

Banks' Reserve Management, Transaction Costs, and the Timing of Federal Reserve Intervention

Leonardo Bartolini,^a Giuseppe Bertola,^b and Alessandro Prati^c

^aFederal Reserve Bank of New York (New York, NY 10045, USA)

^bEuropean University Institute (S. Domenico di Fiesole, I-50016, Italy), U. di Torino, and CEPR

^cInternational Monetary Fund (Washington, DC 20431, USA)

February 2000

Abstract *

We use daily data on bank reserves and overnight interest rates to document a striking pattern in the high-frequency behavior of the U.S. market for federal funds: depository institutions tend to hold more reserves during the last few days of each “reserve maintenance period,” when the opportunity cost of holding reserves is typically highest. We then propose and analyze a model of the federal funds market where uncertain liquidity flows and transaction costs induce banks to delay trading and to bid up interest rates at the end of each maintenance period. In this context, the central bank’s interest-rate-smoothing policy causes a high supply of liquid funds to be associated with high interest rates around reserve settlement days.

JEL G21, E52.

Keywords: reserve requirements, interest-rate smoothing, value of waiting.

Address: Giuseppe Bertola, Dept. of Economics, IUE, I-50016 San Domenico di Fiesole FI
x39-055-4685225, bertola@datacomm.iue.it

* We thank an anonymous referee and seminar participants for helpful comments; S. Hilton, S. Krieger, J. Partlan, and E. Spinner, from the Federal Reserve’s Open Market Desk, for useful discussion and for granting us access to the Federal Reserve Bank of New York’s reserves data archive; J. Clouse, from the Federal Reserve Board, for sharing with us details of his estimates of banks’ reserve demand; and L. Brookins, W. Koeniger, and J. Spataro for research and editorial assistance. G. Bertola thanks the Research Council of the European University Institute for financial support. L. Bartolini thanks the European University Institute for hospitality and financial support. This paper does not necessarily reflect the views of the IMF or of the Federal Reserve System.

1. Introduction

The high frequency behavior of the interbank market for reserves has been intensely scrutinized by research on banking in recent years. One of the goals of this research has been to identify patterns displayed by money-market data, mostly from the U.S., and to examine their consistency with the assumption of rational, profit-maximizing behavior of depository institutions. Our paper contributes to this research by documenting and modeling a remarkable and previously unnoticed pattern in the U.S. market for federal funds: the tendency of U.S. banks to hold--and of the Federal Reserve to provide--the bulk of their reserves during the last days of each “reserve maintenance period” (the bi-weekly period over which banks’ reserves are averaged for the purpose of reserve requirements), when the opportunity cost of holding reserves is typically highest. After documenting this pattern empirically, we rationalize it in the context of a model where banks’ policy of liquidity management interacts with the central bank’s interest-targeting policy, and where transaction costs and liquidity shocks cause both interest rates and bank reserves to rise endogenously at the end of each maintenance period.

The main ingredients of our analysis are banks’ uncertainty on their reserve needs when trading reserves is costly, and the Federal Reserve’s interest-rate-smoothing policy. We show that small costs of trading federal funds and uncertainty on reserve flows induce banks to trade *late* in each maintenance period, so as to conduct federal funds transactions when they have more precise information on their eventual reserve needs. This policy allows individual banks to offset higher reserve-holding costs around period-end by saving on overall transaction costs. The aggregate effect of this policy would only be to cause interest rates to rise at period-end if the supply of reserves remained constant through the maintenance period. If, however, the central bank’s goal is

to control the supply of reserves so as to limit fluctuations of market rates around their target, in equilibrium both interest rates and reserves should rise cyclically around period-end, in accordance with the clear empirical patterns of U.S. federal funds data.

Our modeling perspective is similar to that of several recent studies of interbank markets, including Kopecky and Tucker (1993), Hamilton (1996), and Clouse and Dow (1999), who first pointed to the importance of trading and other fixed costs in the federal funds market. We build on these studies to develop a more complete account of the role of trading costs in explaining the joint behavior of interest rates and reserves. For instance, we extend Kopecky and Tucker's (1993) work by assuming that trading costs are incurred every day (Kopecky and Tucker assume transactions to be costly only on settlement days), so as to be able to genuinely endogenize differences in behavior between days. Conversely, we strip Hamilton's (1996) model of a number of accessory assumptions (such as heterogeneous bank behavior, tighter restrictions on banks' ability to sell rather than buy federal funds, week-end effects, etc.) which blur the analysis of trading costs and lead Hamilton's simulated settlement-day rates below typical non-settlement-day rates, in contrast with U.S. federal funds data. Our emphasis on the tendency of interest rates to rise on settlement days is also similar to that of Clouse and Dow (1999). Clouse and Dow's study, however, hinges on the role of fixed costs incurred when borrowing from the discount window. While analytically interesting, this perspective is empirically more problematic, since discount-window lending has decreased dramatically in the U.S. since the banking crisis of the mid-1980s, while the stylized fact we seek to explain has survived unscathed.¹ More fundamentally, all these studies feature an exceedingly stylized description of Fed intervention. Hamilton (1996), for

¹ See Clouse (1994) and Peristiani (1998) for a discussion.

instance, abstracts from official intervention altogether, assuming aggregate reserves to remain constant over time. Kopecky and Tucker (1993) and Clouse and Dow (1999) let the Fed issue reserves randomly, irrespective of their impact on interest rates. As a result, these studies--and others in the literature--cannot account for the cyclical reserve patterns we identify in U.S. data.

Our paper is structured as follows. Section 2 reviews the main evidence on the high-frequency behavior of U.S. federal funds rates and presents new evidence on the behavior of bank reserves. Section 3 presents a model of the federal funds market, a simplified version of which is solved in Section 4. This simple version of the model yields a number of analytical results and intuition for the more general results that follow. Section 5 solves the more general model in parameterized form, roughly calibrated on U.S. data. Section 6 concludes.

2. The High-Frequency Behavior of Interest Rates and Reserves

U.S. depository institutions must satisfy reserve requirements defined on a period-average basis: both actual reserves and required reserves are computed as averages of daily values over two-week periods (i.e., the “maintenance period” for actual reserves, and an earlier “computation period” for required reserves, each lasting two weeks).² The cyclical character of reserve requirements, along with the no-arbitrage condition that derives from banks’ attempts to minimize the cost of meeting requirements (the so-called martingale hypothesis of federal funds rates)³ has

² We refer to institutions subject to reserve requirements simply as “banks,” although, strictly speaking, these requirements apply (in the United States as elsewhere) to a larger class of depository institutions. See the Appendix for a summary of U.S. reserve requirements.

³ In short: If, at the margin, reserves are held for the purpose of satisfying reserve requirements, overnight interbank rates should be expected to stay constant (aside from negligible discounting) based on information available daily, that is, they should display no predictable within-period

been the target of scrutiny since the U.S. reserve system was reformed in the 1980s. Generally, the martingale hypothesis has been found to fail, the most flagrant violation being the tendency of federal funds rates to rise predictably at the end of each maintenance period (see Campbell, 1987, Hamilton, 1996, and Balduzzi et al., 1998, among others). Relevant evidence is summarized in the top two panels of Figures 1-3, which plot the average and median differences between the federal funds rate prevailing on each day of the maintenance period and the rate prevailing on the first day of each period.⁴ The figures show that average federal funds rates are far from constant within a typical maintenance period. In particular, they tend to rise markedly over the last three days of each period, and the rise recorded on day ten (“settlement day”) is so sharp that the typical value of the rate for this day lies well above the typical value for the other nine (“non-settlement”) days. As discussed by Hamilton (1996), it is important to consider outlier-robust statistics, since an overwhelming proportion of the variance of federal funds rates since the 1980s is accounted for by very few observations. The pattern of average federal funds rates within maintenance periods, however, does not reflect the presence of outliers: as Figures 1-3 show, the mean and the (outlier-

pattern. This is because predictable rate changes would induce banks to cluster demand for reserves in low-rate days, and vice versa, so as to minimize the cost of meeting the requirements, until--in equilibrium--expected rates are brought into equality within each maintenance period.

