Journal of Financial Economics 149 (2023) 235–259

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Market power in wholesale funding: A structural perspective from the triparty repo market☆

Amy Wang Huber

University of Pennsylvania, United States

ARTICLE INFO

Article history:
Received 14 May 2022
Revised 8 April 2023
Accepted 9 April 2023
Available online 6 June 2023

JEL classification:
G11
G12
G21
G23
L13

Keywords:
Triparty repo
Market power
Portfolio allocation
Funding spreads
Intermediary asset pricing

ABSTRACT

I model and structurally estimate the equilibrium rates and volume in the Triparty repo market to study imperfect competition in wholesale funding. Even in this systemically important market, where seemingly homogeneous repos trade, I document persistent rate differences paid by dealers. I characterize the Triparty market as cash-lenders allocating their portfolios among differentiated dealers who set repo rates. I find that cash-lenders' aversion to portfolio concentration and preference for stable lending grant dealers substantial market power: between 2011 and 2017, dealers borrowed at rates that were 26 bps lower than their marginal value from intermediating the borrowed repo funds. Dealers' market power makes the observed wholesale repo rate understate the financing rate available to market participants who rely on repo funding, and offers a novel explanation for funding spreads such as the Treasury cash-futures basis and the Treasury swap spread.

1. Introduction

Funding spreads such as the Treasury cash-futures basis and the Treasury swap spread are large and persistent. Funding spreads exist because intermediation frictions insert a wedge between the wholesale funding rates at which dealers borrow and the rates dealers charge for intermediating their borrowed funds (He and Krishnamurthy, 2018). Balance sheet cost has been advanced as one such intermediation friction (e.g., Jermann 2019, Fleckenstein and Longstaff 2020). In this paper, I argue that an important and complementary friction is imperfect competition in wholesale funding markets. Specifically, I show that dealers' market power in the Triparty repo market causes the observed wholesale repo rate to considerably understate the financing rate available to market participants who rely on dealer-intermediated repo funding. My results illustrate the impact of intermediary competition

☆ I thank the editor, Philipp Schnabl, and an anonymous referee for very helpful comments. Previously circulated under the title "Lender Preference, Borrower Market Power, and the Effect of RRP". I'm grateful for the invaluable guidance from my advisors Arvind Krishnamurthy, Darrell Duffie, Ben Hébert, and Ali Yurukoglu. I thank Juliane Begenau, Jonathan Berk, Greg Buchak, Itamar Drechsler, Wenxin Du, Sam Hanson, Stefan Huber, Sebastian Infante, Richard Kihlstrom, Yi Li (discussant), Hanno Lustig, Matteo Maggiori, Antoine Martin, Daniel Rabetti (discussant), Claudia Robles-Garcia, Amit Seru, Chenzi Xu, Yao Zeng, and Jeff Zwiebel for comments and suggestions. I also thank seminar participants at Columbia University, Cornell University, Dartmouth College, Duke University, the Federal Reserve Bank of New York, Harvard University, Indiana University Bloomington, the London School of Economics and Political Science, the Ohio State University, Princeton University, Rice University, Stanford University, the University of British Columbia, the University of Colorado Boulder, the University of Pennsylvania, the University of Southern California, the University of Texas at Austin, the University of Toronto, and Yale University. All remaining errors are my own.

E-mail address: amyhuber@wharton.upenn.edu

https://doi.org/10.1016/j.jfineco.2023.04.007

on asset prices and offer novel evidence in support of intermediary-based asset pricing.

The $2 trillion Triparty repo market is a key part of the money and bond market, underpinning the working of Treasury and agency mortgage bonds (Copeland et al., 2014; Krishnamurthy et al., 2014). Every day, experienced and sophisticated actors on both sides of the Triparty repo market borrow and lend cash using homogeneous repurchase agreements (repo). Nevertheless, when cash-lenders (e.g., BlackRock) lend to different dealers simultaneously (e.g., Goldman Sachs and Wells Fargo), the rates that cash-lenders accept show persistent cross-dealer differences, suggesting imperfect competition. To quantify the degree of competition, I develop and structurally estimate the first equilibrium model of the Triparty market. My estimates reveal that, although cash-lenders are professional money managers trading with many dealers, their desire to spread out their portfolio and their preference for stable lending grant dealers substantial market power. Between 2011 and 2017, dealers borrowed at rates that were on average 26 bps lower than their estimated marginal value from intermediating the borrowed funds. Many financial market participants rely on dealer-intermediated repo funding to finance their holdings, suggesting that these markdowns represent a significant but hitherto unacknowledged reason why observed wholesale funding rates differ from the financing rates implied in security prices.

I start by documenting three new empirical facts about the Triparty repo market. First, cash-lenders (henceforth, lenders) simultaneously and consistently accept different repo rates for contracts that differ only in the identity of the dealers. Second, dealers' identities drive repo rate dispersion in both the cross-section and time series; in contrast, different lenders that lend to the same dealer do so at rates that are statistically indistinguishable. Third, larger lenders connect to more dealers, not to "rate shop" but to spread out lending, giving smaller shares of their portfolios to each dealer.

These patterns provide new perspectives on how the Triparty market works. First, because Triparty lenders repeatedly lend to a large and fixed set of dealers, lending at persistently different rates is unlikely the result of search frictions but instead reflects that dealers are differentiated. In particular, lenders could have a strong preference for stable investment opportunities and therefore discriminate between dealers who vary in how consistently they use their scarce balance sheet to take on repo loans. Second, although repo contracts are bilaterally determined, the overwhelming importance of between-dealer variation in explaining repo rate dispersion hints at a market where dealers set dealer-specific repo rates for all lenders. As trades in the Triparty market are periodically published,[1] dealer-specific pricing corroborates the notion that dealers have distinct and publicly observable attributes, which the lenders value. Finally, lenders seem to exhibit a size-dependent aversion to concentration, whereby larger lenders are more keen to not concentrate their portfolio in any one dealer. Aversion to concentration compels lenders to spread out their portfolios, possibly at the expense of more profitable lending opportunities. Fears of operational risks such as cyber attacks, which cause operational disruptions and potential fire sales, could trigger this aversion.

These empirical insights motivate me to model the Triparty market using a supply-and-demand framework, where the "good" traded is the cash extended against repo collateral. On the supply side, lenders allocate their portfolio of cash among differentiated repo borrowers (dealers). The lender's utility reflects an aversion to concentration and a non-pecuniary preference for stability in lending opportunities. These two forces determine the lender's optimal lending quantity and his sensitivity to repo rate changes. On the demand side, borrowers set borrower-specific repo rates for all lenders. At her optimum, the borrower offers a rate that could be marked down from her marginal value of intermediation.[2] The ability to build in a markdown is the borrower's market power, and the size of her markdown is a function of the lender's supply elasticity. At a given quantity of repo funding, the less the lender reacts to repo rate changes, the more the borrower can mark down the rate she pays. Thus, the model embeds two forces that affect the lender's behavior and engender imperfect competition: the preference for stability determines how each borrower's borrowing consistency is valued, and the aversion to concentration controls how easily the lender substitutes away from each distinct borrower.

The key to pinning down these forces is understanding how lenders on average respond to rate changes, which I estimate by using offerings at Treasury auctions as an instrumental variable.[3] Over my sample period, a borrower needs to increase the repo rate she pays by 1 bp on average to raise $0.64b in additional funding (about 4%). This estimate is in line with the out-of-sample and recent funding flow to the Overnight Reverse Repo Facility (RRP) following an unexpected rate increase in June 2021.[4] It also signals relative inelasticity in the Triparty market compared to other large, short-term whole-

[1] Many Triparty lenders are required to make regulatory filings. Data from these filings are publicly available; I use these data for my analysis.

[2] For ease of exposition, throughout, I refer to the lender as "he" and the borrower as "she".

[3] The U.S. Treasury Department conducts periodic auctions of Treasury securities. The quantity up for auction influences the amount of repo borrowings sought by dealers because dealers buy securities, which they finance with repo. At the same time, the amount *offered* — not purchased — at each auction is likely driven by the Treasury Department's fiscal concerns and is plausibly exogenous to concurrent preference shocks. The instrument purposely excludes auctions of Treasury bills, which can be purchased by money market funds who are cash-lenders on Triparty.

[4] The RRP is a Federal Reserve (Fed) policy tool that allows Triparty lenders to invest cash with the Fed via repo. On June 17, 2021, the repo rate at the RRP increased unexpectedly from 0 to 5 bps. The volume of repo placed at the RRP increased from $520.9b on June 16, 2021, to $755.8b on June 17, 2021. The $225b overnight increase relative to the total size of Triparty Treasury repo, which was $1628b as of June 9, 2021, implies an elasticity of 2.9% per basis point increase. As the RRP is the lenders' alternative to lending to repo borrowers (see Section 5.2 for more discussion), the lenders' sensitivity to changes in the RRP rate is also the lenders' sensitivity to changes in the borrowers' repo rates.

sale funding markets such as the one for Treasury bills (e.g., Greenwood et al. (2015), Bernanke et al. (2004), Duffee (1996)).[5]

Leveraging this IV-estimated semi-elasticity and other key moments in the data, I estimate my model parameters using indirect inference and maximum likelihood. My parameter estimates accord with the notion that lenders exhibit size-dependent aversion to concentrated portfolios, thus purposely spreading out their lending. This aversion leads to lenders' relatively inelastic responses in volume to repo rate changes, and gives dealers market power in a market of homogeneous goods with observable prices from many counterparties. The magnitude of dealers' market power also depends on lenders' non-pecuniary preferences. Consistent with the conjecture that these preferences capture a taste for stable investment opportunities, the recovered preference values strongly correlate with measures of dealers' borrowing consistency. On average, I calculate that the repo rates paid by dealers during the sample period reflect a 26.2 bps markdown from their marginal value of intermediation. Considering that the spread between the Triparty repo rate and the lenders' outside option is 5.7 bps, dealers command 82% of the 31.9 bps total surplus.

Imperfect competition in the Triparty market induces funding spreads. Dealer-intermediated funding is often the marginal funding that prices assets. In a frictionless world where dealers were a "veil" and passed on funding at cost, there would be no difference between wholesale funding rates and financing rates implied in security prices. Yet Triparty dealers' market power over cash-lenders allows the dealers to pay less than what they charge, thus wedging funding spreads in securities typically financed by repo. Two such funding spreads are the spread between Treasury yields and benchmark rates (Treasury swap spread), and the spread between the implied financing rates in Treasury futures and wholesale repo rates (Treasury cash-futures basis).[6] The magnitude of these two funding spreads is approximately the sum of Triparty dealers' markdowns and riskless arbitrage profits that reflect the opportunity cost of using balance sheet. Hence, in addition to the extensively studied balance-sheet cost,[7] dealers' market power is a key intermediation friction in repo funding and this imperfect competition directly impacts asset prices. Importantly, competition can be altered by policy. Through counterfactual analyses, I show that the Federal Reserve's Overnight Reverse Repo Facility effectively reduces dealers' markdowns by offering the Triparty lenders a competitive outside option.

This paper adds to the intermediary asset pricing literature (see He and Krishnamurthy (2018) for a survey) by studying the role of intermediary competition in asset prices. Relative to the focus on balance sheet cost in studies such as Jermann (2019) and Fleckenstein and Longstaff (2020), this paper surfaces intermediary competition as a complementary angle to interpreting funding spreads involving repo-financed securities. To the extent that the competitive landscape varies in different funding markets, my results provide a micro-foundation for the segmentation in secured versus unsecured funding spreads documented by Siriwardane et al. (2021). More broadly, this paper emphasizes that financing rates used to price assets can differ from observed rates in wholesale funding markets, a point similarly made by van Binsbergen et al. (2021) in regard to the options market.

This study also contributes to the recent stream of research on market power in wholesale funding: by employing three innovative approaches, I provide the first quantification of such market power. First, I tailor a supply-and-demand framework, informed by the new facts that I document, to capture Triparty agents' joint price and quantity decisions. This approach departs from the traditionally sole emphasis on pricing and search frictions in over-the-counter markets (e.g., Duffie et al. (2005), Hendershott et al. (2020)). Second, in contrast to the discrete choice model increasingly employed in finance (e.g., Koijen and Yogo (2018), Benetton (2021), Di Maggio et al. (2022)), my lender's model is bespoke to reflect portfolio allocation considerations that lead to the simultaneous selection of multiple choices at the optimum. The modeling style takes inspiration from Martin and Yurukoglu (2017), Crawford et al. (2018), and Kim et al. (2002) in the industrial organization literature. Third, this paper explores preferences instead of market structures as the source of market power. In the foreign exchange market, dealers' market power owes to the quarter-end regulatory implementation that limits the number of active dealers (Wallen, 2020), and in the European repo market, dealers' market power arises from customers' isolation in a core-periphery network (Eisenschmidt et al., 2021). In contrast, the Triparty market features many large and well-connected counterparties active on both sides of the market. The preferences that grant Triparty dealers market power — for example, the aversion to concentration motivated by operational rather than credit risks — could be applicable in other centrally-cleared derivatives markets, and beyond the financial system in areas such as supply chain management.

This paper examines market power in wholesale funding through an in-depth study of the Triparty repo market. In so doing, I advance the understanding of the Triparty repo market and related short-term money markets. The Triparty repo market has long drawn the attention of scholars. Examples include Krishnamurthy et al. (2014), Copeland et al. (2014), Martin et al. (2014), Hu et al. (2021), Weymuller (2013), Han et al. (2022) and Anbil and

[5] As an example, Greenwood et al. (2015) estimate that a 1-percentage-point decrease in $\frac{\Delta Treasury}{GDP}$ leads to a 38.6 bps decrease in the two-week Treasury yield. The average annual GDP between 2011 and 2017 was \$18.7T. Together, these estimates imply that, over my sample period, a 1 bp change in yield is associated with a \$4.8b change in volume. This response is much higher than in the Triparty market (\$0.6b for 1 bp).

[6] See Du et al. (2022b) and Jermann (2019) for examples of recent studies of the Treasury swap spreads; see Fleckenstein and Longstaff (2020) and Barth and Kahn (2021) for examples of recent studies of the Treasury cash-futures basis.

[7] Certain banking regulations impose surcharge based on the total size of the balance sheet. Because a bank's equity is fixed in the short-run, all types of funding intermediation, which increase total asset, bear opportunity costs (Duffie, 2017). Examples of studies that focus on balance-sheet cost include (Du et al., 2018) and (Du et al., 2022a).

Zeynep (2018).[8] To this vibrant literature, I add new empirical facts, and use these facts to discipline a structural model that generates the first joint determination of rate and volume. The structural model highlights the economic forces that motivate agents, and sheds light on dealer intermediation by quantifying otherwise unobserved market power. Dealers' intermediation of repo funding has profound implications for asset pricing, yet data on this remain elusive. Studies that try to examine and explain repo intermediation have therefore relied on theory or proprietary or indirect data (e.g., Infante (2019), Gorton and Metrick (2012), Gorton et al. (2020)). Recently, Barth and Kahn (2021) used data on cleared bilateral repo, which contain some intermediation activities, but the vast majority of dealer intermediation is uncleared and remains uncaptured.[9] The markdowns I estimate help bridge this data gap by providing a gauge of the intermediation spread using available Triparty data. Finally, because money market funds are important providers of capital (Anderson et al., 2019), studying their interactions with dealers in the Triparty market complements our understanding of not just money market funds' behaviors (e.g., Kacperczyk and Schnabl (2013), Strahan and Tanyeri (2015), Schmidt et al. (2016)) but also the dynamics among many related short-term funding markets (e.g., Aldasoro et al. (2022), Li (2021), Macchiavelli and Zhou (2022), Chernenko and Sunderam (2014)).

In the next section, I provide details on the Triparty market and the data used to study it. In Section 3, I present and discuss the salient empirical observations that motivate my modeling choices. Then, in Section 4, I outline my model of the two sides of the Triparty market. I estimate the lender's model in Section 5, and discuss results and implications in Section 6. In Section 7, I conclude.

2. Triparty repo market and data

In this section, I highlight the distinct features of the Triparty market, outline the role the Triparty market plays in collateral financing, and describe my data.

2.1. The Triparty market and the RRP

Repurchase agreements are contracts between two counterparties to exchange cash against collateral. I refer to the party that provides the cash as the lender and the party that pledges collateral to get cash as the borrower. The posted collateral, often valued at a haircut, is returned to the borrower when the cash is repaid — with interest. I refer to the rate used to determine that interest as the repo rate. Because repo lending is secured, repo contracts can differ depending on what collateral is used and how the collateral is managed. The Triparty market offers a way to homogenize the otherwise complex over-the-counter repo collateralization.

