

Excess Reserves and Monetary Policy Implementation*

Roc Armenter
Federal Reserve Bank of Philadelphia

Benjamin Lester[†]
Federal Reserve Bank of Philadelphia

December 29, 2016

Abstract

In response to the Great Recession, the Federal Reserve resorted to several unconventional policies that drastically altered the landscape of the federal funds market. The current environment, in which depository institutions are flush with excess reserves, has forced policymakers to design a new operational framework for monetary policy implementation. We provide a parsimonious model that captures the key features of the current federal funds market, along with the instruments introduced by the Federal Reserve to implement its target for the federal funds rate. We use this model to analyze the factors that determine rates and volumes under the new implementation framework, and to study the effects of changes in the policy rates and other shocks to the economic environment. We also calibrate the model and use it as a quantitative benchmark for applied analysis, with a particular emphasis on understanding the role of the overnight reverse repurchase agreement facility in supporting the federal funds rate.

J.E.L. codes: E42, E43, E52, E58

Keywords: excess reserves, federal funds market, federal funds rate

*Thanks to Gara Afonso, Andrea Ajello, Marco Cipriani, James Clouse, Huberto Ennis, Urban Jermann, Beth Klee, Antoine Martin, Cyril Monnet, Chris Waller, and Ronald Wolthoff. Thanks also to seminar participants at the Federal Reserve Bank of New York; Federal Reserve Bank of Philadelphia; University of Bern; the System Committee Meeting on Macroeconomics; the Central European University; the 2015 Konstanz Seminar; the Society of Economic Dynamics 2015 conference; the 2015 Summer Workshop on Money, Banking, Payments, and Finance at the Federal Reserve Bank of Saint Louis; the Federal Reserve Board; and the Federal Reserve Bank of Richmond. All errors are our own.

[†]The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

The federal funds market is the first cog in the transmission of monetary policy in the U.S. As such, it has been extensively studied in the academic literature, from the seminal contribution of Poole (1968) to the recent work of Afonso and Lagos (2015). However, in the wake of the extraordinary measures taken in response to the 2007–2008 financial crisis, the Federal Open Market Committee (FOMC) now faces a vastly different federal funds market—one for which past experience and existing theory provide little guidance.

Prior to the financial crisis, most trades in the federal funds market were between depository institutions trying to achieve their optimal level of reserves. In particular, some depository institutions would be borrowing to satisfy reserve requirements, while others would be lending to avoid holding idle excess reserves. In this environment, monetary policy implementation was fairly straightforward: The Open Markets Trading Desk at the Federal Reserve Bank of New York would engage in open market operations, adjusting the supply of reserves available in the federal funds market until the rates traded at the target prescribed by the FOMC.

In response to the Great Recession, the Federal Reserve resorted to a number of unconventional policies that have drastically changed the landscape of the federal funds market. More specifically, in the wake of the large-scale asset purchase programs, most depository institutions found themselves awash with excess reserves. As a result, only a small fraction of trades in the federal funds market are now *between* depository institutions, since virtually none of them need to borrow in order to satisfy reserve requirements. Instead, the market is now dominated by other investors (such as government-sponsored enterprises, or GSEs, and money market funds) looking for some yield on overnight cash balances. This environment poses a challenge for monetary policy implementation, since the size of the open market operations required to raise the federal funds rate is neither feasible nor desirable in the medium term.¹

Recognizing this situation, the Federal Reserve has developed a new framework for implementing the desired target for federal funds rates in the current environment of excess reserves. As detailed by the FOMC in the September 17, 2014, press release, “Policy Normalization Principles and Plans,” the new framework relies on two tools to implement the desired policy rate. First, the committee “intends to move the federal funds rate into the target range [...] by adjusting the interest rate it pays on excess reserve balances.” Second, the committee also “intends to use an overnight reverse repurchase agreement facility [...] to help control the federal funds rate,” though the plan is to use this latter tool “only to the extent necessary.”

Unfortunately, policymakers have limited experience with this new framework and thus a number of important questions remain. First and foremost, will this framework be able to successfully implement monetary policy, particularly as target rates increase? What are the factors that could endanger the Desk’s ability to carry out the FOMC’s directives? In addition, the design of the framework itself is bound to evolve going forward. For example, in the September 17, 2014 press release, the FOMC indicated that it plans to eventually “phase [the overnight reverse repurchase facility] out when it is no longer needed” in order to minimize the Fed’s footprint in short-term money markets. What are the effects of reducing the capacity of the overnight

¹While we focus on the federal funds market, the situation of excess reserves is far from unique to the U.S. Central banks in several advanced economies have also relied on large-scale asset purchases and will face similar challenges when they choose to start on their own path to policy normalization.

reverse repurchase facility, or of removing it altogether? Are there other ways to reduce activity at this facility, encouraging trading within the federal funds market instead?

In order to address the questions raised above, and many more, we develop a tractable theoretical model that captures the key features of the current federal funds market, along with the instruments that the FOMC currently relies upon for monetary policy implementation. We solve the model and use it to identify the factors that affect the federal funds rate, and whether or not this rate will remain within the target range in response to changes in policy or in the economic environment more generally. Then, exploiting the few available moments in the data, we calibrate the model and use it as a quantitative benchmark for applied analysis, with a particular emphasis on understanding the role of the overnight reverse repurchase facility in supporting federal funds rates, and the ramifications of limiting its size or eliminating it entirely.

To capture the basic institutional arrangements in the federal funds market, as it currently operates, we start with a central bank that manages two separate facilities. The first facility pays interest on overnight excess reserves (IOER) to qualified depository institutions (DIs), while the second facility provides a lower, but positive rate of return for overnight reverse repurchase (ON RRP) agreements. The ON RRP facility is available to financial institutions with excess cash (who we call *lenders*) that do not qualify as DIs, i.e., the GSEs and money market funds. Hence, there are gains from trade between lenders and DIs, as they attempt to exploit the arbitrage opportunity between the ON RRP rate and the IOER rate. Consistent with the Federal Reserve's current operational framework, the ON RRP rate is equal to the bottom of the FOMC's target range for the federal funds rate (FFR), while the IOER rate is set at the top of the target range.

In addition to the relevant agents and policy instruments, we attempt to capture the key features of the federal funds market. The first thing to note about the federal funds market is that it is an “over-the-counter” market, where individual participants search for willing counterparties and ultimately borrow and lend at a variety of different interest rates. Second, DIs in the federal funds market earn substantial margins in their trades and thus appear to have some degree of market power.² Lastly, it's important to note that not all DIs are active in the federal funds market because there are nontrivial balance sheet costs associated with accepting deposits; these costs include both the direct expenses that a DI faces from expanding its balance sheet, like FDIC fees, as well as the indirect expenses associated with requirements on capital and leverage ratios.³ Moreover, these balance sheet costs vary substantially across DIs, due to differences in regulation by, e.g., jurisdiction or size.

To capture these features of the federal funds market, we add two ingredients to our model. First, we assume that the market is not perfectly competitive, but rather characterized by search frictions. In particular, we posit a *directed search* model; as we argue later, this model accounts nicely for the features of the federal funds market outlined above, it is flexible enough to match several key moments from the data, and yet it remains tractable enough to accommodate various extensions. Second, we assume that each DI incurs a cost when it accepts a deposit from a lender, and that these costs are heterogeneous across DIs. As a result, a typical equilibrium will exhibit a nondegenerate distribution of interest rates being traded between lenders and DIs, with the measure of DIs participating in the market (and hence the overall degree of competition)

²See, e.g., Bech and Klee (2011), for evidence that DIs have market power, along with estimates of how much.

³See, e.g., Potter (2013, 2014) and Martin et al. (2013b), who cite balance sheet costs as an important barrier to entry for DIs.

endogenously determined.

We provide a full analytic characterization of the equilibrium and study its key properties. We show that, within the context of this model, the new framework is successful at implementing monetary policy: traded rates between lenders and DIs will always lie in the interior of the target range. Moreover, if policymakers move the IOER and ON RRP rates in parallel—keeping the *spread* between the two constant—we show that traded rates move one-to-one with the target range, and that volume in the federal funds market (and demand at the ON RRP facility) remain constant. Intuitively, given the supply of funds and the distribution of balance sheet costs, the spread between the two policy rates pins down the gains from trade between lenders and DIs and hence the relative size of each side of the market in equilibrium. This, in turn, determines where exactly interest rates will lie within the target range, and the volume of activity in the federal funds market vs. the ON RRP facility. We also conduct comparative statics to illustrate how an isolated change in one of the policy rates—or, more generally, changes in the underlying parameters—affect traded rates and volume.

We complement our analytic results with a parsimonious calibration, allowing us to tackle a number of important policy questions that require quantitative answers. We first use data from before December 2015—i.e., before the first rate hike or so-called “liftoff”—to show that the model is flexible enough to replicate several salient features of the data. Then, using this calibration, we show that the model does a good job of matching the data after liftoff occurred. We then explore various policy proposals that are intended to reduce trading activity at the ON RRP facility. These proposals include eliminating the facility altogether, placing a cap on the daily volume, or more indirect options like increasing the spread between the IOER and ON RRP rates, which will encourage more trading activity in the federal funds market and relieve some of the pressure at the ON RRP facility. Our results provide a quantitative estimate of how these policies would affect interest rates and trading volume, and also reveal several additional, more general lessons regarding monetary policy implementation within the current framework. For example, we show that a cap on volume at the ON RRP facility poses a risk to successful monetary policy implementation, and that this risk increases as the target range rises—holding the spread between the IOER and ON RRP rates fixed—but falls as the spread itself widens.

Lastly, we return to our benchmark framework and extend the model to incorporate several relevant considerations into our analysis. First, we recognize that some lenders in the federal funds market can only deposit their reserves at a subset of DIs, either because of regulatory requirements or because some DIs are simply not part of their trading network. Hence, we assume that a fraction of lenders can only deposit funds at a predetermined set of DIs, and explore the ramifications for interest rates, trading volume, and implementation. Second, although in practice the Federal Reserve designated a virtually comprehensive list of counterparties for the ON RRP facility, eligibility is a relevant policy margin. Therefore, we extend the model so that some lenders are not eligible to deposit reserves at the ON RRP facility and explore the implications for reducing or expanding the set of eligible counterparties. Last, in an environment with a cap at the ON RRP facility, the amount of spare capacity or “headroom” at the ON RRP facility is irrelevant as long as the cap is not binding. If there is uncertainty over the supply of funds held by lenders, however, we show that the amount of headroom is indeed important, as it affects both the volume of trade and interest rates in the federal funds market.

The structure of the paper is as follows. Section 2 provides background information on the

federal funds market and monetary policy implementation, and reviews the existing literature on these topics. We describe our benchmark model in Section 3, then characterize the equilibrium and explore its properties analytically in Section 4. Section 5 describes our calibration strategy, explores the quantitative properties of the calibrated model, and tests the predictions of the model against the data since the rate hike from December 2015. Then, in Section 6, we use the calibrated model to study the role of the ON RRP facility, and the effects of various policy proposals. Section 7 contains the extensions described above and Section 8 concludes. All proofs are in the Appendix.

2 Institutional Background and Related Literature

Before we introduce our model, it's helpful to provide a more detailed description of the federal funds market, the operational framework that the FOMC traditionally used to implement monetary policy in the federal funds market, and the new framework that it has chosen to use in the current environment with excess reserves. We close this section with a review of the existing literature on monetary policy implementation in the federal funds market, both before and after the Great Recession.

2.1 The Federal Funds Market and Monetary Policy Implementation

A federal funds transaction is an unsecured loan of U.S. dollars between eligible entities, like depository institutions (DIs) or government-sponsored enterprises (GSEs). The FOMC pursues its mandate of price stability and full employment by setting a target level or range for the effective federal funds rate (FFR).⁴ There is no central repository for federal funds trades; instead, participants arrange transactions directly with each other or through brokers, in what is commonly described as an “over-the-counter” market.

As noted previously, the federal funds market was traditionally dominated by trades between DIs seeking to adjust their reserve holdings to desired levels: institutions with excess reserves would lend to institutions with reserves short of the required level. In this context, the Trading Desk at the Federal Reserve Bank of New York (commonly known as “the Desk”) would implement the targeted rate simply by fine-tuning the supply of reserves in the banking system via open market operations.⁵ As Potter (2013) describes, the Desk was typically able to achieve the FOMC’s policy directives, often without much need to conduct operations.

In recent years, the landscape of the federal funds market has changed dramatically. In response to the financial crisis and the following recession, the FOMC reduced its federal funds target to virtually zero in December 2008 and embarked on a series of asset purchase programs. These expanded the supply of reserves by several multiples and, as a result, all but very few DIs had an extremely high level of excess reserves. Thus, trade between banks dried up and the FFR dropped substantially below the FOMC’s target. In an attempt to put a floor on short-term

⁴The effective FFR had traditionally been computed as a weighted average of interest rates charged in federal funds transactions, obtained from data supplied by brokers. As of March 1, 2016, the effective FFR is calculated as a volume-weighted median of overnight federal funds transactions reported in the FR 2420 Report of Selected Money Market Rates.

