Federal Reserve Bank of New York

Economic Policy Review

July 2000
Volume 6 Number 2

1 Listening to Loan Officers: The Impact of C rcial Credit Standards on Lending and Output
Cara S. Lown, Donald P. Morgan, and Sonali Rohatgi

17 The Timing and Funding of Fedwire Funds Transfers
James McAndrews and Samira Rajan

33 Capital Ratios as Predictors of Bank Failure
Arturo Estrella, Sangkyun Park, and Stavros Peristiani

53 Support for Resistance: Technical Analysis and Intraday Exchange Rates
Carol Osler

The Economic Policy Review is published by the Research and Market Analysis Group of the Federal Reserve Bank of New York. Articles undergo a comprehensive refereeing process prior to their acceptance in the Review. The views expressed in the articles are those of the individual authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

www.ny.frb.org/rmaghome

Federal Reserve Bank of New York Economic Policy Review

July 2000
Volume 6 Number 2

Contents

Articles:

1 **Listening to Loan Officers: The Impact of Commercial Credit Standards on Lending and Output**
Cara S. Lown, Donald P. Morgan, and Sonali Rohatgi

Over most of the last thirty-three years, the Federal Reserve has polled a small number of bank loan officers about their moves to tighten or ease commercial credit standards. Although the Senior Loan Officer Opinion Survey uses a small sample and gathers only qualitative information, it proves to be a useful tool in predicting changes in commercial lending and output. The authors find a strong correlation between tighter credit standards and slower loan growth and output, even after controlling for credit demand and other predictors of lending and output. The analysis also shows that the loan officer reports can help predict narrower measures of business activity, including inventory investment and industrial production.

17 **The Timing and Funding of Fedwire Funds Transfers**
James McAndrews and Samira Rajan

An examination of the Federal Reserve's Fedwire Funds Transfer service reveals that the highest concentration of funds-transfer value occurs in the late afternoon. The authors attribute this activity peak to attempts by banks (and their customers) to coordinate payment timing more closely. By synchronizing payments, banks can take advantage of incoming funds to make outgoing payments—especially during periods of heavy payment traffic. Conversely, during off-peak times, banks must rely more on account balances or overdrafts to fund payments, which increases the cost of making payments. For this reason, banks time their payments to coincide with an activity peak, thereby reinforcing the peak.

33 Capital Ratios as Predictors of Bank Failure
Arturo Estrella, Sangkyun Park, and Stavros Peristiani

The current review of the 1988 Basel Capital Accord has put the spotlight on the ratios used to assess banks' capital adequacy. This article examines the effectiveness of three capital ratios—the first based on leverage, the second on gross revenues, and the third on risk-weighted assets—in forecasting bank failure over different time frames. Using 1988-93 data on U.S. banks, the authors find that the simple leverage and gross revenue ratios perform as well as the more complex risk-weighted ratio over one- or two-year horizons. Although the risk-weighted measures prove more accurate in predicting bank failure over longer horizons, the simple ratios are less costly to implement and could function as useful supplementary indicators of capital adequacy.

53 Support for Resistance: Technical Analysis and Intraday Exchange Rates
Carol Osler

"Support" and "resistance" levels—points at which an exchange rate trend may be interrupted and reversed—are widely used for short-term exchange rate forecasting. Nevertheless, the levels' ability to predict intraday trend interruptions has never been rigorously evaluated. This article undertakes such an analysis, using support and resistance levels provided to customers by six firms active in the foreign exchange market. The author offers strong evidence that the levels help to predict intraday trend interruptions. However, the levels' predictive power is found to vary across the exchange rates and firms examined.

Cara S. Lown, Donald P. Morgan, and Sonali Rohatgi

Listening to Loan Officers: The Impact of Commercial Credit Standards on Lending and Output

- The Federal Reserve's Senior Loan Officer Opinion Survey offers information useful in forecasting commercial loan growth and overall economic activity.
- Statistical analysis reveals a strong correlation between loan officers' reports of tighter credit standards and slowdowns in commercial lending and output.
- Reported changes in credit standards can also help predict narrower measures of business activity, including inventory investment and industrial production.
- The chain of events following a tightening of standards resembles a "credit crunch": Commercial loans plummet, output falls, and the federal funds rate is lowered.

When the Federal Open Market Committee (FOMC) eased monetary policy on October 15, 1998, it noted the "growing caution by lenders and unsettled conditions in financial markets more generally. . . ." Sharply higher spreads on commercial paper and corporate bonds made it clear that the U.S. markets were unsettled, but how could policymakers tell if lenders were growing more cautious? Commercial bank loans are rarely traded, so loan rates are not instantly observable. Moreover, the "price" of commercial bank credit extends beyond the interest rate; bank loan officers set standards that firms must clear even before the rate is negotiated. These standards are decided in thousands of bank offices across the country, so how can the Federal Reserve tell if lenders are growing cautious? For that matter, how can we tell if banks are "throwing caution to the wind" and easing standards? We ask. Once each quarter, participants in the Federal Reserve's Senior Loan Officer Opinion Survey are asked whether their standards for making commercial loans have "tightened" or "eased" since the previous quarter. Loan officers at approximately sixty large domestic banks across the United States participate in the survey.

Although we praise the survey in the end, there certainly are reasons to doubt it. The survey is entirely qualitative, for one: respondents provide opinions, not hard numbers. The small

Cara S. Lown is a research officer, Donald P. Morgan an economist, and Sonali Rohatgi an assistant economist at the Federal Reserve Bank of New York.

The authors thank John Duca, Ken Kuttner, Phil Strahan, Egon Zakrajšek, participants in the Domestic Research/Banking Studies Seminar Series, and two anonymous referees. The views expressed are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

sample size is another concern: with more than 8,000 banks in the United States, can sixty bankers tell us anything useful about aggregate lending? Reporting bias is yet another concern; respondents in general may try to please surveyors, but the bias with this survey may be more severe because the loan officers work at banks that are likely supervised by the Federal Reserve. Loan officers who suspect, albeit wrongly, that their input will be used for supervisory purposes may shade their responses accordingly.

Because of these concerns, this article examines the value of the Senior Loan Officer Opinion Survey in predicting both lending and output.[1] We find that the changes in commercial credit standards reported by loan officers are indeed linked to aggregate loan growth. Commercial lending by U.S. banks slows substantially following reports of tighter standards, even after we control for other factors that might affect growth. Loan officers not only report accurately, they provide us with information that we could not infer from other measures of credit availability, such as loan rates, loan growth itself, or the mix of bank loans and other sources of credit. Changes in credit standards also help to predict economic growth and narrower measures of business activity such as inventory investment, a notoriously unforecastable variable that is closely tied to the banking sector.

Once each quarter, participants in the Federal Reserve's Senior Loan Officer Opinion Survey are asked whether their standards for making commercial loans have "tightened" or "eased" since the previous quarter.

In the end, we estimate a system of equations—a vector autoregression (VAR)—that enables us to isolate and quantify the impact of a shock to credit standards on lending output. Our VAR is an off-the-shelf model of the economy with two additional variables: commercial loans extended by banks and the change in commercial credit standards reported by bank loan officers. A shock to credit standards and its aftermath very much resemble a "credit crunch": Lenders tighten standards very abruptly, but ease up only gradually. Commercial loan volume at banks plummets immediately after the shock and does not bottom out until lenders start to ease standards again. Output also falls shortly after the tightening in standards. The federal funds rate, which we identify with the stance of monetary policy, declines.

In the next section, we describe the Senior Loan Officer Opinion Survey and relate the motivation for the survey to the credit "availability doctrine" of the 1950s and to more recent theories of quantity credit rationing. We then examine the correlation between the changes in commercial credit standards reported by loan officers in the survey and various measures of credit availability, including lending itself. We follow this discussion with a look at the link between standards and economic activity and an analysis of the impact of a "shock" to standards in a multiple-equation framework.

The Survey: Background and Motivation

The Senior Loan Officer Opinion Survey on Bank Lending Practices, as it is officially known, was unveiled in 1967. In its most recent incarnation, the survey includes approximately twenty core questions about the supply of and demand for various types of credit, including commercial credit. Apart from these regular questions, the survey includes ad hoc questions about disruptions and trends in credit markets. The sample includes about sixty domestic banks, usually the largest in each of the twelve Federal Reserve Districts. Banks are added or replaced as needed. "Megamergers" between very large U.S. banks in recent decades, for example, have necessitated frequent changes in the sample. The response rate of lenders is very near 100 percent.[2]

In contrast to the more quantitative survey on commercial loan rates, the Senior Loan Officer Opinion Survey is, as its name suggests, more *qualitative*.[3] Loan officers are essentially asked whether their standards for approving commercial credit have tightened or eased since the quarter before:

> Over the past three months, how have your bank's credit standards for approving loan applications for C&I [commercial and industrial] loans or credit lines—excluding those to finance mergers and acquisitions—changed? 1) tightened considerably, 2) tightened somewhat, 3) remained basically unchanged, 4) eased somewhat, 5) eased considerably.

Except for a hiatus in the 1980s, when the question was dropped, and apart from minor changes in wording and emphasis in earlier years, the basic question-and-answer options have been more or less the same over various eras of the survey.

1967-77. In terms of sample size and constancy, these were the golden years. The sample numbered 121 large U.S. banks, nearly twice the size of the sample today. Sample coverage was also relatively constant, as the bank mergers that cause frequent changes in the sample these days were yet to come. The standards question was virtually identical to the question above. The answer options differed only trivially in their wording: "much" instead of "considerably," for example, and "firmer" instead of "tightened."

1978-84. To account for the growing role of the prime lending rate in allocating bank credit, the question on standards was essentially divided in two in 1978. Lenders were first asked to report changes in standards for approving loans *at* the prime rate. A second question asked about standards for loans at "spreads above" prime. The answer options for each question were not changed. The survey was expanded to include a sample of foreign banks during this period, and the number of domestic banks in the survey was reduced to sixty, about the same as today.

1984-90. Questions on commercial credit standards were dropped altogether during this period. With the deregulation of deposit and other interest rates in the early 1980s, policymakers and their staffs may have presumed that bankers would rely more on unfettered interest rates and less on standards in allocating loans among borrowers.

With more than 8,000 banks in the United States, can sixty bankers tell us anything useful about aggregate lending?

1990 to fourth-quarter 1998. The standards questions were revived because of concerns about a possible credit crunch in the spring of 1990. The question is still divided in two, as in 1978-84, but the division these days is by firm size; lenders are asked to report separately on standards for small firms (with annual sales under $50 million) versus large and middle-sized firms.

The changes in standards reported by loan officers are pieced together in Chart 1. For the 1990s, we use the standards for loans to middle-sized and large firms (as opposed to small firms) on the theory that the former matter more in terms of aggregate lending conditions. The choice is largely immaterial, however, as the correlation between the two series was 0.96. For the years 1978-84, when the question on standards distinguished between loans at prime and loans above prime, we use the average of the responses to the two questions.

Chart 1

Changes in Commercial Credit Standards Reported over Various Periods of the Senior Loan Officer Opinion Survey

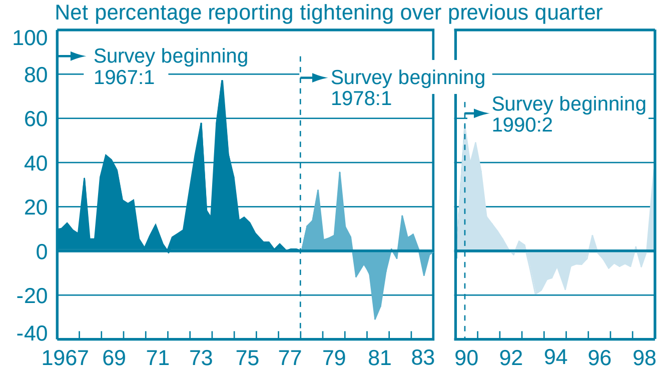

Source: Board of Governors of the Federal Reserve System, Senior Loan Officer Opinion Survey.

Plotted in the chart is the *net* percentage tightening: the number of loan officers reporting tightening standards less the number reporting easing divided by the total number reporting.[4] As Schreft and Owens (1991) noted, loan officers almost *never* reported a net easing of standards over the 1967-77 period; if the reported changes were summed, credit would have been *extremely* tight by the end of the period. This curious tendency to report tighter standards, at least in the early years, raises concerns about reporting bias; bank loan officers may be loath to ever tell the Federal Reserve that they are letting their standards fall.[5] The first substantial easing of standards was not reported until the 1980s. Credit standards were indeed tight in the early 1990s, after the hiatus, suggesting that credit-crunch concerns may have been well founded. The last substantial tightening reported by lenders was in 1998, after the Russian default and financial deterioration in southeast Asia.

The Importance of Standards

Why do monetary policymakers care about credit standards in the first place? Why not simply ask lenders to report loan rates and leave it at that? Because the market for credit may not operate like other markets, where prices do all the adjusting to keep the market cleared. For various reasons, loan rates may be secondary to standards of creditworthiness and other *nonprice* terms in the allocation of bank credit.

During the years leading up to the survey, interest rates were held down by government-imposed ceilings or by the Federal

Reserve's efforts to support Treasury bond prices. The stickiness of loan rates gave rise to the availability doctrine in the 1950s—the idea that the quantity of credit available from banks mattered more (for spending) than the price. While the availability doctrine has waned, modern theories of "quantity" rationing also emphasize the primacy of nonprice terms in the allocation of credit among competing borrowers.[6] The friction holding down interest rates in these theories is not government intervention, but the information and incentive problems that can gum up credit markets: adverse selection and moral hazard. By raising loan rates, lenders may drive off all but the least creditworthy applicants or elicit riskier behavior by borrowers. Rather than raising loan rates to curtail the supply of credit, lenders may tighten their standards and cut off credit to the marginal borrowers that do not meet the higher standards. In essence, credit markets may operate like a trendy night club in New York City: you have to clear the velvet rope before you pay the door charge.

Why do monetary policymakers care about credit standards in the first place? Why not simply ask lenders to report loan rates and leave it at that? Because the market for credit may not operate like other markets, where prices do all the adjusting to keep the market cleared.

Despite this theory, there is surprisingly little evidence that the commercial standards reported by loan officers actually matter for lending and output. Schreft and Owens (1991) noted the frequent breaks in the series and some of the curious features of the reports, but they did not actually test whether standards were informative nevertheless. Duca and Garrett (1995) and McCarthy (1997) investigate whether bankers' willingness to lend affects spending, but they focus on *consumer* credit standards and spending.

The propensity of lenders to always report tightening, especially in the early years, makes one wonder whether lenders are just "talking tough" when they say they are tightening. The link between loan growth and standards reported by senior loan officers can tell us; if their actions match their words, and if the actions of lenders in the survey are representative, then reports of tighter standards should lead to slower commercial loan growth, all else equal.

Credit Standards and Loan Growth

Loan growth does indeed slow at times when loan officers report tightening standards (Chart 2). Following the tightening reported from 1973 to 1975, for example, loan growth slowed and eventually turned negative. The sharp tightening reported during the early 1990s was also followed by much slower loan growth. Lending grew relatively rapidly in the ensuing years as loan officers began to report less tightening and eventually eased standards. The more recent tightening reported in the summer of 1998 also preceded sharply slower loan growth in the first quarter of 1999.

Table 1 confirms the negative relationship between standards and loan growth and shows that the changes in standards reported by loan officers are correlated with several other measures of credit availability as well. The correlation between loan growth and standards has been higher thus far in the 1990s, but it was also significant in the pre-1984 period.[7] In both periods, the change in standards tends to play the leading

Chart 2
Changes in Commercial Credit Standards and Loan Growth

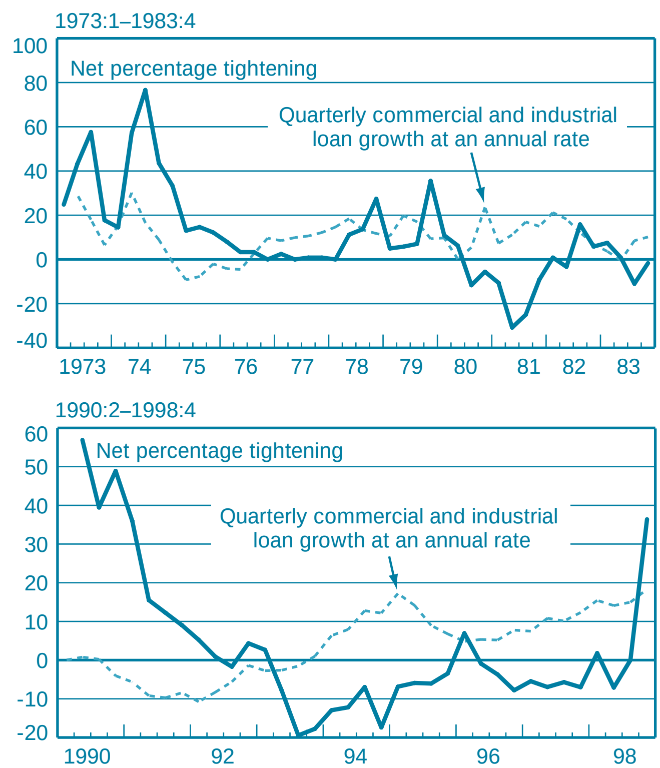

Source: Board of Governors of the Federal Reserve System, Federal Reserve Statistical Release H.8 and Senior Loan Officer Opinion Survey.

role. The loan spread—the difference between commercial loan rates and the federal funds rate—tends to rise following reports of tighter standards, as we would expect if lenders are contracting credit. Here again, standards play the leading role.

Changes in the "mix" of commercial paper and bank loans are also positively correlated with changes in standards. This mix variable was used in Kashyap, Stein, and Wilcox (1993) to identify shifts in the supply of bank loans relative to other sources of short-term credit. If loans become more expensive, owing to a monetary contraction, for instance, this mix tends to rise as the large firms that can borrow in the commercial paper market substitute paper for loans. Increases in the mix are also positively related to tightening credit standards, as Table 1 shows. This correlation is positive at both leads *and* lags, however, so we cannot say for sure whether standards lead the mix or vice versa.

In contrast to the paper-loan mix, which measures relative quantities, the paper-bill spread measures relative prices. Friedman and Kuttner (1992) and others have found that this spread is a particularly good predictor of economic activity, with higher spreads signaling slower future growth. Although researchers have different theories as to why this spread is so informative, one hypothesis is that a rise in the paper-bill spread signals disturbances in credit markets. As Table 1 shows, the spread is positively correlated with commercial bank credit standards. The strongest correlation is contemporaneous, suggesting that bankers and investors in the commercial paper market are reacting to the same news.

The last measure of credit availability in Table 1 comes from another survey, this one of firms. Once a quarter, a sample of small firms belonging to the National Federation of Independent Business is asked whether credit is "easier or harder" to get than it was in the previous quarter. As shown in the last column of the table, the net percentage of firms reporting easier credit availability falls as the net percentage of bankers reporting tightening standards rises.

Although suggestive, these correlations hardly prove that the tighter standards reported by bankers actually reduce the

Table 1

Changes in Commercial Credit Standards Reported by Bank Loan Officers and Measures of Credit Availability

Correlations at Various Leads and Lags

Standards Reported at t[a]	Δ *Loans*$_t$ t<1983:4	Δ *Loans*$_t$ t>1990:2	*Loan Spread*$_t$[b]	Δ *Paper-Loan Mix*$_t$[c]	*Paper-Bill Spread*$_t$[d]	*Reports of Eased Credit*$_t$[e]
-4	-0.31**	-0.80***	0.67***	0.06	0.19*	-0.16
-3	-0.36**	-0.74***	0.70***	0.01	0.13	-0.15
-2	-0.35**	-0.69***	0.60***	0.17	0.21**	-0.33***
-1	-0.17	-0.56***	0.38**	0.32***	0.39***	-0.40***
0	0.20	-0.29*	0.38**	0.14	0.49***	-0.31***
1	0.42***	-0.24	0.14	0.15	0.33***	-0.15
2	0.27*	-0.16	0.08	0.32***	0.19**	-0.08
3	0.21	-0.02	0.02	0.40***	0.15	-0.02
4	0.34**	0.10	0.15	0.22*	0.16	0.10
Memo:						
Sample period	1973:2-1983:4	1990:2-1998:4	1990:2-1998:4	1973:2-1983:4 1990:2-1998:4	1967:1-1983:4 1990:2-1998:4	1974:1-1983:4 1990:2-1998:4

Source: Data sources are in Table A2.

[a] Net percentage of domestic banks reporting a tightening of standards for commercial and industrial loans.

[b] Spread of the commercial loan rate over the federal funds rate.

[c] Mix = 100*(Nonfinancial CP outstanding/(Nonfinancial CP + C&I bank loans)).

[d] Spread = (Nonfinancial CP interest rate)-(T-bill interest rate). The spread was computed using six-month rates until 1971 and three-month rates from 1971 to 1998.

[e] Net percentage of small firms reporting "easier" credit from the previous quarter.

*Statistically significant at the 10 percent level.

**Statistically significant at the 5 percent level.

***Statistically significant at the 1 percent level.

Chart 3
Changes in Commercial Credit Standards and Credit Demand

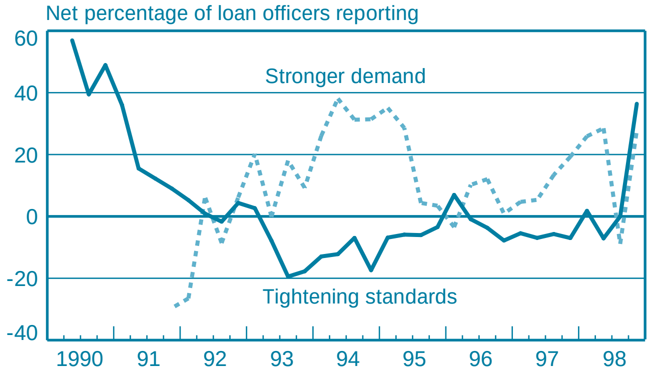

Source: Board of Governors of the Federal Reserve System, Senior Loan Officer Opinion Survey.

supply of bank credit. The problem with these pairwise correlations is that they fail to control for a second important determinant of loan growth: demand. Consider the 1990-91 period: did lending contract during that period because bankers were tightening standards or simply because the recession over that period slowed the demand for loans?

In fact, loan officers do tend to report weaker demand for commercial loans at the same time that they report tightening standards (Chart 3).[8] This correlation makes sense, especially when we consider the business cycle. Credit demand falls during contractions, at the same time that cautious bankers become less willing to lend. Firms demand more credit during expansions, and the good times may also make banks more willing to lend. To isolate the link between lending and credit standards, we use a regression equation to control for these multiple interactions between economic activity and the supply and demand for credit.

Regression Results

We estimate a loan growth equation of the following form:

$$(1) \quad \Delta Loans_t = \alpha + \beta Standards_{t-1} + \gamma \Delta D_{t-1} + \varepsilon_t .$$

The dependent variable, $\Delta Loans_t$, is growth in commercial and industrial loans at U.S.-chartered banks over quarter *t*, expressed at an annual percentage rate.[9] $Standards_{t-1}$ is the net percentage of loan officers reporting tightening standards for approving commercial loans to large and medium-sized firms. Both series are plotted in Chart 2. ΔD_{t-1} is a vector of other variables that may influence loan growth: the lagged dependent variable ($\Delta Loans_{t-1}$), lagged real output growth ($\Delta Output_{t-1}$), and the lagged value of the commercial loan spread ($Loan\ spread_{t-1}$). The summary statistics on these and other variables that we use later are reported in the appendix (Table A1). From 1992 onward, we also have data on the net percentage of loan officers reporting strengthening demand for commercial loans in the previous quarter ($Demand_{t-1}$), the series plotted in Chart 3. Note that in the equation we are regressing loan growth in one quarter on the values of the right-hand-side variables in the previous quarter.

The economy does seem to grow more slowly during periods in which bankers tighten credit standards; four of the past five recessions were preceded by sharply tighter standards.

Since the data are available over relatively short subperiods, to conserve degrees of freedom we use only one lag of each variable. The short lag length is actually conservative, since the correlation between loan growth and standards is higher at longer lags (Table 1).

The equation with demand and standards on the right-hand side is the complete specification since, in theory at least, changes in loan growth should reflect either changes in demand or changes in standards, that is, supply, or both. The question here is whether our survey measure of standards is a reasonable proxy for changes in standards and supply across the economy. If the actions of the loan officers surveyed match their words, and if loan officers across the country act likewise, reports of tighter standards should lead to slower loan growth. In terms of equation 1, we expect $\beta < 0$. We would expect a positive sign on the lagged values of loans, output, and demand to the extent that these variables are good proxies for loan demand. Lagged loan spreads could enter with either sign, depending on whether they reflect mostly supply-side or mostly demand-side factors.

Regression estimates over three distinct sample periods are reported in Table 2. Lagged loan growth is positive and significant in every specification, and the large coefficient indicates considerable momentum in the lending process. Lagged output growth is insignificant over the 1990s, but is significantly negative over 1973 to first-quarter 1984, contrary to expectations.[10] The demand variable enters negatively over the 1990s sample period, but it is statistically insignificant. The loan spread is highly significant and enters negatively in every regression, suggesting that increases in this spread are due more to reductions in loan supply than to increases in the demand for loans.

What the regressions show, most importantly, is that the reports of tighter standards by loan officers are still associated with slower loan growth, even after controlling for other factors that affect loan growth. The standards variable enters negatively, as expected, and is significant at the 5 percent level or lower over every period and specification.

With hindsight, the strong connection between credit standards and loan growth is not really surprising. Loan growth *should* depend on the supply of credit, and we suspect that the supply of loans depends on credit standards. More surprising is the result that standards matter even after we control for changes in loan rates or spreads. This finding supports the notion that bankers allocate loans not by simply raising and lowering rates, but by tightening and loosening other nonprice terms as well—the rationing concept that may have motivated the survey in the first place.

Table 2

Commercial Loan Growth and Credit Standards: Regression Equations over Various Periods

	(1) 1990:3 – 1998:4	(2) 1992:1 – 1998:4	(3) 1973:3 – 1984:1
C	9.203*	10.879**	12.618***
	(5.167)	(4.282)	(3.618)
Δ $Loans_{t-1}$	0.803***	0.798***	0.523***
	(0.080)	(0.112)	(0.102)
Δ $Output_{t-1}$	0.280	—	-0.476**
	(0.307)		(0.219)
$Demand_{t-1}$	—	-0.029	—
		(0.048)	
$Standards_{t-1}$	-0.075**	-0.225***	-0.237***
	(0.036)	(0.067)	(0.057)
$Loan\ spread_{t-1}$	-4.641*	-5.332**	-3.058**
	(2.577)	(2.309)	(1.276)
Adjusted R^2	0.921	0.902	0.595
BG test	1.007	0.942	2.478
Observations	34	28	43

Source: Data sources are in Table A2.

Notes: Reported are regression coefficients and standard errors (in parentheses). The dependent variable is quarterly growth (at an annual rate) in commercial and industrial loans at U.S. banks. In columns 1 and 2, *Standards* is the net percentage of senior loan officers reporting tighter standards on large firms. *Demand* is the net percentage reporting stronger demand by large firms. *Loan spread* is the spread of the average commercial and industrial loan rate over the federal funds rate. For the earlier period (column 3), *Standards* is the average of the net percentage reporting tighter standards for making loans at the prime rate or above. *Loan spread* is the spread of the prime rate over the federal funds rate. The BG test is the Breusch-Godfrey test for first-order autocorrelation. The test statistic is distributed Chi-squared with one degree of freedom.

*Statistically significant at the 10 percent level.

**Statistically significant at the 5 percent level.

***Statistically significant at the 1 percent level.

Credit Standards and Economic Activity

Credit standards may be linked to economic activity for either of two reasons. To the extent that credit availability depends on lenders' standards, a tightening of standards should cause a decline in spending by firms that depend on banks for credit. Tighter standards may also signal other disturbances that cause the economy to slow: lenders may batten down the hatches ahead of the storm. The causal impact of a change in standards and the signal provided by the change both imply a negative correlation between standards and economic activity.

The economy does seem to grow more slowly during periods in which bankers tighten credit standards; four of the past five recessions were preceded by sharply tighter standards (Chart 4). The exception was the 1981-82 recession. Loan officers were loosening standards when that recession began, but they quickly tightened as the economy contracted.

Chart 4

Change in Commercial Credit Standards, GDP Growth, and Recessions

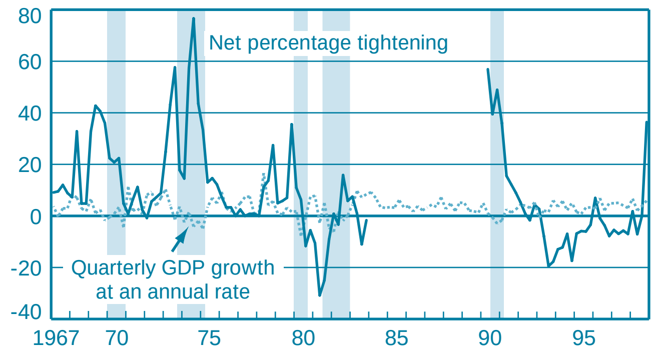

Sources: Board of Governors of the Federal Reserve System, Senior Loan Officer Opinion Survey; Bureau of Economic Analysis, Survey of Current Business.