⁴ We obtained our data from the Federal Reserve Bank of New York. Interest rate data are daily effective federal funds rates, computed as transaction-weighted rates on unsecured overnight loans arranged by the primary funds dealers (currently numbering 39). These are the rates monitored by the Fed for intervention purposes, and represent 1/3 of all transactions taking place in the federal funds market. Reserves data include end-of-day total actual reserves (applied vault cash, plus deposits held at the Fed by institutions subject to reserve requirements) and required reserves. The sample runs from January 1, 1986 (the inception of the Fed’s archive), to July 8, 1998 (the month when reserve requirements were reformed by lagging the maintenance period by four additional weeks with respect to the computation period--see the Appendix for details). To our knowledge, the only prior study using daily reserve data (from a much smaller sample than ours) is Hamilton (1996), which however focuses on issues of monetary transmission rather than of market efficiency, and does not identify cyclical patterns in reserves.

robust) median statistics display a very similar behavior. Furthermore, the pattern cannot be ascribed to the behavior of the market, or of the Fed, in specific sub-periods. To document this claim, we split the full sample, analyzed in Figure 1, into two sub-samples: one running from January 2, 1986, to January 5, 1994; the other running from January 6, 1994, to July 1, 1998. This particular split is motivated by the Fed's switch, in January 1994, to a procedure whereby changes in target rates have been publicly announced and implemented mainly at times of FOMC meetings, a regime change which Bartolini et al. (1998) show to have affected patterns of interest rate *volatility*. As is apparent from the figures, the patterns described above are robust to sample choice, and remain statistically significant in both the pre-1994 and post-1994 sub-samples. Other sample splits that we considered yielded qualitatively identical results.⁵

Besides the rates' sharp rise at period-end, Figure 1 shows a pattern of slightly falling rates over the first days of each period which, as discussed in Section 6, may reflect carry-over of high rates from settlement day to the early days of the following period. The data also display other patterns (for instance, the rate tends to fall slightly on Fridays and rise on Mondays), well known to reflect accounting conventions (e.g., the triple counting of reserves on Fridays; Hamilton, 1996). It will be useful in what follows to abstract from such details, and compare the behavior of interest rates on settlement days with their average behavior on the remaining nine days. The evidence of high settlement-day rates is even more transparent from this perspective. Table 1 shows that mean settlement-day rates exceed those for non-settlement days by 18 basis points. The mean-based test of equality of interest rates between these two groups of days strongly rejects the null of equality.

⁵ The patterns are also independent of our use of volume-weighted rates, which better represent overall market behavior: the cyclical behavior of end-of-day rates is qualitatively identical (see, for instance, Balduzzi et al., 1997).

The data also display periodic patterns of interest rate volatility, reflecting banks' last-minute scramble to build or off-load reserves in the face of unforeseen aggregate liquidity shocks. Such volatility patterns are indirectly relevant to the of the *mean* patterns on which we focus, since evidence of higher settlement-day rates could in principle be spuriously generated by outliers. In Table 1, standard deviation and (outlier-robust) mean absolute median difference statistics indicate that the volatility of interest rates is indeed significantly higher on settlement days than in the rest of each maintenance period.⁶ The evidence is similar across all the subsamples we investigated, and consistent with previous findings (Spindt and Hoffmeister, 1988, Griffiths and Winters, 1995, and Bartolini et al., 1999). We also explored the settlement-vs.-non-settlement behavior of the rates' skewness. The distribution of settlement and non-settlement rates appears to be positively skewed in our full sample, and more markedly so on non-settlement than settlement days (see Table 1). Measures of skewness tend to be very unstable, however (see, for instance, Mood, Graybill, and Boes, 1974, pp. 75-76). Our analysis confirmed this feature, as neither the ranking of settlement vs. non-settlement skewness, nor the sign of the skewness itself, proved robust to sample choice.

The stylized facts we focus on, however, do not depend on the shape of the distribution of interest rates. The non-parametric outlier-robust statistics reported in Table 1 indicate that the gap between settlement- and non-settlement-day rates is a strongly significant and remarkably robust feature of the data. *Median* settlement-day rates exceed those for non-settlement days by 12 basis points, and median-based tests of equality of interest rates between these two groups of days

⁶ The F-test and the Bartlett test are standard. The outlier-robust Brown-Forsythe test is based on an analysis of variance of the absolute median difference.

strongly reject the null of equality of settlement and non-settlement rates.⁷ Non-parametric tests based on the frequency of rate increases on settlement days (occurring in 70 percent of the periods) and non-settlement days (occurring in 41 percent of the periods) yield similar results. In sum, the main testable prediction of simple models of reserve management under periodic requirements is clearly violated: federal funds rates have a predictably higher mean on settlement days; hence, they do not behave as martingales within maintenance periods.

Models that reconcile this evidence with profit-maximizing behavior usually allow for market frictions such as transaction costs, credit lines, bid-ask spreads, etc., or recognize that reserves offer liquidity-provision benefits, besides satisfying reserve requirements. A common approach to modeling this liquidity role of reserves (see, for instance, Campbell, 1987) posits that banks incur increasing costs when liquidity deviates from a target level (determined by the bank's relationship with its nonbank customers, its role in the payment system, and the need to secure positive balances to avoid end-of-day overdraft penalties). Postulating a tractable quadratic form for such liquidity yields, and including foregone profits from interbank lending, a bank's daily cost of holding reserves is:

$$l(x_t, r_t) = x_t r_t + (\alpha/2)(x_t - T)^2, \quad (1)$$

where x_t denotes the level and r_t the opportunity cost of reserve holdings; α indexes the strength

⁷ The Wilcoxon-Mann-Whitney statistic tests the equality of the sum of ranks from two subgroups (here: settlement and non-settlement days) with the same median under the null. The median chi-square test is a rank-based ANOVA test comparing the number of observations in each subgroup falling above and below the overall median. This test's adjusted version corrects for continuity. The Kruskal-Wallis test is a rank-based, one-way ANOVA which extends the Mann-Whitney test. The van der Waerden test is analogous to the Kruskal-Wallis test, although it converts ranks into normal quantiles.

of liquidity concerns; and T denotes the target level of liquidity. Absent other constraints, such as reserve requirements, optimal reserves should satisfy the first-order condition associated with minimizing (1), namely, $x_t = T - r_t / \alpha$. By this condition, as long as $\alpha > 0$ banks should hold relatively fewer reserves on days with predictably high interest rates, and vice versa.

From this perspective, the pattern of excess reserves (i.e., reserves in excess of bi-weekly requirements) documented in the bottom panels of Figures 1-3 and in Table 2 is remarkable. Far from economizing on reserve-carrying costs by holding more reserves when their opportunity cost is low, banks hold more reserves around end-period, when interest rates are typically highest: mean and median excess reserves decline slightly from the first to the sixth day of each period, and rise sharply thereafter, as banks tend to accumulate the bulk of their excess reserves in the last two days of each maintenance period, especially on settlement day.⁸ The market's liquidity, like the federal funds rate, features a slightly declining pattern over the first few days of the maintenance period (with oscillations attributable to day-of-the-week effects), and a sharp increase thereafter. It is particularly insightful to focus again on the contrast between non-settlement and settlement days. While excess reserves equal 2.4 percent of required reserves (or about \$1.3 billion) over a typical whole period, they equal about 7 percent of required reserves on a typical settlement day, but less than 2 percent on non-settlement days. All tests reported in Table 2 reject the null of equality between settlement and non-settlement days, with non-parametric statistics showing that excess reserves--similarly to interest rates--rise on 70 percent of settlement days.

An interesting issue, relevant for the analysis that follows, is that of whether such reserve

⁸ Hence this pattern contrasts with Hamilton's (1996) conjecture that banks tend to build reserves early in each period, as they fear being unable to fill their reserve needs before period-end.

patterns reflect the cyclical behavior of reserves provided to banks by the non-bank sector, or that of reserves supplied by the Fed. Figure 4 and Tables 3-4 settle this issue by documenting the behavior of excess reserves when broken down into their “non-bank” and “official” components. The top two panels of Figure 4 show mean and median reserves injected by the Fed each day of the maintenance period, measured as the change in securities held by the Fed (under outright purchase or under repo). In the first nine days the Fed is a net seller of securities, and thus is withdrawing liquidity from the federal funds market. On day ten, instead, the Fed typically increases market liquidity by purchasing securities. The Fed’s stock of securities rises on 67 percent of settlement days, against 51 percent of non-settlement days (Table 3).⁹

By contrast, the bottom panels of Figure 4 show that excess reserves *net* of Fed intervention display only a mild, statistically insignificant U-shaped pattern within the maintenance period. This is not surprising, since one would not expect non-Fed sources of liquidity such as Treasury payments, changes in household and corporate deposits, etc., to be related to the bi-weekly maintenance cycle. (It is also unsurprising to find that non-Fed liquidity displays some day-of-the-week effects.) Table 4 formally confirms this feature: tests of equality between mean and median excess reserves net of Fed intervention indicate insignificant difference between settlement- and non-settlement- day rates. Only non-parametric tests show net reserves to rise slightly more frequently (59 percent of the time vs. 48 percent of the time) on settlement than on non-settlement

⁹ Note that the Fed’s stock of securities rises on about half (51 percent) of non-settlement days, although the Fed is a net *seller* of securities on these days. Thus, on these days the Fed disposes of its securities (thereby draining liquidity from the market) by increasing the size, rather than the frequency, of its reverse-repo operations. This pattern is not statistically significant and plays no role in our subsequent model, which hinges only on average changes in security holdings, and not on the choice between size and frequency of operations.

days, but even this feature is significantly less marked than for gross reserves.¹⁰

To summarize, this section documents strong evidence of two stylized facts which the analysis of the next sections is called to explain. First, there is a strong positive correlation between the daily behavior of prices (interest rates) and quantities (reserves) in the U.S. market for federal funds, which indicates that endogenous patterns in demand dominate the behavior of this market at this frequency. Second, the cyclical evolution of reserves reflects essentially the behavior of Fed intervention, which therefore must be explicitly accounted for, if one's goal is to provide a reasonably realistic model of the U.S. interbank market for reserves.

