow such payments, including those related to long term debt, to be made.[66] Fiscal rules or bond covenants may prohibit the proceeds of new debt issuance to be used for debt repayment, but because cash is fungible, access to liquidity can be used to fund non-debt expenses, like payroll, thereby freeing up resources to be used for debt service. Finally, by avoiding unnecessary public employee layoffs or tax increases, short-term financing may help to preserve the local tax base from which the funds for long term bond repayment are drawn. Each, or all, of these mechanisms is a potential connection between the kind of short-term assistance offered by MLF and the performance of long-term debt.

Given our new estimate of the willingness to credit-risk share, a parameter which enters a broader welfare calculus regarding the efficiency of macroeconomic stabilization policies, and sizable employment effects, one natural question is when does municipal credit-risk sharing represent a socially

[62]Illinois has issued via the MLF twice, for a total $3.2 billion in loans (of its $9.67b MLF state borrowing cap), while the New York Metropolitan Transit Authority (MTA) has also issued twice, totaling roughly $3.35 billion, equal to its cap.

[63]https://www.reuters.com/article/usa-new-jersey-fed-bonds/new-jersey-picks-muni-market-over-fed-for-4-billion-bond-sale-idUSL1N2HD1ZD.

[64]https://www.spglobal.com/marketintelligence/en/news-insights/latest-news-headlines/new-york-s-mta-sells-over-450m-in-debt-to-fed-s-municipal-liquidity-facility-59983543.

[65]Indeed, cash on hand is a major factor in determining bond ratings. See Moody's July 2020, pages 9-12 (Moody's Investors Service, 2020).

[66]In addition to making payments directly to bondholders, municipalities may be required by bond covenants to maintain a liquid sinking fund sufficient to cover several months of debt service.

efficient policy improvement? From a financial frictions viewpoint, several influential papers have suggested that sub-optimal risk sharing could arise if institutional investors are overly concentrated in locally exempt bonds—a home market bias that distorts the efficient spread of risk (Poterba, 1989, Pirinsky and Wang, 2011, Schwert, 2017, Babina et al., 2021). Other market structure constraints may also result in inefficient muni market pricing (Garrett, 2020). If financial frictions pass through to labor market frictions or revenue sources that have positive fiscal multipliers or high public service externalities (high social value), these too, enter the welfare calculus. When such frictions are also heterogeneously correlated with the mean income of underlying geographies, inequality weights must also be taken into consideration. Due to these complexities and the many models that could rationalize such trade-offs, we leave to future work analysis of the social efficiency of such a policy with so many diverse potential mechanisms, as well as questions related to what issuance would have looked like were the penalty pricing schedule steeper down the ratings distribution, or less punitive overall.

References

Adelino, M., C. Cheong, J. Choi, and J. Y. J. Oh (2021): "Mutual fund flows and capital supply in municipal financing," *Available at SSRN 3966774*.

Adelino, M., I. Cunha, and M. A. Ferreira (2017): "The economic effects of public financing: Evidence from municipal bond ratings recalibration," *The Review of Financial Studies*, 30, 3223–3268.

Amihud, Y. (2002): "Illiquidity and stock returns: cross-section and time-series effects," *Journal of financial markets*, 5, 31–56.

Ang, A., V. Bhansali, and Y. Xing (2014): "The muni bond spread: Credit, liquidity, and tax," *Columbia Business School Research Paper*.

Ang, A., R. C. Green, F. A. Longstaff, and Y. Xing (2017): "Advance refundings of municipal bonds," *The Journal of Finance*, 72, 1645–1682.

Appleson, J., A. F. Haughwout, E. Parsons, et al. (2012): "The untold story of municipal bond defaults," Tech. rep., Federal Reserve Bank of New York. https://libertystreeteconomics.newyorkfed.org/2012/08/the-untold-story-of-municipal-bond-defaults.html.

Armantier, O., E. Ghysels, A. Sarkar, and J. Shrader (2015): "Discount window stigma during the 2007–2008 financial crisis," *Journal of Financial Economics*, 118, 317–335.

Auerbach, A. J., W. G. Gale, B. Lutz, and L. Sheiner (2020): "Fiscal Effects of COVID-19," *Brookings Papers on Economic Activity, Brookings Institution, Washington, DC*.

Babina, T., C. Jotikasthira, C. Lundblad, and T. Ramadorai (2021): "Heterogeneous taxes and limited risk sharing: Evidence from municipal bonds," *The review of financial studies*, 34, 509–568.

Bi, H. and B. Marsh (2020): "Flight to Liquidity or Safety? Recent Evidence from the Municipal Bond Market," *Working Paper (November 24, 2020)*.

Bordo, M. D. and J. V. Duca (2021): "How the New Fed Municipal Bond Facility Capped Muni-Treasury Yield Spreads in the COVID-19 Recession," .

Boyarchenko, N., A. Kovner, and O. Shachar (2022): "It's what you say and what you buy: A holistic evaluation of the corporate credit facilities," *Journal of Financial Economics*, 144, 695–731.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014): "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 82, 2295–2326.

Cattaneo, M. D., N. Idrobo, and R. Titiunik (2019a): *A practical introduction to regression discontinuity designs: Foundations*, Cambridge University Press.

Cattaneo, M. D., M. Jansson, and X. Ma (2018): "Manipulation testing based on density discontinuity," *The Stata Journal*, 18, 234–261.

Cattaneo, M. D., R. Titiunik, and G. Vazquez-Bare (2019b): "Power calculations for regression-discontinuity designs," *The Stata Journal*, 19, 210–245.

Chen, H., L. Cohen, and W. Liu (2022): "Calling all issuers: The market for debt monitoring," Tech. rep., National Bureau of Economic Research.

Chodorow-Reich, G. (2019): "Geographic cross-sectional fiscal spending multipliers: What have we learned?" *American Economic Journal: Economic Policy*, 11, 1–34.

CIPRIANI, M., A. F. HAUGHWOUT, B. HYMAN, A. KOVNER, G. LA SPADA, M. LIEBER, S. NEE, ET AL. (2020a): "Municipal Debt Markets and the COVID-19 Pandemic," Tech. rep., Federal Reserve Bank of New York. https://libertystreeteconomics.newyorkfed.org/2020/06/municipal-debt-markets-and-the-covid-19-pandemic.html.

CIPRIANI, M., G. LA SPADA, R. ORCHINIK, A. PLESSET, ET AL. (2020b): "The Money Market Fund Liquidity Facility," Tech. rep., Federal Reserve Bank of New York. https://libertystreeteconomics.newyorkfed.org/2020/05/the-money-market-mutual-fund-liquidity-facility.html.

CORNAGGIA, J., K. J. CORNAGGIA, AND R. D. ISRAELSEN (2018): "Credit ratings and the cost of municipal financing," *The Review of Financial Studies*, 31, 2038–2079.

FELDHÜTTER, P. (2012): "The same bond at different prices: identifying search frictions and selling pressures," *The Review of Financial Studies*, 25, 1155–1206.

FRANDSEN, B. R. (2017): "Party bias in union representation elections: Testing for manipulation in the regression discontinuity design when the running variable is discrete," in *Regression discontinuity designs*, Emerald Publishing Limited.

FRITSCH, N., J. BAGLEY, AND S. NEE (2021): "Municipal Markets and the Municipal Liquidity Facility," *Working Paper*.

GARRETT, D. G. (2020): "Conflicts of Interest in Municipal Bond Advising and Underwriting," *Working Paper*.

GRAY, C., A. LEIVE, E. PRAGER, K. PUKELIS, AND M. ZAKI (2020): "Employed in a SNAP? The Impact of Work Requirements on Program Participation and Labor Supply," *Working Paper (August 18, 2020)*.

GREEN, D. AND E. LOUALICHE (2020): "State and Local Government Employment in the COVID-19 Crisis," *Journal of Public Economics*, 104321.

GREEN, R. C., B. HOLLIFIELD, AND N. SCHÜRHOFF (2007a): "Dealer intermediation and price behavior in the aftermarket for new bond issues," *Journal of Financial Economics*, 86, 643–682.

——— (2007b): "Financial intermediation and the costs of trading in an opaque market," *The Review of Financial Studies*, 20, 275–314.

GREEN, R. C., D. LI, AND N. SCHÜRHOFF (2010): "Price discovery in illiquid markets: Do financial asset prices rise faster than they fall?" *The Journal of Finance*, 65, 1669–1702.

GÜRKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2007): "The US Treasury yield curve: 1961 to the present," *Journal of monetary Economics*, 54, 2291–2304.

HADDAD, V., A. MOREIRA, AND T. MUIR (2021): "When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed's response," *The Review of Financial Studies*, 34, 5309–5351.

HAHN, J., P. TODD, AND W. VAN DER KLAAUW (2001): "Identification and estimation of treatment effects with a regression-discontinuity design," *Econometrica*, 69, 201–209.

HARRIS, L. E. AND M. S. PIWOWAR (2006): "Secondary trading costs in the municipal bond market," *The Journal of Finance*, 61, 1361–1397.

HYMAN, B. (2018): "Can displaced labor be retrained? evidence from quasi-random assignment to trade adjustment assistance," *Evidence from Quasi-Random Assignment to Trade Adjustment Assistance (January 10, 2018)*.

IVANOV, I. T., T. ZIMMERMANN, AND N. W. HEINRICH (2021): "Limits of Disclosure Regulation in the Municipal Bond Market," *Working Paper*.

JANKOWITSCH, R., A. NASHIKKAR, AND M. G. SUBRAHMANYAM (2011): "Price dispersion in OTC markets: A new measure of liquidity," *Journal of Banking & Finance*, 35, 343–357.

LI, T. AND J. LU (2020): "Municipal Finance During the COVID-19 Pandemic: Evidence from Government and Federal Reserve Interventions," *Working Paper*.

LI, Y., M. O'HARA, AND X. A. ZHOU (2021): "Mutual fund fragility, dealer liquidity provisions, and the pricing of municipal bonds," *Available at SSRN 3728943*.

LUCK, S. AND T. ZIMMERMANN (2020): "Employment effects of unconventional monetary policy: Evidence from QE," *Journal of Financial Economics*, 135, 678–703.

MCCRARY, J. (2008): "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of econometrics*, 142, 698–714.

MOODY'S INVESTORS SERVICE (2020): "Rating Methodology: US Local Government General Obligation Debt," .

MOORE, E. (2017): "Auction-Based Liquidity of Last Resort," *Available at SSRN 2801138*.

NOVY-MARX, R. AND J. D. RAUH (2012): "Fiscal imbalances and borrowing costs: evidence from state investment losses," *American Economic Journal: Economic Policy*, 4, 182–213.

O'HARA, M. AND X. A. ZHOU (2021): "Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis," *Journal of Financial Economics*, 142, 46–68.

PIRINSKY, C. A. AND Q. WANG (2011): "Market segmentation and the cost of capital in a domestic market: Evidence from municipal bonds," *Financial Management*, 40, 455–481.

POTERBA, J. M. (1989): "Tax reform and the market for tax-exempt debt," *Regional Science and Urban Economics*, 19, 537–562.

ROLL, R. (1984): "A simple implicit measure of the effective bid-ask spread in an efficient market," *The Journal of finance*, 39, 1127–1139.

SCHULTZ, P. (2012): "The market for new issues of municipal bonds: The roles of transparency and limited access to retail investors," *Journal of Financial Economics*, 106, 492–512.

SCHWERT, M. (2017): "Municipal bond liquidity and default risk," *The Journal of Finance*, 72, 1683–1722.

SHEINER, L. (2021): "Why is state and local employment falling faster than revenues?" .

WANG, J., C. WU, AND F. X. ZHANG (2008): "Liquidity, default, taxes, and yields on municipal bonds," *Journal of Banking & Finance*, 32, 1133–1149.

WU, S. (2018): "Transaction Costs for Customer Trades in the Municipal Bond Market: What Is Driving the Decline?" *Municipal Securities Rulemaking Board*, 1, 29.

YI, H. L. (2020): "Finance, Public Goods, and Migration," *Working Paper*.

Figures and Tables

Figure 1: Turmoil in Secondary Market Municipal Bond Yields during COVID-19

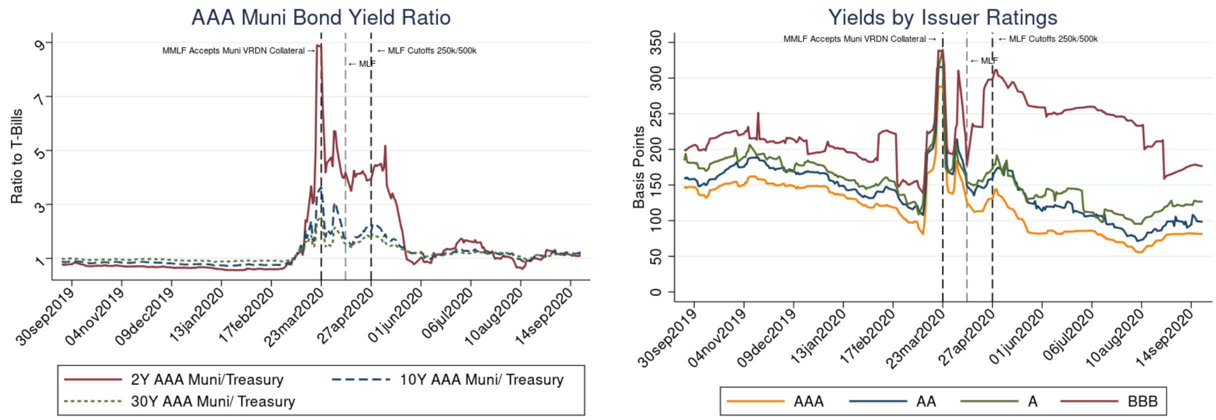

NOTES–This figure plots municipal bond yields by maturity and by rating. Both panels use daily yields (6-month) in the calculation of the dependent variable. Bonds ratings are averaged across Moody's and S&P ratings (take directly from Bloomberg). Yields in the first panel are sourced from the AAA BVAL curve from Bloomberg. Vertical dotted lines mark key events at March 23, 2020 (MMLF Expansion), April 9, 2020 (MLF Announcement), and April 27, 2020 (Expansion of MLF population cutoffs). *Source*: Bloomberg, originally calculated in Cipriani et al. (2020a).

Figure 2: Running Variable Distribution of Issuers and Trades around Lending Cutoffs

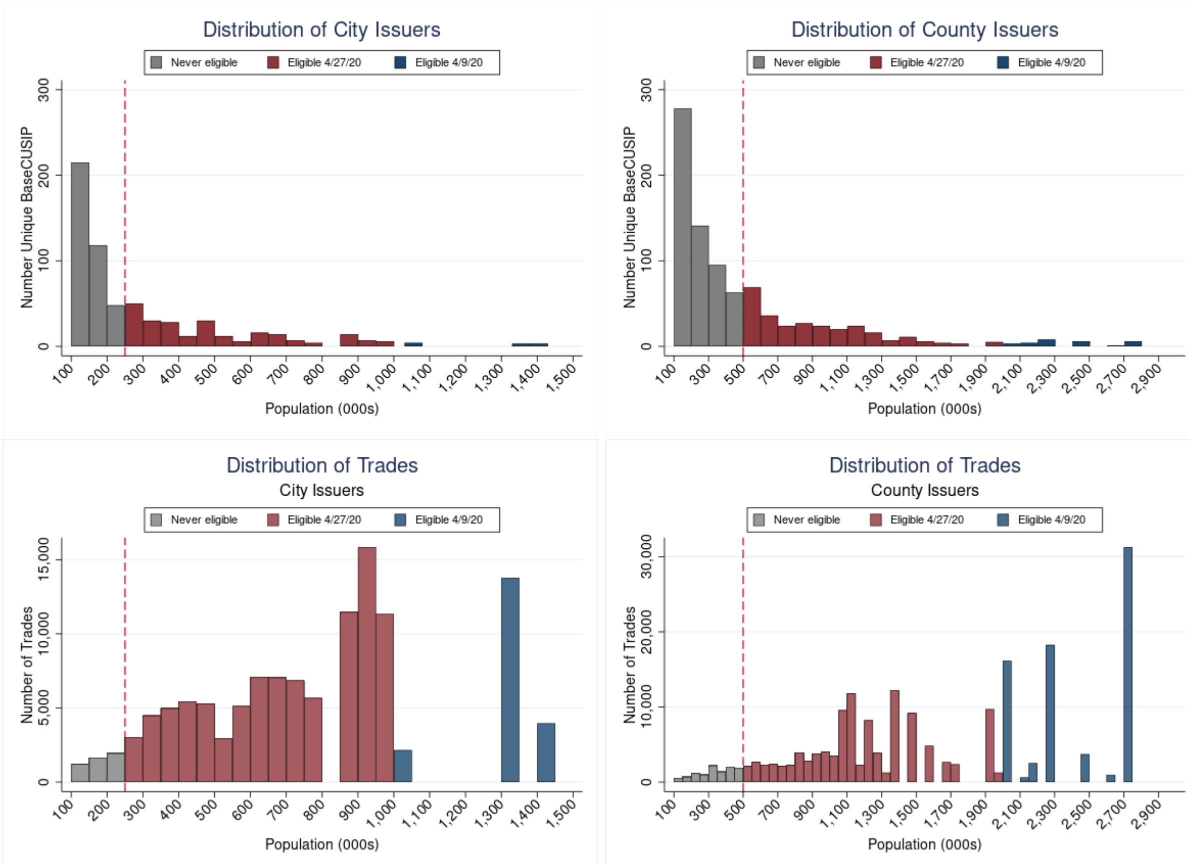

NOTES—This figure shows the distribution of number of issuers and number of trades in our matched sample around the lending eligibility cutoffs. Cities and counties below or equal to 100,000 are excluded for exposition. Trades are restricted to occur prior to March 1, 2020, to demonstrate variation prior to the onset of the pandemic and choice of cutoffs. The dashed line is the cutoff of interest for identification. *Source:* US Census Bureau, MSRB, Bloomberg

Figure 3: Plurality Ratings Distribution for Cities and Counties

(a) All Cities and Counties

(b) Cities and Counties within 100k of Cutoff

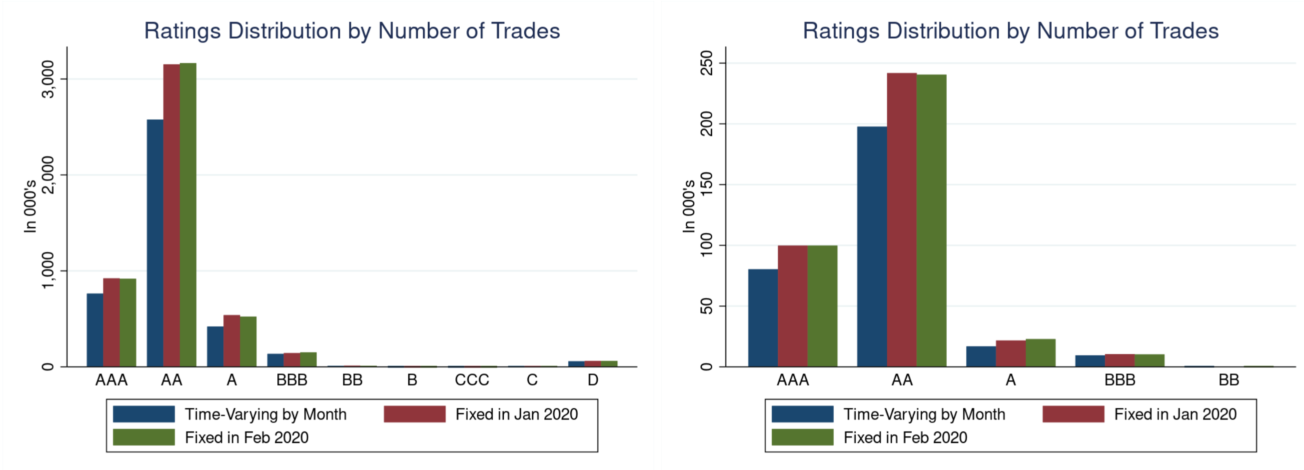

NOTES—This figure shows distributions of trades according to monthly plurality issuer ratings from January 1, 2019 to November 20, 2020. See text for details. *Source:* MSRB, S&P, Moody's, Fitch