Note: The shaded areas indicate periods designated national recessions by the NBER.

Table 3 reports regression equations relating quarterly growth in real GDP to its own lagged value and several other variables. The question is whether $Standards_{t-1}$ provides any additional information, given these other variables. Separate regressions were estimated over the early years of the survey (1967-84) and the 1990s.[11]

The results for the early years show that standards help considerably in predicting GDP growth. Lagged growth by itself explains only about 5 percent of the variation in current growth. Adding $Standards_{t-1}$ to the equation more than doubles the R^2 (column 2), and $Standards_{t-1}$ is statistically significant between the 1 and 5 percent levels. $Standards_{t-1}$ remains significant even when the equation includes two additional variables that have proved to be powerful forecasters: the federal funds rate and the spread between rates on commercial paper and Treasury bills (column 3).[12] Adding these variables to the equation more than triples the adjusted R^2. Nevertheless, knowing whether bank lenders recently tightened or loosened their standards for commercial credit still helps in predicting GDP growth.

These results are all at least as strong over the more recent period: third-quarter 1990 to fourth-quarter 1998. $Standards_{t-1}$ is more significant over this period than it is over the earlier period, and the adjusted R^2 for the equation with $Standards_{t-1}$ is twice as high as it is for the regression without $Standards_{t-1}$ (Table 3, columns 4 and 5). $Standards_{t-1}$ remains significant at 1 percent or better when we add lags of the funds rate and the paper-bill spread (column 3).[13] $Standards_{t-1}$ wins the horse race over the more recent period—in fact, neither the funds rate nor the spread is

Table 3

The Link between Commercial Credit Standards and Output Growth

	1967:2 – 1984:1			1990:3 – 1998:4		
	(1)	(2)	(3)	(4)	(5)	(6)
C	2.255***	3.212***	10.666***	1.439**	2.748***	2.920**
	(0.658)	(0.893)	(1.875)	(0.630)	(0.481)	(1.372)
Δ *Output*$_{t-1}$	0.259**	0.199*	-0.164	0.478**	0.027	0.032
	(0.100)	(0.103)	(0.110)	(0.184)	(0.147)	(0.163)
$Standards_{t-1}$	—	-0.067**	-0.072**	—	-0.084***	-0.080***
		(0.030)	(0.035)		(0.017)	(0.027)
Federal funds rate$_{t-1}$	—	—	0.262	—	—	-0.111
			(0.384)			(0.470)
Federal funds rate$_{t-2}$	—	—	-0.828*	—	—	—
			(0.433)			
Paper-bill spread$_{t-1}$	—	—	-2.029	—	—	0.874
			(1.351)			(3.719)
Adjusted R^2	0.052	0.102	0.371	0.190	0.427	0.389
BG test	0.462	0.288	0.235	1.855	0.766	0.719
Observations	69	68	68	34	34	34

Source: Data sources are in Table A2.

Notes: Reported are regression coefficients and standard errors (in parentheses). The dependent variable is quarterly growth (at an annual rate) in real GDP. *Paper-bill spread* is the spread between interest rates on commercial paper and Treasury bills. Specification tests called for the second lag of the federal funds rate in the column 3 regression. *Standards* is the net percentage of senior loan officers reporting tighter standards on large firms. For the earlier period (columns 1, 2, and 3), *Standards* is the average of the net percentage reporting tighter standards for making loans at the prime rate or above. The BG test is the Breusch-Godfrey test for first-order autocorrelation. The test statistic is distributed Chi-squared with one degree of freedom.

*Statistically significant at the 10 percent level.

**Statistically significant at the 5 percent level.

***Statistically significant at the 1 percent level.

individually significant, and the adjusted R^2 with these variables included in the regression is lower.[14] The fact that standards matter at least as much or more over this period provides no support for the view that the role of banks in economic activity diminished over the 1990s.

Commercial Credit Standards and Business Activity

The regressions in Table 4 show that $Standards_{t-1}$ helps to predict three narrower measures of business activity as well: investment in producers' durables, the change in inventory investment, and industrial production. $Standards_{t-1}$ is significant across the board, with or without the funds rate and the paper-bill spread. The adjusted R^2 actually falls when these additional variables are included, suggesting that they add little information beyond that already contained in standards.

The connection between standards and inventories is especially notable, as inventory investment is notoriously unpredictable (Blinder and Maccini 1991). Inventory investment should vary with interest rates, but researchers have never found a strong link between them. This missing link has vexed business-cycle researchers because fluctuations in

Table 4

Commercial Credit Standards and Business Activity

	Dependent Variable					
	Investment in Producers' Durable Equipment		Change in Business Inventories		Industrial Production	
	(1)	(2)	(3)	(4)	(5)	(6)
C	11.774***	10.657*	16.346***	10.887	4.285***	5.629**
	(2.466)	(6.093)	(4.005)	(13.693)	(0.879)	(2.166)
$\Delta\ Depvar_{t-1}$	-0.158	-0.182	0.518***	0.511***	0.130	0.144
	(0.218)	(0.231)	(0.123)	(0.137)	(0.144)	(0.146)
$Standards_{t-1}$	-0.327***	-0.357**	-0.653***	-0.689**	-0.178***	-0.159***
	(0.116)	(0.149)	(0.180)	(0.287)	(0.042)	(0.046)
$Federal\ funds\ rate_{t-1}$	—	1.162	—	-2.237	—	-0.332
		(1.532)		(3.499)		(0.479)
$Paper\text{-}bill\ spread_{t-1}$	—	-10.565	—	40.148	—	0.484
		(13.371)		(23.694)		(5.115)
Adjusted R^2	0.254	0.219	0.611	0.609	0.491	0.461
BG test	2.251	1.366	0.974	1.954	0.637	1.229
Observations	34	34	34	34	33	33

Source: Data sources are in Table A2.

Notes: Reported are regression coefficients and standard errors (in parentheses). The dependent variables *investment in producers' durable equipment* and *industrial production* are quarterly growth rates (at annual rates); *change in business inventories* is a quarterly change (at an annual rate). The equations in columns 5 and 6 include a second lag of the dependent variable (not reported). *Paper-bill spread* is the spread between interest rates on commercial paper and Treasury bills. *Standards* is the net percentage of senior loan officers reporting tighter standards on large firms. The BG test is the Breusch-Godfrey test for first-order autocorrelation. The test statistic is distributed Chi-squared with one degree of freedom. All equations are estimated from third-quarter 1990 to fourth-quarter 1998.

*Statistically significant at the 10 percent level.
**Statistically significant at the 5 percent level.
***Statistically significant at the 1 percent level.

inventory investment account for a disproportionate share of fluctuations in GDP. In their study of the 1981-82 recession, Kashyap, Lamont, and Stein (1994) suggest that "quantity rationing" by banks—that is, the allocation of loans through non-interest-rate terms—could explain the missing link. They find that during that recession the firms presumably most dependent on banks for credit—those without a lot of cash or a public bond rating—cut their inventory investment by substantially more than did the less bank-dependent firms. In the end, the authors attribute the inventory recession to tight "credit conditions" at banks. Our results seem roughly consistent.[15]

Vector Autoregression Analysis

In contrast to the single equations estimated above, a VAR is a system of equations that lets us better control for the feedback between current and past levels of output, lending, and credit standards. Controlling for this feedback is crucial. Suppose lenders tighten standards in response to weakness in the economy, both past and present. It is really weak output driving up standards in that case, not the other way around. To crack this chicken-and-egg problem, a VAR lets every variable in the system depend on past values of itself and every other variable in the system. Given estimates of these interactions, we can identify the changes in credit standards that were *not* predicted by the other variables in the system. Running these unpredicted "shocks" back through the system traces the impact of a shock to standards to all the other variables, and vice versa.

The core of our VAR, which is relatively standard, includes (in order): log real GDP, log GDP deflator, log commodity prices, and the federal funds rate. Note that these four variables make up a more or less complete model of the economy. Apart from output (real GDP) and prices (GDP deflator), we include commodity prices as a proxy for supply-side disturbances (like oil shocks). The federal funds rate provides a measure of interest rates that is tied particularly closely to the stance of monetary policy.[16] We tack two more variables to the end of the system: the log of commercial bank loans and the change in commercial credit standards. Placing standards last in the VAR tends to minimize their impact on output and the other variables. The VAR includes four lags of each variable above. We estimate the model jointly over the two periods in which we have data on loans and credit standards: first-quarter 1974 to fourth-quarter 1984 and second-quarter 1990 to fourth-quarter 1998.

Chart 5 shows the shock to standards and the subsequent dynamics. The initial increase in the net percentage of loan officers reporting tightening standards is about 8 percentage points. Lenders continue to tighten (at diminishing rates) for about a year after the initial shock. Nearly two years pass before lenders ease standards significantly, and the easing commencing then is relatively gradual compared with the tightening over preceding quarters.[17] Credit crunches appear more abrupt than credit expansions.

Chart 6 shows how shocks to other variables in the system affect credit standards. Standards seem largely independent of the other variables in the system. Shocks to commodity prices and the federal funds rate cause some tightening of standards, but the impact is short and barely significant. These findings suggest that lenders set their standards based largely on their own lending capacity and on their expectations, so that standards appear to be relatively exogenous to the other macroeconomic variables included in the system.

Conclusion

Off and on since 1967, the Federal Reserve has surveyed loan officers at a small sample of large banks about their commercial credit standards. The idea behind the survey is that the availability of bank credit depends not just on interest rates, but on credit standards as well. Notwithstanding the small and changing sample, the checkered pattern of questions, and the sometimes curious responses of lenders, the reports are informative. The changes in standards that they report help to predict both commercial bank lending and GDP, even after controlling for past economic conditions and interest rates. Standards matter even in the 1990s, when capital markets were supposed to have eclipsed the role of banks in the economy. Changes in standards also help to predict narrower measures of business activity, where commercial credit availability from banks seems most crucial. The connection between bank standards and inventories is especially promising, because inventory investment is notoriously unpredictable and heavily bank dependent.

A shock to credit standards and its aftermath very much resemble a "credit crunch." Loan officers tighten standards very sharply for a few quarters, but ease up only gradually: two to three years pass before standards are back to their initial level. Commercial loans at banks plummet immediately after the tightening in standards and continue to fall until lenders ease up. Output falls as well, and the federal funds rate, which we identify with the stance of monetary policy, is lowered. All in all, listening to loan officers tells us quite a lot.

Chart 5

A Shock to Credit Standards and the Response of Other Variables

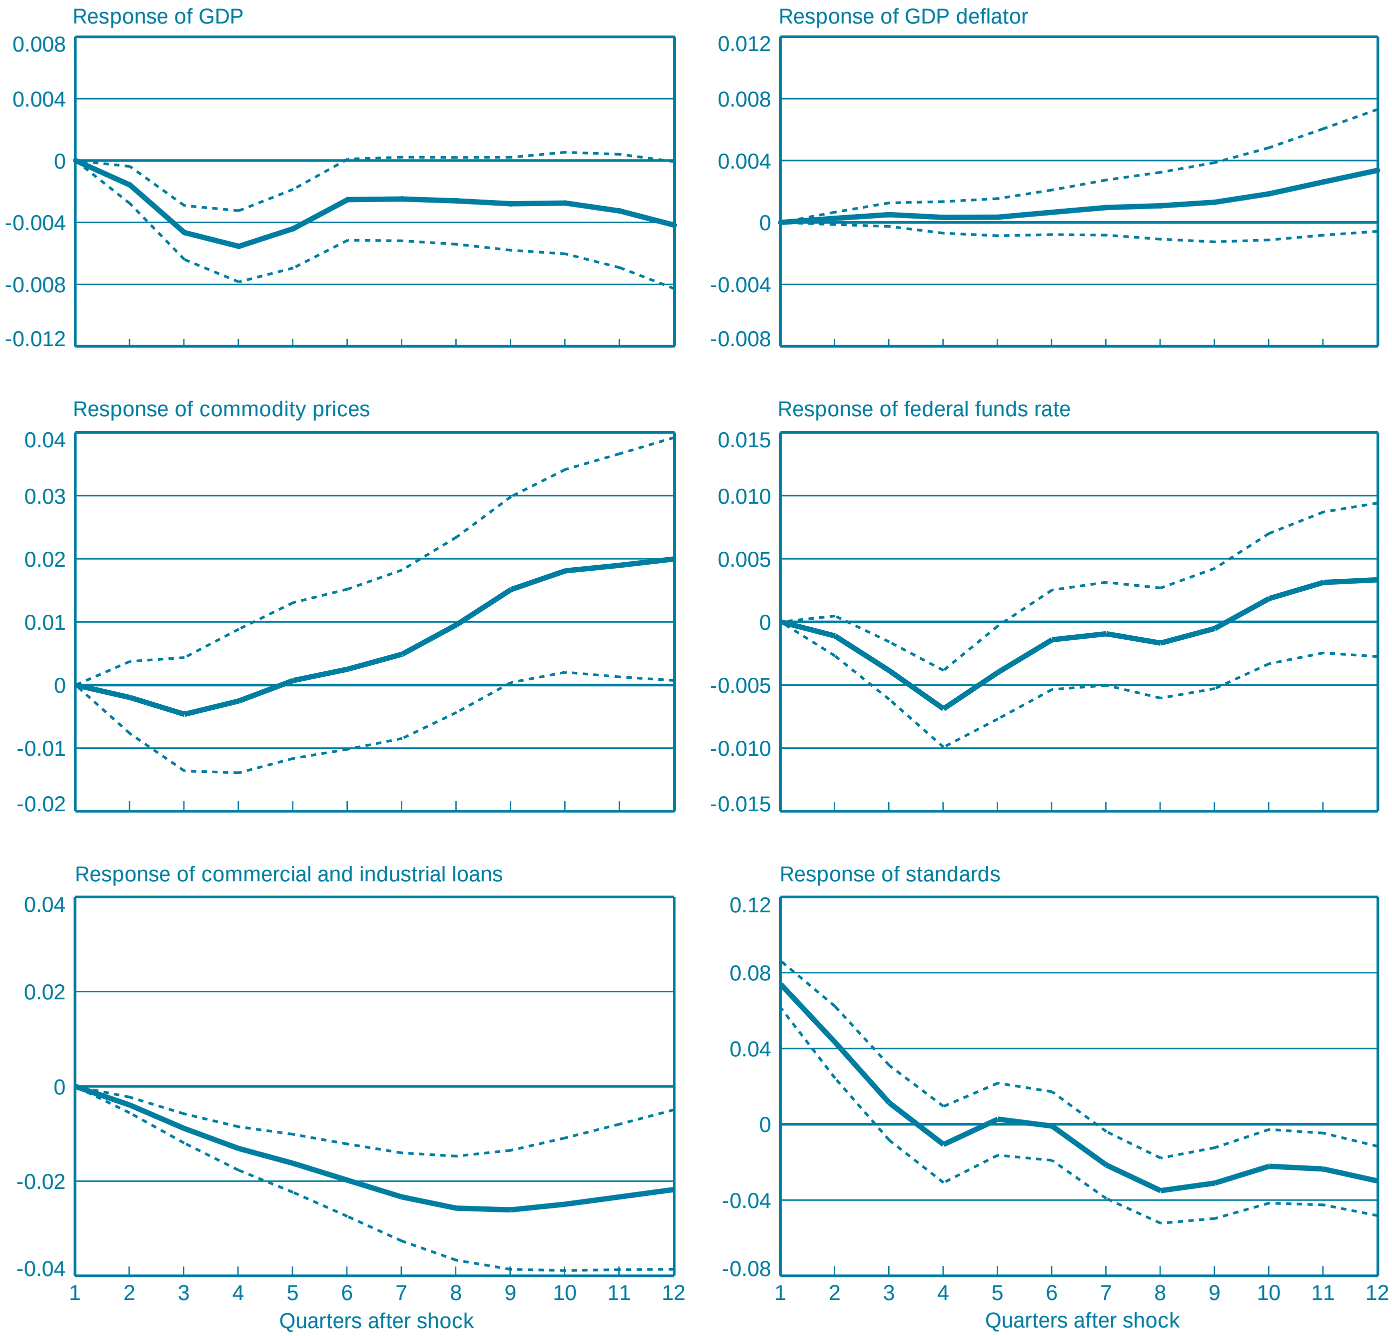

Source: Authors' VAR analysis using data sources in Table A2.

Notes: Plotted is the impact on each variable of a one-standard-deviation tightening in commercial credit standards by banks. The dashed lines indicate 95 percent confidence levels.

Chart 6

Response of Credit Standards to Shocks in Other Variables

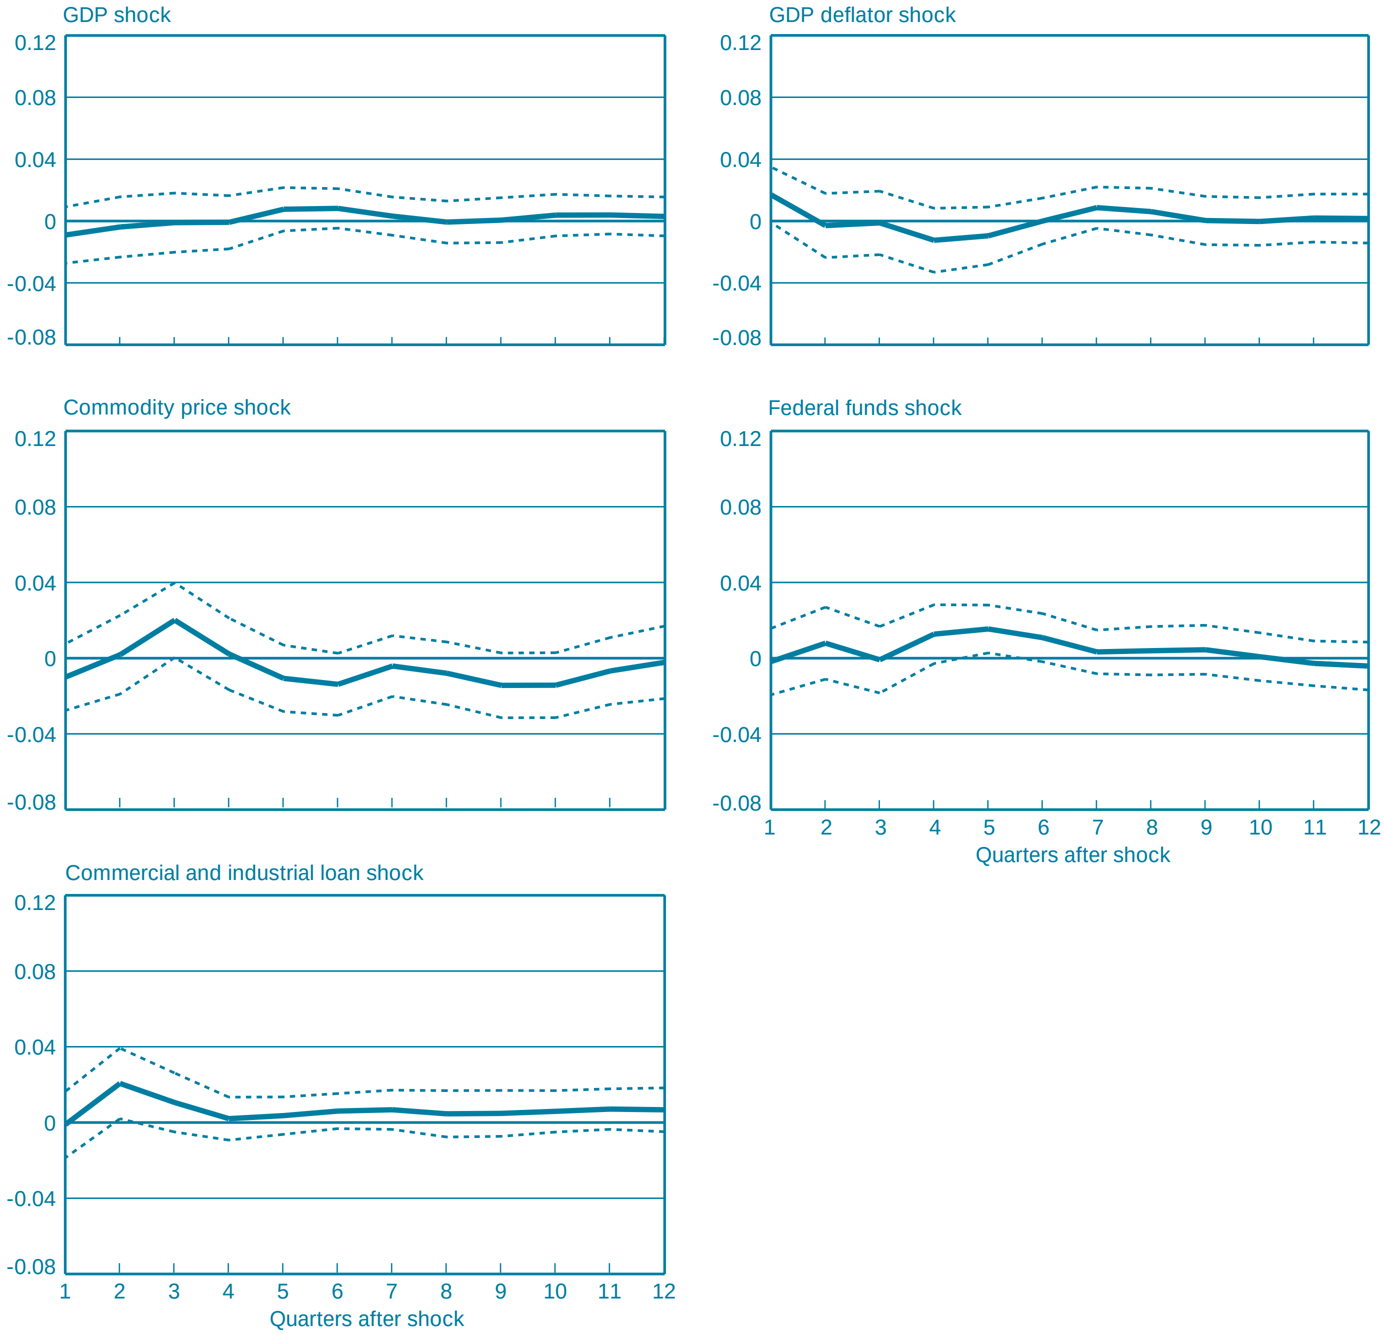

Source: Authors' VAR analysis using data sources in Table A2.

Notes: Plotted is the impact of a one-standard-deviation shock in each variable on commercial credit standards by banks. The dashed lines indicate 95 percent confidence levels.

Appendix: Summary Statistics and Data Sources

Table A1
Summary Statistics

Variable	Definition	Observations	Mean	Standard Deviation	Minimum	Maximum
Δ *Loans*	Quarterly growth in commercial and industrial loans (annual rate, 1973:3-1984:1)	44	9.959	8.479	-9.233	30.128
Δ *Loans*	(1990:3-1998:4)	34	4.253	8.775	-10.964	18.230
Δ *Output*	Quarterly growth in real GDP (annual rate, 1967:2-1984:1)	69	3.025	4.703	-9.300	16.100
Δ *Output*	(1990:3-1998:4)	34	2.626	2.276	-4.000	6.100
Δ *Industrial production*	Quarterly growth in industrial production (annual rate)	34	3.504	3.612	-8.415	9.641
Δ *Business inventories*	Quarterly change in business inventories (annual rate)	34	30.803	28.593	-27.800	91.400
Δ *Producers' durable*	Quarterly growth in investment in producers' durable equipment (annual rate)	34	9.646	9.382	-14.957	34.214
Demand	Net percentage of domestic banks reporting stronger demand over the previous quarter	28	12.809	15.452	-26.500	38.100
Standards	Net percentage of domestic banks reporting tighter standards over the previous quarter. During 1978-83, *Standards* is computed by averaging changes in credit standards on loans at prime and loans above prime (1967:2-1983:4)	67	12.175	18.792	-30.833	76.613
Standards	(1990:3-1998:4)	34	1.656	16.416	-19.450	48.900
Federal funds rate	(1967:2-1984:1)	69	8.352	3.610	3.550	17.790
Federal funds rate	(1990:3-1998:4)	34	4.984	1.294	2.990	8.160
Paper-bill spread	Spread of the nonfinancial commercial paper rate over the secondary market T-bill rate. The spread was computed using six-month rates until 1971 and three-month rates from 1971 to 1998 (1967:2-1984:1)	69	0.842	0.617	0.030	3.510
Paper-bill spread	(1990:3-1998:4)	34	0.435	0.155	0.180	0.910
Loan spread	Spread of the prime rate over the federal funds rate (1973:3-1984:1)	44	1.554	0.994	-1.440	3.720
Loan spread	Spread of the commercial and industrial loan rate over the federal funds rate (1990:3-1998:4)	34	1.836	0.224	1.440	2.320
Δ *Paper-loan mix*	Quarterly growth in the ratio of nonfinancial commercial paper outstanding to the sum of nonfinancial commercial paper outstanding and commercial and industrial bank loans (1973:2-1983:4)	39	11.076	29.200	-45.397	81.706
Reports of eased credit	Net percentage of small firms that borrow money at least once every three months reporting "easier credit" compared with three months ago (1974:1-1983:4)	35	-7.629	6.174	-28.000	1.000
Δ *Paper-loan mix*	(1990:2-1998:4)	31	1.979	10.629	-18.471	24.860
Reports of eased credit	(1990:2-1998:4)	31	-5.452	2.815	-11.000	-1.000

Source: Data sources are in Table A2.

Table A2
Data Sources

Variable	Definition	Data Source
Δ *Loans*	Quarterly growth in commercial and industrial loans (annual rate)	Federal Reserve Board Statistical Release H.8: Assets and Liabilities of Commercial Banks in the U.S. (seasonally adjusted and break-adjusted, in millions of dollars)
Δ *Output*	Quarterly growth in real GDP (annual rate)	Bureau of Economic Analysis, Survey of Current Business (seasonally adjusted annual rate, in billions of dollars)
Δ *Industrial production*	Quarterly growth in industrial production (annual rate)	Federal Reserve Board Statistical Release G.17: Industrial Production and Capacity Utilization (seasonally adjusted, 1992=100)
Δ *Business inventories*	Quarterly change in business inventories (annual rate)[a]	Bureau of Economic Analysis, Survey of Current Business (seasonally adjusted annual rate, in billions of dollars)
Δ *Producers' durable*	Quarterly growth in investment in producers' durable equipment (annual rate)[b]	Bureau of Economic Analysis, Survey of Current Business (seasonally adjusted annual rate, in billions of dollars)
Demand	Net percentage of domestic banks reporting stronger demand over the previous quarter	Board of Governors of the Federal Reserve System, Senior Loan Officer Opinion Survey
Standards	Net percentage of domestic banks reporting tighter standards over the previous quarter	Board of Governors of the Federal Reserve System, Senior Loan Officer Opinion Survey
Federal funds rate	Effective overnight interbank lending rate	Federal Reserve Board Statistical Release H.15: Selected Interest Rates
Paper-bill spread	Spread of the nonfinancial commercial paper rate over the secondary market T-bill rate. The spread was computed using six-month rates until 1971 and three-month rates from 1971 to 1998	Federal Reserve Board Statistical Release H.15: Selected Interest Rates
Loan spread	Spread of the prime rate over the federal funds rate (1973:3-1984:1)	Federal Reserve Board Statistical Release H.15: Selected Interest Rates
Loan spread	Spread of the commercial and industrial loan rate over the federal funds rate (1990:3-1998:4)	Federal Reserve Board Statistical Release E.2: Survey of Terms of Business Lending (for commercial and industrial loan rate). See above for federal funds rate
Δ *Paper-loan mix*	Quarterly growth in the ratio of nonfinancial commercial paper outstanding to the sum of nonfinancial commercial paper outstanding and commercial and industrial bank loans	Federal Reserve Board Statistical Release: Commercial Paper (for nonfinancial commercial paper outstanding). See above for commercial and industrial bank loans
Reports of eased credit	Net percentage of small firms that borrow money at least once every three months reporting "easier credit" compared with three months ago	National Federation of Independent Business, Small Business Economic Survey

[a]This variable is now referred to as "private inventory" in the source data.
[b]This variable is now referred to as investment in "equipment and software" in the source data.

Endnotes

1. Academics interested in the credit effects of monetary policy—the theory that changes in policy affect spending partly through changes in the supply of bank loans—will want to know if the reports on bank credit standards are a reliable proxy for bank loan supply.

