

Transparency, Expectations, and Forecasts

ANDREW BAUER, ROBERT A. EISENBEIS, DANIEL F. WAGGONER, AND TAO ZHA

Bauer is a senior economic analyst in the macropolicy section, Eisenbeis is executive vice president and director of research, Waggoner is a research economist and assistant policy adviser in the financial section, and Zha is a research economist and policy adviser in the macropolicy section, all in the Atlanta Fed's research department. They thank Jinill Kim, Brian Madigan, John Robertson, and Ellis Tallman for critical comments and Cindy Soo and Eric Wang for research assistance. A similar version of this research is also published with the same title as Federal Reserve Bank of Atlanta Working Paper 2006-3.

Many macroeconomists have argued that a central bank should be transparent about its objectives, its views about the economic outlook, and the reasoning behind its policy changes (see Faust and Leeper 2005). In 1994 the Federal Open Market Committee (FOMC) began to release statements accompanying changes in the federal funds rate target. Since then, the degree of specificity of the statements and the guidance provided on the likely course of future policy have evolved significantly.¹ In a recent paper, Woodford (2005) discusses two kinds of central-bank communications: current policy decisions and the central bank's view of likely future policy. He articulates four categories of information—the central bank's view of current economic conditions, current operating targets, strategies guiding policy decision making, and the outlook for future policy—that a central bank might seek to communicate to the public. Woodford argues that these open communications are “beneficial, not only from the point of view of reducing the uncertainty with which traders and other economic decision makers must contend, but also from that of enhancing the accuracy with which the FOMC is able to achieve the effects on the economy that it desires, by keeping the expectations of market participants more closely synchronized with its own.”

This article investigates whether the public's views about the economy's current path and about future policy have been affected by changes in the Federal Reserve's communications policy as reflected in private-sector forecasts of future economic conditions and policy moves. In particular, has private agents' ability to predict the direction of the economy improved since 1994, when the FOMC began to publicly state its views of the economic outlook? If so, on which dimensions has the ability to forecast improved? The analysis focuses on both the short-term and longer-term economic forecasts of key macroeconomic variables—such as inflation, gross domestic product (GDP) growth, and unemployment—and of policy variables such as short-term interest rates. Private agents' current-year and next-year forecasts are used as proxies for the public's short-term and longer-term expectations, and empirical

evidence is presented regarding whether such forecasts have performed better in predicting future economic and policy conditions since 1994.

The private-agent forecasts used in this article are those of individual participants as well as the consensus (average) forecasts contained in the monthly *Blue Chip Economic Indicators* surveys from 1986 to 2004, which include both the pre-FOMC-statement subperiod (1986:01–1993:12) and the post-FOMC-statement subperiod (1994:01–2004:12). We employ the econometric methodology of Eisenbeis, Waggoner, and Zha (2002), which permits us to evaluate the accuracy of forecasts both in cross section and across time and to examine the errors in forecasting key economic variables on both a univariate and a multivariate basis. The latter is important because agents are not simply forecasting one economic variable but rather a set of variables that presumably are interrelated and jointly capture important dimensions of economic performance. Good forecasts on one dimension but poor overall performance may provide some indication of the internal consistency of the forecaster's approach.

This cross-sectional data set enables us to decompose forecast accuracy into two components: the common error that affects all individual participants and the idiosyncratic error that reflects discrepant views across individuals about future economic and policy conditions. According to Woodford (2005), one should expect the idiosyncratic error to become smaller as FOMC open communications become more transparent. But the common error may not change much because it is likely to be affected by factors other than changes in policy transparency, such as unforeseen business cycles.

To preview the main result, we find that since 1994 the idiosyncratic errors for key macroeconomic variables have steadily declined and the expectations of market participants are more closely synchronized to one another. We find no evidence, however, that the common error has become smaller since 1994, especially for the longer-term forecasts.

The Methodology

Let μ_t be an $n \times 1$ vector of economic variables at time t , let y_t be the realized value of these economic variables, and let y_t^i be the i th individual's forecast value of the variables. Assume that y_t is normally distributed with mean μ_t and an economywide (common) covariance matrix Ω_t^R and that y_t^i is normally distributed with mean μ_t and a forecastwise covariance matrix Ω_t^F . (The superscripts R and F stand for "realized" and "forecast," respectively.) The covariance matrix Ω_t^R reflects the aggregate shocks that affect the realized value of μ_t ; the covariance matrix Ω_t^F captures the discrepancy in forecasts across individual participants. The assumption that the mean forecast among individual participants is μ_t is reasonable because previous work has suggested that the Blue Chip Consensus forecast, serving as a proxy for the mean forecast, is close to being an unbiased estimate of μ_t (Bauer et al. 2003). We denote the forecast error for the i th forecaster by $x_t^i = y_t^i - y_t$. Therefore, the individual forecast error x_t^i has mean zero and a variance matrix

$$\Omega_t = \Omega_t^R + \Omega_t^F,$$

which indicates that x_t^i is subject to both idiosyncratic and common shocks.² The standard statistical theory implies that

$$\chi_t^i \equiv x_t^{i'} \Omega_t^{-1} x_t^i \sim \chi^2(n),$$

where $\chi^2(n)$ denotes the χ^2 distribution with n degrees of freedom and χ_t^i is a square error weighted by Ω_t . The above expression shows that the weighted square error

χ_t^i follows the χ^2 distribution with n degrees of freedom. To measure the forecast accuracy for each individual participant, we compute a score value (p value) associated with this χ^2 distribution and call it an “accuracy score.” The score for individual forecaster i at forecast time t is a function of χ_t^i and n :

$$p(\chi_t^i, n) = 1 - \chi_{\text{cdf}}^2(\chi_t^i, n),$$

where $\chi_{\text{cdf}}^2(\chi_t^i, n)$ is the probability that a random observation from the χ^2 distribution with n degrees of freedom falls in the interval $[0, \chi_t^i]$.³

As Eisenbeis, Waggoner, and Zha (2002) point out, the summary measure $p(\chi_t^i, n)$ is a probability that is invariant to the underlying scales-of-error variances. One possible interpretation is that the i th participant’s forecast is closer to the realized value than that of 100 $p(\chi_t^i, n)$ percent of all possible forecasters. Moreover, the score $p(\chi_t^i, n)$ can be compared across forecasters, within a forecast period, and across periods.