Figure 4: Mean Yields within 100k Population of MLF Eligibility Cutoffs

By geography:

(a) Cities:

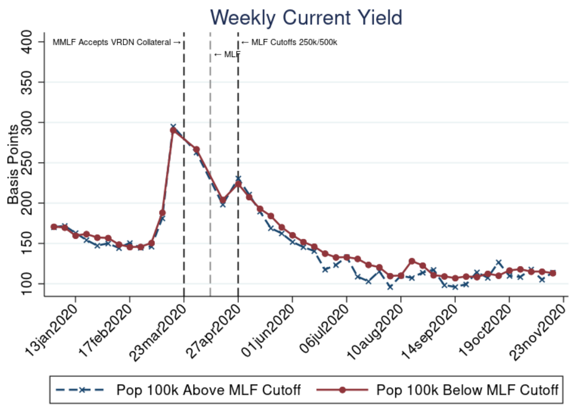

(b) Counties:

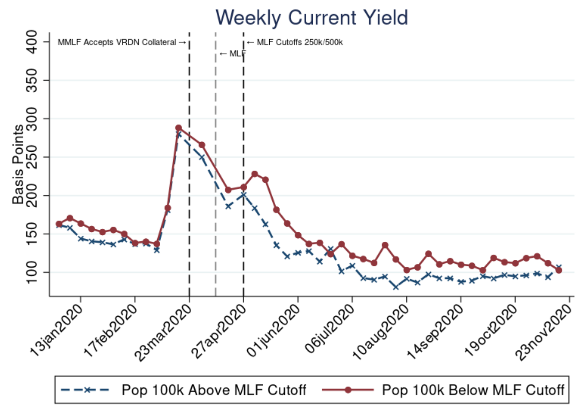

By credit worthiness:

(c) High-Rated (AAA & AA):

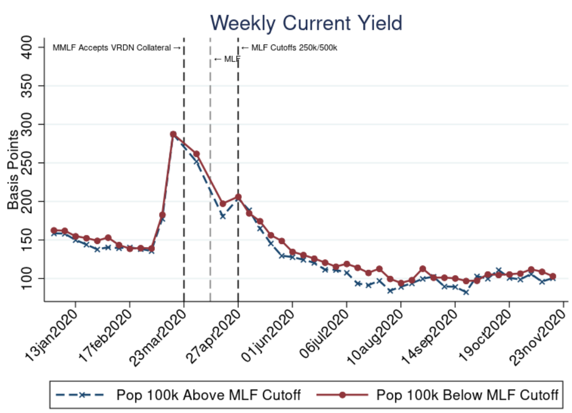

(d) Low-Rated (A & BBB):

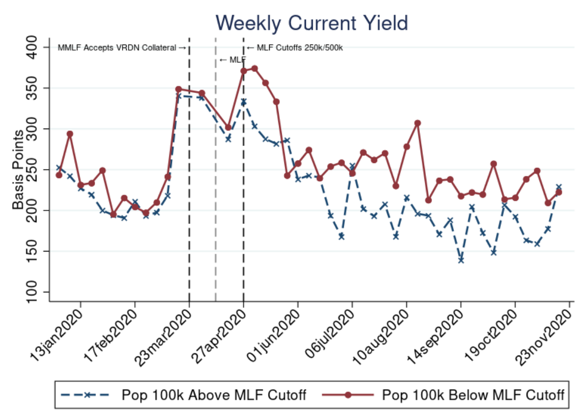

NOTES—Figure shows mean weekly yields (pooled over buyer and seller prices) for eligible issuers with cutoff-relative populations between 0 and 100,000 inclusively (blue dashed series), and issuers with cutoff-relative populations greater than or equal to -100,000 and less than 0 (red solid series). Trades between announcements are pooled into a single period beginning on the announcement day. See text for sample restrictions and definitions. *Source:* MSRB, S&P, Moody's, Fitch

Figure 5: RD Plots of MLF Access Effect on Yields by Credit Worthiness (Post-Period)

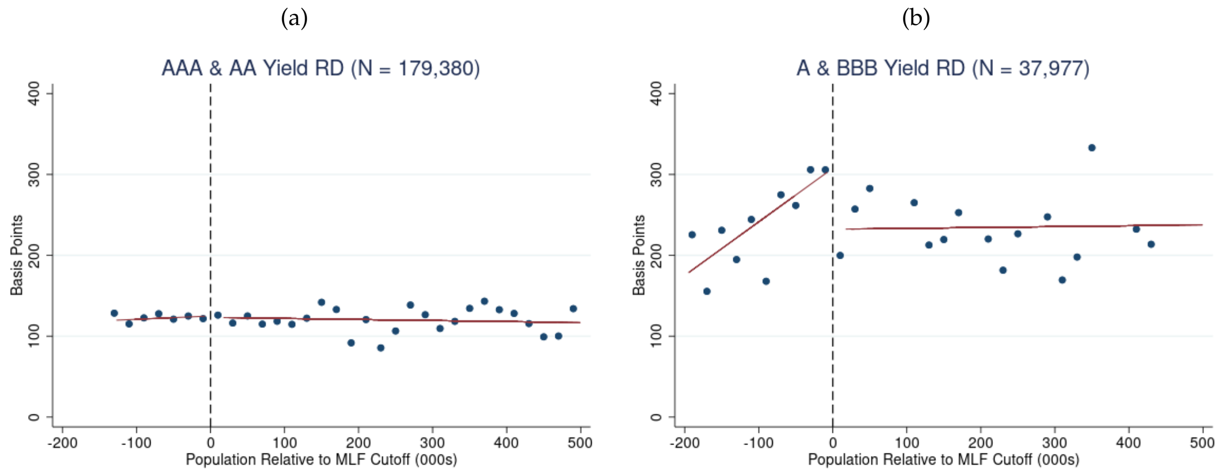

NOTES—These figures show regression slopes and intercepts from Equation 1 in the post period, overlaid on top of equally spaced pre-binned outcome data with a bin size of 20 (x-axis in thousands). Plots are shown over the optimal bandwidth selected using the IMSE-procedure, which produces (potentially) asymmetric optimal bandwidth boundaries for each sample. Plots correspond to Table 2, which includes state fixed effects as controls. Plots thus residualize all yields by state fixed effects (added back to their overall mean) prior to mean-collapsing by bin.

Figure 6: Robustness of A/BBB Yield RD Estimates to Bandwidth Size Around Cutoff

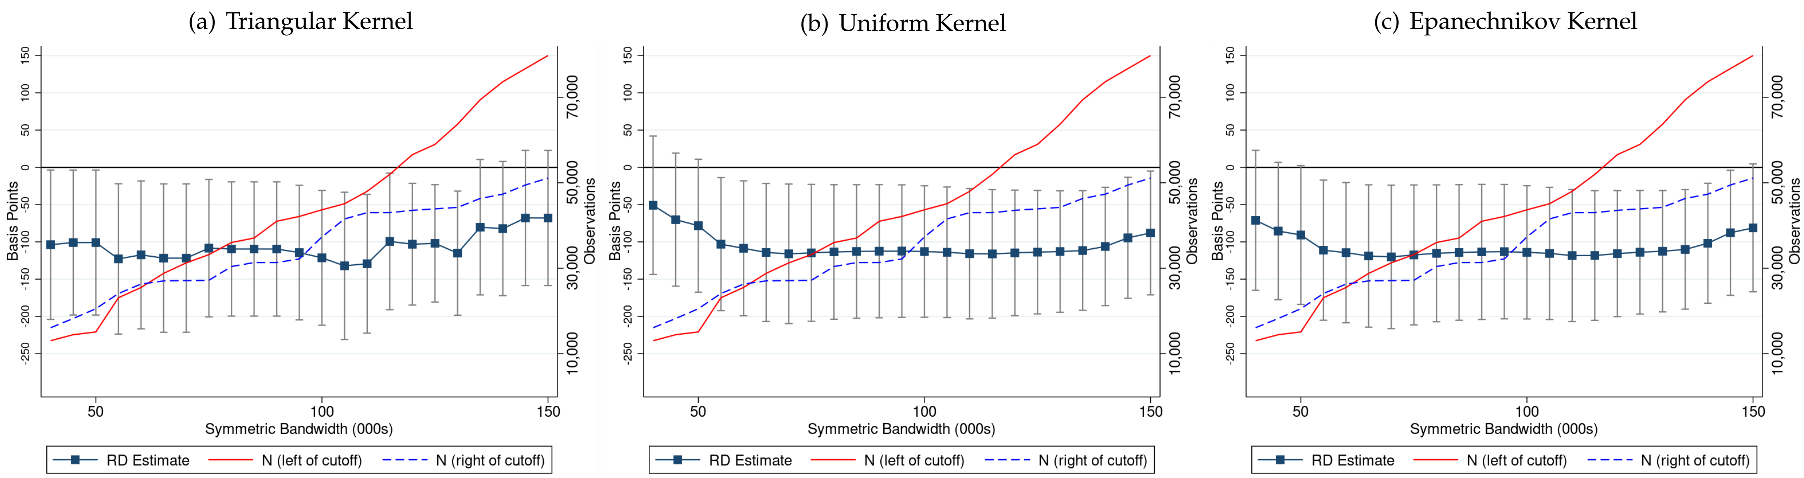

NOTES—In this figure, we show the sensitivity of our main A/BBB yield results to using larger and larger bandwidths. In our preferred specification, we use a data-driven IMSE-optimal asymmetric bandwidth in which the econometrician cannot choose the bandwidth. Here, to demonstrate sensitivity to size, we show results for a *symmetric* bandwidth in population relative to the cutoff using our preferred first-order polynomial, by increasing the bandwidth by 5,000 in population in each iteration. We start at 40,000 in population to ensure at least 10,000 observations on each side of the cutoff. Panel (a) shows results using a triangular kernel (which is also used in our preferred IMSE specification); panel (b) shows results using a uniform kernel; panel (c) shows results using an Epanechinkov kernel.

Figure 7: New Primary Issuance within 100k Population of MLF Eligibility Cutoffs

(a) Cumulative New Issuance, All Ratings   (b) Cumulative New Issuance, Low-Rated (A & BBB)

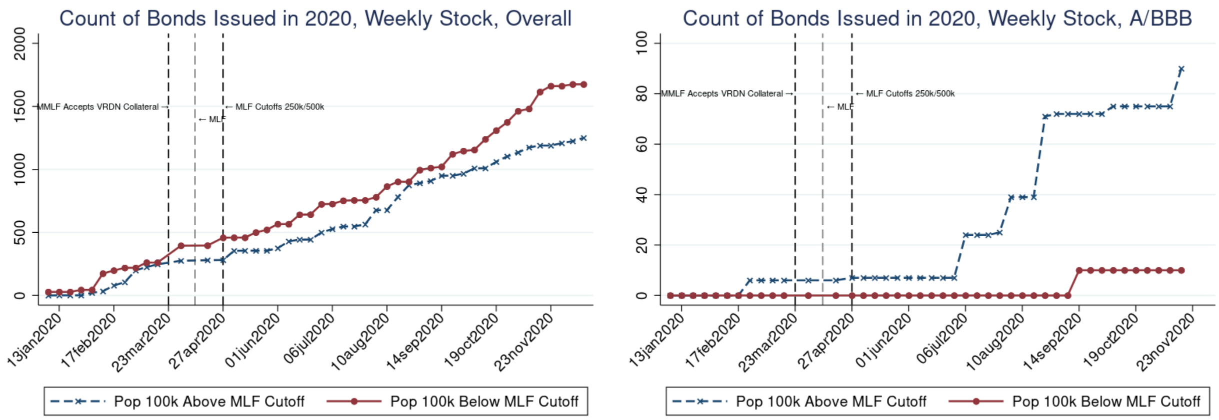

NOTES—Theese figures show new primary issuance as cumulative counts of newly issued bonds by MLF-eligible and ineligible issuers within a symmetric 100,000 population bandwidth. See text for further details. *Source:* MSRB, Mergent, Bloomberg

Figure 8: RD Plots of MLF Access Effect on Public Sector Employment

(a) Δ Overall Employment Post-MLF   (b) Δ Services Employment Post-MLF

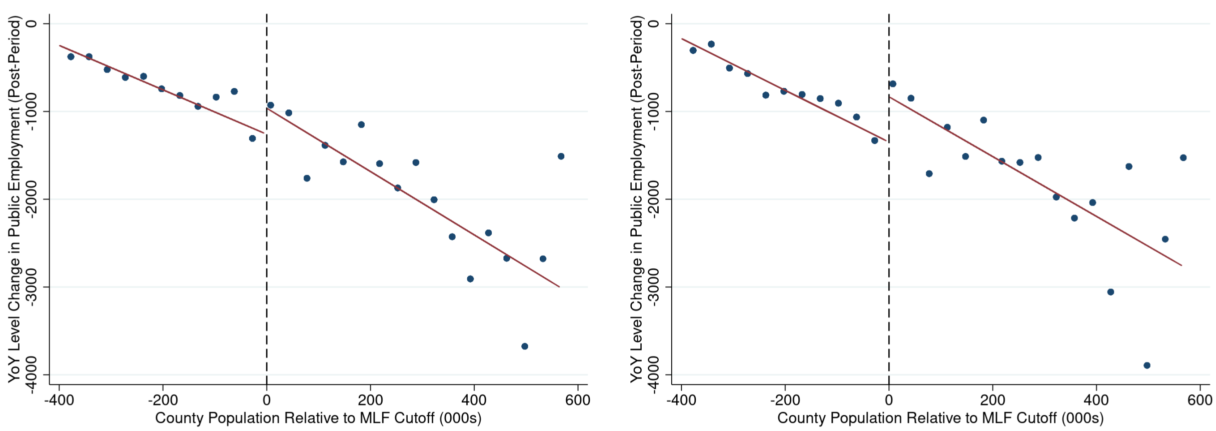

NOTES—Plots show regression slopes and intercepts from Equation 1 in the post-period (May and June 2020 pooled), overlaid on top of equally spaced pre-binned outcome data with a bin size of 35 (x-axis in thousands). Plots are shown over a fixed 1 million population wide bandwidth on the running variable. Plots first residualize employment of state fixed effects (added back to their overall mean) prior to mean-collapsing by bin, and correspond to(column (4) of Table 5.

Figure 9: Differenced RD Effects on Local Government Employment by Month

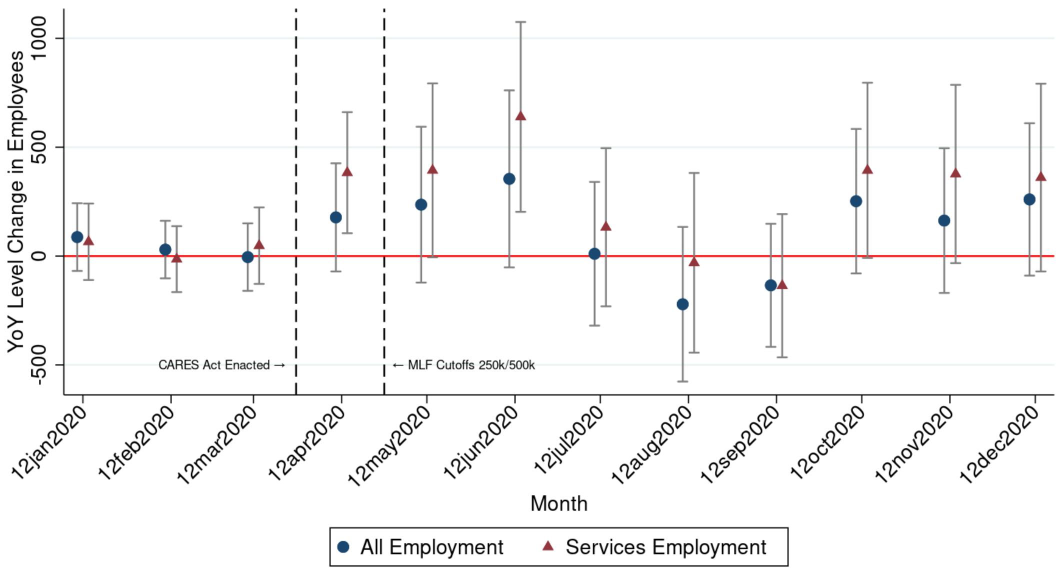

NOTES—Figure plots RD coefficients and 90% confidence intervals estimated using the specification from Table 5 column (4), with QCEW 12th of the month employment data. See Table 5 notes for further details.

Figure 10: Variation in MLF Lending Caps and direct CARES Aid Allocations

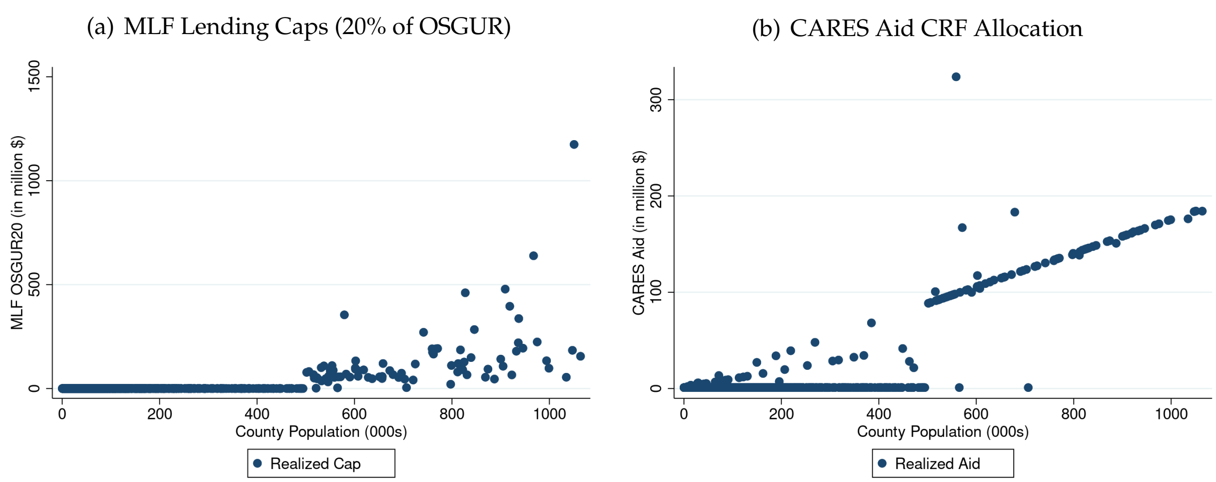

NOTES—Panels use different scales. Observations in panel (a) show MLF lending caps (based on 2017 OSGUR) over the county running variable. Observations in panel (b) show realized direct aid allocations calculated from the CRF population formula (see footnote) plus $1, combined with incomplete data on downstreamed aid from states to counties on the left of the cutoff.