3. The Operating Environment of Depository Institutions

We now present a stylized model of the U.S. market for federal funds to explain the correlation between prices and quantities across settlement and non-settlement days documented above. We cast the model in discrete time by collapsing all transactions a bank undertakes in the course of a day into (at most) a single daily transaction, after which the bank's position is tallied for the purpose of reserves computation. Let x_t denote the reserves a bank holds at the end of day t , and let r_t denote the interest charged on inter-bank loans extended overnight from day t to day $t+1$. The opportunity cost of holding x_t overnight is then $r_t x_t$. Banks also incur a (quadratic) cost when their liquidity deviates from a target level T . Thus, as in Equation (1) and in related studies, a bank's daily cost of holding x_t is $l(x_t, r_t) = x_t r_t + (\alpha/2)(x_t - T)^2$.

Every day, banks receive net overnight deposits from their non-bank customers, denoted

¹⁰ The evidence presented in Figure 4 is also robust to splitting the sample into pre- and post-1994 periods, as well as into alternative sub-samples. Details are available upon request.

ϵ_t . If a bank does not trade in the interbank market, its reserve position at the end of day t is ϵ_t . For simplicity, we assume *required* reserves to remain constant, so that ϵ_t describes a shock to both *actual* and *excess* reserves, the latter being the only relevant reserve variable in our model. Also for simplicity, we let the distribution of ϵ_t be continuous on an infinite support (to imply that there is always a possibility--however remote--of a reserve shortfall), and we assume it to be distributed identically and independently over time. To capture in the simplest possible way the different behavior of the market between the first nine (non-settlement) and the last (settlement) days of each period, we consider a sequence of maintenance periods, each lasting only two days, referring to day one of each period as “non-settlement day” and to day two as “settlement day.” Our theoretical results are qualitatively valid under more general conditions than those just outlined, however. In particular, in Section 6 we briefly discuss how more complex models might explain interest rate and liquidity patterns over the ten days of real-life maintenance periods.

To model the aggregate market cleared by the average federal funds rate, we let ϵ_t be uncorrelated across banks and assume banks to be of infinitesimal size, with unit total mass, so that aggregate reserves originating from the non-bank sector, $\mu \equiv E[\epsilon_t]$, are not random. Disregarding discounting, whose impact is negligible over such short horizons, the banks’ problem is to minimize the expected cost of holding reserves, while keeping average reserves over the period at least equal to a required level, denoted \bar{a} in what follows.

Each day, after the realization of ϵ_t , banks may trade federal funds, borrowing or lending overnight, to achieve an end-day reserve position different from ϵ_t . Whenever a bank trades funds, so that $x_t \neq \epsilon_t$, it must pay a fixed cost which, to simplify tracking of reserve accounts, is assumed payable by means other than cash or Fed deposits. We denote this transaction cost with k

and interpret it broadly, to include a variety of fixed costs associated with trading federal funds, such as brokerage and Chips or Fedwire transaction fees, as well as--and more importantly--the costs of searching for banks with matching liquidity needs and those resulting from the need to split large transactions into many small ones to work around credit lines, minimize exposure to interest rate fluctuations, etc. (see Kopecky and Tucker, 1993, Hamilton, 1996, and Hancock et al., 1999, for a discussion). Our results show that such costs allow our model to replicate the patterns observed in the U.S. federal funds market. In any case, the most revealing sign that transaction costs play an important role in this market is that banks do hold excess reserves in equilibrium, foregoing the option of selling them on the market--at positive returns--at period's end.

We assume penalties on reserve shortfalls to be prohibitively large, abstracting from banks' ability to roll-over a fraction of their reserve imbalances to future periods, or to borrow reserves from the Fed at penalty rates. Finally, although we ignore non-cyclical regulations such as daily overdraft penalties and the requirement for banks to maintain a fixed deposit at the Fed as a "clearing balance," we capture their effects on banks' desire to hold reserves irrespective of periodic requirements by our assumption (1) that banks target a daily reserve balance T .

4. The Timing of Optimal Reserve Accumulation

4.1. A typical bank's problem. We solve our model beginning with the liquidity-management policy of an isolated, price-taking bank, facing reserve requirements, stochastic liquidity shocks, and transaction costs. We then build on this analysis to study the equilibrium of the banking sector as a whole, the behavior of the Fed, and the dynamics of aggregate reserves.

The individual bank we consider must choose, on each day t , whether and how much to

trade in the federal funds market. Since our problem is stationary *across* maintenance periods--though not *within* periods--we can drop the subscript t from our notation, while keeping the distinction between settlement and non-settlement days. Thus, we denote the interest rate by $r_t = r_N$ if t falls on a non-settlement day, and by $r_t = r_S$ if t falls on a settlement day. Recalling that the bank's reserves must average at least \bar{a} over the period, its intertemporal problem is:

$$V_N(\epsilon_N) = \min_{x_N} \left[\left(x_N r_N + (\alpha/2)(x_N - T)^2 + k \cdot 1_{[x_N \neq \epsilon_N]} \right) + E \left[V_S(\epsilon_S, x_N) \right] \right], \quad (2a)$$

$$V_S(\epsilon_S, x_N) = \min_{x_S} \left[\left(x_S r_S + (\alpha/2)(x_S - T)^2 + k \cdot 1_{[x_S \neq \epsilon_S]} \right), \text{ s.t. } x_N + x_S \geq 2\bar{a} \right], \quad (2b)$$

where the expectation in (2a) is taken over the distribution of ϵ_S , while the bank faces no residual uncertainty at trading time on settlement day, since by then ϵ_S is already known. In (2a-b), the indicator function $1_{[\cdot]}$ takes a value of one when the event in square brackets is true, and zero otherwise; and the bank must satisfy reserve requirements on every settlement day.

To develop intuition for the economics of our problem, in this Section we analyze the special case where reserves do not provide liquidity benefits (i.e., we let $\alpha=0$), so that banks hold reserves only to satisfy reserve requirements. The next Section considers the more general case, which preserves much of the intuition provided by the simpler case.

4.2. Settlement-day behavior. Given the reserve position inherited from non-settlement day, x_N , and the day's liquidity shock, ϵ_S , on settlement day the bank follows a simple policy. If $x_N + \epsilon_S < 2\bar{a}$, it must borrow at least $2\bar{a} - x_N - \epsilon_S$ to meet reserve requirements. Since excess reserves have no value, the bank will borrow exactly this amount. If $x_N + \epsilon_S \geq 2\bar{a}$, instead, the bank may choose either to trade or to remain idle. If the bank decides to trade, it will still find it optimal to meet the requirements exactly, rather than carry unremunerated reserves on its books. In this case,

the bank's net borrowing $2\bar{a} - x_N - \epsilon_S$ is negative, reflecting its decision to *lend* funds overnight.

When a bank finds itself with excess reserves (i.e., $x_N + \epsilon_S > 2\bar{a}$), however, it may prefer not to trade. This is the case if the interest flow it would earn from a loan does not compensate it for the cost k , that is, if $-(2\bar{a} - x_N - \epsilon_S)r_S \leq k$ or $\epsilon_S \leq 2\bar{a} - x_N + k/r_S$. (We break ties in favor of no trading.) Thus, a bank facing trading costs will hold excess reserves on average, as it never has a reserve shortfall but sometimes holds excess reserves. Formally:

Result 1: If $k > 0$, i.e., if transactions are costly, then an individual bank's excess reserves are positive over the maintenance period, on average.

Proof: On day S , the bank borrows $2\bar{a} - x_N - \epsilon_S$ if either $\epsilon_S < 2\bar{a} - x_N$ or $\epsilon_S > 2\bar{a} - x_N + k/r_S$, bringing its daily reserves to $2\bar{a} - x_N$ and its average reserves to \bar{a} . If, instead, $2\bar{a} - x_N \leq \epsilon_S \leq 2\bar{a} - x_N + k/r_S$, the bank remains idle on day S , and its day S reserves are $\epsilon_S \geq 2\bar{a} - x_N$. Hence, average reserves over the period weakly exceed \bar{a} . They do so strictly if and only if $\text{Prob}(2\bar{a} - x_N \leq \epsilon_S \leq 2\bar{a} - x_N + k/r_S) > 0$ and $E[\epsilon_S + x_N - 2\bar{a} | 2\bar{a} - x_N \leq \epsilon_S \leq 2\bar{a} - x_N + k/r_S] > 0$, which can be true only if $k/r_S > 0$. \square