2. See Schreft and Owens (1991) for more on the history of and revisions to the survey. For current and recent surveys, see http://www.federalreserve.gov/boarddocs/snloansurvey.

3. The Survey of Terms of Business Lending collects quantitative information on commercial loan rates and other lending terms at banks.

4. Weighting the responses over the 1990s by the extent of change ("somewhat," versus "considerably") did not change the picture or the results, nor did using a diffusion index. Integrating the changes reported by lenders over time did not work as well as any of the other measures.

5. This apparent bias toward reporting tightening in these early years could reflect bankers reporting standards relative to some long-term notion. Alternatively, bankers may not have reported easier standards for fear of scrutiny by regulators. Bankers need not have feared the regulator's club, however, since the responses of individual bankers are viewed as highly confidential and would not be shared with supervisory personnel except under extreme circumstances.

6. The availability doctrine has waned since the deregulation of interest rates and the accord that relieved the Federal Reserve of an obligation to support bond prices. See Blanchard and Fischer (1989) for references to the availability doctrine (p. 486) and quantity rationing (pp. 479-88, 492-3).

7. The peak correlation between loan growth and standards in both periods is at six quarters (not shown in the table).

8. Since 1992, loan officers have been asked how the demand for commercial and industrial loans has changed over the preceding three months (apart from normal seasonal variation). The multiple-choice answers enable them to identify demand as substantially stronger, moderately stronger, about the same, moderately weaker, or substantially weaker.

9. We include only loans at U.S. banks, as the survey responses discussed above were from loan officers at domestic banks. Loan officers at branches and agencies of foreign banks are questioned separately. We use nominal loan growth in the analysis as a proxy for what we would like to use—the real value of new credit extensions, which is unavailable. The results are quite similar when real loan growth is used as the proxy instead. For a discussion of the advantages and disadvantages of each measure, see Bernanke and Lown (1991, p. 209).

10. The negative sign over this period suggests a countercyclical demand for bank credit. Although on the surface this finding appears to be contrary to expectations, previous researchers have obtained similar results while exploring other issues, arguing that firms may have a greater need for financing as the economy begins to slow (see, for example, Bernanke and Blinder [1992]).

11. The regressions for the earlier period go back farther than the loan growth regressions because we have a longer time series on GDP growth.

12. Positive shocks to the funds rate are thought to reflect tighter monetary policy (Bernanke and Blinder 1992), while a higher paper-bill spread may signal policy shocks as well as other, adverse shocks to financial markets (Friedman and Kuttner 1992). Specification tests called for the second lag of the funds rate.

13. The second lag of the funds rate was insignificant and was not necessary to reduce autocorrelation, so we dropped it from the regression.

14. Friedman and Kuttner (1998) also found that the spread did not forecast well in the 1990s. In addition, when the regression is extended to include the mix variable—the ratio of commercial paper to commercial paper plus bank loans—this variable is not significant. However, the mix variable is significant in explaining GDP growth over the earlier time period.

15. Eckstein and Sinai (1986) go further. They contend that all six of the recessions between 1957 and 1982 were caused by "credit crunches."

16. See Christiano, Eichenbaum, and Evans (1996) and Bernanke and Mihov (1998).

17. These fluctuations make sense, as lenders are reporting changes in standards; a change in one direction eventually requires an opposite and equal change to return to the initial level. Lenders seem to jerk the tail hard, and they relax their grip very gradually.

References

Bernanke, Ben S., and Alan S. Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." American Economic Review 82, no. 4 (September): 901-21.

Bernanke, Ben S., and Mark Gertler. 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission." Journal of Economic Perspectives 9, no. 4 (fall): 27-48.

Bernanke, Ben S., and Cara S. Lown. 1991. "The Credit Crunch." Brookings Papers on Economic Activity, no. 2: 205-47.

Bernanke, Ben S., and Ilian Mihov. 1998. "Measuring Monetary Policy." Quarterly Journal of Monetary Policy 113 (August): 869-902.

Blanchard, Olivier, and Stanley Fischer. 1989. Lectures in Macroeconomics. Cambridge, Mass.: MIT Press.

Blinder, Alan S., and Alan Maccini. 1991. "Taking Stock: A Critical Assessment of Recent Research on Inventories." Journal of Economic Perspectives 5, no. 1 (winter): 73-96.

Board of Governors of the Federal Reserve System. Various years. "Senior Loan Officer Opinion Survey on Bank Lending Practices." Washington, D.C.

Bram, Jason, and Sydney Ludvigson. 1998. "Does Consumer Confidence Forecast Household Expenditure? A Sentiment Index Horse Race." Federal Reserve Bank of New York Economic Policy Review 4, no. 2 (June): 59-78.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 1996. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." Review of Economics and Statistics 80, no. 1 (February): 16-34.

Duca, John, and Bonnie Garrett. 1995. "Credit Availability, Bank Consumer Lending, and Consumer Durables." Federal Reserve Bank of Dallas Working Paper no. 9514.

Eckstein, Otto, and Allen Sinai. 1986. "The Mechanisms of the Business Cycle in the Post-War Era." In Robert J. Gordon, ed., The American Business Cycle: Continuity and Change, 39-122. Chicago: University of Chicago Press.

Friedman, Benjamin M., and Kenneth N. Kuttner. 1992. "Money, Income, Prices, and Interest Rates." American Economic Review 82 (June): 472-92.

———. 1998. "Indicator Properties of the Paper-Bill Spread: Lessons from Recent Experience." Review of Economics and Statistics 80, no. 1 (February): 34-44.

Kashyap, Anil, Owen Lamont, and Jeremy Stein. 1994. "Credit Conditions and the Cyclical Behavior of Inventories." Quarterly Journal of Economics 109, no. 3 (August): 565-93.

Kashyap, Anil, Jeremy Stein, and David Wilcox. 1993. "Monetary Policy and Credit Conditions: Evidence from the Composition of External Finance." American Economic Review 83 (March): 78-99.

McCarthy, Jonathan. 1997. "Debt, Delinquencies, and Consumer Spending." Federal Reserve Bank of New York Current Issues in Economics and Finance 3, no. 3 (February).

Schreft, Stacey L., and Raymond E. Owens. 1991. "Survey Evidence of Tighter Credit Conditions: What Does It Mean?" Federal Reserve Bank of Richmond Economic Review 77, no. 2 (March-April): 29-34.

James McAndrews and Samira Rajan

The Timing and Funding of Fedwire Funds Transfers

- The dollar value of payments made over the Fedwire Funds Transfer service reaches its highest level between 4 p.m. and 5 p.m. each day.
- This peak in payment activity likely reflects efforts by banks to synchronize their outgoing payments with the large payment inflows they expect to receive in the late afternoon.
- By using the incoming transfers to fund outgoing payments, banks avoid the more costly alternatives of drawing down their account balances at the Federal Reserve or using overdraft credit.
- To support the banks' funding strategy, policymakers might establish formal "synchronization periods" and encourage banks to concentrate payments during these periods.
- The resulting increase in payment coordination could further reduce financing costs and minimize the number and duration of overdrafts.

The timing of payments across the Fedwire Funds Transfer service exhibits a regular pattern over the course of the day, with payment activity peaking in the late afternoon.[1] This pattern can be explained, in part, by the fact that banks derive benefits from coordinating the timing of their payment activity. Many payments made by banks during the day are offsetting. By synchronizing payments, banks can take advantage of incoming funds to make outgoing payments. The afternoon peak in activity reflects, to some extent, banks' coordination of payment timing in an attempt to tap this funding source.

A full explanation of the timing of funds transfers recognizes two factors that affect banks' intraday liquidity management. First, the timing of banks' payment activity reflects underlying customer demand. For example, settlement of financial transactions customarily takes place in the late afternoon, which tends to cause a demand for payments late in the day. Second, such timing also reflects a bank's response to customer demand for prompt payment. When responding to this demand, banks incur costs that take up expensive liquidity resources—either deposits at, or overdrafts from, the Federal Reserve System.

The liquidity cost of making a payment varies with the amount of coordination involved in payment timing. During periods of heavy payment traffic, a bank can, to a greater extent, fund an outgoing payment with incoming payments. Conversely, during off-peak times, a bank must rely more on account balances or overdrafts to fund payments, which increases the cost of making

James McAndrews is a research officer at the Federal Reserve Bank of New York. Samira Rajan, formerly an assistant economist at the Bank, is now a master's degree candidate in public policy at Harvard University's John F. Kennedy School of Government.

The authors thank Paolo Angelini, Shawn Cole, Ed Green, Darryll Hendricks, Jeff Stehm, and John Wenninger for helpful discussions, as well as two anonymous referees. They also thank the staffs of the Federal Reserve Bank of New York's Wholesale Payments Product Office and Credit Risk Management Division for providing data. The views expressed are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

a payment. As a result, banks are induced to time their payments to coincide with an activity peak, thereby reinforcing the peak. Such behavior can lead to the observed aggregate patterns during periods of light as well as heavy payment activity.

In this article, we measure banks' alternative funding sources for Fedwire funds transfers throughout the day, using a data set that includes all banks' Fedwire funds transfers and Federal Reserve System deposits. This approach allows us to gauge the importance of incoming payments as a source of funding. We find that incoming payments used by banks to offset outgoing payments that are entered within the same minute account for 25 percent of the value of these transfers during normal activity periods and as much as 40 percent during peak periods.

This level of payment coordination is impressive. However, economic analyses suggest that activity coordination by subjects in similar environments typically falls short of the level that would allow the subjects to benefit fully from such coordination.[2] Accordingly, with many thousands of banks participating in Fedwire, there is reason to believe that the banks would prefer even greater coordination of payment activity. Furthermore, greater synchronization of payments would lead to a decrease in daylight overdrafts extended by the central bank. With these considerations in mind, we also examine a policy that might allow banks to coordinate their payment activity even more effectively: the creation of activity periods that would serve as "focal times" for entering payments.

Our study proceeds as follows. In the next section, we review the intraday pattern of Fedwire funds transfers. We then offer possible explanations for this pattern by examining a model of payment timing. Next, we measure the different sources of Fedwire funding during the day. Finally, we discuss the implications of our findings for various policy issues, including the expansion of the operating hours of the Fedwire Funds Transfer service and the facilitation of payment coordination.

The Timing of Fedwire Payments

Fedwire is a real-time gross settlement (RTGS) system, in which payment requests are processed and settled by the Federal Reserve System as soon as they are initiated by banks. The number of funds transfers sent per minute varies over the course of the day in a fairly predictable pattern. The average minute-by-minute patterns of the number of

Chart 1

Average Daily Fedwire Funds Transactions by Time of Day

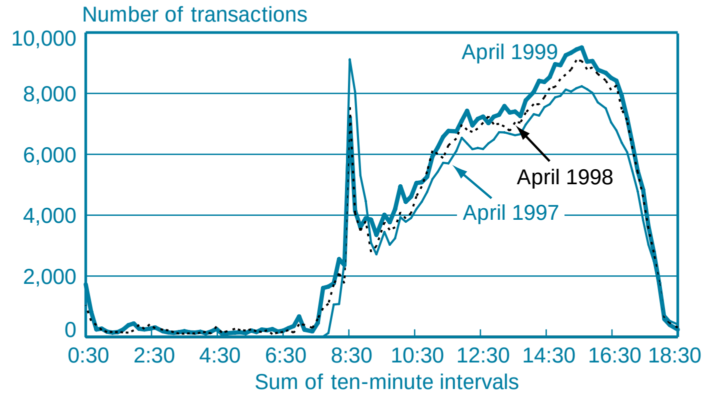

Source: Federal Reserve Bank of New York.

transfers for April 1997, 1998, and 1999 appear in Chart 1.[3] We see that a flurry of payments occurs at 8:30 a.m., which used to be the opening time for the Fedwire Funds Transfer service. After this flurry, the number of transfers sent per minute falls to a much lower level around 9:30 a.m. From that trough, the number of transfers grows fairly steadily throughout the day, reaching a peak from 2:30 to 4:30 in the afternoon. Transaction volume declines rapidly after 4:30 p.m. and approaches zero transfers per minute at the close of the service at 6:30 p.m.

The very largest payments are even more concentrated late in the day. The patterns of payments above the ninety-ninth percentile and those below it are shown in Chart 2. The chart indicates that for much of the day, there is a fairly low level of the largest-value payments. After a sharp increase following 4:30 p.m., once the Clearing House Interbank Payments System (CHIPS) has closed, the number of such payments falls considerably after 5:30 p.m.[4] Because the largest-value payments constitute in dollar terms the bulk of the value transferred by the Funds Transfer service, the patterns of these payments strongly influence the patterns of value exchanged per minute throughout the day. Chart 3 confirms that the value exchanged is more heavily concentrated in the period around 4:30 p.m. than is the number of funds transfers. Hence, in terms of the number of transfers, the dollar value of payments, and the number of largest-value payments, we can place the peak period for the Fedwire Funds Transfer service at 2:30 to 5:30 p.m., with the peak in value transfer occurring between 4 and 5 p.m.

Chart 2

Distribution of Fedwire Funds Transactions by Size of Payment

March 18, 1999

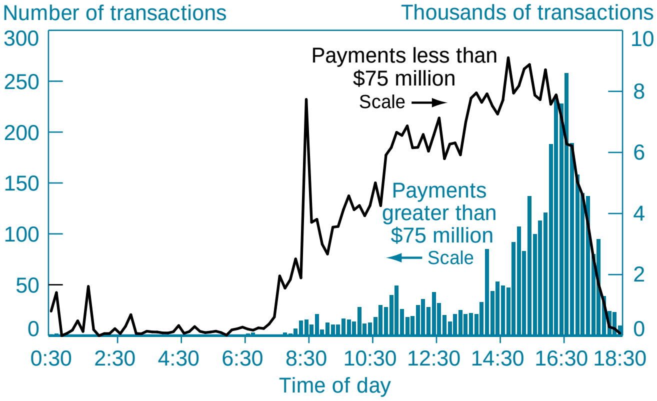

Source: Federal Reserve Bank of New York.

Note: $75 million is the ninety-ninth percentile for payment size on March 18, 1999.

Chart 3

Value of Average Daily Fedwire Funds Transactions by Time of Day

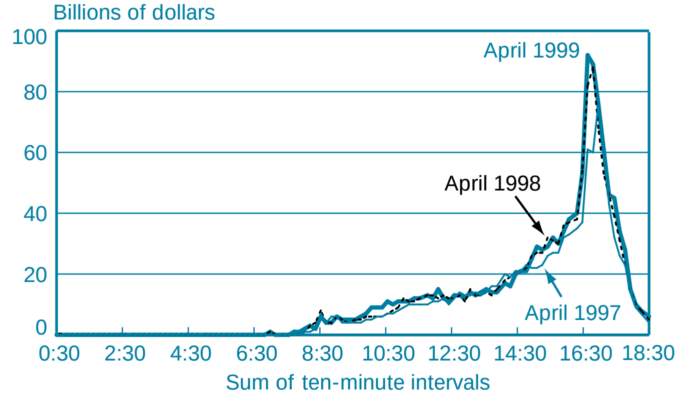

Source: Federal Reserve Bank of New York.

Durability of Payment Patterns

The Fedwire Funds Transfer service expanded its hours of operation from ten to eighteen hours in December 1997, so that it is now open from 12:30 a.m. until 6:30 p.m. eastern time. The change was made mainly to accommodate potential earlier settlement of foreign exchange trades. However, neither the timing of activity peaks nor the timing of any other payment patterns has been significantly affected by the lengthening of the Fedwire day.

The primary difference in payment patterns before and after December 1997 is the decrease in the number of payments made at 8:30 a.m.: there has been a decline equal to about 0.5 percent of the number of payments made between 8:30 and 9:30 a.m. Yet the percentage of funds transfers made between 12:30 a.m. and 8:30 a.m. remains roughly 1 percent, so some of the activity that took place at the 8:30 a.m. opening now takes place prior to that time. Overall, there has been a slight increase in the share of the value of payments completed by noon: for the period April-November 1998, 13.75 percent of the value was completed by that time, compared with 13.30 percent for the same period in 1997.

Some evidence suggests that the afternoon peak is higher today than it was prior to the implementation of the pricing of daylight overdrafts in 1987. Richards (1995), for example, notes that the share of value transferred by noon dropped about 5 percent in the year following the imposition of overdraft fees. In addition, a report by the Federal Reserve Bank of New York (1987) of large-value funds transfers during a single day in 1986 shows a less concentrated pattern of payment activity during the day. Using data from the report, we compare the percentage of the day's payments completed during various times of the day in 1986—prior to the imposition of overdraft fees—with the timing of payments in 1999 (Chart 4). We see that a larger share of the day's payments

Chart 4

Percentage of Dollar Volume Completed over Fedwire Funds Transfer Service

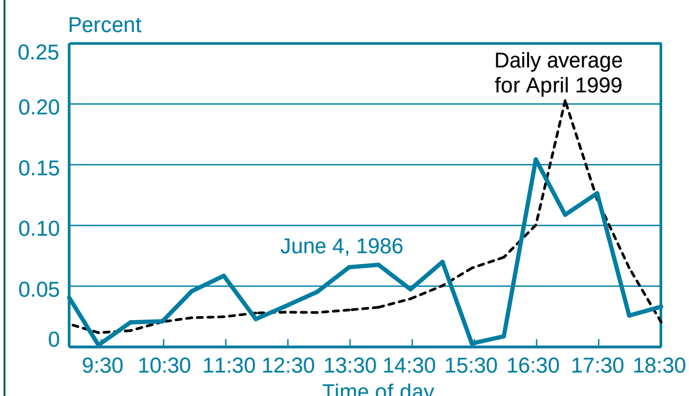

Source: Federal Reserve Bank of New York.

Notes: Data for 1986 payments include transfers in excess of $1 million. April 1999 data include transfers in excess of $1 million in 1986 dollars.

was completed earlier in the day in 1986 than in 1999 (although after 5:30 p.m., payments were made more quickly in 1999). At the same time, it is clear that in 1986 there was a substantial concentration of payments in the late afternoon. In short, the evidence confirms that payment traffic has long been characterized by a late afternoon peak.

Liquidity Externalities and the Coordination of Payments

Why are payments, especially the largest ones, concentrated in the late afternoon? As noted, this phenomenon may result from the timing of payment requests by customers and from the payments generated by the banks' own financial activity, which may be concentrated at the end of the day so that banks can settle financial market trades. In addition, banks themselves may time the submission of payments to coincide with the incoming payments that they expect to receive late in the day.[5] To explain this latter possibility, we first describe the funding sources for a bank's payments.

Sources of Payment Funding

Banks face a budget constraint when making payments: those made in a real-time gross settlement system run by a central bank typically are made by transfer of deposit account balances held at the central bank. Although a bank may have other assets, RTGS systems generally require that funds be in an account in the system at payment time, so that the systems do not have to rely on other forms of bank assets.[6] Account balances, then, serve as one source of funds by which a bank can make payments. However, account balances at central banks usually pay low interest rates, which creates an incentive for banks to minimize the amount of funds on account there.[7]

In the Fedwire Funds Transfer service, as in many other RTGS systems, banks transfer their account balances to make payments. Of course, one could reasonably ask, what if a bank's account balance falls to zero? For banks that are allowed to incur daylight overdrafts, that form of credit from the central bank is an additional source of funds that can be used for payments.[8] Finally, if a bank receives a payment from another system participant, that payment replenishes its account balance and allows the bank to make outgoing payments. A recent report on RTGS systems described these funding sources as: "(a) balances maintained on account with the central bank, (b) incoming transfers from other banks, [and] (c) credit extensions from the central bank."[9]

Before we discuss liquidity externalities in an RTGS system, we should look more closely at these three sources of funding. In particular, we consider incoming transfers from other banks. As noted earlier, when a bank exhausts its account balances at a particular time, it can make additional payments (without borrowing) if it receives incoming transfers from other banks. But because banks receive incoming payments and make outgoing payments throughout the day, it is important to examine the extent to which banks use incoming payments to fund the outgoing ones. We adopt the view that incoming payments arriving *at roughly the same time* as offsetting outgoing payments serve as a source of funding for the outgoing payments. Conversely, we also adopt the view that incoming transfers that "sit" in the receiver's account for a long period of time do not fund specific payments. If incoming payments sit in such an account, then we consider payments made long after the bank has received funds as being made by the transfer of balances maintained at the central bank.

Coordination of Payment Timing

It may be surprising to learn that the synchronous receipt of incoming transfers is a legitimate source of funding for a bank. This possibility exists whenever banks exchange payments throughout the day. For example, assume that Bank A owes Bank B $100, Bank B owes Bank C $75, and Bank C owes Bank A $50. If these payments took place at different times (in this sequence), Bank A's balance, for example, would fall by $100 in the first period and then would rise by $50 in the third. However, if these payments took place simultaneously, Bank A, which owes $100, would see its deposit balance fall by only $50 because it would receive a $50 payment from Bank C. In this way, the receipt of the incoming transfer from Bank C allows Bank A to "fund" its $100 payment—half with its own deposit balance and half with incoming funds. Although the end-of-day balance for all of the banks would be the same in either scenario, the uncoordinated timing of payments requires the banks either to incur larger overdrafts for a longer period or to maintain higher levels of

Why are payments, especially the largest ones, concentrated in the late afternoon?

The Effect of Synchronization on the Changes in Bank Balances

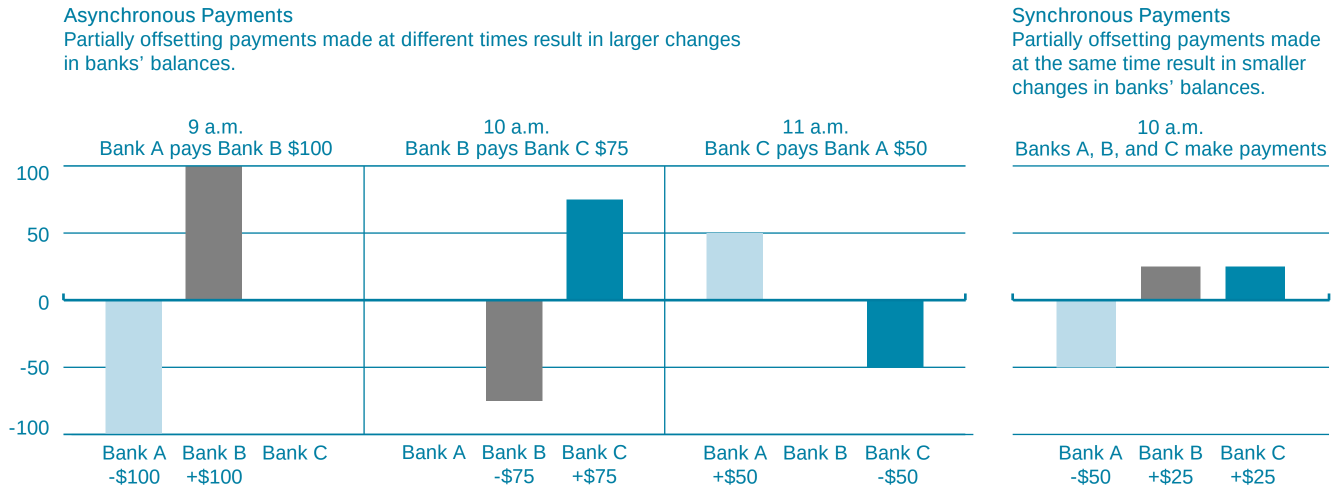

deposits to avoid overdrafts, relative to the synchronous timing of payment. If Banks A, B, and C could all coordinate the timing of their payments, each would have a lower funding cost than it would have had it been the first bank to pay. The exhibit presents the effect of synchronization on the change in the balances of the banks as they make these payments.

It is important to note that, in this regard, what is true for banks is also true for their customers. Banks impose limits on their customers' overdrafts and charge fees for the use of overdraft credit. Customers, like banks, try to seek the lowest cost funding for their payments. The timing of payments among bank customers therefore can lead to similar benefits for them. In particular, as customers receive payments, they can send payments using the incoming funds to avoid (or to limit the size of) overdrafts. In this way, payment coordination can reduce the customer's costs of making payments. We saw in Chart 4 that a noticeable peak existed in Fedwire payments prior to the imposition of overdraft fees by the Federal Reserve System. That pattern likely reflects, to some extent, the coordination of customer payments as well as the underlying timing of other late-in-the-day customer demand, such as for making settlement payments in the financial markets.

Although the exhibit illustrates the benefit of coordinating payment timing, the difficulty of achieving such a synchronized pattern is considerable because the timing of payments in some respects resembles a coordination game.[10] Banks can benefit by entering payments simultaneously to Fedwire, but they typically do not know when their counterparties might send offsetting payments. Hence, there is the potential for miscoordination. For instance, one bank enters a payment expecting to receive, but in fact does not get, an offsetting payment. Or two banks each delay sending their payments, as one expects the other to send its payment first. In these examples, coordination could be achieved simply by establishing conventions, such as sending payments regularly at a particular time, day after day. The 4:30 p.m. peak in payment activity might represent such a convention. When banks repeatedly send payments to one another day after day,

Synchronization of payments . . . allows banks to tap incoming transfers from other banks as a key source of funding.

the repetition in payment patterns can in some cases lead to successful coordination among a bank and its counterparties.

By concentrating payments in a short period, banks can work to resolve the coordination problem. They can then delay sending customer requests during the day, provided that the delay is not too costly, if they anticipate that other banks will make their payments later in the day (either because of customer requests or because other banks are also anticipating that their counterparties will send payments later in the day). As more and more banks behave in this manner, a peak period of payment activity will emerge during which banks receive payments more frequently than they do at other times. With these incoming payments, each receiving bank will see its Fedwire balance increase, enabling it to make its own payments and in turn replenishing the balances of the banks to which it sends funds. Synchronization of payments thus allows banks to tap incoming transfers from other banks as a key source of funding.

It is possible, however, that the amount of synchronization is less than ideal. The Fedwire Funds Transfer service has many thousands of participating banks, and each day their payment flows are at least slightly different from the previous day's flows. In this environment, a bank may be unaware of incoming funds that may be arriving from a bank with which it rarely exchanges payments. The two banks therefore might not coordinate the timing of their payments as successfully as they might have if they had full information or if they exchanged payments regularly. Furthermore, economic analyses of similar environments suggest that the participants rarely can coordinate well enough to take advantage of the full benefits of coordination, even with full information. Often, the participants coordinate less fully than they would prefer. Repetition of the situation tends to increase the amount of coordination achieved, while the inclusion of more participants tends to decrease the amount. Although none of these analyses has been repeated as frequently as the number of times in which a day's Fedwire payments occur, none has involved as many participants as there are Fedwire banks. Therefore, the amount of payment coordination among banks is conceivably less than desirable.

Measurement of the Different Payment Funding Sources during the Day

We now consider what practical application these observations hold for the Fedwire Funds Transfer service. To accomplish this, we begin by choosing appropriate measures of the different funding sources. Then, using Federal Reserve System data, we can assess the degree to which banks participating in Fedwire use these sources to make payments and we can track that usage at different times of the day. Our goal is to confirm that during the peak activity period, banks fund a larger share of their payments with incoming transfers from other banks than they do at any other time of the day.

We measure the sources of funding available to banks as follows, beginning with the extension of daylight funds overdrafts. To assess fees for banks' use of daylight credit, the Federal Reserve measures overdrafts using the Daylight Overdraft Reporting and Pricing System.[11] Only those overdrafts outstanding at fifty-nine seconds after the minute are included in the overdraft fee calculations (Box A describes the calculation of daylight overdraft charges). We adopt a similar method for measuring overdrafts as a funding source for bank payments: we measure the extension of daylight funds overdrafts in terms of the amount by which a bank's balance falls below zero (or below its negative balance of the previous minute) at the end of a minute, measured on a minute-by-minute basis throughout the day for all banks.[12] In other words, this source of funding measures the amount by which a bank's payments during a minute cause its account balance to fall into (or further into) a negative position.

Our measure of incoming transfers of other banks depends on the time of receipt of the transfer. If the incoming transfer quickly offsets an outgoing one, we consider the incoming transfer to be a source of funding for the outgoing payment. More specifically, our measure of this source of funding is the value of incoming payments that offset outgoing payments within a minute. We adopt this definition because of its relationship to the Federal Reserve's method of measuring overdrafts when assessing fees. As described above, our measure of payments made by overdrafts is based on the amount of overdrafts outstanding at the end of the minute. For that reason, we choose to measure incoming payments that offset outgoing payments made *within the same minute* as those that fund outgoing payments. Those incoming payments either prevent the extension of an overdraft that will be included in the bank's fee calculation or prevent a reduction in the bank's maintained account balance (exact definitions of the variables appear in the appendix).[13] In other words, this source of funding is the value of the payments a bank makes during a minute that, because of funds received during that minute, do not reduce its account balance.