Figure 11: RD Effects on Local Government Employment by Ratings and Month

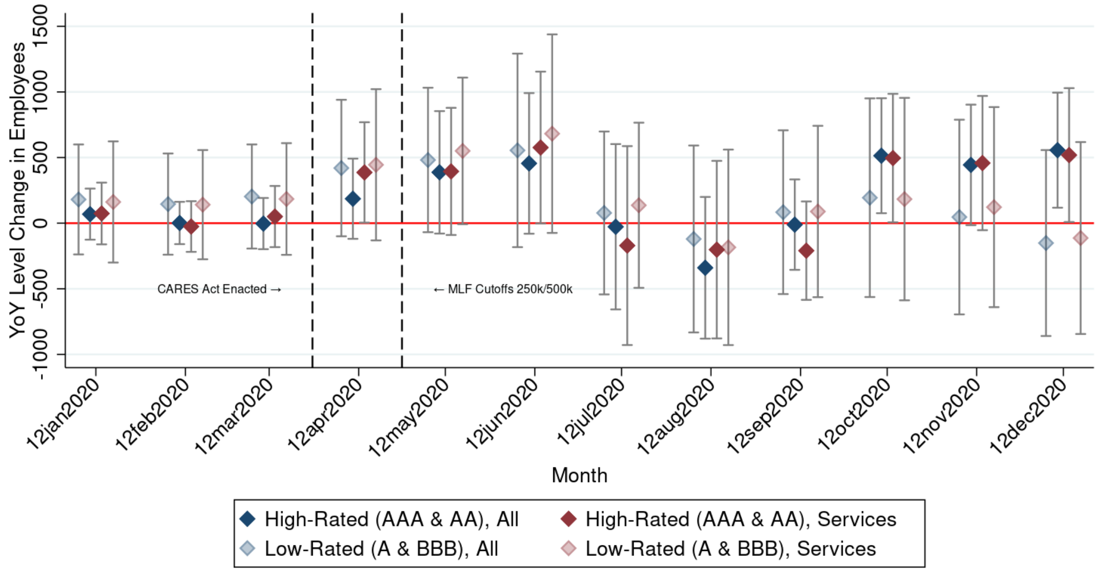

NOTES—-Figure plots RD coefficients and 90% confidence intervals estimated using the specification from Table 5 column (3) with QCEW 12th of the month employment data, separately estimated by January 2020 plurality ratings. Here we use the specification omitting state fixed effects so not to over-absorb differences in ratings that may differentially load on to states). See Table 5 notes for further details.

Figure 12: Illiquidity and Credit Risk Evolution in the 2020 Municipal Bond Market

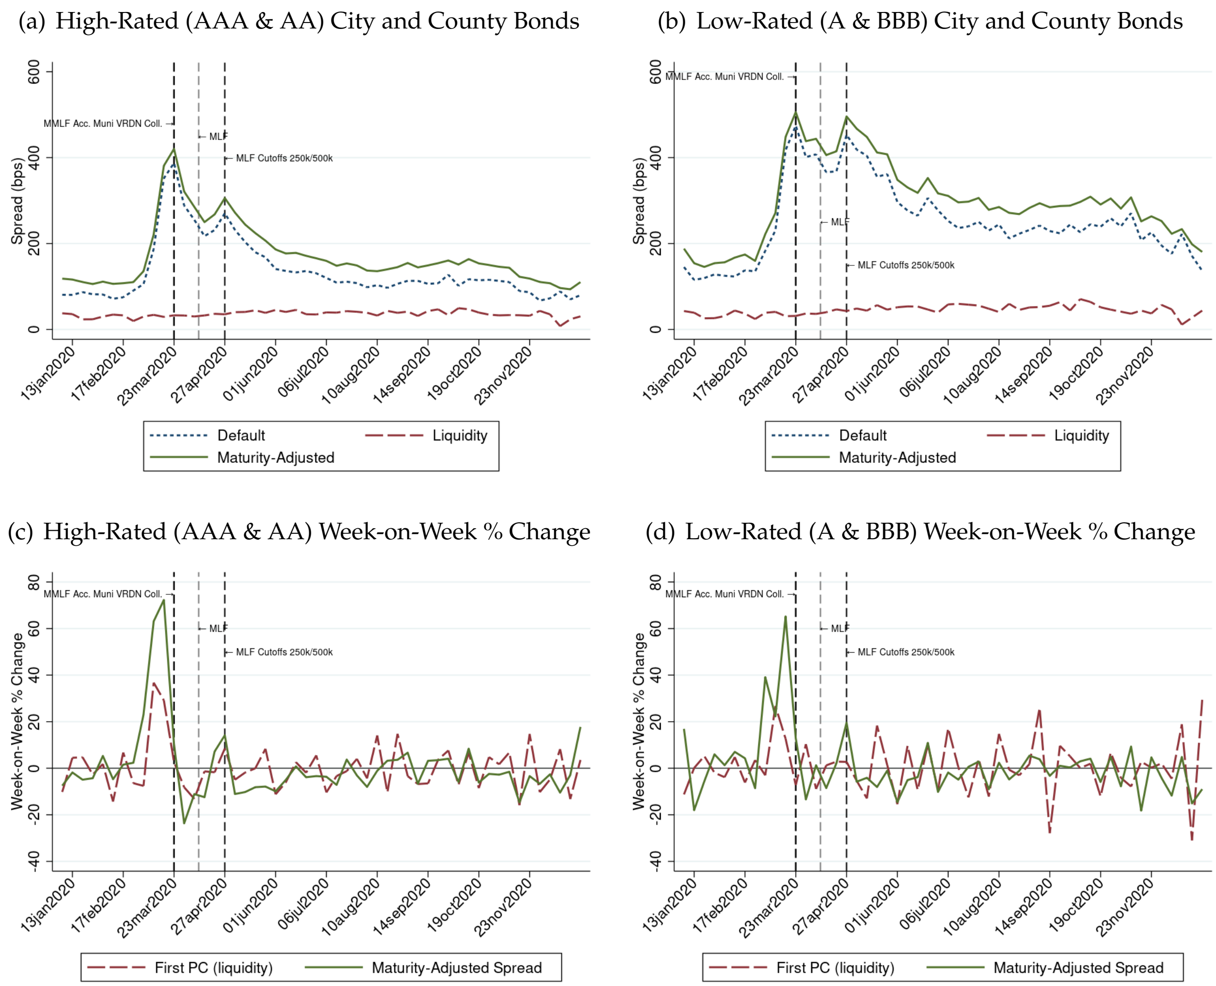

NOTES—Panels (a) and (b) show the weekly average of the default and liquidity components of city and county bond yields by rating group for each week in the 2020 calendar year (details in Appendix A.4.2). Panels (c) and (d) further show week-on-week changes in both the first principal component liquidity measure ($\lambda_{bt} - \lambda_{p01}$) and ($\hat{y}_{bist}$), demonstrating the relative speed in growth of overall yields and illiquidity respectively.

Table 1: RD Balance Table in Placebo Period

	Discontinuity	Standard Error	Control Mean	N (IMSE)	Bandwidths [l,r]
Coupon Rate (b.p.)	-23.32	17.36	435.32	99,376	[- 114, + 993]
Security Price (per 100 par)	-1.98*	1.11	109.95	111,063	[- 156, + 830]
Current Yield (b.p.)	-6.12	16.20	185.81	70,567	[- 88, + 587]
Δ Yield (Feb20-Jan20)	-0.03	0.05	-0.10	66,117	[- 187, + 1,459]
Δ Yield YoY (Jan20-Jan19)	-0.23	0.20	-0.94	54,748	[- 495, + 877]
Δ Yield YoY (Feb20-Feb19)	0.04	0.11	-1.01	53,251	[- 495, + 1,429]
Amount Outstanding (MM)	-18.84	190.89	473.60	90,163	[- 172, + 426]
Maturity Size (MM)	11.86	268.45	707.74	89,923	[- 172, + 410]
Tenor of Bond (Years)	0.36	0.90	12.88	110,907	[- 166, + 696]
Remaining Duration of Bond (Years)	-0.19	0.85	8.26	97,529	[- 133, + 686]
Market Share of Issuer	0.03	0.10	0.18	92,770	[- 165, + 467]
Number of Securities by Issuer	-29.24	43.26	227.60	129,779	[- 187, + 780]
Par Traded of Bond (1000s)	37.78	35.76	92.27	95,764	[- 118, + 741]
S&P Ratings (1-7 scale)	0.08	0.14	5.60	89,886	[- 189, + 511]
Moodyś Ratings (1-7 scale)	0.13	0.18	5.72	87,394	[- 196, + 493]
Fitch Ratings (1-7 scale)	-0.03	0.17	5.69	75,032	[- 490, + 1,101]
Time of Day of Trade (minute)	10.67**	4.86	768.28	101,011	[- 170, + 572]

NOTES—Table presents balance tests of covariates for MSRB active trades during the pre-period for the pooled sample. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). The Control Mean denotes the mean of that characteristic immediately to the left of relative population 0. Each regression uses IMSE-optimal bandwidths calculated separately for each side of the cutoff and for each outcome, and a triangular kernel to weight observations. Sample sizes vary depending on the bandwidth used. See Appendix A.3 for further details on variable definitions and construction. Standard errors are clustered by population relative to the cutoff (the running variable). *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.1$. *Source:* MSRB, Bloomberg, S&P, Moody's, Fitch

Table 2: RD Estimates of MLF Access on Secondary Market Yields

	Discontinuity	Standard Error	Control Mean	N (IMSE)	Bandwidths [l,r]
a. Pooled Post:					
Overall Yield	-19.42	19.47	156.71	189,076	[- 113, + 652]
City Only	-27.91	25.48	171.75	92,349	[- 77, + 472]
County Only	-16.04	24.11	134.83	54,574	[- 109, + 446]
High-Rated	-1.75	7.80	124.99	179,380	[- 130, + 685]
Low-Rated	-75.43**	34.11	307.69	37,977	[- 206, + 827]
b. Pooled Pre (Placebo):					
Overall Yield	-13.26	12.15	194.07	70,567	[- 88, + 587]
City Only	-9.34	12.14	196.35	32,499	[- 65, + 393]
County Only	-21.61	19.17	191.88	28,325	[- 125, + 527]
High-Rated	-6.09	7.73	177.31	67,350	[- 107, + 618]
Low-Rated	-25.14	21.64	263.29	12,576	[- 166, + 885]

NOTES—Table presents RD estimates of current yields (in basis points) for MSRB-active trades during each sample period. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). The Control Mean denotes the mean of that characteristic immediately to the left of relative population 0. Each regression uses IMSE-optimal bandwidths calculated separately for each side of the cutoff and for each outcome, and a triangular kernel to weight observations. Sample sizes vary depending on the bandwidth used. State and calendar month fixed effects are included in all regressions. Standard errors are clustered by population relative to the cutoff (the running variable). See text for further sample restrictions. *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.1$.

Table 3: Robustness of Main Yields Results (A/BBB) to RD Kernel and Bandwidth Choice

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
a. Pooled Post:									
Overall Yield	-10.58	**-19.42**	-15.45	-4.93	-13.64	-9.61	-9.11	-21.83	-14.47
City Only	-7.34	**-27.91**	-30.97	2.32	-30.58	-39.92	-4.43	-29.11	-30.20
County Only	-20.78	**-16.04**	-13.52	-23.72*	-20.66	-10.23	-22.19	-14.84	-12.37
High-Rated	-0.58	**-1.75**	-1.70	1.89	1.67	1.70	-0.22	-1.42	-1.05
Low-Rated	-72.56*	**-75.43****	-75.28**	-61.60*	-86.35**	-93.87**	-72.57*	-76.89**	-89.49**
b. Pooled Pre (Placebo):									
Overall Yield	-5.34	**-13.26**	-8.26	-2.21	-10.60	-8.18	-4.96	-15.03	-8.01
City Only	-0.08	**-9.34**	-17.72	3.56	-8.70	-8.53	0.44	-8.20	-17.58
County Only	-15.94	**-21.61**	-22.98	-12.77	-22.87	-21.54	-15.46	-21.41	-22.40
High-Rated	-3.37	**-6.09**	-5.29	-0.86	-1.74	-6.58	-3.08	-6.12	-6.56
Low-Rated	-20.45	**-25.14**	-25.69	-3.48	-36.17	-43.17	-19.59	-25.24	-28.54
Bandwidth	Fixed	**IMSE**	IMSE	Fixed	IMSE	IMSE	Fixed	IMSE	IMSE
Kernel	TRI	**TRI**	TRI	UNI	UNI	UNI	EPN	EPN	EPN
Polynomial Degree	1	**1**	2	1	1	2	1	1	2
N (Low-Rated Post)	34,189	**37,977**	72,306	34,189	24,512	60,352	34,189	31,144	65,872

NOTES—Table presents RD estimate sensitivity of current yields (in basis points) for MSRB-active trades during each sample period to different RD kernel and bandwidth choices. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). IMSE: Integrated Mean-Squared Error; TRI: Triangular; UNI: Uniform; EPN: Epanechnikov. State and calendar month fixed effects are included in all regressions. Standard errors are clustered by population relative to the cutoff (the running variable), suppressed here for exposition. See text for further sample restrictions. *** p≤0.01, ** p≤0.05, * p≤0.1.

Table 4: RD Estimates of MLF Access on New Primary Issuance

	Discontinuity	SE	Control Mean	N (Fixed)	Bandwidths [l,r]
a. Pooled Post:					
Prob(New Issuance), Overall (27apr-20nov)	0.08**	0.03	0.10	89,882	[- 400, + 600]
Prob(New Issuance), A & BBB (27apr-20nov)	0.22**	0.11	-0.06	8,190	[- 400, + 600]
Total Issuance (MM), Overall (27apr-20nov)	0.76	24.78	81.66	89,882	[- 400, + 600]
Total Issuance (MM), A & BBB (27apr-20nov)	18.86	46.14	43.16	8,190	[- 400, + 600]
b. Pooled Pre (Placebo):					
Prob(New Issuance), Overall (01jan-23mar)	0.07**	0.03	0.02	45,108	[- 400, + 600]
Prob(New Issuance), A & BBB (01jan-23mar)	0.08	0.08	-0.07	3,901	[- 400, + 600]
Total Issuance (MM), Overall (01jan-23mar)	23.93	26.67	25.84	45,108	[- 400, + 600]
Total Issuance (MM), A & BBB (01jan-23mar)	8.58	7.53	2.30	3,901	[- 400, + 600]

NOTES—Table presents RD estimates of the probability of new primary issuance and total amount of new issuance (millions) in each period. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). The Control Mean denotes the mean of that characteristic immediately to the left of relative population 0, which can be negative due to the projected RD polynomial. Sample sizes vary depending on the bandwidth used. Standard errors are clustered by population relative to the cutoff (the running variable). See text for further sample restrictions. *** p≤0.01, ** p≤0.05, * p≤0.1.

Table 5: RD Estimates of MLF Access on Public Sector Employment

	Emp. (1)	Emp. (2)	Δ Emp. (3)	Δ Emp. (4)	% Δ Emp. (5)	% Δ Emp. (6)	N (fixed-bwdth)
a. Pooled Post:							
Overall Employment	-477	323	325	297	1.19	1.18	945
	(1,016)	(1,854)	(239)	(223)	(0.96)	(0.83)	
– Goods Employment	-42	-49	1	2	2.61	3.93	248
	(63)	(37)	(5)	(5)	(3.85)	(4.27)	
– Services Employment	-412	-666	422*	517**	1.61	1.69**	711
	(1,134)	(2,001)	(238)	(242)	(1.00)	(0.85)	
b. Pooled Pre (Placebo):							
Overall Employment	-828	189	53	58	0.23	0.26	946
	(1,042)	(1,995)	(84)	(82)	(0.41)	(0.37)	
– Goods Employment	-44	-51	-1	-1	-0.38	0.51	248
	(64)	(38)	(3)	(3)	(2.23)	(2.76)	
– Services Employment	-762	-896	41	26	0.10	0.04	712
	(1,128)	(2,188)	(98)	(93)	(0.50)	(0.41)	
Month FEs		X					
State FEs		X		X		X	
Control Mean (post): Overall Employment	18,835	13,346	-1,717	-1,259	-8.06	-8.13	
Control Mean (post): Goods Employment	162	127	-7	-5	-6.44	-7.16	
Control Mean (post): Services Employment	18,506	14,054	-1,674	-1,349	-7.96	-8.58	

NOTES—Table presents RD estimates of employment measure in column header and sample period. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). The Control Mean denotes the mean of that characteristic immediately to the left of relative population 0. Each regression uses a fixed bandwidth with a triangular kernel to weight observations (see Section 4 for details). Percents in columns (5) and (6) range from 0 to 100. Standard errors are conservatively clustered by population relative to the cutoff (the running variable); as such, inference is drawn across unique issuers. *** p≤0.01, ** p≤0.05, * p≤0.1.

Table 6: Spreads Decomposition Summary Statistics, City and County Issuers

	Pre				Outbreak/Peak				Interventions/Post			
	Mean (1)	Wgt. Mean (2)	SD (3)	# Obs (4)	Mean (5)	Wgt. Mean (6)	SD (7)	# Obs (8)	Mean (9)	Wgt. Mean (10)	SD (11)	# Obs (12)
A. Liquidity Measures												
BAS	0.62	0.54	0.58	353	0.72	0.61	0.67	650	0.67	0.55	0.68	4,182
IRC (%)	0.57	0.51	0.55	6,711	0.63	0.56	0.62	5,394	0.63	0.57	0.61	47,528
Roll	0.86	0.85	0.94	6,711	1.30	1.15	1.33	5,394	1.08	1.07	1.13	47,528
Dispersion	0.28	0.26	0.27	6,711	0.35	0.36	0.33	5,394	0.33	0.32	0.30	47,528
Amihud (% per $1m)	23.1	9.87	35.3	6,711	22.7	11.2	34.5	5,394	23.7	12.6	34.8	47,528
First PC	1.26	1.05	1.24	6,711	1.55	1.35	1.36	5,394	1.46	1.28	1.33	47,528
B. AAA/AA Sample												
Yield (%)	1.72	1.78	0.90	5,511	2.29	2.51	1.10	4,728	1.70	1.86	1.07	38,937
Tax-Adj. Yield (%)	2.87	2.85	1.09	5,511	3.96	4.25	1.82	4,728	2.92	3.07	1.67	38,937
Maturity-Adj. Spread (bps)	11.1	10.4	10.0	5,511	28.3	30.1	17.9	4,728	21.0	21.3	15.7	38,937
Liquidity Spread (bps)	2.97	2.45	3.00	5,511	3.07	2.58	2.63	4,728	3.68	3.13	3.42	38,937
Default Spread (bps)	8.12	7.97	9.40	5,511	25.2	27.6	17.8	4,728	17.3	18.2	15.4	38,937
Liquidity/Mat-Adj. (%)	26.8	23.5			10.9	8.55			17.6	14.7		
Default/Mat-Adj. (%)	73.2	76.5			89.1	91.5			82.4	85.3		
C. A/BBB Sample												
Yield (%)	2.36	2.98	1.22	922	2.79	3.22	2.16	487	2.69	3.57	1.41	6,980
Tax-Adj. Yield (%)	3.44	3.74	1.48	922	4.32	5.18	2.65	487	4.30	5.62	2.09	6,980
Maturity-Adj. Spread (bps)	16.3	19.1	14.2	922	31.8	38.1	26.2	487	34.6	46.7	20.3	6,980
Liquidity Spread (bps)	3.35	3.05	3.14	922	3.67	2.20	3.32	487	4.75	4.34	4.05	6,980
Default Spread (bps)	12.9	16.0	13.9	922	28.2	35.9	25.9	487	29.8	42.4	20.1	6,980
Liquidity/Mat-Adj. (%)	20.6	16.0			11.5	5.78			13.7	9.28		
Default/Mat-Adj. (%)	79.4	84.0			88.5	94.2			86.3	90.7		

NOTES—This table details decomposition summary statistics at the bond-week level for three representative periods: *Pre-Crisis (January 6 - March 1)*, *Outbreak/Peak (March 2 - March 22)*, and *Post-Intervention (March 23 - Dec 31)*. *Yield (%)* is yield to maturity. *Tax-Adjusted Yield* is the yield adjusted for its tax-status and maturity in Equation 2, but not yet benchmarking to treasury yields. *Tax-Adjusted Spread (%)* is the tax (maturity) adjusted yield relative to the risk-free Treasury curve ($\hat{y}_{bist}$). *Default Spread (bps)* and *Liquidity Spread (bps)* are the respective components γ_{bist} and $\hat{l}_{bt}$ of the tax (and maturity) adjusted spread. *Default/Tax-Adj (%)* and *Liquidity/Tax-Adj (%)* are the ratio of the Default and Liquidity Spreads to the Tax (and maturity) Adjusted Spread. Weighted and unweighted means refer to whether measures in the regression are weighted by par volume traded (see Appendix A.4.2 for further details).