4.3. Non-settlement-day behavior. Given the information available at trading time on day N , the bank can compute the expected loss for settlement day, $E[V_S(\epsilon_S, x_N)]$, as a function of x_N as:

$$\begin{aligned} E[V_S(\epsilon_S, x_N)] &= E\left[\min_{x_S} \left(x_S r_S + k \cdot 1_{[x_S \neq \epsilon_S]}\right), \text{ s.t. } x_N + x_S \geq 2\bar{a} \mid x_N\right] \\ &= \left(\text{Prob}(\epsilon_S < 2\bar{a} - x_N) + \text{Prob}(\epsilon_S > 2\bar{a} - x_N + k/r_S)\right) \left[k + (2\bar{a} - x_N)r_S\right] \\ &\quad + \text{Prob}(2\bar{a} - x_N \leq \epsilon_S \leq 2\bar{a} - x_N + k/r_S) E[\epsilon_S r_S \mid 2\bar{a} - x_N \leq \epsilon_S \leq 2\bar{a} - x_N + k/r_S] \\ &= \left(F(2\bar{a} - x_N) + 1 - F(2\bar{a} - x_N + k/r_S)\right) \left(k + (2\bar{a} - x_N)r_S\right) + r_S \int_{2\bar{a} - x_N}^{2\bar{a} - x_N + k/r_S} \epsilon \, dF(\epsilon) \end{aligned}$$

$$\equiv \tilde{V}_S(x_N), \quad (3)$$

Substituting (3) into (2a) allows us to characterize the bank's optimal policy on non-settlement day, and highlights that a fixed transaction cost does not impair the martingale property of federal funds rates (which, in our context, requires the rate to remain constant over each maintenance period). Intuitively, since banks would react to a predictable pattern in interest rates within a maintenance period by opening unbounded positions to generate infinite profits, a finite transaction cost cannot prevent a bank from arbitraging even minuscule differences in rates between days. Formally, this property becomes apparent by considering, for example, the case $r_N > r_S$. In this case, a bank would try to lend an unboundedly large amount of reserves on day N ($x_N \rightarrow -\infty$), and borrow a corresponding amount of reserves (plus $2\bar{a}$) on day S ($x_S = 2\bar{a} - x_N \rightarrow \infty$). As the terms involving cumulative distributions in (3) go to zero, the firm's expected non-settlement loss, $V_N(\epsilon_N)$, goes to minus infinity with $x_N(r_N - r_S) + 2k + 2\bar{a}r_S$ (since $(r_N - r_S) > 0$), as the term $x_N(r_N - r_S)$ swamps any finite transaction cost. Conversely if $r_S > r_N$. Hence:

Result 2: When $\alpha=0$, and irrespective of whether $k>0$ or $k=0$, a bank's optimization problem has a bounded solution only if interest rates are constant: $r_S = r_N$.

As long as the only role of reserves is to satisfy reserve requirements, as assumed in this section, (lump-sum) transaction costs do not relax the required equality of settlement and non-settlement-day rates.¹¹ Transaction costs, however, do affect banks' optimal reserve holdings. If

¹¹ Strict equality of interest rates reflects our assumption of no discounting within maintenance periods. With discounting, the exact equilibrium condition is $r_N = r_S / (1 + r_N)$.

$k=0$, a bank is indifferent as to the timing of its reserve holdings: its total cost over the period is $r_S 2\bar{a} = r_N 2\bar{a}$, irrespective of when reserves are held. If $k>0$, instead, the bank's position on day N affects its trading on day S , as captured by the dependence of day S 's inaction range on the reserves inherited from day N . The bank is then no longer indifferent to when it holds reserves, even with constant interest rates. In particular, it benefits from waiting for information on eventual reserve needs to arrive on day S , so as to economize on the transaction costs. Formally:

Result 3: If reserves have no liquidity-provision role ($\alpha=0$) and the interest rate is constant over the maintenance period ($r_N = r_S$), then a bank never trades on non-settlement day: $x_N = \epsilon_N$.

Proof: Consider a candidate policy whereby the bank trades on day N with positive probability, and compare its cost with that of an alternative policy which postpones all the candidate policy's transactions for day N to day S . The two policies feature identical average reserves over the period, hence both satisfy reserve requirements if the candidate policy does. The opportunity cost of funds is also the same for both policies, since the same rate applies to funds traded on either day. Total transaction costs are the same across the two policies along all paths where the candidate policy involves no trade on either day N or on day S , but the alternative policy saves k along paths where the candidate policy requires trading on both days. Thus, the alternative policy dominates the candidate policy at least weakly (and strictly if the probability of trading on both days is positive, as is the case if the support of the distribution of ϵ_t is unbounded below). Hence, optimal behavior entails inaction on day N . \square

Recalling that equality of interest rates is necessary for this section's model to yield a well-

defined solution, Result 3 completes the description of the bank's optimal policy: on day N , the bank lets its reserve account fully absorb the liquidity shock, holding $x_N = \epsilon_N$ at day's end; on day S the bank brings reserves exactly in line with requirements ($x_S = 2\bar{a} - x_N = 2\bar{a} - \epsilon_N$) if either $\epsilon_S < 2\bar{a} - x_N$ or $\epsilon_S > 2\bar{a} - x_N + k/r_S$, and holds that day's random liquidity ($x_S = \epsilon_S$) otherwise.

4.4. Reserve patterns. We can now compare the expected level of reserves on settlement day to that for non-settlement day, $E[x_N] = E[\epsilon_N] = \int_{-\infty}^{\infty} \epsilon dF(\epsilon)$. We shall do so under the regularity assumption that $\bar{a} \geq E[\epsilon]$, i.e., that the banking sector as a whole is required to hold reserves at least weakly larger than exogenous flows would imply. Clearly, this calls for the Fed to stand ready to provide reserves as needed by the banking sector. Based on our previous analysis, we then show that expected reserve holdings on settlement days exceed those on non-settlement days:

Result 4: If $\bar{a} \geq E[\epsilon]$, $k/r_S > 0$, and $\alpha = 0$, then $E[x_S] > E[x_N]$.

$$\mathbf{Proof:} \quad E[x_S] = E[x_S | x_S \neq \epsilon_S] \cdot \Pr(x_S \neq \epsilon_S) + E[x_S | x_S = \epsilon_S] \cdot \Pr(x_S = \epsilon_S) \quad (4a)$$

$$= E[2\bar{a} - \epsilon_N] \cdot \Pr(x_S \neq \epsilon_S) + E[\epsilon_S | \epsilon_S \geq 2\bar{a} - \epsilon_N] \cdot \Pr(x_S = \epsilon_S) \quad (4b)$$

$$> E[2\bar{a} - \epsilon_N] \cdot \Pr(x_S \neq \epsilon_S) + E[\epsilon_S] \cdot \Pr(x_S = \epsilon_S) \quad (4c)$$

$$\geq E[\epsilon_N] \cdot \Pr(x_S \neq \epsilon_S) + E[\epsilon_S] \cdot \Pr(x_S = \epsilon_S) \quad (4d)$$

$$= E[\epsilon_N] \quad (4e)$$

$$= E[x_N], \quad (4f)$$

where (4b) reflects the optimal policy studied above; (4c) holds because an idle bank never has a reserve deficiency, so that the unconditional distribution of ϵ_S is truncated below at $2\bar{a} - x_N$; (4d) holds because reserve requirements bind on average, i.e., $\bar{a} \geq E[\epsilon_N]$; and (4f) holds because in

equilibrium, with constant interest rates, banks never trade on non-settlement day. \square

Thus, under the mild regularity conditions we have imposed on the definition of reserve requirements and on the distribution of reserve shocks (aimed, essentially, to assure that reserve requirements bind weakly on average, and with a positive probability in all states), banks *strictly* prefer to hold more reserves on settlement day if $r_S = r_N$. Intuitively, extending this section's model so as to allow for a well-defined (as opposed to vertical) daily trade-off between prices and quantities should preserve, for small differences between settlement and non-settlement rates, banks' preference for postponing acquisition of reserves to settlement day. Before undertaking this extension in Section 5, however, we outline how the model of this section can be closed by industry aggregation and to include the Fed's intervention.

4.5. Industry equilibrium, interest targeting, and Fed intervention. To study the market's equilibrium and the associated behavior of the Fed, we must aggregate reserve demand across banks. As discussed above, we abstract from aggregate interest rate variability by assuming that the market is populated by a continuum of banks such as that studied in Sections 4.1-4.4, and by letting each bank receive an idiosyncratic ϵ_t shock on each day.¹²

These assumptions have two related implications. First, for the banking sector as a whole, there is no uncertainty on aggregate liquidity, implying non-stochastic equilibrium interest rates.¹³

¹² Strictly speaking we should identify shocks using both time- and bank-specific subscripts, as in ϵ_{it} . We can avoid this notational complication, provided that ϵ_t is understood to be bank-specific.

¹³ See Hamilton (1996) for a similar assumption. Bartolini et al. (1998) make the polar assumption that there is only a common, market-wide shock, reflecting that paper's focus on interest rate volatility rather than on the rates' mean behavior.

Recalling that we interpret exogenous liquidity shocks as customer-originated transactions, and that their average $\int_{-\infty}^{\infty} \epsilon dF(\epsilon)$ equals μ with probability one, the assumption that liquidity shocks are identically distributed across days implies that liquidity net of Fed intervention displays no maintenance-period-related pattern, consistently with the evidence documented in Section 2.

Second, the assumption of bank-specific shocks implies that the empirical frequency of the time-series behavior of an individual bank's reserves corresponds to the steady-state (ergodic) cross-sectional distribution of reserves among banks: since banks are identical ex-ante, the probability of each of them being in a particular state must equal the fraction of time a bank spends, on average, in that particular state between $t=0$ and $t=\infty$ (i.e., as $t \rightarrow \infty$, the bank's initial condition becomes irrelevant, and banks are mixed according to their steady-state distribution).