> *It is possible . . . that the amount of [payment] synchronization is less than ideal. A bank may be unaware of incoming funds that may be arriving from a bank with which it rarely exchanges payments.*

After accounting for the payments made by the extension of overdraft credit from the Federal Reserve System and those made by the receipt of incoming transfers from other banks, we assign the remaining payments to banks' maintained account balances at the Federal Reserve. In other words, this source of

Box A

Calculation of Daylight Overdraft Charges[a]

As of April 14, 1994, each depository institution using the Fedwire Funds Transfer and Book-Entry Securities services is charged a fee based on the level of daylight overdrafts it incurs. A daylight overdraft is a negative account balance that occurs during the operating day.

Before describing the calculation of these fees, we note that all daylight overdrafts incurred by a depository institution are subject to a net debit cap. The cap represents the maximum dollar amount of uncollateralized daylight overdrafts that an institution can incur.[b] There are several categories that institutions may fall into that govern the amount of the cap, and the Federal Reserve System monitors their account balances to ensure that cap violations do not occur frequently.

The Federal Reserve follows three steps when calculating an institution's daylight overdraft fee on a particular day:

- First, the average per-minute overdraft incurred by the institution on that day is computed. To do this, the Federal Reserve uses the Daylight Overdraft Reporting and Pricing System to record all negative end-of-minute balances (fifty-nine seconds after the minute). These negative balances are added for the institution for all the minutes of the day in which it has had an overdraft (positive end-of-minute balances are not used to offset negative balances). This sum is divided by the number of minutes in a standard Fedwire day to arrive at the average daily overdraft. Since the expanded operating hours began, a standard Fedwire day runs from 12:30 a.m. to 6:30 p.m. eastern time, for a total of 1,081 minutes.
- Second, the average daily overdraft is multiplied by the fee that the Federal Reserve imposes on daylight overdrafts. Currently, this effective rate equals 15 basis points—or 18/24—an annualized rate of 36 basis points. This effective rate is the annualized rate multiplied by the fraction of the day during which Fedwire operates. To determine the effective daily rate, the Federal Reserve multiplies this number by 1/360. The fee multiplied by the average daily overdraft yields the gross overdraft charge.
- Third, institutions have a deductible, which is a level of overdrafts that they can incur without having to pay a fee. It allows an institution some flexibility in its liquidity management. The deductible is equal to 10 percent of the institution's qualifying capital for daylight overdrafts. The value of the deductible is subtracted from the gross overdraft charge to yield the daily charge to an institution. To determine the value of the deductible, the Federal Reserve multiplies the deductible by a daily effective rate, as in the calculation in the previous bullet. However, there is one difference in the calculations: although the annual rate by which the threshold is valued is also 36 basis points, the fraction of the day is multiplied by 10/24, rather than by 18/24.

After ascertaining each of the above parameters, the Federal Reserve multiplies the average per-minute overdraft by the effective daily rate charged for overdrafts. The value of the institution's deductible is then subtracted from this gross daily charge to arrive at the daily overdraft charge assessed.

The Federal Reserve calculates this daily overdraft charge for each day and totals the charges over a two-week reserve maintenance period. If the sum of the daily overdraft charges incurred during these two weeks is less than $25, the fee is waived.

[a]This section is based on Board of Governors of the Federal Reserve System (1998).

[b]An institution may choose to increase its capacity for daylight overdrafts by pledging collateral, but this collateral is applied to overdrafts related to book-entry securities only. Overdrafts related to funds transfers may not be collateralized.

funding is the value of all payments made during a minute that result in a reduced, but positive, balance in a bank's account. The sum of all these sources of funding equals the sum of payments sent in each minute.

Our measures of the different sources of intraday funds are shown in Chart 5, which depicts the average amounts of each funding source for all Fedwire funds transfers for March 18, April 5, May 13, and June 17, 1999.[14] The outside line of the chart indicates the gross payments made by minute of the day. The interior lines denote the amount of payments made with the three possible sources of funding. It is clear that the utilization of each source varies over the day. In particular, we see a considerable increase in the funding of payments by incoming payments of other banks (arriving in the same minute) during the late afternoon peak. This effect was anticipated by our model (and by our discussion), which suggests that banks coordinate payment timing during the peak afternoon period to take advantage of this funding source.

The shares of the various sources of funding throughout the day are depicted in Chart 6. Early in the day, nearly all

Chart 5

Contributions of Funding Sources of Fedwire Funds Transfers

Average of Four Days

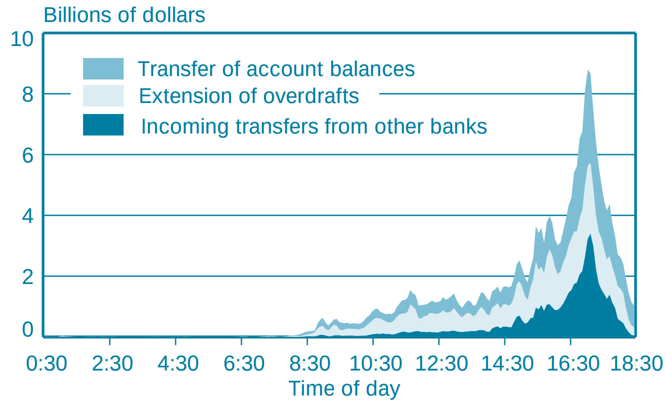

Source: Federal Reserve Bank of New York.

Chart 6

Shares of Funding Sources of Fedwire Funds Transfers

Average of Four Days over Half-Hour Intervals

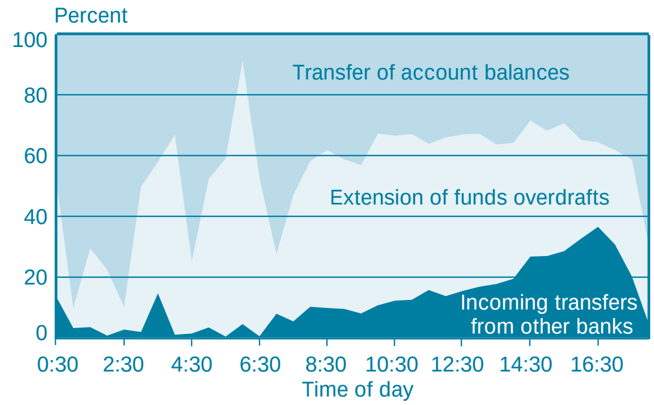

Source: Federal Reserve Bank of New York.

Note: Because few payments are made between 12:30 and 8:30 a.m., the variation in the shares of funding sources during that period of the day is driven by a small number of payments.

payments are made by the transfer of maintained balances or by the use of funds overdrafts extended by the Federal Reserve. As the day progresses, these sources continue to predominate. Finally, as the afternoon payment peak gets under way, incoming payments from other banks that offset outgoing payments within the minute become an important component of payment funding. When payments are highly concentrated, as they are between 4:30 and 5:30 p.m., this (inexpensive) source of funding is the most available and the most utilized. For example, between 4:30 and 5:30 p.m., 16 percent more payment value is funded by incoming payments within the minute than is funded between 2:30 and 4:30 p.m. Overall, 35.6 percent of funds transfers are funded by the movements of maintained balances, 39.0 percent are funded by the extension of funds overdrafts, and 25.4 percent are funded by incoming payments within the minute.

Chart 7 displays the value of the incoming payments that offset outgoing payments within the minute across the four sample days, illustrating both the pattern of funding and the stability of that pattern across the sample days. The correlation between the series in Chart 7 averages .907, indicating that payment activity is highly predictable.

Of course, our measure of the payments funded by incoming funds might be considered conservative. For example, a bank that receives an incoming payment three or five minutes after making a large payment may still be satisfied that the payment was accomplished with less expense than it would have been if an offsetting payment had not arrived until hours later. To gauge how a longer period might affect our measure of funding, we compare the amounts of incoming payments that offset outgoing payments within a fifteen-minute period and within a one-minute period (Chart 8). Again, we see a strong pattern: during the peak period, offsetting payments are matched (in time) more effectively than at any other time during the day. This pattern leads to lower payment costs during the peak period, which in turn reinforces the payment pattern.

Chart 7

Value of Payments Matched by Incoming Transfers

Thirty-Minute Moving Average

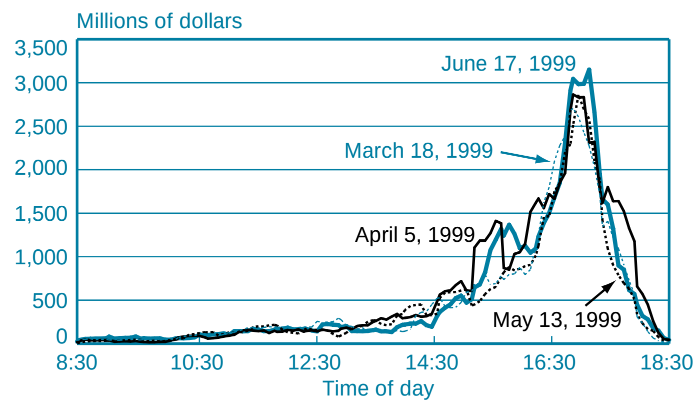

Source: Federal Reserve Bank of New York.

Chart 8

Payments Funded by Incoming Transfers Compared at Different Intervals

Average of Four Days

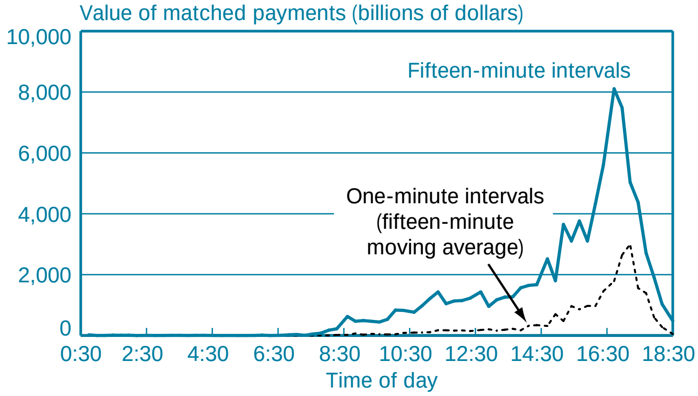

Source: Federal Reserve Bank of New York.

Sensitivity of the Offsetting of Incoming and Outgoing Payments to the Concentration of Payments

The calculation of an elasticity measure offers another way to examine the sensitivity of the value of incoming payments as a funding source to the concentration of payments. Box B displays a fitted relationship between the percentage of payments within a minute that are matched by incoming payments and the percentage of the day's payments that occur in that minute. We find that a quadratic equation fits the data better than a linear relationship does. Using the fitted relationship, we see that the elasticity of the percentage of payments made by incoming payments that offset outgoing payments within the minute to the concentration of payments at the median minute of activity across the four sample days is 0.25. The elasticity initially rises as the concentration of payments rises, and averages approximately 0.55 between 4:30 and 5:30 p.m. This elasticity implies that if 1 percent of payments were transferred from minutes of median payment

Box B

Sensitivity of Funding by Incoming Payments to Payment Concentration

Variables

$Incoming\ \ payments_t$: The percentage of the value of minute t's payments that are offset by incoming payments.

$Amount\ \ of\ \ payments_t$: The percentage of the value of the day's payments that are conducted in minute t.

Fitted Equation

$$Incoming\ \ payments_t = constant + bAmount\ \ of\ \ payments_t + c(Amount\ \ of\ \ payments_t)^2 + \varepsilon .$$

The parameters $constant$, b, and c are to be estimated using the four days of activity used in the construction of Charts 6-8 in the text, and ε is an error term. The estimated equation is given by:

$$Incoming\ \ payments_t = \underset{(.001)}{.0564} + \underset{(1.78)}{76.85}Amount\ \ of\ \ payments_t - \underset{176.1.}{3743.11}(Amount\ \ of\ \ payments_t)^2$$

The standard errors of the parameter estimates are in parentheses. All of the estimated parameters are significant at the 1 percent level. The F value for the equation is 1415 and the adjusted R^2 is .39.

The fitted equation is increasing for all levels of $Amount\ \ of\ \ payments_t$ between [0, .01026]. The average of $Amount\ \ of\ \ payments_t$ during the peak hour between 4:30 and 5:30 p.m. is .00521. In fewer than twenty minutes out of the 4,324 observations does $Amount\ \ of\ \ payments_t$ exceed .01026.

This equation suggests that there is a strongly positive relationship between the degree to which payments offset within a minute, and the concentration of payments within that minute, throughout most of the range of the sample. This equation leads to an elasticity of $Incoming\ \ payments_t$ with respect to $Amount\ \ of\ \ payments_t$ equal to 0.55 during the peak period of payment activity. The elasticity is positive in the concentration of payments over the sample, as shown in the chart below.

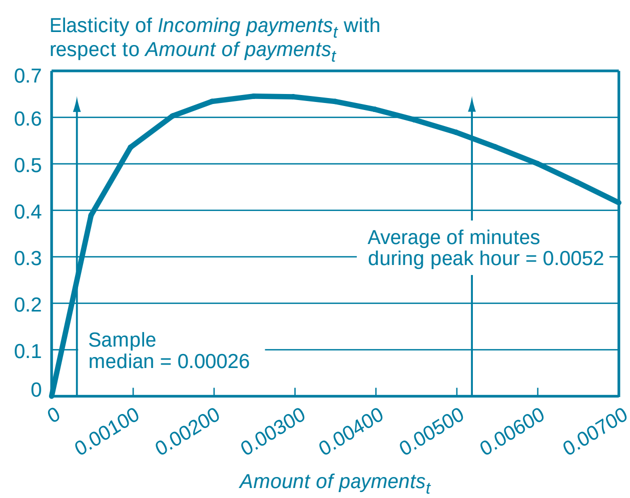

volume to minutes of the hour from 4:30 to 5:30 p.m., there would be a net gain of approximately 0.30 percent in the proportion of payments funded by the matching of offsetting payments. That is, while 0.25 percent of the matching of offsetting payments would be lost from the median minutes of payment activity, 0.55 percent would be gained during the peak hour, for a net increase of 0.30 percent.

It is important to recognize that this relationship is fitted "within-sample"—that is, it is not a forecast of what would happen should more concentration of payments occur. Instead, it records statistically the relationship between the concentration, or synchronization, of payments, and the amount of funding by incoming payments that accompanies that synchronization in the sample days. Within the sample, the positive relationship between concentration of payments and the matching of offsetting payments suggests that by synchronizing payments, one has an effective way to tap this source of funding.

Discussion

The realization that the concentration of payments occurring late in the day may reflect the resolution of a coordination problem suggests that the pattern is stable and durable. With regard to the recently extended early morning hours of the Fedwire Funds Transfer service, one of the difficulties faced in encouraging more payments to be made during these hours is how to raise the expectations of banks that many other banks will also enter payments then. This is a chicken-and-egg problem: until many payments are actually made early in the morning, any individual early payment will rely more on overdrafts or account balances and therefore will be more expensive, at the margin, than if it was made later in the day.[15] This general situation is likely to be faced by new systems planning to operate at that time, such as the one proposed by CLS Bank or the new CHIPS system, which would require that Fedwire payments be made during the early morning hours. For example, CLS Bank proposes settling matched foreign exchange trades at the same time across different currencies (see Roscoe [1998] or Board of Governors of the Federal Reserve System [1999] for a description). Similarly, the new CHIPS system plans to fund an intraday matching system partially with funds sent to a special account early in the morning (see Nelson [1998] for a description). Participants in such arrangements may have to hold additional account balances or utilize more overdraft funding to make their early morning payments than would be necessary if their payments were designed to take place during the peak in activity.

The bunching of payments in the afternoon is in accord with our theoretical model: banks are induced to coordinate payments to take advantage of potentially offsetting funds. The coordination of payment activity—the synchronization of payments—reflects banks' expectations that, as a larger number of payments are entered, more of them will be offsetting and more may, in part, be funding the settlement of other payments. This effect results in a greater use of incoming payments to fund outgoing payments, which in turn would tend to lessen the reliance on account balances or overdrafts to make payments. It saves costs for the banks involved and may help to explain the strong peak in payment activity.

The decreased reliance on other sources of payment funding, including overdrafts from the Federal Reserve, that accompanies synchronization of payments not only can lower costs for commercial banks, but can also reduce the risk of exposure by the Federal Reserve. As banks synchronize their payments more closely, the duration of overdrafts outstanding would be reduced (and the amount of overdrafts would likely fall as well). This relationship is clear if one imagines the extreme case of all payments being made at the same time: overdrafts would be at a minimum level in that the simultaneous entry of all payments would make maximum use of offsetting payments.[16] The reduction in the duration, and possibly the amount, of overdrafts that accompanies the increased concentration of payment timing would therefore reduce the risk of failure to which the Federal Reserve is exposed during the extension of an overdraft loan to a bank.

As we have noted, the degree of payment synchronization might be less than the ideal amount that banks would choose if they could coordinate their payment activity successfully. If that is so, a greater effort to coordinate payment timing may result in a greater share of the day's payments being funded by the synchronized matching and offsetting of payments. As our model suggests, this could be an underutilized source of payment funding. If greater synchronization could be achieved, payments could be made at lower cost and the risk of exposure by the Federal Reserve could be lessened.

Encouraging Payment Synchronization

If the amount of payment coordination is less than ideal, policies that encourage greater coordination of payment timing might be useful. Of course, such policies should not be coercive, but should instead provide more opportunities for banks to coordinate payment timing.

To overcome a lack of payment coordination by banks, the central bank could attempt to guide their expectations by creating a short period—a focal time—in which banks could expect that incoming payments would be entered by other banks. An example of such a policy might be the establishment of two ten- or twenty-minute "synchronization periods." During these periods, only overdrafts outstanding at the end of the period would be entered into a bank's overdraft fee calculation. Banks would not be charged for any overdrafts that they incurred within the synchronization periods and repaid prior to the end of the periods.[17] For example, these periods might operate late in the morning and then early in the afternoon peak.[18] This policy could increase banks' expectations that many payments, including incoming ones, would be entered during the synchronization periods. If the policy was successful, a greater percentage of the day's payments would be made within short periods and therefore would be offsetting within the periods. Less reliance would then be placed on other sources of funding, including the extension of overdrafts. In particular, the average duration of overdrafts would decline and overall overdrafts, including those made within the synchronization periods, would fall.

Potential Problems . . . and Solutions

A policy such as the one just described could conceivably pose some problems—yet those problems are not without solutions:

- *The high degree of payment bunching at the end of the day might increase uncertainty, and could be deleterious to the smooth functioning of the federal funds market near the end of the day.* The act of timing the synchronization periods in the late morning or early in the afternoon peak could mitigate any problems of delayed resolution of uncertainty. Moreover, as long as the delay in anticipation of the periods results in earlier payments on average, the periods would help the participants to overcome the problem of payment delay.
- *If the synchronization period is too short, the successful coordination of so many payments could impose a greater burden on the system's equipment.* This problem represents a resource cost, in terms of computers as well as telecommunications links, of handling a large number of payments in a short period of time. A solution to this problem may be to increase the length of the synchronization periods, or perhaps to employ a peak-load pricing system that would accurately recover the additional costs incurred by the banks that utilize these periods.
- *The central bank would be extending overdrafts during the synchronization periods but not assessing fees on them*. Such overdrafts would be made for a slightly different purpose and they would not last as long as overdrafts extended at other times of the day. In addition, one could view the intrasynchronization-period overdrafts as a cost of achieving the coordination that may lower the overall level and duration of daylight overdrafts. Nonetheless, firm overdraft caps would be necessary to prevent banks from borrowing excessively during the synchronization periods. In addition, overdrafts that exceed some threshold could be required to be collateralized. In this way, if large overdrafts accumulated at any time during the day, the central bank's risk exposure would be securely capped.
- *Depending on the durability of the existing pattern of payments, the synchronization periods might be relatively ineffective, attracting few additional payments*. The cost of implementing the synchronization periods is low. Moreover, we would expect that any increase in the amount of payment coordination would require some time to achieve, as banks adjust to the changing opportunities provided by the periods and the behavior of their counterparties.

Conclusion

Our review of the timing of Fedwire funds transfers suggests that it is reasonable to expect the observed peak in payment activity. It is likely that some payment requests are coordinated to be entered during a peak period of activity late in the afternoon. This pattern is consistent with the outcome of a coordination game among the banks (and among their customers): as banks synchronize their payments more closely, their need for account balances or explicit overdrafts to make payments diminishes. This activity makes payments sent during the peak less expensive, at the margin, than payments sent at other times of the day.

By measuring the funding sources of payments made in the Fedwire Funds Transfer service, we found that approximately 25 percent of a day's payments are funded by incoming payments that offset payments made by banks within the same minute. This source of funding is more readily available during the late-afternoon activity peak, when large-value payments are more closely synchronized; such activity accounts for approximately 40 percent of payment funding at that time.

Payment synchronization benefits banks through the reduced costs of making synchronized payments, but it also has other benefits, as it tends to reduce the amount and duration of overdrafts from the Federal Reserve System. The extension of an overdraft creates a slight risk for the Federal Reserve: should a borrowing bank fail while an overdraft is outstanding, the Federal Reserve would have to seek repayment in bankruptcy court. For this reason, the Federal Reserve has adopted policies to reduce overdrafts (see Board of Governors of the Federal Reserve System [1998] for details). The synchronization of payments is another potential tool for the Federal Reserve to reduce both overdrafts and their duration.

However, the synchronization process, in its current form, may be less than ideal for the Fedwire system's participants. With that in mind, and with the goal of reducing the extent and duration of overdrafts, we considered a policy initiative that could assist banks in synchronizing their payments. The initiative would create synchronization periods in the late morning and early in the current activity peak. During these periods, banks could run intrasynchronization-period overdrafts and not face any charges for them. Banks could be encouraged to enter more payments during these periods, which would lead to reduced payment funding costs and a decreased reliance on overdraft funding during the day.

Finally, many countries recently have adopted real-time gross settlement systems for large-value payments.[19] RTGS systems offer many advantages in managing risk and in linking payment flows with securities markets and other payment systems in a timely fashion. It is important, therefore, to understand better the economic incentives and behavior of participants in an RTGS system. We have focused on the issue of how the cost of liquidity in an RTGS system is affected by the timing of a bank's payment activity, but many other issues remain to be investigated. With better availability of data—and with a range of system designs now operating across countries—the potential for further research into these systems is greater than ever.

Appendix: Definition of Variables

Dollar value of payments made by Bank i to Bank j in minute t:

p_{ij}^{t}.

Funds balance of Bank i at the end of minute t:

$$B_i^t = B_i^0 - \sum_{s=0}^{t} (p_{ij}^s) + \sum_{s=0}^{t} (p_{ji}^s),$$

where B_i^0 is the balance in the bank's account at the Federal Reserve at the start of the day.

Extension of daylight funds overdrafts:

$$D_t = \sum_i \{ min\{ |B_i^t - 0|, |B_i^t - B_i^{t-1}| \} : B_i^t < 0 \text{ and } B_i^t < B_i^{t-1} \}.$$

Dollar value of payments funded by offsetting incoming payments within the minute:

$$I_t = GrossPayments\left[1 - \left\{1 \Big/ \left(½\left(\sum_i \left|\sum_j p_{ij}^t - \sum_j p_{ji}^t\right|\right)\right)\right\}\right].$$

We decompose the funding of gross payments in minute t into the following sources:

$$GrossPayments = D_t + I_t + Allotherpayments.$$

The last category, $Allotherpayments$, consists of those payments funded by the transfer of balances maintained in banks' accounts (for more than a minute) at the Federal Reserve.

Endnotes

1. Fedwire is a large-value payment system owned and operated by the Federal Reserve System. Two services are associated with Fedwire: the Funds Transfer service and the Book-Entry Securities service. In this article, we focus on the Funds Transfer service, which allows a depository institution to transfer funds from an account held at a Federal Reserve Bank to the account of any other Fedwire Funds Transfer service participant.

2. See van Huyck, Battalio, and Beil (1990) for an example of such an analysis.

3. Throughout this article, the dates chosen for the various calculations were governed by the availability of data at the time the calculations were performed. Here, April data are used because April was the most recent full month for which data were available.

4. CHIPS, operated by CHIPCo, is a large-value deferred netting system that settles at 4:30 p.m. Banks face constraints within CHIPS on their net debit positions and may be uncertain as to whether a particular payment can be settled over CHIPS. If a payment does not satisfy the constraints during the operating hours of CHIPS, banks tend to send the payment over Fedwire when CHIPS closes.

5. Angelini (1998) and Kobayakawa (1997) consider alternative models of payment timing in a real-time gross settlement system.

6. Of course, a bank can sell other assets and add to its account balances at the central bank.

7. Reserve balances at the Federal Reserve are charged a zero interest rate and are determined by reserve requirements on the amount of certain deposits at the participating bank. Participating banks receive earning credits on their required clearing balances at the Federal Reserve; these credits are not transferable to third parties.

8. Banks using the Fedwire Funds Transfer service must stay within their "debit cap," which limits the amount by which they can overdraft their account. Some banks have a zero debit cap, which means that they cannot overdraft their accounts at all.

9. Bank for International Settlements (1997). The report included a fourth category: "(d) borrowing from other banks through the money markets," which we include as part of (a).

10. A coordination game is a social situation in which there are gains to the participants from coordinating their actions—such as everyone in the United States driving on the right-hand side of the road or adhering to a uniform calendar of holidays. A model of a payment timing decision that results in a coordination game is available from James McAndrews.

11. See Board of Governors of the Federal Reserve System (1998) for a description of the measurement of overdrafts for assessing overdraft fees.

12. Note that we are measuring only funds-related overdrafts. Daylight overdrafts created in a transaction involving book-entry securities are not considered. See Box A for details.

13. To calculate the amount of offsetting payments, we use an intraminute "netting ratio," which is the ratio of gross payments sent in a given period to the net change in balances required to make those payments within the period. In our earlier example, only \$50 had to be transferred from Bank A to settle all \$225 worth of payments. There, the netting ratio is (225/50) = 4.5, indicating the dollar's worth of payments being made per dollar of deposit funds. A high netting ratio indicates that there is a high degree of offset among the payments being made during the period. We then measure the amount of offsetting payments as $(\textit{gross payments})*(1-(1/\textit{netting ratio}))$. See the appendix for a more complete explanation.

14. The days represent a sample from a set of days for which we have collected bank balance and overdraft data. For each month between March and June 1999, two days were randomly chosen as days for which data were collected.

15. Stehm (1998) points out this issue when reviewing early morning payment activity.

16. We performed a simulation of a multibank payment system with a random distribution of payments across banks. Moving from a situation in which banks spread payments evenly across several periods to a situation in which payments are all made at the same time, the overdrafts of the system fall, and are of reduced duration as payments are concentrated in fewer periods.

17. This policy can be interpreted as a special case of time-varying pricing for overdrafts. That is, banks would not be charged for overdrafts incurred during a synchronization period but would be assessed fees for overdrafts at other times, including the overdrafts outstanding at the close of the synchronization period.

Endnotes (Continued)

18. Our analysis suggests that attempts to alter the timing of the afternoon peak drastically would not be very effective. Placing the synchronization periods at these particular times—in the late morning or early in the afternoon peak—would encourage banks to send more payments then.

19. See Bank for International Settlements (1997).

References

Angelini, Paolo. 1998. “An Analysis of Competitive Externalities in Gross Settlement Systems.” Journal of Banking and Finance 22, no. 1 (January): 1-18.

Bank for International Settlements. 1993. Central Bank Payment and Settlement Services with Respect to Cross-Border and Multi-Currency Transactions. Basel, Switzerland.

_____. 1997. Real-Time Gross Settlement Systems. Basel, Switzerland.

Board of Governors of the Federal Reserve System. 1992. “Notice” in the Federal Register, October 7.

_____. 1994. “Notice” in the Federal Register, February 24.

_____. 1998. “Overview of the Federal Reserve System’s Payments Risk Policy,” May.

_____. 1999. “Application to the Board of Governors of the Federal Reserve System by CLS Services Ltd,” August 6.

Federal Reserve Bank of New York. 1987. “A Study of Large-Dollar Payment Flows through CHIPS and Fedwire,” December.

Federal Reserve Bulletin. 1999. “Highlights of Domestic Open Market Operations during 1998.” Vol. 85, no. 4 (April): 217-35.

Freeman, Scott. 1996a. “The Payments System, Liquidity, and Rediscounting.” American Economic Review 86, no. 5: 1126-38.

_____. 1996b. “Clearinghouse Banks and Banknote Over-Issue.” Journal of Monetary Economics 38, no. 1 (August): 101-15.

Green, Edward J. 1997. “Money and Debt in the Structure of Payments.” Bank of Japan Monetary and Economic Studies 15, no. 1 (May): 63-87.

Kobayakawa, Shuji. 1997. “The Comparative Analysis of Settlement Systems.” Centre for Economic Policy Research Discussion Paper no. 1667.

McAndrews, James, and William Roberds. 1999. “Payment Intermediation and the Origins of Banking.” Federal Reserve Bank of New York Staff Reports, no. 85.