The Triparty market derives its name from its institutional setup. On Triparty, every transaction involves a third agent, who is the clearing bank that handles the logistics of cash and collateral transfers. All Triparty borrowers and lenders maintain accounts with the same clearing bank.[10] Once a borrower and a lender agree on the terms of a repo, the clearing bank places the pledged collateral in a segregated account and monitors the value of the collateral. Triparty repo are also general collateral, that is, contracts specify only the class of collateral, e.g., Treasury securities, but not the exact securities used, e.g., CUSIPs of specific five-year on-the-run Treasury securities. These features make Triparty repo contracts comparable within a collateral class, and make Triparty repo convenient for funding.[11]

The Triparty market is a crucial step in the intermediation process that channels cash through the financial system to market participants looking to finance their holdings. Cash-rich individuals and corporations place cash in vehicles such as money market funds (MMFs). MMFs keep a stable fraction of their assets under management (AUM) in overnight cash for liquidity, and they see lending to dealers using overnight repo as an attractive option for investing this cash. Dealers, in turn, intermediate their borrowed cash to clients in the broader financial market who rely on repo to finance their securities. Every day, over $2 trillion is injected into the secured funding market through the Triparty market.

The wholesale repo rates in the Triparty market influence the price of many assets. To start, Triparty repo rates affect the pricing of Treasury securities because repo is the predominant way in which financial market participants purchase Treasury securities with leverage. By using the security as collateral, repo allows the purchase of any security to be done with as little capital as the haircut. Repo is particularly suited to finance safe, government-backed securities because the haircut on these securities is modest and stable.[12] In fact, over 90% of the collateral pledged in Triparty repo are Treasury and agency mortgage-backed securities (MBS), and over 80% of bilateral repo use Treasury securities as collateral (Baklanova et al., 2017). The funding condition on the Triparty market therefore directly shapes the funding condition for Treasury securities. More-

[8] Krishnamurthy et al. (2014), Copeland et al. (2014), and Martin et al. (2014) investigate the role of the Triparty market in the 2007-09 Financial Crisis. Hu et al. (2021) and Weymuller (2013) study how haircuts are determined, especially for repo involving risky collateral such as equities. Han et al. (2022) and Anbil and Zeynep (2018) examine the role of relationships and regulations in Triparty activities.

[9] Eisenschmidt et al. (2021) use European regulatory data, which cover all repo transactions made by 38 dealer banks including bilateral repo trades, to study the segmentation between a competitive inter-dealer market and a sparsely connected dealer-customer OTC market. Because the European repo market is decentralized, with each customer — be it lender or borrower — transacting with only one or a few dealers, the intermediation frictions that cause lending and borrowing rates to be different in Europe are potentially different from those in the U.S.

[10] The sole Triparty clearing bank in the U.S. is Bank of New York Mellon. J.P. Morgan used to also provide Triparty clearing service for about 15% of the market. It discontinued its service in 2017.

[11] In contrast, repos done outside of the Triparty market allow the borrower and lender to maintain accounts at different custodial banks, and in some cases, collateral is not physically transferred from the borrower's custodial bank to the lender's custodial bank. Moreover, when the repo contract is not general collateral, the borrower needs to stipulate the CUSIPs of all securities used as collateral at the time of the transaction.

[12] For example, over 95% of the Treasury repo in my sample has a haircut of 2%.

over, Triparty repo rates indirectly impact trillions of dollars of financial contracts indexed to the Secured Overnight Financing Rate (SOFR) because SOFR, the dollar interest rate benchmark, uses Triparty repo trades in its construction.

The Federal Reserve (Fed) has long been a keen observer of the Triparty repo market. In September 2013, in anticipation of implementing monetary policy changes in an abundant reserves environment, the Fed set up an overnight, fixed-rate, full-allotment reverse repo facility (RRP) on the Triparty market.[13] The RRP gives a wide array of Triparty cash lenders[14] the ability to lend to the Fed in the form of overnight repo at a pre-announced interest rate. When the RRP was first set up in September 2013, access to the facility was capped at $500 million per eligible lender. This cap was subsequently raised six times, eventually reaching $30 billion per eligible lender by September 2014, at which point the cap no longer seemed binding for any lender. September 2014 was also when the Fed stated in the FOMC's Policy Normalization Principles and Plans that "the Committee intends to use the RRP facility as a tool to help control the federal funds rate during the normalization of the stance of monetary policy". This policy bolstered lenders' confidence in the Fed's commitment to the RRP. In effect, by September 2014, Triparty lenders had an attractive alternative to lending to repo borrowers, in the form of the RRP.

2.2. Money market funds and data

To study the Triparty market, I use monthly filings with the Securities and Exchange Commission (SEC) made by MMFs domiciled in the United States. MMFs are the largest class of cash lenders on the Triparty market, accounting for 40% to 60% of all repo transactions. Other Triparty lenders include security lenders, pension funds, insurance companies, and various municipalities with temporary excess cash (Copeland et al., 2012).

Money market funds are fixed income mutual funds that invest in high quality, short-term securities. MMFs are established and operated by sponsors such as BlackRock and State Street, who typically run a family of MMFs that differ in their investment universe, e.g., government-security-only funds vs. prime (or general purpose) funds. Because MMFs issue demand deposits, investors regard MMFs as a higher-yielding cash alternative, and park about $3 trillion with MMFs. However, because MMF investments are not insured by the government, and many MMF investors are large-scale institutional investors, MMFs are particularly prone to runs (Kacperczyk and Schnabl, 2013; Schmidt et al., 2016). In fact, the money market fund industry suffered a run following the default of Lehman Brothers in September 2008.

In the wake of the 2007-09 Financial Crisis, the money market fund industry saw three important changes. First, there was an increased focus on liquidity. Not only did many MMFs voluntarily increase holdings of liquid assets such as repo (Strahan and Tanyeri, 2015), but regulations also strengthened requirements on portfolio liquidity and maturity. Second, stability became a central tenet. Low volatility in yield is paramount to prime MMFs that converted to reporting floating Net Asset Values (NAVs) after the crisis. Yet even constant NAV MMFs tout stability — along with liquidity — as key advantages to investing in MMFs.[15]

The third change to the money market fund industry was that MMFs were required to file monthly N-MFP reports starting in October 2010. These filings are snapshots of an MMF's entire portfolio as of the last business day of each month.[16] For each repo contract that the MMF has, information is available on the counterparty, the amount, the repo rate,[17] the maturity date, and the collateral type and value.

I obtain all N-MFP reports between 2011 and 2017 from the SEC EDGAR data base and collapse the filings by sponsor.[18] Sponsors enter into repo contracts on behalf of all the funds in the family and then distribute the investment across funds (Copeland et al., 2014). Consequently, to analyze the equilibrium rate and volume determination, I consider all the funds in the same fund family as one entity.[19]

To manage portfolio liquidity and maturity, MMFs keep a steady fraction of their AUM in overnight cash. As there are limited options to invest cash overnight, Triparty repo, which is typically overnight, forms an important part of MMFs' overnight portfolio. I explicitly focus on Triparty repos that are overnight in duration.[20]

Triparty repo activities are concentrated in a relatively small set of agents. As Appendix Fig. A1 illustrates, over 85% of the activities in the N-MFP filing data are done by 18 MMFs and 20 dealers. My analysis thus focuses on these agents. My final data set is a MMF-dealer-month panel of repo transactions from January 2011 through December 2017 on 84 month-ends, for a total of 14,571 observations.

I supplement the data from N-MFP filings with the Federal Reserve's releases on RRP activities, TreasuryDirect's reporting of Treasury securities auctions, and credit default swap (CDS) pricing from Markit and Bloomberg.

[13] https://www.newyorkfed.org/markets/opolicy/operating_policy_130920.html

[14] https://www.newyorkfed.org/markets/rrp_counterparties

[15] https://www.fidelity.com/learning-center/investment-products/mutual-funds/what-are-money-market-funds

[16] Although N-MFP filings are done monthly, these reports are likely representative of the MMF's repo activities throughout the month. Anderson et al. (2019) use proprietary data available from the Federal Reserve and show that repo activities are stable throughout the month, with the exception of window-dressing activities on the last day of the quarter.

[17] Before April 2016, repo rates were not separately reported. I parse the title of each contract to obtain rates where available. To address potential misreporting issues, I winsorize repo rates at 1% and 99%.

[18] My sample starts in 2011 because the first three months of reporting had many issues as MMFs learned to comply with new reporting requirements. My sample ends in 2017 because I am interested in the effect of imperfect competition in an abundant reserve regime. The Fed started shrinking its balance sheet in 2018.

[19] To illustrate, the standard deviation of rates obtained by funds within a fund family is 0 at the median and less than 1 basis point at the 73^{rd} percentile.

[20] The N-MFP reports the maturity date without reporting the start date. I therefore identify overnight contracts as those that mature on the first business day of the following month. This approach is universal in papers that use N-MFP filing data. See Aldasoro et al. (2022) for a discussion of potential shortcomings.

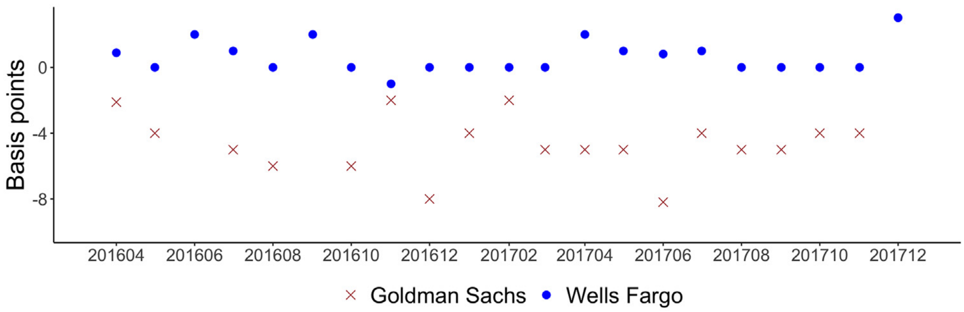

Fig. 1. Select repo rates of BlackRock MMF's lending.
Notes: This figure plots the repo rates accepted by BlackRock MMF for lending to Goldman Sachs and Wells Fargo via overnight repo collateralized by Treasury securities with a 2% haircut. The repo rates are reported as gross rates less the daily median repo rate and are stated in basis points. Two outliers are omitted: the repo rate by Goldman Sachs was 20 bps below the median in September 2016, and 12 bps below the median in December 2017.

3. Empirical patterns on Triparty

In this section, I present three facts of the Triparty market and discuss possible reasons for these empirical observations.

Fact 1: MMFs simultaneously and consistently accept different repo rates from different dealers.

Given repo contracts that differ only in the identity of the dealer, MMFs simultaneously lend to multiple dealers and do so at consistently different rates. To compare repo rates across dealers, I focus on a subsample of overnight repo contracts that are collateralized only by Treasury securities with a 2% haircut,[21] which retains 75% of all MMF-dealer-month observations.[22] Prior research has documented differences in repo rates between contracts backed by different collateral, even though the haircut on collateral theoretically adjusts for the quality of the underlying assets (Weymuller, 2013). I therefore compare repo rates only across repos that are backed by the same collateral. This choice is not restrictive because Triparty repo differentiates collateral not by CUSIP but by asset class. By focusing my analysis of repo rates on Treasury-backed overnight repo contracts that all have a 2% haircut, I retain a sizeable subsample whose repo rate is not affected by differences in duration, collateral, or haircut, leaving dealer identity the only other factor that differentiates repo.

In this subsample of homogeneous repos, MMFs are seen to simultaneously accept different repo rates from different dealers. As an example, Fig. 1 plots the repo rate that BlackRock MMF received from lending to Goldman Sachs and Wells Fargo in 2016 and 2017.[23] Although Goldman Sachs and Wells Fargo consistently borrowed at different rates, BlackRock nonetheless lent to both dealers month after month.

This observation is surprising because the difference in repo rates cannot be due to differences in the contractual terms of the repo. Moreover, although the creditworthiness of dealers can often be a source of price dispersion, especially in OTC markets, in 2016 and 2017, the average short-term (6M) CDS rate of Goldman Sachs was in fact 12 bps above that of Wells Fargo. This observation is all the more surprising when considering that the same three agents transacted every month. How could BlackRock not be aware that it is paying different rates to Goldman Sachs and Wells Fargo on a persistent basis? Informational friction cannot explain such a persistent difference.

More generally, MMFs in my sample on average lend to 10 dealers at a time, and the difference between the highest and the lowest rate *simultaneously* accepted by MMFs is on average 4 bps. This dispersion is, again, present in the context of repeated interaction. Indeed, the AR(1) persistence of whether a MMF-dealer pair trades is 85% ($R^2 = 0.72$). Between sophisticated financial institutions that repeatedly interact with each other, repos that differ only in the identity of the dealer trade at persistently different rates. These observations spur me to further investigate patterns in rate dispersion and in MMFs' lending.

Fact 2: Dealer identity drives repo rate dispersion.

The dispersion in apparently homogeneous repo contracts is driven by dealer identities. Using the overnight Treasury repo subsample, I show that dealer identities explain the preponderance of the repo rate variation in both the cross-section and time series.

In Fig. 2, I examine the dispersion of repo rates in the cross-section by regressing repo rates in each month on MMF or dealer fixed effects.[24] The plotted R^2 from these

[21] I restrict the sample to contracts with a collateral-to-principal ratio of 102% ± 0.1%.

[22] I start with 14,571 MMF-dealer-month observations. Focusing on transactions that use Treasury securities as collateral decreases the sample to 11,486 observations. Further restricting the sample to a 2% haircut leaves me with 10,938 observations.

[23] To remove the effect of general interest rate trends, I restate the transacted repo rate as the deviation to the volume-weighted median on that day. Doing so both conforms to the convention in the repo market and minimizes the impact of outliers. All published repo-based indices are calculated as the volume-weighted median. Examples include the Secured Overnight Financing Rate (SOFR), the Triparty repo index, and the DTCC GCF repo index.

[24] Comparing the explanatory power of models with different fixed effects is the focus of the figure and tables in this subsection. To avoid fitting fixed effects over one or two data points, I exclude observations in which a borrower or a lender has fewer than three transactions in

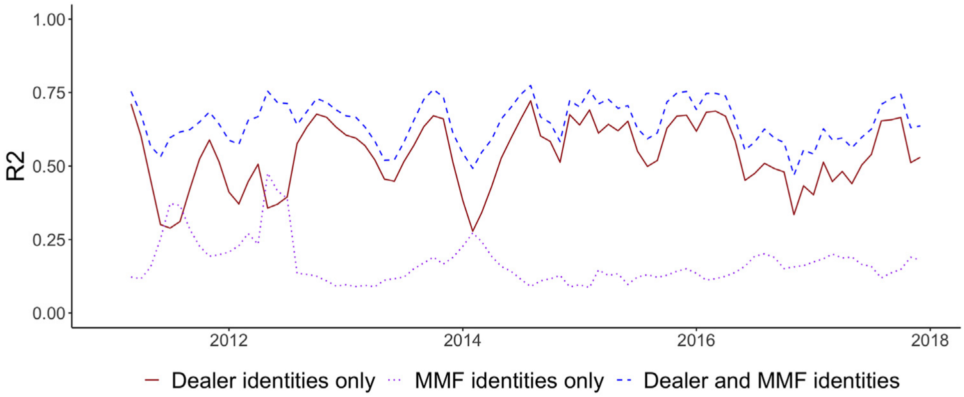

Fig. 2. Decomposition of cross-sectional variation in rate dispersion.
Notes: This figure plots the 3-month rolling average of the R^2 from monthly cross-sectional regressions of repo rates on MMF and dealer fixed effects. Repo rates are measured as gross repo rates less the daily median repo rate, and are for overnight repo collateralized by Treasury securities with a 2% haircut. The solid red line is the R^2 from regressing repo rates on dealer fixed effects; the dotted purple line is the R^2 from regressing repo rates on MMF fixed effects; and the dashed blue line is the R^2 from regressing repo rates on both dealer and MMF fixed effects. The sample period is January 2011 through December 2017. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

regressions show that dealer identities alone explain about 50% of the variation, while MMF identities explain much less. Even if MMF identities did not affect dispersion, I would expect MMF fixed effects to have some explanatory power as long as the sorting of dealer to MMF is not completely symmetric. I thus regress repo rates in each month on both dealer and MMF fixed effects. Once dealer identities are controlled for, adding MMF fixed effects does not improve the R^2 by much. I formally test whether the additional MMF fixed effects are jointly 0, and I cannot reject the null at the 10% significance level in 72 of the 84 months.[25]

Not only is repo rate dispersion mostly driven by between-dealer variation, but within a dealer, variation cannot be explained by MMF or MMF-dealer pair characteristics. In Table 1, I test the following characteristics: the amount of Treasury-backed overnight repo lending between a MMF-dealer pair (*Pair Treasury repo volume*), the share — or importance — of this pair's lending volume to the MMF's overall Treasury-backed overnight repo lending (*Pair vol as percent of MMF*), the share of this pair's lending volume to the dealer's overall Treasury-backed overnight repo borrowing (*Pair vol as percent of borrower*), the MMF's total Treasury-backed overnight repo lending (*MMF total Treasury repo vol*), and the number of dealers to whom the MMF lends (*MMF number of counterparty*). None of these characteristics statistically significantly explains the rate dispersion, judging by either the individual coefficients' statistical significance or the improvement in R^2 from including these regressors (difference between R^2 (proj model) and R^2 (full model)). Together, the results from Fig. 2 and Table 1 indicate that while different dealers systematically pay different rates for repo funding, different MMFs lending to the same dealer do not receive statistically different rates.