⁵See Akhtar (1997) for a detailed description of open market operations.

interest rates, the Federal Reserve started to pay interest on excess reserves to DIs.⁶

The federal funds rate and other money market rates, however, consistently traded below the IOER rate. The reason is that the IOER rate is only available to DIs holding balances at the Federal Reserve, and not to key participants like GSEs. These participants—most notably the Federal Home Loan Banks—account for the majority of the “supply” of funds in the market, as they try to place their holdings at a DI that can then earn the IOER rate.⁷

This leads us to the final piece of the institutional landscape: the ON RRP facility, which the Desk introduced in September 2013 with an expanded list of counterparties and a maximum allotment.⁸ In a reverse repurchase, the Desk sells a security to a counterparty with an agreement to buy the security back at a pre-specified date and price, with the interest rate computed from the difference between the original purchase price and the (higher) repurchase price. The interest rate available at the ON RRP facility constitutes the relevant outside option for institutions that are ineligible for the IOER rate, and thus helps support the level of the federal funds rate.

It is worth noting that the list of eligible counterparties at the ON RRP also includes key money market funds, which are actually not eligible to transact federal funds. Instead, these funds lend to DIs through the eurodollar markets. Federal funds and eurodollar deposits are near-perfect substitutes for DIs, since eurodollar deposits have also been exempt from reserve requirements since 1991, as documented in Bartolini et al. (2008) and others. Indeed, the Federal Reserve Bank of New York has started publishing an estimate of the overnight bank funding rate that includes transactions in eurodollar markets by DIs.⁹

On December 16, 2015, using the two instruments described above, the FOMC decided to increase rates for the first time since the Great Recession, targeting a range for the FFR between 25 and 50 basis points. As described in the “Decisions Regarding Monetary Policy Implementation,” the ON RRP and IOER rates were set at the bottom and the top, respectively, of the target range.¹⁰ The Desk was authorized to operate the ON RRP facility in amounts limited only by the value of Treasury securities held outright in the System Open Market Account (SOMA)—a full-allotment facility for all practical purposes.

The so-called “lift off” was successful, with the effective FFR trading at 37 basis points the day after the FOMC decision, and staying near the middle of the target range since then, with only the exceptions of month- and quarter-end dates. Demand at the ON RRP facility was elevated in the weeks following the FOMC meeting, but stabilized soon after. Eurodollar rates have also remained tightly connected with federal funds rates, showing that the higher rates are being transmitted through money markets.

⁶The so-called “floor” or “corridor” system has been successfully implemented by a number of central banks, including the European Central Bank, the Bank of Japan, the Sveriges Riksbank, the Bank of Canada, and the Bank of England. It has also been extensively studied in the literature; see, among others, Ennis and Keister (2008), Whitesell (2006), Berentsen and Monnet (2008), and Berentsen et al. (2015).

⁷See Bech and Klee (2011) and Afonso et al. (2013a,b) for a description of the key participants in the federal funds market after the financial crisis.

⁸Prior to December 2015, the ON RRP facility operated with a cap on total allotment. If the demand for ON RRP exceeds the maximum allotment, the Desk uses an auction to set the interest rate on ON RRP, which can then be lower than the initial rate announced. See Frost et al. (2015) for a complete discussion of the design of the ON RRP facility.

⁹The overnight bank funding rate can be obtained at <https://apps.newyorkfed.org/markets/autorates/obfr>. See <https://www.newyorkfed.org/markets/obfrinfo> for additional information.

¹⁰See <https://www.federalreserve.gov/newsevents/press/monetary/20151216a1.htm>.

2.2 Related Literature

There is a long tradition of developing models to analyze monetary policy implementation in an environment with scarce reserves. The original contribution by Poole (1968) posited a downward-sloping demand curve for reserves and analyzed how the Federal Reserve could target the desired FFR by manipulating the supply of reserves. While this approach abstracted from the actual trading mechanisms in place, it proved very useful and became the workhorse model for the federal funds market, being further developed by Ho and Saunders (1985) and Hamilton (1996), among many others.

Recently, a second generation of models has been developed to capture the actual micro-mechanics of trading in the federal funds market—in particular, its over-the-counter nature, which was first emphasized by Furfine (1999) and Ashcraft and Duffie (2007). One prominent example is Afonso and Lagos (2015), who develop a random search model to capture the idea that trade in this market is bilateral in nature and trading partners often take time to locate.¹¹ They use their model to explore intra-day trading dynamics in the federal funds markets and the determinants of the FFR. Ennis and Weinberg (2013) also develop a search and matching model of the federal funds market to study the “stigma” associated with the use of the discount window and its implications for the demand of reserves. Bianchi and Bigio (2014) embed a simplified model of over-the-counter trading in a macro model designed to study the bank lending channel of monetary policy. Explicit models of over-the-counter trading have also been used to study monetary policy implementation outside the U.S.; see, among others, Berentsen and Monnet (2008) and Bech and Monnet (2014).

These second-generation models provide a more accurate description of trading and monetary policy implementation in the federal funds market *as it was* before the Federal Reserve’s unconventional policies. To date, very few papers have attempted to model the federal funds market *as it is now*. Examples include Martin et al. (2013a,b), Ennis (2014), and Williamson (2015), who study the macroeconomic effects of monetary policy in an environment with excess reserves. In these papers, the federal funds market—or the interbank market more generally—is treated as a competitive market, which allows the researchers to preserve tractability in the context of a general equilibrium model. Thus, these models are designed to capture the consequences of monetary policy but mainly abstract from the challenges of its implementation.

Our model, on the other hand, is focused exclusively on understanding the determinants of interest rates in this new environment. To the best of our knowledge, the only other paper with a similar focus is Bech and Klee (2011), who study the role that GSEs play in the federal funds market and, in particular, the reasons why the FFR has traded below the IOER rate. Like many of the second-generation models, we utilize a model with search frictions to capture the

¹¹There are several reasons why search-based models are useful for studying the federal funds market. First, modeling trade as bilateral captures the reality that participants in this market trade directly with one another, and not through a central repository. Second, and most important, search-based models are well-suited to capture a number of important features of the data. For example, Ashcraft and Duffie (2007) document significant heterogeneity in both the time it takes banks to trade in the federal funds market and in the terms at which they trade; such observations are difficult to rationalize within the context of a standard Walrasian paradigm, where trade occurs instantaneously at a single price. They also document a systematic relationship between time to trade, terms of trade, and *need* to trade (as captured by a bank’s reserve holdings relative to its optimal holdings) which, they conclude, “are consistent with the thrust of search-based OTC financial market theory.” (Ashcraft and Duffie, 2007, p. 221)

over-the-counter nature of the federal funds market. However, unlike these models, we utilize a model of *directed* (instead of random) search.¹² This choice is motivated by several factors. First, this paradigm seems to capture several salient features of trading activity in the federal funds market. For example, given the repeated daily interactions between, e.g., GSEs and DIs, it seems unlikely that market participants have no ex ante information about which DIs typically pay higher or lower interest rates on overnight loans. A virtue of our directed search model is that the interest rate that a DI is willing to pay is known to all market participants, and those that offer relatively high interest rates will, in equilibrium, attract more overnight loans. Put differently, our model captures the bilateral and stochastic nature of meetings in the federal funds market without severing the link between interest rates and allocations. Second, models of directed search tend to have several attractive technical features. For one, they offer a framework in which sellers compete with one another in a setting that lies between the extreme cases of monopoly and Bertrand competition. Moreover, directed search models are highly tractable, they can easily incorporate various types of (observed or unobserved) heterogeneity, and the equilibrium in these models tend to be constrained efficient.¹³

3 Environment

We consider a two-period economy that is populated by three types of agents. First, there is a measure λ of non-depository financial institutions which we will refer to as “lenders,” as this is representative of their primary role in the model. Second, there is a measure of depository institutions, normalized to one, which we will refer to as “DIs” for brevity. Finally, there is a central bank which, for obvious reasons, we will call the “Fed.” All agents are risk neutral and do not discount between $t = 1$ and $t = 2$.

The Fed. The Fed is an institution that operates two facilities. The first is the overnight reverse repurchase, or ON RRP, facility, where an agent can use cash to purchase a security from the Fed at $t = 1$, and then resell the security to the Fed at $t = 2$ at a pre-specified price. We denote the net rate of return on this investment by r , and assume this is chosen by the Fed. Both lenders and DIs have access to the ON RRP facility.

The Fed also accepts cash deposits at $t = 1$, which earn a net interest rate of $R > r$ at $t = 2$. We refer to R as the interest on excess reserves, or IOER, rate, and assume that this rate is also chosen by the Fed. Importantly, only DIs have access to this second facility.

Lenders. At $t = 1$, each lender is endowed with one unit of excess cash. As noted previously, lenders cannot deposit their reserves directly at the Fed to earn the IOER rate, R . However, there is an interbank market where they can lend their cash to a DI, who can then deposit it at the Fed to earn R , retaining some of the return for its own profit. The interbank market is

¹²For seminal contributions to the directed (or competitive) search literature, see Moen (1997), Burdett et al. (2001), and the references therein. For a recent example that uses directed search to model financial markets, see Lester et al. (2015).

¹³In contrast, random search models typically require that any heterogeneity is observable, in order to avoid issues with bargaining under private information. Also, random search models are generically inefficient, as congestion externalities are not internalized when prices are determined by Nash bargaining except in a knife-edge case; see Hosios (1990).

a frictional one, though: Not all lenders will match with DIs, and not all DIs will match with a lender. We discuss the probability that a lender matches with a DI, and the interest rate that they earn if they do match successfully, in greater detail below. Those lenders who do not match in the interbank market can access the ON RRP facility and earn the interest rate r .¹⁴

Depository Institutions. Each DI, which we index by j , can accept up to one unit of cash from a lender. However, doing so imposes a “balance sheet cost” on the DI, which we denote by c_j . DIs are heterogeneous with respect to their balance sheet costs: We denote by $G(c_j)$ the distribution of costs across DIs, and assume that $G(\cdot)$ is continuous and strictly increasing on the support $c_j \in [0, \infty)$. Given its balance sheet cost, each DI decides whether or not to enter the interbank market and, if it enters, it chooses an interest rate that it will pay to borrow a unit of cash, which we denote by ρ_j . Figure 1 summarizes the possible transactions that can occur between lenders, DIs, and the Fed, along with the interest rates that are paid in each type of transaction.¹⁵




Figure 1: Summary of Possible Transactions

¹⁴Later, we consider what happens if policymakers eliminate the ON RRP facility, or impose a cap on the volume of trade that is allowed, so that lenders no longer have the option to earn the interest rate r with certainty.

¹⁵As discussed in Section 2, some of the financial institutions that would be labeled “lenders” in our model do not technically qualify to make federal funds transactions; instead, these lenders’ deposits are executed as eurodollar transactions. From the point of view of the model, the distinction is irrelevant: Lenders approach the DIs directly and their trades are formalized as a federal funds or a eurodollar transaction without impacting the terms of trade. We return to this point in Section 5, when it becomes necessary to map the available data to the model in order to calibrate the model.

Matching in the Interbank Market. Once DIs have made entry decisions and posted interest rates, each lender observes the interest rates that have been posted and chooses one to approach. We will often refer to the set of DIs that have chosen a particular interest rate as a “submarket.” Conditional on choosing a submarket, a lender may or may not be paired with a DI—there are *matching frictions* in the interbank market. In particular, suppose a measure d of DIs post a particular interest rate and a measure ℓ of lenders choose to pursue that rate. Then, letting $q = \ell/d$ denote the market tightness or “queue length” in that submarket, the probability that each DI receives a deposit is $1 - e^{-q}$ and, symmetrically, the probability that each lender matches with a DI is $\frac{1 - e^{-q}}{q}$.¹⁶

Summary of Timeline and Payoffs. Figure 2 depicts the sequence of events. A DI who is matched with a lender at $t = 1$ incurs the balance sheet cost c_j and deposits the reserves at the Fed. Then, at $t = 2$, the DI earns the IOER rate R and pays the lender the promised interest rate ρ_j . Hence, a DI with balance sheet cost c_j that posts interest rate ρ_j and matches with a lender will earn a net profit of $R - c_j - \rho_j$. An unmatched DI, on the other hand, earns zero. Meanwhile, a lender earns a payoff of ρ_j if the lender successfully matches with a DI who has posted an interest rate ρ_j , and otherwise earns r from an ON RRP transaction with the Fed.




Figure 2: Timeline

4 Equilibrium

In this section, we fully characterize the equilibrium in the benchmark model and explore its properties. We start by deriving the optimal behavior of DIs and lenders, which allows us to formally define an equilibrium. We then establish existence and uniqueness. Finally, we discuss several important properties of equilibria, including how rates, trading volume, and take-up at the ON RRP facility respond to changes in policy.

¹⁶The Poisson matching function is standard in this literature and, as we show later, does a nice job in matching the data. However, almost all of our analytical results extend to more general matching functions; we formalize this claim in Appendix A.2.

4.1 Optimal Strategies and Definition of Equilibrium

Decision-making occurs at three different stages. First, DIs have to decide whether or not to enter the interbank market given their balance sheet cost, c_j . Second, those DIs that enter have to choose an interest rate, ρ_j . Finally, given the interest rates that have been posted, lenders have to choose which one to approach. In order to describe optimal behavior at each stage, we work backwards.