In this setting, it is easy to characterize the Fed's interest-targeting policy: intervention, added to liquidity originating from the non-bank sector, must be such as to satisfy banks' daily reserve demand at the target rate. That is, denoting its target for day t by r_t^* , the Fed must provide m_t so that:

$$m_N = \int_{-\infty}^{\infty} x_N(\epsilon, r_N^*, r_S^*) dF(\epsilon) - \int_{-\infty}^{\infty} \epsilon dF(\epsilon), \quad (5a)$$

$$m_S = \int_{-\infty}^{\infty} x_S(\epsilon_N, \epsilon, r_N^*, r_S^*) dF(\epsilon) - \int_{-\infty}^{\infty} \epsilon dF(\epsilon), \quad (5b)$$

where $x_i(\cdot)$ is an individual bank's demand for end-day reserves for day t , given past liquidity shocks and the market-determined interest rates. Now, in the model of this section, $r_N = r_S$. Hence the target rate should be a constant r^* over the two days. In turn, with equal settlement and non-settlement interest rates, banks remain idle on non-settlement day. Therefore:

$$x_N(\epsilon_N, r^*, r^*) = \epsilon_N, \quad (6)$$

so that, from (5a) and (6):

$$\int_{-\infty}^{\infty} x_N(\epsilon, r^*, r^*) dF(\epsilon) = \int_{-\infty}^{\infty} \epsilon dF(\epsilon) = \mu \quad (7)$$

is the liquidity of the banking sector at the end of non-settlement day, and $m_N = 0$.

On settlement day, banks borrow (lend, if negative) $2\bar{a} - \epsilon_N - \epsilon_S$ if either $\epsilon_N + \epsilon_S < 2\bar{a}$ or $\epsilon_N + \epsilon_S > 2\bar{a} + k/r^*$, and remain idle otherwise. Therefore:

$$x_S(\epsilon_N, \epsilon_S, r^*, r^*) = 2\bar{a} - \epsilon_N, \quad \text{for } \epsilon_N + \epsilon_S < 2\bar{a} \text{ or } \epsilon_N + \epsilon_S > 2\bar{a} + k/r^*, \quad (8a)$$

$$x_S(\epsilon_N, \epsilon_S, r^*, r^*) = \epsilon_S, \quad \text{otherwise.} \quad (8b)$$

So that, combining (5b) and (8b), and using Result 4,

$$m_S = \int_{-\infty}^{\infty} x_S(\epsilon_N, \epsilon, r^*, r^*) dF(\epsilon) - \int_{-\infty}^{\infty} \epsilon dF(\epsilon) = E[x_S] - \mu > 0 = m_N. \quad (9)$$

Intuitively, to smooth interest rates over the period, the Fed must issue more liquidity on settlement days than on non-settlement days, a prediction matching the evidence discussed in Section 2. What happens, however, if the Fed is unwilling to fully accommodate the increase in reserve demand on settlement days? To answer this question, we now consider a slightly more general model, which allows for differences in equilibrium rates across days.

5. Patterns in Interest Rates and Reserves with Liquidity Yields

While the simple model of the previous section allows analytic results and conveys the intuition behind our problem, it enforces a tight no-arbitrage condition whereby interest rates must remain constant over each maintenance period and banks do not trade at all on non-settlement

days. For our model to capture realistic patterns in interest rates and reserves, we now reinstate a liquidity-provision role for reserves by assigning positive values to the parameter α in (1). Since deviations from the target level of reserves T are now costly, banks will not open *unbounded* reserve positions in response to *finite* differences in rates across days. Equilibrium interest rates can then differ across days in response to liquidity shocks. As noted in Section 2, a liquidity-provision role for reserves also gives to banks an incentive to hold less reserves on days with higher interest rates. The resulting model can still feature high reserves on high-interest-rate days if the transaction-saving motive identified in the previous section is not more than offset by this liquidity-targeting motive. As we show below, this is indeed the case in equilibrium when banks' behavior is allowed to interact with a realistic description of the Fed's behavior, whereby the Fed tolerates small fluctuations of interest rates around their target level. The resulting model can then capture the positive correlation between interest rates and reserves we identified in U.S. data.

5.1. Settlement-day behavior. On day S , a bank must choose whether to trade or not, a choice it makes by comparing the loss from trading with the loss from remaining idle, knowing that if it does trade, then it would trade the amount minimizing its total loss. This optimal holding of reserves *conditional on trading* is obtained by differentiating (2b) with respect to x_S , upon setting $k \cdot 1_{[x_S \neq \epsilon_S]} = k$, which yields $x_S^* = T - \frac{r_S}{\alpha}$. Since this unconstrained optimum may fail to satisfy reserve requirements, settlement-day reserves conditional on trading equal $\max\{x_S^*, 2\bar{a} - x_N\}$.

Of course, a bank may choose not to trade on day S to save the cost k . We now show that the optimal policy for day S is defined by a range of inaction whose boundaries depend on the inherited reserves x_N : the bank does not trade if $k \cdot 1_{[x_S \neq \epsilon_S]} = k$; otherwise it trades, so as to set

reserves at $\max\{x_S^*, 2\bar{a}-x_N\}$. We characterize this policy as a function of x_N over three ranges:

i) $x_N \geq 2\bar{a}-x_S^* + \sqrt{\frac{2k}{\alpha}}$: in this case x_S^* satisfies the requirement $x_N+x_S^* > 2\bar{a}$. Hence, if the bank trades, its loss is minimized at $x_S^*r_S + (\alpha/2)(x_S^*-T)^2 + k = r_S T - \frac{r_S^2}{2\alpha} + k$. By remaining idle, instead, the bank's loss is $\epsilon_S r_S + (\alpha/2)(\epsilon_S - T)^2$. Inaction is then optimal if $x_S^* - \sqrt{\frac{2k}{\alpha}} \leq \epsilon_S \leq x_S^* + \sqrt{\frac{2k}{\alpha}}$; the bank borrows (lends, if negative) $x_S^* - \epsilon_S$ otherwise.

ii) $2\bar{a}-x_S^* + \sqrt{\frac{2k}{\alpha}} > x_N \geq 2\bar{a}-x_S^*$: in this case, the lower bound of the inaction band that was optimal in case **i)**, $x_S^* - \sqrt{\frac{2k}{\alpha}}$, falls below the required level $2\bar{a}-x_N$. The bank must then acquire reserves even for ϵ_S in the range $\left[x_S^* - \sqrt{\frac{2k}{\alpha}}, 2\bar{a}-x_N\right]$, where it would have preferred to remain idle. As it is forced to pay the cost k for these values of ϵ_S , the bank will choose to hold x_S^* . Its policy is to remain idle if $2\bar{a}-x_N \leq \epsilon_S \leq x_S^* + \sqrt{\frac{2k}{\alpha}}$, and to borrow (lend, if negative) $x_S^* - \epsilon_S$ otherwise.

iii) $2\bar{a}-x_S^* > x_N$: now reserves inherited from day N are too low for x_S^* to satisfy the requirement $x_N+x_S^* \geq 2\bar{a}$. Hence, if the bank trades, it can at best settle on the corner solution $x_S = 2\bar{a}-x_N$, which just meets the requirement. (The problem's convexity ensures that the constrained optimum lies at the corner.) Since x_S^* cannot be attained, the bank may well decide not to trade at all. To study this decision, let $\xi(x_N) \equiv 2\bar{a}-x_N-x_S^* \equiv 2\bar{a}-x_N - \left(T - \frac{r_S}{\alpha}\right)$, which is positive in this range and is the level of reserves the bank is required to hold in excess of the unconstrained optimum. If the bank trades, then it holds $x_S = 2\bar{a}-x_N \equiv x_S^* + \xi(x_N)$ and its loss is $(x_S^* + \xi(x_N))r_S + (\alpha/2)(x_S^* + \xi(x_N) - T)^2 + k = r_S T - \frac{r_S^2}{2\alpha} + \frac{\alpha \xi(x_N)^2}{2} + k$. Since the loss from inaction is $\epsilon_S r_S + (\alpha/2)(\epsilon_S - T)^2$, the bank stays idle if $2\bar{a}-x_N \leq \epsilon_S \leq x_S^* + \sqrt{\frac{2k}{\alpha} + \xi(x_N)^2}$ and holds $2\bar{a}-x_N$ otherwise.

5.2. Non-settlement day. Stepping back to day N , we first considering the expected loss for day S , based on information available at the end of day N . To do so, recall definition (1) and write:

$$\begin{aligned}
\mathbb{E}[V_S(\epsilon_S, x_N)] &= \mathbb{E}\left[\min_{x_S} \left(l(x_S, r_S) + k \cdot \mathbf{1}_{[x_S \neq \epsilon_S]} \right), \text{ s.t. } x_N + x_S \geq 2\bar{a} \mid x_N \right] \\
&= \left(\text{Prob}(\epsilon_S < L_S(x_N)) + \text{Prob}(\epsilon_S > U_S(x_N)) \right) \left[k + l(\max\{x_S^*, 2\bar{a} - x_N\}, r_S) \right] \\
&\quad + \text{Prob}(L_S(x_N) \leq \epsilon_S \leq U_S(x_N)) \mathbb{E}\left[l(\epsilon_S, r_S) \mid L_S(x_N) \leq \epsilon_S \leq U_S(x_N) \right] \\
&= \left(F(L_S(x_N)) + 1 - F(U_S(x_N)) \right) \left(k + l(\max\{x_S^*, 2\bar{a} - x_N\}, r_S) \right) + \int_{L_S(x_N)}^{U_S(x_N)} l(\epsilon, r_S) dF(\epsilon) \\
&\equiv \tilde{V}_S(x_N), \tag{10}
\end{aligned}$$

where the boundaries of the inaction range, $L_S(x_N)$ and $U_S(x_N)$, are those studied above for cases **i)-iii)**, and are summarized in Table 5.