Dealer identity is also the principal driver of repo rate variation in the time series. Given the OTC nature of repo transactions, prior studies such as (Han et al., 2022) focus on differences in MMF-dealer pairs to explain rate variations. In Table 2, I compare the goodness of fit between models that include only dealer fixed effects and models that include pair fixed effects. Moving from Model (1) to Model (2), including all the pair fixed effects improves the R^2 by about 0.06, yet this is achieved with 251 more regressors. Therefore, the Akaike information criteria (AIC) ranks these two models similarly, and the Bayesian information criteria (BIC) prefers the more parsimonious model with only dealer fixed effects. The same pattern holds when comparing Models (3) and (4), both of which include year fixed effects. In short, dealer identities are of first-order importance in explaining Triparty repo rate dispersion.

Fact 3: Larger MMFs connect to more dealers to spread out lending.

MMFs construct their overnight cash portfolios in a systematic and size-dependent way. As an example, Fig. 3 compares the repo lendings done on October 31, 2016, by two MMFs. The larger MMF (BlackRock) lent to more dealers and lent smaller shares of the portfolio to each dealer. Table 3 shows that this relationship between portfolio size and both the extensive and the intensive margins of the portfolio holds across MMFs.

In Model (1) of Table 3, the number of dealers to whom an MMF lends increases by about 3 as the size of MMF's overnight cash portfolio doubles.[26] This finding is perhaps intuitive. An MMF incurs some fixed costs, such as setting up master repurchase agreements, before it can lend to a

a month. Doing so reduces the sample for Fig. 2 and Tables 1 and 2 to 10,453 observations.

[25] Multiple hypothesis testing is corrected using the method described in Holm (1979).

[26] Overnight cash portfolios are measured as the sum of all overnight repo lending, inclusive of repo placed in RRP.

Table 1
Within dealer characteristics and rate dispersion.

	Deviation of Treasury repo rate from median					
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Pair Treasury repo volume	0.013					0.052
	(0.022)					(0.041)
Pair vol as percent of MMF		0.049				−0.661
		(0.366)				(0.419)
Pair vol as percent of dealer			0.289			0.508
			(0.245)			(0.355)
MMF total Treasury repo vol				−0.007		−0.013
				(0.007)		(0.013)
MMF number of counterparty					−0.013	−0.011
					(0.011)	(0.018)
Dealer + Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Num. obs.	10,453	10,453	10,453	10,453	10,453	10,453
R^2 (full model)	0.227	0.227	0.227	0.227	0.227	0.228
R^2 (proj model)	0.000	0.000	0.000	0.000	0.000	0.002

Standard errors in parentheses.
Notes: In this table, repo rates are regressed on MMF characteristics and MMF-dealer pair characteristics, as well as dealer fixed effects and month (time) fixed effects. Repo rates are for lending between MMF-dealer pairs via overnight repo collateralized by Treasury securities with a 2% haircut, and they are measured as the deviation from the daily median. "Pair Treasury repo volume" is the amount of overnight repo lending collateralized by Treasury securities with a 2% haircut on the day of the observation between an MMF-dealer pair. "Pair vol as percent of MMF" is the pair's volume as a percentage of the MMF's total lending via overnight repo collateralized by Treasury securities with a 2% haircut. "Pair vol as percent of dealer" is the same ratio against the total borrowing of the dealer. "MMF total Treasury repo vol" is the MMF's total amount of lending via overnight repo collateralized by Treasury securities with a 2% haircut, on the day of the observation. "MMF number of counterparty" is the number of dealers to which the MMF lends via overnight repo collateralized by Treasury securities with a 2% haircut on the day of observation. The sample period is January 2011 through December 2017. Standard errors are clustered by dealer. *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

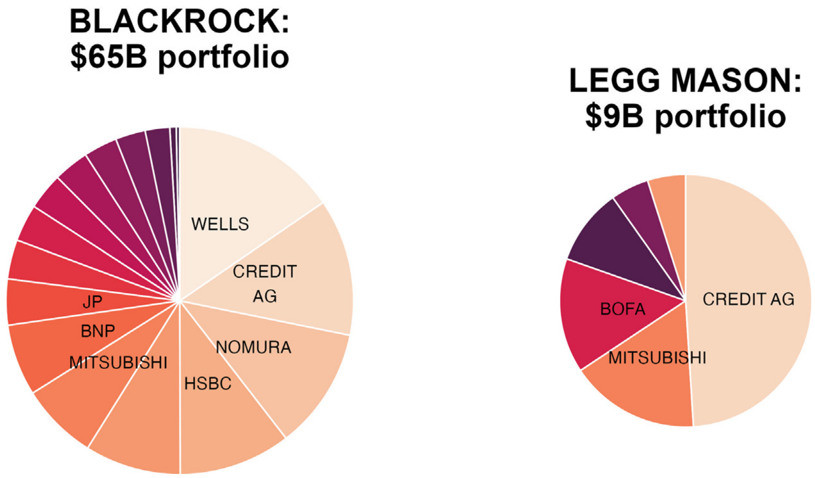

Fig. 3. Select MMF repo portfolios on October 31, 2016.
Notes: This figure plots the repo lending to dealers by BlackRock and Legg Mason on October 31, 2016. The size of the pie corresponds to BlackRock and Legg Mason's overnight repo lending volume, as labeled. The size of each slice represents the share of the portfolio lent to different dealers.

dealer,[27] and larger MMFs can afford to establish lending relationships with more dealers. But why do MMFs want to lend to more dealers?

MMFs are not establishing more lending relationships to better "rate shop". If MMFs willingly incur fixed costs because lending to more dealers allows MMFs to find more attractive rates, then I would expect the portfolios of more connected MMFs to change more frequently and significantly.[28] However, as Fig. 4 shows, more-connected MMFs maintain portfolios that are just as stable, if not more stable, over time. I measure the similarity between MMFs' overnight repo portfolios across time using cosine similarity: $CosSim(\mathbf{x}, \mathbf{y}) = \frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\| \|\mathbf{y}\|}$, where $\mathbf{x}$ and $\mathbf{y}$ denote a given MMF's portfolio at time t and $t+n$, respectively. The closer the cosine similarity is to 1, the more stable the portfolio.

[27] Master repurchase agreements (MRAs) typically set out the protocols for margin maintenance and default; thus, day-to-day, cash-lenders and cash-borrowers only need to decide on the repo rate and volume.

[28] Fact 2 revealed that having more connections does not grant an MMF preferential pricing vis-à-vis any one dealer. Hence, if there is a price-related advantage to forming more connections, it is likely the ability to shift lending to whichever dealer offers the best rate on a given day.

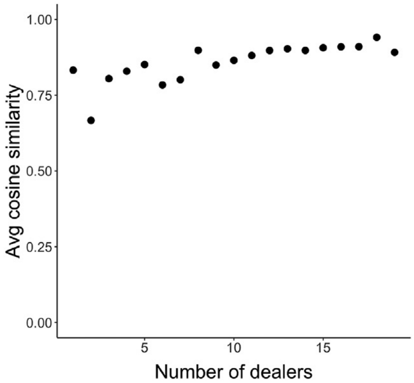

(a) **Between t and $t+1$**

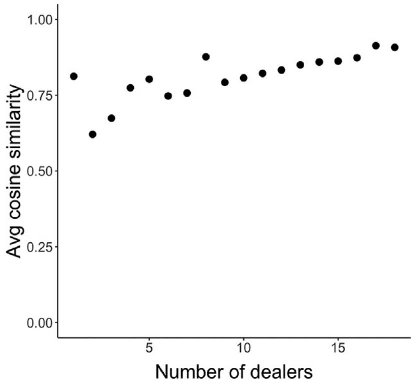

(b) **Between t and $t+6$**

Fig. 4. Stability in MMF overnight portfolios.
Notes: This figure plots the average cosine similarity between MMFs' overnight portfolios at time t and $t+n$ against the number of dealers in the portfolio at time t. Cosine similarity between portfolio $\mathbf{x}$ and $\mathbf{y}$ is defined as $CosSim(\mathbf{x},\mathbf{y}) = \frac{\mathbf{x}\cdot\mathbf{y}}{\|\mathbf{x}\|\|\mathbf{y}\|}$. Panel (a) compares portfolios at time t and at time $t+1$, Panel (b) compares portfolios at time t and at time $t+6$. The sample period is January 2011 through December 2017.

Table 2
Model fit with dealer vs. relationship FE.

	Deviation of Treasury repo rate from median			
	Model 1	Model 2	Model 3	Model 4
Dealer FE	Yes	No	Yes	No
MMF-Dealer Pair FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes
N	10,453	10,453	10,453	10,453
Num of FE	20	271	26	277
R^2	0.18	0.24	0.19	0.25
AIC (in 1000s)	41.53	41.22	41.38	41.07
BIC (in 1000s)	41.68	43.19	41.58	43.08

Notes: This table reports the goodness of fit for regressions of repo rates on dealer fixed effects or MMF-dealer pair fixed effects. Repo rates are for lending between MMF-dealer pairs via overnight repo collateralized by Treasury securities with a 2% haircut, and they are measured as the deviation from the daily median. Goodness-of-fit measures are R^2, the Akaike information criterion (AIC), and the Bayesian information criterion (BIC). The sample period is January 2011 through December 2017.

Panel (a) of Fig. 4 shows that, irrespective of the number of dealers MMFs lend to at time t, MMFs maintain very similar portfolios at time $t+1$. In fact, for portfolios that start with 10 or more dealers, the mean cosine similarity is about 0.9, slightly higher than the 0.8 mean of portfolios that start with fewer than 10 dealers. Panel (b) of Fig. 4 shows that the stability in portfolios is not a quirk of short-term stickiness, but that portfolios at time t and time $t+6$ also exhibit high degrees of similarity across MMFs.

Instead, by connecting to more dealers, MMFs are able to lend a smaller share of the portfolio to each dealer, achieving a distributed portfolio. Model (2) of Table 3 shows that as MMFs double in size, the median portfolio share lent to dealers decreases by 6.7% on average. Unlike the mean share of the portfolio, the median share of the portfolio need not change as the number of dealers in the portfolio changes. Rather, this declining median share signals MMFs' desire to reduce exposure to any one dealer. Indeed, in Models (3) and (4), we see that the maximum share and the minimum share also decrease with portfolio size, and by similar magnitude, reflecting a deliberate effort to spread out the portfolio.

Table 3
MMF size and portfolio composition.

	Number of dealers	Median portfolio share	Max portfolio share	Min portfolio share
	Model 1	Model 2	Model 3	Model 4
Constant	1.403	0.266***	0.399***	0.188***
	(1.021)	(0.037)	(0.033)	(0.042)
Log(MMF size)	3.022***	−0.067***	−0.074***	−0.055***
	(0.290)	(0.013)	(0.012)	(0.014)
Num. obs.	1467	1467	1467	1467
R^2	0.556	0.524	0.369	0.439

Standard errors in parentheses.
Notes: This table reports regressions of the extensive and intensive margins of MMFs' portfolio on MMFs' overnight portfolio size and a constant. The dependent variable in Model (1) is the number of dealers to which a MMF lends. The dependent variables in Models (2) through (4) are, respectively, the median, the maximum, and the minimum share of an MMF's portfolio lent to dealers. The sample period is January 2011 through December 2017. Standard errors are clustered by MMF. *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

Table 4
Dealer repo activities on quarter-ends.

	Dealer repo volume (vol_{jt})		Log(dealer repo volume)		Dear rate
	Model 1	Model 2	Model 3	Model 4	Model 5
QE	0.843	0.716	0.087	0.078	−0.037
	(0.733)	(0.872)	(0.087)	(0.083)	(0.297)
QE * Dealer EU	−7.754***	−7.716***	−0.500***	−0.496***	−0.086
	(1.205)	(1.227)	(0.120)	(0.112)	(0.348)
QE * Dealer JP	−0.814	−0.761	−0.137	−0.134	0.292
	(2.711)	(2.540)	(0.227)	(0.199)	(0.611)
QE * Dealer UK	−4.394***	−4.331***	−0.326**	-0.322**	0.097
	(1.659)	(1.477)	(0.150)	(0.130)	(0.344)
QE * Dealer US	−1.658	−1.548	−0.157	−0.150	0.138
	(1.396)	(1.424)	(0.112)	(0.107)	(0.383)
Avg change: EU, UK	−5.231	−5.308			
Avg change: CA, JP, US	0.019	−0.054			
Dealer HQ FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	No	Yes	No	Yes	Yes
Num. obs.	1415	1415	1415	1415	1415
R^2	0.098	0.251	0.078	0.258	0.109

Standard errors in parentheses.
Notes: This table reports regressions of the dealer's overnight repo volume and rate on indicators of quarter-ends and the dealer's headquarters jurisdiction. The dependent variable is the dealer's repo volume in Models (1) and (2), the log of the dealer's repo volume in Models (3) and (4), and the dealer's repo rate in Model (5). Headquarters jurisdictions are Canada (CA), the European Union (EU), Japan (JP), the United Kingdom (UK), and the United States (US). "Avg change: EU, UK" is calculated as $(QE * DealerEU + QE * DealerUK + QE * 2)/2$. "Avg change: CA, JP, US" is calculated as $(QE * DealerJP + QE * DealerUS + QE * 3)/3$. The dealer's repo rate is defined as the volume-weighted average of repo rates between a dealer and all lenders in overnight repo collateralized by Treasury securities with a 2% haircut. It is reported as the gross rate less the daily median, in basis points. The sample period is January 2011 through December 2017. Standard errors are robust to heteroskedasticity. *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

Discussion of empirical patterns

Faced with dealers that borrow at different repo rates (Fact 2), MMFs lend to multiple dealers simultaneously (Fact 1) and show a size-dependent tendency to spread out their portfolios (Fact 3). These facts about the Triparty market suggest at least two economic forces that work in tandem.

First, MMFs knowingly accepting different rates indicates that dealers' identities differentiate repo lending in a way that rationalizes differences in pecuniary returns. How do lenders differentiate among dealers? Conversations with over a dozen industry participants point to the importance of consistency. MMFs have a strong preference for stable repo investments because low volatility in yield helps them market themselves as higher-yielding cash-alternatives. Dealers, however, may not be able to consistently borrow because repo is balance-sheet intensive and can come at a high opportunity cost. MMFs therefore likely favor dealers who devote a consistent amount of their balance sheet space to repo intermediation. I test this hypothesis in Section 6 and find that MMFs' estimated non-pecuniary preferences for dealers indeed exhibit strong correlations with measures of borrowing consistency. The presence of this preference makes lenders perceive lending to different dealers as differentiated investment opportunities that warrant differential rates.