Optimal Search by Lenders. Once interest rates have been posted, each lender must choose the submarket (or mix between submarkets) that offers the maximum expected payoff, taking into account both the interest rate being offered in that submarket *and* the probability of being matched. In particular, the expected payoff from a lender choosing a submarket with interest rate ρ_j and queue length q_j is

$$u(\rho_j, q_j) = \left[\frac{1 - e^{-q(c_j)}}{q(c_j)} \right] \rho_j + \left[1 - \frac{1 - e^{-q(c_j)}}{q(c_j)} \right] r. \quad (1)$$

The first term captures the probability that the lender is matched with a DI, in which case the lender earns ρ_j , while the second term captures the probability that the lender fails to match, in which case the lender will approach the ON RRP facility and earn the rate r .

Let U denote the maximum expected payoff that a lender can obtain, or what we will call the “market utility.” In equilibrium, then, any submarket with $q_j > 0$ must satisfy

$$u(\rho_j, q_j) = U. \quad (2)$$

That is, in equilibrium, any DI that is able to attract lenders must deliver an expected payoff equal to the market utility: For example, if a DI posts a relatively low interest rate, lenders must be compensated with a high probability of being matched (i.e., a short queue length), and vice versa. Using (1) to solve (2) yields

$$\rho_j = r + \left[\frac{q_j}{1 - e^{-q_j}} \right] (U - r). \quad (3)$$

Optimal Interest Rate Posting by DIs. The lenders’ indifference condition in equation (3) lays bare the trade-off facing DIs: They can post a low interest rate and match with low probability, or they can attract a longer queue by posting a higher interest rate, in which case they will match with higher probability. Taking this trade-off as given, a DI with balance sheet cost c_j who has entered the interbank market solves the following profit maximization problem:

$$\begin{aligned} \max_{\rho_j, q_j} \quad & [1 - e^{-q_j}] (R - c_j - \rho_j) \\ \text{sub to} \quad & \rho_j = r + \left[\frac{q_j}{1 - e^{-q_j}} \right] (U - r), \end{aligned} \quad (4)$$

where the market utility U is taken parametrically by each DI. From the objective function, (4), it’s clear that a DI’s expected profits are equal to the product of the probability of being

matched, $1 - e^{-q_j}$, and the revenue from accepting a deposit, $R - c_j - \rho_j$, taking as given the positive relationship between interest rates and queue lengths.

One can substitute the constraint into (4), which yields an objective function that is strictly concave over a single choice variable, q_j . Hence the first order condition delivers the optimal queue length for any c_j and any market utility U :

$$q_j \equiv q(c_j; U) = \log \left(\frac{R - c_j - r}{U - r} \right). \quad (5)$$

From (3), then, the optimal interest rate for a DI with balance sheet cost c_j , given U , is

$$\rho_j \equiv \rho(c_j; U) = r + \log \left(\frac{R - c_j - r}{U - r} \right) \left[\frac{(R - c_j - r)(U - r)}{R - c_j - U} \right]. \quad (6)$$

Optimal Entry by DIs. A DI enters the interbank market if, and only if, its expected profits from doing so are nonnegative.¹⁷ One can easily show that a DI's profits are decreasing in c_j for any U , so that the optimal entry decision is determined by a cutoff rule: for any $U > r$, there exists a unique $\bar{c} > 0$ such that profits are nonnegative if, and only if, $c_j \leq \bar{c}$. Substituting (5) and (6) into (4) and solving reveals that this cutoff satisfies

$$\bar{c}(U) = R - U. \quad (7)$$

Market Clearing. The previous analysis describes the optimal decisions by lenders and borrowers, taking as given the market utility U . The final condition requires that markets clear:

$$\int_0^{\bar{c}(U)} q(c_j; U) dG(c_j) = \lambda. \quad (8)$$

In words, (8) requires that aggregating the queue lengths (or expected number of lenders per DI) across the active DIs yields the total measure of lenders in the market, λ .

Definition of Equilibrium. Given the results above, an equilibrium is a market utility U^* , a cutoff $\bar{c} = \bar{c}(U^*)$, queue lengths $q(c_j) = q(c_j; U^*) > 0$ and interest rates $\rho(c_j) = \rho(c_j; U^*) \in (r, R)$ for all $c_j < \bar{c}$ such that

1. Lenders are indifferent between all active DIs; i.e., (2) is satisfied.
2. Given lenders' search behavior, those DIs that enter the market choose interest rates to maximize profits; i.e., (6) is satisfied.
3. DIs enter the market if, and only if, it is profitable to do so; i.e., (7) is satisfied.
4. Markets clear; i.e., (8) is satisfied.

¹⁷In other words, we are assuming that a DI will stay out of the interbank market if it would not attract any lenders by entering. One could motivate this assumption by assuming that there was a cost $\epsilon > 0$ associated with posting an interest rate, where ϵ was arbitrarily close to zero.

4.2 Characterization of Equilibrium.

In this section, we establish that there exists a unique equilibrium, and offer a complete characterization. As we will show, the equilibrium allocation—i.e., the entry decision of DIs and the subsequent matching probabilities—are completely determined by the *spread* between the IOER and ON RRP rates. Intuitively, this spread is the key policy decision because it determines (i) the gains from trade between lenders and DIs, and (ii) the share of these gains that are appropriated by DIs, through its effect on the ratio of lenders to DIs or “market tightness.” For a given spread, we show that the *level* of policy rates simply shifts the distribution of interest rates up or down in a linear fashion.

Formally, let $s = R - r$ denote the spread between the IOER and ON RRP rates. Then, using (7), we can rewrite the market clearing condition (8) as

$$\int_0^{\bar{c}} \log \left(\frac{s - c_j}{s - \bar{c}} \right) dG(c_j) = \lambda. \quad (9)$$

Equation (9) reveals that characterizing the equilibrium boils down to solving one equation in one unknown, \bar{c} . Moreover, since \bar{c} only depends on s , and not on the specific values of R and r , so too do the queue lengths. In particular, the equilibrium queue length at a DI with balance sheet cost $c_j \leq \bar{c}$ is

$$q(c_j) = \log \left(\frac{s - c_j}{s - \bar{c}} \right). \quad (10)$$

It is in this sense that allocations only depend on policy through the spread between R and r . Given the cutoff and queue lengths, the remaining equilibrium objects follow immediately:

$$\rho(c_j) = r + \frac{q(c_j)}{1 - e^{-q(c_j)}}(s - \bar{c}) \text{ for all } c_j \leq \bar{c}, \text{ and} \quad (11)$$

$$U = r + s - \bar{c}. \quad (12)$$

Equations (11) and (12), respectively, illustrate that the spread also determines the “markup” that is charged by a DI with balance sheet costs c_j , $\rho(c_j) - r$, as well as the lender’s share of the surplus, $U - r$.

In the Appendix, we establish that there exists a unique solution to equation (9). In particular, if $s > 0$, then $\bar{c} > 0$ and there is a strictly positive measure of active DIs. If $s = 0$, of course, there are no gains from trade and the market shuts down, i.e., $\bar{c} = 0$. The following proposition summarizes.

Proposition 1. *There exists an equilibrium and it is unique. The cutoff $\bar{c} > 0$ if, and only if, $s > 0$.*

4.3 Properties of Equilibria

In this section, we exploit the characterization above to derive the relationship between DIs’ balance sheet costs, the interest rates they post, the queue lengths they attract, and the corresponding profits they earn. Then, we conduct comparative statics to explore how equilibrium outcomes respond to changes in the policy rates, r and R .

Costs, Rates, Queues, and Profits. Given an equilibrium cutoff $\bar{c} < s$, notice immediately that $q'(c_j) = -1/(s - c_j) < 0$ for all $c_j \leq \bar{c}$, so that

$$\rho'(c_j) = [1 - e^{-q(c_j)}(1 + q(c_j))]q'(c_j) < 0.$$

Hence, DIs with lower balance sheet costs post higher rates and attract longer queue lengths in equilibrium. As a result, a DI's profits $\Pi(c_j) = [1 - e^{-q(c_j)}][R - c_j - \rho(c_j)]$ are also strictly decreasing in c_j , with $\Pi(\bar{c}) = 0$.

Given the monotonic relationship between balance sheet costs and posted rates, the maximum and minimum rates that are offered in equilibrium must be $\rho(0)$ and $\rho(\bar{c})$, respectively. As we establish in Proposition 2, the former lies strictly below R and the latter lies strictly above r . Hence, absent any other frictions, the average traded rate will surely lie within the target range (r, R) .

Proposition 2. *For any $R > r \geq 0$,*

$$r < \rho(\bar{c}) < \rho(0) < R. \tag{13}$$

To understand the first inequality in (13), note that lenders have the option to receive a rate $\rho_j > r$ with *some* probability (and receive r otherwise). Hence, in order to generate any demand, the DI with balance sheet cost \bar{c} also has to offer a rate $\rho(\bar{c}) > r$. Since $\rho(c_j) > r$ for all $c_j \leq \bar{c}$, we see that lenders always appropriate some of the surplus from DIs—since lenders can search for another DI, the market is somewhat competitive. To understand the last inequality in (13), consider a DI with zero balance sheet costs. Since search frictions imply that queue lengths are not perfectly elastic, this DI has incentive to offer a rate $\rho(0) < R$. Hence, since $\rho(c_j) < R$ for all $c_j \leq \bar{c}$, we see that DIs are also assured of retaining some of the joint surplus—since lenders face frictions in searching for a better counterparty, DIs have some market power.

Comparative Statics. In the previous analysis, we established that the cutoff, \bar{c} , and the queue lengths at each type of DI, $q(c_j)$, are completely determined by s , λ , and $G(c_j)$. Hence, so long as the spread remains constant, changes in R and r alone have no effect on the aggregate market tightness, the number of lenders that match, and hence on the volume of lenders that use the ON RRP facility. Moreover, from (11), we know that increasing both the ON RRP and IOER rates by the same amount simply shifts the entire distribution of interest rates up one-for-one. However, a policy change that effects the spread $s = R - r$ will have implications for DIs' entry decisions, queue lengths, and the volume of trade.

To see this, consider the effects of a marginal increase in R , holding r constant. The first effect comes from a change in the *level* of interest rates: in response to an increase in the IOER rate, DIs will have more incentive to attract deposits, and hence they will raise the rates they offer. The second effect comes from a change in the *spread*: in response to an increase in the gains from trade, previously inactive DIs will have incentive to enter the market. With more DIs per lender, the measure of lenders that match with a DI increases, and the volume of lenders at the ON RRP facility falls. Moreover, as lenders become relatively more scarce, DIs will offer even better interest rates in response to an increase in competition.

To illustrate these effects analytically, let $\hat{\rho}$ denote the average rate that is traded between DIs and lenders, and let μ^ℓ denote the fraction of lenders that trade with a DI, so that $1 - \mu^\ell$

is the fraction of lenders at the ON RRP facility. Using the previous results, and aggregating across submarkets, one can show that

$$\hat{\rho} = r + \frac{s - \bar{c}}{\mu^\ell} \quad (14)$$

and

$$\mu^\ell = \frac{G(\bar{c})}{\lambda} \int_0^{\bar{c}} [1 - e^{-q(c_j)}] dG(c_j). \quad (15)$$

Lemma 3. *Holding $r \geq 0$ fixed, an increase in $R > r$ causes the average traded rate to rise and the fraction of lenders who use the ON RRP facility to fall. In particular,*

$$\mu^\ell \frac{\partial \hat{\rho}}{\partial R} = \int_0^{\bar{c}} e^{-q(c_j)} \left\{ \frac{[1 - e^{-q(c_j)}] g(c_j)}{\lambda \mu^\ell} \right\} dc_j \geq 0, \quad (16)$$

and

$$\frac{\partial \mu^\ell}{\partial R} = \frac{G(\bar{c})(s - \bar{c})}{\lambda} \left\{ \mathbb{E} \left[\left(\frac{1}{s - c_j} \right)^2 \right] - \mathbb{E} \left[\frac{1}{s - c_j} \right]^2 \right\} \geq 0. \quad (17)$$

The effect of a change in R on $\hat{\rho}$ has an intuitive interpretation. The first term in the integrand in (16) is the probability that each DI in submarket j doesn't match—or the “excess demand” of DIs in that submarket—while the second term is the fraction of total trades that occur in submarket j . Recall that an increase in R causes an increase in the potential gains from trade between DIs and lenders. Therefore, roughly speaking, equation (16) states that the share of these gains that accrue to the lenders, via higher interest rates, is equal to the trade-weighted excess demand for loans across DIs. In other words, when DIs are desperate to match, an increase in R is mostly passed along to lenders via higher rates.

Similar comparative statics can be derived for a change in r , holding R constant. In particular, since

$$\frac{\partial \hat{\rho}}{\partial R} + \frac{\partial \hat{\rho}}{\partial r} = 1$$

and $\partial \hat{\rho} / \partial R < 1$, it follows that $\partial \hat{\rho} / \partial r > 0$. However, since an increase in r decreases the spread s , in this case \bar{c} falls. One can show that, as a result, $\partial \mu^\ell / \partial r < 0$ and take-up at the ON RRP facility increases; we omit the derivation here in the interest of space.