Bauer et al. (2003) show how to estimate the covariance matrices Ω_t^R and Ω_t^F . The matrix Ω_t^R can be estimated as the sample covariance matrix of the Blue Chip Consensus forecast errors across time under the assumption that Ω_t^R is the same across years for each month but varies across months within a year. Thus, the variances on the diagonal of Ω_t^R become smaller as t approaches the end of the year because more information becomes available to forecast economic conditions for the current year. The covariance matrix Ω_t^F can be estimated as the sample covariance matrix of forecast errors across individual forecasters; this covariance varies both across months and across years.⁴ The estimate of Ω_t , denoted by $\hat{\Omega}_t$, is the sum of the estimates of Ω_t^R and Ω_t^F . Given this estimate, the weighted-square error can be calculated as

$$\hat{\chi}_t^i = x_t^{i'} \hat{\Omega}_t^{-1} x_t^i.$$


At each time t , the average accuracy score is

$$\hat{p}_t(n) = \frac{1}{N_t} \sum_{i=1}^{N_t} p(\hat{\chi}_t^i, n),$$

where N_t is the number of individual forecasters at time t . One can also calculate the cross-sectional distribution of accuracy scores; the process is described in detail in the sidebar on page 6.

-
1. Kohn and Sack (2003) characterize several distinct periods of increasing transparency in FOMC statements: statements on changes in the discount rate (1989–93), statements on changes in the federal funds rate (1994–98), statements including policy tilt (1998–99), and statements including assessment of the balance of risks (2000–04). In May 2003 a further refinement was added to separately state the committee’s views on the risks to inflation and growth. And, finally, in August 2003 the committee provided explicit guidance on the likelihood that policy would remain accommodative.
 2. In future research, we intend to relax the assumptions that the Blue Chip Consensus forecast is equal to μ_t and idiosyncratic shocks are independent of common shocks.
 3. If the assumptions used are valid, the distribution of accuracy scores from 1986 to 2004 should be uniform. We have verified that such a distribution is more or less uniform, taking into account small-sample uncertainty.
 4. Other estimates can also be constructed using model-based methods.

Figure 1
Blue Chip Average of Individual Scores for the Current Year




Vintage Data and Forecast Errors

The monthly *Blue Chip Economic Indicators* report the forecasts of key macroeconomic variables for the current and next years. We study the annual average forecasts of five key variables: the three-month Treasury bill (T-bill) rate, the consumer price index (CPI) inflation rate, real gross national product (GNP) for 1986 to 1995 or real gross domestic product (GDP) from 1996 to 2004, the unemployment rate, and the long-term bond yield (the corporate bond yield from 1986 to 1995 or the ten-year Treasury note yield from 1996 to 2004). The three-month T-bill rate, the CPI inflation rate, the unemployment rate, and the long-term bond yield are monthly variables while real GNP/GDP is a quarterly variable. This frequency difference is important to note when evaluating forecasts. (See Appendix 1 for a description of and sources for these data.)

More information becomes available about the actual current-year data as the end of the year approaches, and therefore the forecast errors for both the current and next years get smaller. For example, the forecasters participating in the December Blue Chip survey will have monthly data on the three-month T-bill rate and the long-term bond yield through November, data on the unemployment rate through October or November, and data on the CPI inflation rate through October. However, since GNP/GDP data are released quarterly, forecasters will have information regarding GNP/GDP only through the third quarter of the year. The weighted-square error $\hat{\chi}_i$ is designed to avoid the influence of different amounts of available data so that the errors are comparable across time.

To gauge forecast errors, the realized values of each variable at a given time must be used. The values of some variables are revised over time by the agencies responsible for reporting those variables. In particular, real GNP/GDP is reported quarterly and revised twice. Every year additional benchmark revisions may be made in July to past GDP data. Hence, the information reported is actually the continuously changing estimates of many key economic variables' final values. Finally, sometimes the definition of GDP is changed and the series is completely revised.

Figure 2
Blue Chip Average of Individual Scores for the Next Year



Such revisions raise the question, What vintage data should one use to evaluate forecast errors? From a macropolicy perspective, one could argue that the focus should be on the “best” estimate of the final value of the variable of interest. Often, however, that value is not known for several years, and sometimes the difference between even a preliminary estimate and its nearest neighbor estimates can be very large. For example, the advanced estimate for real GDP for the first quarter of 2005 was 3.1 percent. This number was revised upward by the Bureau of Economic Analysis (BEA) to 3.4 percent and finally to 3.8 percent as more data on the performance of the economy became available. Policymakers might have inferred that the economy was growing below trend according to the first number but above trend based on the final estimate. Such differences could have significantly different implications for policy. For this reason, we would argue that the focus should be on forecast methods that best approximate the final number rather than the initial estimate. Also, a priori knowledge of the expected performance of a model or forecasting method can help policymakers decide how to weigh the evidence when significant differences exist between the initial releases of data and forecasts.

For the purposes of this study, for the current-year forecasts, we use vintage data available at the end of January following the current year; for the next-year forecasts, we use data available at the end of January following the next year. This study uses vintage data so that its results will be comparable with those of previous studies. It also provides a comparison between the average Blue Chip Consensus score using vintage and final data, using January 2005 for the final data.

Accuracy Scores

This section looks at the distribution of scores at each month and examines whether the distribution has changed over time, especially from the prestatement subperiod to the poststatement period. The technical details of how to characterize the cross-sectional distribution of scores are provided in the sidebar on page 6.