At the same time, MMFs lend to multiple dealers simultaneously, suggesting that MMFs have an aversion to concentration and purposely spread out their lending.[29] At first blush, this sounds perplexing. Triparty repo carries little credit risk: these contracts are largely overnight and are backed by high-quality, bankruptcy-remote collateral valued with conservative haircuts (Hu et al., 2021). Even in the depth of the 2007-09 Financial Crisis, there was no run or default on Triparty (Krishnamurthy et al., 2014). Nevertheless, MMFs may want to limit their dollar exposure to any one dealer to minimize operational risks. Tail risks such as cyber attacks could prevent a dealer from returning its repo loan on time. If MMFs take possession of the posted collateral, regulations require that MMFs liquidate the collateral immediately.[30] Fire sales could materially compromise the value of the collateral if the total amount to be sold is large. Worries over fire sales are consistent with the pattern of portfolio diversification that intensifies with size: even small shares of a large portfolio could be difficult to quickly liquidate. Although no such operational risks have come to pass in recent memory, concerns about them could still be powerful forces.

One manifestation of MMFs' aversion to concentration is that on quarter-ends, when some dealers window dress and cut back on repo borrowing, MMFs do not redistribute lending to non-window-dressing dealers. As Table 4 shows, on quarter-ends, the 10 dealers governed by regulations in the European Union (EU) and the United Kingdom (UK) cut back on repo borrowing,[31] both in dollar terms (Mod-

[29] There are regulated counterparty caps on certain assets that MMFs own (e.g., commercial paper). However, such caps do not apply to repo. Repo is treated as a "look-through" asset, so it is as if MMFs are holding the underlying Treasury collateral, which bears no cap.

[30] Regulations limit the tenor of securities held by MMFs to 397 days. Securities posted as repo collateral exceed this tenor requirement.

[31] UK banks reported their balance sheet size as quarter-end snapshots through 2015. During calendar-year 2016, UK banks reported month-

els (1) and (2)) and in percentage measures (Models (3) and (4)). On average, an EU or UK dealer reduces its repo borrowing by about \$5 billion on quarter-ends, even after controlling for the average repo borrowing in a given quarter (Model (2)). Yet the repo borrowing by dealers in Canada (CA), Japan (JP), and the US barely changes on quarter-ends. MMFs' apparent reluctance to shift their lending to CA, JP, or US dealers is not because these dealers offer dramatically lower repo rates on quarter-ends (Model (5)). Rather, this behavior suggests that MMFs are averse to lending too much of the portfolio to any one dealer. This aversion, coupled with MMFs' tendency to lend to the same dealers over time, means that MMFs may not nimbly respond to rate changes because shifting volumes may push MMFs against their concentration limits. This, in effect, gives dealers monopsony power over the MMFs that lend to them. Differently sized MMFs exhibit different levels of aversion. Therefore, dealers who borrow from different subsets of MMFs face supply curves of differing elasticity, providing a second reason for differential rates.

Hence, both MMFs' preference for stable investment and their aversion to concentration could contribute to the observed empirical patterns. To account for the effect of these forces in the equilibrium, I next build and estimate a structural model of the Triparty market. Such a model must account for the patterns in the data. A search model is unlikely to be useful, as the dispersion is observed in the context of repeated interactions between the same set of agents. A model that relies purely on linear utility from pecuniary returns is also likely insufficient, as the agent's optimal choice in a linear utility is to concentrate everything in a single best choice, which is at odds with the observed pattern of lending to multiple dealers. The first-order importance of dealers' identities in explaining rate dispersion, and the striking simultaneity in MMFs' lending, therefore, lead me to develop a model that features lenders with possibly concave utilities responding to posted, borrower-specific pricing.

4. Model

I now develop a model for borrowing and lending overnight cash via repo in the Triparty market. My aim is to construct a model that captures the key economic forces generating the distinct patterns in the data. The model should provide a reasonable representation of the market, so that it can be used to quantify the degree of competition.

The model has two types of agents interacting in the supply and demand of repo funding. On the supply side, lenders, e.g., MMFs, allocate their overnight cash with possible aversion to portfolio concentration and non-pecuniary preferences for borrowers. On the demand side, borrowers, i.e., dealers, act as monopsonies and set the repo rates used to attract funding.

end snapshots. Starting in 2017, UK banks reported quarter averages. UK banks' quarter-end window-dressing activities decreased starting in 2016.

4.1. The lender's problem

Let i index lenders and j borrowers. Lender i has a portfolio of investments with one-day maturity. At each time t, he chooses the share of this overnight portfolio going to each of the J borrowers, x_{ijt}. The share of the portfolio not lent out, x_{izt}, goes to the lender's outside option: safe investments that mature overnight and for which the lender harbors no concentration aversion.[32]

$$U(\mathbf{x_{it}};\omega,\alpha) = \max_{\mathbf{x_{it}}} \sum_{j=1}^{J} \frac{\omega_{ijt} R_{jt}}{\alpha_{it}} \{\exp(\alpha_{it} x_{ijt}) - 1\} + R_{zt} x_{izt},$$

$$\text{s.t.} \sum_{j=1}^{J} x_{ijt} + x_{izt} = 1, x_{i1t}, \ldots, x_{iJt} \geq 0.$$

The lender's utility is quasi-linear: linear in his portfolio allocation to the outside option, which earns a gross return of R_{zt}. His utility from lending to one of the J borrowers is possibly concave in the shares lent, with the degree of the curvature controlled by an aversion to concentration parameter $\alpha_{it} \leq 0$. The utility from lending to a borrower further depends on the gross return, R_{jt}, set by the borrower and taken as given by the lender, and the lender's non-pecuniary preference for that borrower, $\omega_{ijt} \geq 0$.

From the lender's first-order condition (FOC), the optimal share of the portfolio lent to borrower j is

$$x^*_{ijt} = \frac{\log(R_{jt}) + \log(\omega_{ijt}) - \log(R_{zt})}{-\alpha_{it}}. \tag{1}$$

The optimal share increases in the attractiveness of borrower j relative to the outside option (R_{zt}), where borrower j's attractiveness is a function of the repo rate she pays (R_{jt}) and the non-pecuniary preference she garners (ω_{ijt}).[33] At a given R_{jt}, R_{zt}, and ω_{ijt}, different lenders will allocate different shares based on their concentration aversion (α_{it}).

The concentration aversion, α_{it}, controls how quickly the lender's utility diminishes when lending to a given borrower. It therefore determines how distributed lender i's portfolio is and how i reacts to repo rate changes. Consider the extreme case of $\alpha_{it} \to 0$: lender i's utility becomes linear, and he would concentrate all of his lending into one single best repo investment. As α_{it} becomes more negative, the utility becomes more concave, compelling the lender to spread out his lending. Consequently, the lender lends concurrently to multiple borrowers, reflecting an aversion to concentration. Intuitively, if the lender is averse to lending too much to any one borrower, then when one of the borrowers raises her rate, the lender will not consolidate his lending to take advantage of this rate increase. The concentration aversion parameter, α_{it}, is thus intimately tied

[32] Examples of the lender's outside option include the RRP and Treasury bills. See Section 5.2.

[33] The FOC emphasizes the trade-off between returns from borrower j and the outside option, and assumes the between-borrower cross-elasticity to be 0. This modeling assumption is supported by the minimal substitution between borrowers observed on quarter-ends, when a subset of borrowers experience funding shocks. See *Discussion of empirical patterns* in Section 3.

to the lender's semi-elasticity. This relation becomes apparent by differentiating the lender's FOC (Eq. 1) with respect to the log of the repo rate. The optimal response in share to a (percent) rate change is exactly $\frac{\partial x^*_{ijt}}{\partial \log(R_{jt})} = -\frac{1}{\alpha_{it}}$. That is, if a borrower doubles the repo rate she offers, the lender who is lending to her would increase his lending by $-\frac{1}{\alpha_{it}}$ of his portfolio.

As documented in Fact 3, there is an empirical relationship between a lender's portfolio size and his aversion to concentration. I therefore parameterize α_{it} as

$$\alpha_{it} = \beta_0 + \beta_1 \cdot \sqrt{y_{it}},$$

where y_{it} is the size of the lender's overnight cash portfolio.[34,35] I take y_{it} as exogenous, as the overnight cash portfolio serves to meet an MMF's liquidity needs and tends to be a stable fraction of the MMF's overall AUM.

The attractiveness of lending to borrower j depends on the non-pecuniary preference ω_{ijt}. This preference affects both whether and how much i lends. The marginal utility of lending the first dollar to borrower j is $\frac{\partial U}{\partial x_{ijt}}\Big|_{x=0} = \omega_{ijt}R_{jt}$. Given that the lender's cash could otherwise earn a return of R_{zt}, lending to j occurs if and only if $\omega_{ijt}R_{jt} > R_{zt}$. Therefore, ω_{ijt} differentiates borrowers and determines to whom the lender lends. Moreover, because the utility from lending depends on an ω_{ijt}-scaled R_{jt}, differences in ω_{ijt} govern how much different borrowers can get from the same lender. I parameterize ω_{ijt} as

$$\begin{aligned}
&\omega_{ijt} = \chi_{ijt} \cdot (\nu_{ijt} + \epsilon_{jt}); \qquad (2)\\
&\chi_{ijt} \sim Bernoulli(Logistic(\rho_{ij} + \delta \log(y_{it}))),\\
&\nu_{ijt} \sim 1 + Gamma(shape = k, scale = \psi_j/k),\\
&\epsilon_{jt} \sim LogNormal(\frac{-\sigma^2}{2}, \sigma^2).
\end{aligned}$$

χ_{ijt} is a binary random variable that determines whether lender i has a nonzero preference for borrower j. Its mean depends on borrower-lender pair-specific parameters, ρ_{ij}. This parameterization is motivated by the high persistence in trading, and is necessitated by the fact that trading relationships in Triparty were established long before my sample began. χ_{ijt} further depends on the size of the lender's overnight cash portfolio, y_{it}, through δ. This provision allows lenders to lend to more borrowers as lenders grow in size.

If the lender has a nonzero preference for a borrower, then the intensity of the preference, ν_{ijt}, is drawn from a gamma distribution whose mean depends on public and borrower-specific ψ_j.[36] Thus, in reduced-form, ψ_j captures the systematic variations in ω_{ijt}, conditional on lending. Perceptions of a high non-pecuniary preference could prompt the lender to either lend at a lower rate or lend more at a given rate. I speculate that variations in ψ_j are driven by a preference for consistent borrowing. This interpretation is supported by the results in Table 7, which show a strong correlation between the ψ_j parameters I recover and measures of borrower consistency (see Section 6 for more discussion).

Finally, the model explicitly accounts for possible borrower-time-specific shocks to the lender's preference, which are known to market participants but not to the econometrician. These "supply shocks", ϵ_{jt}, if present, threaten the ordinary least squares (OLS) identification of the relationship between rate and volume, because these shocks affect the observed lending volume without having observable proxies that one can use to control for their effect.

4.2. The borrower's problem

At each t, borrower j maximizes her profit by choosing the gross repo rate R_{jt} that she pays to all lenders for repo funding:

$$\max_{R_{jt}} \left[S_{jt}(Q_{jt}) - R_{jt}\right] \cdot Q_{jt}(R_{jt}),$$

where $Q_{jt}(R_{jt}) = \sum_i [\mathbf{E}[x_{ijt}(R_{jt})] \cdot y_{it}]$ is the total quantity of funds borrower j expects to obtain at rate R_{jt}, and $S_{jt}(Q_{jt})$ is the average value of intermediating funds at Q_{jt}.

Triparty borrowers demand repo funds because these funds can be intermediated to generate value. For example, the funds could be lent out (via repo again) to a borrower's clients, such as hedge funds, that do not have access to the Triparty market but need to finance security holdings using repo. Alternatively, the repo funds could be used to finance a borrower's own security holdings, such as those purchased during a Treasury security auction. The value that a borrower requires to intermediate her repo funding, S_{jt}, could depend on the total amount of funds that she obtains. Importantly, this intermediation value reflects the pure economic benefit accruing to the borrower and is thus *net* of any opportunity cost from taking up balance sheet space. Balance sheet costs also matter for asset prices (see Duffie and Krishnamurthy (2016) and Du et al. (2022a)). In fact, the full cost of repo funding a borrower's clients face should be the sum of the modeled value of intermediation and the unmodeled balance sheet cost. Section 6 explores this relationship in depth.

Differentiating the borrower's problem with respect to the repo rate, the FOC yields that borrower j's optimal repo rate is

$$R^*_{jt} = \underbrace{S'_{jt} \cdot Q_{jt} + S_{jt}}_{\text{marginal value of intermediation at Q}} - \underbrace{\frac{Q_{jt}}{Q'_{jt}}}_{\text{markdown}}. \qquad (3)$$

The optimal rate offered by the borrower could embed a markdown from her marginal value of intermediation.

[34] The results in Table 3 suggest that the relationship between portfolio construction and portfolio size is concave. The square root transformation of y_{it} captures this concavity while ensuring positivity. The estimation results are robust to alternative concave transformations; see Appendix Section Appendix A.

[35] Concentration aversion (α_{it}) is modeled to not depend on borrower attributes because this aversion is likely motivated by fears of operational disruption, and collateral pledged by borrowers is comparable and would not be differentiated in fire sales. Independence of α_{it} from borrower attributes is helpful though not necessary for identification.

[36] The choice of gamma ensures positive preferences and gives flexibility in fitting the data. If *shape* parameter k is large, the gamma distribution approximates Normal; if k is small, then the gamma distribution approximates Exponential.

The magnitude of the markdown is therefore a measure of the borrower's market power. This markdown is a direct function of the funding supply the borrower faces, $Q_{jt}(R_{jt})$. At a given quantity Q_{jt}, if the supply is highly responsive to rate changes, then Q'_{jt} would be large and the markdown would be small. Conversely, the borrower can build in a large markdown if the lenders' response to her repo rate changes is small. To measure the borrower's ability to set a markdown — or the extent of her market power — I therefore must find the lenders' concentration aversion (α_{it}) and non-pecuniary preferences (ω_{ijt}, captured by ψ_j).

4.3. Model discussions

The lender's and the borrower's problems together describe the equilibrium of the Triparty market. In trying to capture the most salient features of the data, the model deliberately de-emphasizes certain considerations, which I discuss below.

Concentration aversion among borrowers. The model's inclusion of the lender's aversion to concentration is motivated by the fact that MMFs (lenders) maintain systematically distributed portfolios. This observation alone could also stem from dealers' (borrowers') having an aversion to concentration, possibly out of a desire for a diversified funding base. Yet even if present, concentration aversion among dealers is unlikely a dominant force. Aversion to concentration means that, on the margin, the utility of one more dollar invested with the same counterparty is not as high as previous dollars. If dealers harbored strong concentration aversion, then they would optimally attract marginal new lenders with better rates to achieve diversification. The empirical data show no statistically significant distinctions in rates received by MMFs lending to the same dealer (Fact 2 in Section 3). I therefore choose to emphasize lenders' concentration aversion in the economic model.

The fact that this aversion is pronounced in MMFs but not dealers suggests that it is likely caused by considerations unique to MMF. One such consideration is MMFs' inability to hold long-duration assets, which leads to worries over operational risks and fire sales; see also *Discussion of empirical patterns* in Section 3.

The role of private information. Information in this model is mostly public. The only private information that the lender has is the realization of the intensity of his non-pecuniary preference (ν_{ijt} in ω_{ijt}). These realizations affect equilibrium volume but not rates because the rate is set by the borrower using the mean level of the preference, ψ_j, which is public knowledge.[37] The lender's preferences in this model thus contrast with those in models of bank lending relationships, which tend to reflect an informational advantage (e.g., through better monitoring) that often leads to preferential rates (e.g., Darmouni (2020)).

The model here downplays the role of relationships and private information for two reasons. First, the amount of private information is likely limited because the advantage of superior information is muted when lending is standardized, secured against high-quality collateral, and renewed daily. Second, even if private relational information exists, it does not seem to be of first-order importance. The fact that each dealer borrows at a distinct price from all MMFs suggests that dealers' public attributes determine pricing. I therefore intend for the lender's non-pecuniary preferences in the model to reflect the borrower's public attributes in reduced-form. An example of the borrower's public attributes is her borrowing consistency; see *Discussion of empirical patterns* in Section 3. My modeling choice underscores that market power in the Triparty market does not come from information asymmetry.

[37] The borrower *will* take into account borrower-time specific preference shocks, which are unobservable to the econometrician but known to both the borrower and the lender.

5. Estimation

In this section, I outline the estimation strategy that separately quantifies the two key parameters necessary to compute the borrower's markdown: α_{it} and ω_{ijt} (captured by ψ_j). I discuss sources of variation, measurement of R_{zt}, the instrumental variable analysis, and the indirect inference and maximum likelihood estimation approach I use to estimate all model parameters.