A Internet Appendix (IA)

A.1 Defining County Exposure to MLF

Figure A.1: Three Cases for Employment RD Analysis

(a) County-Only: Salt Lake County, UT

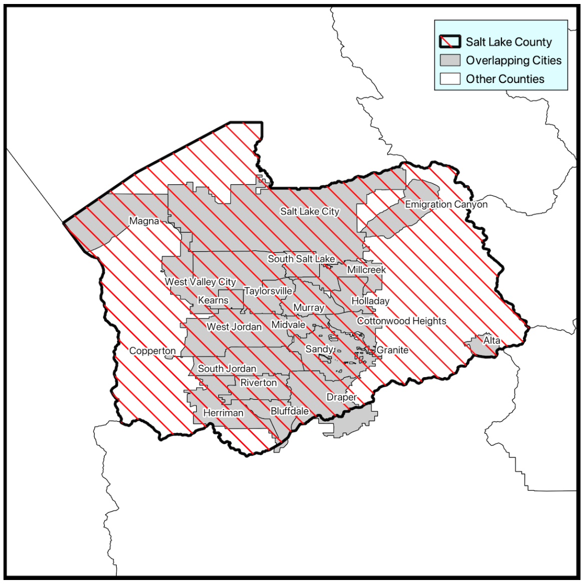

(b) City-Only: Lincoln, NE

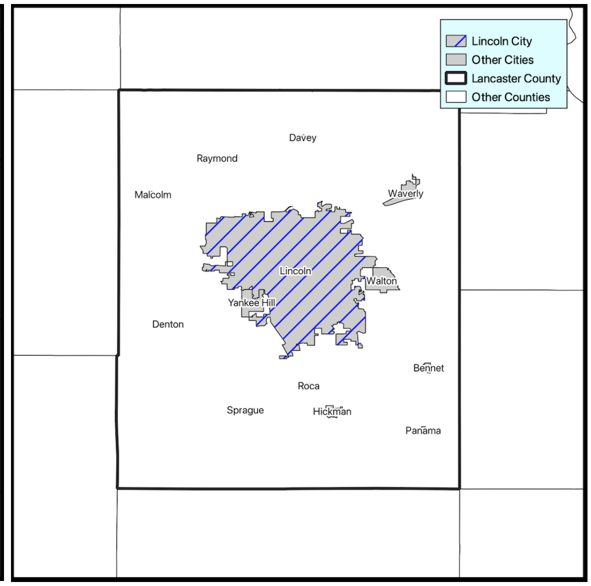

(c) Both: K.C. & Jackson County, MO

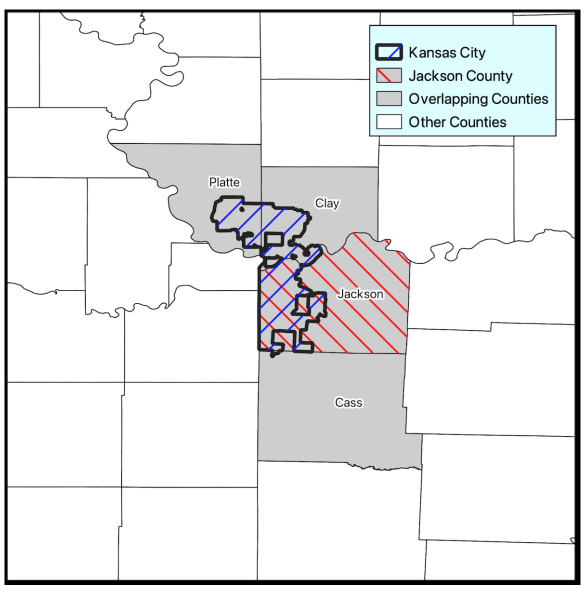

NOTES—Figure shows three representative types of counties in which employment is measured at county-level as the QCEW total of public sector employees across both cities and counties. In panel (a), Salt Lake County, UT, is MLF-eligible while none of its overlapping cities (nor unincorporated areas) is eligible; in panel (b), Lincoln, NE is eligible, but none of the other cities nor the county containing it (Lancaster County) are eligible; in panel (c), both Kansas City, MO, and Jackson County, MO, are eligible, while the remaining 3 counties spanned by Kansas City would be considered type (b), city-eligible only. Not shown is the final case (d), where no cities nor the underlying county are eligible.

A.2 Robustness and Sensitivity

A.2.1 Dynamic Yields Estimates: RD vs. Difference-in-Differences

Our preferred strategy uses an RD design rather than a difference-in-differences (DID) approach leveraging policy timing because the initial April 9 MLF announcement was contemporaneous with a $2.3 trillion suite of confounding policies including support for the Payroll Protection Program (PPP), Main Street Lending program, expansions to corporate credit facilities (PMCCF and SMCCF) as well as the Term Asset-Backed Securities Loan Facility (TALF). Focusing on cutoffs specific to the MLF allows us to better isolate the effects of one policy from others during this period. However, an alternative DID specification is also available if the April 27 MLF cutoff expansion is considered the DID "event", as it is insulated from contemporaneous policies when not studied in conjunction with other staggered policies. Below we compare estimates from the two strategies, which we also use as an opportunity to briefly discuss dynamic effects on yields beyond the first few months after the MLF eligibility expansion.

In panel (a) of Figure A.2, we report coefficients looping through our main RD specification for low-rated issuers, estimated bi-weekly across the extent of our sample frame. The specification is identical to that used throughout Table 2, but omits calendar month fixed effects which are absorbed by the sampling restriction of keeping each 2-week period respectively. In panel (b), we run an analogous DID equation at the bi-weekly level, but restrict treatment and control units to be within an 100,000 symmetric bandwidth of MLF cutoffs (similar to Figure 4), and weight observations closer to the cutoff with a triangular kernel. For panel (b), we estimate the following equation:

$$Y_{n(bi)t} = \alpha D_i + \sum_{\tau \neq -1} [\delta_\tau \times \mathbb{1}_{\{t-Apr27=\tau\}} + \beta_\tau \times \mathbb{1}_{\{t-Apr27=\tau\}} \times D_i] + \mathbf{X}_{bit}^{'}\gamma + \varepsilon_{n(bi)t} \tag{6}$$

where D_i is an indicator variable that takes 1 if an issuer's population exceeds the expanded cutoff by at most 100,000 in population, and 0 if an issuer's population is at most 100,000 less than the cutoff. The vector $\mathbf{X}$ includes only state fixed effects, τ is event time relative to the omitted two-week period just prior to April 27, and β_τ are the coefficients of interest.

Figure A.2: Dynamic RD vs. DID Estimates for Low-Rated Yields Sample

(a) Dynamic RD Estimates (Low-Rated A & BBB)

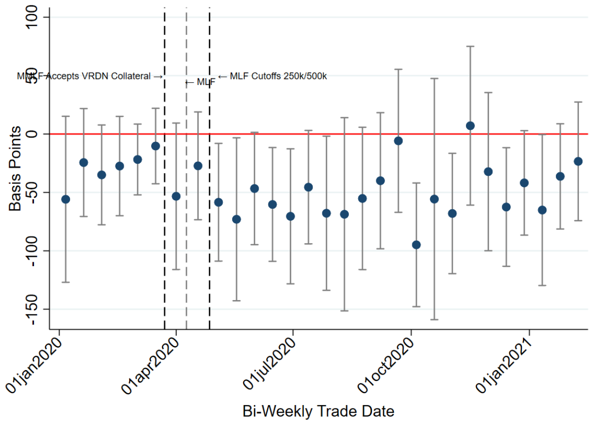

(b) Dynamic DID Estimates (Low-Rated A & BBB)

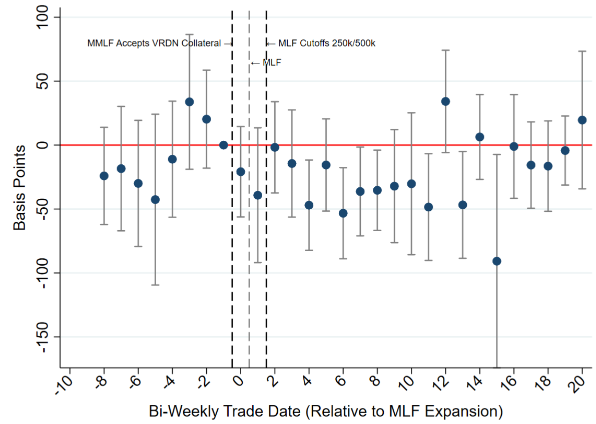

NOTES—Estimates and 90% confidence intervals shown looping over bi-weekly intervals of the support, and collapsing all observations between announcement dates into one estimate (consistent with main results).

The two approaches appear to deliver qualitatively similar results: a mildly upward sloping set of estimates prior to the crisis (though not statistically different from zero), exhibits a trend reversal immediately following the crisis and Federal interventions, which is then followed by considerable mean reversion by the end of the sample frame—a pattern which is consistent with recovery in the municipal market by the end of the 2020 calendar year. While effects are slightly muted using a DID strategy relative to RD (perhaps because observations are not limited to an optimal bandwidth and thus add greater weight further away from the cutoff), many of the negative point estimates in the critical first several months following the MLF expansion remain statistically different from zero in the DID plot as well. These amount on average to a similar point estimate in the pooled post-period as estimated using the RD in Table 2. To ensure as well-identified comparisons as possible, i.e. closer to the cutoff, we opt for the considerable advantages of the RD approach.

A.2.2 Yields Sensitivity to CARES Act Notch in CRF Aid Formula

In Appendix A.2.5 we confirm that cutoffs were not chosen based on the underlying variation in issuers. Yet, this does not rule out the possibility that these heuristic population cutoffs could be binding for other government policies, confounding our interpretation of the intervention. As discussed in the context of employment effects, we identify one potential confounding policy: the March 27 CARES act included a Coronavirus Relief Fund (CRF) provision that "local governments serving a population of at least 500,000, as measured in the most recent census data, may elect to receive assistance directly from Treasury. Such direct local assistance allocations reduce the allocation that is made to the state government (keeping the state allocation constant)."[67] In other words, localities over 500,000 in population may have had quicker or more direct access to CRF aid, relative to localities just under that cutoff that had to rely on downstreaming from their underlying state allocation cap.[68] Since allocation caps were fixed to the right of this cutoff, but unrestricted to the left, the amounts received may have been differentially higher or lower to the right of the cutoff.

As discussed, this poses a potential problem for counties in our analysis, which are subject to the April 27 MLF population cutoff (also using 500,000 as an eligibility threshold). In the context of yields, we are served by two additional tests to help us tease apart CARES from MLF effects, which were not available in our employment analysis. We first note that visually, we have substantial variation in the month that elapsed between the CARES and MLF revision announcement. During this inter-announcement interval, yields appeared to trend almost identically (shown in Figure 4), whereas they diverged earlier among employment outcomes. Second, only one of the two cutoffs (counties) were aligned exactly at 500,000, whereas the other cutoff (cities) was not. We thus show in Figure A.3 the variation underlying our main results for low-rated (A & BBB) yields, excluding counties from the analysis and thus only focusing on the 250,000 *city* population cutoff which is unimpaired by CARES aid—or at least, only affected far away from the cutoff in a way in which the RD polynomial can fully control for.

Figure A.3: Mean Yields within 100k of MLF Population Cutoff, Low-Rated (A & BBB) Cities Only

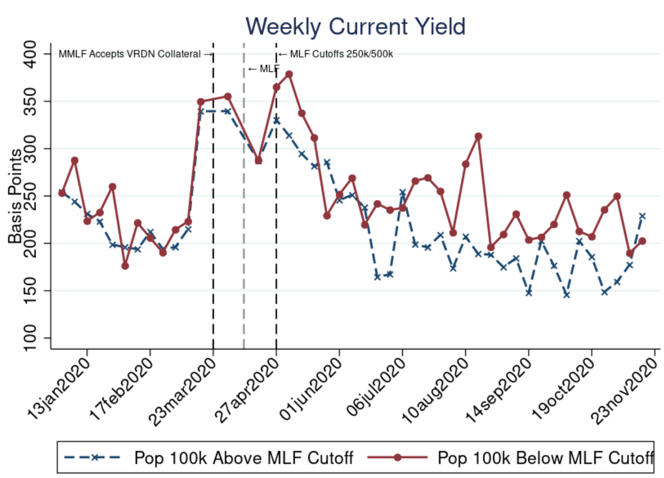

NOTES—Figure shows mean weekly yields (pooled over buyer and seller prices) for low-rated city eligible issuers with cutoff-relative populations between 0 and 100,000 inclusively (blue dashed series), and issuers with cutoff-relative populations greater than or equal to -100,000 and less than 0 (red solid series). Trades between announcements are pooled into a single period beginning on the announcement day. See text for sample restrictions and definitions. *Source:* MSRB, S&P, Moody's, Fitch

[67] https://crsreports.congress.gov/product/pdf/R/R46298.

[68] CRF aid allocation caps for localities greater than 500,000 in population are calculated as the product of the total state allocation (based on population) weighted by locality population share.

This additional placebo test shows a very similar pattern to our main results including counties, and our yield RD estimates for low-rated cities are similarly robust and statistically significant. While the results for employment are less conclusive, we can more confidently conclude that *investors* seem not to have responded to CARES act aid, whereas they did respond to an MLF emergency liquidity option.

A.2.3 Sensitivity to Controls

We report sensitivity of our main RD yields estimates to controls and sample restrictions in Table A.1, where column (7) is our preferred specification.

Table A.1: RD Effects on Yields, Sensitivity to Controls and Restrictions

	(1)	(2)	(3)	(4)	(5)	(6)	**(7)**	(8)	(9)	(10)	(11)	(12)	(13)
a. Pooled Post:													
Current Yield (Overall)	-17.66	-10.02	-11.29	-11.42	-18.75	-16.36	**-19.42**	-20.68	-19.69	-20.13	-20.62	-21.26	-20.39
City Only	-22.00	-13.11	-13.56	-13.60	-25.38	-24.41	**-27.92**	-30.52	-28.08	-27.07	-26.67	-25.73	-23.26
County Only	-16.85	-20.90	-23.04	-23.00	-13.64	-12.04	**-16.04**	-13.68	-16.59	-14.67	-14.40	-16.27	-21.06
High-Rated (AAA & AA)	8.67	4.14	4.20	4.29	-1.30	0.22	**-1.74**	-3.17	-1.94	-1.57	-2.16	-7.93	-6.99
Low-Rated (A & BBB)	-132.40**	-123.17***	-122.60***	-122.39***	-73.20**	-73.09**	**-75.43****	-77.05**	-75.99**	-65.61*	-67.62**	-60.27*	-56.14*
b. Pooled Pre (Placebo):													
Current Yield (Overall)	-9.41	-6.12	-6.78	-6.93	-14.62	-13.01	**-13.26**	-13.52	-14.66	-13.51	-13.62	-11.71	-11.26
City Only	-1.01	-1.18	-1.31	-1.26	-8.30	-8.11	**-9.34**	-9.52	-11.13	-8.26	-8.09	-5.47	-5.10
County Only	-42.47	-24.49	-26.48	-26.81	-24.05	-22.29	**-21.61**	-19.84	-19.46	-20.23	-20.42	-22.40	-24.29
High-Rated (AAA & AA)	1.21	1.54	1.67	1.93	-5.90	-4.96	**-6.09**	-6.60	-5.67	-5.68	-5.86	-7.13	-5.73
Low-Rated (A & BBB)	-77.80**	-57.08**	-55.54**	-55.58**	-22.19	-23.53	**-25.14**	-25.95	-31.68	-16.35	-19.06	-21.98	-20.79
Fed. Tax Adjust			X	X	X								
St. & Fed. Tax Adjust				X	X								
State FE					X	X	**X**	X	X	X	X	X	X
Month FE							**X**	X	X	X	X	X	X
Revenue/GO Bond								X					
Day/Week of Trade									X				
Maturity Size										X			
Amount Outstanding											X		
Tenor Length												X	
Remaining Duration													X
Sample Restrictions		X	X	X	X	X	**X**	X	X	X	X	X	X

NOTES—Sample restrictions include 1% winsorizing and outlier filtering. See Table 2 notes and text for further details. *** p≤0.01, ** p≤0.05, * p≤0.1.

Column (2) applies the NBER-TAXSIM deduction-adjusted top federal tax rate for a wealthy household to all exempt bonds, and assumes the marginal investor buys bonds *outside* of their resident state. Since local coupon payments are exempt from state tax liabilities, it is optimal in most states for individual investors to invest in within-state issuers. Columns (3) and (4) therefore apply both federal and state effective tax rates to double-exempt bonds, and federal rates to federal-exempt only bonds, assuming the marginal investor resides in the state. That is, for this specification, we apply $1/(1-\tau_f^{deduct}-\tau_s^{deduct})$ to tax-exempt yields, which accounts for both local and state tax deductions.[69]

From columns (4) onward, we include state fixed effects that absorb 2020 state-tax adjustments (the coefficient difference between (4) and (5) is thus coming entirely from Federal-exempt only bonds that are eventually absorbed by month fixed effects). While perfect balance among the low-rated A and BBB group requires conditioning on these state fixed effects (to partial out volatility from the smaller sample), we interpret this as forcing crisper comparisons due to divergence in early 2020 (when low-rated issuer states were less balanced and trading volume low). These concerns however, should be minimal given that unconditional raw mean trajectories in the symmetric 100,000 population bandwidth (Figure 4) demonstrated that pre-trends were relatively parallel (especially in the run up to the first Federal announcement). All other ratings groups are otherwise unconditionally balanced, as indicated by the lack of point estimate statistical significance in the placebo period.

A.2.4 Placebo RD Plots Prior to Crisis

Next we show placebo RD estimate plots associated with prior regression output, using our preferred specification that includes state and month fixed effects to absorb volatility given the small sample size involved in cutting results by credit ratings bin. Figure A.4 shows yield RD plots for the placebo period by creditworthiness, and Figure A.5 shows employment RD plots for the placebo period using our difference in discontinuity strategy. These plots show relative smoothness across the cutoff, consistent with the tables showing no statistically significant RD effect prior to the policies being implemented.

[69] This is similar to using a state-deducted formula such as $1/(1-\tau_f)(1-\tau s)$, which does not adjust for local deductions. See, for example, Schwert (2017).