The first panel of Figure 1 shows the time-series paths of average scores and standard deviations of scores for the current year. The first panel of Figure 2 shows

Characterizing the Distribution of Accuracy Scores

The distribution of accuracy scores can be summarized by the first four moments. The method for calculating the mean or average score $\hat{p}_t(n)$ is shown in the text. The other three moments—standard deviation, skewness, and kurtosis—can be calculated as follows:

$$\hat{\sigma}_t(n) = \left[\frac{1}{N_t} \sum_{i=1}^{N_t} \left(p(\hat{\chi}_t^i, n) - \hat{p}_t(n) \right)^2 \right]^{\frac{1}{2}},$$

$$\hat{s}_t(n) = \frac{\frac{1}{N_t} \sum_{i=1}^{N_t} \left(p(\hat{\chi}_t^i, n) - \hat{p}_t(n) \right)^3}{\hat{\sigma}_t(n)^3}, \text{ and}$$

$$\hat{u}_t(n) = \frac{\frac{1}{N_t} \sum_{i=1}^{N_t} \left(p(\hat{\chi}_t^i, n) - \hat{p}_t(n) \right)^4}{\hat{\sigma}_t(n)^4},$$


where σ stands for the standard deviation, s the skewness, and u the kurtosis.

similar paths for the next year. The measure of standard deviation is often used to approximate the volatility of the public's expectations or forecasts at each point in time. As the first panel of Figure 1 shows, both the average score and the standard deviation of scores fluctuate over time. No noticeable differences exist in the degree of fluctuation before and after 1994, nor are there differences for any subperiods after 1994. No trend appears in which the average score has increased or the standard deviation of scores has decreased since 1994. The figures clearly display periods when forecasters made big errors, such as missing the onset of the recessions in 1990 and 2001. In addition, while the average scores increased in 2004, so did the standard deviations of the scores. Similarly, the average scores dropped significantly in 1995 primarily because the definition of the GDP series changed. In January 1996 the BEA changed the measurement of GDP to a chain-weighted system, but the forecasts made before January 1996 might be based on the non-chain-weighted series. Interestingly, this change seems to have had relatively less effect on the longer-term forecast errors (the second panel of Figure 2).

The average score for the next year (Figure 2) shows no improvement since 1994 and in fact appears to have drifted lower since 1996. The standard deviation of scores since 2001 has drifted steadily upward. The pattern of the drift in the standard deviation is similar to that just prior to and coming out of the 1990–91 recession. As discussed further in the next section, these lower scores after 1996 are most likely associated with the nature of the business cycle and a surge of unexpected productivity growth in the late 1990s.


The second panels of Figures 1 and 2 display the skewness and kurtosis of accuracy scores. Skewness measures the asymmetry of the score distribution. The more negative this measure is, the more scores spread out toward 0 percent. Conversely, the more positive this measure is, the more scores spread out toward 100 percent. Kurtosis measures the likelihood that the score distribution has extreme outliers that may affect the average score. The bigger the value of this measure is, the more likely the presence of outliers in the score distribution is. For the current-year forecasts, the skewness and kurtosis have remained stable except for a few periods. The 1995 spike is the result of the redefinition of GDP, and the small spikes around 2001 are associated with the recent recession. For the next-year forecasts, again, no clear pattern or trend is apparent in which skewness and kurtosis have changed since 1994. Two spikes in skewness and kurtosis correspond to the Asian financial crisis and the recent recession.

Figure 3
**Blue Chip Consensus Scores and the Averages of the
 Five Top and Bottom Forecaster Scores**



Further information about the distributional changes of accuracy scores is provided in Figure 3, which displays the time-series paths of accuracy scores of the Blue Chip Consensus forecast and the average of the top and bottom five forecasts for each month. The consensus forecast is of particular interest because its score is on average the highest (see Appendix 2 for details) and because it performs better than any single individual forecaster over the sample. Again, Figure 3 demonstrates that these scores have had no tendency to improve over time since 1994. In fact, the

Figure 4
Cross-Sectional Standard Deviations of Three-Month Treasury Bill Forecasts




scores of consensus forecasts appear to be slightly lower after 1996 than before, especially for the next-year forecast. Moreover, the drop in the consensus scores around the recent recession and again following September 11, 2001, suggests that events and exogenous shocks affected forecast performance much more than FOMC statements did. The drop in the scores toward the end of 1995 is attributable to the redefinition of GDP. The average scores for the five top and the five poorest forecasters suggest that the data have fat tails, with most of the forecasts being clustered at the high end with a few really poor performers on the bottom.

All these findings suggest that the individual participant's forecast performance relative to other participants has not improved between the prestatement and post-statement periods. Although the accuracy score is a powerful summary measure of forecasting performance, it is a nonlinear function of the square forecast errors weighted by the overall covariance matrix Ω_t . Separating Ω_t and forecast errors for further analysis would be informative. In the next section, we examine whether the covariance matrix Ω_t^F has changed over time and study the sources of forecast errors that do not depend on Ω_t^F and Ω_t^R .⁵

Transparency and Sources of Forecast Errors

Kohn and Sack (2003) and Woodford (2005) argue that the contents of FOMC statements have become more transparent since 1994. To evaluate this argument, it is important to determine whether the expectations of market participants as reflected in the forecasts of key economic variables have become more synchronized in the poststatement period than in the prestatement subperiod. If the statement contains useful information, then one might expect an overall improvement in forecast accuracy, ceteris paribus, or at least more agreement among forecasters (that is, a tighter distribution of idiosyncratic errors). A positive answer may provide evidence about the effects of the FOMC statements on the private sector's agreement on the direction of the future economy.

Figure 5
Cross-Sectional Standard Deviations of CPI Forecasts



We also examine the sources of forecast errors by directly decomposing the mean square error (MSE) into the idiosyncratic component that reflects the discrepancy in individual participants from the Blue Chip surveys and the common component that is associated with unanticipated aggregate shocks and affects all participants. The technical details of this decomposition are provided in the sidebar on page 19.