5.1. Sources of variation

From the lender's FOC (Eq. 1), I know that both α_{it} and ω_{ijt} can affect how much lender i lends to borrower j (x_{ijt}) at a given repo rate (R_{jt}). It is possible to separate their effects because α_{it} captures differences across lenders and ω_{ijt} centers on means that vary across borrowers. Specifically, comparing lending to the same borrower by different lenders can provide information about the relative magnitude of α_{it}. Similarly, comparing lending received by two different borrowers from the same set of lenders can reveal the relative magnitude of ψ_j.

Cross-sectional comparisons alone are insufficient to pin down the level of these parameters. Instead, I also leverage the direct relationship between α_{it} and lenders' semi-elasticity. As discussed in Section 4.1, lender i's portfolio allocation response to borrower j's repo rate change depends on the lender's α_{it}. Hence, the supply schedule each borrower faces is a function of all the individual α_{it} associated with the lenders that lend to her. Estimating lenders' supply semi-elasticity can garner information about the average level of α_{it}, which in turn provides information about the levels of ψ_j.

From estimated α_{it} and ψ_j, I can calculate borrowers' markdowns. The borrower's unobserved marginal value of intermediation is then the sum of the estimated markdown and the observed borrower repo rates (Eq. 3).

5.2. Measuring R_{zt}, the lender's outside option

In the lender's FOC (Eq. 1), the lending decision directly depends on the difference between a borrower's offered repo rate (R_{jt}) and the return on the lender's outside option (R_{zt}). An outside option is a safe, overnight investment, for which the lender harbors no concentration aversion. Two plausible outside options are the Fed's RRP and the Treasury bills. I measure R_{zt} as the higher of the RRP rate or the 1-day Treasury bill yield.

The RRP is a fixed-rate, full-allotment facility that allows lenders to lend to the Fed via repo (see also Section 2.1). The rate offered by the RRP provides a good measure of the lender's outside option. MMFs can also invest in Treasury securities that have a maturity of less than one year. Similar to placing repo at the RRP, buying Treasury securities is investing with the U.S. government and is an outside option to lending to Triparty dealers. However, there is no reported overnight Treasury yield. I thus impute a 1-day Treasury yield by adjusting for the term structure using the 1-day and the 1-month overnight indexed swap (OIS).

I generate the time series of R_{zt} as the 1-day Treasury bill yield before September 2013, when the RRP was introduced, and the RRP rate thereafter; the RRP rate is always higher than the 1-day Treasury bill yield in my sample. In my sample, the correlation between the 1-day Treasury bill yield and the median Triparty repo rate is 0.78 before the introduction of the RRP and 0.13 thereafter, supporting my choice of R_{zt}.

The introduction of the RRP in 2013 was followed by a year of testing, during which the RRP had a constraining counterparty cap (see Section 2.1). In the estimation, I purposely leave out this testing period of September 2013 to September 2014, as the counterparty cap makes the true marginal outside option difficult to ascertain. I also exclude all quarter-end months because many regulations are enforced only on quarter-ends. Numerous studies have focused on quarter-ends to study how regulations distort markets (e.g., Du et al. (2018), Wallen (2020)). My study aims to reveal the extent of imperfect competition even outside of quarter-ends. The final estimation sample therefore consists of 48 month-ends from 2011 through 2017.

5.3. Instrumental variable

Lenders' supply semi-elasticity is an essential ingredient in my estimation strategy. However, I cannot measure the relationship between rate and volume using OLS due to preference shocks that are unobserved by the econometrician (ϵ_{jt} in Eq. 2). For example, if negative preference shocks occur, a borrower will be seen to offer a high repo rate but will attract only a modest amount of funds, biasing the true relationship to 0.

I therefore estimate the inverse supply semi-elasticity by using an instrumental variable that shocks the borrowers' funding needs. The U.S. Treasury Department periodically auctions marketable debt securities of various maturities. Dealers bid, make markets, and take speculative positions around Treasury auctions (Fleming and Rosenberg, 2008), and they typically finance their Treasury holding with repo. Consequently, the amount of Treasury securities auctioned likely correlates with how much borrowers want to borrow. Using Treasury auctions as an instrument for borrowers' demand for funding, I can find by how much borrowers need to raise their repo rates to attract additional funding.

Specifically, I construct the instrument as the amount of *non-bill* Treasury securities *offered* at auction such that they *settle* on the same days as MMFs' reporting dates. On these *settlement* dates, titles transfer and dealers must finance their acquisitions. Repo volumes on settlement days are therefore mostly directly impacted by Treasury auctions. To avoid potential endogeneity between repo rates and how much dealers decide to purchase, I focus on the amount of Treasury securities *offered* for sale, which likely reflects the Treasury Department's fiscal needs and is plausibly exogenous to borrower-specific preference shocks.[38] Finally, I include only auctions of Treasury securities with maturities of 1 year or more. *Non-bill* securities cannot be purchased by MMFs. Auctions of these securities thus do not change the lenders' trade-offs.[39]

Table 5 summarizes the instrument-induced inverse semi-elasticity. All regressions in the table are run at the borrower-time level, as borrowers set borrower-time-specific repo rates. Because identification relies on shocks that impact borrowers at each point in time, standard errors are clustered by time (month).

Models (1) and (2) of Table 5 report the first-stage impact of Treasury auction offers on repo volume, as in $Vol_{jt} = \beta_{1st} TreasuryOffer_t + BorrowerFE + YearFE + e_{1st,jt}$. Model (1) shows a strong correlation between the amount of Treasury securities offered in auctions and the amount of Triparty repo funding obtained by borrowers. As there may be macroeconomic shocks that affect both the Triparty repo market and the Treasury Department's decision to issue debt, I add year fixed effects in Model (2). Thus, the preferred instrument of Model (2) relies on auction variations within the calendar year, which typically reflect tax revenue fluctuations.[40] The magnitude of the volume response reduces from 46.8 to 16.3 but is still significant at the 5% level. As I measure Treasury auction offers in trillions of dollars, the estimated coefficient implies that a \$40 billion, or one-standard-deviation, increase in the amount offered in Treasury auctions is associated with an average increase of \$0.65 billion in a borrower's repo funding.

Model (3) of Table 5 reports the repo rate response to instrumented volume change, as in $\log(R_{jt}) - \log(R_{zt}) = \beta^{IV}\widehat{Vol}_{jt} + BorrowerFE + YearFE + e_{IV,jt}$.[41] The estimated coefficient shows that to raise \$1 billion more in repo funding, a borrower needs to increase her repo rate by 1.57 bps. In other words, a 1 bp increase in the repo rate is associated with a \$0.64 billion increase in funding. For

[38] Another potential endogeneity comes from macroeconomic shocks that could simultaneously affect Treasury's auction offers and borrowers' purchases. To allay this concern, year fixed-effects are included in the estimation.

[39] In unreported robustness checks, I confirm that the amount of non-bill securities offered in Treasury auctions is not significantly correlated with either the level or the volatility of MMF AUM flows.

[40] The year fixed effects are, specifically, separate indicator variables for 2011, 2012, 2015, 2016, and 2017, and one indicator variable for the first 6 non-quarter-end months in 2013 and the last 2 non-quarter-end months in 2014. The two calendar months included in 2014 are the two calendar months missing from 2013, which completes the fiscal year. Robustness checks using separate fixed effects for 2013 and 2014 show similar results but are more noisily estimated.

[41] I obtain borrower-time-specific repo rates (R_{jt}) by volume-weighting the observed borrower-lender pair repo rates for Treasury-backed repo with a 2% haircut.

Table 5
Inverse semi-elasticity using Treasury auction IV.

	1st stage: vol_{jt}		IV: $R_{jt} - R_{zt}$	Alt. 1st	Alt. IV	OLS
	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Non-bill Treasury offer amount	46.79***	16.29**				
	(14.66)	(6.76)				
Treasury offer * borrower share				241.64**		
				(97.20)		
Vol_jt (fit)			1.57**		1.57**	
			(0.67)		(0.61)	
Vol_jt						0.01
						(0.01)
Dealer FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	No	Yes	Yes	Yes	Yes	Yes
Cluster-robust F-stat		5.81		6.18		
Anderson-Rubin 95% CI			(0.5, 8.6)		(0.6, 7.9)	
Num. obs.	821	821	821	821	821	821

Standard errors in parentheses.
Notes: This table reports the instrumental variable estimations of the inverse semi-elasticity borrowers face. Models (1) and (2) are first-stage regressions of the borrower's overnight repo volume on the amount of non-bill Treasury securities offered at auctions and settled on the same day as MMF N-MFP reporting dates. Model (3) regresses the difference between the borrower's repo rate and the outside option, on the borrower's overnight repo volume, as instrumented using Model (2). Models (4) and (5) are similar to Models (2) and (3) but use as the instrument the product of Treasury auction offers (as defined above) and each borrower's average share of the total Triparty overnight repo volume. Model (6) regresses the borrower's repo rate on volume without using an instrument. The borrower's repo rate is defined as the volume-weighted average of repo rates between a borrower and all lenders in overnight repo collateralized by Treasury securities with a 2% haircut. The outside option is defined as the imputed 1-day Treasury bill yield before September 2013 and the rate on the RRP thereafter. The estimation period is January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter endss. Standard errors are clustered by month (frequency of observation). *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

the average borrower, this is 3.6% of her funding. This estimate is comparable to recent events in the Triparty market. When the Fed unexpectedly raised the RRP rate by 5 bps on June 17, 2021, the RRP saw an overnight inflow of $225 billion from a base of $1628 billion in Treasury-backed Triparty repo,[42] implying a semi-elasticity of 2.9% per 1 bp change. At the same time, this estimate is higher than comparable estimates for the Treasury bills market in Greenwood et al. (2015),[43] Duffee (1996),[44] and (Bernanke et al., 2004),[45] suggesting that the supply on Triparty is more inelastic.

The first-stage specification in Model (2) of Table 5 features a market-wide instrument that applies to all Triparty borrowers, $TreasuryOffer_t$. The IV estimate in Model (3) is therefore the average rate response to the average induced volume. Borrowers may have heterogeneous volume responses to Treasury auction offers. If I knew the borrower-specific participation rate in Treasury auctions, I could refine my instrument to capture individual shocks that are the product of Treasury auction offers and individual auction participation. Unfortunately, these data are not publicly available. In Models (4) and (5), I run a version of this heterogeneous-response IV by using a borrower's average repo share as a proxy for her auction participation. Specifically, I calculate each borrower's share of the overall Triparty repo volume at each point in time and take the time series average to arrive at a time-invariant borrower share. There are two assumptions behind using the product of Treasury auction offerings and borrower repo share as an instrument. First, borrowers that are more active in repo are also more likely to participate in Treasury auctions. Second, because this share is time-invariant, it is not correlated with errors in the IV regression. The estimated inverse semi-elasticity from Model (5) is 1.566, very similar in magnitude to the point estimate of 1.573 in Model (3).

The precision of the IV estimation depends on the strength of the instrument. The cluster-robust effective F-statistic of the instrument in Model (2) of Table 5 is 5.8, below the rule-of-the-thumb threshold of 10. To better understand the implications of using a potentially weak instrument on the IV inference, I compute the Anderson-Rubin confidence interval. This confidence interval has the correct coverage regardless of the strength of the instrument and is efficient in just-identified models with a single instrument (Andrews et al., 2019), as is the case here. The 95% Anderson-Rubin confidence interval for the IV coefficient estimate in Model (3) is (0.5, 8.6). See Appendix Fig. A2 for more details. This interval is bounded away from the imprecise and near-zero OLS estimate in

[42] The closest public release of the size of the Triparty market is as of June 9, 2021: https://www.newyorkfed.org/data-and-statistics/data-visualization/tri-party-repo.

[43] Greenwood et al. (2015) use an IV approach on a 1983–2007 sample. They find that a one-percentage-point decrease in $\frac{\Delta Treasury}{GDP}$ leads to a 38.6 bps decrease in the two-week Treasury yield. The average annual GDP between 2011 and 2017 was $18.7T. Together, these figures imply that a $1b increase in the supply of Treasury bills increases the yield by 0.21 bps.

[44] Using data for each January from 1983 to 1994, (Duffee, 1996) estimates that a 1% increase in 1-month Treasury bills outstanding increases the yield by 1.012 bps. The average Treasury bill outstanding over my sample period is $1.6 trillion, of which roughly 30% is due within a month. This implies that, again, a $1b increase in 1-month Treasury bills outstanding increases the yield by 0.21 bps.

[45] Bernanke et al. (2004) uses Japanese purchase of Treasury securities to estimate that a $1b reduction in Treasury securities outstanding decreases the yield on 3-month Treasury by 0.18 bps and on 2-year Treasury by 0.55 bps.

Model (6), suggesting that the instrument is useful. At the same time, this interval is very wide in the other direction. In other words, I am reasonably confident that the instrumented semi-elasticity is not zero; however, I am much less certain that the true value is not larger. A larger estimate would mean that borrowers need to raise their rates even more in order to induce additional volume, implying an even more inelastic supply.

5.4. Estimation approach

I estimate the parameters of the lender's model using a mixture of indirect inference (Gourieroux et al., 1993) and maximum likelihood.

Applying indirect inference, I choose parameters such that the data simulated by these parameters would generate moments matching those generated from the original data. I include three groups of moments that summarize the distinct data patterns discussed so far. First, because the size-dependent concentration aversion parameter, α_{it}, controls the lender's response to rate changes, my moments include the IV coefficient on $\widehat{Vol_{jt}}$ from Model (3) of Table 5. This coefficient is a direct function of α_{it} — $\beta_{IV} = \frac{1}{T}\frac{1}{J}\sum_{t\in T}\sum_{j\in J}(\sum_{i\in x_{ijt}>0} - \frac{y_{it}}{\alpha_{it}})^{-1}$ — and can inform the average level of α_{it}. The parameter β_1 in α_{it} governs the dependence of α_{it} on the lender's portfolio size. I therefore include as a moment the coefficient from regressing the lender's median portfolio share on portfolio size (Model (2) of Table 3). Next, because ψ_j in ω_{ijt} reflects borrower-specific influences on portfolio allocation, I include as moments each borrower's conditional average lender share and unconditional average probability to borrow.[46] Finally, as σ^2 and k (*shape*) determine the variance of ϵ_{jt} and ν_{ijt}, respectively, they determine how much variation in the observed data can be explained by the included model parameters. I use the R^2 from regressing portfolio shares on lender portfolio size and borrower fixed effects,[47] and on lender portoflio size and borrower-time fixed effects[48] to learn about these two parameters.

The χ_{ijt} in ω_{ijt} controls whether a borrower can attract any lending. I recover the parameters of χ_{ijt} by maximizing the proportion of correctly predicted lending occurrences between each pair at each time. Given the logistic transformation of the underlying parameters, the estimation of pair-specific ρ_{ij} poses a potential incidental parameter problem. I apply the analytical bias correction suggested by Hahn and Newey (2004) to address this concern. The difference between the bias-corrected estimates and the simple maximum likelihood estimates are small because the estimation sample is large ($T = 48$ for most pairs).

The parameters of my model are over-identified. I weigh the moments using the inverse of the variance-covariance matrix for moment conditions calculated in bootstrapped samples. The bootstraps are done in blocks of time (month) clusters, in accordance with the IV regression.

6. Results and discussions

In this section, I present estimation results and discuss the derived dealers' (borrowers') markdowns. I consider in particular the relationship between these markdowns and various funding spreads, and the role of policy in shaping intermediary competition.

6.1. Parameter estimates

Table 6 summarizes all parameter estimates, along with their time-clustered block-bootstrapped confidence intervals.[49] Data simulated from these parameter estimates generate both targeted and untargeted moments that closely match those found in the original data, suggesting that the estimated model provides a reasonable representation of actuality. Appendix Section Appendix B discusses the model fit in detail.

I use the estimated β_0 and β_1 to calculate α_{it}, which has a median of -0.042, a mean of -0.043, and an interquartile range of (-0.053, -0.032). Because the possible support of α_{it} is $[-\infty, 0)$, the fact that α_{it} is bounded away from 0 shows that lenders do exhibit an aversion to portfolio concentration. Based on calculated α_{it}, dealers on average need to raise repo rates by 31 bps to increase their funding by 100%. The negative β_1 indicates that the aversion to concentration increases with size. This result coheres with the notion that fire sale risks induce the lender to spread out his portfolio, and larger lenders are more sensitive because — should operational disruptions force the lender to take possession of the posted collateral — even a small share of a large portfolio could be costly to liquidate.