Nelson, Norman R. 1998. “Private Sector Clearing and Payment Systems.” In Payments Systems in the Global Economy: Risks and Opportunities. Proceedings of the Federal Reserve Bank of Chicago’s 34th Annual Conference on Bank Structure and Competition.

Richards, Heidi Willman. 1995. “Daylight Overdraft Fees and the Federal Reserve’s Payment System Risk Policy.” Federal Reserve Bulletin 81, no. 12 (December): 1065-77.

Roscoe, David L. III. 1998. “Continuous Linked Settlement.” In Payments Systems in the Global Economy: Risks and Opportunities. Proceedings of the Federal Reserve Bank of Chicago’s 34th Annual Conference on Bank Structure and Competition.

Stehm, Jeff. 1998. “The Eighteen-Hour Payment Day.” Back-Office Focus 37 (April): 5-10.

van Huyck, J. B., R. C. Battalio, and R. O. Beil. 1990. “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure.” American Economic Review 80, no. 1: 234-48.

Arturo Estrella, Sangkyun Park, and Stavros Peristiani

Capital Ratios as Predictors of Bank Failure

- The current regulatory framework for determining bank capital adequacy is under review by the Basel Committee on Banking Supervision.
- An empirical analysis of the relationships between different capital ratios and bank failure suggests that two simple ratios—the leverage ratio and the ratio of capital to gross revenue—may merit a role in the revised framework.
- The leverage ratio and the gross revenue ratio predict bank failure about as well as more complex risk-weighted ratios over one- or two-year horizons. Risk-weighted ratios tend to perform better over longer horizons.
- The simple ratios are virtually costless to implement and could supplement more sophisticated measures by providing a timely signal of the need for supervisory action.

Capital ratios have long been a valuable tool for assessing the safety and soundness of banks. The informal use of ratios by bank regulators and supervisors goes back well over a century (Mitchell 1909). In the United States, minimum capital ratios have been required in banking regulation since 1981, and the Basel Accord has applied capital ratio requirements to banks internationally since 1988. The Basel Committee on Banking Supervision (1999) is currently engaged in an effort to improve the Basel Accord and, once again, capital ratios are being discussed as part of the proposed solution. In this article, we examine some of the roles that capital ratios play in bank regulation and we argue that, to be successful in any of these roles, capital ratios should bear a significant negative relationship to the risk of subsequent bank failure. We then present empirical evidence of those relationships.

We focus here on three types of capital ratios—risk-weighted, leverage, and gross revenue ratios. For each ratio, we examine what makes it actually or potentially useful for bank regulation and we ask whether it is indeed significantly related to subsequent bank failure. Perhaps not surprisingly, we find that all three ratios are strongly informative about subsequent failures. Our analysis suggests that the most complex of the ratios—the risk-weighted ratio—is the most effective predictor of failure over long time horizons. However, perhaps somewhat surprisingly, we also find that the risk-weighted ratio does not consistently outperform the simpler ratios, particularly with short horizons of less than two years. Over the

Arturo Estrella is a senior vice president, Sangkyun Park an economist, and Stavros Peristiani a research officer at the Federal Reserve Bank of New York.

The authors thank Beverly Hirtle, Jim Mahoney, Tony Rodrigues, Phil Strahan, two anonymous referees, and participants in a workshop at the Federal Reserve Bank of New York for helpful comments and suggestions. The authors also thank Gijoon Hong for excellent research support. The views expressed are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

shorter time periods, the leverage and gross revenue ratios can play a crucial role as timely backstop thresholds that would trigger regulatory action if breached. They also have the advantage of being less costly to calculate and report. In this context, the trade-off between regulatory burden and predictive accuracy may not favor the risk-based measures.

In the next section, we develop the conceptual arguments in favor of applying capital ratios in bank regulation. We then proceed to use the empirical evidence on U.S. bank failures to evaluate the effectiveness of the ratios in predicting bank failures.

The Role of Capital Ratios in Bank Analysis and Supervision

Although bank regulators have relied on capital ratios formally or informally for a very long time, they have not always used the ratios in the same way. For instance, in the days before explicit capital requirements, bank supervisors would use capital ratios as rules of thumb to gauge the adequacy of an institution's level of capital. There was no illusion that the simple ratios used (for example, capital to total assets or capital to deposits) could provide an accurate measure of the appropriate capital level for a bank, but large deviations of actual capital ratios from supervisory benchmarks suggested the need for further scrutiny.

When capital ratios were introduced formally in regulation in 1981 (see Gilbert, Stone, and Trebing [1985]), they were applied in a different way. The regulatory requirement set a minimum level of capital that the institution had to hold. The degree to which the requirement was binding depended significantly on the type of institution because, then as now, there was substantial diversity among banking institutions. Indeed, several classes of institutions were initially defined and accorded different treatment by the regulation. Basically, the requirements were most binding for less than a couple of dozen large banks, whereas smaller banks were generally already in compliance with the more stringent requirements.

The Basel Accord of 1988 attempted to deal with the diversity in institutional activities by applying different credit risk weights to different positions and by including in the base for the capital ratio a measure of the off-balance-sheet exposures of the bank. Despite these calibrations, the intent was not to determine an exact appropriate level of capital for the bank, but rather to provide a more flexible way of determining the minimum required level (Basel Committee on Banking Supervision 1988).

Another significant regulatory development in the United States was the Federal Deposit Insurance Corporation Improvement Act of 1991 (FDICIA), which introduced the concept of "prompt corrective action." The degree of supervisory intervention in specific banks is now guided by a formula driven largely by the Basel ratios and by a simple leverage ratio. Banks are classified as "adequately capitalized" if they meet the Basel requirements plus a leverage ratio requirement, but additional distinctions are made among levels of capital. For example, a bank is "well capitalized" if it holds a certain buffer above the adequate levels.

Our goal . . . is to suggest that [the leverage and gross revenue] ratios contain valuable and virtually costless information, and therefore have a role in an overall framework for regulatory capital.

In contrast, a bank that falls under a specific level, set somewhat below the minimum adequate level, is determined to be "critically undercapitalized" and must be shut down by supervisors. This is a different concept of a minimum requirement from the one used in earlier regulation in that failure to comply results in the closure of the institution. Rather than a minimum safe operating level, which the earlier rules had tried to identify, the new cutoff point is a backstop level, below which the bank is no longer considered to be viable.

The June 1999 Basel capital proposal goes beyond the ratios based on accounting data that we have discussed so far. The proposal contemplates (1) the use of external credit ratings as determinants of the weights to be applied to various asset categories, (2) the use, for the same purpose, of internal bank credit ratings based on the firm's own judgment, and (3) the extended recognition of various forms of credit risk mitigation. These features constitute a difference in kind, not simply magnitude, as compared with the accounting-based ratios on which we focus in this article. Ideally, we would like to be able to compute ratios based on the new proposal to examine their power to predict failure, but the required information simply does not exist at this time. We should note, however, that we do not argue here that the ratios that we do examine should substitute for any of the foregoing Basel proposals. Our goal instead is to suggest that some of those ratios contain valuable and virtually costless information, and therefore have a role in an overall framework for regulatory capital.

The preceding discussion alludes to a number of distinctions between approaches to benchmarks based on capital ratios, and it may be helpful to spell these out. In some cases, a ratio is intended as a minimum acceptable level, whereas in other cases, the ratio may identify an appropriate level of capital for the bank. This distinction between a minimum and an "optimum" level is discussed in detail in Estrella (1995).

Another distinction is between adequate levels and backstop levels, such as in the 1991 FDICIA legislation. In one case, there is a certain level of comfort for bank supervisors, while in the other case, the bank is no longer considered viable. It is possible that a particular ratio may be more suited for one of these two cases than for the other.

Closely related is the distinction between the value of a bank in liquidation and the value of a bank as a going concern. For instance, one of the motivations for the 1991 legislation was that the net value of a bank tends to decrease when the bank ceases to be a going concern and moves into liquidation mode (see, for example, Demsetz, Saidenberg, and Strahan [1996]). Thus, the level of capital that is adequate for regulatory and supervisory purposes may differ between banks operating normally and banks in the process of liquidation. These distinctions are demonstrated in the following simple graph.

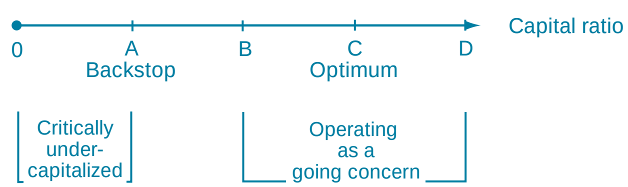

The optimum level, defined in various ways in economic research (see discussions in Estrella [1995] and Berger, Herring, and Szegö [1995]), is shown as point C in the graph. Theoretically, this is the level that maximizes some objective function for bank owners, but in practice this exact level is very difficult to ascertain with any precision. Nevertheless, there is an informal range around this level, say from point B to point D, over which capital may be generally considered adequate for a going concern. That is, capital is high enough to allow regulators, shareholders, and depositors to sleep at night, but not so high that the total cost of capital to the firm outweighs its benefits. Finally, point A identifies the backstop level at which the bank is no longer viable and must be shut down to prevent losses to depositors and to the public.

The Relationship between Capital Ratios and Bank Failure

The relationship between the level of capital and subsequent failure is clear in the case of a backstop level as defined above. At this level, the bank is either a de facto failure or is in imminent danger of falling into that category. Therefore, regulators must choose a backstop level that is highly correlated with failure in the very short run; that is, the level should be associated with a fairly high probability of failure. Regulators will generally select a positive level for the backstop rather than the level of technical insolvency at which the net worth of the bank is zero. One reason is that the valuation of the bank is not known precisely until liquidation. There is no assurance that a liquidated bank will be valued at the accounting net worth, although this type of uncertainty could signify that the actual value of the bank could be either higher or lower than the accounting value. A second reason is that, for a going concern, there is generally a "charter value"—an intangible value that disappears with the closure of the institution. Hence, even if the accounting valuation were perfectly accurate in the first sense, the mere liquidation of the institution could lead to a loss in net value.

Regulators must choose a backstop level that is highly correlated with failure in the very short run; that is, the level should be associated with a fairly high probability of failure.

This potential loss in the value of the firm in liquidation also helps explain why capital levels in general should be significantly related to bank failure. The charter value of the bank produces a strong incentive to the owners of the bank to manage it as a going concern. If the bank fails, one consequence is the dissipation of charter value—value that the owners could capture by selling their stakes if the institution were viable. Thus, owners have an interest in maintaining a level of capital that is consistent with a low probability of failure. Needless to say, regulators and supervisors also tend to favor low probabilities of failure.

To summarize, with reference once more to the graph, the backstop level at point A corresponds to a fairly high probability of failure but represents enough capital to deal with uncertainties relating to the value of the firm in liquidation. In contrast, values above point B correspond to probabilities of failure that are sufficiently low to satisfy the requirements of owners, regulators, and others.

Useful Features of Capital Ratios

A capital ratio is constructed from two components. The numerator is a measure of the absolute amount (for example, the dollar value) of capital of the firm and is inversely related to the probability of failure. The denominator is a proxy for the absolute level of risk of the bank. By taking the ratio we are able to gauge whether the absolute amount of capital is adequate in relation to some indicator of absolute risk. Basically, a large bank needs a larger amount of capital than a small bank, ceteris paribus, and a riskier bank needs more capital than a less risky bank, ceteris paribus. Absolute risk is probably roughly proportional to scale, so that measures of scale are generally good proxies for absolute risk. The three ratios we examine in this paper represent various approaches to measuring scale and absolute risk.

We will define the ratios more precisely in the next section, but we provide here some preliminary discussion of how each deals with scale and risk. Let us assume, as is the case in our empirical sections, that the numerator is the same measure of capital for all ratios, which allows us to focus on the alternative denominators. In the case of the leverage ratio, the denominator is the total assets of the bank. This measure, which has a long history, assumes implicitly that the capital needs of a bank are directly proportional to its level of assets. For some broad classes of banks, this may not be a bad assumption. However, if we take the example of two banks, only one of which has substantial and risky off-balance-sheet activities, the use of the leverage ratio may produce misleading relative results.

A leverage ratio requirement may also affect the asset allocation of banks that are constrained by the requirement. Constrained banks are likely to reduce low-risk assets such as Treasury securities, which are easily marketable, as opposed to less marketable assets such as loans. Nevertheless, a clear advantage of the leverage ratio is simplicity. It is virtually costless to administer and very transparent.

In 1988, the Basel Accord introduced the concept of risk-weighted assets as the denominator of the capital ratio. This measure contains a component representing off-balance-sheet exposures and also adjusts for differentials in credit risk according to type of counterparty and type of instrument. As such, the Basel ratio represents a well-known example of a risk adjustment to the basic scale of the denominator.

Risk weighting effectively requires financial institutions to charge more capital for riskier assets, discouraging them from holding risky assets. By responding to the risk-reducing incentives, banks can increase the risk-weighted ratio without raising capital. On the other hand, failure to respond would result in a low risk-weighted ratio. Thus, if risk weights accurately reflect the riskiness of assets, the risk-weighted ratio should better distinguish between risky and safe banks and should be a more effective predictor of bank failure than simple ratios. Inaccuracy is unavoidable, however. Because each loan is unique, it is difficult to evaluate the credit risk of bank assets. In addition, the business of banking is subject to significant sources of risk other than credit risk, such as interest rate risk, operational risk, and reputational risk. Weighting assets can weaken the relationship between the capital ratio and these other risks—operational risk in particular.

It is not certain a priori that the risk-based capital ratio is meaningfully superior to simple ratios in capturing the overall risk of banks.

Furthermore, the financial sector is so dynamic that new products are introduced continuously. Even a well-designed risk-weighting scheme may soon become obsolete as new instruments provide means of economizing on regulatory capital. Considering these difficulties, it is not certain a priori that the risk-based capital ratio is meaningfully superior to simple ratios in capturing the overall risk of banks. Regulatory capital arbitrage under risk-based capital requirements could even produce harmful economic effects. For instance, since lending to risky borrowers belongs in the highest risk-weight category, the incentive to economize capital might induce banks to reduce lending to those borrowers that do not have alternative financing sources.[1] Economic activity may contract as a result. In addition, it is costly to administer risk-based capital requirements, especially since both monitoring and reporting burdens may be heavy.

Our third ratio—not currently part of the regulatory framework but suggested, for example, by Shepheard-Walwyn and Litterman (1998)—uses the gross revenue of the bank as the measure of scale. Like total assets, gross revenue is easily obtainable from the financial statements of the firm and thus is virtually costless to administer. Unlike assets, however, gross revenue includes components associated with off-balance-sheet activities. Moreover, gross revenue contains a crude “risk adjustment” in that riskier projects are likely to be undertaken only if they provide larger revenues, at least ex ante. Thus, gross revenue may reflect the riskiness of bank assets better than total assets, though in principle not as well as risk-weighted assets.

A potential drawback is that gross revenue also captures factors other than risk. For example, banks engaging heavily in fee-generating activities, which may carry only a limited amount of risk, will report large revenue. Gross revenue may also be more sensitive to business cycles than total assets, although this is not entirely clear and is largely an empirical question. This measure has not been subjected to the test of actual usage, but gross revenue seems to be less susceptible to regulatory capital arbitrage than other measures. For instance, it may be difficult for banks to reduce gross revenue without hurting profits or general investor perceptions. As for simplicity, gross revenue is, like assets, a standard accounting concept. Thus, the gross revenue ratio is as simple and transparent as the leverage ratio.

Capital Ratios and the Likelihood of Failure

To assess the predictive efficacy of capital ratios, our analysis utilizes standard measures defined by the existing capital adequacy rules. The measure of capital applied in the numerator of all three ratios is *tier 1 capital*, defined to include common stock, common stock surplus, retained earnings, and some perpetual preferred stock. The *risk-weighted capital ratio* is defined as the ratio of tier 1 capital to risk-weighted assets. The definition of the *leverage ratio* is tier 1 capital divided by total tangible assets (quarterly average). The *gross revenue ratio* is tier 1 capital divided by total interest and noninterest income before the deduction of any expenses.

Our database includes all FDIC-insured commercial banks that failed or were in business between 1989 and 1993. The sample period ends in 1993 because for the most part there were just a handful of bank failures after this period. Because risk-weighted capital measures were not implemented and reported until after 1990, it is difficult to estimate meaningful risk-weighted ratios in the early and mid-1980s. To compute the various capital ratios, we used information from the Consolidated Reports of Condition and Income (Call Reports) produced by the Federal Financial Institutions Examination Council. The Federal Reserve Board provides a formal algorithm for calculating risk-weighted ratios for 1991 and subsequent years. Risk-weighted capital ratios for 1988, 1989, and 1990 were estimated based on the *Capital Adequacy Guidelines* published by the Board of Governors of the Federal Reserve System.

Table 1 presents summary statistics for the three different measures of capital adequacy for the period 1988-92. Looking at the top panel of the table, we observe that the mean and median leverage ratios for our sample of banks during this period are fairly stable at around 9 and 8 percent, respectively. Since these statistics are based on unweighted data, they are influenced heavily by the large number of small banks that tend to have higher capital ratios. The average capital ratios weighted by assets (not shown in table) are lower. The table also helps to highlight that the gross revenue measure (middle panel) varies more widely across years, reflecting its close relationship with economic conditions. Relatively high gross revenue ratios in 1991 and 1992 can be explained by reduced banking revenue caused by an economic downturn. Both the mean and the median of the risk-weighted capital ratio (bottom panel) were substantially higher than the required ratio (4 percent). The standard deviation, however, was large, suggesting that many banks had difficulty in meeting the capital requirement.

Table 1
Summary Statistics

Year	Number of Observations	Mean	Median	Standard Deviation	Minimum	Maximum
	Leverage Ratio					
1988	13,299	0.094	0.082	0.077	-0.512	0.998
1989	12,903	0.096	0.083	0.076	-0.440	0.995
1990	12,388	0.094	0.082	0.072	-0.549	0.998
1991	11,941	0.094	0.082	0.070	-0.438	0.998
1992	11,473	0.096	0.085	0.068	-1.663	0.997
Total	62,004	0.095	0.083	0.073	-1.663	0.998
	Gross Revenue Ratio					
1988	13,299	1.146	0.866	3.712	-4.938	300.110
1989	12,903	1.228	0.816	13.192	-4.228	1,345.000
1990	12,388	1.032	0.819	2.239	-4.124	135.240
1991	11,941	1.211	0.864	15.051	-1.088	1,601.330
1992	11,473	1.253	1.004	6.683	-0.729	679.500
Total	62,004	1.173	0.871	9.595	-4.938	1,601.330
	Risk-Weighted Capital Ratio					
1988	13,299	0.186	0.142	0.264	-0.607	12.383
1989	12,903	0.195	0.144	0.608	-0.739	52.089
1990	12,388	0.179	0.136	0.298	-0.524	9.534
1991	11,941	0.208	0.139	3.040	-0.439	330.902
1992	11,473	0.193	0.147	0.487	-1.584	34.249
Total	62,004	0.192	0.141	1.390	-1.584	330.902

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; authors' calculations.

Table 2

Measures of Correlation

Year	Leverage Ratio–Gross Revenue Ratio	Leverage Ratio–Risk-Weighted Capital Ratio	Gross Revenue Ratio–Risk-Weighted Capital Ratio
	Pearson Correlation Coefficient		
1988	0.410	0.749	0.284
1989	0.216	0.442	0.179
1990	0.496	0.740	0.344
1991	0.151	0.179	0.020
1992	0.221	0.537	0.567
Total	0.194	0.210	0.069
	Spearman's Rank Correlation		
1988	0.930	0.825	0.840
1989	0.932	0.849	0.865
1990	0.921	0.849	0.859
1991	0.911	0.824	0.833
1992	0.874	0.783	0.788
Total	0.917	0.830	0.841

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; authors' calculations.

In Table 2, we present measures of correlation for all three capital adequacy ratios. While the Pearson correlation coefficients (top panel) are statistically significant, one may surmise from their magnitude that these capital measures are not consistently correlated over time. However, looking at the bottom panel of the table, which shows large and significant rank correlation estimates, we conclude that most of the large fluctuations in the parametric measure of correlation are caused by the presence of outliers. Although the rank correlation is high, these capital ratios are far from perfectly correlated. Thus, each capital ratio may provide some independent information about capital adequacy.

Distribution of Bank Failure

A good measure of capital adequacy should be related very closely to bank failure. The first phase of our analysis investigates this issue by looking at the distribution of bank failures with respect to the alternative capital ratios. Table 3 presents one-year bank failure rates for various levels of the leverage ratio at the end of the preceding year. The table covers all failed and surviving banks during the period 1989-93. We excluded from the analysis all banks that were acquired during the period because many of these mergers involved problem target banks. In its final form, the data set is an unbalanced panel of banks, in which a bank is observed until the time of failure or until the end of 1993. To be specific, a bank that survived between 1989 and 1993 is counted five times as a nonfailure, and a bank that failed in 1991 is counted twice as a nonfailure (1989 and 1990) and once as a failure (1991). In the next subsection, we will also present a parametric model of survival that gives a more precise account of the conditional distribution of failure.

In the top panel of Table 3, we use an absolute scale to tally failures (observations of banks that failed within a year of the reported capital ratio) and nonfailures (observations of banks that did not fail within a year) for individual capital ratio ranges and cumulatively up to a given cutoff point. For noncumulative data, each range is bounded above by the cutoff point of the row and bounded below by the cutoff point of the previous row. The bottom panel of Table 3 uses a relative scale for the leverage ratio by classifying banks according to percentiles. The absolute scale is helpful for examining the failure experience at specific ranges of the ratio. In contrast, by dividing the data set into percentile classes of equal size, ranked by the ratio, the relative scale facilitates a uniform comparison of the different capital ratios.

A good measure of capital adequacy should be related very closely to bank failure. . . . Our analysis investigates this issue by looking at the distribution of bank failures with respect to the alternative capital ratios.

As the column headed "Failure Rate for Row" indicates, the proportion of failed observations (number of failures divided by the total number of banks in the leverage ratio class) on an absolute scale (top panel) was more than 80 percent for institutions with negative leverage ratios. The proportion of failing bank observations decreases monotonically and rapidly with the leverage ratio; the relative frequency drops below 10 percent in the leverage ratio range of 4 to 5 percent and below 1 percent in the 6 to 7 percent range. The proportion is quite small (0.1 percent or lower) for bank observations with leverage ratios higher than 7 percent. In relative terms, the bottom panel of Table 3 shows that the proportion of failures is very high (74.7 percent) for banks in the lowest 1 percentile leverage ratio range but quickly drops below 10 percent in the 3 to 4 percentile class. The sharp drop-off in the proportion of failures is indicative of a successful measure.

Table 3

Distribution of Bank Failures by Leverage Ratios

Cutoff Percentile	Cutoff Point	Failures 1989-93	Nonfailures 1989-93	Failure Rate for Row (Percent)	Cumulative Proportion of Nonfailures (Type II Error) (Percent)	Cumulative Proportion of Failures (Type I Error) (Percent)
			Absolute Scale			
—	0	231	51	81.9	0.1	63.2
—	1.0	100	62	61.7	0.3	47.3
—	2.0	90	95	48.6	0.5	33.0
—	3.0	76	194	28.1	0.9	20.9
—	4.0	45	367	10.9	1.8	13.7
—	5.0	31	628	4.7	3.2	8.8
—	6.0	25	1,799	1.4	7.3	4.8
—	7.0	17	5,136	0.3	19.1	2.1
—	8.0	8	8,175	0.1	37.8	0.8
—	9.0	0	7,767	0.0	55.6	0.8
—	10.0	3	5,858	0.1	69.0	0.3
—	11.0	0	3,940	0.0	78.1	0.3
—	12.0	0	2,702	0.0	84.3	0.3
—	Infinity	2	6,869	0.0	100.0	0.0
			Relative Scale			
1	0.97	330	112	74.7	0.3	47.5
2	2.95	166	277	37.5	0.9	21.0
3	4.03	46	397	10.4	1.8	13.7
4	4.78	22	420	5.0	2.8	10.2
5	5.20	13	430	2.9	3.7	8.1
6	5.51	7	436	1.6	4.7	7.0
7	5.75	8	435	1.8	5.7	5.7
8	5.92	3	439	0.7	6.8	5.3
9	6.06	3	440	0.7	7.8	4.8
10	6.18	2	441	0.5	8.8	4.5
25	7.22	18	6,180	0.3	22.9	1.6
50	8.55	5	11,063	0.0	48.3	0.8
75	10.46	3	11,065	0.0	73.6	0.3
100	Infinity	2	11,508	0.0	100.0	0.0

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Noncumulative data are for the range defined by cutoffs in the current and the previous row. Cumulative data are aggregated up to the cutoff point.

In addition to reporting the frequency of failure for specific ranges, Table 3 presents cumulative frequencies. The cumulative proportion of nonfailures represents the number of surviving observations up to that leverage ratio cutoff point, divided by the aggregate number of nonfailing observations. In contrast, the cumulative proportion of failures represents the total number of failures for bank observations having a leverage ratio greater than or equal to the leverage ratio cutoff value, divided by the total number of failures.[2] Looking at the cumulative proportion of nonfailures, we find that only 0.5 percent of nonfailures would be classified under prompt corrective action as critically undercapitalized (that is, showing a leverage ratio of less than 2 percent).[3] In comparison, 33 percent of the failures did not fall in the critically undercapitalized region (67 percent did).

We may interpret these cumulative proportions using simple statistical hypothesis-testing terminology. In this context, the null or testable hypothesis is that the bank will fail within one year; the alternative hypothesis is that the bank will not fail over the same period. Acceptance of the null hypothesis, in turn, would be associated with some appropriate action on the part of the supervisory authority—for instance, closure of the bank. Accepting the null hypothesis when it is actually false (known as Type II error) is equivalent to closing a bank that would have survived beyond one year, which in Table 3 corresponds to the proportion of nonfailed bank observations. Similarly, the cumulative proportion of failures is analogous to the so-called Type I error, that is, the decision not to close an institution that failed within one year. Consider, for example, the 2 percent closure rule for critically undercapitalized banks, using the figures reported in the previous paragraph. The Type II error is only 0.5 percent (0.5 percent of nonfailures were statistically misclassified). In contrast, the Type I error for observations with a leverage ratio greater than 2 percent is 33 percent (that is, 33 percent of the failures were statistically misclassified). Note that there is a trade-off in general between the probabilities of Type I and Type II errors. It is impossible to reduce both simultaneously by shifting the cutoff ratio.

The gross revenue ratio classifies failures and nonfailures about as accurately as the leverage ratio.

Although it would be difficult for bank supervisors to frame any practical regulatory goals based solely on these statistical errors, sound regulatory policies should help to promote some balance between these cumulative proportion errors of failure and nonfailure. As the lower panel of Table 3 suggests, the two cumulative ratios are approximately equal around the seventh percentile cutoff, which is equivalent to the 5.75 percent leverage ratio cutoff point.[4] In addition, it is interesting to note that current FDICIA capital adequacy guidelines for well-capitalized banks, which require a 5 percent leverage ratio, would have generated Type I and Type II errors of 3.2 percent and 8.8 percent, respectively.

Bank failures are correlated about as strongly with gross revenue ratios as with leverage ratios (Table 4). As in the case of leverage ratios, the proportion of failing observations declines quite rapidly with the gross revenue ratio, and failures are highly concentrated at low gross revenue ratios. The top panel may be somewhat difficult to interpret because the levels of the gross revenue ratio tend to be less familiar than the levels of standard capital ratios. Nonetheless, our results illustrate that the likelihood of failure is quite small for depository institutions that maintain a gross revenue ratio greater than 60 percent. Interestingly, the bottom panel reveals that the cumulative proportion of failed banks (Type I error) is approximately equal to the cumulative proportion of nonfailures (Type II error) around the 60 percent gross revenue ratio threshold. Overall, a comparison of the bottom panels of Tables 3 and 4 suggests that the gross revenue ratio classifies failures and nonfailures about as accurately as the leverage ratio. The two panels show very similar failure rates, Type I errors, and Type II errors in each percentile class.