The MSE is the average of square errors across individual forecasters. Arguably, both the idiosyncratic and common errors may show a decreasing trend if the statement contains useful information and forecasters gain better understanding of the economy over time, especially after 1994. To the extent that the common error is affected by exogenous aggregate shocks and the distribution of the shocks is not constant, no clear inference may exist about the size of the common error. However, we hypothesize that the more important impact is likely to be seen for the idiosyncratic component, in that the idiosyncratic errors should be tighter—that is, greater agreement should be evident among the forecasters. The empirical results presented below confirm this hypothesis.

The degree of synchronization among market participants' expectations is measured by the cross-sectional standard deviations of all the variables, which are equal to square roots of the diagonal elements of Ω_t^e . Figures 4–8 report the cross-sectional standard deviation of each of the five macroeconomic variables considered in this study. These charts clearly show that the trend for these variables has been downward, and the standard deviations tend to be smaller after 1994 than before 1994. These findings suggest that individual participants' forecasts have indeed been more synchronized since 1994 in terms of both their overall view of the economy and the interest rate variable most closely tied to policy.

5. The reader may recall that by assumption Ω_t^e does not change from one year to another. We intend to relax this assumption in future research.

Figure 6
Cross-Sectional Standard Deviations of GDP Forecasts


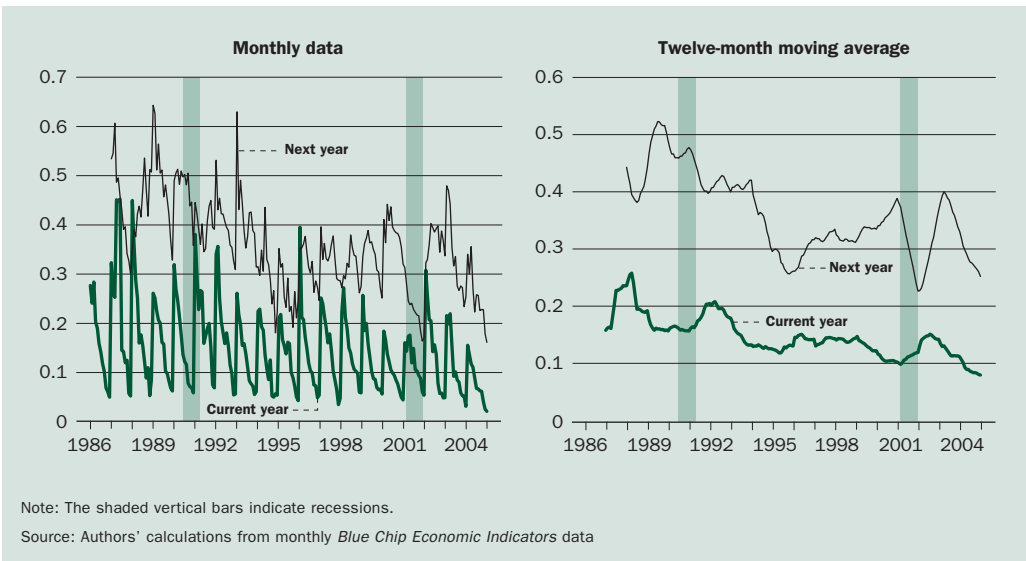



Figure 7
Cross-Sectional Standard Deviations of Unemployment Rate Forecasts



Figures 9–14 show the time-series paths of decompositions for each of the five key variables as well as all the variables jointly. One uniform result seen in the first panel of each figure is that the time path of idiosyncratic errors shows a pattern of steady decline as well as a seasonal pattern for the current-year forecasts. Within the current year, the individual participant's forecast error becomes much smaller as December approaches. The seasonal pattern is much less obvious for the next-year forecasts (the second panel of each figure) partly because the uncertainty about the economy during the coming year is still large even if one tries to forecast as of

Figure 8

Cross-Sectional Standard Deviations of Ten-Year Treasury Note Forecasts

December in the current year. For both the current-year and next-year forecasts, a clear pattern of smaller idiosyncratic errors emerges after 1994. Again, these results are consistent with the hypothesis that individual forecasts have been more synchronized since 1994.


Patterns of common errors are distinctively different from those of idiosyncratic ones, and the difference seems to be associated with business cycles unrelated to the FOMC statements. One can see from Figures 9–14 that the common errors in the current-year forecast are large relative to the idiosyncratic errors whereas the common errors are dominant in the next-year forecasts. But there is no apparent pattern that the common errors are smaller after 1994 than before.

According to the first panel of Figure 9, unusually large common errors for the current-year forecasts of the short-term interest rate occur in 2001. These errors are associated with the unexpected sharp decline of the federal funds rate. The large common errors of longer-term (next-year) forecasts seem to be associated with missing the turning point of the federal funds rate in the early 2000s and failing to predict the unchanged rate in 2002 and 2003 (the second panel of Figure 9).

For CPI inflation, except for two unusually large common errors before 1994, the common errors of the current-year forecasts have similar patterns before and after 1994 (the first panel of Figure 10). The common errors for the next-year forecasts tend to be larger in the period after 1996 than before (the second panel of Figure 10), and no tendency is apparent that these errors have become smaller than before 1994.

Typically, as the end of the year approaches, both idiosyncratic and common errors become smaller for the current-year forecasts. But unusually large common errors of the current-year forecasts of real GNP/GDP develop toward the end of 1995, caused mainly by the definition change of the GDP series. When divided by the diminishing variances of forecast errors, these errors are amplified, accounting for the steep drop of accuracy scores toward the end of 1995 (see the first panel of Figure 3). In the first panel of Figure 11, the errors are not divided by the variances of forecast errors and thus are not as visually dramatic as in Figure 3. The substantial,


Figure 9
Mean Square Errors of Three-Month Treasury Bill Forecasts



persistent common errors of the next-year forecasts in the late 1990s are consistent with the sustained increase in productivity growth being largely unexpected by the public, while the federal funds rate did not change much.