5 Calibration

In the previous section, we derived analytically the key properties of equilibrium and the implications for monetary policy implementation. A truly helpful model of the federal funds market, however, should be able to fit the data and provide guidance for quantitative questions as well. In this section, we calibrate the structural parameters of the model to match a few key moments of the data from the federal funds market in the fourth quarter of 2015, before the FOMC decided to raise rates. We use our calibration to illustrate that the model is sufficiently flexible to match the data, and to provide a more complete picture of the implied relationships between policy rates, balance sheet costs, and the distributions of posted and traded rates. Then, we use

our calibrated model to forecast interest rates and trading volumes when the target range for the FFR is raised to $1/4 - 1/2$ percent, as it was in December 2015. We find that our model’s quantitative predictions are very close to the actual data, which provides some justification for the quantitative policy analysis we do in Section 6.

5.1 The Model and the Data

Identifying Agents and Trades. We first describe how we map the agents and trades in our model to their counterparts in the data. The first key step is to identify the institutions in the real world that have access to the IOER and ON RRP rates, i.e., to identify those institutions that correspond to the “DIs” and “lenders” in our model. This is fairly straightforward: Essentially all DIs other than the GSEs qualify to earn the IOER rate, while the ON RRP facility is open to both GSEs and to prime money market funds.¹⁸

The second step is to identify the types of trades that occur in the real world that correspond to the trades that occur between lenders and DIs in our model. As we discussed in Section 2, not all institutions that have access to the ON RRP facility can participate in the federal funds market; instead, these institutions typically structure their loans as eurodollar contracts. Unfortunately, data regarding the rates paid between DIs and ON RRP counterparties in eurodollar contracts is not publicly available. To get around this, we assume that whether a loan is classified as a federal funds contract or a eurodollar contract has no impact on the terms of trade. This enables us to map the traded rates in the model to those captured in the data as federal funds transactions. We take comfort that several studies have described federal funds and eurodollars as near-perfect substitutes, and differences between the FFR and broader eurodollar rates as minimal.¹⁹

Targets. We base our calibration on the data corresponding to last quarter of 2015, prior to the FOMC’s decision to raise rates on December 16th. Throughout the period we consider, the IOER rate was set to 25 basis points and the ON RRP rate was set to 5 basis points. Hence, we set $R = 0.25$ and $r = 0.05$. Moreover, though allotment at the ON RRP was officially capped at \$300 billion, this cap was not binding over the period (since September 2014 the cap has been slack at all but quarter-end days). Hence, our assumption of unlimited allotment at the ON RRP facility is appropriate.

To capture dispersion in balance sheet costs, we calibrate $G(c)$ using a Gamma distribution with mean and standard deviation denoted μ_c and σ_c , respectively. In order to map the model’s predictions to dollars, we also need to define what a “unit of cash” is, in dollars. To do so, we let x denote the quantity of dollars held by each lender.

Hence, the parameters left to calibrate are $\{\lambda, x, \mu_c, \sigma_c\}$, which we discipline as follows. First, we match the average effective FFR and ON RRP take-up during the last quarter of 2015 up to

¹⁸The ON RRP facility is also open to several DIs, including all primary dealers. However, since they qualify to earn the IOER rate, DIs have little use for the ON RRP facility—as the Federal Reserve Bank of New York reports, GSEs and money market funds account for more than 98 percent of all bids at the ON RRP facility.

¹⁹See Bartolini et al. (2008), Demiralp et al. (2006), and Cipriani and Gouny (2015). The Federal Reserve Bank of New York has also started publishing in March 2016 the overnight bank funding rate, which does include Eurodollar transactions by DIs: Differences with the effective FFR are very small and short-lived.

December 15th, which are 12 basis points and \$90 billion, respectively.²⁰ Over the same period, the ON RRP facility averaged just over 45 bids per day, which are 31 percent of the (143) lenders with access to the facility. We also target a range of traded rates of 4 basis points, which encompasses more than 90 percent of the traded volume.²¹ These four targets are sufficient to pin down the four parameters left to calibrate. The measure of lenders is set at $\lambda = .375$.²² The parameters governing the distribution of balance sheet costs are set at $\mu_c = .15$ and $\sigma_c = .07$, which suggests that balance sheet costs are substantial and quite dispersed.

5.2 The implied distribution of rates

To flesh out the implications of our chosen parameter values and complete the description of equilibrium, we now use our calibrated model to illustrate the relationship between trading (or search) frictions, balance sheet costs, and the distribution of rates that are offered and traded. In the left panel of Figure 3, we plot the distribution of interest rates that are posted by DIs in equilibrium (light blue) and the distribution of interest rates that DIs actually pay (dark blue). In the right panel, we plot the distribution of interest rates chosen or “asked” by lenders (light blue), and the distribution of interest rates they actually receive (dark blue).

Note that, as we established in Proposition 2, the lowest posted rate is bounded away from r and the highest posted rate is bounded away from R . This is because matching frictions endow DIs with some market power, so that each side captures some of the gains of trade. Also note that, as we reported in Section 4.3, DIs posting higher interest rates are more likely to receive deposits, since they attract larger queue lengths. As a result, the average interest rate paid by DIs exceeds the average posted rate. On the flip side, lenders who choose DIs with high interest rates are less likely to trade and more likely to end up at the ON RRP facility earning an interest rate of r . Hence, lenders are ex-post heterogeneous despite being ex-ante identical.²³

Finally, note that the distributions of posted—and, to a lesser extent, traded—rates inherit the overall shape of the distribution of balance sheet costs among active DIs. Through the lens of the model, then, detailed data on traded rates would be potentially informative about both the level and dispersion of the underlying balance sheet costs. However, there are limits to what

²⁰The fact that interest rates are greater than r and take-up at the ON RRP is positive should further convince the reader that a search-based model is useful for studying this market. A frictionless, competitive market would have trouble generating these two facts simultaneously: if there was excess supply of funds from lenders, so that take-up at the ON RRP was positive, then the FFR would be driven to r ; if there was excess demand, so that the FFR was greater than r , then there would be no deposits at the ON RRP. We illustrate this point in greater detail in Appendix A.3.

²¹All data exclude end-month dates and are publicly available at the Federal Reserve Bank of New York website, <http://www.ny.frb.org/markets/>. Also see Cipriani and Cohn (2015) for additional details on the distribution of traded rates in the federal funds market. The Federal Reserve Bank of New York also reports the standard deviation, maximum, and minimum of the federal funds rates traded per day. However, we are hesitant to use these data points because both the maximum and minimum rates are clearly outliers. For example, the maximum traded rate is above the IOER rate—occasionally double the IOER rate—reflecting the stray DI that has trouble meeting reserve requirements. As a result, the standard deviation is not very informative about the underlying distribution either.

²²By construction, the cash per lender x matches the average dollar value of a bid at the ON RRP facility, about \$3.5 billion.

²³In Section 7, we provide two extensions that introduce some ex-ante heterogeneity in lenders that may also explain some of the dispersion in traded rates.

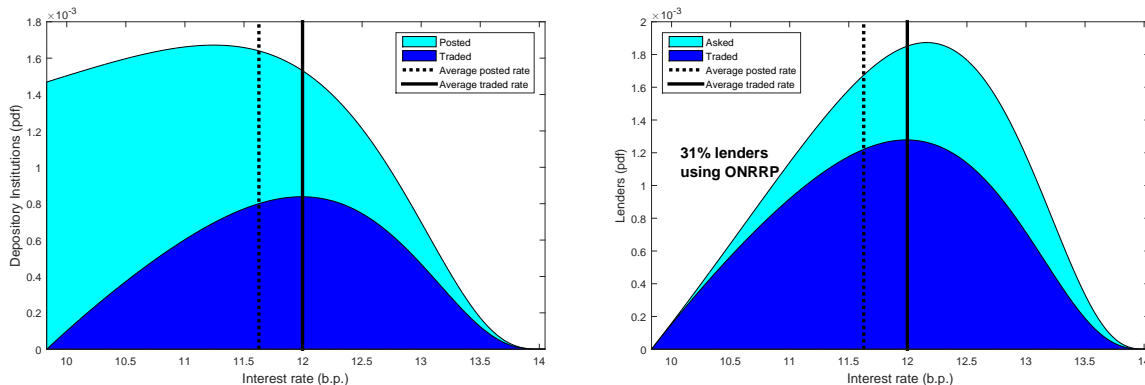


Figure 3: Posted, Asked, and Traded Interest Rates

can be learned from these rates, as they only reflect the balance sheet costs of *active* DIs, and not those of potential entrants with $c > \bar{c}$.²⁴

5.3 Liftoff

On December 16, 2015 the FOMC decided to raise the target range for the FFR to 1/4 to 1/2 percent. It did so by setting the IOER at the top of the target range, 50 basis points (bps), and by increasing the spread from 20 to 25 basis points, so that the ON RRP rate was set at the bottom of the target range. These changes in rates provide a natural test of our calibrated model.²⁵

Table 1 summarizes the key moments in the data before and after liftoff. After raising the target range, the effective FFR rose to 37 basis points, while activity at the ON RRP facility fell—both in volume and the number of bids—though the ongoing changes in the banking and money-market funds industry may also have played a role in decreasing demand at the ON RRP facility.²⁶ The range of traded rates, measured as the spread between the 5th and 95th percentiles, increased slightly to 6 basis points.

Table 1 also reports the model’s predictions, keeping all the structural parameters as in the benchmark calibration but adjusting the policy rates, i.e., the IOER and ON RRP rates. The

²⁴The distribution of balance sheet costs among DIs that are initially inactive determines the strength of the entry margin which is, in turn, an important determinant of the model’s predictions for several policy experiments like, say, an increase in the IOER.

²⁵Of course, these changes are arguably too small to constitute a definitive test of the model’s predictions. Ideally, one would like to compare the model’s predictions after multiple rate hikes, along with changes to other structural parameters, such as the distribution of balance sheet costs, the measure of lenders, or the total supply of funds, as all of these variables are likely to impact monetary policy normalization in the future. Other than liftoff, the only source of variation in rates are the small-scale tests that the Federal Reserve Bank of New York conducted in the late Fall of 2014. The model is able to fit the observed responses in the FFR very well but less so with the demand at the ON RRP facility—though it is likely that the tests were too short (less than 10 business days) to wash out the large daily variation in ON RRP take-up. See the working paper version, Armenter and Lester (2015), for further details.

²⁶We report the averages over the first half of 2016, excluding month-end dates.

model exactly matches the data before liftoff by construction, as these are the moments we used for calibration. The model’s predictions for liftoff are very close to the data. The effective FFR is predicted to be at 36 basis points, within one basis point from the average observed in the data. The model slightly underestimates the fall in participation at the ON RRP facility both in volume and bids, the former falling by \$4 billion more than predicted by the model. The range of traded rates expanded by a modest amount, only slightly more in the data than in the model. Note that the drop in ON RRP demand and expansion in the range of traded rates are exclusively driven by the increase in the spread between the IOER and the ON RRP rate in the model.

Table 1: Data and Model Before and After Liftoff

| | Before liftoff | | After liftoff | |
|--------------------|-----------------------|----------|----------------------|----------|
| | Data | Model | Data | Model |
| ON RRP rate | 5 b.p. | | 25 b.p. | |
| IOER | 25 b.p. | | 50 b.p. | |
| FFR | 12 b.p. | 12 b.p. | 37 b.p. | 36 b.p. |
| ON RRP take-up | \$90 bn. | \$90 bn. | \$74 bn. | \$78 bn. |
| ON RRP bids (%) | 31% | 31% | 25% | 27% |
| Range traded rates | 4 b.p. | 4 b.p. | 6 b.p. | 5 b.p. |

See text for description of data calculations.

6 Quantitative Policy Analysis

In our benchmark model, we established that an ON RRP facility provides an effective floor for interest rates, and hence is a useful tool for ensuring that interest rates remain in the target range. There are, however, concerns that a large volume of activity at the ON RRP facility could have undesirable effects on the financial industry, and perhaps even negative implications for financial stability.²⁷ Recognizing these risks, the FOMC stated that it intends to rely on the ON RRP facility “only to the extent necessary” and will “phase [the overnight reverse repurchase facility] out when it is no longer needed.”²⁸

However, since policymakers have little experience with the current implementation framework, the role of the ON RRP facility in supporting interest rates—and the effect of either reducing its size or eliminating it altogether—remains highly uncertain. In this section, we use our calibrated model to quantify the role of the ON RRP facility, and to study various alternatives to limit demand at the facility. We first consider what happens to market rates and trading

²⁷The minutes from the June 2014 FOMC meeting noted that “a relatively large ON RRP facility had the potential to expand the Federal Reserve’s role in financial intermediation and reshape the financial industry in ways that were difficult to anticipate.” See Frost et al. (2015) for a detailed discussion of potential secondary effects associated with an ON RRP facility.