Figure A.4: RD Placebo Plots of MLF Access Effect on Yields by Credit Worthiness

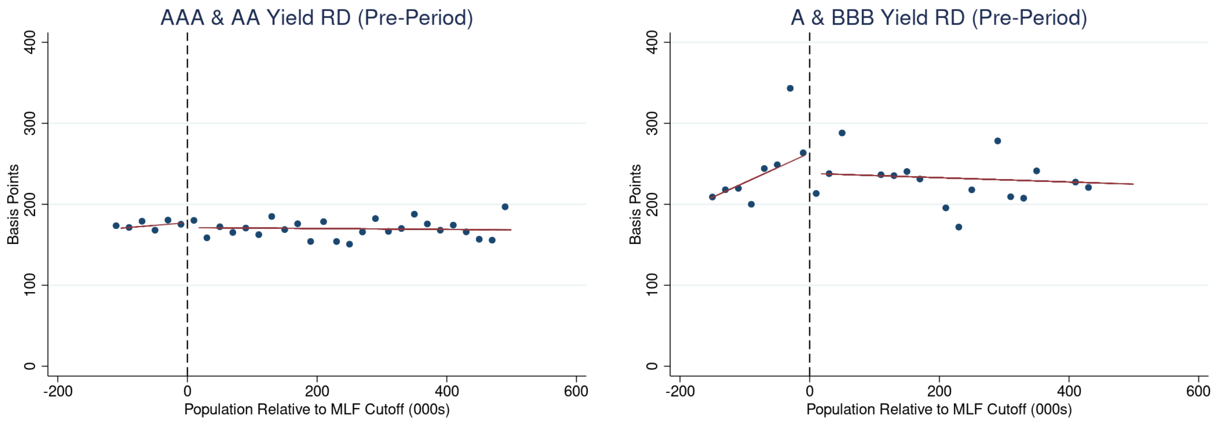

NOTES—Plots show regression slopes and intercepts from Equation 1 in the pre-period, overlaid on top of equally spaced pre-binned outcome data with a bin size of 20 (x-axis in thousands). Plots are shown over the optimal bandwidth selected using the IMSE-procedure, which produces (potentially) asymmetric optimal bandwidth boundaries for each sample. Plots correspond to Table 2, which includes state and calendar month fixed effects as controls. Plot thus residualizes all yields by state fixed effects (added back to their overall mean) prior to mean-collapsing by bin.

Figure A.5: RD Placebo Plots of MLF Access Effect on Public Sector Employment

(a) Δ Overall Employment, Pre-MLF (Placebo)   (b) Δ Services Employment, Pre-MLF (Placebo)

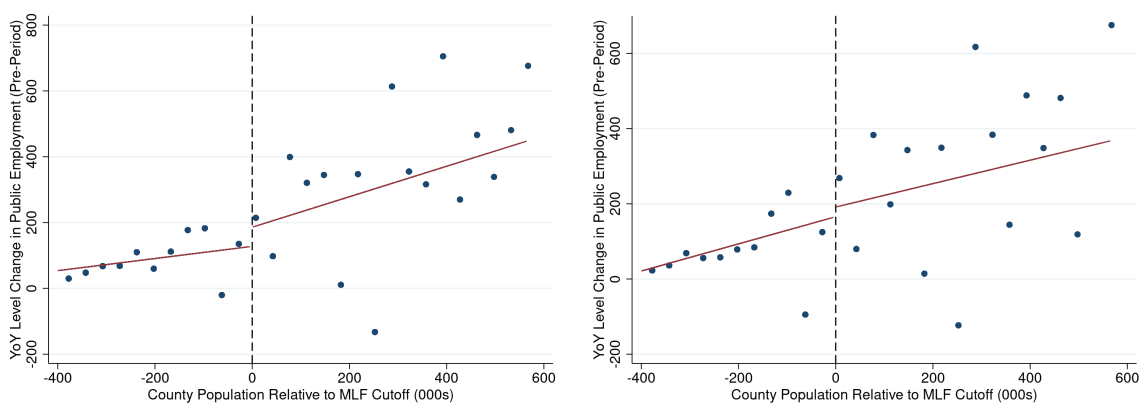

NOTES—Panels (a) and (b) show year-on-year differences in public employment for the months of January and February, 2020, relative to January and February, 2019. Results are shown for the differenced RD strategy discussed in the text, pooling all issuers in one sample. This RD does not yet implement an IMSE-optimal bandwidth. The baseline number of local employees in a given locality in 2019 was about 750 employees. Source: QCEW, May 26, 2021, Q4 release.

A.2.5 Manipulation Tests

In Figure A.6, we show formal McCrary (2008) running variable manipulation tests, using the Cattaneo et al. (2018) method of local polynomial density estimation with robust standard errors, where bandwidths for density tests are data-driven and thus do not have to span the whole support of the histogram.

Consistent with cutoffs being chosen using round-number heuristics and the MLF's regulatory population eligibility being lagged by 1 to 2 years, we find no evidence of cutoff manipulation by policymakers, nor cutoff-targeting based on the underlying distribution of of issuers. We also provide p-values from a density manipulation test associated with a discrete running variable (Frandsen, 2017), should lumpy population observations be interpreted as discrete rather than continuous. The Frandsen (2017) test of equality of projected intercepts at the cutoff produces to a p-value of 0.967 for counties, and 0.388 for cities, significantly away from conventional p-values that would reject equality.

Figure A.6: Manipulation Test for City and County Population Running Variables

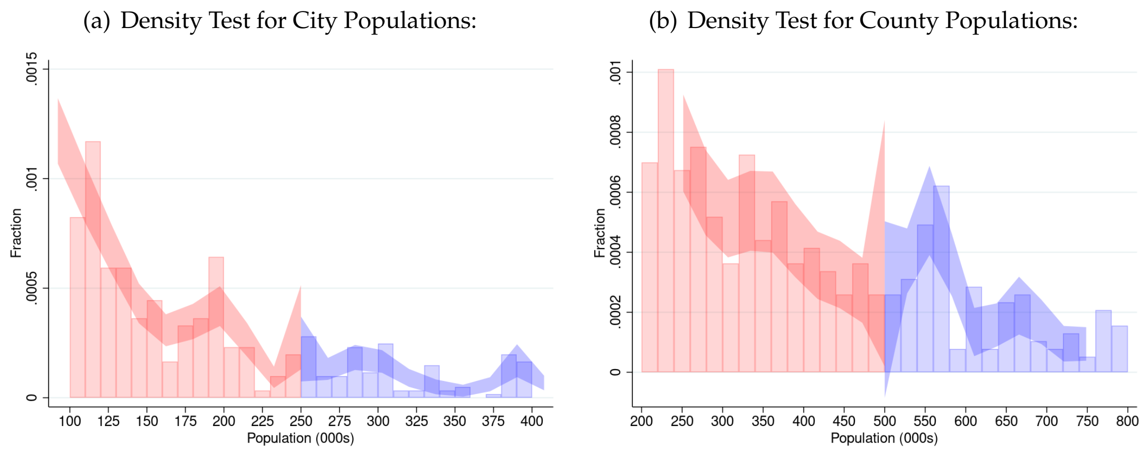

NOTES—Figure separately plots histograms of unique cities and counties on either side of their respective MLF population cutoffs, and estimates 95% confidence intervals from local polynomial fits on each side of cutoff using the Cattaneo et al. (2018) method. Significant overlap in confidence intervals signifies a passing test. City manipulation density test is overlaid on top of a histogram with binwidths of 10,000 while county manipulation density test is overlaid on top of histograms with binwidths of 25,000 (due to there being fewer counties than cities). Fractions are small because majority of cities and counties fall below 100,000 and 200,000 in population respectively.

A.2.6 Yields Trends Disaggregated by Rating

In Figure A.7 below, we show more detailed variation behind our aggregated low-rated yields results (from Figure 4 of the main draft), disaggregated by A and BBB issuers separately.

Figure A.7: Mean Yields within 100k Population of MLF Cutoffs, Separately for A and BBB

(a) A-Rated Issuers in Jan 2020: (b) BBB-Rated Issuers in Jan 2020:

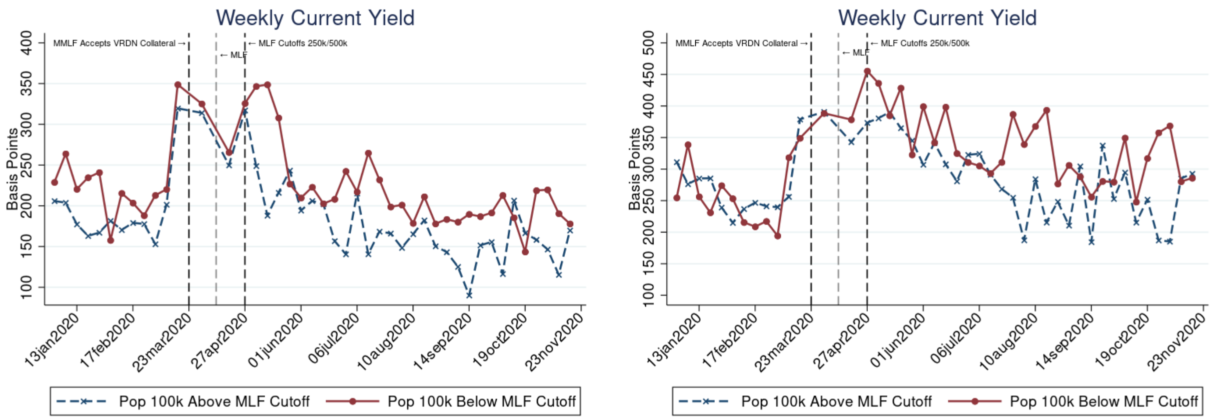

NOTES—Figure shows mean yields for two subgroups: eligible issuers with cutoff-relative populations between 0 and 100,000 inclusively (blue dashed series), and issuers with cutoff-relative populations greater than or equal to -100,000 and less than 0 (red solid series). See text for sample restrictions and subgroup definitions.

A.2.7 Kernel Weights

In Figure A.8, we present kernel weights generated through our optimal RD procedure, shown for exposition for our main yields estimates results for low-rated issuers.

Figure A.8: Triangular Kernel Weights for Asymmetric Optimal Bandwidth, Effects on Yields

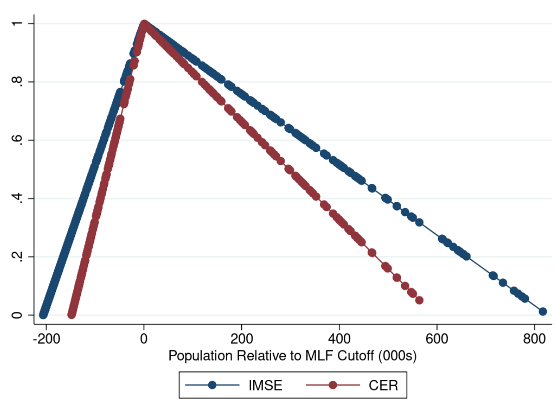

NOTES—Figure shows calculated triangular weights using asymmetric IMSE-optimal bandwidth on each side of the cutoff for main yields estimate, as well as an alternative weighting scheme based on the regression coverage error (CER), used for robustness. Weights and bandwidth are associated with the yields dependent variable in the pooled city and county sample in the post-period.

A.2.8 Trade Seasonality and Composition

In Figure A.9, we present the seasonality of new issuance for a standard expansion year like 2019, which demonstrates that effects on primary issuance are unlikely due to seasonal mean reversion in issuance.

Figure A.9: 2019 Seasonality in Primary Issuance

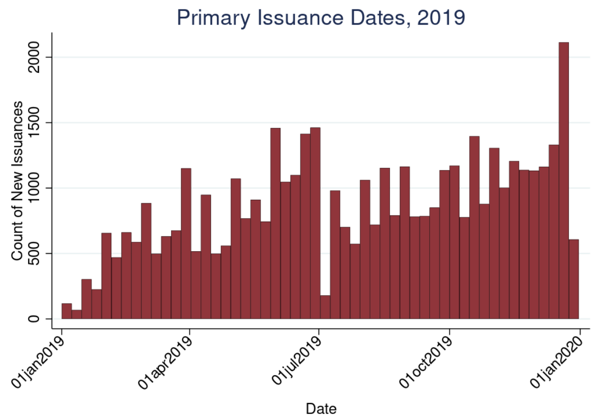

NOTES—Figure shows all CUSIPs tagged as new primary issuance based on initial offering date by calendar week in 2019. Source: Mergent

In Figure A.10, we present the share of A/BBB trading bonds 100,000 in population above and below the pooled cutoff that are general obligation (left) and revenue (right) bonds respectively. This plot shows that while initial trading patterns support a spike in revenue bond trends when issuers were revealed to be ineligible, this trend does not sustain. In conjunction with earlier work demonstrating balance, composition differences across the cutoff immediately following the MLF announcement and unlikely to be the main factor driving these results.

Figure A.10: Composition Sensitivity: GO and RB A/BBB Trends

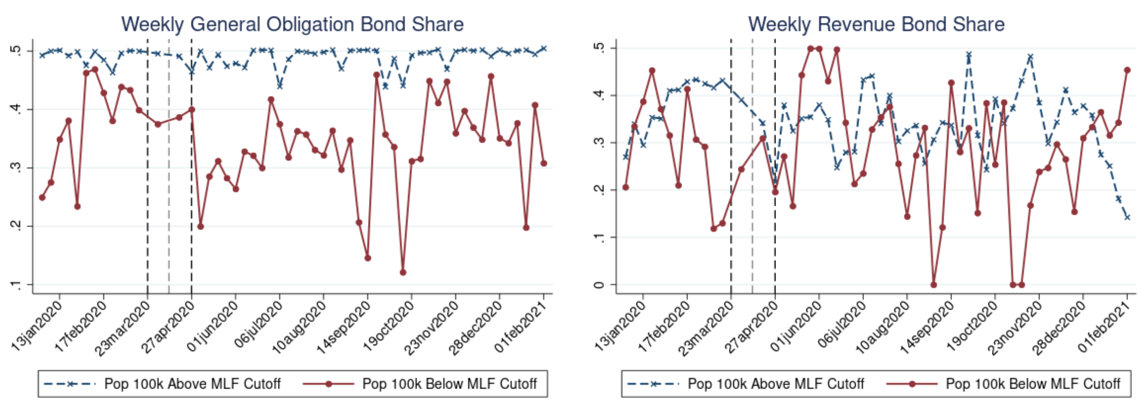

NOTES—Figure shows shares of all trading CUSIPs for issuers 100,000 above and below the cutoffs.

A.2.9 Power Calculation for City Running Variable

In Table A.2 below, we provide power tests for various effect sizes to assess whether we can detect reliable effect sizes (using the common convention of high power exceeding 0.80) for the 21 cities that are above the 250,000 cutoff (along the city running variable).

Table A.2: Power Tests on City Running Variable (Cutoff = 250,000)

Level of Employment	$\tau = 1,000$	$\tau = 2,000$	$\tau = 3,000$	$\tau = 4,000$
Employment, City Only Treated	0.10	0.27	0.53	0.78
– Services Employment, City Only Treated	0.11	0.29	0.56	0.81
Employment, City & County Treated	0.30	0.81	0.99	1.00
– Services Employment, City & County Treated	0.22	0.67	0.95	1.00
YoY Level Change in Employment	$\tau = 100$	$\tau = 200$	$\tau = 300$	$\tau = 400$
Employment, City Only Treated	0.09	0.22	0.42	0.65
– Services Employment, City Only Treated	0.07	0.14	0.25	0.40
Employment, City & County Treated	0.16	0.47	0.81	0.97
– Services Employment, City & County Treated	0.12	0.33	0.63	0.87
YoY Percent Change in Employment	$\tau = 1\%$	$\tau = 2\%$	$\tau = 3\%$	$\tau = 4\%$
Employment, City Only Treated	0.10	0.25	0.48	0.72
– Goods Employment, City Only Treated	0.05	0.05	0.06	0.06
– Services Employment, City Only Treated	0.06	0.11	0.18	0.28
Employment, City & County Treated	0.20	0.61	0.92	0.99
– Goods Employment, City & County Treated	0.05	0.07	0.09	0.12
– Services Employment, City & County Treated	0.11	0.28	0.55	0.79

NOTES—We run power tests of our RD estimator for our city sample with a symmetric bandwidth of 250,000 in population using the *rdpow* and *rdsampsi* packages (Cattaneo et al. (2019b)). We omit the goods employment rows in the first two panels because our detected RD effects are lower in magnitude than for services and overall employment. Please see Cattaneo et al. (2019b) for more details. To perform the conventional power calculations above, we assume there is ignorable "smoothing bias" on either side of the cutoff. We can therefore interpret these calculations as upper bounds.

While we conclude from these results that only the largest effect sizes can be reliably estimated with full power, for full transparency, we provide the corresponding RD estimation results in Table A.3 with this caution in mind. To this end, we generate city employment as the county-weighted employment across counties spanned by both eligible and ineligible cities. If one were to ignore power concerns and take these results at face value, we are unable to detect any positive effects using this sample.

Table A.3: City-Level RD Estimates of MLF Access on Public Sector Employment

	Emp. (1)	Emp. (2)	Δ Emp. (3)	Δ Emp. (4)	% Δ Emp. (5)	% Δ Emp. (6)	N (fixed-bwdth)
a. Pooled Post:							
Overall Employment	-948	-1,045	-52	-49	-1.21	-0.56	29,609
	(694)	(781)	(94)	(94)	(0.89)	(0.72)	
– Goods Employment	-87**	-113***	-1	0	-4.55	-2.29	5,370
	(35)	(41)	(3)	(3)	(5.00)	(4.35)	
– Services Employment	1	-597	-98	-62	-1.01	0.37	19,669
	(826)	(943)	(124)	(120)	(1.03)	(0.74)	
b. Pooled Pre (Placebo):							
Overall Employment	-740	-867	54**	52**	1.04***	1.14***	29,643
	(725)	(819)	(24)	(23)	(0.40)	(0.33)	
– Goods Employment	-83**	-110***	-1	-1	0.22	-3.06	5,260
	(35)	(42)	(2)	(2)	(3.12)	(2.56)	
– Services Employment	281	-229	69**	65**	0.86**	1.20***	19,606
	(880)	(966)	(28)	(28)	(0.42)	(0.39)	
Month FEs		X					
State FEs		X		X		X	
Control Mean (post): Employment	10,067	9,487	-779	-730	-6.64	-7.09	
Control Mean (post): Goods Employment	154	134	-5	-4	-5.28	-4.92	
Control Mean (post): Services Employment	8,555	8,022	-701	-653	-6.94	-7.75	

NOTES—Table presents RD estimates of employment measure in column header and sample period. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). The Control Mean denotes the mean of that characteristic immediately to the left of relative population 0. Each regression uses a fixed bandwidth with a triangular kernel to weight observations (see Section 4 for details). Percents in columns (5) and (6) range from 0 to 100. Standard errors are conservatively clustered by population relative to the cutoff (the running variable); as such, inference is drawn across unique issuers. *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.1$.

A.2.10 Full Baseline Summary Statistics

In Table A.4, we provide a full list of baseline summary statistics prior to March 23, including the entire calendar year of 2019. These do not restrict to issuers within 100,000 of the cutoff, and instead are meant to report endogenous differences between the samples of eligible and non-eligible issuers were they to be compared naively, for example in a difference-in-differences analysis.