The common errors in forecasting the unemployment rate for the current year appear to be somewhat smaller after 1994 than before, but those errors for the next year have similar patterns before and after 1994 (Figure 12). The large common errors for the next-year forecasts have much to do with business cycles and with the errors in predicting output growth.


Figure 10
Mean Square Errors of CPI Forecasts



No clear patterns exist in which the common forecast errors of the long-term bond yield have become smaller since 1994 (Figure 13). In particular, the errors around the recent recession are relatively large in magnitude. Interestingly, a noticeable drop in the idiosyncratic errors in both the current-year and next-year forecasts occurs after 1987, when Alan Greenspan became chairman and the effects of the stock-market problems dissipated.

Figure 14 summarizes the decomposition of the MSE for the five variables combined. For the current-year forecasts, the seasonal pattern is evident, as explained


Figure 11
Mean Square Errors of GDP Forecasts



early in this article. For the next-year forecasts, the large common errors occurred in the periods around the last two recessions. The persistent and volatile common errors since 1994 are mainly caused by the correlation effect among forecast errors across variables because the forecast errors for individual variables other than GNP/GDP do not share these features. Overall no evidence indicates that the public's forecasts of key macroeconomic variables have improved since 1994, following the FOMC's efforts to increase transparency.

The table (on page 18) reports the average of percentages of the MSE that are attributed to the idiosyncratic component and the common component. Two meth-


Figure 12
Mean Square Errors of Unemployment Rate Forecasts



ods are used to compute the average percent contributions. The first is to calculate the percent contributions of idiosyncratic and common errors for each period and then average them over all the periods. This method helps eliminate outliers of extremely large errors, so the results may not conform to the patterns in the charts. The top panel of the table reports these results.

The second method is to accumulate the forecast errors of both types throughout the entire sample and then calculate the percent contributions of idiosyncratic and common errors (see the bottom panel of the table). This method is likely to be influenced by outliers but will be consistent with the patterns shown in the charts.

Figure 13
Mean Square Errors of Ten-Year Treasury Note Forecasts




Note: The shaded vertical bars indicate recessions.

Source: Authors' calculations from monthly *Blue Chip Economic Indicators* data

In the top panel of the table, the idiosyncratic errors for the current-year forecasts, except for GNP/GDP, contribute much more to the total errors than the common errors do despite the fact that the common errors are much larger at times. But for all the variables jointly, the common errors become more important. This result implies that while predicting a single variable may be relatively easy, predicting a set of economic variables may be more difficult.⁶ For the longer-term (next-year) forecasts, the picture is completely different: The common errors are clearly a driving force for almost all variables (except for CPI), individually and jointly.

Figure 14
Mean Square Errors of All Variables Forecasts



Compared to the results in the top panel of the table, the results in the bottom panel give a more dominant role to the common errors, partly because the common errors are much larger than the idiosyncratic errors in some periods. All in all, the common errors clearly play a dominant role in overall forecast errors.

6. One might also infer that different models are being used and that these models perform better on some variables than others, but in aggregate significant differences exist among the forecasts.

Table

Decomposition of the Mean Square Error

	All variables	3-month T-bill	CPI	GDP	Unempl. rate	10-year T-note
By average percent contribution to error in each period						
Current-year forecasts (1986–2004)						
Idiosyncratic component	44.5	57.0	69.7	43.3	64.0	58.7
Common component	55.5	43.0	30.3	56.7	36.0	41.3
Next-year forecasts (1986–2003)						
Idiosyncratic component	30.0	40.0	52.7	41.0	36.6	48.5
Common component	70.0	60.0	47.3	59.0	63.4	51.5
By percent contribution of total error across sample						
Current-year forecasts (1986–2004)						
Idiosyncratic component	31.9	30.9	40.6	28.0	39.6	32.0
Common component	68.1	69.1	59.4	72.0	60.4	68.0
Next-year forecasts (1986–2003)						
Idiosyncratic component	22.1	15.1	38.6	20.1	24.7	32.1
Common component	77.9	84.9	61.4	79.9	75.3	67.9

This finding suggests that unexpected shocks, which of course are also not anticipated in the FOMC statements, are dominant factors in affecting forecast performance, and improvements in policy transparency would be unlikely to make the forecast errors smaller except on the margins.⁷ Another possibility is that clearer patterns may show up as more observations become available; the FOMC only began in August 2003 to provide explicit guidance on the likely path of future policy and state-contingent economic conditions in the future. Given the data available today, however, we find no empirical evidence of significant improvement in the common forecast errors over the period in which the FOMC attempted to clarify its views of the economy or the likely course for future policy. This finding does not necessarily suggest that the movement toward transparency has been a failure. It may simply indicate that no new information was provided in the statements that had not already been inferred by market participants. Given the unpredictable nature of business cycles, moreover, the common error may be mostly affected by factors other than monetary policy transparency.

Vintage Data versus Final Data

One could argue that whenever forecast errors for a particular period are evaluated, final data available at that time should be used. The reason is obvious: From a policy perspective, being able to accurately predict initially released data that are subsequently revised may lead to policy errors, especially when turning points are imminent or when the revisions may substantially alter one's view of the economy. However, when policy formulation relies heavily upon model forecasts, it is important that those forecasts capture, as well as possible, the true underlying paths for key economic variables. If they do not, then the risk of serious policy errors may be increased. Furthermore, deciding how to choose the vintage data at various points in time is completely arbitrary, and no statistical or economical foundation exists to