²⁸All quotations are from “Policy Normalization Principles and Plans,” Federal Reserve System Press Release, September 17, 2014.

volumes when the ON RRP facility is removed entirely. Then, we analyze two alternative solutions for reducing the volume of activity at this facility: setting a limit on the daily volume of orders that can be accepted, and increasing the spread between the IOER and ON RRP rates to encourage more trading within the federal funds market.

6.1 Eliminating the ON RRP facility

We start by using our model to study the effects of completely phasing out the ON RRP facility, which allows us to quantify this facility’s role in supporting federal funds rates. In what follows, we use our calibrated model to consider a counterfactual exercise in which the IOER is set at the top of the target range and the ON RRP facility is absent.²⁹ We compute the effective FFR predicted by the model and check whether it remains inside the target range, i.e., within 25 basis points of the IOER.

Figure 4 displays the model’s predictions for the FFR for an array of IOER values; in one case the ON RRP rate is set 25 bps below the IOER rate, and in the other case the ON RRP facility is removed entirely. The grey area indicates the target range. The main takeaway is that the ON RRP facility is indispensable for monetary policy implementation, even when the target range is fairly low. The IOER does have a pull on rates when the ON RRP facility is absent, but it is not enough to keep the effective FFR within the target range when the upper bound exceeds 75 basis points. Indeed, the counterfactual effective FFR drops further and further below the target range, with an economically meaningful gap opening up fairly quickly.

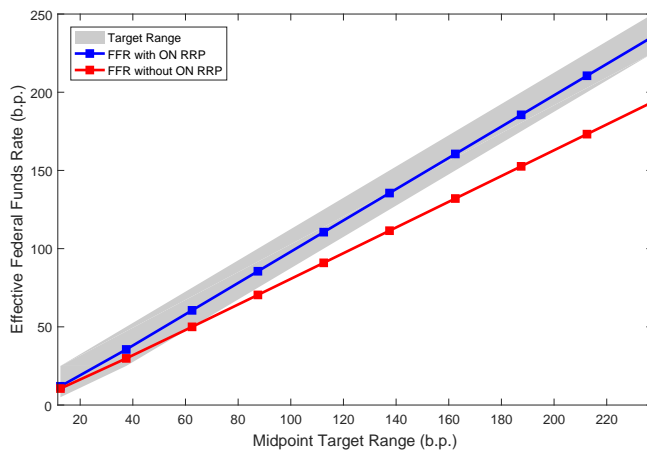


Figure 4: Contribution of the ON RRP Facility

It is worth noting that, as the IOER increases in the counterfactual exercise, DIs return

²⁹This exercise requires us to take a stand on the alternative investment available to a lender in the absence of the ON RRP facility. GSEs are very reluctant to lend outside the federal funds market, so when they fail to find a DI, their funds are likely to sit unremunerated at their balance accounts with the Fed. Similarly, clearing accounts for money market funds are typically non-remunerated, though this may change if overnight rates become substantially higher. In any case, we assume that the outside option for lenders, absent the ON RRP facility, stays at zero—the worst-case scenario for implementation and a very likely one in the medium term.

to the federal funds market (i.e., \bar{c} rises), increasing trading volume and providing additional support to rates. Indeed, the terms of trade for lenders steadily improves: the share of the surplus received by lenders who successfully match is 60 percent when the IOER is set to 50 basis points, 70 percent when the IOER reaches 100 basis points, and approaches 80 percent when the IOER reaches 200 basis points. In this sense, the FF market shows notable resilience, though not enough to keep the rates within the target range.

6.2 A cap on ON RRP activity.

The analysis above suggests that there may exist a tension between policymakers' intention to limit the size or "footprint" of the ON RRP facility and their desire to effectively implement a target policy rate. Eliminating the ON RRP facility altogether would be a rather blunt resolution of this tension, so we next consider an intermediate solution: A cap on the volume of activity, or a "maximum allotment," that is allowed at the ON RRP facility. Indeed, as we noted earlier, daily volume at the ON RRP facility was capped at \$300 billion prior to December 2015.

To study the effects of such a cap within the context of our model, let ϑ denote the demand for deposits at the ON RRP in equilibrium, and let κ denote the maximum allotment chosen by policymakers. If demand at the ON RRP facility exceeds the maximum allotment, so that $\kappa < \vartheta$, we assume that each lender is able to deposit $\frac{\kappa}{\vartheta}$ units of cash at the posted ON RRP rate, r , while earning zero on the remainder.³⁰ Otherwise, if $\kappa \geq \vartheta$, lenders are able to deposit all of their cash at the ON RRP facility, as in our benchmark model.

In Appendix A.4, we provide a full analytical characterization of the equilibrium when the ON RRP facility is capped. Trivially, if the cap is not binding, then the equilibrium characterization in our benchmark model is unchanged. If the cap is binding, the equilibrium remains remarkably similar to the equilibrium in our benchmark model, with one key difference: the relevant outside option for lenders is not the repo rate r , but rather the *effective repo rate*

$$\tilde{r} = \left(\frac{\kappa}{\vartheta}\right) r. \quad (18)$$

As a result, the key parameter for equilibrium allocations is not the spread s , but rather the *effective spread* $\tilde{s} = R - \tilde{r}$.³¹

From (18), it's clear that a tighter cap drives down the effective repo rate, and hence the FFR. Indeed, if the cap is sufficiently small, it's possible for the FFR to fall out of the target range. A natural question, then, is how likely this is to occur for a given value of κ , i.e., how imposing a particular cap affects the robustness of monetary policy implementation. We study this question below.

³⁰Alternatively, one could imagine that each lender gets to deposit their unit of cash with probability $\frac{\kappa}{\vartheta}$ if $\kappa < \vartheta$ and probability 1 otherwise; given our specification of linear utility, the two assumptions are identical. In practice, the Desk awards funds by auction. Modeling explicitly the allocation mechanism would not change the substance of our results, so long as a lender's effective return at the ON RRP facility, defined in (18), is a strictly increasing function of $\frac{\kappa}{\vartheta}$ for $\frac{\kappa}{\vartheta} \in (0, 1)$, with $\lim_{\frac{\kappa}{\vartheta} \rightarrow 0} \tilde{r} = 0$ and $\lim_{\frac{\kappa}{\vartheta} \rightarrow 1} \tilde{r} = r$.

³¹The characterization is more involved than in our benchmark model, however, because the (exogenously specified) repo rate r is replaced by the (endogenously determined) effective repo rate \tilde{r} . This adds one additional "fixed point" condition to the equilibrium characterization: Given an effective spread $\tilde{s} = R - \tilde{r}$, one can determine the entry cutoff \bar{c} , and thus the volume at the ON RRP facility $\vartheta(\bar{c})$, which then must be consistent with the effective repo rate $\tilde{r} = (\kappa/\vartheta(\bar{c})) r$.

We frame our discussion in terms of the supply of funds in the market. In reality, the supply of funds in the federal funds market is a key source of uncertainty—a “known unknown”—and a natural source of vulnerability for implementation in an environment with excess reserves. In our model, the equilibrium FFR is strictly decreasing in the measure of lenders, λ . Moreover, if the ON RRP facility is capped at $\kappa < \infty$, one can show that there exists a critical value, λ_1 , such that the cap binds for all $\lambda \geq \lambda_1$; and that there exists a second critical value, $\lambda_2 > \lambda_1$, such that the average traded rate is less than r for all $\lambda \geq \lambda_2$.³²

These analytic results confirm the policymakers’ concerns that an increase in the supply of funds could threaten implementation, but the underlying question is inherently quantitative: How much more funds can be absorbed before the average federal funds rate steps outside the target range? To answer this question, we return to the calibrated model from Section 5, and set the target range for the FFR to 25 to 50 basis points. We will consider three values of κ : the pre-liftoff level, \$300 bn, and plus and minus \$150 billion. We will assume that changes in the supply of funds are due, in equal parts, from an increase in the numbers of lenders and an increase in the amount of cash per lender. It turns out that this is not innocuous in our model: A surge in the supply of funds driven entirely by the number of lenders puts more downward pressure to the FFR than a comparable supply surge driven exclusively by the amount of cash per lender. The reason is that search frictions constrain the number of matches made but, conditional on being matched, a DI and a lender can transact as large as an amount as they wish.³³ We discuss in the next subsection how our quantitative results are robust to other assumptions along this dimension.³⁴

Table 2 reports, for each of the cap values considered, the amount that the supply of funds would need to increase in order for (i) the cap at the ON RRP facility to bind and (ii) the FFR to fall below the floor. Notice that, in all three cases, there is a substantial “buffer” between the increase in the supply of funds needed for the ON RRP cap to bind and the amount needed for the FFR to breach the floor, which affords policymakers some additional room to maneuver.

With a cap of \$300 billion, it would take about \$560 bn of additional supplied funds for the ON RRP facility to become over-subscribed, which is more than twice the initial amount of spare capacity at the ON RRP facility. Hence, our model predicts that the federal funds market will be able to absorb a large fraction of a surge in the supply of funds. Two margins are key for the resilience of the federal funds market in the model. First, lenders and active DIs can scale up the contracts, absorbing most of the increase in the amount of cash per lender. Second, matching in the market improves due to the entry of additional DIs set to make a profit.

Staying with a cap of \$300 bn, we also see from Table 2 that it takes a substantial increase in the supply of funds for the FFR to breach the floor, \$1,629 billion, which is about double the initial supply of funds. As discussed earlier, once the cap is binding, additional increases in the supply of funds bring substantial downward pressure on rates as the effective repo rate falls. However, there are some offsetting effects that explain the high threshold for a floor breach. As the effective repo rate falls, some DIs find it profitable to become active, improving the lenders’ chances of matching in the market. This decreases demand for the ON RRP facility, relieving some of the pressure on both the cap and the FFR.

Not surprisingly, a larger cap implies that a larger surge in the supply of funds can be

³²For a formal statement and proof of these results see the working paper version.

³³Recall that balance sheet costs are per unit of cash, so total balance sheet costs do scale up with the amount

Table 2: Thresholds: Implementation robustness

Rates (percentage points): $r = .25, R = .50$

| ON RRP Cap | Constant supply | | Supply thresholds | |
|------------|------------------------|--|--------------------------|----------------|
| | Spare ON RRP capacity | | ON RRP cap binds | Floor breached |
| \$150 bn. | \$60 bn. | | + \$200 bn. | + \$1,194 bn. |
| \$300 bn. | \$210 bn. | | + \$559 bn. | + \$1,629 bn. |
| \$450 bn. | \$360 bn. | | + \$884 bn. | + \$2,005 bn. |

Rates (percentage points): $r = 1.5, R = 1.75$

| ON RRP Cap | Constant supply | | Supply thresholds | |
|------------|------------------------|--|--------------------------|----------------|
| | Spare ON RRP capacity | | ON RRP cap binds | Floor breached |
| \$150 bn. | \$60 bn. | | + \$200 bn. | + \$426 bn. |
| \$300 bn. | \$210 bn. | | + \$559 bn. | + \$838 bn. |
| \$450 bn. | \$360 bn. | | + \$884 bn. | + \$1,113 bn. |

accommodated without compromising monetary policy implementation. It is worth noting, though, that larger caps provide proportionally less room for monetary policy implementation. For example, tripling the initial spare capacity at the ON RRP facility, from a cap of \$150 billion to \$450 billion, less than doubles the largest supply increase that can be absorbed before the FFR breaches the floor. In any case, the fact that the thresholds are quite large suggests that even a modest cap on ON RRP facility should not interfere with monetary policy implementation, at least when the target range is relatively low.

Table 2 also displays the results from repeating the previous exercise with a higher target range of 1.5 percent to 1.75 percent. The increase in the supply of funds needed to equate ON RRP demand to its cap is unchanged: This follows immediately from our analytic results showing that demand at the ON RRP depends only on the spread when the cap is not binding. However, note that implementation is a bit more fragile at higher rates, since a substantially smaller increase in the supply of funds can cause the effective FFR to breach the floor. The reason is that, once the ON RRP facility is oversubscribed, higher IOER and ON RRP rates imply that the effective repo rate falls faster with each additional dollar of demand. For example, if demand at the ON RRP facility is twice the cap, the effective repo rate is 12 basis points below the floor of a target range of 25 – 50 basis points, but a whopping 75 basis points under the floor if the target range is 150 – 175 basis points.

6.3 Changing the Spread

Our model suggests an alternative, indirect way to reduce the volume of trade at the ON RRP facility: increasing the spread between the IOER and ON RRP rates, with the former set above

of the transaction.

³⁴Additional robustness checks are available from the authors upon request.

the top of the trading range. According to our model, this would induce more trading in the federal funds market—as more DIs enter to exploit the gains from trade—and hence fewer lenders would end up at the ON RRP facility.