It is now easy to compute $\tilde{V}_S(x_N)$ by numerical integration for any given distribution of the shock ϵ_t , and to find the optimal reset point for day N by a standard grid search over x_N . The solution for x_N can then be checked against the analytic first-order conditions for optimal trading yielded by (2a):

$$r_N + \alpha(x_N^* - T) + \tilde{V}'_S(x_N^*) = 0. \tag{11}$$

The inaction range on day N is also easy to obtain numerically, by comparing expected losses from trading with expected losses from inaction. The bank's policy on day N is then the following: *do not trade if $L_N \leq \epsilon_N \leq U_N$; trade, re-setting reserves at x_N^* , otherwise.* L_N and U_N satisfy the conditions:

$$\tilde{V}_S(L_N) + l(L_N, r_N) = \tilde{V}_S(x_N^*) + l(x_N^*, r_N) + k, \tag{12a}$$

$$\tilde{V}_S(U_N) + l(U_N, r_N) = \tilde{V}_S(x_N^*) + l(x_N^*, r_N) + k. \tag{12b}$$

This solution is shown in Figure 5 for an illustrative set of parameters (discussed below). In the figure, the total expected cost as of day N , $V_N(x_N)$, is plotted as the sum of the expected loss for day S , $\tilde{V}_S(x_N)$, and the current cost of holding reserves on day N , $l(x_N, r_N)$. The former is decreasing in x_N , as the bank foresees holding less reserves on day S when holding more reserves on day N ; the latter is increasing in x_N , as the opportunity and liquidity costs of holding reserves on day N rise with x_N . Both are convex in x_N . Their sum, $V_N(x_N)$, is U-shaped and is minimized at x_N^* . Interestingly, and intuitively, this optimal reset point exceeds the daily average reserve requirement (\$3 million in the figure), because overshooting requirements decreases the likelihood of having to transact again on day S . The figure also identifies the optimal range of inaction, whose bounds--as required by (12a-b)--are the points L_N and U_N at which $V_N(x_N)$ equals the expected loss at x_N^* , $V_N(x_N^*)$, plus the cost k .

5.3. Implications for Fed intervention in equilibrium. It is now easy to characterize the aggregate demand for reserves, and then to close the model by incorporating the Fed's daily target rates. Following the previous analysis, the $[F(U_N) - F(L_N)]$ banks whose liquidity falls in the $[L_N, U_N]$ inaction range do not trade on day N and hold ϵ_N . The remaining banks (with mass $[F(L_N) + 1 - F(U_N)]$) demand $x_N^* - \epsilon_N$. Therefore, the aggregate net demand of reserves on day N is:

$$X_N^*(r_N, r_S) = \int_{-\infty}^{L_N} [x_N^* - \epsilon_N] dF(\epsilon) + \int_{U_N}^{\infty} [x_N^* - \epsilon_N] dF(\epsilon), \quad (13)$$

which depends on r_N and r_S through L_N , U_N , and x_N^* . In equilibrium, the Fed must provide $m_N^* = X_N^*(r_N^*, r_S^*)$ if its target rates for the two days are r_N^* and r_S^* , respectively.

We proceed to illustrate the model's solution, and to assess the extent to which our model--despite its many simplifying assumptions--delivers reasonably realistic results. While we try to choose plausible parameters, we are of course aware of the crude nature of our exercise.¹⁴

We fix the interest rate on day N , r_N , at 5 percent (on an annual basis, like all other rates mentioned in this section), and solve the model for a variety of possible levels of the rate on day S , r_S . We assume the liquidity shocks to be normally distributed, and choose parameters drawing on Clouse and Elmendorf's (1997) study of the federal funds market. From that study we draw estimates of total required reserve balances of \$18 billion; of the reserves banks would optimally hold for liquidity purposes (the empirical analog of $(T-r/\alpha)$), about \$6 billion; of the interest semi-elasticity of reserve balances, which Clouse and Elmendorf set (mostly judgementally) at 0.1 percent; and of the standard deviation of *aggregate* reserve shocks, \$1-1½ billion per day. With 6,000 banks in Clouse and Elmendorf's sample, these figures suggest parameters for our model in the range of \$3 million for \bar{a} and T , and of $1E-10$ for α . Ceteris paribus, these parameters imply that each additional interest basis point decreases a typical bank's reserves by $10^{10} \times 0.01 / 360 = 1389$ dollars; and that a bank's optimal reserves in the absence of reserve requirements (the analog of $(T-r/\alpha)$) should be in the one- to two-million dollar range, for interest rates of about 5 percent.

The volatility of aggregate shocks (absent in our model) provides no information on the cross-sectional variability of bank-level liquidity; with a standard deviation of aggregate liquidity

¹⁴ Clearly, our simple calibration cannot capture the historical evolution, nor the cross-sectional features, of the market we study. For instance, while we hold required reserves constant, these have been on a declining trend in recent years, reflecting banks' growing use of "sweep" practices that reduce required reserves by transferring liabilities subject to requirements (e.g., checking accounts) overnight to liabilities not subject to requirements (e.g., money-market accounts). We also abstract from differences in behavior among banks, such as those between small banks (typically net lenders in the market) and large banks (typically net borrowers).

shocks of about \$0.2 million per bank, however, a standard deviation of end-day bank-level reserve shocks of $\sigma=\$0.5$ million seems reasonable. We then set $\mu = \bar{a}$ ($=T$), assuming that--in the long run--banks could optimally manipulate their non-bank customer and portfolio positions so as to bring average reserve balances near their required level. Finally, we set k conditional on the other parameters, so that the model predicts excess reserves of about 2 percent of required reserves, as in our estimates of Section 2; this is the case for a transaction cost k in the order of \$90. This amount, while larger than explicit costs of trading (e.g., Fedwire/Chips and brokerage fees), seems small relative to that of other resources--such as expensive managerial time--banks must dedicate to each reserve transaction, rather than to other profitable tasks.

The model's dynamic properties are summarized in Figure 6, which displays on its vertical axis the aggregate daily demand of reserves, computed as the average from a Montecarlo run with 20,000 maintenance periods. On average over the two days, reserves exceed the requirement by about 1.8 percent, and barely decrease as the settlement-day rate r_S (on the horizontal axis) rises from slightly below to slightly above the non-settlement day rate $r_N=5$ percent. Realistically, aggregate excess reserves are positive on both days, but markedly lower on the first day. The distance between the (descending and ascending) lines plotting mean daily reserves corresponds to the Fed's intervention needed to secure market clearing at the value of r_S plotted on the horizontal axis. Thus, for instance, the value $r_S=5$ percent yields the solutions for m_N^* and m_S^* when the Fed fully accommodates banks' demand for reserves at a constant target rate (since r_N is also set at 5 percent). For our choice of parameters, to smooth interest rates completely, the Fed would have to ensure that reserves lie at the points marked by round dots in the figure, and increase its supply of reserves on settlement days by about 2.2 percent over its non-settlement-days' supply.

Solutions for higher values of r_S illustrate equilibria in which the Fed lets demand for liquidity in high-demand (typically, settlement) days partly spill over into a small rise in interest rates. This realistically captures the Fed's reluctance to provide liquidity elastically at a fixed target rate, and its preference for enforcing a small corridor of rates, the width of which reflects its tolerance for day-to-day expected interest rate fluctuations. For instance, Feinman (1993) estimates that only expected deviations from the target rate greater than 20 basis point have, historically, triggered Fed intervention with almost-unit probability. Feinman's roughly linear intervention-hazard function (valued at about 0.20 at a zero-expected-deviation of the market rate from its target) corresponds to a band of tolerance less than 20 basis points wide.

Our analysis in Section 2 provides concurring--if crude--evidence to the same extent, pointing to the Fed's tolerance of a systematic spread in rates within maintenance periods of about 15 basis points. In our simulations, a spread of 15 basis points between high-demand (settlement) and low-demand (non-settlement) days requires the Fed to provide 1.7 percent more reserves on settlement days than on non-settlement days, at the levels marked by square dots in Figure 6. This gap falls for higher settlement-day rates, as a larger interest premium trims the inaction band's width. Should the Fed decide to let the higher demand of reserves on day S spill over *only* to interest rates (hence allowing no adjustment in reserves), it would have to accept settlement-day rates 64 basis points higher than non-settlement-day rates.¹⁵

¹⁵ Qualitatively similar results emerge for different choices of parameters. An interesting case is that of a smaller α , for which our numerical results approach the analytical results of Section 4. In the limiting case of $\alpha=0$, a flat interest-rate pattern ($r_S = r_N = 5$ percent) would be needed to prevent arbitrage. In this case, holding the other parameter at the values discussed above, excess settlement-day reserves would be 3.2 percent above their non-settlement-day level (with period-average excess reserves of about 3 percent of the required level). Increasing k to \$190 brings period-average excess reserves up to 8 percent of their required level.