Decomposition of the Mean Square Error

Let the estimate of μ_t be

$$\hat{\mu}_t = \frac{1}{N_t} \sum_{i=1}^{N_t} y_t^i.$$

Note that $\hat{\mu}_t$ is also the Blue Chip Consensus forecast. The weighted mean square error at time t can be decomposed as

$$\frac{1}{N_t} \sum_{i=1}^{N_t} x_t^{i'} x_t^i = \frac{1}{N_t} \sum_{i=1}^{N_t} \left[(y_t^i - \hat{\mu}_t) - (y_t - \hat{\mu}_t) \right] \left[(y_t^i - \hat{\mu}_t) - (y_t - \hat{\mu}_t) \right]$$

$$= \frac{1}{N_t} \sum_{i=1}^{N_t} (y_t^i - \hat{\mu}_t)' (y_t^i - \hat{\mu}_t) + \frac{1}{N_t} \sum_{i=1}^{N_t} (y_t - \hat{\mu}_t)' (y_t - \hat{\mu}_t),$$

where the first term on the right-hand side is the MSE attributed to the idiosyncratic component and the second term is the MSE attributed to the common component. The cross term is zero because

$$\frac{1}{N_t} \sum_{i=1}^{N_t} (y_t^i - \hat{\mu}_t)' (y_t - \hat{\mu}_t) = (\hat{\mu}_t - \hat{\mu}_t)' (y_t - \hat{\mu}_t) = 0.$$

guide such decisions. The public know that data such as GDP are often revised and sometimes thoroughly revised. They take such unpredictable outcomes into account and make their forecasts as accurately as possible on average.

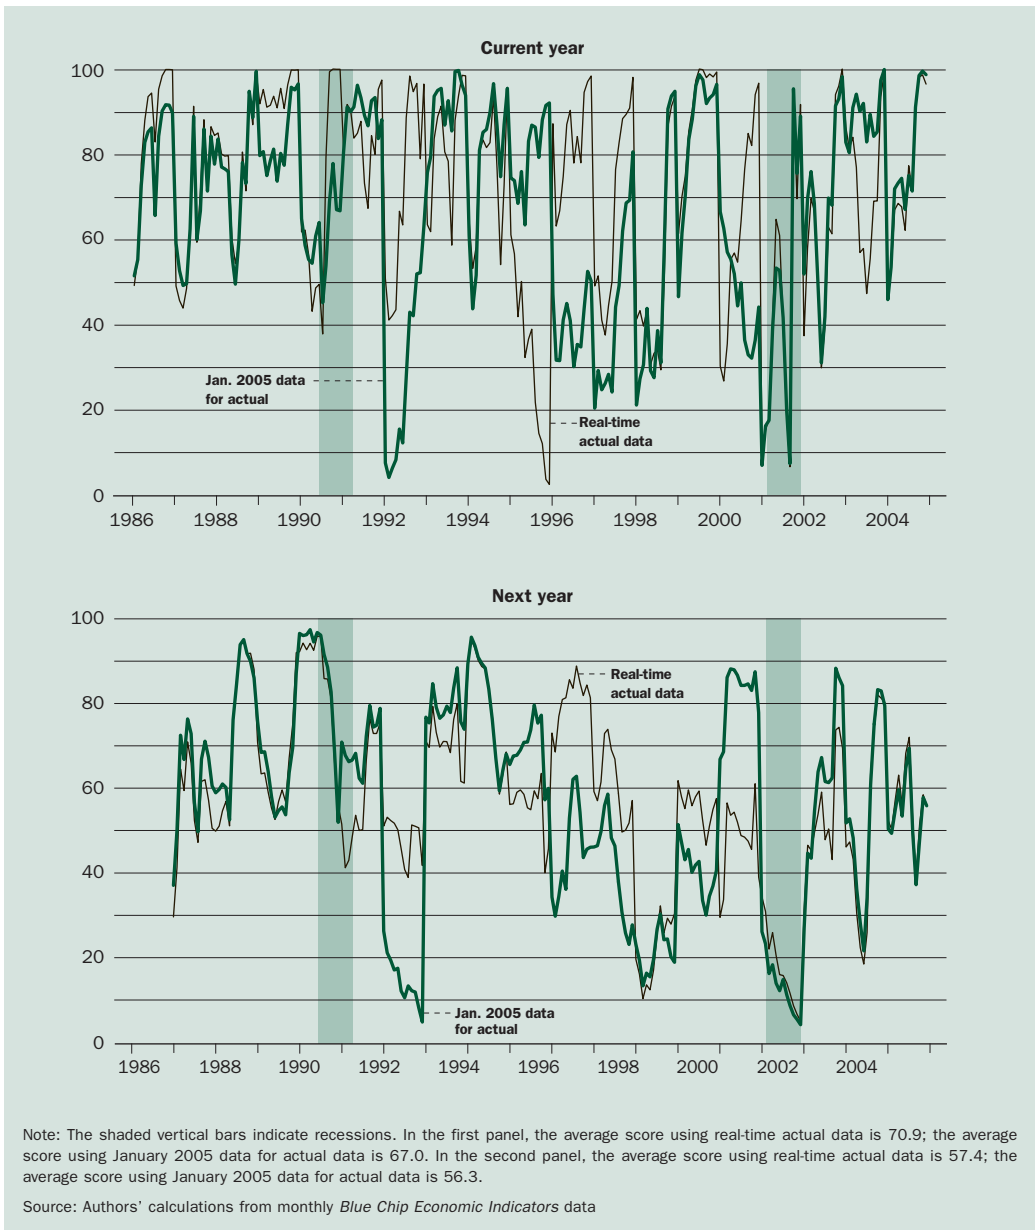
In this section, we use the revised and most current data available at the beginning of 2005 to recompute the forecast errors. Figure 15 displays the Blue Chip Consensus accuracy scores with the vintage data and the final data for both the current-year and next-year forecasts. The average current year score using vintage data is 70.9 while the average current-year score using final data is 67, just 3.9 points lower. For the next-year forecast, the average scores using vintage data and final data are very similar: 57.4 using vintage data and 56.4 using final data. During several periods (1992, 1995–96, and 1998) the next-year forecast scores are lower using final data, but several periods (1994, 1999, and 2002) have higher scores. These results indicate that future data revisions are random enough that they do not introduce a bias that significantly affects forecast scores on average. More important, the findings also suggest that the data revisions do not pose significant risks for policymakers.