Table A.4: Full Baseline Summary Statistics, January 1, 2019 to March 23, 2020

	MLF Eligible		MLF Ineligible		MLF Eligible - Ineligible
	Mean/SD	# Observations	Mean/SD	# Observations	Δ/SE
	(1)	(2)	(3)	(4)	(5)
A. MSRB-Bloomberg Trade-Level Data					
Coupon Rate (b.p.)	430.3	669,311	384.2	844,395	46.2***
	[135.1]		[116.6]		(0.21)
Security Price (per 100 par)	108.0	671,659	105.8	844,500	2.23***
	[8.95]		[7.45]		(0.014)
Current Yield (b.p.)	203.0	620,560	209.6	825,352	-6.68***
	[83.2]		[80.3]		(0.14)
Δ Yield (Feb20-Jan20)	-0.066	281,596	-0.14	153,899	0.078***
	[1.21]		[0.95]		(0.0033)
Δ Yield YoY (Jan20-Jan19)	-0.89	262,725	-1.11	124,523	0.22***
	[4.28]		[1.17]		(0.0090)
Δ Yield YoY (Feb20-Feb19)	-1.00	242,935	-1.07	110,336	0.074***
	[1.56]		[1.07]		(0.0045)
Amount Outstanding (MM)	2328.8	671,659	144.5	844,500	2184.3***
	[2982.3]		[224.1]		(3.65)
Maturity Size (MM)	3542.0	671,659	215.7	844,500	3326.2***
	[4209.3]		[305.5]		(5.15)
Tenor of Bond (Years)	14.8	671,464	13.2	844,500	1.55***
	[7.64]		[6.98]		(0.012)
Remaining Duration of Bond (Years)	9.22	671,464	8.72	844,500	0.51***
	[7.02]		[6.74]		(0.011)
Market Share of Issuer	0.95	671,659	0.064	844,500	0.89***
	[1.06]		[0.11]		(0.0013)
Number of Securities by Issuer	306.4	671,659	123.0	844,500	183.3***
	[210.6]		[107.5]		(0.28)
Par Traded (1000s)	290.9	671,659	95.4	844,500	195.5***
	[1915.2]		[579.9]		(2.42)
S&P Ratings (1-7 scale)	5.67	575,172	5.83	624,205	-0.16***
	[0.86]		[0.65]		(0.0014)
Moody's Ratings (1-7 scale)	5.78	557,662	5.88	474,590	-0.099***
	[0.90]		[0.92]		(0.0018)
Fitch Ratings (1-7 scale)	5.65	424,063	5.63	149,058	0.018***
	[1.03]		[0.79]		(0.0026)
Time of Day of Trade (minute)	770.6	671,659	776.1	844,500	-5.53***
	[131.7]		[132.2]		(0.22)
B. QCEW Month-County Loc. Gov. Emp. Data					
Δ Employment	627.8	228	19.9	5,250	607.9***
	[861.7]		[99.5]		(57.1)
Δ Goods Employment	9.77	130	0.20	2,402	9.57***
	[35.4]		[8.89]		(3.11)
Δ Service Employment	1509.8	228	23.1	5,250	1486.7**
	[10905.4]		[414.9]		(722.2)

NOTES—Table presents summary statistics for key variables of the analysis using a longer period prior to the first Federal intervention in the municipal market, along with a two-sided t-test with heteroskedastic-robust standard errors). See Appendix A.3 for further details on variable definitions and construction.

A.3 Additional Background and Data Construction Details

A.3.1 Primary Issuance during the COVID-19 Pandemic

Figure A.11: State and Local Government New Issuance Shortfalls during COVID-19

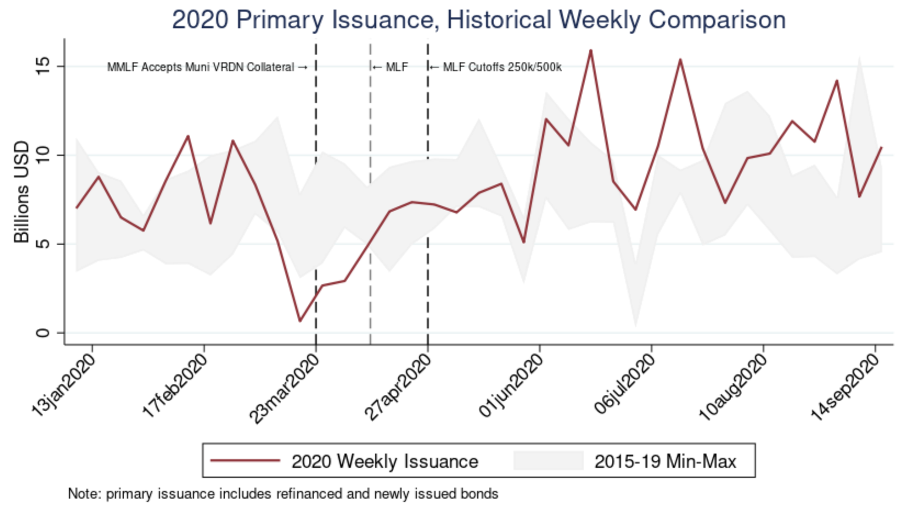

NOTES—Figure shows weekly total primary issuance in the municipal bond market throughout the 2020 calendar year, compared to the five-year historical maximum and minimum issuance in the same week. Weekly means are aligned beginning-of-week on Mondays (when most Federal interventions are announced), and may include refinanced bonds which appear as new CUSIPs. *Source:* Bloomberg, originally calculated in Cipriani et al. (2020a).

A.3.2 MLF Pricing Grid

Figure A.12: MLF Pricing Grid

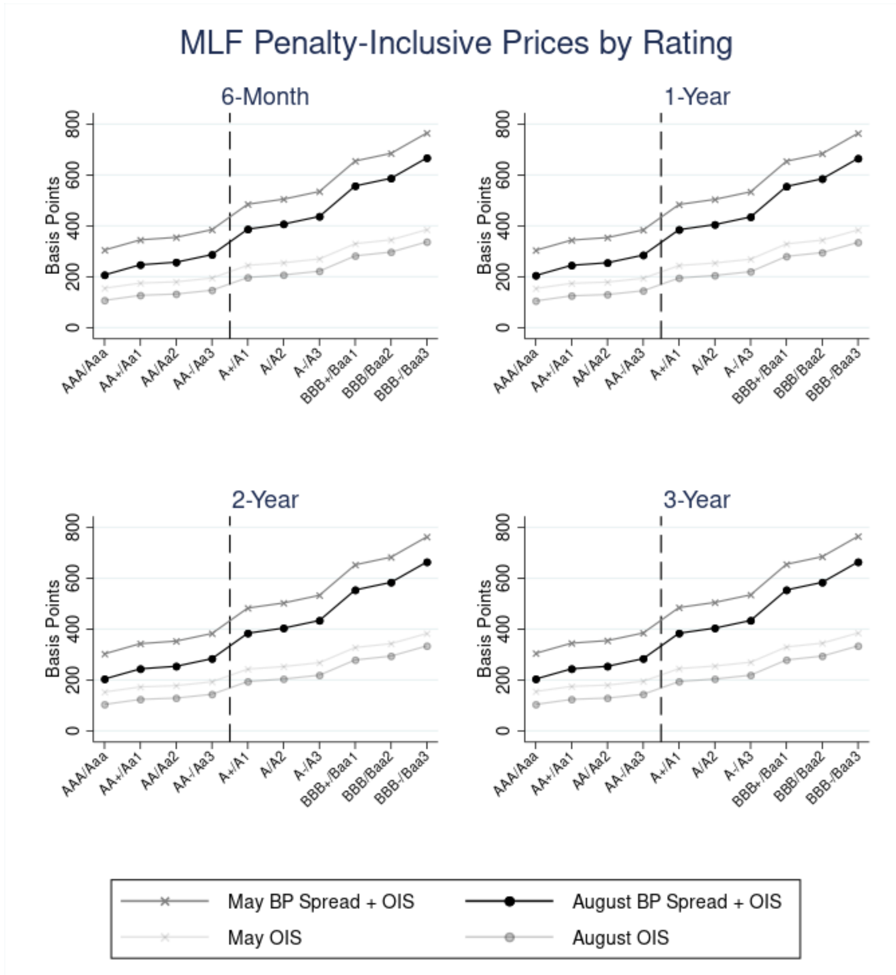

NOTES—Figure shows penalty pricing grid by issuer ratings when it was announced on May 11, 2020, and revised on August 11, 2020, for different loan term lengths. OIS = overnight index swap rate, meant to represent prevailing private market interest rates, here pulled on May 27, 2020, and August 12, 2020, respectively. Dotted line indicates distinction between "high" and "low" rated issuers as defined in the paper. Pricing schedules are similar with slight variations for different tenors. See MLF Pricing Term Sheet for further details.

A.3.3 Matching MSRB Issuer Names to Census Populations

Our goal is to assign a single Census population to any city, county, or state municipal bond issuer that has traded on the secondary market at any point since Jan. 1, 2019, including all MLF eligible and ineligible issuers. We start with the term-sheet referenced U.S. Census Bureau datasets that include 2018 city populations greater than 50,000, and 2019 county populations greater than 50,000.[70] This provides populations for larger issuers, however we also desire to match populations to smaller issuers

[70] U.S. Census Bureau, "Annual Estimates of the Resident Population: April 1, 2010 to July 1, 2018" for cities, as of April 6, 2020; and U.S. Census Bureau, "Population, Population Change, and Estimated Components of Population Change: April 1, 2010 to July 1, 2019" for counties, as of April 6, 2020.

so that they may contribute to the RD running variable further away from the cutoff. Toward this end, we complement our analysis with the full population Census file "sub-est2019_all.csv", a place-level dataset with both 2019 and 2018 populations for localities of *all* populations.[71] This file contains a place description such as the city or county's name (which is always followed by the type suffix, such as "CITY" or "COUNTY"), state, census place code, and geographic code level (*sumlev*), allowing us to isolate the administrative name and its geographic level.

We desire the dataset be uniquely identified by city (county) name and city (county) state, so reshape wide the 2018 (2019) populations by *sumlev*, which produces a maximum of 6 potential population measures for each geographic code, corresponding to the number of unique values *sumlev* can take:

1. Minor Civil Divisions (MCDs)[72]
2. MCD "parts"
3. County place parts
4. Incorporated places (cities, towns, boroughs, villages)
5. Consolidated cities
6. Consolidated city parts

The data separately document localities whose "part" or "balance" spills over into another locality, which we subsequently drop (i.e. we drop (2) and (6)).[73] We then designate a rule to choose among these candidate populations. Most places have one or two population measures, but they will not have all six. If all non-missing populations are the same, or if there is only one population, we use that population. If any of the populations are different (0.1% of localities (27 observations; 3 of which ultimately do not merge to MSRB-Bloomberg)), we flag and omit them the analysis.

Finally, we go through a multi-step process to merge MSRB issuer names (from their security descriptions) to Census place names. This involves flagging duplicate MSRB place names within state which arise erroneously from our cleaning algorithm, and remapping them based on their original security description to their appropriate population. For example, "SPRINGFIELD TOWNSHIP, NJ" in the Census may correspond to "SPRINGFIELD, NJ" in MSRB trade data; by removing the township

[71] The sub-est2019_all.csv file can be found here: https://www2.census.gov/programs-surveys/popest/datasets/2010-2019/cities/totals/.

[72] States in New England, New York, and Wisconsin, all classify towns as MCDs.

[73] A "balance" is a consolidated city minus the semi-independent incorporated places located within the consolidated city (overlapping service populations which would be double counted if we did not drop them). Incorporated places can cross both county and Minor Civil Division (MCD) boundaries. In such cases a separate record indicates the population estimates for the part of a place in each of its parent counties or MCDs. For such records, the place name is sometimes followed by the designation "pt" (which stands for part), allowing us to isolate these cases.

suffix, we can match these. But doing this universally risks false positives in other cases: for example, the Census might contain "HEMPSTEAD VILLAGE, NY", "HEMPSTEAD TOWN, NY", and "HEMPSTEAD CITY, NY" (all separate localities), while MSRB may only contain one HEMPSTEAD, NY. Removing suffixes here would lead to duplicate matches. We thus only remove such suffixes when the resulting merge is one-to-one, causing a loss rate of 4% of unmatched trades.

We also consider special treatment for matching fully consolidated cities (e.g. City and County of San Francisco) which have only one set of issuers at either the city or county level, and partially consolidated city-counties (e.g. Miami City vs. Dade County) which may have different revenue sources. We follow the MLF term sheet in assigning these cases to city and/or county eligibility lists based on the revenue sources of underlying issuers. Lastly, we manually flag potential counties that issue on behalf of cities ("downstreamers") based on their detailed issuer and security description. Among issuers with a reliable MSRB issuer name and actively trading post-2019, our final cleaning procedure results in 6,143 city (1,880 county) BaseCUSIPs (unique at the issuer level) that match a population to MSRB data, and 361 city (41 county) issuers that do not retrieve a match.

MSRB Data Cleaning

Since January 1998, MSRB has required registered dealers to report all municipal bond transactions. The trade record includes information about CUSIP, date and time of the trade, price and yield, maturity, coupon, and a flag whether the dealer bought from a customer, sold to a customer, or whether it is an interdealer trade. In January 2015, MSRB started to publicly disseminate those transactions with up to a 15 minutes delay.

We first keep unique observations at the CUSIP—trade_date—RTRS_Control_Number level, using the RTRS_PUBLISH_DATE and RTRS_PUBLISH_TIME variables to ensure that duplicates are not arising from missing data. Then, in addition to the sampling restrictions applied to our RD analysis, for the yield decomposition sample, we further apply the following conditions:

1. Delete CUSIPs with:
   (a) missing coupon and maturity information for all trades
   (b) a listed coupon greater than 20%
   (c) a listed maturity over 100 years
   (d) fewer than 10 trades in the entire sample
2. Delete transactions where:
   (a) the price is less than 50 (i.e., 50% of face value)

(b) the price is greater than 150 with a short time to maturity

(c) trade date is after the maturity date of the bond

Mergent Data Cleaning

The key characteristics from Mergent that we use include CUSIP, issuer name, offering amount, source of funds and use of proceeds, bond credit rating as rated by S&P, Moody's, and/or Fitch, coupon type (fixed, variable, or zero), the tax status of the coupon payments, callability and first call date, insurance status and the identity of the insurer, and pre-refunding status and timing. The Mergent data ratings provide a longer time series of ratings data relative to the Bloomberg ratings data used in the RD analysis. A few of the variables require some adjustments. Specifically:

Coupon Type: Mergent's variable *coupon_code* indicates the coupon type. However, for fixed-rate bonds issued at a discount or at a premium, it only indicates OID and OIP respectively. Most of these bonds are fixed-rate, nevertheless for these bonds we also use the VARRATE table to determine the coupon type.

Ratings: There are also duplicate observations at the CUSIP—rating_type—rating_date level, sometimes with different ratings values. Rating agencies submit revisions of ratings that are then posted in the Mergent database as a new rating instead of a revision to a current rating. The data provider is in the process of correcting that in its database, so for such observations we use ratings data from Bloomberg.

A.3.4 NRSRO Ratings Concordance and Plurality Ratings

Each of S&P, Moody's, and Fitch, maintain separate ratings systems for long and short-term bonds. S&P and Moody's also use separate systems for short-term municipal note ratings. We map of each these to 8 aggregate ratings bins that are guided by S&P's convention: AAA, AA, A, BBB, BB, B, C, D. To do so, we use the following concordances, which were developed manually in consultation with a number of sources.[74] The resulting columns ending in *_agg* form the basis of our plurality ratings, calculated as the plurality across long-term, short-term, and muni-note bonds based on their "aggregate" ratings.

[74] These include S&P reference, Moody's reference, and Fitch reference. Some disaggregated ratings are identified from MSRB trades rather than the NRSRO's themselves, however have natural mappings.

Figure A.13: Ratings Concordances to Aggregated Credit Rating Bins

Fitch, Long-Term Ratings			
rank	*fitch_rating_lt*	*fitch_rating_lt_agg*	*fitch_rating_desc*
1	AAA+	AAA	Prime
2	AAAe	AAA	Prime
3	AAApre	AAA	Prime
4	AAA	AAA	Prime
5	AAA-	AAA	Prime
6	AA+e	AA	High Grade
7	AA+	AA	High Grade
7	AAe	AA	High Grade
8	AA	AA	High Grade
9	AA-e	AA	High Grade
10	AA-	AA	High Grade
11	A+e	A	Upper Medium Grade
12	A+	A	Upper Medium Grade
12	Ae	A	Upper Medium Grade
13	A	A	Upper Medium Grade
14	A-	A	Upper Medium Grade
15	BBB+	BBB	Lower Medium Grade
16	BBB	BBB	Lower Medium Grade
17	BBB-	BBB	Lower Medium Grade
18	BB+	BB	Non-Investment Grade Speculative
19	BB	BB	Non-Investment Grade Speculative
20	BB-	BB	Non-Investment Grade Speculative
21	B+	B	Highly Speculative
22	B	B	Highly Speculative
23	B-	B	Highly Speculative
24	CCC	C	Extermeley Speculative
25	DDD	D	In Default
26	DD	D	In Default
27	D	D	In Default

Fitch, Short-Term Ratings		
rank	*fitch_rating_st*	*fitch_rating_st_agg*
1	F1+e	AA
2	F1+	AA
3	F1	A
4	F2	BBB
5	F3	BBB
6	B	B
7	C	C
8	D	D
9	/	D

Moody's, Long-Term Ratings			
rank	*moody_rating*	*moody_rating_lt_agg*	*moody_rating_desc*
1	#Aaa	AAA	Prime
2	Aaa	AAA	High grade
3	Aa1	AA	High grade
4	Aa2	AA	High grade
5	Aa3e	AA	High grade
6	Aa3	AA	High grade
7	A1	A	Upper medium grade
8	A2	A	Upper medium grade
9	A3	A	Upper medium grade
10	Baa1	BBB	Lower medium grade
11	Baa2	BBB	Lower medium grade
12	Baa3	BBB	Lower medium grade
13	Ba1	BB	Non-investment grade speculative
14	Ba2	BB	Non-investment grade speculative
15	Ba3	BB	Non-investment grade speculative
16	B1	B	Highly speculative
17	B2	B	Highly speculative
18	B3	B	Highly speculative
19	Caa1	CCC	Substantial risks
20	Caa2	CCC	Extremely speculative
21	Caa3	CCC	Default imminent with little prospect for recovery
22	Ca	C	Default imminent with little prospect for recovery
23	C	D	In default
24	/	D	In default

Moody's, Short-Term Ratings		
rank	*moody_rating_st*	*moody_rating_st_agg*
1	P-1	A
2	P-2	BBB
3	P-3	BBB

Moody's, Muni-Note Ratings		
rank	*moody_rating_muninotes*	*moody_rating_muninotes_agg*
1	MIG1	A
2	VMIG1	A
3	MIG2	A
4	VMIG2	A
5	VMIG3	BBB
6	MIG3	BBB
7	SG	D

S&P, Long-Term Ratings			
rank	*sp_rating_lt*	*sp_rating_lt_agg*	*sp_rating_desc*
1	AAA+	AAA	Prime
2	AAA	AAA	Prime
3	AAA-	AAA	Prime
4	AA+	AA	High Grade
5	AA	AA	High Grade
6	AA-	AA	High Grade
7	A+	A	Upper Medium Grade
8	A	A	Upper Medium Grade
9	A-	A	Upper Medium Grade
10	BBB+	BBB	Lower Medium Grade
11	BBB	BBB	Lower Medium Grade
12	BBB-	BBB	Lower Medium Grade
13	BB+	BB	Non-Investment Grade Speculative
14	BB	BB	Non-Investment Grade Speculative
15	BB-	BB	Non-Investment Grade Speculative
16	B+	B	Highly Speculative
17	B	B	Highly Speculative
18	B-	B	Highly Speculative
19	CCC+	CCC	Substantial Risks
20	CCC	CCC	Extermeley Speculative
21	CCC-	CCC	Default Imminent with little prospect for recovery
22	CC+	CC	Default Imminent with little prospect for recovery
23	CC	CC	Default Imminent with little prospect for recovery
24	C-	C	Default Imminent with little prospect for recovery
25	SD	D	In Default
26	D	D	In Default