The estimated ψ_j (capturing ω_{ijt}) has a median of 25 bps, a mean of 27 bps, and an interquartile range of (18 bps, 34 bps). One way to interpret these estimates is to recall that ψ_j reflects the lender's marginal utility from lending the *first* dollar. All else being equal, the dealer at the 25th percentile would have to offer 16 bps more than the dealer at the 75th percentile to attract the same first dollar of lending.[50]

Yet ψ_j is not merely ranking dealers based on rates because dealers seldom stop after getting that first dollar. For example, although Goldman Sachs pays among the lowest rates, its estimated ψ_j is far from being the highest. This is because Goldman Sachs cannot borrow much at the low rate it offers. In contrast, dealers such as Wells Fargo are

[46] Unconditional probabilities are necessary to inform ψ because for some observations the borrower's repo rate (R_{jt}) is less than the return on the outside option (R_{zt}). To rationalize these observations, not only would χ_{ijt} need to take on a value of 1 (as opposed to 0) but ψ_j would also need to be sufficiently large.

[47] $x_{ijt} = b_{1,\sigma}\log(y_{it}) + BorrowerFE + e_{\sigma,ijt}$, where $\log(y_{it})$ absorbs the effect from α_{it} and *BorrowerFE* absorbs the effect from ψ_j

[48] $x_{ijt} = b_{1,k}\log(y_{it}) + BorrowerMonthFE + e_{k,ijt}$, where $\log(y_{it})$ absorbs the effect from α_{it} and *BorrowerMonthFE* absorbs the effect from ψ_j and ϵ_{jt}

[49] Estimates of the pair-specific ρ_{ij} are omitted from Table 6 for brevity but are available upon request.

[50] In the model, ω_{ijt} (which has conditional mean of ψ_j) enters the lender's utility as a multiplier for gross repo rates. Here, I suggest an additive increase heuristically because gross repo rates are close to 1 and the first-order condition is based on the logs of R_{jt} and ω_{ijt}.

Table 6
Estimates of model parameters.

Parameter	Estimate	95% CI
Concentration aversion		
α_{it}		
β_0	0.30	(-1.09, 1.1)
β_1	-9.48	(-10.26, 9.17)
Non-pecuniary preference		
ω_{ijt}		
ν_{ijt}		
ψ_j		
Barclays	25.70	(24.02, 29.42)
BNP Paribas	31.82	(28.42, 35.46)
Bank of America	29.76	(27.72, 32.20)
Citi	19.60	(17.97, 22.60)
Crédit Agricole	39.25	(37.61, 45.51)
Credit Suisse	18.56	(16.95, 21.45)
Deutsche Bank	34.93	(32.67, 39.14)
Goldman Sachs	15.54	(14.58, 17.84)
HSBC	18.22	(16.46, 19.94)
JP Morgan	16.37	(14.60, 18.67)
Mitsubishi	21.08	(19.79, 24.42)
Natixis	45.27	(42.36, 51.26)
Nomura	58.37	(53.29, 64.55)
Nova Scotia	17.27	(15.40, 19.70)
Royal Bank of Canada	14.92	(13.61, 17.82)
Royal Bank of Scotland	25.25	(22.21, 27.97)
Société Générale	28.54	(26.73, 32.32)
Sumitomo	39.69	(34.29, 48.75)
UBS	13.99	(11.10, 17.59)
Wells Fargo	33.23	(31.32, 36.48)
k (shape)	1.51	(1.43, 1.68)
ϵ_{jt}		
σ^2	6.10	(4.59, 12.4)
χ_{ijt}		
δ	0.63	(0.5, 0.75)

Notes: This table reports the parameter estimates in the lender's problem. β_0, β_1 are parameters of $\alpha_{it} = \beta_0 + \beta_1 \cdot \sqrt{y_{it}}$, and are scaled by 1×10^{-3}. ψ_j, k are parameters of $\nu_{ijt} \sim 1 + Gamma(shape = k, scale = \psi_j/k)$; the random variable defined by *Gamma* is scaled by 1×10^{-4}. σ^2 is the parameter of $\epsilon_{jt} \sim LogNormal(\frac{-\sigma^2}{2}, \sigma^2)$; ϵ_{jt} is scaled by 5×10^{-5}. δ is the parameter of $\chi_{ijt} \sim Bernoulli(Logistics(\rho_{ij} + \delta \log(y_{it})))$. Estimates of the pair-specific ρ_{ij} are omitted for brevity but are available upon request. With the exception of δ, parameters in this table are estimated using indirect inference, where the 95% confidence intervals are block bootstrapped by time (month). δ is estimated using maximum likelihood, with the Hahn and Newey (2004) analytical bias correction applied for potential incidental parameter problems. The model estimation period is January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter-ends.

estimated to have higher ψ_j than Goldman Sachs because they can borrow a lot more by paying a slightly higher rate. In short, ψ_j conveys how preferred a dealer is perceived to be by lenders, and this preference can be reflected in a combination of rates and volume.

One possible explanation for lenders exhibiting dealer-specific preferences is that dealers vary in how consistently they take on repo loans. I test this hypothesis by proxying a dealer's borrowing consistency with its average coefficient of variation in volume vis-à-vis lenders: $CoefVar_j = mean_j(\frac{SD_{ij}(vol_{ijt})}{mean_{ij}(vol_{ijt})})$. The lower the coefficient of variation, the more consistent a dealer is in using its balance sheet to take on repo loans. Fig. 5 shows that the estimated ψ_j is strongly and negatively correlated with the dealer's average coefficient of variation. In Table 7, I explore the correlation between estimated ψ_j and the average and median of the dealer's coefficient of variation (Models (1) and (2)), and between estimated ψ_j and the dealer's creditworthiness as measured in CDS rates (Models (3)–(5)). The conventional CDS contract varies by jurisdiction.[51] Even after controlling for jurisdictions, that is, comparing estimated ψ_j with CDS rates among dealers within the same jurisdiction, CDS rates still do not appear to be a significant predictor of ψ_j (Models (4) and (5)). In contrast, measures of dealer consistency are strongly correlated with estimated ψ_j.

6.2. The dealer's markdown

Having estimated α_{it} and ω_{ijt} in the lender's problem, I can now calculate the dealer's markdown in Eq. 3. In the cross-section, the time-series average of each dealer's markdown has an interquartile range of (23 bps, 39 bps; see Table 8). Because the S_{jt} in the model, representing the dealer's value of repo intermediation, is net of balance sheet costs, these markdowns are the economic rents dealers extract. Considering that the mean dealer repo volume is about \$18b per day, these estimates imply a midspread of mean annual profits that ranges from \$41m to \$70m.

The cross-sectional variation in markdowns arises from the interaction between α_{it} and ω_{ijt}. The preference for investment stability differentiates lending to different dealers, making each dealer a local monopsony over its funding needs. How much rent a dealer can then extract depends on how easily lenders can substitute away and is thus linked to lenders' concentration aversion. As shown in Table 8, dealers differ in their demand consistency for repo loans and thus draw varying degrees of non-pecuniary preferences from the lenders (ψ_j, capturing ω_{ijt}). Dealers further differ in the set of lenders with whom they trade and therefore face supply curves of varying elasticity ($f(\alpha_{it})$). Goldman Sachs, for example, commands a mean non-pecuniary preference (ψ_j) that ranks at the 10th percentile, a reasonable value partially reflecting the fact that lenders are willing to lend to Goldman Sachs at low rates. Yet because Goldman faces a group of lenders that have low aversion to concentration, its repo funding supply is highly elastic. The combination of these two forces results in Goldman Sachs having one of the lowest markdowns.

Remarkably, the low rate that Goldman Sachs pays is not the result of wedging in a large markdown; rather, it is indicative of Goldman Sachs' high opportunity cost from repo intermediation. Regulations, especially those imposed in the wake of the 2007-09 Financial Crisis, require banks to maintain a minimum capital ratio against all assets. Because a bank's equity is fixed in the short run, repo intermediation activities, which increase the bank's total asset size, compete with other potential trades for balance sheet space (Duffie, 2017). Goldman Sachs likely has more lucrative trades that it can do with its limited balance sheet space, making its post-balance-sheet-cost value of

[51] The most common CDS terms are no restructuring (XR) in the U.S. and Canada, modified restructuring (MM) in the EU, and full-restructuring (CR) in Japan.

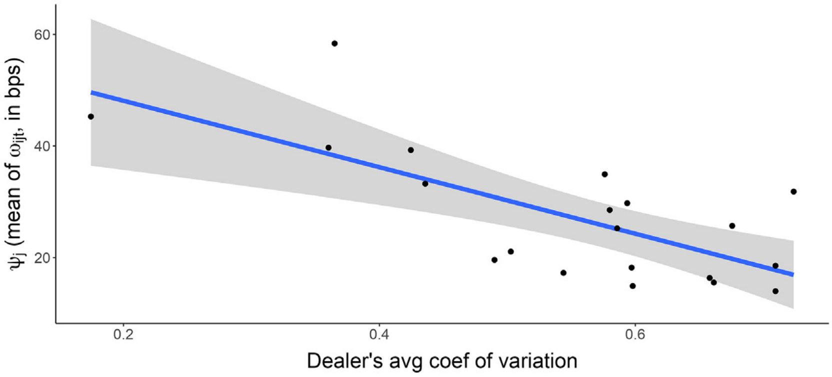

Fig. 5. Correlation between estimated dealer preference and coefficient of variation in volume.
Notes: This figure plots estimated ψ_j against the dealer's average coefficient of variation. The blue line represents the fitted value from regressing ψ_j on the dealer's average coefficient of variation, and the shaded regions represent the 95% heteroskedasticity-robust confidence bands. The dealer's average coefficient of variation is the by-dealer average of the dealer-lender coefficient of variation in repo volume throughout the model estimation period. The model estimation period is January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was first introduced and was in testing, and excluding months that fall on quarter-ends. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 7
Estimated ψ_j versus borrowing consistency and CDS.

	Estimated ψ_j				
	Model 1	Model 2	Model 3	Model 4	Model 5
Constant	59.985***	53.011***	18.151***		
	(8.399)	(7.164)	(5.747)		
Average coef of variation	-59.492***				
	(14.281)				
Median coef of variation		-48.771***			
		(12.686)			
Average CDS: last 3 days of month			0.225	0.338	
			(0.137)	(0.239)	
Average CDS					0.321
					(0.232)
Dealer HQ FE	No	No	No	Yes	Yes
Num. obs.	20	20	20	20	20
R^2	0.501	0.413	0.104	0.486	0.480

Standard errors in parentheses.
Notes: This table reports the regression of the estimated preference parameter, ψ_j (capturing ω_{ijt}), on measures of dealers' borrowing consistency and creditworthiness. "Average coef of variation" is the average of a dealer's coefficients of variation in volume vis-à-vis all lenders: $\overline{CoefVar}_j = mean_j(\frac{SD_{ij}(vol_{ijt})}{mean_{ij}(vol_{ijt})})$. "Median coef of variation" is the median of a dealer's coefficients of variation. "Average CDS on last 3 days of month" is the average dealer credit default swap rate on the last 3 business days of each month during the model estimation period. "Average CDS" is the average of a dealer's CDS rate over the model estimation period. CDS rates are for 6M debt for all dealers except for Mitsubishi and the Royal Bank of Scotland, which only have 5Y CDS. CDS rates are for contracts in local currency (except for the Canadian banks, whose CDS are in U.S. dollars), and follow the most common CDS convention, which is no restructuring (XR) in the U.S. and Canada, modified restructuring (MM) in the EU, and full-restructuring (CR) in Japan. The model estimation period is January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter-ends. Standard errors are robust to heteroskedasticity. *, **, and *** denote significance at the 10%, 5%, and 1% confidence levels, respectively.

repo intermediation (S_{jt}) rather modest. This situation rationally leads Goldman to price its repo aggressively and to borrow sparingly and opportunistically. Although Goldman pays among the least for repo funding, it does not earn much rent because of its limited trading volume. This point is reminiscent of the model of Berk and Green (2004), in which mutual fund managers' skills are reflected not in excess returns alone but in the fund-size-adjusted alpha.

6.3. Markdowns and funding spreads

Taking the median dealer markdown at each point in time to be representative, I next construct the time series of dealer markdowns through the estimation period. Fig. 6 shows that the average of the median markdown over my estimation period is 26.2 bps. Compared to the 5.7 bps average spread between the median Triparty repo rate and

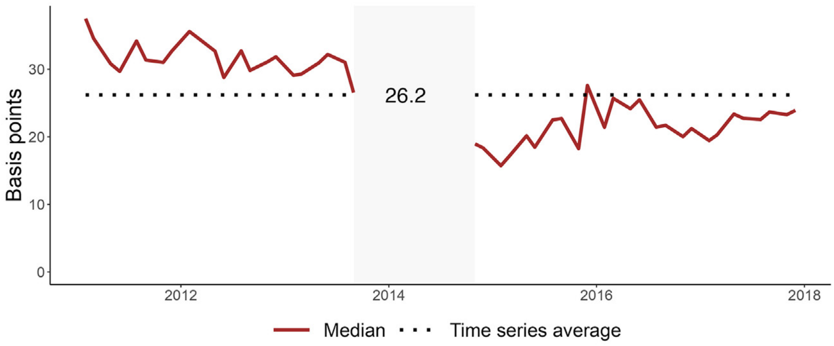

Fig. 6. Triparty dealer markdown.
Notes: This figure plots the time series variation in the median dealer markdown over the model estimation period (solid red line). The dotted black line indicates the average of this value over the whole sample. The shaded area corresponds to September 2013–September 2014, when the RRP was in testing. The model estimation period is January 2011 to December 2017, excluding September 2013–September 2014 and months that fall on quarter-ends. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 8
Dealers' average non-pecuniary preference, elasticity, and markdown.

Dealer	Preference (ψ_j)	Elasticity (%)	Markdown (bps)
Bank of America	25.7	2.5	35.7
Barclays	31.8	2.7	33.2
BNP Paribas	29.8	2.5	39.0
Citi	19.6	4.1	23.9
Crédit Agricole	39.3	2.3	42.7
Credit Suisse	18.6	3.5	25.4
Deutsche Bank	34.9	2.5	32.3
Goldman Sachs	15.5	6.6	14.9
HSBC	18.2	4.0	24.2
JP Morgan	16.4	4.0	22.7
Mitsubishi	21.1	4.6	23.8
Natixis	45.3	2.0	54.5
Nomura	58.4	1.4	70.9
Nova Scotia	17.3	5.1	23.4
Royal Bank of Canada	14.9	5.9	16.9
Royal Bank of Scotland	25.2	3.6	22.9
Société Générale	28.5	3.1	33.7
Sumitomo	39.7	2.1	50.8
UBS	14.0	4.7	18.5
Wells Fargo	33.2	2.3	40.6

Notes: This table reports three time-series averages for each dealer in the sample. "Preference" is the estimated ψ_j, which is the dealer-specific mean of lenders' non-pecuniary preference, ω_{ijt}. "Elasticity" is the percentage of repo volume a dealer would attract if it raised its repo rate by 1 bp. "Markdown" is the difference between the estimated dealer's marginal value of intermediation and the repo rate it pays. The analysis in this table is based on parameters estimated in the model estimation period of January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter-ends.

the lenders' outside option, dealers extract 82% of the (26.2 + 5.7 =) 31.9 bps total surplus in the Triparty market.[52]

The presence of markdowns means that dealers pay less for obtaining funding than what they charge their customers for intermediating funding. The higher, markdown-incorporated rate of dealer-intermediated repo funds is likely the marginal rate that prices securities. Imperfect competition in the Triparty market thus contributes to large and persistent funding spreads involving repo-financed securities such as Treasurys. Take the Treasury cash-futures basis as an example. The Treasury futures contract is priced by the cost of buying an equivalent physical Treasury security and holding that security till the futures' maturity. The financing rate implied in Treasury futures is therefore the repo rate on the funds used to purchase and hold the physical Treasury. There would be no funding spread if dealers intermediated repo funds at cost. Instead, the implied financing rate in Treasury futures consistently exceeds the wholesale repo rate. What explains the wedge between these two rates? Studies of this and other Treasury funding spreads have advanced regulation-induced balance-sheet costs as the cause (e.g., Fleckenstein and Longstaff (2020) and Jermann (2019)). Market power in the wholesale funding market offers a complementary explanation.