One would expect, perhaps, a greater disparity between the two scores given that additional revision errors are unpredictable. However, an important advantage of using the final data is that one can avoid the distorted GDP forecast errors caused by the 1995 data revision. By comparing the first panels of Figures 6 and 11, one can see that the distortion is completely eliminated when the final data are used to measure the forecast accuracy. Still, when the 1995 period is excluded, the difference between the current-year scores using vintage and final data increases from 3.9 to 7.7. Looking more closely at the source of this difference, we find that it can be attributed mostly to the GNP/GDP forecast error.

Figure 16 displays the decompositions of forecast errors for GNP/GDP using the final data as realized values. A comparison of this figure with Figure 11 reveals some notable differences in the breakdown in the composition for both the current-year and next-year forecasts. In the first panel of Figure 11, we see larger overall errors in 1992 and in the 1996–2004 period that are due to increases in the common component


7. This interpretation is consistent with the results of Stock and Watson (2003) and Sims and Zha (2006).

Figure 15
Blue Chip Consensus Scores: Current versus Real-Time Actual Data



of the forecast error. Consequently, a greater proportion of the error each period is due to the common component. The average contribution of the common component to the overall error rises to 73.9 percent from 56.7 percent. In addition, the overall error in 1995 using vintage data (which resulted from the changing to chain-weighted GDP) is no longer present. For the next-year forecasts in the second panel of Figure 11, we again see that the overall error has increased but to a considerably more modest degree. The overall forecast error prior to the 1990–91 recession is less using final data but is greater (on aggregate) for the 1996–2000 period. But once again, this increase

Figure 16
Mean Square Error (Using January 2005 Data as Actual Data) of GDP Forecasts



in overall error is attributable to the common component. The average contribution of the common component rises to 61.8 percent from 59 percent.

Our findings suggest that using final data or vintage data may make little difference when evaluating forecasts. The results show that the average Blue Chip Consensus score is modestly affected for current-year forecasts and almost unchanged for next-year forecasts. In addition, the decrease in score for current-year and next-year forecasts results from an increase in the common component of the forecast error and does not affect the idiosyncratic component. Therefore, the effect of a switch to final data

for evaluating individual forecasts scores should be roughly equal across forecasts. The use of final data eliminates the need for arbitrarily choosing among different vintages.

Conclusion

In 1994 the FOMC began to release statements after each meeting. The amount of policy information released in the statements has increased and changed over time. The findings from Kohn and Sack (2003) and Ehrmann and Fratzscher (2004) suggest that financial markets are sensitive to the information revealed in these statements. While knowing whether the statements have affected markets is important, understanding whether the statements are providing strong signals concerning the FOMC's views about the future path of the economy or economic policy is also important. That is, has the public's ability to forecast future economic and financial conditions improved since 1994? This question is important because one hopes that transparency, if appropriately communicated, enhances market participants' ability to forecast (Woodford 2005).

This article analyzes the forecast errors across a large section of forecasters and for a set of five key macroeconomic variables. The analysis finds evidence that the individuals' forecasts have been more synchronized since 1994, implying the possible effects of the FOMC's transparency. On the other hand, we find little evidence that the common forecast errors, which are the driving force of overall forecast errors, have become smaller since 1994. In fact, common forecast errors have increased and have become more volatile on several dimensions. These common errors seem to be associated with business cycles and other economic shocks. Transparent monetary policy may not necessarily enhance the public's ability to predict business cycles.

On the other hand, it is possible that we do not have a long-enough sample to observe the effects of transparency because the FOMC just began in August 2003 to provide more explicit guidance on the likely path of future policy and its contingency on future economic conditions. We hope that our findings will generate more research on this important topic.

Appendix 1

Data Description

Three-month Treasury bill rate: 1986–2004. Secondary market, monthly average. Source: Board of Governors of the Federal Reserve System.

Consumer price index: 1986–2004. CPI-U (all urban consumers). Source: U.S. Department of Labor, Bureau of Labor Statistics.

Gross national/domestic product: 1986–95, not chained; 1996–2004, chained. Source: U.S. Department of Commerce, Bureau of Economic Analysis.

Unemployment rate: 1986–2004. All workers sixteen years or older. Source: U.S. Department of Labor, Bureau of Labor Statistics.

Corporate bond yield: 1986–95. Aaa, monthly average. Source: Moody's Investors Service Inc.

Ten-year Treasury note yield: 1996–2004. Constant maturity, monthly average. Source: Board of Governors of the Federal Reserve System.

Appendix 2

Scores and Ranks for Individual Forecasters

In this appendix, the following table shows the average scores for all the individual forecasters who have continued to participate in the surveys in recent years. The table also includes the consensus forecast and the Bayesian vector autoregressive (BVAR) model. The BVAR model is often used in the empirical literature as a benchmark for model compari-

son (Robertson and Tallman 1999, 2001), and reporting the real-time forecasting performance of this model is of particular interest to academic researchers. For completeness, we also report other forecasters' scores toward the end of the table. The years in which each forecaster participated in the Blue Chip surveys are also reported in the table.

Table

Overall Performance: Score

Forecaster Name	Overall		Current year		Next year		Participation	
	Avg. score	Std. dev.	Avg. score	Std. dev.	Avg. score	Std. dev.	Current year	Next year
BC—average of top 10	82.24	16.86	86.45	15.99	77.81	16.66	228	216
BC—consensus	64.36	23.49	70.92	24.07	57.43	20.77	228	216
Macroeconomic Advisers, LLC	62.58	27.71	71.57	26.25	53.10	26.06	227	215
Schwab Washington Research Group	62.04	28.26	69.97	27.11	53.64	27.07	197	186
Atlanta BVAR	59.69	31.19	69.21	29.54	49.64	29.75	228	216
U.S. Trust Company	59.25	27.15	64.61	26.25	49.96	26.25	227	131
ClearView Economics	59.23	28.94	66.69	27.72	50.10	27.99	66	54
Banc of America Corporation	59.22	27.10	63.28	27.82	54.87	25.68	204	190
Northern Trust Company	58.75	28.01	63.34	27.27	53.17	27.95	222	183
Wayne Hummer & Company	55.89	27.27	58.05	27.61	53.58	26.78	228	214
Moody's Investors Service	55.04	28.03	65.77	28.63	42.35	21.34	78	66
Perna Associates	54.61	26.31	60.90	28.35	47.82	22.08	167	155