S&P, Short-Term Ratings		
rank	*sp_rating_st*	*sp_rating_st_agg*
1	A-1+	A
2	A-1	A
3	A-2	BBB
4	A-3	BBB
5	B	B
6	C	C
7	/	D
8	D	D

S&P, Muni-Note Ratings		
rank	*sp_rating_muninotes*	*sp_rating_muninotes_agg*
1	SP-1+	A
2	SP-1	A
3	SP-2	BBB
4	SP-3	BB
5	D	D

A.3.5 List of A and BBB Issuers within 100k of Cutoff

Table A.5: Low-Rated (A and BBB) Issuers within 100k of MLF Cutoff

BaseCUSIP	Issuer Name	State	Issuer Type	Number of Cusips	Above100
650367	Newark	NJ	Local Government	89	1
119677	Buffalo	NY	Local Government	81	1
889278	Toledo	OH	Local Government	73	1
86607C	Summit County	OH	County Government	58	1
46360R	Irvine	CA	Local Government	37	1
46360T	Irvine	CA	Local Government	35	1
660393	North Las Vegas	NV	Local Government	34	0
825434	Shreveport	LA	Local Government	31	0
549310	Lucas County	OH	County Government	29	0
463805	Irving	TX	Local Government	28	0
759861	Reno	NV	Local Government	27	1
79307T	St Paul	MN	Local Government	25	1
70643U	Pembroke Pines	FL	Local Government	24	0
10047	Akron	OH	Local Government	22	0
85233S	St Louis	MO	Local Government	22	1
555542	Macon Bibb County	GA	Local Government	21	0
270764	East Baton Rouge Parish	LA	County Government	20	0
79488C	Salinas	CA	Local Government	19	0
759829	Reno	NV	Local Government	16	1
743787	Providence	RI	Local Government	15	0
133402	Cameron County	TX	County Government	14	0
86607D	Summit County	OH	County Government	13	1
537374	Little Rock	AR	Local Government	10	0
534310	Lincoln	NE	Local Government	9	1
43615F	Hollywood	FL	Local Government	9	0
613549	Montgomery County	OH	County Government	8	1
650366	Newark	NJ	Local Government	8	1
55553N	Macon Bibb County	GA	Local Government	8	0
690278	Overland Park	KS	Local Government	7	0
66041H	North Las Vegas	NV	Local Government	7	0
79164T	St Louis	MO	Local Government	6	1
743940	Providence	RI	Local Government	5	0
702521	Pasco County	FL	County Government	4	1
29634D	Escondido	CA	Local Government	4	0
344610	Fontana	CA	Local Government	2	0
607715	Modesto	CA	Local Government	2	0
928844	Volusia County	FL	County Government	2	1
873477	Tacoma	WA	Local Government	2	0
696712	Palmdale	CA	Local Government	1	0
759830	Reno	NV	Local Government	1	1
51672	Aurora	IL	Local Government	1	0
768861	Riverside	CA	Local Government	1	1
35895	Anne Arundel County	MD	County Government	1	1

NOTES—Data contain 43 unique issuers with 831 unique outstanding bonds or notes. Each bond contributes additional variation to RD estimates, however inference only leverages differences between issuers (standard errors clustered by relative population). 20 issuers are above the MLF cutoff, while 23 are below. Irvine, Newark, North Las Vegas, Reno, and St. Louis are listed twice, but these reflect different issuers within those cities (for example, economic development corporations may issue their own independent debt).

A.4 Additional Results

A.4.1 RD Effects on Credit Downgrade Probability

Here we explore the possibility that ratings agencies disproportionately downgraded issuers that narrowly miss the cutoff (or do not downgrade issuers if supported by the MLF), and investors either price MLF access similarly to NRSROs, or in response to them. While this is a very specific type of credit risk measure, we expect it to be broadly correlated with default risk. Toward this end, we first construct a balanced month-by-bond panel, and count how many times a bond received a downgrade from any of the three major NRSROs relative to its rating in the previous month. Analogous to earlier event study plots, in Figure A.14 we focus on the arbitrary symmetric bandwidth 100,000 in population around each cutoff, and report cumulative downgrades separately by the underlying issuer's plurality rating in January 2020.

Figure A.14: Credit Rating Downgrades within 100k Population of MLF Eligibility Cutoffs

(a) Cumulative Downgrades, All Ratings

(b) Cumulative Downgrades, Low-Rated (A & BBB)

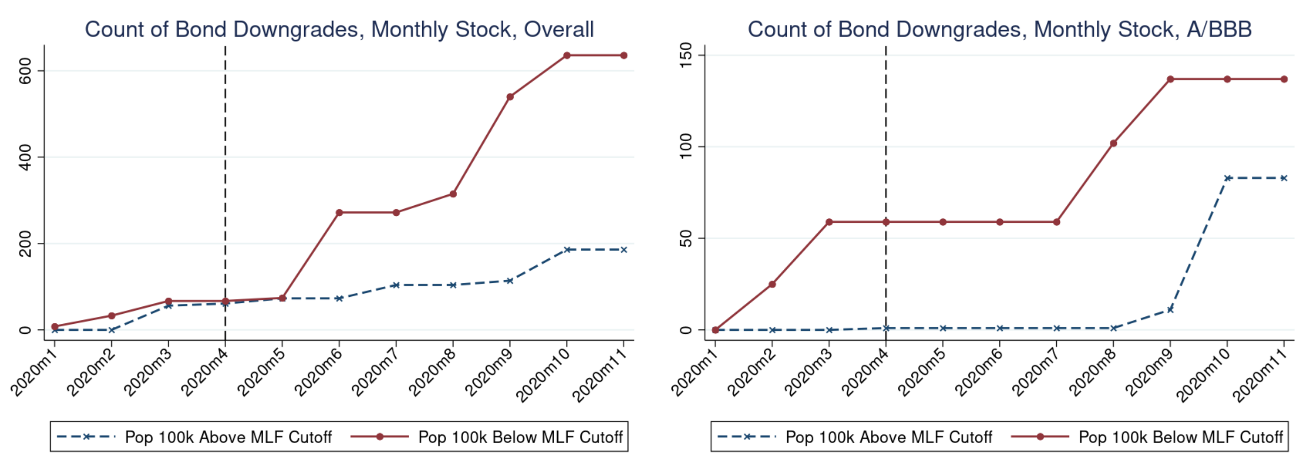

NOTES—Panels (a) and (b) count the number of times a bond is downgraded relative to the prior month for eligible issuers with cutoff-relative populations between 0 and 100,000 inclusively (blue dashed series), and issuers with cutoff-relative populations greater than or equal to -100,000 and less than 0 (red solid series).

Panel (a) shows that *independent* of MLF eligibility, while the number of downgraded bonds moved in tandem throughout April and May, by June of 2020, issuers that were revealed to be ineligible and therefore without an emergency liquidity option (red solid series), were downgraded at a faster rate than MLF-eligible issuers' bonds. The number of MLF-ineligible issuer bond downgrades rose from roughly 100 to 600 by November of 2020, whereas downgrades were rarer for issuers that had MLF optionality.[75] Panel (b) shows a noisier but related pattern among low-rated issuers: in August, low-rated ineligible

[75] This pattern also persists when looking at the cumulative share of total bonds that are newly issued, which accounts for concerns that effects may be mechanically related to a greater number of issuers below the cutoff.

issuers began to experience more downgrades (rising to roughly 150 downgrades by November), and accounted for about 25% of the aggregate downgrade pattern—a large share relative to the number of total issuers and bonds that are in the sample. One interpretation for why there may have been downgrades across the ratings distribution, but these were only priced into yields on the low end, is that underlying mal-priced credit risk may be particularly exposed during a crisis for issuers that are closer to the default margin. Both credit rating agencies (NRSROs) and private investors may have responded to the newly revealed risks differently, depending on whether a given issuer had an emergency credit-risk sharing option, such as access to the MLF.

While these patterns are striking, such rare events are difficult to detect statistically due to the majority of muni bonds exhibiting no change in ratings over the analysis period. In Table A.6 below, we formalize this notion by estimating whether a bond was downgraded (relative to the previous month) in the post-MLF or pre-MLF period, and how many bonds on average were downgraded during this period. Similar to our analysis of primary issuance, here we collapse the data to two observations per bond—one in the pre-MLF period and one in the post-MLF period.

Table A.6: RD Estimates of MLF Access on Credit Rating Downgrades

	Discontinuity	Standard Error	Control Mean	N (Fixed-bwdth)
a. Pooled Post:				
Pr(Downgrade), Overall (may20-nov20)	-0.02	0.04	0.06	142,155
Pr(Downgrade), A & BBB (may20-nov20)	0.06	0.15	0.15	12,132
Number of Downgrades, Overall (may20-nov20)	-0.03	0.04	0.06	142,155
Number of Downgrades, A & BBB (may20-nov20)	0.06	0.15	0.15	12,132
b. Pooled Pre (Placebo):				
Pr(Downgrade), Overall (jan20-feb20)	-0.02	0.01	0.02	120,786
Pr(Downgrade), A & BBB (jan20-feb20)	-0.05	0.04	0.05	10,313
Number of Downgrades, Overall (jan20-feb20)	-0.02	0.01	0.02	120,786
Number of Downgrades, A & BBB (jan20-feb20)	-0.05	0.04	0.05	10,313

NOTES—Table presents RD estimates of the probability that a bond was downgraded in the period of interest (post-MLF or pre-MLF), and the number of downgrades each bond may have received in the same period. Each row corresponds to a separate regression with that characteristic as the dependent variable. The discontinuity measures the jump in the regression function at population eligibility 0 (relative to cutoff). The Control Mean denotes the mean of that characteristic immediately to the left of relative population 0. Sample sizes vary depending on the bandwidth used. Standard errors are clustered by population relative to the cutoff (the running variable). See text for further sample restrictions. *** $p \leq 0.01$, ** $p \leq 0.05$, * $p \leq 0.1$.

While the table reveals that in the post-period, bonds of all ratings have negative coefficients associated with a lower downgrade rate for eligible issuers, these cannot be detected statistically against the vast majority of bonds that experience no change in ratings. For low-rated issuers, we in fact report positive point estimates, which suggests that the number of downgrades among MLF eligible issuers represented a larger share of total bonds than downgrades among ineligible issuers. Nevertheless, neither is statistically different from zero, and both are estimated with noise. We thus cautiously interpret these results in conjunction with the overall plots, as providing modest evidence that MLF-eligible issuers had lower downgrade rates than observationally equivalent ineligible issuers.

A.4.2 Liquidity-Credit Risk Decomposition Details

In this section we discuss the various liquidity measures that enter our spreads decomposition. As with our main cities and counties sample, we source municipal bond transaction-level prices and yields from MSRB, but here extend our sample back to January 2013, which is the earliest date for which we have transaction data that includes all three trade type indicators (i.e., D for interdealer trade, P for a dealer purchase from a client, and S for a dealer selling to a client) to properly estimate liquidity measures. Each trade record includes the trade time, bond price, yield, par value traded, and an indicator for whether the trade was a customer purchase from a broker-dealer, customer sale to a broker-dealer, or an interdealer trade. We follow the cleaning procedure of MSRB transactions data that is common in the literature (see e.g., Green et al., 2010), and exclude bonds with missing coupon and maturity information (which effectively removes variable rate bonds), a listed coupon greater than 20%, or a listed maturity over 100 years. These are likely to be data errors.

To mitigate the effect of outliers, we further remove CUSIPs that are traded less than 10 times over the entire sample, trades that: occur on weekends and SIFMA holidays, have a price less than 50 or greater than 150, a when-issued flag equal to "Y", a non-missing Brokers Broker Indicator, occur within 90 days of issuance, and any trade date after the maturity date of the bond. Next, we use bond characteristics from Bloomberg and Mergent's Municipal Bond Securities Database, including maturity, offering amount, coupon rate, state of issuance, issue series, issuance date, whether it was a negotiated or competitive sale, whether the bond is general obligation, revenue, insured or callable, and ratings from S&P, Moody's, and Fitch, to further filter the data. We exclude variable rate bonds, insured bonds, AMT bonds, and bonds that are only subject to federal taxes (for the decomposition), as these characteristics will add noise to spreads. Table A.7 shows how each of these filters reduces the data to our final baseline decomposition sample:

Table A.7: Decomposition Sample Construction

	No. of Trades	No. of Bonds
Full MSRB (01/2013-12/2020)	76,151,987	1,650,345
Clean MSRB	59,907,262	653,692
MSRB-Mergent merge	59,905,179	653,590
Exclude variable rate	58,748,919	648,723
Exclude insured bonds	45,575,328	476,877
Exclude AMT bonds	44,884,556	469,367
Exclude trades occurring [0,90 days] of issuance	41,056,328	462,386
Exclude bonds with time-to-maturity $<$ one year	38,888,523	453,988
Exclude insular areas	37,990,711	452,711
Keep State City and County issuers	8,052,129	157,259
City/County Baseline sample (2020 cal. year)	1,089,760	73,556

NOTES—Observation counts here are reported at the trade-level. As such, these counts will not align perfectly with those in Table 6 in which we report mean-collapsed bond-week observations after calculating liquidity measures on the disaggregated Baseline sample data.

Using this transactions sample, we estimate five measures of liquidity spread at the bond-week level. First, we calculate each liquidity measure at the bond-day level and then we collapse to bond-week level by taking the mean of the daily observations. Then, we use this set of liquidity measures as an input to generate the first principal component of liquidity (our main metric of interest):

1. Amihud price impact
2. Effective Bid-Ask Spread
3. Imputed round-trip cost
4. Roll measure
5. Price dispersion

Each measure captures a different aspect of trading costs and is based on traded market prices rather than quotes. While quotes might be available through Bloomberg or other sources, they are not binding, and often only hold for small quantities or can be stale, due to the over-the-counter structure of the municipal bond market, where prices are the result of bilateral negotiations between clients and dealers. This market structure could also lead to situations when a bond is traded at significantly different prices, at approximately the same time. Yet, the assumption that underlies all these measures is that the actual trading is done mostly within the bounds of the quotes.

1. Amihud: Amihud (2002) is a low-frequency measure of the price impact of a trade, i.e., how much the price moves per a \$1 million trade. The Amihud measure for bond b on day t is calculated as:

$$Amihud_{b,t} = \frac{1}{N_t}\sum_{j=1}^{N_t}\frac{|r_j|}{Q_j} = \frac{1}{N_t}\sum_{j=1}^{N_t}\frac{\left|\frac{P_j - P_{j-1}}{P_{j-1}}\right|}{Q_j} \tag{7}$$

where N_t is the number of trades on day t, P_j is the price of the bond at trade j, and Q_j is the par amount of trade j. We require at least two transactions on a given day to estimate the Amihud measure. Weekly estimates of the Amihud measure are obtained by taking the mean of CUSIP-day estimates in a week, and shown in Figure A.15.

Figure A.15: Amihud Evolution in Municipal Bond Markets

(a) Amihud since Great Recession

(b) Amihud in 2020 Calendar Year

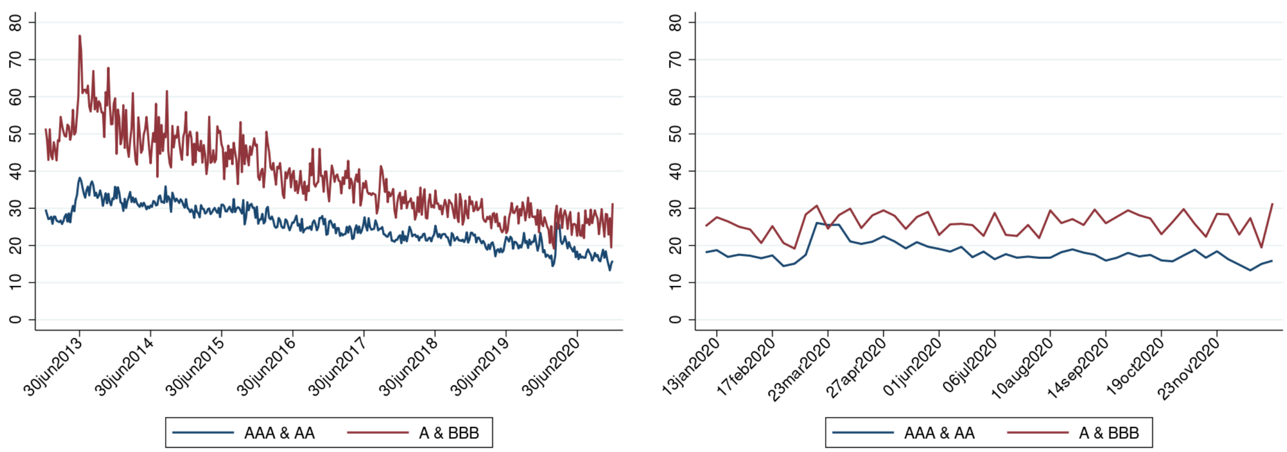

NOTES—Figure shows the weekly mean of the Amihud measure (the mean of CUSIP-days) over time for city and county bonds by January plurality ratings.

2. Effective Bid-Ask Spread (BAS): The effective bid-ask spread is meant to capture dealer compensation for the information advantage possessed by informed traders in a context of few buyers and sellers at desired price ranges. We calculate the effective BAS for each bond as the difference between the size-weighted buy price within the week and the size-weighted sell price within the week, divided by the midpoint of the two prices (adjusting for heterogeneous mark-ups by price levels). Specifically:

$$BAS_{bt} = \frac{\sum_{n\in bt} w^B_{nbt}p^B_{nbt} - \sum_{n\in bt} w^S_{nbt}p^S_{nbt}}{0.5\times\left(\sum_{n\in bt} w^B_{nbt}p^B_{nbt} + \sum_{n\in bt} w^S_{nbt}p^S_{nbt}\right)}\,. \tag{8}$$

Figure A.16 plots the weekly time-series of the average effective bid-ask spread (BAS) across all city and county bonds, separately by January 2020 plurality ratings.

Figure A.16: Bid-Ask Spread Evolution in Municipal Bond Markets

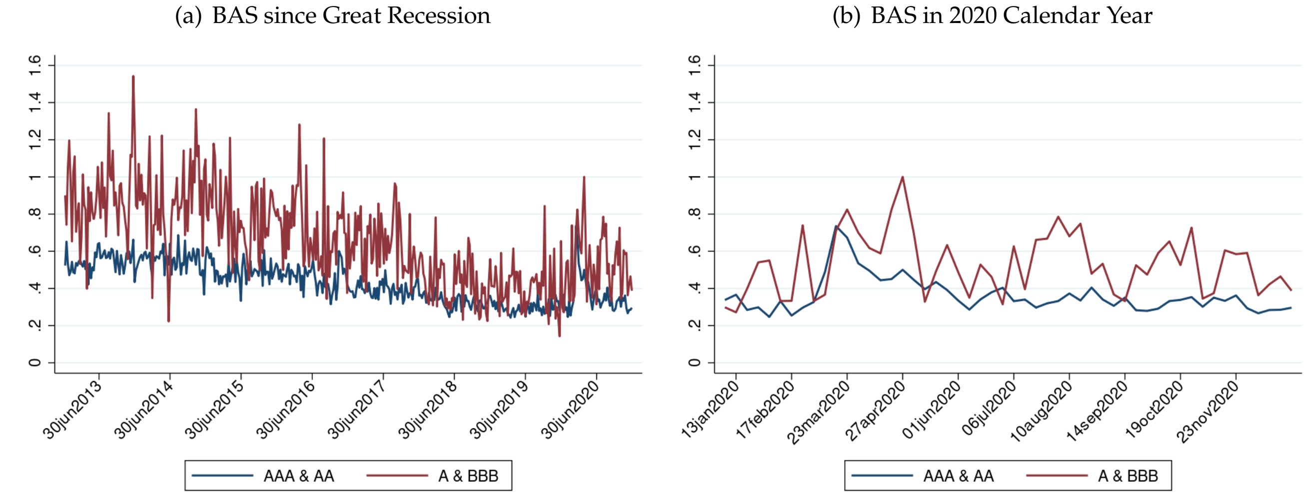

NOTES—Figure shows the average effective bid-ask spread over time for city and county bonds. We calculate the effective bid-ask at the bond-week level as the volume-weighted difference between transacted buy price and sell price from the client's perspective, adjusted by the midpoint level of the two prices. We show results separately by January 2020 issuer plurality ratings, as discussed in the text.