Figure 5 illustrates the effect of widening the spread between the IOER and ON RRP rates; we plot the distribution of traded rates under our benchmark calibration, before and after increasing the IOER rate R by 5 basis points, from 25 to 30 basis points. There are several things to note. First, DIs pass to lenders, on average, 3.5 basis points out of the 5 basis points increase in the IOER rate—a relatively large pass-through given that the average traded rate, at 12 basis points, is closer to the repo rate than to the IOER. Second, note that there is a significant increase in the dispersion of rates; in response to the hike in the IOER rate, additional DIs enter, the market becomes more competitive, and interest rates shift up. Lastly, since there are more DIs per lender, note that the fraction of lenders that ultimately end up depositing their funds at the ON RRP facility falls by a bit more than 10 percent.³⁵




Figure 5: Increasing the IOER Rate

To further explore the relationship between the spread and the target range, in Table 3 we recalculate the thresholds for the increase in the supply for different two spreads—35 basis points and 50 basis points, achieved by raising the IOER rate alone—when the target range is 150–175 basis points. Recall that, for this target range, the results for the baseline spread of 25 basis points are included in the second panel of Table 2. This exercise helps us to quantify when, and how much, an increase in the spread relieves pressure on the ON RRP facility.

Table 3 shows that a larger spread renders implementation more robust, as a bigger increase in the supply of funds can be absorbed before implementation is threatened. The reason is that, as we discussed earlier, widening the spread between IOER and ON RRP increases the gains from trade between parties, inducing some DIs to enter the federal funds market. Lenders see their terms of trade improve—increasing rates—and have an easier time finding a DI, thus

³⁵In the working paper version, Armenter and Lester (2015), we also illustrate the effects increasing the ON RRP rate, r , while keeping the IOER rate constant. We show that rates increase, though the pass through from DIs to lenders is smaller. Meanwhile, since the spread $s = R - r$ shrinks, the dispersion of rates decreases, while the fraction of lenders that ultimately end up depositing their funds at the ON RRP facility rises.

Table 3: Thresholds: Larger spreads

| <i>Rates (percentage points): $r = 1.5, R = 1.85$</i> | | | |
|--|-----------------------|--------------------------|----------------|
| Constant supply | | Supply thresholds | |
| ON RRP Cap | Spare ON RRP capacity | ON RRP cap binds | Floor breached |
| \$150 bn. | \$60 bn. | + \$263 bn. | + \$559 bn. |
| \$300 bn. | \$210 bn. | + \$662 bn. | + \$1,005 bn. |
| \$450 bn. | \$360 bn. | + \$1,004 bn. | + \$1,370 bn. |

| <i>Rates (percentage points): $r = 1.5, R = 2$</i> | | | |
|---|-----------------------|--------------------------|----------------|
| Constant supply | | Supply thresholds | |
| ON RRP Cap | Spare ON RRP capacity | ON RRP cap binds | Floor breached |
| \$150 bn. | \$60 bn. | + \$312 bn. | + \$770 bn. |
| \$300 bn. | \$210 bn. | + \$715 bn. | + \$1,234 bn. |
| \$450 bn. | \$360 bn. | + \$1,078 bn. | + \$1,630 bn. |

reducing demand at the ON RRP facility. Hence, a larger increase in the supply of funds can be accommodated before the ON RRP facility is oversubscribed and rates drop below the bottom of the target range.

In terms of quantitative predictions, our model suggests that the framework’s capacity to absorb additional funds by lenders increases by economically meaningful amounts as the spread widens. Averaging across Table 3, a ten basis points increase in the IOER delivers about \$200 billion additional absorption capacity before implementation is threatened. Even for a tight cap of \$150 billion, the threshold increases by more than \$100 billion. There are, however, limits to this approach: If we keep increasing the spread via the IOER, the effective FFR may trade *above* the top of the target range. For example, when the IOER and ON RRP rates are set at 2.0 percent and 1.5 percent, respectively, the effective FFR is just below 1.75 percent, so it could easily exceed the top of the target range if the supply of funds decreased.

7 Extensions

An attractive feature our framework is that it can be easily extended along several dimensions without compromising its tractability. In this section, we sketch three particularly relevant extensions to our basic model: one in which some lenders can only match with a restricted set of DIs, one in which some lenders do not have access to the ON RRP facility, and one in which there is aggregate uncertainty about the supply of funds in the market. All proofs can be found in the Appendix A.5.

7.1 Counterparty restrictions

Some lenders may only be able or willing to deposit their funds at a subset of DIs. Regulation, for example, requires that certain financial institutions (like GSEs) only deposit funds at DIs that are deemed sufficiently “safe.” Other institutions may only be willing to lend to counterparties that have been previously vetted or are part of a set of pre-existing relationships (or network). In any case, a natural question is whether counterparty restrictions can introduce additional dispersion in traded rates and/or hinder the implementation of monetary policy.

To focus our analysis squarely on this extension, we will assume that there is a fixed measure of DIs and all of them have zero balance sheet costs. To capture counterparty restrictions, suppose there is a measure d_s of “safe” DIs and a measure d_n of “non-safe” DIs. Moreover, assume that there is a measure λ_s of lenders that are required to deposit funds at a safe DI, while a measure λ_n of lenders can deposit funds at any DI.

In the proposition below, we establish that there is a unique equilibrium, and it can be one of two types, depending on the abundance of type s and n lenders, relative to the abundance of safe and non-safe banks. To be more precise, let

$$q_s = \frac{\lambda_s}{d_s}, \quad \text{and} \quad q_n = \frac{\lambda_n}{d_n}.$$

If $q_s > q_n$, then market segmentation arises: Type n lenders seek out and match with exclusively non-safe DIs, while type s lenders (by construction) seek out exclusively safe DIs. Intuitively, this type of equilibrium can only be supported if there is a relative abundance of type s lenders; otherwise, it is not sufficiently profitable for safe banks to serve only type s lenders, and they choose instead to set a rate that attracts type n lenders, too. Indeed, if $q_n \geq q_s$, then all DIs set the same rate and safe banks match with both types of lenders. In this latter type of equilibrium, all DIs earn the same profits and all lenders receive the same expected utility and thus the counterparty restrictions are essentially irrelevant. The following proposition formalizes the result.

Proposition 4. *Assuming $0 < q_j < \infty$ for $j \in \{s, n\}$, there exists a unique equilibrium. If $q_s > q_n$, then safe DIs serve the type s lenders exclusively at rate*

$$\rho_s = r + \frac{q_s e^{-q_s}}{1 - e^{-q_s}}(R - r)$$

while non-safe DIs trade only with type n lenders at rate

$$\rho_n = r + \frac{q_n e^{-q_n}}{1 - e^{-q_n}}(R - r) > \rho_s.$$

If $q_s \leq q_n$, then a strictly positive fraction of type n lenders match with safe DIs, and all lenders pursue the same rate

$$\rho_p = r + \frac{q_p e^{-q_p}}{1 - e^{-q_p}}(R - r)$$

where $q_p = (\lambda_s + \lambda_n) / (d_s + d_n)$.

Note that, in the equilibrium with market segmentation, lenders with counterparty restric-

tions will receive lower rates and match less often. As a result, these lenders will constitute a disproportionately high fraction of the take-up at the ON RRP facility. It is thus possible that some of the dispersion in traded rates is due to counterparty restrictions. However, also note that the results regarding monetary policy implementation stand: As long as the cap on ON RRP take-up is not binding, the rates offered by both safe and non-safe DIs will remain between the ON RRP rate and the IOER rate.

In the second type of equilibrium—where safe banks attract both types of lenders—the equilibrium rates, profits, and payoffs are identical to those that we would have obtained had we ignored the counterparty restrictions and aggregated lenders and DIs types. This suggests that counterparty restrictions can be safely ignored as long as they are not widespread, or if the subset of DIs deemed safe is sufficiently large.

7.2 Limited Access to ON RRP Facility

In the previous sections, we analyzed two of the primary choices that policymakers face when designing an ON RRP facility: the interest rate and the cap on volume. However, there is a third margin that is potentially important as well—namely, the list of institutions eligible to use the ON RRP facility. We did not make this margin explicit in the benchmark model simply because, at the introduction of the ON RRP facility, the Desk at the Federal Reserve Bank of New York designated a virtually comprehensive list of counterparties, including all the major lenders in the federal funds market, as well as several money market funds. We now study whether there was any basis for allowing such broad access to the ON RRP facility, or if allowing only a limited set of ON RRP counterparties would have sufficed.

To address this issue, consider our benchmark model with a unit measure of DIs with zero balance sheet costs. Again, these simplifying assumptions help us to focus exclusively on the new ingredients we are adding to the model. In particular, suppose that there is a measure of lenders, λ_0 , who do not have access to the ON RRP facility; if these “ineligible” lenders fail to deposit their funds at a DI, they earn zero overnight interest. There is also a measure λ_1 of lenders who do have access to the ON RRP rate; if these “eligible” lenders fail to deposit their funds, they earn the overnight interest rate r .

As we establish in the proposition below, there are two types of equilibrium that exist in disjoint regions of the parameter space. When there are sufficiently few ineligible lenders, the equilibrium exhibits market segmentation: A fraction of DIs post a relatively high rate and attract only eligible lenders, and the remaining DIs post a relatively low rate and attract only ineligible lenders. In equilibrium, the DIs are indifferent because the queue length of eligible lenders is longer than the queue length of ineligible lenders. Hence, DIs serving only eligible lenders match with high probability, but earn less profit per match. Meanwhile, since they have different outside options, eligible and ineligible lenders place different values on the probability of a match and the interest rate being offered—the standard requirement for the single-crossing property—and hence are content to search only in their designated submarket.

On the other hand, if there are sufficiently many ineligible lenders, then DIs have no incentive to offer rates that would attract eligible lenders. In particular, since eligible lenders have a better outside option, there simply aren’t sufficient gains from trade between an eligible lender and a DI; the profit from trading with ineligible lenders is simply too high. As a result, eligible lenders are shut out of the interbank market altogether. The following proposition summarizes.

Proposition 5. *If $r < Re^{-\lambda_0}(1 + \lambda_0)$, then there exists an equilibrium in which a fraction μ_0 of DIs post*

$$\rho_0 = \frac{Rq_0e^{-q_0}}{1 - e^{-q_0}} \geq r, \text{ with } q_0 = \frac{\lambda_0}{\mu_0};$$

the remaining $1 - \mu_0$ DIs post

$$\rho_1 = r + \frac{(R - r)q_1e^{-q_1}}{1 - e^{-q_1}} > \rho_0, \text{ with } q_1 = \frac{\lambda_1}{1 - \mu_0} > q_0;$$

and the fraction μ_0 is the unique solution to

$$(1 - e^{-q_0}) \rho_0 = (1 - e^{-q_1}) \rho_1.$$

Otherwise, if $r \geq Re^{-\lambda_0}(1 + \lambda_0)$, then all DIs post

$$\rho_0 = \frac{Rq_0e^{-q_0}}{1 - e^{-q_0}} < r, \text{ with } q_0 = \lambda_0.$$

Proposition 5 establishes that allowing broad access to the ON RRP facility makes it easier to contain the average traded rate within the target range. In particular, if the fraction of funds being supplied by ineligible lenders becomes sufficiently large, then the average traded rate can fall below the ON RRP rate, and eligible lenders will avoid this market entirely. Hence, our results suggest that the decision to designate a large list of eligible counterparties at the ON RRP facility will support successful implementation of monetary policy as rates begin to rise.

It is also worth highlighting the differential effects of increasing r and R in the context of this extension. In particular, as the ON RRP rate rises, the measure of DIs trading with eligible lenders shrinks. As a result, the interest rate that eligible lenders receive rises, but not one-for-one: As fewer DIs target eligible lenders, the market tightness q_1 rises, which puts *downward* pressure on the interest rate ρ_1 . Indeed, a hike in the ON RRP rate can shut out eligible lenders altogether from the market and thus drive traded rates below the floor. In contrast, when the IOER rate rises, the fraction of DIs trading with eligible lenders increases, and this creates additional *upward* pressure on the rates eligible lenders receive in equilibrium.

7.3 Aggregate uncertainty

In our model, the amount of spare capacity or “headroom” at the ON RRP facility is irrelevant, so long as the cap is not binding. In reality, however, the volume of bids at the ON RRP facility is not deterministic, as it is in our model, but rather it is both stochastic and notoriously volatile. As a result, the amount of headroom may indeed be important: Since it’s an important component of the effective repo rate, the amount of headroom determines the lenders’ outside option, and thus entry by DIs, the rates that they post, and the volume of trade that occurs in the interbank market.

To capture these considerations, suppose there is aggregate uncertainty about excess cash balances. In particular, suppose there is a measure λ of lenders and each lender has x units of cash, which is drawn from a cumulative distribution function $F(x)$ with mean $\mathbb{E}[x] = 1$. To study the simplest possible case, also assume that lenders match with DIs before the realization

of x , i.e., lenders establish a connection with a DI before they know the exact value of their excess cash holdings.