6. Concluding Remarks

A simple model of the U.S. market for federal funds can explain the tendency of U.S. banks to hold more reserves--and of the Federal Reserve to provide them--on days when the cost of holding reserves is cyclically highest: uncertainty over reserve needs and small costs of trading induce banks to trade mainly when they are endowed with the most precise information on their reserve needs--particularly on settlement days--even if this strategy entails higher costs of carrying reserves overnight. In turn, upward pressure on interest rates around period-end combines with the Fed's goal to smooth interest rates over time to induce the central bank to provide reserves mostly around the end of each reserve maintenance period.

Our model is clearly stylized, and focussed on only a few key features of the U.S. market for federal funds. We designed it with the main goal of replicating the observed correlation between prices and quantities across settlement and non-settlement days. While doing so, we left aside other features of this market--such as week-end effects, heterogeneous bank behavior, and patterns of interest rate volatility. Existing research provides a convincing explanation for these phenomena, which appear analytically unrelated to the role of the transaction strategies on which we focus. Our modeling strategy also disregarded the intriguing--if less significant--pattern of declining reserve holdings and interest rates during the first few days of typical maintenance periods. However, we conjecture that a multi-period model of liquidity management in the same spirit of our model should be able to explain this pattern. If banks behave as our model predicts, they would hold unusually high reserves on settlement days, and dispose of these excess reserves in the following period, actually unwinding their positions only when reached by a liquidity shock sufficiently large to justify the payment of the transaction cost. Thus, transaction costs would lead

to serially correlated demand for reserves, in analogy to models of investment or hiring and firing in models with costly capital and labor adjustment. To the extent that the Fed accommodates this demand only imperfectly, any such cyclical behavior on the market's demand side would also lead to a cyclical behavior of interest rates.

Other aspects of the federal funds market deserve closer scrutiny and are related to this paper's goals and modeling perspective. Our model, for instance, implies that banks prefer to transact late in the maintenance period, when more relevant information is available, suggesting that the volume of *gross* transactions should rise around end-period. It would be interesting to verify this implication, but daily data on gross interbank transactions are not available. Recently, Dell'Ariccia and Garibaldi (1999) have also shown that search costs in credit markets may help explain some features of the dynamics of gross credit flows at the quarterly frequency. This view, applied to our study's higher frequency, points to a channel that would strengthen our results: if settlement-day transaction costs fall as trade volume rises (e.g., because of the associated decline in "search" costs), then banks' tendency to cluster trading on these days should be exacerbated.

Finally, future research should offer a more detailed analysis of market imperfections and institutional features which we captured under the gross heading of "transaction costs." We suspect, for instance, that indivisibilities in federal funds trades and rigidities induced by Fed regulations (especially penalties for overnight overdrafts) may also induce banks to defer federal funds trading to the late portion of each maintenance period. A formal analysis of these rigidities is more difficult than that of transaction costs, and is best deferred to further research. In general, however, a more detailed and realistic description of constraints and frictions faced by market participants should yield results similar to those of this paper, as long as a bias towards "late" trading--whether stemming from transaction costs, or other sources--is preserved.

Appendix. Summary of Reserve Requirements in the United States

Under the Garn-St. Germain Depository Institutions Act of 1982, most commercial banks, mutual savings banks, savings and loan institutions, credit unions, and branches and agencies of foreign banks operating in the United States are subject to reserve requirements. Banks with reservable liabilities of less than \$2 million are not subject to reserve requirements; slightly larger banks are subject to requirements, although they are only required to report quarterly. A regime of almost-contemporaneous two-week computation and maintenance periods was established in 1984. (Until then, one-week computation periods determined reserves to be held in one-week maintenance periods lagged by two weeks.) In this regime, required reserves are computed over two-week “computation periods” beginning on Tuesday and ending on the second following Monday; actual reserves must be held over “maintenance periods” beginning on Thursday, two days after the beginning of the computation period, and ending two weeks later on “settlement Wednesday.” (The levels prevailing on the most recent business day are imputed to week-ends and holidays.) In July 1998 the maintenance period was lagged by an additional four weeks with respect to the previous regime. Reservable liabilities consist of average transaction deposits (including checking accounts and demand deposits, but excluding money-market deposits and savings accounts); reserves include average Federal Reserve deposits and vault cash, the latter applied with a one-period lag. Excess Fed balances (but not excess vault cash), up to the greater of \$50,000 or four percent of required reserves, may be carried over to the next maintenance period. Four percent of deficiencies may also be carried over to the next period, while further deficiencies are charged a penalty equal to the discount rate plus 2 percent.

References

- Balduzzi, Pierluigi, Giuseppe Bertola, and Silverio Foresi (1997), "A Model of Target Changes and the Term Structure of Interest Rates," *Journal of Monetary Economics* 39:2,
- Balduzzi, Pierluigi, Giuseppe Bertola, Silverio Foresi, and Leora Klapper (1998), "Interest Rate Targeting and the Dynamics of Short-Term Rates," *Journal of Money, Credit, and Banking* 30, 26-50.
- Bartolini, Leonardo, Giuseppe Bertola, and Alessandro Prati (1998), "Day-to-Day Monetary Policy and the Volatility of the Federal Funds Interest Rate," European University Institute, Working Papers in Economics No. 98/35.
- Campbell, John Y. (1987), "Money Announcements, the Demand for Bank Reserves, and the Behavior of the Federal Funds Rate within the Statement Week," *Journal of Money, Credit, and Banking* 19, 56-67.
- Clouse, James A. (1994), "Recent Developments in Discount Window Policy," in *Federal Reserve Bulletin* 80:11, 965-977.
- Clouse, James A., and James P. Dow, Jr. (1999), "Fixed Costs and the Behavior of the Federal Funds Rate," *Journal of Banking and Finance* 23, 1015-29.
- Clouse, James A., and Douglas W. Elmendorf (1997), "Declining Required Reserves and the Volatility of the Federal Funds Rate," Finance and Economics Discussion Series No. 1997-30, Federal Reserve Board, Washington, DC.
- Dell'Araccia, Giovanni, and Pietro Garibaldi (1999), "Gross Credit Flows," Bocconi University, mimeo; presented at the 1998 Summer Meeting of the Econometric Society, Berlin.
- Feinman, Joshua N. (1993), "Estimating the Open Market Desk's Daily Reaction Function," *Journal of Money, Credit, and Banking* 25, 231-427.
- Goodfriend, Marvin (1991), "Interest Rates and the Conduct of Monetary Policy," *Carnegie-Rochester Series on Public Policy* 34, 7-30.
- Griffiths, Mark D., and Drew B. Winters (1995), "Day-of-the-Week Effects in Federal Funds Rates: Further Empirical Findings," *Journal of Banking and Finance* 19, 1265-1284.
- Hamilton, James D. (1996), "The Daily Market for Federal Funds," *Journal of Political Economy* 104, 26-56.
- Hancock, Diana, David B. Humphrey, and James A. Wilcox (1999), "Cost Reductions in Electronic Payments: The Roles of Consolidation, Economies of Scale, and Technical Change," *Journal of Banking and Finance* 23, 391-421

- Ho, Thomas S. Y., and Anthony Saunders (1985), "A Micro Model of the Federal Funds Market," *Journal of Finance* 40, 977-988.
- Kopecky, Kenneth J., and Alan L. Tucker (1993), "Interest Rate Smoothness and the Nonsettling-day Behavior of Banks," *Journal of Economics and Business* 45 297-314
- Lasser, Dennis J. (1992), "The Effect of Contemporaneous Reserve Accounting on the Market for Federal Funds," *Journal of Banking and Finance* 16, 1047-1056.
- Peristiani, Stavros (1998), "The Growing Reluctance to Borrow at the Discount Window: An Empirical Investigation," *Review of Economics and Statistics* 80, 611-620.
- Meulendyke, Ann-Marie (1989), *U.S. Monetary Policy and Financial Markets*, Federal Reserve Bank of New York, December, 124-147.
- Mood, Alexander M., Franklin A. Graybill, and Duane C. Boes (1974), *Introduction to the Theory of Statistics*, New York: McGraw-Hill.
- Rudebusch, Glenn D. (1995), "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure," *Journal of Monetary Economics* 35, 245-274.
- Spindt, Paul A., and Ronald J. Hoffmeister (1988), "The Micromechanics of the Federal Funds Market: Implications for Day-of-the-Week Effects in Funds Rate Variability," *Journal of Financial and Quantitative Analysis* 23:4, 401-416.
- Stigum, Marcia (1990), *The Money Market*, 3rd edition, Homewood, IL: Dow Jones-Irwin.