Finally, Table 5 shows the distribution of bank failures for the tier 1 risk-weighted capital ratio. In general, the distribution of failures against tier 1 risk-weighted capital ratios is comparable to that for the other capital ratios. However, the table also reveals a number of small differences between the tier 1 risk-based measure and the leverage ratio. Current FDICIA rules specify that a well-capitalized bank must maintain, as minimum levels, a 6 percent tier 1 risk-weighted capital ratio, a 10 percent total (tier 1 plus tier 2) risk-weighted capital ratio, and a 5 percent leverage capital ratio.[5] Note that the failure rate at the 6 to 7 percent tier 1 capital range is 5.2 percent. In comparison, the failure rate for well-capitalized banks with 5 to 6 percent leverage ratios is only 1.4 percent (Table 3). This pair-wise comparison suggests that the 5 percent leverage ratio threshold is more binding than the 6 percent tier 1 risk-based requirement. Having said that, however, we should note that the stringency in the risk-weighted ratios is best captured by the total (tier 1 plus tier 2) ratio. Although the distribution table for the total risk-weighted measure is not included in this article, we find that the failure rate at the 10 to 11 percent range is only 0.4 percent, suggesting that

Table 4

Distribution of Bank Failures by Gross Revenue Ratios

Cutoff Percentile	Cutoff Point	Failures 1989-93	Nonfailures 1989-93	Failure Rate for Row (Percent)	Cumulative Proportion of Nonfailures (Type II Error) (Percent)	Cumulative Proportion of Failures (Type I Error) (Percent)
			Absolute Scale			
—	0	231	51	81.9	0.1	63.2
—	10	102	76	57.3	0.3	47.0
—	20	93	160	36.8	0.7	32.2
—	30	75	299	20.1	1.3	20.2
—	40	42	488	7.9	2.5	13.5
—	50	36	772	4.5	4.2	7.8
—	60	13	1,755	0.7	8.3	5.7
—	70	14	3,634	0.4	16.6	3.5
—	80	13	5,431	0.2	29.0	1.4
—	90	5	5,945	0.1	42.6	0.6
—	100	1	5,431	0.0	55.1	0.5
—	110	2	4,526	0.0	65.5	0.2
—	120	0	3,499	0.0	73.5	0.2
—	Infinity	1	11,576	0.0	100.0	0.0
			Relative Scale			
1	8.85	323	119	73.1	0.3	48.6
2	25.17	148	295	33.4	0.9	25.0
3	34.56	60	383	13.5	1.8	15.4
4	42.61	24	418	5.4	2.8	11.6
5	47.93	20	423	4.5	3.8	8.4
6	51.97	7	436	1.6	4.8	7.3
7	54.83	4	439	0.9	5.8	6.7
8	57.16	3	439	0.7	6.8	6.2
9	59.09	3	440	0.7	7.8	5.7
10	60.87	2	441	0.5	8.8	5.4
25	75.30	19	6,179	0.3	22.9	2.4
50	94.27	11	11,057	0.1	48.3	0.6
75	120.24	3	11,065	0.0	73.6	0.2
100	Infinity	1	11,509	0.0	100.0	0.0

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Noncumulative data are for the range defined by cutoffs in the current and the previous row. Cumulative data are aggregated up to the cutoff point.

the total risk-based measure may be the most binding of all the FDICIA capital adequacy ratios.

As expected, the performance of capital ratios deteriorates somewhat when we move from a one-year to a two-year horizon, that is, when we focus on failures occurring between one and two years after the capital ratio is observed. Tables 6-8 summarize the second-year failure rates and cumulative distribution of second-year failures and nonfailures for firms that survive the first year. The three capital ratios still provide a fairly clear signal, as evidenced by the sharp drop in the failure

Table 5

Distribution of Bank Failures by Risk-Weighted Capital Ratios

Cutoff Percentile	Cutoff Point	Failures 1989-93	Nonfailures 1989-93	Failure Rate for Row (Percent)	Cumulative Proportion of Nonfailures (Type II Error) (Percent)	Cumulative Proportion of Failures (Type I Error) (Percent)
			Absolute Scale			
—	0	231	52	81.6	0.1	63.2
—	1.0	69	39	63.9	0.2	52.2
—	2.0	59	46	56.2	0.3	42.8
—	3.0	60	73	45.1	0.5	33.3
—	4.0	55	140	28.2	0.8	24.5
—	5.0	35	203	14.7	1.3	18.9
—	6.0	33	261	11.2	1.9	13.7
—	7.0	25	454	5.2	2.9	9.7
—	8.0	17	775	2.1	4.7	7.0
—	9.0	7	1,251	0.6	7.5	5.9
—	10.0	10	2,217	0.4	12.6	4.3
—	11.0	5	3,061	0.2	19.6	3.5
—	12.0	8	3,492	0.2	27.6	2.2
—	Infinity	14	31,579	0.0	100.0	0.0
			Relative Scale			
1	1.50	330	112	74.7	0.3	47.5
2	4.31	158	285	35.7	0.9	22.3
3	5.89	51	392	11.5	1.8	14.2
4	6.87	27	415	6.1	2.8	9.9
5	7.55	10	433	2.3	3.8	8.3
6	8.03	8	435	1.8	4.7	7.0
7	8.44	2	441	0.5	5.8	6.7
8	8.77	4	438	0.9	6.8	6.1
9	9.05	1	442	0.2	7.8	5.9
10	9.28	5	438	1.1	8.8	5.1
25	11.42	15	6,183	0.2	22.9	2.7
50	14.66	10	11,058	0.1	48.3	1.1
75	19.86	2	11,066	0.0	73.6	0.8
100	Infinity	5	11,505	0.0	100.0	0.0

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Noncumulative data are for the range defined by cutoffs in the current and the previous row. Cumulative data are aggregated up to the cutoff point.

Table 6

Distribution of Bank Failures by Leverage Ratios: Two-Year Failure Horizon

Cutoff Percentile	Cutoff Point	Failures 1989-93	Nonfailures 1989-93	Failure Rate for Row (Percent)	Cumulative Proportion of Nonfailures (Type II Error) (Percent)	Cumulative Proportion of Failures (Type I Error) (Percent)
			Absolute Scale			
—	0	24	15	61.5	0.0	94.8
—	1.0	28	19	59.6	0.1	88.8
—	2.0	43	36	54.4	0.2	79.6
—	3.0	44	107	29.1	0.5	70.1
—	4.0	60	227	20.9	1.2	57.2
—	5.0	69	428	13.9	2.4	42.4
—	6.0	71	1,391	4.9	6.4	27.1
—	7.0	57	4,001	1.4	17.9	14.8
—	8.0	32	6,627	0.5	37.0	8.0
—	9.0	9	6,285	0.1	55.1	6.0
—	10.0	6	4,714	0.1	68.6	4.7
—	11.0	6	3,242	0.2	78.0	3.4
—	12.0	5	2,190	0.2	84.3	2.4
—	Infinity	11	5,462	0.2	100.0	0.0
			Relative Scale			
1	3.11	154	198	43.8	0.6	66.9
2	4.22	63	289	17.9	1.4	53.3
3	4.93	44	308	12.5	2.3	43.9
4	5.31	25	327	7.1	3.2	38.5
5	5.59	23	329	6.5	4.2	33.5
6	5.80	14	338	4.0	5.1	30.5
7	5.97	13	339	3.7	6.1	27.7
8	6.10	10	342	2.8	7.1	25.6
9	6.22	10	342	2.8	8.1	23.4
10	6.33	8	344	2.3	9.1	21.7
25	7.29	39	4,891	0.8	23.2	13.3
50	8.60	31	8,771	0.4	48.4	6.7
75	10.49	12	8,791	0.1	73.7	4.1
100	Infinity	19	9,135	0.2	100.0	0.0

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Noncumulative data are for the range defined by cutoffs in the current and the previous row. Cumulative data are aggregated up to the cutoff point.

Table 7

Distribution of Bank Failures by Gross Revenue Ratios: Two-Year Failure Horizon

Cutoff Percentile	Cutoff Point	Failures 1989-93	Nonfailures 1989-93	Failure Rate for Row (Percent)	Cumulative Proportion of Nonfailures (Type II Error) (Percent)	Cumulative Proportion of Failures (Type I Error) (Percent)
			Absolute Scale			
—	0	24	15	61.5	0.0	94.8
—	10	30	25	54.5	0.1	88.4
—	20	51	82	38.3	0.4	77.4
—	30	50	183	21.5	0.9	66.7
—	40	65	311	17.3	1.8	52.7
—	50	69	494	12.3	3.2	37.8
—	60	64	1,183	5.1	6.6	24.1
—	70	49	2,545	1.9	13.9	13.5
—	80	25	3,998	0.6	25.4	8.2
—	90	10	4,628	0.2	38.8	6.0
—	100	5	4,429	0.1	51.5	4.9
—	110	3	3,840	0.1	62.6	4.3
—	120	3	2,988	0.1	71.2	3.7
	Infinity	17	10,023	0.2	100.0	0.0
			Relative Scale			
1	26.19	130	222	36.9	0.6	72.0
2	36.63	69	283	19.6	1.5	57.2
3	44.64	58	294	16.5	2.3	44.7
4	50.08	32	320	9.1	3.2	37.8
5	53.72	21	331	6.0	4.2	33.3
6	56.48	24	328	6.8	5.1	28.2
7	58.90	10	342	2.8	6.1	26.0
8	60.93	15	337	4.3	7.1	22.8
9	62.56	6	346	1.7	8.1	21.5
10	63.99	8	344	2.3	9.1	19.8
25	78.24	49	4,881	1.0	23.1	9.2
50	97.39	19	8,783	0.2	48.4	5.2
75	123.78	8	8,795	0.1	73.7	3.4
100	Infinity	16	9,138	0.2	100.0	0.0

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Noncumulative data are for the range defined by cutoffs in the current and the previous row. Cumulative data are aggregated up to the cutoff point.

Table 8

Distribution of Bank Failures by Risk-Weighted Capital Ratios: Two-Year Failure Horizon

Cutoff Percentile	Cutoff Point	Failures 1989-93	Nonfailures 1989-93	Failure Rate for Row (Percent)	Cumulative Proportion of Nonfailures (Type II Error) (Percent)	Cumulative Proportion of Failures (Type I Error) (Percent)
	Absolute Scale					
—	0	24	16	60.0	0.0	94.8
—	1.0	18	10	64.3	0.1	91.0
—	2.0	22	11	66.7	0.1	86.2
—	3.0	32	27	54.2	0.2	79.4
—	4.0	39	68	36.4	0.4	71.0
—	5.0	34	125	21.4	0.7	63.7
—	6.0	49	156	23.9	1.2	53.1
—	7.0	46	306	13.1	2.1	43.2
—	8.0	58	546	9.6	3.6	30.8
—	9.0	38	974	3.8	6.4	22.6
—	10.0	37	1,784	2.0	11.6	14.6
—	11.0	15	2,533	0.6	18.9	11.4
—	12.0	10	2,880	0.3	27.2	9.2
—	Infinity	43	25,308	0.2	100.0	0.0
	Relative Scale					
1	4.62	150	202	42.6	0.6	67.7
2	6.30	80	272	22.7	1.4	50.5
3	7.15	41	311	11.6	2.3	41.7
4	7.73	34	318	9.7	3.2	34.4
5	8.22	25	327	7.1	4.1	29.0
6	8.58	14	338	4.0	5.1	26.0
7	8.88	14	338	4.0	6.1	23.0
8	9.15	10	342	2.8	7.0	20.9
9	9.35	10	342	2.8	8.0	18.7
10	9.55	7	345	2.0	9.0	17.2
25	11.52	32	4,898	0.6	23.1	10.3
50	14.66	28	8,774	0.3	48.4	4.3
75	19.73	12	8,791	0.1	73.7	1.7
100	Infinity	8	9,146	0.1	100.0	0.0

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Noncumulative data are for the range defined by cutoffs in the current and the previous row. Cumulative data are aggregated up to the cutoff point.

rates for individual ranges as the ratio increases. However, the failure rates for adequately capitalized bank observations are now considerably greater. In particular, the failure rate for observations in the 4 to 5 percent leverage ratio class is 13.9 percent, as compared with the 4.7 percent one-year rate. Similarly, the failure rate is now 21.4 percent for observations in the 4 to 5 percent risk-weighted ratio range, as compared with a one-year rate of 14.7 percent. Overall, in the metric of a second-year horizon, the three capital ratios perform quite similarly, although the likelihood of failure is somewhat harder to estimate than in the case of a one-year horizon.[6]

Qualitative Forecasts and the Probability of Failure

Our simple frequency distribution analysis shows that the three alternative measures of capital adequacy perform equally well in identifying failure. In this section, we employ parametric models of bank failure to examine more formally the conditional relationship between the likelihood of failure and the capital ratios. The simplest way to analyze bank failure is to use a qualitative response model. In this model, the dependent variable takes discrete outcomes (in our case, failure or nonfailure). We first estimate the likelihood of failure using a discrete logit model. Estimating the model over the entire panel may lead to biased estimates because the typical logit specification assumes that the event of failure is independent over time. To avoid the apparent time-dependency in the observations, we have estimated the logit model cross-sectionally for each year from 1989 to 1993. In addition to these cross-sectional regressions, we analyze our sample of banks using a proportional hazard model. This model of survival will enable us to better estimate the conditional likelihood of failure over time.

The primary objective of the cross-sectional qualitative choice model is to evaluate how consistently these alternative capital ratios predict failure over time. In this framework, the dependent variable is the probability of failure in a given year, and the explanatory variables are the leverage ratio, the gross revenue ratio, and the risk-weighted ratio. Although many other balance-sheet and income-statement explanatory variables are relevant in predicting bank failure, we focus on the three capital ratios because our main purpose is not to build a failure-prediction model but instead to compare the effectiveness of various capital ratios.[7]

Table 9 reports the results of cross-sectional logit regressions for each year between 1989 and 1993. Overall, the logit analysis shows that all three alternative capital ratios predict fairly accurately failures occurring within one year. When each capital ratio is entered separately in the regression (models 1-3), the model coefficients are, without exception, statistically significant at the 1 percent level. Looking at the concordance ratios, we observe that the logit models based solely on capital ratios can accurately predict failures.[8] The predictive performance of these capital measures is fairly robust over time. Among the three capital ratios, the leverage ratio generally achieves the highest pseudo-R^2 and concordance ratio.[9] The difference in these forecasting efficiency measures among the alternative capital ratios, however, is very small. When all three capital ratios are included together in the logit regression (model 4), the gross revenue ratio appears to have the highest significance overall. Not surprisingly, the sign and magnitude of the regression coefficients in model 4 are less stable across the different years of estimation because of the high degree of collinearity between the three capital measures. Consequently, the interpretation of the logit coefficients is quite difficult in this joint model. As Table 2 shows, however, one advantage of the gross revenue ratio is that it is relatively less correlated with the other two competing capital ratios, meaning that it has the potential to add, on average, more information in the joint regression.

Our simple frequency distribution analysis shows that the three alternative measures of capital adequacy perform equally well in identifying failure.

The relative performance of the risk-weighted ratio improves when the time horizon is extended to between one and two years (Table 10). The risk-weighted ratio outperforms the leverage ratio by small margins in terms of both the pseudo-R^2 and the concordance ratio. On the other hand, the gross revenue ratio performs about as well as the risk-weighted ratio, especially when all three ratios are included.

Based on these regression results, simple capital ratios (the leverage ratio and the gross revenue ratio) appear to predict bank failure as well as the risk-weighted capital ratio, especially over short time horizons. A noteworthy finding is the strong performance of the gross revenue ratio in regressions that include all three variables. One explanation for the strong significance of the gross revenue measure may be that the ratio, in contrast to the others, draws independent information about financial flows from both balance sheets and income statements.

Table 9

Logit Regressions

Dependent Variable: Failure in Less Than One Year

1989	Model 1	Model 2	Model 3	Model 4
Intercept	-0.0878	-0.2646	-0.3497	-0.0901
	(0.5450)	(0.0591)	(0.0126)	(0.5345)
Leverage ratio	-77.8819			-74.4450
	(0.0001)			(0.0001)
Gross revenue ratio		-7.2188		0.0093
		(0.0001)		(0.9588)
Risk-weighted ratio, tier 1			-46.5865	-2.0587
			(0.0001)	(0.6595)
Pseudo-R^2	0.1190	0.1120	0.1101	0.1191
Concordant (percent)	98.0	97.7	97.0	98.1
Discordant (percent)	1.5	1.7	2.1	1.5
Tie (percent)	0.4	0.6	0.9	0.4
Failures	195			
Nonfailures	13,104			

1990	Model 1	Model 2	Model 3	Model 4
Intercept	0.3984	0.2650	0.1679	0.3967
	(0.0179)	(0.1007)	(0.2992)	(0.0182)
Leverage ratio	-96.0482			-49.5560
	(0.0001)			(0.0194)
Gross revenue ratio		-10.0654		-5.0353
		(0.0001)		(0.0258)
Risk-weighted ratio, tier 1			-58.8834	0.7287
			(0.0001)	(0.7317)
Pseudo-R^2	0.1350	0.1330	0.1269	0.1359
Concordant (percent)	97.6	96.7	97.8	97.3
Discordant (percent)	1.1	1.2	1.1	1.1
Tie (percent)	1.2	2.1	1.1	1.6
Failures	161			
Nonfailures	12,742			

1991	Model 1	Model 2	Model 3	Model 4
Intercept	-0.3688	-0.2871	-0.4797	-0.2754
	(0.0260)	(0.0781)	(0.0034)	(0.0939)
Leverage ratio	-74.3724			-0.4529
	(0.0001)			(0.9353)
Gross revenue ratio		-8.2146		-8.0113
		(0.0001)		(0.0001)
Risk-weighted ratio, tier 1			-46.6516	-0.9220
			(0.0001)	(0.7826)
Pseudo-R^2	0.0790	0.0756	0.0648	0.0757
Concordant (percent)	97.5	97.5	97.4	97.5
Discordant (percent)	1.5	1.3	1.5	1.3
Tie (percent)	1.0	1.1	1.1	1.1
Failures	122			
Nonfailures	12,266			

1992	Model 1	Model 2	Model 3	Model 4
Intercept	0.5121	0.4550	0.2586	0.5875
	(0.0166)	(0.0242)	(0.2099)	(0.0057)
Leverage ratio	-87.2859			-7.2337
	(0.0001)			(0.3267)
Gross revenue ratio		-8.8321		-7.9533
		(0.0001)		(0.0001)
Risk-weighted ratio, tier 1			-52.4554	-1.8505
			(0.0001)	(0.5221)
Pseudo-R^2	0.0832	0.0770	0.0665	0.0781
Concordant (percent)	96.0	92.7	91.9	92.8
Discordant (percent)	2.4	3.2	4.0	3.1
Tie (percent)	1.6	4.1	4.1	4.1
Failures	114			
Nonfailures	11,827			

1993	Model 1	Model 2	Model 3	Model 4
Intercept	-2.4270	0.0234	-2.3277	0.0534
	(0.0001)	(0.9416)	(0.0001)	(0.8761)
Leverage ratio	-40.6257			2.4996
	(0.0001)			(0.3609)
Gross revenue ratio		-7.9371		-7.8714
		(0.0001)		(0.0001)
Risk-weighted ratio, tier 1			-25.8946	-2.0740
			(0.0001)	(0.2988)
Pseudo-R^2	0.0192	0.0290	0.0157	0.0293
Concordant (percent)	91.4	92.9	93.8	92.9
Discordant (percent)	4.8	2.2	3.4	2.2
Tie (percent)	3.8	5.0	2.8	5.0
Failures	42			
Nonfailures	11,431			

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Numbers in parentheses are p-values. Pseudo-R^2 is defined in endnote 9. See also Estrella (1998).

This regression finding provides evidence that the gross revenue ratio can effectively supplement more complicated capital ratios.

Thus far, we have focused on the capacity of the capital measures to predict failure over shorter time horizons. One would expect that the efficacy of these regulatory capital ratios might deteriorate if we evaluate their forecasting ability beyond the one- or two-year horizon. Peek and Rosengren (1997) point out that most banks that failed during the New England banking crisis of 1989-93 were well capitalized two years before failure. Similarly, Jones and King (1995) argue that between 1984 and 1989 most troubled banks would not have been classified as undercapitalized under the FDICIA rules. Those studies suggest that prompt corrective action rules mandated by FDICIA would have been ineffective in dealing with banking problems during those periods.

Table 10

Logit Regressions

Dependent Variable: Failure between One and Two Years

1990	Model 1	Model 2	Model 3	Model 4
Intercept	-0.1870	-0.4087	-0.5030	-0.2442
	(0.2954)	(0.0177)	(0.0034)	(0.1774)
Leverage ratio	-62.1593			-22.2474
	(0.0001)			(0.0437)
Gross revenue ratio		-5.7019		-0.6953
		(0.0001)		(0.4567)
Risk-weighted ratio, tier 1			-36.5074	-19.7745
			(0.0001)	(0.0001)
Pseudo-R^2	0.0437	0.0425	0.0449	0.0466
Concordant (percent)	86.7	87.1	88.8	88.8
Discordant (percent)	10.4	10.0	8.6	8.8
Tie (percent)	2.9	2.8	2.6	2.4
Failures	167			
Nonfailures	12,550			

1991	Model 1	Model 2	Model 3	Model 4
Intercept	-0.9504	-0.6484	-0.9654	-0.6917
	(0.0001)	(0.0010)	(0.0001)	(0.0007)
Leverage ratio	-50.6460			18.8294
	(0.0001)			(0.0001)
Gross revenue ratio		-5.9608		-4.6201
		-0.0001		(0.0001)
Risk-weighted ratio, tier 1			-31.9536	-19.3007
			(0.0001)	(0.0002)
Pseudo-R^2	0.0191	0.0278	0.0252	0.0299
Concordant (percent)	86.1	87.2	89.7	88.4
Discordant (percent)	10.9	9.6	8.1	8.4
Tie (percent)	3.0	3.2	2.2	3.2
Failures	125			
Nonfailures	12,205			

1992	Model 1	Model 2	Model 3	Model 4
Intercept	-0.6818	-0.8511	-0.5561	-0.6623
	(0.0016)	(0.0001)	(0.0079)	(0.0027)
Leverage ratio	-56.1702			19.9661
	(0.0001)			(0.0805)
Gross revenue ratio		-5.5291		-0.4750
		(0.0001)		(0.5797)
Risk-weighted ratio, tier 1			-37.8934	-47.2949
			(0.0001)	(0.0001)
Pseudo-R^2	0.0236	0.0242	0.0302	0.0305
Concordant (percent)	88.1	87.4	89.4	88.4
Discordant (percent)	9.3	9.7	8.2	8.5
Tie (percent)	2.6	2.9	2.4	3.1
Failures	119			
Nonfailures	11,702			

1993	Model 1	Model 2	Model 3	Model 4
Intercept	-2.4512	-1.7406	-2.1743	-1.6986
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Leverage ratio	-41.4685			13.1137
	(0.0001)			(0.0419)
Gross revenue ratio		-5.2671		-4.3610
		(0.0001)		(0.0001)
Risk-weighted ratio, tier 1			-28.6207	-14.0894
			(0.0001)	(0.0741)
Pseudo-R^2	0.0048	0.0091	0.0072	0.0097
Concordant (percent)	79.0	85.0	83.2	85.5
Discordant (percent)	11.6	8.1	9.4	7.8
Tie (percent)	9.4	6.9	7.4	6.7
Failures	43			
Nonfailures	11,292			

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Numbers in parentheses are p-values. Pseudo-R^2 is defined in endnote 9. See also Estrella (1998).

Despite the evidence that the performance of capital ratios is not very good at more distant horizons, our analysis suggests that these measures are actually able to disseminate useful signals long before the event of failure. For one, we find that failing banks begin to show signs of weakness (that is, become undercapitalized) two to three years before they are closed by supervisors. The chart presents the time-profile of the three capital ratios for failed banks, plotted according to the number of quarters before failure. The figure also includes analogous measures for a control sample of nonfailed banks. The control group consists of randomly chosen banks located in the same state and having an asset size similar to that of the banks in the failed group.

As the chart shows, the median capital ratios for the group of failed banks are consistently lower than the median ratios for the control sample of surviving banks. The shaded area in each panel of the figure represents the critical region for a one-sided test of equality. When the median capital ratio for the control group (dashed line) is in the shaded area, we cannot reject the hypothesis that the median capital ratios for the two groups are the same at the 1 percent level. For the most part, the median capital ratio for the control group of nonfailed banks is outside the shaded critical region, suggesting that all three capital ratios are fairly good predictors of failure even as far back as two to three years.

Another simple but interesting way to test the long-run effectiveness of the capital ratios in predicting failure is hazard analysis. Although the hazard specification is closely related to binary models such as logit or probit models, it offers a better way to analyze the apparent time-dependency in the conditional probability of failure. More specifically, the dependent variable in hazard analysis is the probability that an institution will fail given that it has not failed until that point of time.[10] Thus, in contrast to the cross-sectional logit model that examines failure over shorter horizons, the proportional hazard specification analyzes the conditional likelihood farther into the future. To simplify our analysis, Table 11 examines two scenarios of survival. The top panel of the table evaluates the efficacy of capital ratios in forecasting the probability of failure from the first quarter of 1988. In this case, the implied dependent variable is the duration of time from the first quarter of 1988 until the bank fails or until the fourth quarter of 1993 for nonfailing banks (so-called censored observations). The explanatory variables in the hazard models (models 1-4) consist of the competing capital adequacy ratios as of the first quarter of 1988. Thus, in contrast to the yearly logit regression, which estimates the effectiveness of capital ratios in forecasting failure within one year or between one and two years, the hazard regressions evaluate the early warning capacity of the capital measures from the first quarter of 1988. To account for the economic downturn in 1990, the bottom panel of Table 11 also estimates the probability of bank failure from the first quarter of 1990.

Capital Ratios before Failure

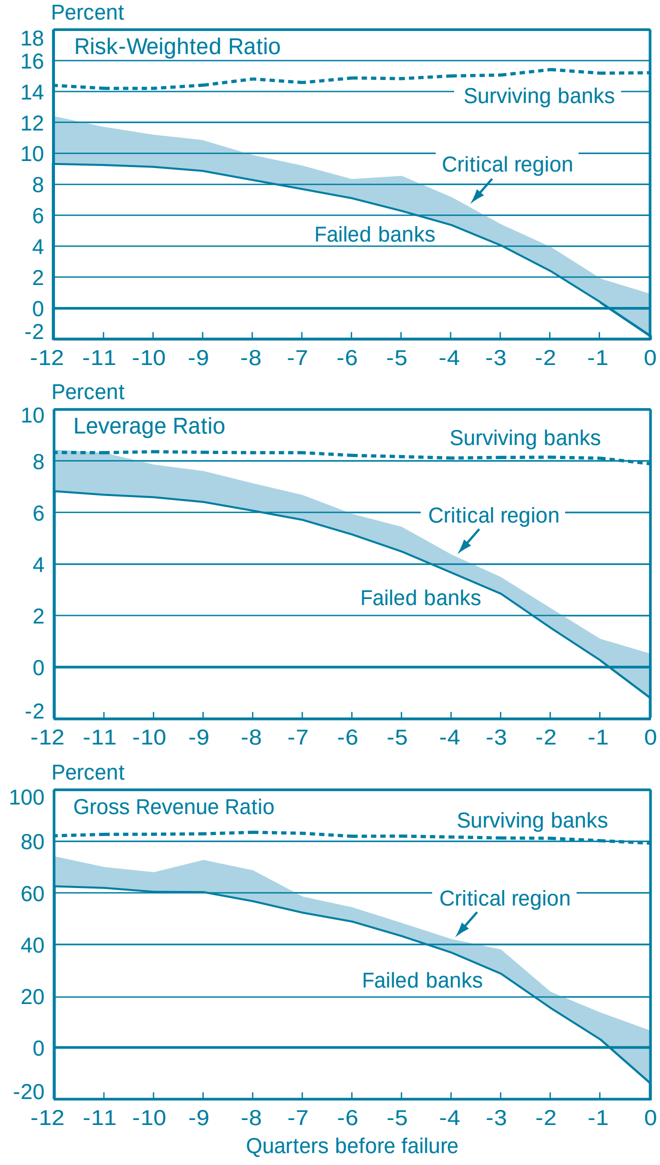

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: The shaded area represents a 1 percent critical region of equality for failed and surviving banks. When the dashed line is outside the shaded area, the population median of surviving banks is statistically greater than the population median of failed banks.

Table 11

Cross-Sectional Proportional Hazard Analysis

Capital as of 1988:1

	Model 1	Model 2	Model 3	Model 4
Leverage ratio	-22.7113			7.3990
	(0.0001)			(0.0207)
Gross revenue ratio		-1.8274		0.0263
		(0.0001)		(0.7664)
Risk-weighted ratio, tier 1			-13.3307	-18.5383
			(0.0001)	(0.0001)
Pseudo-R^2	0.0320	0.0240	0.0470	0.0620
Model χ^2	280.789	213.624	418.078	549.112
Failures	475			
Nonfailures (censored)	8,189			

Capital as of 1990:1

	Model 1	Model 2	Model 3	Model 4
Leverage ratio	-10.3256			21.5610
	(0.0001)			(0.0001)
Gross revenue ratio		-1.3772		-1.7109
		(0.0001)		(0.0001)
Risk-weighted ratio, tier 1			-10.4816	-17.7188
			(0.0001)	(0.0001)
Pseudo-R^2	0.0310	0.0430	0.0550	0.0740
Model χ^2	269.647	382.794	487.820	660.746
Failures	326			
Nonfailures (censored)	8,348			

Sources: Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income; Board of Governors of the Federal Reserve System, National Information Center database; authors' calculations.