It is straightforward to establish that the model is linear in x if the ON RRP facility has no cap, implying that only the expected level of cash balances is relevant. However, if the ON RRP facility has a cap, then the value of the cap can impact traded rates even in states of the world in which it ultimately does not bind. To see this, consider the effective repo rate introduced in Section 6, where the dependence on the cap κ and the realization of x is made explicit:

$$\tilde{r}(x, \bar{c}) = \min \left\{ 1, \frac{\kappa}{x\vartheta(\bar{c})} \right\} r. \quad (19)$$

Lenders now compare their *expected* payoff across submarkets,

$$\mathbb{E}_x [\tilde{u}(\rho_j, q_j)] = \left[\frac{1 - e^{-q(c_j)}}{q(c_j)} \right] \rho_j + \left[1 - \frac{1 - e^{-q(c_j)}}{q(c_j)} \right] \mathbb{E}_x [\tilde{r}(x, \bar{c})x].$$

The term $\mathbb{E}_x [\tilde{r}(x, \bar{c})x]$ captures the lenders' uncertainty regarding the amount of headroom at the ON RRP facility. Since \bar{c} is determined before x is realized, the probability that the cap is binding is

$$\xi(\bar{c}) \equiv 1 - F \left(\frac{\kappa}{\vartheta(\bar{c})} \right),$$

so that the *expected* effective repo rate can be written

$$\mathbb{E}_x [\tilde{r}(x, \bar{c})x] = (1 - \xi(\bar{c}))r + \xi(\bar{c}) \frac{\kappa}{\vartheta(\bar{c})} r. \quad (20)$$

The DIs' problem remains linear in x , so the analytic expressions for queues and rates apply simply by replacing the effective repo rate \tilde{r} with the *expected* effective repo rate in (20), yielding

$$\tilde{q}(c_j; U) = \log \left(\frac{R - c_j - r + \xi(\bar{c}) \left(1 - \frac{\kappa}{\vartheta(\bar{c})} \right) r}{U - r + \xi(\bar{c}) \left(1 - \frac{\kappa}{\vartheta(\bar{c})} \right) r} \right). \quad (21)$$

The remainder of the equilibrium characterization follows almost exactly the same steps required to characterize equilibrium when the cap is binding. The key difference here is that the lenders' relevant outside option—the expected effective repo rate—depends on the amount of headroom as long as there is some chance that the cap is binding, i.e., $\xi(\bar{c}) > 0$. Indeed, it is now even possible that the average traded rate drops below the repo rate *and* the cap does not bind. It is thus important that the cap is set to leave enough headroom such that lenders are sufficiently confident that the ON RRP facility will be available at the statutory rate of r if they fail to match.

Before concluding, we offer some final thoughts on the effects of an increase in uncertainty over the level of excess cash balances. Clearly the key moment is the tail risk of large excess cash balances, perhaps in response to a “flight to quality” triggered by some financial turbulence. This would increase the probability that the ON RRP cap binds, reduce the expected effective repo rate, and bring downward pressure on traded rates. The potential for entry, however,

would provide a countervailing force, as the drop in rates would attract some additional DIs and lenders would find it easier to match.

8 Conclusion

We have constructed a new model of the federal funds market in the wake of the unconventional policies introduced during the Great Recession. This model captures many relevant features of the current federal funds market, along with the operational framework that the Federal Reserve has chosen to implement monetary policy. In characterizing the equilibrium of this model and studying its properties, our analysis reveals a number of new insights about the features of the economic environment that will determine the FFR, and whether this rate will fall within the FOMC’s target range, i.e., whether monetary policy implementation will be successful. The model is sufficiently parsimonious that the mechanisms underlying our results are transparent and easy to understand, yet flexible enough to fit the data reasonably well. We calibrate the model and use it as a quantitative benchmark for applied analysis, with an emphasis on understanding the role of the overnight reverse repurchase agreement facility in supporting the FFR.

On the whole, our findings are quite positive: Taken together, the IOER and ON RRP facility can be used to keep rates in any target range of choice, which appears to be a top priority. A cap on the ON RRP facility could potentially hinder implementation, though our best estimates suggest that the FFR would only drop out of the target range in extreme circumstances.

Going forward, the FOMC has already stated that it “intends to reduce the Federal Reserve’s securities holdings in a gradual and predictable manner primarily by ceasing to reinvest repayments of principal on securities held” in the Federal Reserve’s balance sheet, and expects to “cease or commence phasing out reinvestments after it begins increasing the target range for the FFR.”³⁶ Given the large size of the Federal Reserve’s balance sheet, we should thus expect the current situation of excess reserves to endure for several years. Eventually, though, the supply of reserves will dwindle and become scarce again. To this end, a natural next step would be to develop a model that can encompass both environments—with excess and scarce reserves—and thus address questions regarding the transition between the two. We leave this for future research.

³⁶ “Policy Normalization Principles and Plans,” Board of Governors of the Federal Reserve System Press Release, September 17, 2014.

A Appendix

A.1 Benchmark Model

Proof of Proposition 1

Let $\Psi(\bar{c})$ denote the left hand side of (9). The result follows immediately from the observations that $\Psi(\bar{c})$ is continuous and strictly increasing over the domain $\bar{c} \in (0, s)$ —since $\Psi'(\bar{c}) = \frac{G(\bar{c})}{s-\bar{c}} > 0$ —with $\lim_{\bar{c} \rightarrow 0} \Psi(\bar{c}) = 0 < \lambda$ and $\lim_{\bar{c} \rightarrow s} \Psi(\bar{c}) > \lambda$ for any finite λ . ■

Proof of Proposition 2.

From (3), we have that

$$\lim_{c \rightarrow \bar{c}} \rho(c_j) = U.$$

Moreover, since $\bar{c} < s$, (12) implies that $U > r$ in any equilibrium. Thus, $\rho(\bar{c}) > r$. Clearly no DI would set $\rho > R$ as it would have strictly negative expected profits by (4). If $\rho(0) = R$ then a DI with zero balance sheet costs has zero profits. Since $q(0) > 0$ by (10), by continuity there exists $\tilde{\rho} < R$ with a strictly positive queue length $\tilde{q} > 0$ and strictly positive expected profits by (4). Hence $\rho(0) \geq R$ is not optimal. ■

Proof of Lemma 3

Since $\frac{\partial \bar{c}}{\partial R} = \mu^d$, it follows that

$$\mu^\ell \frac{\partial \hat{\rho}}{\partial R} = 1 - \mu^d - \frac{\partial \mu^\ell}{\partial R} \frac{s - \bar{c}}{\mu^\ell}. \quad (22)$$

Given the definition of $q(c_j)$, the market clearing condition implies that

$$\int_0^{\bar{c}} \frac{1}{s - c_j} dG(c_j) = \frac{(1 - \mu^d)G(\bar{c})}{s - \bar{c}}. \quad (23)$$

Differentiating μ^ℓ yields

$$\frac{\partial \mu^\ell}{\partial R} = \frac{1}{\lambda} \int_0^{\bar{c}} \left[\frac{\mu^d}{s - c_j} - \frac{\bar{c} - c_j}{(s - c_j)^2} \right] dG(c_j). \quad (24)$$

Plugging (23) into (24) yields

$$\frac{\partial \mu^\ell}{\partial R} = \frac{\mu^d(1 - \mu^d)G(\bar{c})}{\lambda(s - \bar{c})} - \frac{1}{\lambda} \int_0^{\bar{c}} \left[\frac{\bar{c} - c_j}{(s - c_j)^2} \right] dG(c_j). \quad (25)$$

Plugging (25) into (22) yields

$$\begin{aligned}\mu^\ell \frac{\partial \hat{\rho}}{\partial R} &= 1 - \mu^d - \left[\frac{\mu^d(1 - \mu^d)G(\bar{c})}{\lambda(s - \bar{c})} \right] \frac{s - \bar{c}}{\mu^\ell} + \frac{1}{\lambda} \int_0^{\bar{c}} \left[\frac{\bar{c} - c_j}{(s - c_j)^2} \right] dG(c_j) \frac{s - \bar{c}}{\mu^\ell} \\ &= \frac{1}{\lambda} \int_0^{\bar{c}} \left[\frac{\bar{c} - c_j}{(s - c_j)^2} \right] dG(c_j) \frac{s - \bar{c}}{\mu^\ell} > 0\end{aligned}\quad (26)$$

since the first two terms cancel out. Using the definition of $q(c_j)$, (26) is equivalent to (16).

To derive (17), note that (24) can be written

$$\frac{\partial \mu^\ell}{\partial R} = \frac{G(\bar{c})}{\lambda} \int_0^{\bar{c}} \left[\frac{\int_0^{\bar{c}} \frac{\bar{c} - c_i}{(s - c_i)} \frac{dG(c_i)}{G(\bar{c})}}{s - c_j} - \frac{\bar{c} - c_j}{(s - c_j)^2} \right] \frac{dG(c_j)}{G(\bar{c})}.$$

Let $X_j = 1/(s - c_j) \geq 0$, so that $1 - X_j(s - \bar{c}) = (\bar{c} - c_j)/(s - c_j)$. Then

$$\begin{aligned}\frac{\partial \mu^\ell}{\partial R} &= \frac{G(\bar{c})}{\lambda} \{ \mathbb{E}[1 - X_j(s - \bar{c})] \mathbb{E}[X_j] - \mathbb{E}[(1 - X_j(s - \bar{c}))X_j] \} \\ &= \frac{G(\bar{c})(s - \bar{c})}{\lambda} \{ \mathbb{E}[X_j^2] - \mathbb{E}[X_j]^2 \} \geq 0\end{aligned}$$

by Jensen's inequality. ■

A.2 General Matching Functions

Let q denote the ratio of lenders to DIs in a submarket. Let $\mu : q \mapsto [0, 1]$ denote the probability that a DI matches given queue length q . By aggregate consistency, the probability a lender matches in this submarket must be $\mu(q)/q$. As is standard, we assume that $\mu(q)$ is strictly increasing, differentiable, and concave, with $\lim_{q \rightarrow 0} \mu(q) = 0$, $\lim_{q \rightarrow \infty} \mu(q) = 1$, $\lim_{q \rightarrow 0} \mu(q)/q = 1$, and $\lim_{q \rightarrow \infty} \mu(q)/q = 0$.

Given this matching technology, the DIs problem can be written

$$\max_{q_j, \rho_j} \quad \mu(q_j)[R - c_j - \rho_j] \quad (27)$$

$$\text{s.t.} \quad \frac{\mu(q_j)}{q_j}(\rho_j - r) + r = U. \quad (28)$$

Repeating the analysis in the text, it is straightforward to establish that, in any equilibrium, (7) holds, i.e., $\bar{c} = R - U$. Taking first order conditions and letting $s = R - r$, one can show that the queue length in each submarket must satisfy

$$\mu'(q_j) = \frac{s - \bar{c}}{s - c_j} \quad (29)$$

and, given these queue lengths, the cutoff \bar{c} must satisfy

$$\int_0^{\bar{c}} q_j dG(c_j) = \lambda. \quad (30)$$

Rates, then, follow from (28):

$$\rho_j = r + \frac{q\mu'(q)}{\mu(q)}(s - c_j). \quad (31)$$

Given this characterization, one can easily show that Propositions 1 and 2 still hold.

A.3 A Frictionless Market

Consider the equilibrium in our model if the federal funds market were a frictionless, competitive market. The “supply” of funds is simply λ , for any interest rate greater than r , and zero otherwise. The “demand” curve satisfies $Q = G(R - \rho)$, i.e., the measure of DIs with $c \leq R - \rho$.

In Figure 6, we plot the two relevant cases, with ρ^* denoting the unique equilibrium interest rate. In the left panel, we plot the case when there is “excess demand” for deposits by DIs, i.e., when the measure of DIs with balance sheet cost less than s , $G(s)$, is greater than the measure of lenders, λ . In the right panel, we plot the case when there is excess supply, i.e., when $G(s) < \lambda$.

As we noted in the text, this exercise illustrates why a search-based model is helpful for confronting the data. In a Walrasian model, either $\rho^* > r$ but the volume at the ON RRP is zero (as in the first case described above) *or* $\rho^* = r$ and volume at the ON RRP $\lambda - G(s)$ is strictly positive (as in the second case described above). In our search-based model, we have *both* interest rates above r and volume at the ON RRP above zero, as in the data.