Table 1

Federal Funds Rate: Summary statistics and tests of equality between settlement and non-settlement days
(January 2, 1986 - July 1, 1998)

Statistic	Non-settlement days	Settlement days	Equality Test	Degrees of freedom	Value	Probability
Mean difference from day 1 (basis points)	-7.5	10.3	<i>t</i> -test	2812	7.44	0.000
Median difference from day 1 (basis points)	-6.0	6.0	Wilcoxon/Mann-Whitney	1	7.59	0.000
			Median chi-square	1	32.80	0.000
			Adj. median chi-square	1	32.124	0.000
			Kruskal-Wallis	1	57.62	0.000
			van der Waerden	1	70.95	0.000
Fraction of increases (percent)	40.5	69.3	Pearson chi-square	1	98.12	0.000
			Likelihood ratio chi-square	1	99.20	0.000
Standard Deviation of daily changes (basis points)	28.0	47.4	F-test	(2490,322)	2.86	0.000
			Bartlett	1	204.58	0.000
Mean absolute median difference of daily changes (basis points)	12.8	31.1	Brown-Forsythe	(1,2812)	135.68	0.000
Skewness of daily changes	9.69	5.41				

Table 2

Excess Reserves: Summary statistics and tests of equality between settlement and non-settlement days
(January 2, 1986 - July 1, 1998)

Statistic	Non-settlement days	Settlement days	Equality Test	Degrees of freedom	Value	Probability
Mean difference from day 1 (percent of required reserves)	1.8	7.5	<i>t</i> -test	3138	21.82	0.000
Median difference from day 1 (percent of required reserves)	1.6	6.8	Wilcoxon/Mann-Whitney	1	14.74	0.000
			Median chi-square	1	87.24	0.000
			Adj. median chi-square	1	86.15	0.000
			Kruskal-Wallis	1	217.32	0.000
			van der Waerden	1	272.66	0.000
Fraction of increases (percent)	48.9	69.3	Pearson chi-square	1	48.75	0.000
			Likelihood ratio chi-square	1	50.01	0.000

Table 3

Securities bought by the Fed (outright or under repo): Summary statistics and tests of equality between settlement and non-settlement days (January 2, 1986 - July 1, 1998)

Statistic	Non-settlement days	Settlement days	Equality Test	Degrees of freedom	Value	Probability
Mean difference from day 1 (percent of required reserves)	-3.9	1.5	<i>t</i> -test	3138	8.63	0.000
Median difference from day 1 (percent of required reserves)	-3.1	1.5	Wilcoxon/Mann-Whitney	1	8.25	0.000
			Median chi-square	1	36.61	0.000
			Adj. median chi-square	1	35.90	0.000
			Kruskal-Wallis	1	68.07	0.000
			van der Waerden	1	78.26	0.000
Fraction of increases (percent)	51.1	67.2	Pearson chi-square	1	30.07	0.000
			Likelihood ratio chi-square	1	30.76	0.000

Table 4

Excess reserves net of intervention: Summary statistics and tests of equality between settlement and non-settlement days (January 2, 1986 - July 1, 1998)

Statistic	Non-settlement days	Settlement days	Equality test	Degrees of freedom	Value	Probability
Mean difference from day 1 (percent of required reserves)	5.7	6.1	<i>t</i> -test	3138	0.68	0.499
Median difference from day 1 (percent of required reserves)	4.9	5.4	Wilcoxon/Mann-Whitney	1	0.52	0.606
			Median chi-square	1	0.17	0.681
			Adj. median chi-square	1	0.12	0.725
			Kruskal-Wallis	1	0.27	0.606
			van der Waerden	1	0.26	0.611
Fraction of increases (percent)	47.8	58.8	Pearson chi-square	1	14.10	0.000
			Likelihood ratio chi-square	1	14.14	0.000

Table 5

Optimal policy for settlement day

Inherited level of x_N	Inaction band for ϵ_S		Reset point, $x_S =$
	Lower bound $L_S(x_N)$	Upper bound $U_S(x_N)$	$\text{Max}\{x_S^*, 2\bar{a} - x_N\}$
<i>high</i> x_N : $x_N \geq 2\bar{a} - x_S^* + \sqrt{\frac{2k}{\alpha}}$	$x_S^* - \sqrt{\frac{2k}{\alpha}}$	$x_S^* + \sqrt{\frac{2k}{\alpha}}$	$x_S^* \equiv T - r_S/\alpha$
<i>intermediate</i> x_N : $2\bar{a} - x_S^* + \sqrt{\frac{2k}{\alpha}} > x_N \geq 2\bar{a} - x_S^*$	$2\bar{a} - x_N$	$x_S^* + \sqrt{\frac{2k}{\alpha}}$	x_S^*
<i>low</i> x_N : $2\bar{a} - x_S^* > x_N$	$2\bar{a} - x_N$	$x_S^* + \sqrt{\frac{2k}{\alpha} + (2\bar{a} - x_S^*)^2}$	$2\bar{a} - x_N$

Fig. 1: Federal funds rate and excess reserves (Jan. 2, 1986 - July 1, 1998)







Fig. 2: Federal funds rate and excess reserves (Jan. 6, 1994 - July 1, 1998)







Fig. 3: Federal funds rate and excess reserves (Jan. 2, 1986 - Jan. 5, 1994)

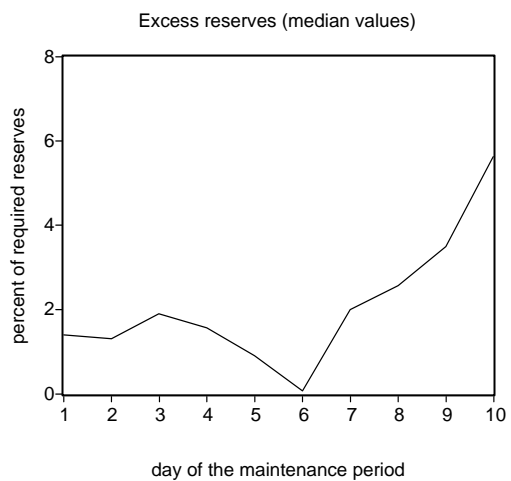

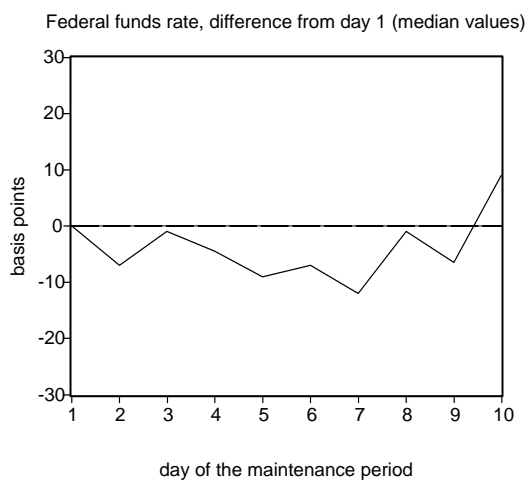
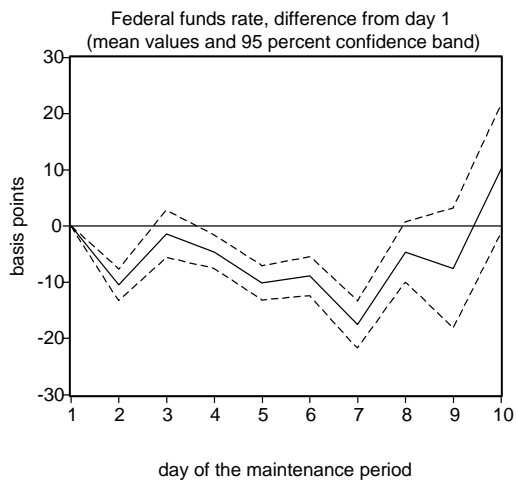


Fig. 4: Securities bought by the Fed, and excess reserves net of intervention (Jan. 2, 1986 - July 1, 1998)

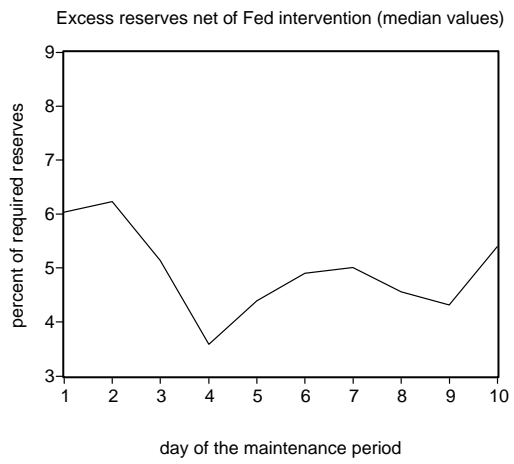


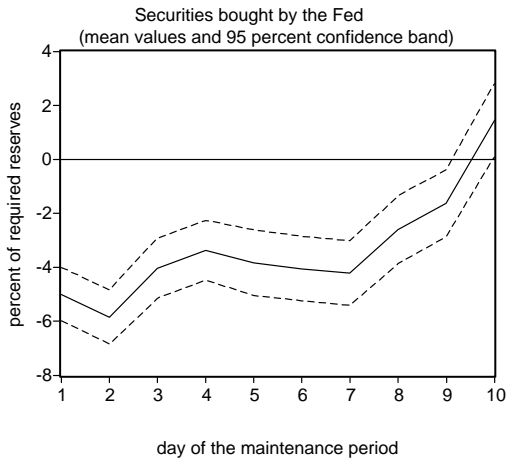


FIGURE 5: Optimal policy on non-settlement day.
 $r_N=5.00$, $r_S=5.15\%$ p.a., $k=\$90$, $T=\$3.00\text{m}$, $a=\$3.00\text{m}$
 $\mu=\$3.00\text{m}$, $\sigma=\$0.50\text{m}$, $\alpha=1.00\text{E}-010$




FIGURE 6: Interest rate differentials and optimal reserves.
 $r_N=5.00\%$ p.a., $k=\$90$, $T=\$3.00\text{m}$, $a=\$3.00\text{m}$
 $\mu=\$3.00\text{m}$, $\sigma=\$0.50\text{m}$, $\alpha=1.00\text{E}-010$