Notes: Numbers in parentheses are p-values. Pseudo-R^2 is defined in endnote 9. See also Estrella (1998).

It is clear from the estimated hazard that capital ratios continue to be fairly good predictors of failure even over longer time horizons. When each capital ratio is entered individually in the hazard regression (models 1-3), we find that all three capital ratios are again statistically significant at the 1 percent level. As the pseudo-R^2 statistics indicate, the explanatory power of these capital measures is lower than that obtained with a one-year horizon (Table 9). This finding is not surprising, because the controls are now asked to forecast the likelihood of failure over a longer duration, sometimes as long as six years.

The risk-based measure shows a relatively high pseudo-R^2 in the hazard models separately estimating the effect of each capital ratio and also shows high statistical significance in the models including all three capital ratios. The good performance of the risk-based capital ratio is more pronounced in the analysis using a longer time horizon (top panel). The statistical significance of the gross revenue ratio is comparable to that of the risk-weighted ratio in model 4 of the bottom panel, using a shorter time horizon. This finding is consistent with the result of the logit analysis, which shows that the relative performance of the risk-weighted ratio improves over a longer time horizon.

The risk-weighted ratio takes into account the riskiness of assets, and the gross revenue ratio reflects the asset risk to the extent that riskier assets have higher expected returns. The results in a longer time horizon are more consistent with these expectations. Risk weighting is an attempt to reflect heterogeneous return variances across assets. In a short time horizon, however, differences in return variances across assets may not be significant. For example, the probability that default occurs within a month may be very low even for a risky loan that is highly likely to default within three years. Thus, a possible explanation for the improved performance of the risk-weighted capital ratio over a longer time horizon is that the realization of differences in asset return variances takes time. This possibility also implies that in a short time horizon, risk weighting can overstate differences in asset return variances and hence reduce the accuracy of the risk-weighted ratio as a measure of capital adequacy.

Conclusion

This article compares the effectiveness of different types of capital ratios in predicting bank failure. An important result of our study is that simple ratios—specifically the leverage ratio and the ratio of capital to gross revenue—predict bank failure about as well as the more complex risk-weighted ratio over one- or two-year time horizons. This finding suggests that bank regulators may find a useful role for the simple ratios in the design of regulatory capital frameworks, particularly as indicators of the need for prompt supervisory action. Risk-weighted ratios, in contrast, tend to perform better over longer horizons.

Our intention, however, is not to argue against the use of more sophisticated measures of capital adequacy in regulation. On the contrary, we suggest that simple capital ratios may not be well suited for the determination of optimum levels of bank capital. However, we show that simple capital ratios contain useful information and are virtually costless to compute. Thus, it may be possible to derive substantial benefits from the use of simple ratios—for instance, as supplementary or backstop requirements—even when more sophisticated measures are available for use in formulating the primary requirements.

Endnotes

1. If banks prefer riskier assets (moral hazard), they might choose riskier borrowers within the highest risk-weight category. This effect, however, is unlikely to be large enough to offset the primary effect of reducing assets in the highest risk-weight category.

2. Note that the proportions of failures and nonfailures are cumulated in opposite orders. For instance, the cumulative proportion of nonfailures for the leverage ratio class of 2 percent is 0.5 percent. This proportion is the total number of surviving banks up to and including that class (51+62+95=208), divided by the aggregate number of surviving banks (43,643). In contrast, the cumulative proportion of failures for this same leverage ratio class is 33.0 percent. This value is equal to the cumulative number of bank failures for all banks with a leverage ratio greater than 2 percent (76+45+31+25+17+8+3+2=131), divided by 628, the total number of failures.

3. Technically, the criterion for critically undercapitalized banks uses tangible equity as a measure of capital, instead of tier 1, as in the leverage ratio. To economize on data reporting and to make results more comparable within the article, we base our illustrations on Table 3, which is based on the leverage ratio. Tangible equity ratios produce similar results.

4. Equality of Type I and Type II errors is an interesting illustrative benchmark, but regulators can clearly choose different levels of this trade-off to suit their goals and preferences.

5. Tier 2 includes loan-loss reserves and a number of convertible and subordinated debt instruments. Banks are allowed to use loan-loss reserves up to a maximum of 1.25 percent of risk-weighted assets.

6. If p is the estimated proportion (failure rate), a measure of the variance of the estimate is given by $p(1-p)/n$, where n is the number of observations. This variance is larger when p is closer to ½ and n is smaller, both of which apply in the case of second-year rates as compared with one-year rates.

7. Early warning models use various balance-sheet and income-statement variables to predict bank failure (see, for example, Cole, Cornyn, and Gunther [1995], Cole and Gunther [1995], and Thompson [1991]). Capital adequacy is highly significant in those models. Nevertheless, high correlation among variables reflecting financial strength makes it difficult to infer the significance of individual variables.

8. The concordance ratio is calculated based on the pair-wise comparison of failure probabilities estimated by a logit model. The estimated probability for each failure is compared with those for nonfailure ($m \times [n-m]$ pairs when there are m failures out of n observations). A pair is counted as concordant if the estimated probability is higher for the failed one and discordant in the opposite case. Thus, a high concordance ratio indicates that the logit model accurately classifies failure and nonfailure.

9. The pseudo-R^2 is defined as in Estrella (1998) by $1-(\log L_u/\log L_c)^{-2\log L_c/n}$, where L_u is the value of the unconstrained likelihood, L_c is the value of the likelihood with only a constant term in the model, and n is the number of observations.

10. Because bank failure is a terminal event, the probability of bank failure at time τ given that it has not failed until that point in time or hazard rate is $h(\tau, x) = f(\tau, x)/(1-F(\tau, x))$, where $F(\tau, x)$ is the cumulative probability of failure up to time τ. The proportional hazard specification assumes that the hazard function is separable, that is, $h(\tau, x) = h_0(\tau)\exp[x\beta]$, where x is a vector of explanatory variables and $h_0(\tau)$ is the baseline hazard function.

References

Basel Committee on Banking Supervision. 1988. "International Convergence of Capital Measurement and Capital Standards." Bank for International Settlements, July.

———. 1999. "A New Capital Adequacy Framework," June.

Berger, Allen N., Richard J. Herring, and Giorgio P. Szegö. 1995. "The Role of Capital in Financial Institutions." Journal of Banking and Finance 19 (June): 393-430.

Cole, Rebel A., Barbara G. Cornyn, and Jeffery W. Gunther. 1995. "FIMS: A New Monitoring System for Banking Institutions." Federal Reserve Bulletin 81, no. 1 (January).

Cole, Rebel A., and Jeffery W. Gunther. 1995. "Separating the Likelihood and Timing of Bank Failure." Journal of Banking and Finance 19, no. 6 (September): 1073-89.

Demsetz, Rebecca S., Marc R. Saidenberg, and Philip E. Strahan. 1996. "Banks with Something to Lose: The Disciplinary Role of Franchise Value." Federal Reserve Bank of New York Economic Policy Review 2, no. 2 (October): 1-14.

Estrella, Arturo. 1995. "A Prolegomenon to Future Capital Requirements." Federal Reserve Bank of New York Economic Policy Review 1, no. 2 (July): 1-12.

———. 1998. "A New Measure of Fit for Equations with Dichotomous Dependent Variables." Journal of Business and Economic Statistics 16 (April): 198-205.

Gilbert, Alton, Courtenay Stone, and Michael Trebing. 1985. "The New Bank Capital Adequacy Standards." Federal Reserve Bank of St. Louis Review 67, no. 5 (May): 12-20.

Jones, David S., and Kathleen Kuester King. 1995. "The Implementation of Prompt Corrective Action: An Assessment." Journal of Banking and Finance 19 (June): 491-510.

Mitchell, Wesley C. 1909. "The Decline in the Ratio of Banking Capital to Liabilities." Quarterly Journal of Economics 23 (August): 697-713.

Peek, Joe, and Eric S. Rosengren. 1997. "How Well-Capitalized Are Well-Capitalized Banks?" New England Economic Review, September-October: 41-50.

Shepheard-Walwyn, Tim, and Robert Litterman. 1998. "Building a Coherent Risk Measurement and Capital Optimisation Model for Financial Firms." Federal Reserve Bank of New York Economic Policy Review 4, no. 3 (October): 171-82.

Thompson, James B. 1991. "Predicting Bank Failures in the 1980s." Federal Reserve Bank of Cleveland Economic Review, first quarter: 9-20.

Carol Osler

Support for Resistance: Technical Analysis and Intraday Exchange Rates

- Among the technical trading signals supplied to customers by foreign exchange trading firms are "support" and "resistance" levels. These levels indicate points at which an exchange rate trend is likely to be interrupted or reversed.

- A rigorous test of the levels specified by six trading firms during the 1996-98 period reveals that these signals were quite successful in predicting intraday trend interruptions.

- Although all six firms were able to identify turning points in exchange rate trends, some firms performed markedly better than others. As a group, the firms predicted turning points in the dollar-yen and dollar-pound exchange rates more accurately than turning points in the dollar-mark exchange rate.

- In addition, the predictive power of the support and resistance levels appeared to last at least five business days after they were first communicated to customers.

Early in the morning of each business day, the major foreign exchange trading firms send their customers lists of technical trading signals for that day. Timely technical signals are also supplied by major real-time information providers. These signals, which are based primarily on prior price and volume movements, are widely used by active foreign exchange market participants for speculation and for timing their nonspeculative currency transactions. In fact, 25 to 30 percent of foreign exchange traders base most of their trades on technical trading signals (Cheung and Chinn 1999; Cheung and Wong 1999). More broadly, technical analysis is used as either a primary or secondary source of trading information by more than 90 percent of foreign exchange market participants in London (Allen and Taylor 1992) and Hong Kong (Lui and Mole 1998).

The technical trading signals provided to customers vary over time and across technical analysts, but the vast majority of the daily technical reports include "support" and "resistance" levels. According to technical analysts, support and resistance levels are points at which an exchange rate trend is likely to stop and may be reversed. For example, a firm publishing a support level of \$1.50/£ would claim that the dollar-pound exchange rate is likely to stop falling if it reaches \$1.50/£. If the firm also provided another support level of \$1.45/£, the firm would claim that if the exchange rate passes through \$1.50/£, it is likely to stop falling at \$1.45/£.

Carol Osler is a senior economist at the Federal Reserve Bank of New York.

The author thanks Franklin Allen, Alain Chaboud, and Charles Jordan for thoughtful comments, Gijoon Hong for excellent research assistance, and two anonymous referees. The views expressed are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

Despite the almost universal use of support and resistance levels in short-term exchange rate forecasting, the ability of these trading signals to predict intraday trend interruptions has never been rigorously evaluated. This article undertakes such a test, using actual support and resistance levels published daily by six firms from January 1996 through March 1998. The firms include commercial banks, investment banks, and real-time information providers based in the United States and abroad. I examine the value of three currencies relative to the U.S. dollar: the German mark, the Japanese yen, and the British pound. Support and resistance levels for these exchange rates are tested against indicative exchange rate quotes sampled at one-minute intervals between 9 a.m. and 4 p.m. New York time.

These tests strongly support the claim that support and resistance levels help predict intraday trend interruptions for exchange rates. All six of the firms studied were able to identify points where intraday trends were likely to end. However, some firms were better than others at identifying such points.

Despite the almost universal use of support and resistance levels in short-term exchange rate forecasting, the ability of these trading signals to predict intraday trend interruptions has never been rigorously evaluated.

For most firms, the predictive power of support and resistance levels lasted at least five business days beyond the levels' publication date. Despite their overall success at identifying points of trend interruptions, none of the firms correctly assessed the relative likelihood of trend interruptions at the different levels. These results are consistent across firms and are sustained over a number of sensitivity analyses.

The statistical tests are based on the bootstrap technique (Efron 1979, 1982), a nonparametric method frequently used to evaluate technical trading strategies (Brock et al. 1992; Levich and Thomas 1993). To implement the tests, I compare the behavior of exchange rates upon reaching published support and resistance levels with the behavior upon reaching 10,000 sets of arbitrarily chosen support and resistance levels. If the outcome associated with the actual levels exceeds the average outcome for the arbitrary levels in a high proportion of months, I conclude that the published levels have significant predictive power.

To complement the analysis of these signals' predictive power, I also analyze the signals themselves. I show that support and resistance levels provided by individual firms tend to be fairly stable from day to day. Their range varies very little over time. Firms do not agree extensively with each other on the relevant signals.

The specific conclusion that exchange rates tend to stop trending at support and resistance levels has no precedent in the academic literature. The closest point of comparison is a study by Lo et al. (2000), which finds that the conditional distribution of financial prices is sensitive to the presence of a broad variety of technical trading signals, consistent with the results presented here.

The finding that support and resistance levels are able to predict trend interruptions is consistent with other studies of the usefulness of technical trading rules when applied to currencies. Filter rules were found to be profitable as early as 1984 (Dooley and Shafer 1984), less than a decade into the floating rate period, and this finding has been confirmed repeatedly (Sweeney 1986; Levich and Thomas 1993). Moving-average crossover rules have also been tested frequently on exchange rates, with similar results (Levich and Thomas 1993; Menkhoff and Schlumberger 1995). More recently, Chang and Osler (1998) find that a trading strategy based on the head-and-shoulders chart pattern is profitable for dollar exchange rates vis-à-vis the mark and the yen, although not for four other dollar exchange rates.

This study differs from those earlier studies in four notable ways. First, the technical trading signals used here are intended to anticipate trend reversals, rather than trend continuations. Second, this study uses a type of trading signal that is actively used by market participants. Third, it uses trading signals that were produced by market participants. Other academic studies of technical analysis have typically constructed technical trading signals of their own. Finally, this study uses data sampled at one-minute intervals throughout the New York trading day, while most earlier studies have used data sampled at daily or lower frequencies.

The two existing studies of support and resistance levels—Curcio et al. (1997) and Brock et al. (1992)—test the hypothesis that prices tend to move rapidly once the levels are breached. Curcio et al. find that the hypothesis is not true on average for currencies, but may hold true during periods of strong trending. Brock et al. find that the hypothesis is true for daily movements of the S&P 500 stock index, but the profits may not be sufficient to offset transaction costs. The hypothesis that prices will trend once a trading signal is breached is not unique to support and resistance levels and is not examined here.

Technical Analysis

Technical analysts claim that they can predict financial price movements using an information set limited to a few variables, such as past prices. Many of the major technical indicators were described as early as 1930 by Shabacker, who based his conclusions on observations of U.S. stock prices. By now, technical indicators are widely used in major financial markets around the world, including foreign exchange and futures markets. There are two magazines devoted exclusively to the topic, each of which has more than 40,000 subscribers. To learn about technical analysis, one can consult myriad manuals, software, and on-line sources. Alternatively, one can take courses on technical analysis.

Casual observation and conversations with market participants indicate that support and resistance levels are the most widely used technical indicators in the foreign exchange market. This conclusion is also suggested by the fact that support and resistance levels are the only indicators provided by all six of the technical services covered in this research. In fact, some services provide no technical indicators at all other than support and resistance levels.

Support and Resistance Levels Defined

Before delving further into the analysis, it is important to explore the definition of support and resistance levels provided by technical analysts themselves. According to one major technical analysis manual, "support is a level or area on the chart under the market where buying interest is sufficiently strong to overcome selling pressure. As a result, a decline is halted and prices turn back again. . . . Resistance is the opposite of support" (Murphy 1986, p. 59).

A review of technical analysis manuals reveals that there is little disagreement among analysts on this definition (Arnold 1993; Edwards and Magee 1997; Hardy 1978; Kaufman 1978; Murphy 1986; Pring 1991; Sklarew 1980). For example, Pring states: "Support and resistance represent a concentration of demand and supply sufficient to halt a price move at least temporarily" (p. 199). Likewise, Arnold (1993) observes: "A support level is a price level at which sufficient demand exists to at least temporarily halt a downward movement in prices" (p. 67).[1]

To identify the support and resistance levels relevant for the coming day, practicing technical analysts consult a variety of information inputs. These include visual assessments of recent price performance, simple numerical rules based on recent price performance, inference based on knowledge about order flow, and market psychology.

The simplest approach to visual assessment is to look at recent minima and maxima: "Usually, a support level is identified beforehand by a previous reaction low," and "a resistance level is identified by a previous peak" (Murphy 1986, p. 59). According to Pring (1991), one could also identify support and resistance levels by drawing a trendline, or "channel," in which recent peaks are connected by one line and recent troughs are connected by another: "A good trendline represents an important support and resistance zone" (p. 105).

To identify the support and resistance levels relevant for the coming day, practicing technical analysts consult a variety of information inputs.

One numerical rule used to infer support and resistance levels is the "50 percent rule," which asserts that a major market move will often be reversed by about 50 percent in the first major correction (Pring, p. 187). Fibonacci series, which are widely used, suggest that 38.2 percent and 61.8 percent retracements of recent rises or declines are common.

Market insiders sometimes identify support and resistance levels using private information or information circulated informally in the market about certain market participants. For example, if a technical analyst learned in conversation that Japanese exporters are selling at 100, he or she would report a resistance level at ¥100.00/$. Similarly, if a trader knew that his or her own firm had a large order at DM1.50/$, he or she might expect unusual price behavior at that point.

Simple market psychology is also used to help identify support and resistance levels. According to Murphy (1986): "Traders tend to think in terms of important round numbers . . . as price objectives and act accordingly. These round numbers, therefore, will often act as psychological support or resistance levels." As I will demonstrate later, published support and resistance levels are round numbers that end in 0 or 5 much more often than they would if they were chosen at random.

Some of the firms in the sample provided explanations for their chosen support and resistance levels. All of the approaches listed above are well represented among those explanations.

Properties of the Support and Resistance Database

The data examined here include support and resistance levels for the mark, yen, and pound in relation to the U.S. dollar, published daily by six firms from January 1996 through March 1998. Two of the firms did not report support and resistance levels for the pound. In total, there are approximately 23,700 support and resistance values (combined) for the mark, 22,800 for the yen, and 17,700 for the pound.

The six providers of technical analysis include commercial banks, investment banks, and news services. Some operate in the United States and others operate abroad. The commercial and investment banks provide the information free of charge to their customers, hoping that customers will be encouraged to direct more business toward them. Some news services charge for the information. Since all the providers hope that the usefulness of their signals will generate additional business, they have every incentive to maximize accuracy. In the analysis that follows, the firms are assigned numbers to preserve anonymity.

The data examined here include support and resistance levels for the mark, yen, and pound in relation to the U.S. dollar, published daily by six firms from January 1996 through March 1998.

Table 1

Average Total of Support and Resistance Levels per Reporting Day

Firm	Overall	German Mark	Japanese Yen	British Pound
1	10.0	10.0	10.0	10.0
2	6.0	6.0	6.0	6.0
3	17.8	18.1	18.0	17.5
4	9.0	9.2	8.8	—
5	4.2	4.8	3.8	4.0
6	2.5	2.7	2.3	—

Source: Author's calculations.

Table 2

Average Distance between Support and Resistance Levels

Firm	Overall	German Mark	Japanese Yen	British Pound
1	30	36	29	26
2	61	75	52	55
3	54	58	49	54
4	57	64	49	—
5	162	156	184	144
6	42	47	36	—

Source: Author's calculations.

Note: Distances are measured in points, or 0.0001 marks/dollar, 0.01 yen/dollar, and 0.0001 dollars/pound.

On any given day, technical indicators were likely to be received from about five of the six firms. Firms failed to report for reasons such as vacations, sickness, and equipment problems. For individual firms, the average number of support and resistance levels (combined) that were listed per reporting day per currency ranges from two to eighteen (Table 1).

The support and resistance levels were quite close together for some firms and quite far apart for others. As shown in Table 2, the average distance between levels varied from 30 to 162 points on average (where a point is the smallest unit used in quoting an exchange rate).

Table 3

Average Gap between Current Spot Rates and Outermost Support and Resistance Levels

Firm	German Mark	Japanese Yen	British Pound
1	3.45	2.56	2.38
2	3.71	2.38	2.63
3	8.32	6.73	7.15
4	5.52	3.79	—
5	4.96	5.13	4.52
6	1.24	0.74	—
Daily ranges			
Average	0.45	0.33	0.37
Maximum	3.23	3.88	3.28

Source: Author's calculations.

Note: Distances are measured in units of 100 points, or 0.01 marks/dollar, 1.0 yen/dollar, and 0.01 dollars/pound.

Firms also varied dramatically in the range over which they chose to present support and resistance levels applicable to a given day (Table 3). For Firm 6, the outermost support and resistance levels were typically only about 100 points away from current spot rates, while for Firm 3 the outermost support and resistance levels were typically more than 700 points away. The final two rows of the table show the average and maximum daily exchange rate moves away from their opening rates over the period. For all firms, the outermost support and resistance levels were substantially farther away from the opening rates than this average move. The correlation between daily exchange rate ranges and the gap between opening rates and outermost support and resistance levels was not statistically significant for any firm-currency pair.

Use of Round Numbers

More than 70 percent of the support and resistance levels in the sample end in 0, and a full 96 percent end in either 0 or 5 (Table 4). These proportions greatly exceed the proportions we would observe if levels were chosen randomly, which would be 10 or 20 percent, respectively. Levels ending in 00 or 50 were also disproportionately represented. This may be a manifestation of the psychological interpretation of support and resistance levels mentioned earlier. It is interesting to note that Goodhart and Figliuoli (1991) observed that round numbers were also disproportionately represented in bid-ask spreads for major currencies.

Continuity

To analyze the extent to which support and resistance levels published by a given firm vary from day to day, I counted the number of support and resistance levels shared across days for a given firm, and compared it with the maximum number of levels that could have been shared. That maximum depends on the number of support and resistance levels provided on the two days: if the firm provided three support levels on the first day and four on the second day, the number of shared support levels, or matches, could not possibly exceed three. The maximum number also depends on the size of the exchange rate move from the first to the second day: if the exchange rate falls substantially between days one and two, then some of the support levels provided on day one might be irrelevant on day two. A shared level, or "match," was defined as a pair of support and resistance levels on contiguous days that differed by less than 5 points.

On average, about three-quarters of the still-applicable support and resistance levels from one day would be used again the next day (Table 5). This average masks a clear division of

Table 4
Support and Resistance Levels Ending in Round Numbers
Percent

	Support and Resistance Levels Ending in			
	00	00 or 50	0	0 or 5
Natural frequency	1.0	2.0	10.0	20.0
Firm 1	12.1	17.9	65.8	96.3
Firm 2	13.2	22.4	58.1	84.4
Firm 3	12.5	16.8	82.4	96.6
Firm 4	7.8	15.7	52.9	92.2
Firm 5	49.4	66.4	97.1	99.4
Firm 6	22.9	42.9	74.3	100.0
All firms	13.5	19.6	70.1	95.5

Source: Author's calculations.

Table 5
Continuity in Support and Resistance Levels

Firm	Overall	German Mark	Japanese Yen	British Pound
1	64.4	62.4	62.9	67.9
2	54.0	51.0	56.2	54.5
3	*91.4*	*89.8*	*91.7*	*92.6*
4	*81.9*	*81.2*	*82.7*	—
5	*80.8*	*77.5*	*85.0*	*81.3*
6	56.5	56.0	57.4	—
All firms	77.8	76.5	78.2	79.1

Source: Author's calculations.

Notes: The table shows the percentage of support and resistance levels shared across adjacent days. A pair of levels on adjacent days is defined as shared if the levels differ by at most 5 points. Numbers for firms showing particularly high continuity are italicized.

the firms into two groups. Firms 3, 4, and 5 showed the strongest continuity: more than three-quarters of their still-applicable support and resistance levels were used again the next day. For the remaining three firms, the corresponding proportions were lower, ranging from about one-half to two-thirds. These results do not change qualitatively if a match is defined as two levels within 2 points of each other.

Agreement across Firms

Firms do not agree extensively on the relevant support and resistance levels for a given day. To examine the extent of agreement across firms, I first counted the number of matches across each of the fifteen pairs of firms. For each day, the number of actual matches was then compared with the number of possible matches for that day. For a given day, the number of possible matches among support (resistance) levels was taken to be the minimum number of support (resistance) levels provided across the two firms.

On average, roughly 30 percent of all possible matches were realized as actual matches under the basic definition of a match (a maximum difference of 5 points), as shown in Table 6. Across firm pairs, the frequency of agreement varied from 13 to 38 percent (Table 7). Firm 5 stands out as the least likely to agree with its peers. Firms 1 and 2 stand out as agreeing particularly frequently with each other. Among the other firms, no strong patterns are distinguishable.

Table 6
Overall Agreement on Support and Resistance Levels across Firms

	Overall	German Mark	Japanese Yen	British Pound
Average possible matches per day	112.0	47.9	43.0	21.1
Average actual matches per day	33.5	14.3	12.3	6.9
Average actual matches per day as a percentage of possible matches	29.9	29.9	28.6	32.7

Source: Author's calculations.

Notes: The table shows the number of times all firms' support and resistance levels actually match as a percentage of the total number of possible matches. A match is defined as a pair of support and resistance levels that differ by at most 5 points.

Table 7
Pairwise Agreement across Firms on Support and Resistance Levels
All Currencies

	Firm			
Firm	2	3	4	5
1	*38*	33	27	17
2		35	31	13
3			32	25
4				23
5				

Source: Author's calculations.

Notes: The table shows the number of times a pair of firms' support and resistance levels actually match as a percentage of the number of possible matches. A match is defined as a pair of support and resistance levels that differ by at most 5 points. Numbers representing firm pairs for which agreement falls at or below 23 percent (mean overall agreement of 30 percent minus one standard deviation) are in bold. The italicized number represents firm pairs for which agreement falls at or above 37 percent (mean overall agreement of 30 percent plus one standard deviation).

If a match is defined as two levels within 2 points of each other, then roughly 18 percent of possible matches are actually realized. The same Firm 5 still stands out as the least likely to agree with its peers; the only strong agreement appears to be between Firms 3 and 6.[2]

Exchange Rate Data and Methodology

This section presents the exchange rate data, some important definitions, and the statistical methodology used to test the ability of support and resistance levels to predict intraday trend interruptions.

Exchange Rate Data

The exchange rate data comprise indicative bid-ask rates posted on Reuters, captured at one-minute intervals from 9 a.m. to 4 p.m. New York time. The prices for a given minute were taken to be the last quote made prior to that minute.

The analysis in Goodhart, Ito, and Payne (1996) suggests that these indicative quotes are likely to correspond closely to actual transaction prices. The major divergences between quotes and actual prices seem most likely to occur at times of large, rapid price movements. During the sample period, these divergences often occurred at times of macroeconomic data announcements from the United States, which tended to happen at 8:30 a.m., before the exchange rate data used here begin. Recent research by Danielsson and Payne (1999) finds that quotes may differ from actual transaction prices in other, potentially important ways. This point is discussed in greater detail below.

The construction of the exchange rate data set was driven primarily by the need to capture as closely as possible the price sequence that would be observed by traders operating in the market. This explains why an interval of one minute was selected, rather than the more common interval of five minutes (for example, see Andersen and Bollerslev [1998]). It also explains why prices were not taken as an average of the immediately preceding and following quotes, another common technique in the literature (for example, see Andersen and Bollerslev [1998]): traders operating in real time could not know the immediately following quote.

The construction of the exchange rate data set was driven primarily by the need to capture as closely as possible the price sequence that would be observed by traders operating in the market.

The starting time for the data was chosen as 9 a.m. New York time for two reasons. First, by 9 a.m., the support and resistance levels in the data set have been transmitted to customers, including those from New York firms. Second, by 9 a.m., the reaction to macroeconomic data announcements from all the countries involved—including the United States, where, as noted, major announcements generally occurred at 8:30 a.m. New York time—would largely be over (see Andersen and Bollerslev [1998]).

The data end at 4 p.m. because very little trading takes place between then and the beginning of the next trading day in Asia. The 4 p.m. cutoff was also chosen because, in the underlying tick-by-tick exchange rate data set, quotes are not captured after 4 p.m. on Fridays.

Some Definitions

The exchange rate was defined as hitting a support (resistance) level if the bid (ask) price fell (rose) to within 0.01 percent of that level (see the chart for an illustration). Because the 0.01 percent figure is somewhat arbitrary, more than one definition was tried: the gap was also set at 0.00 percent and 0.02 percent in alternative tests. A trend interruption was defined as follows: once the exchange rate hit a support (resistance) level, the trend was interrupted if the bid (ask) price exceeded (fell short of) the support (resistance) level fifteen minutes later. Since the cutoff at fifteen minutes is also somewhat arbitrary, an alternative of thirty minutes was examined as well.