Indeed, funding spreads reflect the presence of frictions in dealer intermediation. One friction is the opportunity cost of using balance sheet space, which has been substantial since the 2007-09 Financial Crisis. This friction exists in all dealer-intermediated funds. For securities typically financed using repo funds, there is also the friction of monopsony rent, which the dealers extract from cash lenders. In Table 9, I study the composition of funding spreads by looking at two examples: the Treasury cash-futures basis, as discussed above, and the Treasury swap spread, which is the difference between the yield of Treasury securities and a maturity-matched unsecured wholesale funding rate, e.g., LIBOR or OIS. Specifically, I compare the magnitude of these two funding spreads with the estimated markdowns in the Triparty market and measures of balance sheet cost.

There is no standard measure of balance sheet cost, but it could be reflected in profits from risk-less arbitrage opportunities. I explore two such arbitrages.[53] First is the

[52] This split of surplus should not be interpreted to mean that if the Fed raises its policy rate by 1 bps, the passthrough would be 100% − 82% = 18%. Changes in the policy rate could also change the dealer's marginal value of intermediation and consequently the dealer's optimal rate setting and markdown.

[53] For these two arbitrages to provide a proxy of balance sheet cost, the key assumption is that their profits do not also reflect dealer mar-

Table 9
2011–2017 average measures of market power, balance sheet cost, and funding spread.

	Measure of market power	Measures of balance sheet cost	Measures of funding spread
Triparty dealer markdown	26.21		
IOER-EFFR spread		12.79	
USD-EUR 3M CIP basis		12.16	
Treasury swap spread			32.65
Treasury cash-futures basis			47.63

Notes: This table reports the averages of annualized measures in basis points. "Triparty dealer markdown" is the median of the estimated markdowns. "IOER-EFFR spread" is the difference between interest on excess reserve and the effective federal funds rate. "USD-EUR 3M CIP basis" is the difference in the 3-month interest earned on synthetic dollars at OIS (EONIA) versus interest paid on direct dollars at IBOR (LIBOR). Synthetic dollars are obtained by exchanging dollars for euros at the spot rate and converting back to dollars at the forward rate. "Treasury swap spread" is the difference between the Treasury yield and OIS (indexed to federal funds rate), averaged across the 5-year, 10-year, 20-year, and 30-year tenor. "Treasury cash-futures basis" is the difference between the funding rate implied in 5-year Treasury futures and the repo rate. "Triparty dealer markdown" is estimated for month-ends from January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter-ends. "Treasury cash-futures basis" is the average of the mean basis in years 2011 through 2017, as in Table 3 of Fleckenstein and Longstaff (2020). "IOER-EFFR spread" and "USD-EUR 3M CIP basis" are calculated using daily series from January 1, 2011, to December 31, 2017. "Treasury swap spread" is calculated using daily series from September 28, 2011, to December 31, 2017, when OIS rates are available.

IOER-EFFR arbitrage. The bank holding companies (BHCs) to which the dealers in my sample belong have access to the federal funds market and maintain accounts at the Federal Reserve. The interest on excessive reserve (IOER), paid by the Fed, can be earned risk-free if the BHC borrows extra reserve in the Fed funds market at the effective federal funds rate (EFFR). Therefore, any activity that a BHC does should earn at least the IOER-EFFR spread. This spread is on average 13 bps between 2011 and 2017.[54]

Another measure of balance-sheet cost is the deviation from the covered interest-rate parity (CIP) (Du et al., 2018; 2022a). CIP deviations happen across many currency pairs and tenors. Following Du et al. (2022b), I use the 3-month EUR-USD basis to approximate balance sheet costs and I use OIS to measure the risk-free interest that could be earned. Different from Du et al. (2022b), I use LIBOR to measure the cost of borrowing unsecured funds to conduct this arbitrage, which reflects funding value adjustments, as emphasized in Andersen et al. (2019).[55] Hence, if a dealer has excess balance sheet space, the dealer could raise 1 USD at its unsecured funding rate, LIBOR, and convert that 1 USD to 1 EUR at the spot exchange rate. The dealer would earn the risk-free OIS indexed to EONIA. At the end of 3 months, the dealer could convert the proceeds back to USD at the forward rate and pocket the profit after paying back the borrowed funds. On average, such a trade would have generated a profit of 12 bps between 2011 and 2017.

The magnitude of Treasury funding spreads suggests that the intermediation friction in repo funding comprises both balance sheet cost and dealers' market power. Between 2011 and 2017, the balance sheet cost was approximately 12 bps, and the estimated median Triparty markdown was 26 bps. Together, these figures suggest that the financing cost implied in repo-financed Treasury should exceed the wholesale repo rate by about 35 bps. This prediction is borne out in the Treasury-swap spread, where, averaging across the 5-year, 10-year, 20-year, and 30-year tenor, the yield of Treasury securities exceeded that of maturity-matched OIS by 33 bps between 2011 and 2017.[56,57] This prediction is also borne out in the Treasury cash-futures basis. Fleckenstein and Longstaff (2020) show that the financing rate implied in Treasury futures exceeded repo rates by about 50 bps between 2011 and 2017.[58]

In addition to the two discussed here, there are many other types of funding spreads, such as the basis between option-implied funding rates and wholesale funding rates (van Binsbergen et al., 2021). Many of these funding spreads involve securities that rely on dealer-intermediated funding, yet these spreads show limited correlations with each other (Siriwardane et al., 2021). Imperfect competition provides a natural micro-foundation for the observed low correlation: differences in the competitive landscape in different funding markets may well lead to divergences in funding spreads.

ket power. Neither of these two arbitrages relies on repo funding. At a minimum, profits from these two arbitrages exclude the kind of dealer market power studied in this paper.

[54] Although dealers cannot conduct IOER-EFFR arbitrage themselves because they are not depository institutions, the internal setup of BHCs allows the reserve desk and the trading desks to communicate and dynamically adjust balance sheet positions.

[55] Andersen et al. (2019) show that unsecured borrowing incurs debt overhang costs, which can be approximated with LIBOR-OIS or similar credit spreads. The relevant CIP deviation that gives balance-sheet costs can thus be equivalently calculated as the LIBOR-OIS CIP deviation or OIS-OIS CIP deviation adjusted for the LIBOR-OIS spread.

[56] This spread is, by convention, relative to OIS, or the unsecured federal funds rate. Over my sample period, the federal funds rate is on average 6.5 bps above Triparty repo rate index (median), indicating an even larger spread between the Treasury yield and the repo rate.

[57] The yield on Treasury bonds was in fact lower than OIS before the 2007-09 Financial Crisis. This does not necessarily mean that dealer market power is a recent phenomenon. As Du et al. (2022b) argue, the switch in the sign of the spread reflects a regime shift in dealers' Treasury inventory.

[58] Barth and Kahn (2021) provide an alternative measure of the Treasury cash-futures basis by using replicating portfolios of Treasury bills. The average for 5-year futures across the 1st, 2nd, and 3rd roll from 2010 to 2020 is about 21 bps. One possible reason for the discrepancy is the spread between Treasury bills and repo rates. Depending on the maturity of the Treasury bill, the spread with the overnight Triparty repo rate is between -4 bps for 1-month Treasury bills and 17 bps for 12-month Treasury bills.

6.4. The effect of the RRP on markdowns

The Federal Reserve instituted the Overnight Reverse Repo Facility in anticipation of increasing its policy interest rate. The RRP is thought to have helped the Fed successfully raise the interest rate four times between 2015 and 2017, during a period when the Fed's usual tool — reserve supply adjustment — was made obsolete by the abundance of reserves. Did this monetary policy tool also affect dealers' market power? I explore this question through counterfactual analyses detailed in Internet Appendix Section IA-A. Specifically, I ask, what would have happened to the Triparty repo rate and dealers' market power if the RRP were not established?

To conduct the policy counterfactual, I first parameterize and estimate the borrower's problem introduced in Section 4.2. The parameterization features a marginal value of intermediation that decreases in quantity. I estimate the parameters by using the 2016 Money Market Fund Reform as an exogenous shock to the lenders' supply of funds. The estimated parameters reflect a demand sensitivity that sees the borrower reducing the repo rate she offers by about 0.6 bps for every additional billion dollars of funding absorbed.

Next, I consider the effect of the RRP on dealers' market power. The RRP gives Triparty lenders an alternative to lending to dealers. Compared to the lender's other possible outside option, the RRP is very attractive: the RRP rate was on average 12 bps higher than the 1-day Treasury yield between 2014 and 2017. Imagine that the RRP were not established and the alternative to lending to Triparty dealers would instead earn the historical 1-day Treasury yield (Panel (a) of Internet Appendix Figure IA2). The counterfactual median Triparty repo rate would then be 8 bps lower than historical (Panel (b)). Importantly, the counterfactual median markdown would be 4 bps larger (Panel (c)), and this increase would happen at a time of increased volume lent to borrowers (Panel (d)), leading to even more economic rent extracted by the dealers. In other words, although the RRP is a policy tool intended to raise the level of Triparty repo rates, it also constrained dealers' market power and consequently narrowed funding spreads. The result of this counterfactual highlights one tangible way in which policy makers can effectively alter the competitive environment of the Triparty market.

7. Conclusion

I document new facts about the Triparty repo market that shed light on the nature of competition. These facts motivate me to describe the equilibrium of this vital wholesale funding market as cash-lenders allocating their portfolios among differentiated dealers (cash-borrowers) who set repo rates. My estimated model reveals that Triparty dealers enjoy substantial market power. This market power causes the observed wholesale repo rate to diverge from the unobserved, dealer-intermediated financing rates market participants face. Imperfect competition thus contributes to large funding spreads in repo-financed securities, such as the Treasury cash-futures basis and the Treasury swap spread.

More broadly, this study is a first step in a quantitative investigation of intermediary competition and its impact on asset prices. For months since May 2022, the rate offered by the Federal Reserve at the Overnight Reverse Repo Facility (RRP) in fact exceeded the median repo rate on the Triparty market. Such a phenomenon is difficult to rationalize in a perfectly competitive market, where the RRP ought to set the floor for repo rates in the Triparty market. The market power of financial intermediaries, which I explore in this paper, offers a plausible explanation. In general, if intermediaries are central to asset prices, and intermediaries, by definition, interface with many agents, then studying the competitive landscape intermediaries face is indispensable to a more complete understanding of the financial market.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code for this paper can be found at https://doi.org/10.17632/c652ftkj88.1.

Appendix A. Functional form of α_{it}

Results in Table 3 point to a strong empirical relationship between MMFs' overnight cash portfolio size and the distribution of portfolio shares. Importantly, this relationship exhibits concavity in portfolio size, evident in the log transformation applied to portfolio size (y_{it}). I therefore parameterize the concentration aversion parameter, α_{it}, which controls the distribution of portfolio shares, as $\alpha_{it} = \beta_0 + \beta_1 \sqrt{y_{it}}$. Square rooting the portfolio size captures the essence of concavity while ensuring that no values are negative, as would be generated under a log transformation because there are values of y_{it} that are less than 1.

There may be other plausible transformations. One such candidate is $\log(1 + y_{it})$, which is also concave and circumvents the negativity concern. However, as shown in Appendix Table A1, such a transformation would compress the range of transformed portfolio sizes, reducing the power of the estimation by flattening variation. An alternative transformation that produces a range comparable to the original log transformation is $\sqrt[2.5]{y_{it}}$. This transformation is rather uncommon. Therefore, in the absence of conclusive evidence that logs are the only appropriate transformation, I adopt square roots as the preferred functional form.

Reassuringly, all three transformations that maintain the spirit of concavity generate similar estimates of α_{it}. In Appendix Table A1, I report the mean α_{it} estimated using $\log(1 + y_{it})$, $\sqrt[2.5]{y_{it}}$, and $\sqrt{y_{it}}$. Applying the mean α_{it} to the average portfolio size yields a back-of-the-envelope

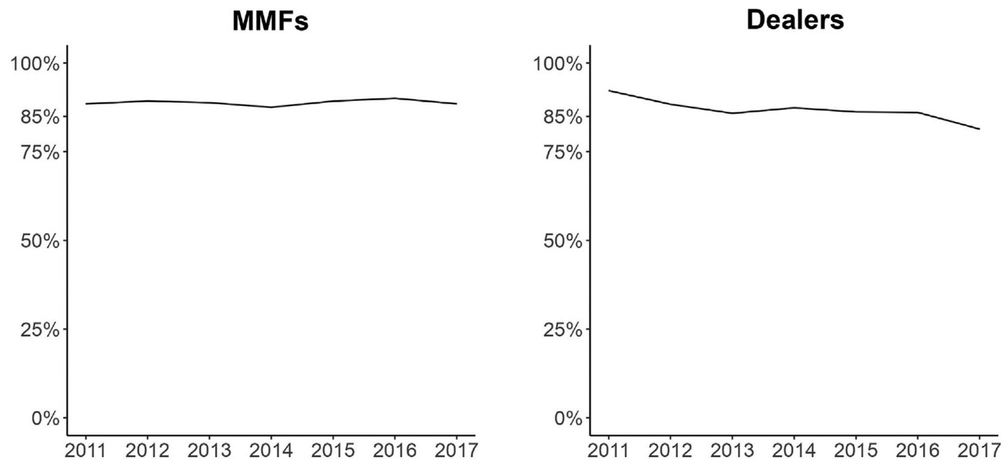

Fig. A1. Percentage of overnight repo market represented by top 18 MMFs and top 20 dealers in the sample.
Notes: This figure plots, on the left, the share of overnight repo done by the top 18 money market fund families relative to all overnight repo done by money market funds that filed N-MFP reports between January 2011 and December 2017. Plotted on the right is the share of overnight repo done by the top 20 dealers relative to all dealers based on money market funds' N-MFP reports from January 2011 to December 2017.

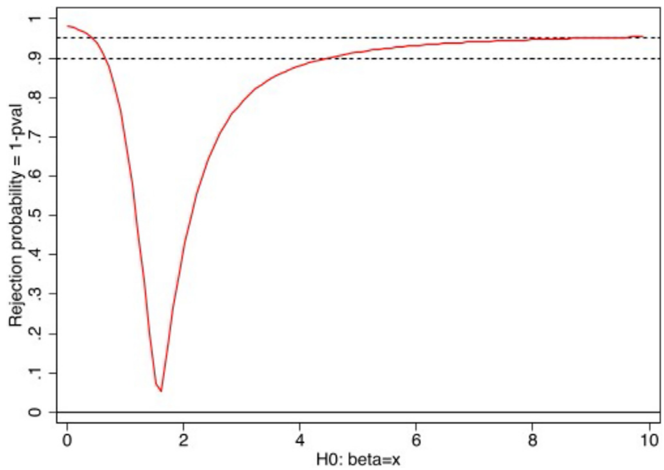

Fig. A2. Anderson-Rubin test of instrumental variable estimate.
Notes: This figure plots the Anderson-Rubin rejection probability for the null hypothesis that the true β_{IV} is equal to a given value on the x-axis. β_{IV} is estimated from $\log(R_{jt}) - \log(R_{zt}) = \beta_{IV}\widehat{vol}_{jt} + BorrowerFE + YearFE + e_{IV,jt}$ in the model estimation period of January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was first introduced and was in testing, and excluding months that fall on quarter-ends. The horizontal dashed lines are at y = 0.9 and y = 0.95. The points where the solid red line crosses the dashed lines represent, respectively, the end points of the Anderson-Rubin 90% confidence interval and 95% confidence interval for the null hypothesis. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

estimate of the inverse semi-elasticity that is comparable to the IV-estimate in Table 5. All three transformations imply an inverse semi-elasticity that is similar to the IV point estimate of 1.57, and all are well within the 95% Anderson-Rubin confidence interval.

Appendix B. Model fit

The estimated model parameters generate simulated data that closely match both targeted and untargeted moments in the original data, suggesting good model fit.

I use 44 moments in the indirect inference estimation. These moments can be grouped into three categories. In Appendix Fig. A3, I assess the fit in these three categories in turn.

The two regression coefficients in Panel (a) of Appendix Fig. A3 are among the most important moments. β_{IV} is the IV coefficient from Model (3) of Table 5, and it pins down the levels of lender preferences (α_{it} and ω_{ijt}). β_{median} is the coefficient on portfolio size for the median portfolio share (Model (2) in Table 3), and it informs the size-dependency of α_{it}. Estimated model parameters simulate data that generate coefficients (Matched moments) similar to the original point estimates (Data moments). In particular, the Matched moments are within the 95% confidence interval of the Data moments.