Figure 6: Frictionless market

A.4 Characterization of Equilibrium with a Cap

In this section, we show how to characterize equilibrium with a cap κ at the ON RRP facility, given the assumption that each lender is endowed with one unit of cash; it is straightforward to extend the analysis to the case when each lender has x units of cash. In any candidate equilibrium with cutoff \bar{c} and queue lengths $q(c_j)$ for $c_j \leq \bar{c}$, the measure of lenders that do not match with a DI is

$$\int_0^{\bar{c}} \left[1 - \frac{1 - e^{-q(c_j)}}{q(c_j)} \right] q(c_j) dG(c_j).$$

It then follows immediately that the supply of cash at the ON RRP facility is given by

$$\vartheta(\bar{c}) = \int_0^{\bar{c}} \left\{ q(c_j) - \left[1 - e^{-q(c_j)} \right] \right\} dG(c_j). \quad (32)$$

In this candidate equilibrium, the expected utility of a lender in a submarket with interest rate ρ_j and queue length q_j can then be written

$$\tilde{u}(\rho_j, q_j) = \left[\frac{1 - e^{-q(c_j)}}{q(c_j)} \right] \rho_j + \left[1 - \frac{1 - e^{-q(c_j)}}{q(c_j)} \right] \tilde{r}(\bar{c}),$$

where

$$\tilde{r}(\bar{c}) = \min \left\{ 1, \frac{\kappa}{\vartheta(\bar{c})} \right\} r. \quad (33)$$

The characterization of the remaining equilibrium objects follows the analysis in Section 4 very closely. First, lenders are indifferent between all active DIs; i.e., $\tilde{u}(\rho_j, q_j) = U$ for all $c_j \leq \bar{c}$. One can then solve for the profit-maximizing queue length at each DI, which yields

$$\tilde{q}(c_j; U) = \log \left(\frac{R - c_j - \tilde{r}(\bar{c})}{U - \tilde{r}(\bar{c})} \right). \quad (34)$$

Then, imposing $\tilde{q}(\bar{c}; U) = 0$ implies that

$$\bar{c}(U) = R - U, \quad (35)$$

as in the benchmark model, while the optimal interest rate set by each active DI can be written

$$\tilde{\rho}(c_j; U) = \tilde{r}(\bar{c}) + \log \left(\frac{R - c_j - \tilde{r}(\bar{c})}{U - \tilde{r}(\bar{c})} \right) \left[\frac{(R - c_j - \tilde{r}(\bar{c}))(U - \tilde{r}(\bar{c}))}{R - c_j - U} \right]. \quad (36)$$

Finally, the market clearing condition is essentially unchanged:

$$\int_0^{\bar{c}} q(c_j) dG(c_j) = \lambda. \quad (37)$$

An equilibrium, then, is a market utility U^* , a cutoff $\bar{c} = \bar{c}(U^*)$, queue lengths $q(c_j) = \tilde{q}(c_j; U^*)$, and interest rates $\rho(c_j) = \tilde{\rho}(c_j; U^*)$ for all $c_j \leq \bar{c}$, a volume of trade at the ON RRP facility $\vartheta = \vartheta(\bar{c}(U^*))$, and an effective interest rate $\tilde{r} = \tilde{r}(\bar{c}(U^*))$ such that (32)–(37) are satisfied.

A.5 Extensions

Proof of Proposition 4

We proceed by guess-and-verify the cases of market segmentation and pooling. The results are the same if we let the DIs set different rates for different types of lenders.

Market Segmentation. Consider first an equilibrium where safe DIs serve the λ_s lenders exclusively, while not safe DIs serve the remaining λ_n lenders exclusively. In this conjectured equilibrium, a DI of type $i \in \{n, s\}$ solves

$$\begin{aligned} \max \quad & (1 - e^{-q_i})(R - \rho_i) \\ \text{s.t.} \quad & \left(\frac{1 - e^{-q_i}}{q_i}\right)\rho_i + \left(1 - \frac{1 - e^{-q_i}}{q_i}\right)r = U_i. \end{aligned}$$

The solution to this problem is

$$e^{-q_i}(R - r) = U_i - r, \quad (38)$$

yielding the DI profits

$$\Pi_i = (R - r)[1 - e^{q_i}(1 + q_i)] \quad (39)$$

and the type i lenders a payoff

$$U_i = r + (R - r)e^{-q_i}. \quad (40)$$

Aggregate consistency requires

$$q_s = \frac{\lambda_s}{d_s} \quad (41)$$

$$q_n = \frac{\lambda_n}{d_n}. \quad (42)$$

Finally, for this to be an equilibrium, it must be that

$$U_n \geq U_s \quad (43)$$

$$\Pi_s \geq \Pi_n. \quad (44)$$

Both of these are satisfied if, and only if,

$$\frac{\lambda_s}{d_s} \geq \frac{\lambda_n}{d_n}. \quad (45)$$

Pooling. Consider now an equilibrium where type s lenders only go to type s DIs (as they must), but type n lenders mix across the two types of DIs. In particular, let σ denote the probability that a type n lender goes to a type s DI.

Notice immediately that, if type n lenders are mixing, they must be getting equal expected payoff at both types of DIs. Moreover, since every lender in the queue is identical, this means that it $U_s = U_n$ at all DIs. Given the DI's first-order condition, this means that $q_s = q_n$, and hence $\rho_s = \rho_n$. All that needs to be determined is σ , which must satisfy

$$q_s = q_n \Leftrightarrow \frac{\lambda_s + \sigma\lambda_n}{d_s} = \frac{(1 - \sigma)\lambda_n}{d_n} \quad (46)$$

$$\Leftrightarrow \sigma = \frac{\lambda_n d_s - \lambda_s d_n}{\lambda_n d_s + \lambda_n d_n}. \quad (47)$$

Note that $\sigma < 1$ everywhere, but $\sigma \geq 0$ if and only if

$$\frac{\lambda_s}{d_s} \leq \frac{\lambda_n}{d_n}. \quad (48)$$

Proof of Proposition 5

The emergence of separating equilibrium with ex ante heterogeneous agents in directed search models is well known, and hence we only sketch the equilibrium construction here. Since DIs are homogeneous, they must earn equal profits in equilibrium. Let Π denote these profits. Then one can formulate the (dual) problem as lenders of type $i \in \{0, 1\}$ solving

$$\begin{aligned} \max_{\rho_i, q_i} \quad & \left(\frac{1 - e^{-q_i}}{q_i} \right) \rho_i + \left(1 - \frac{1 - e^{-q_i}}{q_i} \right) r_i \\ \text{sub to} \quad & \Pi = (1 - e^{-q_i}) (R - \rho_i), \end{aligned} \quad (49)$$

with $r_0 = 0$ and $r_1 = 1$. The solution to this programming problem is

$$\Pi = (R - r_i) [1 - e^{-q_i} (1 + q_i)] \quad (50)$$

for $i \in \{0, 1\}$, which makes it clear that the two types of agents optimally choose separate submarkets, since $r_0 \neq r_1$ immediately implies that $q_0 \neq q_1$. Combining (50) with the market clearing conditions

$$q_0 = \frac{\lambda_0}{\mu_0} \quad (51)$$

$$q_1 = \frac{\lambda_1}{1 - \mu_0}, \quad (52)$$

equilibrium boils down to a single equation in μ_0 :

$$R \left[1 - e^{-\frac{\lambda_0}{\mu_0}} \left(1 + \frac{\lambda_0}{\mu_0} \right) \right] = (R - r) \left[1 - e^{-\frac{\lambda_1}{(1-\mu_0)}} \left(1 + \frac{\lambda_1}{(1-\mu_0)} \right) \right]. \quad (53)$$

The left-hand side of (53) is continuous and decreasing in μ_0 , converges to R as $\mu_0 \rightarrow 0$, and is equal to

$$\underline{\Pi}_0 = R [1 - e^{-\lambda_0} (1 + \lambda_0)] \quad (54)$$

at $\mu_0 = 1$. The right-hand side of (53) is continuous and increasing in μ_0 , is equal to

$$\underline{\Pi}_1 = (R - r) [1 - e^{-\lambda_1} (1 + \lambda_1)] < R \quad (55)$$

at $\mu_0 = 0$, and converges to $R - r$ as $\mu_0 \rightarrow 1$. Hence, a unique μ_0 exists if, and only if, $\underline{\Pi}_0 < R - r$, or equivalently $r < R e^{-\lambda_0} (1 + \lambda_0)$.

If this condition is violated, then DIs earn a profit from ineligible lenders greater than $R - r$ even when all other DIs target ineligible lenders (i.e., when $\mu_0 = 1$). In this case, $q_0 = \lambda_0$,

$\Pi = \underline{\Pi}_0$, and

$$\rho_0 = \frac{Rq_0 e^{-q_0}}{1 - e^{-q_0}} < r \tag{56}$$

when $r > Re^{-\lambda_0}(1 + \lambda_0)$.

References

- Afonso, Gara, Alex Entz, and Eric LeSeur**, “Who is Lending in the Fed Funds Market?,” Technical Report, Federal Reserve Bank of New York December 2013. Liberty Street Economics.
- , – , **and** – , “Who’s Borrowing in the Fed Funds Market?,” Technical Report, Federal Reserve Bank of New York December 2013. Liberty Street Economics.
- **and Ricardo Lagos**, “Trade Dynamics in the Market for Federal Funds,” *Econometrica*, January 2015, *83* (1), 263–313.
- Akhtar, Mahmūd A**, “Understanding Open Market Operations,” Technical Report, Federal Reserve Bank of New York 1997.
- Armenter, Roc and Benjamin Lester**, “Excess reserves and monetary policy normalization,” Technical Report, Federal Reserve Bank of Philadelphia 2015.
- Ashcraft, Adam B. and Darrell Duffie**, “Systemic Illiquidity in the Federal Funds Market,” *American Economic Review*, May 2007, *97* (2), 221–225.
- Bartolini, Leonardo, Spence Hilton, and Alessandro Prati**, “Money Market Integration,” *Journal of Money, Credit and Banking*, 2008, *40* (1), 193–213.
- Bech, Morten L. and Cyril Monnet**, “A Search-based Model of the Interbank Money Market and Monetary Policy implementation,” December 2014. Unpublished manuscript.
- **and Elizabeth Klee**, “The Mechanics of a Graceful Exit: Interest on Reserves and Segmentation in the Federal Funds Market,” *Journal of Monetary Economics*, 2011, *58* (5), 415–431.
- Berentsen, Aleksander and Cyril Monnet**, “Monetary Policy in a Channel System,” *Journal of Monetary Economics*, September 2008, *55* (6), 1067–1080.
- , **Sebastien Kraenzlin, and Benjamin Muller**, “Exit Strategies and Trade Dynamics in Repo Markets,” Technical Report, Swiss National Bank 2015. Working paper 2015-9.
- Bianchi, Javier and Saki Bigio**, “Banks, Liquidity Management, and Monetary Policy,” September 2014. Federal Reserve Bank of Minneapolis, Research Department Staff Report 503.
- Burdett, Kenneth, Shouyong Shi, and Randall Wright**, “Pricing and Matching with Frictions,” *Journal of Political Economy*, 2001, *109* (5), 1060–1085.
- Cipriani, Marco and Jonathan Cohn**, “The FR 2420 Data Collection: A New Base for the Fed Funds Rate,” Technical Report, Federal Reserve Bank of New York April 2015. Liberty Street Economics.
- **and Julia Gouny**, “The Eurodollar Market in the United States,” Technical Report, Federal Reserve Bank of New York May 2015. Liberty Street Economics.

- Demiralp, Selva, Brian Preslopsky, and William Whitesell**, “Overnight Interbank Loan Markets,” *Journal of Economics and Business*, 2006, 58 (1), 67 – 83.
- Ennis, Huberto M.**, “A Simple General Equilibrium Model of Large Excess Reserves,” July 2014. Federal Reserve Bank of Richmond, Working paper WP 14-14.
- Ennis, Huberto M. and John A. Weinberg**, “Over-the-counter Loans, Adverse Selection, and Stigma in the Interbank Market,” *Review of Economic Dynamics*, 2013, 16 (4), 601 – 616.
- **and Todd Keister**, “Understanding Monetary Policy Implementation,” *Economic Quarterly*, 2008, 94 (3), 235–263. Federal Reserve Bank of Richmond.
- Frost, Joshua, Lorie Logan, Antoine Martin, Patrick McCabe, Fabio Natalucci, and Julie Rémache**, “Overnight RRP Operations as a Monetary Policy Tool: Some Design Considerations,” Staff Reports 712, Federal Reserve Bank of New York February 2015.
- Furfine, Craig H.**, “The Microstructure of the Federal Funds Market,” *Financial Markets, Institutions & Instruments*, 1999, 8 (5), 24–44.
- Hamilton, James D.**, “The Daily Market for Federal Funds,” *Journal of Political Economy*, February 1996, 104 (1), 26–56.
- Ho, Thomas S.Y. and Anthony Saunders**, “A Micro Model of the Federal Funds Market,” *The Journal of Finance*, 1985, 40 (3), 977–988.
- Hosios, Arthur J.**, “On the Efficiency of Matching and Related Models of Search and Unemployment,” *The Review of Economic Studies*, 1990, 57 (2), 279–298.
- Lester, Benjamin, Guillaume Rocheteau, and Pierre-Olivier Weill**, “Competing for Order Flow in OTC Markets,” *Journal of Money, Credit and Banking*, 2015, 47 (S2), 77–126.
- Martin, Antoine, Jaime McAndrews, and David Skeie**, “Bank Lending in Times of Large Bank Reserves,” June 2013. Federal Reserve Bank of New York Staff Reports, n. 497.
- **, James McAndrews, Ali Palida, and David R Skeie**, “Federal Reserve Tools for Managing Rates and Reserves,” September 2013. Federal Reserve Bank of New York Staff Reports, no. 642.
- Moén, Espen**, “Competitive Search Equilibrium,” *Journal of Political Economy*, 1997, 105, 385–411.
- Poole, William**, “Commercial Bank Reserve Management in a Stochastic Model: Implications for Monetary Policy,” *Journal of Finance*, 1968, 23 (5), 769–791.
- Potter, Simon**, “Recent Developments in Monetary Policy Implementation,” December 2013. Speech, Federal Reserve Bank of New York.
- **, “Interest Rate Control during Normalization,”** 2014. Speech, Federal Reserve Bank of New York.

Whitesell, William, “Interest Rate Corridors and Reserves,” *Journal of Monetary Economics*, September 2006, 53 (6), 1177–1195.

Williamson, Stephen, “Interest on Reserves, Interbank Lending, and Monetary Policy,” August 2015. Working paper, Federal Reserve Bank of Saint Louis.