For brevity, a trend interruption will be referred to frequently as a "bounce," and the ratio of times the exchange rate bounces to the number of actual hits will be referred to as the "bounce frequency." In formal terms, the goal of the test described below is to ascertain statistically whether the bounce frequencies for the published support and resistance levels are high, as claimed by technical analysts.

Statistical Methodology

The statistical test evaluates whether or not published support and resistance levels are able to identify points of likely trend interruptions, as claimed by technical analysts. The central or "null" hypothesis is that the published levels have no special ability to identify such points. I begin with a summary of the methodology and then present details.

Hypothetical Exchange Rate Paths

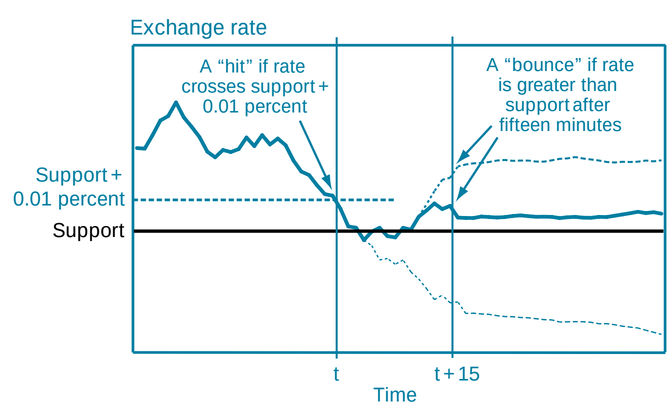

Source: Author's calculations.

Summary

The statistical methodology used to test this null hypothesis is a specific application of the bootstrap technique (Efron 1979, 1982). To apply this technique, I first calculate bounce frequencies for each firm for each month in the sample. I then build a statistical representation of what bounce frequencies for the published support and resistance levels would look like if the null hypothesis were true. In the present context, this representation is constructed by first creating 10,000 sets of artificial support and resistance levels for each day. For each of these artificial sets of support and resistance levels, I then calculate bounce frequencies for each month, using the criteria for hits and bounces listed above.

At this point, I have twenty-eight bounce frequencies for each firm, one for each month of the sample, and twenty-eight average bounce frequencies for the artificial support and resistance levels. In the final step of the test, I determine the number of months in which the bounce frequency for a given firm exceeds the average bounce frequency for artificial levels. If this number of months is quite high, I conclude that the published support and resistance levels have some ability to predict intraday trend interruptions. Additional details on this methodology are presented below.

Calculating Artificial Support and Resistance Levels

For each day, the artificial support and resistance levels are chosen at random from exchange rates within a certain range of the day's opening rate. The range for each month is based on the exchange rate's actual behavior, as follows: for a given month, I calculate the gap between the opening rate and intraday highs and lows for each day. The absolute value of the largest of these gaps is used as the range for calculating artificial support and resistance levels for that month.

For each day, twenty artificial support and twenty artificial resistance levels are calculated using the following algorithm:

$$R_{ti} = O_t + b_{ti}\,range\ ,$$

$$S_{ti} = O_t - a_{ti}\,range\ .$$

Here, t represents time, $S_{ti}(R_{ti})$ is the ith artificial support (resistance) level, O_t is the day's opening rate, a_{ti} and b_{ti} are random numbers generated from a uniform distribution over [0,1], and $range$ is the range for that month. These levels are then rounded off so that they have the same number of significant digits to the right of the decimal point as actual quoted exchange rates.[3]

The Statistical Test

The statistical test is based on comparing the bounce frequencies for the published support and resistance levels (B^P) with the average bounce frequencies for the artificial levels (B^A), month by month. To understand the test intuitively, suppose published support and resistance levels provide no more information than arbitrary levels. In this case, B^P should not consistently be higher or lower than B^A. However, if published support and resistance levels can predict points of likely trend interruptions, as claimed by technical analysts, then B^P should usually be higher than B^A.

This idea can be formalized into a rigorous statistical test. The comparison for each month can be viewed as a "Bernoulli trial," in which a random event occurs with a certain probability. The random event here would be $B^P > B^A$. Under the null hypothesis that published levels have no special predictive power, the likelihood of that event is 50 percent. Over the entire twenty-eight-month sample, if it is true that published levels are not informative, the chance that $B^P > B^A$ for any given number of months will conform to the binomial distribution. This distribution is symmetrical around a single peak at fourteen months, where the probability is about 15 percent.

> *Consistent with the market's conventional wisdom, exchange rates bounced quite a bit more frequently after hitting published support and resistance levels than they would have by chance.*

To understand how to use this distribution, suppose we find that $B^P > B^A$ in twenty of the twenty-eight months of the sample. We might naturally ask: Would it be unusual to get such an extreme outcome if the published levels are truly not informative? More concretely, what is the likelihood, if the published levels are not informative, of finding $B^P > B^A$ in twenty *or more* of the twenty-eight months of the sample? This likelihood is the area under the tail of the distribution to the right of the number 20. This is a very small number: in fact, it is 1.8 percent.

The likelihood of finding $B^P > B^A$ in twenty or more of the twenty-eight months, under the assumption that published levels are not informative, is called the "marginal significance level" associated with the number 20.[4] If the marginal significance level of some result is smaller than 5 percent, it is consistent with standard practice in the literature to conclude that the published numbers are better than arbitrary numbers at predicting trend interruptions. Such a result is said to be "statistically significant."

To summarize our example: it would be extremely unusual to find that $B^P > B^A$ in twenty or more of the twenty-eight months if the published support and resistance levels were truly not informative. In fact, we would realize such an outcome only 1.8 percent of the time. Since 1.8 percent falls below the common critical value of 5 percent, we would conclude that the predictive power of the published levels exceeds that of the arbitrary levels to a statistically significant degree.

Results

Consistent with the market's conventional wisdom, exchange rates bounced quite a bit more frequently after hitting published support and resistance levels than they would have by chance. Exchange rates bounced off arbitrary support and resistance levels 56.2 percent of the time on average.[5] By contrast, they bounced off the published levels 60.8 percent of the time on average (Table 8). Looking more closely, we find that in all sixteen firm-currency pairs, average bounce frequencies for published levels (across the entire sample period) exceeded average bounce frequencies for artificial levels.

The month-by-month breakdown shows that for most firm-currency pairs bounce frequencies for the published levels exceeded average bounce frequencies for artificial levels in twenty or more months. As noted above, these outcomes would be extremely unlikely if the support and resistance levels were truly not informative. More rigorously, the marginal significance levels indicate that the results are statistically significant at the 5 percent level for all but three firm-currency pairs.

The firms' ability to predict turning points in intraday trends seems to have been stronger for the yen and weaker for the mark and the pound. On average, bounce frequencies for published support and resistance levels exceeded those for arbitrary levels by 4.2 percentage points for the mark, 5.6 percentage points for the yen, and 4.0 percentage points for the pound. This relative ranking was maintained fairly consistently for individual firms.

Although all six firms seem to have the ability to predict exchange rate bounces, their performance varied considerably. The bounce frequencies of the best and worst firms differ by 4.0 percentage points on average. At one extreme, Firm 1's support and resistance levels for the yen had a bounce frequency 9.2 percentage points higher than that of the arbitrary levels.

Differences across firms are evaluated statistically in Table 9. Firm 1 is clearly the best overall: it had the highest bounce frequency for two of the three currencies and the second-highest bounce frequency for the third currency. Furthermore, the differences between Firm 1 and the other firms are statistically significant at the 5 percent level in seven of the thirteen possible firm-to-firm comparisons and are statistically significant at the 10 percent level in another comparison. Firm 5 did quite well for the mark, but did not do noticeably well for the other two currencies. No firm was consistently worst.

Table 8

Ability of Support and Resistance Levels to Predict Interruptions of Intraday Exchange Rate Trends

	Artificial Levels	Levels Published by Firm 1	2	3	4	5	6
		Bounce Frequency (Number of Hits)					
German mark	54.9	60.1 (6,291)	56.6 (4,102)	58.0 (8,111)	58.5 (3,570)	62.0 (1,262)	59.3 (2,296)
Japanese yen	57.3	66.5 (4,558)	63.6 (3,874)	62.3 (6,271)	60.7 (2,679)	61.6 (859)	62.6 (1,396)
British pound	56.3	63.0 (5,409)	58.8 (3,920)	59.6 (6,056)		60.0 (1,039)	
		Months $B^P > B^A$/Total Months (Marginal Significance)					
German mark		24/28 (0.000)	17/28 (0.172)	21/28 (0.006)	20/27 (0.010)	20/26 (0.005)	19/28 (0.044)
Japanese yen		26/28 (0.000)	24/27 (0.000)	22/28 (0.002)	23/27 (0.000)	15/23 (0.105)	20/27 (0.010)
British pound		27/28 (0.000)	19/28 (0.044)	24/28 (0.000)		16/26 (0.163)	

Source: Author's calculations.

Notes: The table compares the ability of published support and resistance levels to predict intraday trend interruptions with the distribution of predictive ability for 10,000 sets of arbitrary support and resistance levels. The measure of predictive ability is based on the "bounce frequency," or the number of times the exchange rate stopped trending after reaching support or resistance levels compared with the total number of times the rate actually reached such levels.

The table shows the bounce frequency for published and artificial support and resistance levels, the number of hits, the number of months in which the bounce frequency for published levels (B^P) exceeds the bounce frequency for artificial levels (B^A), and the marginal significance of this number of months under the null hypothesis that published support and resistance levels are not informative. "Total months" varies across firms and currencies because occasionally a firm contributed too few support and resistance levels to have any hits at all.

Table 9

Differences in Firms' Ability to Predict Exchange Rate Bounces

	Firm A				
Firm B	1	2	3	4	5
German mark					
2	3.5**				
3	2.1**	-1.4			
4	1.5	-2.0	-0.6***		
5	-2.0**	-5.5**	-4.1***	-3.5	
6	0.8	-2.7	-1.3**	-0.7	2.8
Japanese yen					
2	2.8*				
3	4.1***	1.3			
4	5.7***	2.9	1.6		
5	4.9	2.0	0.7	-0.9	
6	3.8	1.0	-0.3	-1.9	-1.0
British pound					
2	4.2***				
3	3.5**	-0.8*			
5	3.1*	-1.2	-0.4		

Source: Author's calculations.

Notes: The table compares different firms' ability to predict intraday trend interruptions in exchange rates. The measure of predictive ability is based on the "bounce frequency," or the number of times the exchange rate stopped trending after reaching support or resistance levels compared with the total number of times the rate actually reached such levels. The table presents the difference between bounce frequencies (measured as Firm A minus Firm B).

* Statistically significant at the 10 percent level.
** Statistically significant at the 5 percent level.
*** Statistically significant at the 1 percent level.

Robustness

These results are robust to changes in the test methodology.[6] They are not changed qualitatively if a hit is defined more broadly or more narrowly (as described earlier) or if one looks thirty minutes rather than fifteen minutes beyond a hit. The results are also unchanged if one splits the sample into morning and afternoon sessions (where the morning session is defined to include positions entered before noon).

Interestingly, the results change somewhat if the sample is split in half chronologically. During the first half of the sample period, when volatility was fairly low by historical standards, bounce frequencies were statistically significant for published levels in all but one case. In the second half, when volatility returned to more normal levels, firms' bounce frequencies still exceeded those for the artificial levels, but the differences were no longer statistically significant in half the cases.[7] This outcome is consistent with the market's conventional wisdom that rates tend to "range trade" in periods of low volatility, thus making this type of trend reversal more common.

Quotes versus Transaction Prices

At this point, it is possible to discuss more fully the potential implications of the differences noted by Danielsson and Payne (1999) between exchange rate quotes and actual transaction prices. The first important difference they note is that quotes tend to be more volatile than actual transaction prices. This should not be critical here, because these results concern the direction of price changes, not their magnitude.

Second, Danielsson and Payne (1999) find that quotes tend to be negatively autocorrelated while transaction prices are not. In theory, this could affect the absolute frequency of bounces presented in Table 8. Fortunately, the important qualitative conclusions of the paper are based on the *difference* between bounce frequencies for published and simulated levels, rather than the absolute size of those bounce frequencies. Furthermore, the negative autocorrelation in quote data may largely have dissipated by the end of the fifteen-minute horizon of interest. The reason is that if the exchange rate is required to reach the actual level rather than some nearby level to achieve a hit, bounce frequencies in the simulated support and resistance levels fall slightly short of 50 percent. If negative autocorrelation at the fifteen-minute horizon were an issue, that proportion would presumably exceed 50 percent.

Could an analyst using support and resistance levels published today have any success predicting intraday trend reversals one week from today?
The answer seems to be yes.

Duration of Predictive Power

Could an analyst using support and resistance levels published today have any success predicting intraday trend reversals one week from today? The answer seems to be yes. Five days after

Table 10

Ability of Support and Resistance Levels to Predict Interruptions of Intraday Exchange Rate Trends after Five Trading Days

	Artificial Levels	Levels Published by Firm 1	2	3	4	5	6
		Bounce Frequency (Number of Hits)					
German mark	53.8	58.0 (4,078)	59.0 (2,464)	57.5 (5,902)	56.9 (2,930)	58.3 (1,004)	55.5 (1,148)
Japanese yen	55.4	62.5 (3,531)	60.5 (2,112)	61.2 (4,392)	59.5 (2,373)	59.7 (703)	64.3 (635)
British pound	54.0	57.5 (3,661)	59.0 (2,400)	56.5 (4,641)		57.4 (622)	
		Months B^P>B^A/Total Months (Marginal Significance)					
German mark		22/27 (0.001)	22/27 (0.001)	21/27 (0.003)	16/26 (0.163)	16/23 (0.047)	16/25 (0.115)
Japanese yen		21/27 (0.003)	20/27 (0.010)	22/27 (0.001)	19/26 (0.014)	12/22 (0.416)	16/24 (0.076)
British pound		19/27 (0.026)	17/27 (0.124)	17/27 (0.124)		13/19 (0.084)	

Source: Author's calculations.

Notes: The table shows the results of using published support and resistance levels to predict intraday trend interruptions five business days after the levels' publication date. The measure of predictive ability is based on the "bounce frequency," or the number of times the exchange rate stopped trending after reaching support or resistance levels compared with the total number of times the rate actually reached such levels.

The table shows the bounce frequency for published and artificial support and resistance levels, the number of hits, the number of months in which the bounce frequency for published levels (B^P) exceeds the bounce frequency for artificial levels (B^A), and the marginal significance of this number of months under the null hypothesis that published support and resistance levels are not informative. "Total months" varies across firms and currencies because occasionally a firm contributed too few support and resistance levels to have any hits at all.

their publication, bounce frequencies for our six firms still exceeded those from arbitrary levels for all firms and currencies, and the differences were statistically significant in nine of the sixteen cases (Table 10). Not surprisingly, the published levels were not quite as useful at predicting intraday trend reversals five days after publication as they were on their actual publication day. On average, five days after publication, rates bounced at published levels 1.7 percentage points less frequently than they did on the actual publication day.[8]

The Power of Agreement

If many analysts agree that a particular level is likely to be important, does this imply that the level is more likely than others to be important? I addressed this question by comparing the predictive power of support and resistance levels provided by more than one firm ("agreed levels") on a given day with the predictive power of support and resistance levels provided by only one firm.

As shown in Table 11, the bounce frequencies associated with agreed levels are quite close to the bounce frequencies associated with levels provided by just one firm. Although the agreed levels tend to have higher bounce frequencies, the differences are generally not statistically significant. The one difference found to be statistically significant implies that agreed levels have *less* predictive power than other levels. Overall, these results suggest that, if agreed levels do provide additional predictive power, the benefit is too small to be of much practical importance.[9]

Table 11

Is There Power in Agreement?

	Narrow Agreement			Broad Agreement		
	German Mark	Japanese Yen	British Pound	German Mark	Japanese Yen	British Pound
Agreed levels	58.0	65.9	63.1	59.1	65.2	61.7
Other levels	59.0	63.7	60.6	58.6	63.5	60.3
Months agreed > other/total months	10/28	15/27	17/28	11/28	14/27	14/28
Marginal significance	0.96	0.35	0.17	0.91	0.50	0.57

Source: Author's calculations.

Notes: The table compares the predictive ability of support and resistance levels on which two or more firms agree ("agreed levels") with the predictive ability of support and resistance levels provided by only one firm. The measure of predictive ability is based on the "bounce frequency," or the number of times the exchange rate stopped trending after reaching support or resistance levels compared with the total number of times the rate actually reached such levels. If agreed levels were better able to predict intraday trend interruptions, the numbers would be positive and statistically significant. Two levels were in "narrow agreement" if they were within 2 points of each other; they were in "broad agreement" if they were within 5 points of each other.

Reliability of Estimated "Strengths"

Three of the firms regularly provided estimates of the "strength" of their published support and resistance levels. For example, levels could be categorized as having strength numbers "1," "2," or "3," with 3 being the strongest. The strength of a particular level can be interpreted as a crude measure of the likelihood that an exchange rate that arrives at the level will actually bounce off it.

Were the estimated strengths of support and resistance levels meaningful? To answer this question, I examined the relative frequency of bounces off support and resistance levels in three strength categories: (1) least strong, (2) somewhat strong, and (3) strongest. Unfortunately, in many months there were few observations in strength category 3, so the only reliable comparison was between categories 1 and 2.

Results for this comparison are shown in Table 12, where the reported differences would be positive and statistically significant if the strength categories were meaningful. In fact, the reported strength levels seem to have no consistent correspondence with the actual frequency with which exchange rates bounced off support and resistance levels. All but two of the differences are negative, and the three that are statistically significant are negative. In short, published estimates of the strength of the levels do not seem to be useful.

Table 12

The Meaning of Reported Strength Ratings

Comparison of Strengths 1 and 2	German Mark	Japanese Yen	British Pound
Firm 1	-2.3**	-4.1***	-6.5
	10/27	8/26	12/28
	(0.04)	(0.01)	(0.17)
Firm 2	-2.7**	-5.5	-3.9*
	10/27	10/25	10/26
	(0.04)	(0.11)	(0.08)
Firm 3	2.4	0.2	-0.1
	14/25	11/24	11/23
	(0.35)	(0.27)	(0.34)

Source: Author's calculations.

Notes: The table evaluates whether support and resistance levels considered somewhat strong by their publishers actually predict intraday trend interruptions better than those considered least strong. The measure of predictive ability is based on the "bounce frequency," or the number of times the exchange rate stopped trending after reaching support or resistance levels compared with the total number of times the rate actually reached such levels. Strength 1 corresponds to the support and resistance levels at which trend interruptions are least likely; strength 2 corresponds to support and resistance levels at which trend interruptions are more—but not most—likely.

For each firm listed on the left side of the table, the first row of numbers represents the difference between the predictive ability of support and resistance levels of the two different strengths. If the reported strength levels were reliable, then the numbers would be positive and significant. The first number in each second row represents the months in which the bounce frequency for strength 2 actually exceeded the bounce frequency for strength 1; the second number in each row (following the slash) represents the number of months in which the comparison was valid; the third row of numbers gives the marginal significance of the second row under the null hypothesis that there is no difference between the two sets of numbers.

* Statistically significant at the 10 percent level.
** Statistically significant at the 5 percent level.
*** Statistically significant at the 1 percent level.

Conclusion

This article has examined the predictive power of support and resistance levels for intraday exchange rates, using technical signals published by six active market participants from January 1996 through March 1998. The statistical tests, which use the bootstrap technique (Efron 1979, 1982), cover support and resistance levels for three currency pairs: dollar-mark, dollar-yen, and dollar-pound.

The results indicate that intraday exchange rate trends were interrupted at published support and resistance levels substantially more often than would have occurred had the levels been arbitrarily chosen. This finding is consistent across all three exchange rates and across all six firms studied. The predictive power of published support and resistance levels varies considerably across firms and across exchange rates. It lasts at least one week. The strength estimates published with the levels are not meaningful. These results are highly statistically significant and are robust to alternative parameterizations.

The predictive power of support and resistance levels has many possible sources, some of which are discussed in Osler (2000). Central bank intervention has been cited as a possible source of the predictive power of other technical trading strategies (Szakmary and Mathur 1997; LeBaron 1999). However, central bank intervention seems unlikely to be an important source of the predictive power of support and resistance levels since there was no reported intervention for the mark and the pound during the sample period. Other possible explanations include clustered order flow, which receives support in Osler (2000), and self-fulfilling prophecies.

The ability of support and resistance levels to predict trend reversals suggests that the intraday currency markets may not be fully efficient. To investigate this possibility, it would be natural to examine whether traders could profit from these predictable bounces on a fairly consistent basis. If it were indeed profitable to trade on these readily available technical signals, there would seem to be some incentive for rational traders to trade the profits away. This would be an appropriate subject for future research. It might also be appropriate to examine the claim of technical analysts that trends typically are sustained once support and resistance levels are "decisively" crossed.

Endnotes

1. Support and resistance levels are related to but not identical to trading ranges. A trading range has just one support level and one resistance level. The firms examined here usually provided multiple support levels and multiple resistance levels each day.

2. These results are available from the author upon request.

3. That is, all artificial support and resistance levels for the mark and the pound had the form x.xxxx00, while all artificial support and resistance levels for the yen had the form xxx.xx00000.

4. For some firms, there were few support and resistance levels in some months, and thus few hits and bounces. These months were excluded from the sample for those firms.

5. If intraday exchange rates followed a random walk, the tendency to bounce would, in the abstract, be about 50 percent. The tendency to bounce in the actual data exceeds this benchmark for two reasons. First, changes in the actual and the simulated data have a fairly strong negative first-order autocorrelation, as noted by Goodhart and Figliuoli (1991). Second, to "bounce," the exchange rate must first reach a level a little above (below) the actual support (resistance) level, and then remain above (below) the actual support (resistance) level for a certain interval. Thus, the exchange rate can continue trending slightly after officially hitting the level yet still be considered as having "bounced."

6. Results from these sensitivity tests are available from the author upon request.

7. The standard deviation of daily exchange rate changes rose by one-third on average between the first and second halves of the sample period. In the first half, these standard deviations were 0.199, 0.216, and 0.260 for the mark, yen, and pound, respectively. In the second half, the corresponding standard deviations were 0.252, 0.362, and 0.277 (all figures E+3).

8. The reader may also be interested to know whether the tendency of support and resistance levels to be selected as round numbers or as local highs/lows has any influence on the levels' predictive power. In Osler (2000), I examine whether round numbers or local minima/maxima (both of which are known to be sources of published support and resistance levels) have predictive power for exchange rate bounces. I find that they do, from which I conclude that at least some of the predictive power in the published levels comes from the firms' tendency to choose these types of numbers. I also show that the size of the typical move following a hit differs substantially between the published levels of some firms and the artificial levels. I conclude from this that round numbers and local minima/maxima do not incorporate as much information about intraday trend reversals as do some published support and resistance levels.

9. It would be desirable here to weight the advising firms by their order flow. However, order information is very closely guarded by the firms in question. Furthermore, some of the firms do not actually take orders.

References

Allen, Helen, and Mark P. Taylor. 1992. "The Use of Technical Analysis in the Foreign Exchange Market." Journal of International Money and Finance 11, no. 3 (June): 304-14.

Andersen, Torben, and Tim Bollerslev. 1998. "Deutsche Mark–Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies." Journal of Finance 53 (February): 219-65.

Arnold, Curtis M. 1993. Timing the Market: How to Profit in Bull and Bear Markets with Technical Analysis. Chicago: Probus Publishing Company.

Brock, W., et al. 1992. "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns." Journal of Finance 48 (December): 1731-64.

Chang, P. H. Kevin, and C. L. Osler. 1998. "Methodical Madness: Technical Analysis and the Irrationality of Exchange Rate Forecasts." Economic Journal 109, no. 458 (October): 636-61.

Cheung, Yin-Wong, and Menzie Chinn. 1999. "Traders, Market Microstructure, and Exchange Rate Dynamics." Unpublished paper, University of California at Santa Cruz, January.

Cheung, Yin-Wong, and Clement Yuk-Pang Wong. 1999. "Foreign Exchange Traders in Hong Kong, Tokyo, and Singapore: A Survey Study." Advances in Pacific Basin Financial Markets 5: 111-34.

Curcio, Richard, et al. 1997. "Do Technical Trading Rules Generate Profits? Conclusions from the Intraday Foreign Exchange Market." Unpublished paper, London School of Economics.

Danielsson, Jón, and Richard Payne. 1999. "Real Trading Patterns and Prices in Spot Foreign Exchange Markets." Unpublished paper, London School of Economics, March.

Dooley, Michael P., and Jeffrey Shafer. 1984. "Analysis of Short-Run Exchange Rate Behavior: March 1973 to November 1981." In David Bigman and Teizo Taya, eds., Floating Exchange Rates and the State of World Trade and Payments, 43-70. Cambridge, Mass.: Ballinger Publishing Company.

Edwards, Robert, and John Magee. 1997. Technical Analysis of Stock Trends. 5th ed. Boston: John Magee.

Efron, B. 1979. "Bootstrap Methods: Another Look at the Jackknife." Annals of Statistics 7: 1-26.

———. 1982. The Jackknife, the Bootstrap, and Other Resampling Plans. Philadelphia: Society for Industrial and Applied Mathematics.

Goodhart, C., and L. Figliuoli. 1991. "Every Minute Counts in Financial Markets." Journal of International Money and Finance 10, no. 1 (March): 23-52.

Goodhart, C., T. Ito, and R. Payne. 1996. "One Day in June 1993: A Study of the Working of the Reuters 2000-2 Electronic Foreign Exchange Trading System." In Jeffrey Frankel, Gianpaolo Galli, and Alberto Giovannini, eds., The Microstructure of Foreign Exchange Markets, 107-79. National Bureau of Economic Research Conference Report Series. Chicago: University of Chicago Press.

Hardy, C. Colburn. 1978. The Investor's Guide to Technical Analysis. New York: McGraw-Hill.

Kaufman, P. 1978. Commodity Trading Systems and Methods. New York: Ronald Press.

LeBaron, Blake. 1999. "Technical Trading Rule Profitability of Foreign Exchange Intervention." Journal of International Economics 49 (October): 125-43.

Levich, R., and L. Thomas. 1993. "The Significance of Technical Trading Rule Profits in the Foreign Exchange Market: A Bootstrap Approach." Journal of International Money and Finance 12 (October): 451-74.

Lo, Andrew W., et al. 2000. "Foundations of Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation." NBER Working Paper no. 7613, March.

Lui, Yu-Hon, and David Mole. 1998. "The Use of Fundamental and Technical Analyses by Foreign Exchange Dealers: Hong Kong Evidence." Journal of International Money and Finance 17, no. 3 (June): 535-45.

Menkhoff, Lukas, and Manfred Schlumberger. 1995. "Persistent Profitability of Technical Analysis on Foreign Exchange Markets?" Banca Nazionale del Lavoro Quarterly Review 193 (June): 189-215.

References (Continued)

Murphy, John J. 1986. Technical Analysis of the Futures Market: A Comprehensive Guide to Trading Methods and Applications. New York: Prentice Hall.

Osler, C. L. 2000. "Are Currency Markets Efficient? Predictable Trend Reversals in Intraday Exchange Rates." Unpublished paper, Federal Reserve Bank of New York, February.

Pring, M. 1991. Technical Analysis Explained: The Successful Investor's Guide to Spotting Investment Trends and Turning Points. 3rd ed. New York: McGraw-Hill.

Shabacker, R. W. 1930. Stock Market Theory and Practice. New York: B. C. Forbes Publishing Company.

Sklarew, Arthur. 1980. Techniques of a Professional Commodity Chart Analyst. New York: Commodity Research Bureau.

Sweeney, R. J. 1986. "Beating the Foreign Exchange Market." Journal of Finance 41 (March): 163-82.

Szakmary, Andrew, and Ike Mathur. 1997. "Central Bank Intervention and Trading Rule Profits in Foreign Exchange Markets." Journal of International Money and Finance 16, no. 4 (August): 513-35.

THE FEDERAL RESERVE BANK OF NEW YORK
ECONOMIC POLICY REVIEW

Single-copy subscriptions to the *Economic Policy Review* (ISSN 0147-6580) are free. Multiple copies are available for an annual cost of $12 for each additional subscription. Checks should be made payable in U.S. dollars to the Federal Reserve Bank of New York and sent to the Public Information Department, 33 Liberty Street, New York, NY 10045-0001 (212-720-6134). Single and multiple copies for U.S. subscribers are sent via third- and fourth-class mail. *Economic Policy Review* subscribers also receive the Bank's Annual Report.

Library of Congress Card Number: 77-646559

All Federal Reserve Bank of New York research publications—the *Economic Policy Review*, *Current Issues in Economics and Finance*, *Staff Reports*, *Research Update*, and the annual publications catalogue—are available at the Research Group's web site.

Visitors to the site can also subscribe to a free **Electronic Alert Service**. Subscribers to **Electronic Alert** receive timely e-mail notifications announcing new publications posted on the site.

www.ny.frb.org/rmaghome