Panel (b) of Appendix Fig. A3 compares the R^2 of two regressions in the original data with those in the

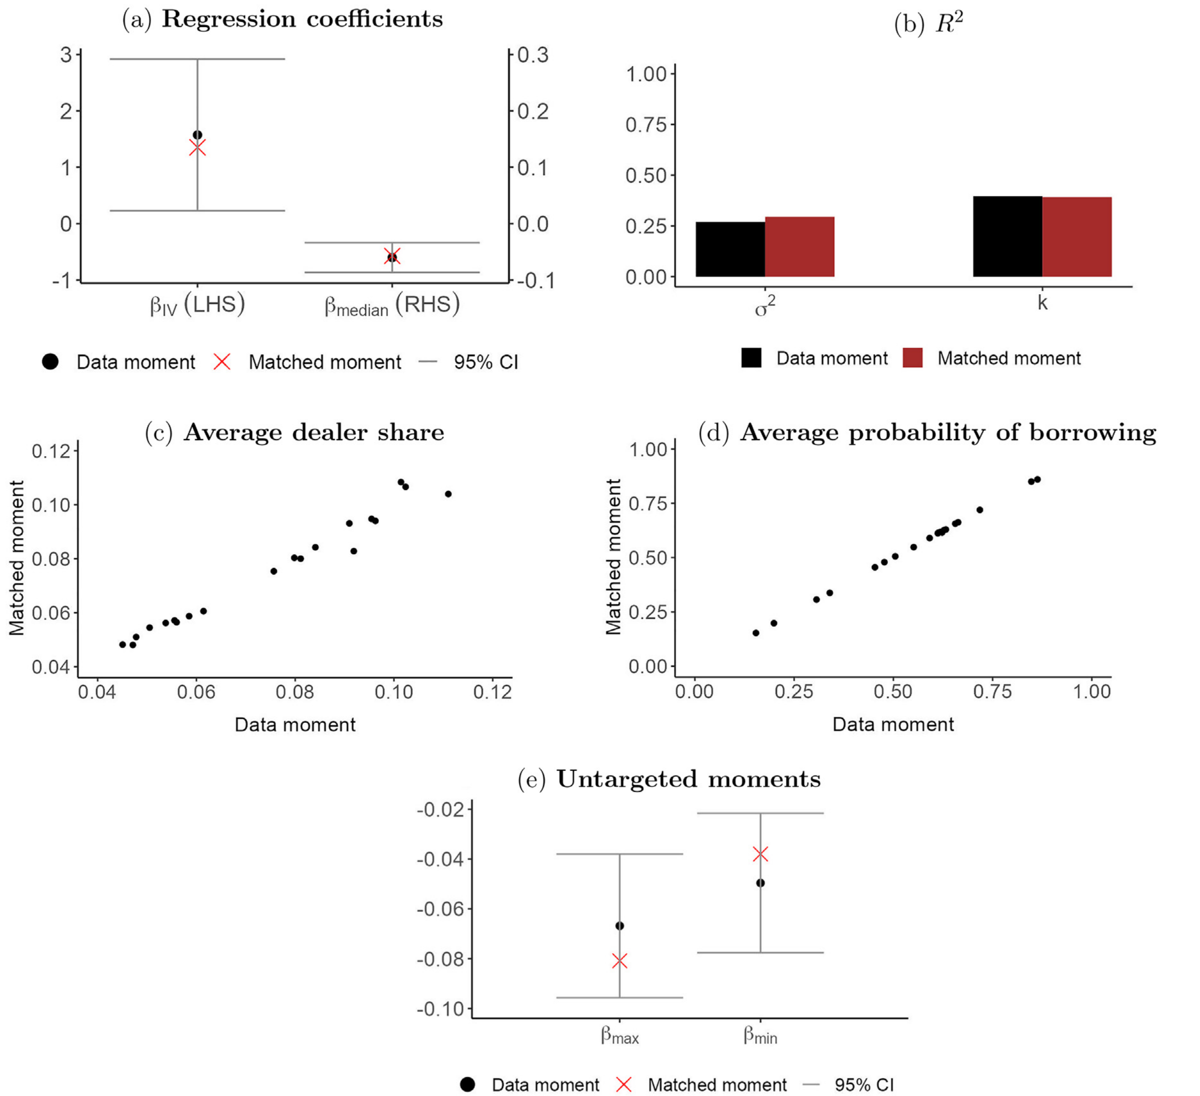

Fig. A3. Data moments vs. matched moments from estimated parameters.
Notes: This figure plots comparisons between data moments and matched moments. Data moments are calculated using the original sample data. Matched moments are the average of moments calculated in 50 sets of data simulated from estimated model parameters. The moments in Panel (a) are β_{IV} from Model (3) of Table 5 and β_{median} from Model (2) of Table 3. The moments in Panel (b) are the R^2 from the two regressions used to inform σ^2 and k (shape). The moments in Panel (c) are each dealer's average portfolio share, and the moments in Panel (d) are each dealer's average probability of borrowing from lenders. The moments in Panel (e) are $\beta_{\max}$ from Model (3) of Table 3 and $\beta_{\min}$ from Model (4) of Table 3. The model estimation period is January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter-ends.

parameter-simulated data. These two R^2s are most informative of two nuisance parameters: k (shape) of ν_{ijt} and σ^2 of ϵ_{jt}. The R^2s in the simulated data do not show significant deviations from those in the original data.

Panels (c) and (d) of Appendix Fig. A3 compare moments based on two types of dealer-specific averages, both used to inform ψ_j. Panel (c) contains each dealer's average portfolio share, and Panel (d) contains each dealer's average probability of borrowing from lenders. If moments from the simulated data match those from the original data, then the points in these two panels would fall along the 45-degree diagonal line. Points in Panel (d) align closely to the diagonal line. Points in Panel (c) exhibit more dispersion but do not show any systematic bias.

The targeted moments focus on the mean of the distribution. As a further check of model fit, I assess the match to two untargeted moments that reflect other parts of the data distribution. Specifically, I consider how well the estimated parameter can reproduce the behaviors of the maximum and the minimum portfolio shares in relation to portfolio size. These two moments come from Models (3) and (4) of Table 3, and their behaviors are not tied to β_{median} in Panel (a). As Panel (e) of Appendix Fig. A3 illustrates, although these two moments were not targeted in the estimation, they can be matched within the 95%

Table A1
Alternative concave transformations of y_{it} in α_{it}.

Transformation	Range	Mean estimate of α_{it}	Implied inverse semi-elasticity
$\log(y_{it})$	7.28		
$\log(1+y_{it})$	4.89	-0.0313	1.22
$\sqrt[2.5]{y_{it}}$	6.93	-0.0322	1.25
$\sqrt{y_{it}}$	11.75	-0.0434	1.69

Notes: This table compares different transformations of lenders' portfolio size (y_{it}). "Range" is for the transformed y_{it}. "Mean estimate of α_{it}" reports the mean α_{it} calculated using estimated β_0 and β_1 based on different transformations of y_{it}. "Implied inverse semi-elasticity" is the inverse semi-elasticity calculated using the mean estimate of α_{it}, the mean lender portfolio size, and the mean number of lenders lending to dealers. The model estimation period is January 2011 to December 2017, excluding September 2013–September 2014, when the RRP was in testing, and excluding months that fall on quarter-ends.

confidence interval by the estimated parameters. Taken together, the comparisons in Appendix Fig. A3 show that the estimated model offers a reasonable representation of the data.

Appendix C. Additional Figures and Tables

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jfineco.2023.04.007

References

Aldasoro, I., Ehlers, T., Eren, E., 2022. Global banks, dollar funding, and regulation. J Int Econ 137, 1036–1069. doi:10.1016/j.jinteco.2022.103609.
Anbil, S., Zeynep, S., 2018. The Regulatory and Monetary Policy Nexus in the Repo Market. FEDS Working Paper.
Andersen, L., Duffie, D., Song, Y., 2019. Funding value adjustments. J Finance 74 (1), 145–192. doi:10.1111/jofi.12739.
Anderson, A., Du, W., Schlusche, B., 2019. Arbitrage Capital of Global Banks. Technical Report.
Andrews, I., Stock, J.H., Sun, L., 2019. Weak instruments in IV regression: theory and practice. Annu Rev Econom.
Baklanova, V., Caglio, C., Cipriani, M., Copeland, A., 2017. The Use of Collateral in Bilateral Repurchase and Securities Lending Agreements. Technical Report. Federal Reserve of New York.
Barth, D., Kahn, J., 2021. Hedge Funds and the Treasury Cash-Futures Disconnect. Technical Report.
Benetton, M., 2021. Leverage regulation and market structure: astructural model of the u.k. mortgage market. J Finance doi:10.1111/jofi.13072.
Berk, J., Green, R., 2004. Mutual fund flows and performance in rational markets. Journal of Political Economy 112 (6), 1269–1295. URL: http://www.jstor.org/stable/10.1086/424739.
Bernanke, B., Reinhart, V., Sack, B., 2004. Monetary policy alternatives at the zero bound: an empirical assessment. Brookings Pap Econ Act 2, 1–100.
van Binsbergen, J.H., Diamond, W.F., Grotteria, M., 2021. Risk-free interest rates. J financ econ doi:10.1016/j.jfineco.2021.06.012.
Chernenko, S., Sunderam, A., 2014. Frictions in shadow banking: evidence from the lending behavior of money market funds. Review of Financial Studies 27, 1717–1750.
Copeland, A., Duffie, D., Martin, A., McLaughlin, S., 2012. Key mechanics of the U.S. tri-party repo market. Federal Reserve Bank of New York Economic Policy Review 18, 17–28.
Copeland, A., Martin, A., Walker, M., 2014. Repo runs: evidence from the tri-Party repo market. J Finance 69 (6), 2343–2380. doi:10.1111/jofi.12205.
Crawford, G.S., Lee, R.S., Whinston, M.D., Yurukoglu, A., 2018. The welfare effects of vertical integration in multichannel television markets. Econometrica 86 (3), 891–954. doi:10.3982/ECTA14031.
Darmouni, O., 2020. Informational frictions and the credit crunch. J Finance 75 (4), 2055–2094. doi:10.1111/jofi.12900.
Di Maggio, M., Egan, M., Franzoni, F., 2022. The value of intermediation in the stock market. J Financ Econ 145 (2, Part A), 208–233. doi:10.1016/j.jfineco.2021.08.020.
Du, W., Hébert, B., Huber, A.W., 2022. Are intermediary constraints priced? Review of Financial Studies.
Du, W., Hébert, B., Li, W., 2022. Intermediary Balance Sheets and the Treasury Yield Curve. Technical Report.
Du, W., Tepper, A., Verdelhan, A., 2018. Deviations from covered interest rate parity. J Finance doi:10.1111/jofi.12620.
Duffee, G.R., 1996. Idiosyncratic variation of treasury bill yields. J Finance 51 (2), 527–551. doi:10.1111/j.1540-6261.1996.tb02693.x.
Duffie, D., 2017. Post-crisis bank regulations and financial market liquidity. Working Paper. Banca d'Italia.
Duffie, D., Gârleanu, N., Pedersen, L.H., 2005. Over-the-counter markets. Econometrica 73 (6), 1815–1847. doi:10.1111/j.1468-0262.2005.00639.x.
Duffie, D., Krishnamurthy, A., 2016. Passthrough Efficiency in the Fed's New Monetary Policy Setting. Fed Symposium, Kansas City. URL: https://people.stanford.edu/akris/sites/default/files/jh-september-8-2016-bw.pdf.
Eisenschmidt, J., Ma, Y., Zhang, A., 2021. Monetary Policy Transmission in Segmented Markets. Technical Report.
Fleckenstein, M., Longstaff, F.A., 2020. Renting balance sheet space: intermediary balance sheet rental costs and the valuation of derivatives. Rev Financ Stud 33 (11), 5051–5091. doi:10.1093/rfs/hhaa033.
Fleming, M., Rosenberg, J., 2008. How Do Treasury Dealers Manage Their Positions? Federal Reserve Bank of New York Staff Reports.
Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. J Financ Econ 104 (3), 425–451. doi:10.1016/j.jfineco.2011.03.016. Market Institutions, Financial Market Risks and Financial Crisis.
Gorton, G.B., Metrick, A., Ross, C.P., 2020. Who ran on repo? AEA Papers and Proceedings 110, 487–492. doi:10.1257/pandp.20201100.
Gourieroux, C., Monfort, A., Renault, E., 1993. Indirect inference. Journal of Applied Econometrics 8, S85–S118. URL: http://www.jstor.org/stable/2285076.
Greenwood, R., Hanson, S.G., Stein, J.C., 2015. A comparative-advantage approach to government debt maturity. Journal of Finance 70 (4), 1683–1722. doi:10.1111/jofi.12253.
Hahn, J., Newey, W., 2004. Jackknife and analytical bias reduction for nonlinear panel models. Econometrica 72 (4), 1295–1319. doi:10.1111/j.1468-0262.2004.00533.x.
Han, S., Nikolaou, K., Tase, M., 2022. Trading relationships in secured markets: evidence from triparty repos. Journal of Banking & Finance 139, 106486. doi:10.1016/j.jbankfin.2022.106486.
He, Z., Krishnamurthy, A., 2018. Intermediary asset pricing and the financial crisis. Annual Review of Financial Economics 1–37.
Hendershott, T., Li, D., Livdan, D., Schurhoff, N., 2020. Relationship trading in over-the-counter markets. J Finance 75 (2), 683–734. doi:10.1111/jofi.12864.
Holm, S., 1979. A simple sequentially rejective multiple test procedure. Scand. J. Stat. 6 (2), 65–70. URL: http://www.jstor.org/stable/4615733.
Hu, G.X., Pan, J., Wang, J., 2021. Tri-party repo pricing. Journal of Financial and Quantitative Analysis 56 (1), 337–371. doi:10.1017/S0022109019000863.
Infante, S., 2019. Liquidity windfalls: the consequences of repo rehypothecation. J Financ Econ 133 (1), 42–63. doi:10.1016/j.jfineco.2019.02.004.
Jermann, U., 2019. Negative swap spreads and limited arbitrage. Rev Financ Stud 33 (1), 212–238. doi:10.1093/rfs/hhz030.
Kacperczyk, M., Schnabl, P., 2013. How safe are money market funds? Quarterly Journal of Economics 128 (3), 1073–1122. doi:10.1093/qje/qjt010. Cited by: 105.
Kim, J., Allenby, G.M., Rossi, P.E., 2002. Modeling consumer demand for variety. Marketing Science 21 (3), 229–250. doi:10.1287/mksc.21.3.229.143.
Koijen, R.S.J., Yogo, M., 2018. An equilibrium model of institutional demand and asset prices. Journal of Political Economy. URL: https://poseidon01.ssrn.com/delivery.php?ID=0111060090671181230980051151151030240380390000650030341030841050870951110851220710780101180050340100991130661040160020821251210390350930090461180651270900240911160710660711201160680640801211050110271270920040840.
Krishnamurthy, A., Nagel, S., Orlov, D., 2014. Sizing up repo. J Finance 69 (6), 2381–2417. doi:10.1111/jofi.12168.
Li, Y., 2021. Reciprocal lending relationships in shadow banking. J financ econ 141 (2), 600–619. doi:10.1016/j.jfineco.2021.04.004.
Macchiavelli, M., Zhou, X., 2022. Funding liquidity and market liquidity: the broker-dealer perspective. Manage Sci 68 (5), 3379–3398. doi:10.1287/mnsc.2021.4053.

Martin, A., Skeie, D., Thadden, E.-L.v., 2014. Repo runs. Rev Financ Stud 27 (4), 957–989. doi:10.1093/rfs/hht134.
Martin, G., Yurukoglu, A., 2017. Bias in cable news: persuasion and polarization. American Economic Review 107 (9), 2565–2599.
Schmidt, L., Timmermann, A., Wermers, R., 2016. Runs on money market mutual funds. American Economic Review 106 (9), 2625–2657. doi:10.1257/aer.20140678.
Siriwardane, E., Sundarem, A., Wallen, J., 2021. Segmented Arbitrage. Working Paper.
Strahan, P.E., Tanyeri, B., 2015. Once burned, twice shy: money market fund responses to a systemic liquidity shock. The Journal of Financial and Quantitative Analysis 50 (1/2), 119–144. URL: http://www.jstor.org/stable/43862245.
Wallen, J., 2020. Markups to Financial Intermediation in Foreign Exchange Markets. Technical Report.
Weymuller, C.-H., 2013. Leverage and Reputational Fragility in Credit Markets. Technical Report.