3. Imputed Round-trip Cost (IRC): Similar to the effective BAS, imputed round-trip costs also capture the degree of liquidity in the market, but also include direct transaction costs such as commissions or fees. Following Feldhütter (2012), we identify imputed round-trip trades for a given bond on a given day if there are exactly two or three trades for a given volume that occur within fifteen minutes. Those trades are likely a part of a pre-matched arrangement in which a dealer has matched a buyer and seller. In an imputed round-trip trade, the highest price is assumed to be an investor buying from a dealer, while the lowest price assumed to be an investor selling to a dealer. The investor round-trip cost for bond b and day t is then calculated as the highest minus the lowest price. Specifically:

$$IRC_{bt} = \frac{P_{max} - P_{min}}{P_{min}} \tag{9}$$

where P_{max} is the highest price and P_{min} is the lowest price in an imputed round-trip trade. Weekly IRC estimates at the bond-level are obtained by taking the mean of daily estimates in a given week, and are shown in Figure A.17.

Figure A.17: Imputed Round-Trip Costs in Municipal Bond Markets

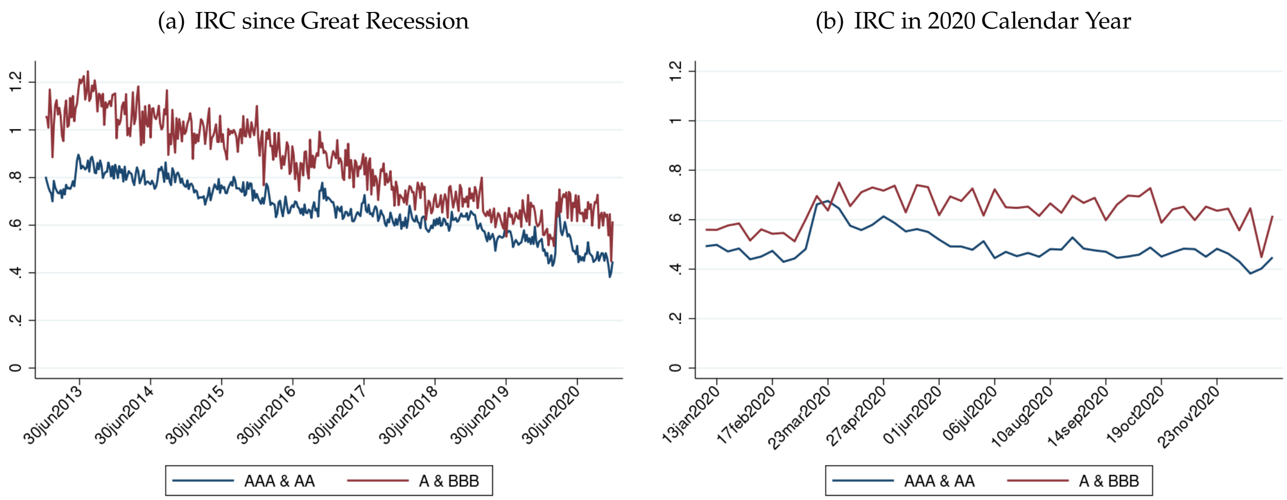

NOTES—Figure shows the average imputed round-trip cost over time for city and county bonds, separately by January 2020 plurality ratings.

4. Roll Measure: Roll's measure (Roll (1984)) provides an alternative construction of the BAS, but requires two strong assumptions. First, the market is informationally efficient. Second, the probability distribution of observed price changes is stationary (at least for short intervals). Under these assumptions, if trading costs are zero, a change in price will occur if and only if unanticipated information is received by market participants. There will be no serial dependence in successive price changes (aside from that generated by serial dependence in expected returns). However, if trading is costly, the dealer will be compensated by the bid-ask spread. When information arrives, both the bid and the ask prices move to different levels such that their average is the new equilibrium value. Thus, the bid-ask average fluctuates randomly in an efficient market. Observed market price changes are no longer independent because recorded transactions occur at either the bid or the ask, not at the average, and a negative serial dependence in observed price changes should be anticipated.

In our context, we calculate the Roll measure for bond b and day t as:

$$Roll_{bt} = 2 \times \sqrt{-Cov\left(\Delta P_n, \Delta P_{n-1}\right)} \tag{10}$$

where P_n is the price at trade n and the measure is set to zero when the covariance between successive price movements is positive. We estimate the Roll measure on each day there is at least one trade, using a trailing 30-day window. Weekly estimates of the Roll measure are obtained by taking the mean of the daily estimates in a week.

Figure A.18: Roll Measure in Municipal Bond Markets

(a) Roll Measure since Great Recession

(b) Roll Measure in 2020 Calendar Year

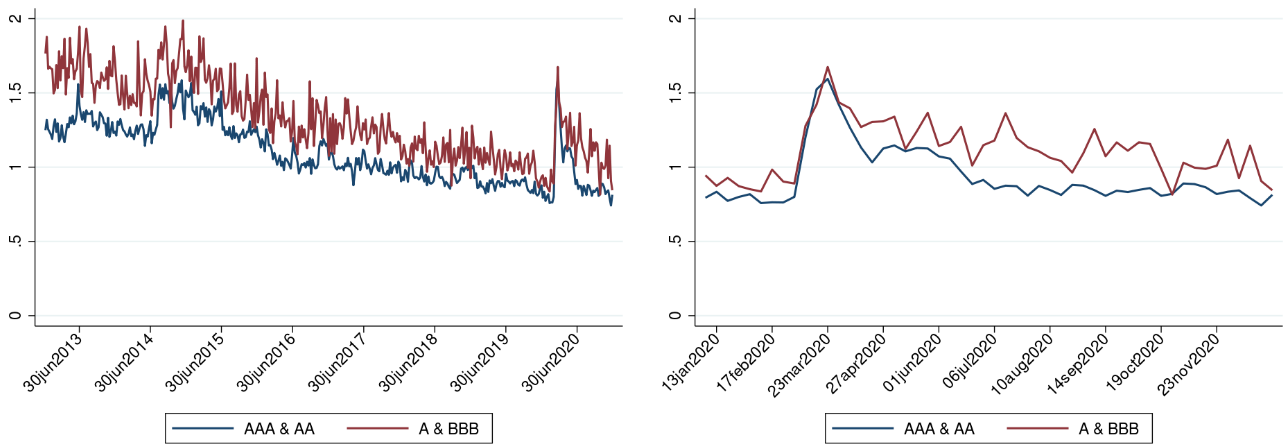

NOTES—Figure show the weekly mean Roll Measure over time for city and county bonds, separately by January 2020 plurality ratings.

5. Price Dispersion: We also consider an alternative measure following Jankowitsch et al. (2011), who use the wedge between expected and traded prices to estimate liquidity. We calculate the dispersion of traded prices around a market "consensus valuation". While this measure is similar to the other liquidity measures, it is not perfectly correlated and seems to contribute additional information about the market liquidity. So, for each bond b on day t we calculate:

$$Dispersion_{bt} = \sqrt{\left(\sum_{n=1}^{N_{bt}} Q_{bn} \sum_{n=1}^{N_{bt}} (P_{bn} - M_{bt})^2 Q_{bn}\right)^{-1}} \tag{11}$$

where N_{bt} is the number of trades of bond b on day t, P_{bn} is the price of the bond at trade n, Q_{bn} is the par amount of trade n, and M_{bt} is the market "consensus valuation" for bond b, which is calculated as the daily volume-weighted average price of the bond. Weekly estimates of the price dispersion measure at the bond-week level are obtained by taking the mean of the daily estimates in a week, and are shown in Figure A.19.

Figure A.19: Price Dispersion Measure in Municipal Bond Markets

(a) Price Dispersion Measure since Great Recession (b) Price Dispersion Measure in 2020 Calendar Year

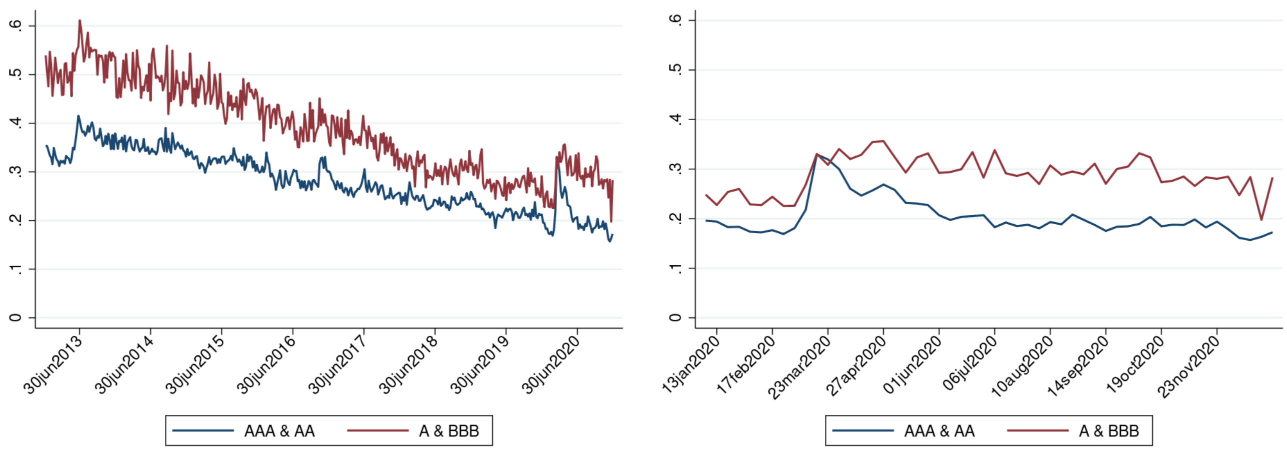

NOTES—Figure shows the average weekly Price Dispersion measure over time for city and county bonds, separately by January 2020 plurality ratings.

Choice of Liquidity Measures: In the main text we choose to present results using a liquidity factor that is calculated based on only four illiqudity measures, excluding the bid-ask spread. Given the relationship between transaction costs and trade size, we would have liked to compare buy and sell prices of trades with similar sizes. As Green et al. (2010) has observed, dealers often provide liquidity to institutional traders by purchasing large blocks of bonds and selling smaller blocks off to retail investors, or to regional dealers with retail distribution capability. This results in asymmetry of the trade size of buys and sales. A reasonable condition such as requiring at least one buy and one sale, each with a trade size greater than $100,000 results in a very restrictive sample, as shown in Figure A.20. Relaxing this restriction indeed results in a similar coverage of bonds as the other illiquidity measures, but we prefer not to use a measure that "mixes" transaction costs of large and small trades without properly "controlling" for the trade size (as the IRC does for example). In either case, our results are robust to including or excluding BAS, and to including BAS without the >$100,000 restriction.

To illustrate the number of bonds that contribute to each measure in a given week, Figure A.20 shows the number of distinct CUSIPs that are in our Baseline sample each week, both overall, and by measure.

Figure A.20: CUSIP Coverage of Each Illiquidity Measure

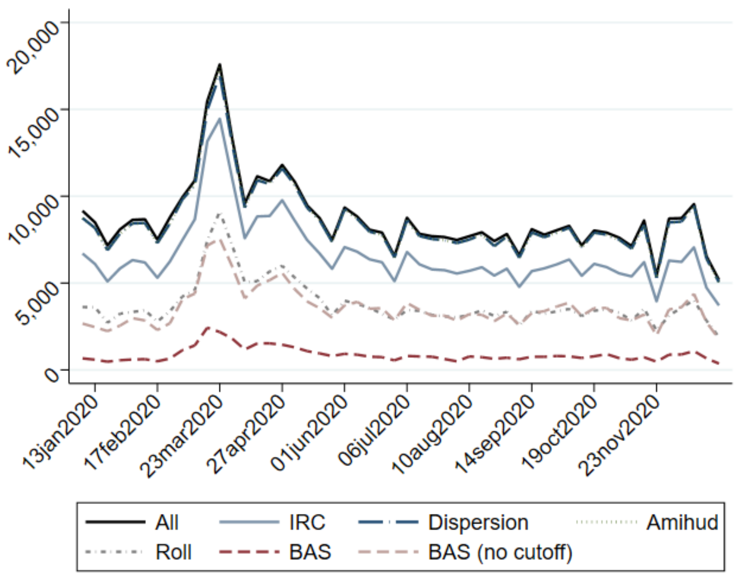

Summary of Liquidity Measures and Liquidity First Principal Component Method: The different illiquidity measures above capture the expected transaction costs and the risk of these costs changing. Historically, as expected, we observe that lower rated bonds (A and BBB) are less liquid than higher rated bonds (AAA and AA).

Table A.8 summarizes the distributive properties of the four illiquidity measures used for the analysis. While the measures are correlated, each individual measure still contains additional distinct information, as observed in Table A.9 which reports the correlation between measures. We also observe that transaction costs are high, though liquidity in the market has been improving over the past decade. We normalize each liquidity measure by subtracting its mean and dividing by its standard deviation over the sample period (January 2013 - December 2020):

$$\lambda_{bt} = \frac{L^k_{bt} - \mu_b}{\sigma_b} \,,$$

where L^k_{bt} is one of the liquidity measures k (Amihud, BAS, IRC, Roll, or Price Dispersion) for bond b on week t. We then consider the first principal component of the four liquidity measures (excluding BAS) as the liquidity factor λ_{bt}. When we present the results for the calendar year 2020, we restrict the sample only after following the procedure described above. That is, the weekly liquidity factors are based on normalization that is based on a longer time series, from January 2013 to December 2020.

Figure A.21 plots the first principal component of the four normalized illiquidity measures. Panel

(a) shows the time series of average illiquidity variable λ_t from 2013 to 2020, exhibiting a long term improvement in liquidity in the municipal bond market consistent with evidence in the literature (e.g., Wu, 2018). Panel (b) takes a closer look at the liquidity during 2020. We see that liquidity worsened in early March during COVID-19 market dislocation across all rating groups. While AAA/AA liquidity improved post-MLF announcement and other Federal interventions that were introduced between mid-March and the end of April 2020, liquidity of the lower rated bonds (A and BBB) has continued to be a problem throughout 2020.

Table A.8: Illiquidity Measures Summary Statistics. This table reports summary statistics for the bond-week illiquidity measures.

	Mean	SD	p1	p10	p50	p90	p99	# Obs.	# Bonds
Amihud (% per $1m)	26.8	48.9	0	0	7.63	75.7	256.7	2,117,126	156,925
BAS	0.44	0.61	0	0.016	0.18	1.33	2.69	187,215	154,322
IRC (%)	0.69	0.70	0.030	0.090	0.42	1.80	2.81	1,670,714	156,933
Roll	1.14	1.21	0	0	0.76	2.83	5.22	907,800	130,840
Dispersion	0.29	0.34	0	0	0.13	0.85	1.36	2,123,759	66,884

Table A.9: Correlations between Illiquidity Measures. This table reports correlations between illiquidity measures on a weekly basis.

	Amihud	BAS	IRC	Roll	Dispersion
Amihud	1	0.2360	0.4220	0.3642	0.4704
BAS	0.2360	1	0.7296	0.4170	0.8203
IRC	0.4220	0.7296	1	0.4987	0.8372
Roll	0.3642	0.4170	0.4987	1	0.5227
Dispersion	0.4704	0.8203	0.8372	0.5227	1

Figure A.21: First Principal Components This figure plots the first principal component of illiquidity measures on a weekly basis, calculated separately for cities and counties by January 2020 plurality ratings.

(a) First PC since Great Recession

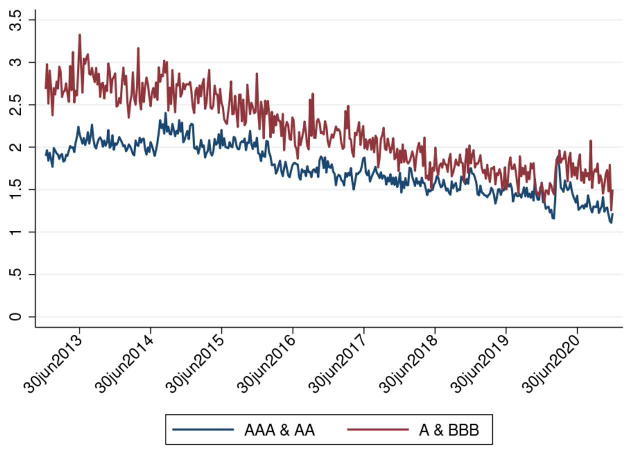

(b) First PC in 2020 Calendar Year

Figure A.22: Illiquidity and Credit Risk Evolution in the 2020 Municipal Bond Market, Including States

(a) States + CC, AAA/AA

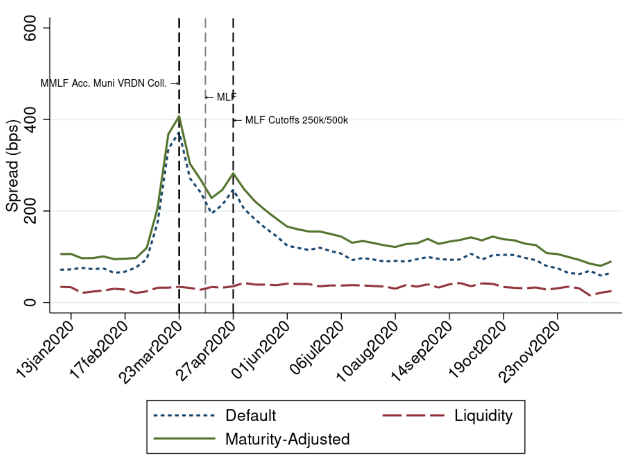

(b) States + CC, A/BBB

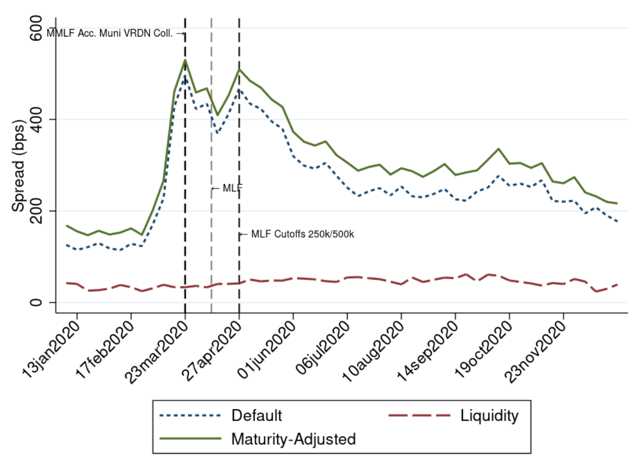

NOTES—Panels (a) and (b) show the weekly averages of the default and liquidity components of city, county, and state bond yields by ratings group for each week in the 2020 calendar year.