Appendix 2 (continued)

Forecaster Name	Overall		Current year		Next year		Participation	
	Avg. score	Std. dev.	Avg. score	Std. dev.	Avg. score	Std. dev.	Current year	Next year
Merrill Lynch	54.50	27.41	58.36	28.97	50.32	25.02	206	190
Wells Capital Management	53.58	28.71	59.83	28.19	46.83	27.81	161	149
National Association of Home Builders	53.56	26.06	58.77	26.69	47.93	24.21	176	163
Nomura Securities	52.55	28.87	55.77	29.82	48.57	27.41	63	51
National City Bank of Cleveland	52.01	26.08	56.75	26.56	46.93	24.61	224	209
DuPont	51.68	25.60	57.06	28.14	46.00	21.23	228	216
Georgia State University	51.67	27.39	51.72	28.64	51.62	26.07	223	211
Fannie Mae	51.43	28.00	59.67	29.13	41.81	23.35	84	72
DaimlerChrysler AG	51.34	29.13	58.94	29.24	43.35	26.85	226	215
Standard & Poors	51.25	30.43	58.86	30.09	42.78	28.63	120	108
Eggert Economic Enterprises	50.79	25.90	50.12	27.56	51.48	24.09	225	215
Siff, Oakley, Marks Inc.	50.66	28.19	56.56	27.41	44.77	27.78	197	197
Evans, Carrol and Associates	50.43	29.77	58.01	30.35	42.86	27.21	202	202
Bank One	49.82	31.39	56.87	31.75	42.21	29.22	205	190
Bear Stearns & Company Inc.	49.67	29.96	53.11	30.39	43.95	28.59	98	59
BC—average of individual scores	48.13	16.08	51.84	16.22	44.21	15.00	228	216
La Salle National Bank	47.47	29.73	54.13	32.16	40.22	24.97	158	145
Prudential Securities	47.07	31.41	47.40	33.06	46.57	28.88	175	117
Prudential Financial	47.01	26.68	50.54	28.97	43.31	23.55	201	192
Goldman Sachs & Company	46.28	27.19	59.47	25.49	30.49	19.85	79	66
National Association of Realtors	46.10	29.24	51.08	29.08	40.10	28.56	64	53
Conference Board	45.08	29.38	52.22	31.03	37.46	25.46	224	210
Chamber of Commerce, USA	44.97	27.68	48.35	28.34	41.20	26.50	214	192
General Motors Corporation	44.30	28.03	46.05	29.42	42.42	26.40	162	150
Econoclast	43.29	27.13	42.32	30.94	44.32	22.44	227	215
Eaton Corporation	43.04	28.51	40.92	30.07	45.37	26.62	127	115
Turning Points (Micrometrics)	43.04	27.86	41.15	29.28	45.04	26.19	185	174
Comerica	42.41	25.44	43.88	29.34	40.84	20.41	178	166
UCLA Business Forecast	42.12	30.19	45.32	32.23	38.75	27.55	227	215
Motorola Inc.	42.02	28.76	50.83	31.73	31.91	20.92	102	89
JPMorgan Chase	40.92	27.12	47.57	29.12	33.40	22.55	104	92
Kellner Economic Advisers	40.79	23.00	41.86	24.65	39.55	21.02	91	79
Genetski.com	40.46	32.50	50.61	32.88	29.53	28.37	154	143
Wachovia Securities	40.39	27.19	44.60	31.05	35.69	21.33	98	88
Federal Express Corporation	39.92	26.15	41.80	28.76	37.63	22.60	65	53
DRH-WEFA	39.02	27.09	48.32	28.48	27.99	20.65	77	65
Morgan Stanley & Company	35.95	29.39	38.27	31.92	32.30	24.75	85	54
Inforum—University of Maryland	35.72	26.46	33.15	27.15	38.46	25.47	222	208
Deutsche Banc Alex Brown	30.71	28.22	31.86	26.76	29.08	30.33	91	64
Naroff Economic Advisors	29.96	28.67	33.36	32.65	25.86	22.59	70	58
Ford Motor Company	25.80	25.12	27.32	26.09	23.69	23.73	103	74
BC—average of bottom 10	7.50	6.35	6.12	6.32	8.96	6.05	228	216

REFERENCES

- Bauer, Andy, Robert A. Eisenbeis, Daniel F. Waggoner, and Tao Zha. 2003. Forecast evaluation with cross-sectional data: The Blue Chip Surveys. Federal Reserve Bank of Atlanta *Economic Review* 88, no. 2:17–31.
- Ehrmann, Michael, and Marcel Fratzscher. 2004. Central bank communication: Different strategies, same effectiveness? European Central Bank, unpublished paper.
- Eisenbeis, Robert A., Daniel F. Waggoner, and Tao Zha. 2002. Evaluating *Wall Street Journal* survey forecasters: A multivariate approach. *Business Economics* 37, no. 3:11–21.
- Faust, Jon, and Eric M. Leeper. 2005. Forecasts and inflation reports: An evaluation. Paper presented at the Sveriges Riksbank conference “Inflation Targeting: Implementation, Communication and Effectiveness,” Stockholm, June 10–12.
- Kohn, Donald L., and Brian P. Sack. 2003. Central bank talk: Does it matter and why? Board of Governors of the Federal Reserve System Finance and Economics Discussion Series No. 2003-55, November.
- Robertson, John C., and Ellis W. Tallman. 1999. Vector autoregressions: Forecasting and reality. Federal Reserve Bank of Atlanta *Economic Review* 84, no. 1:4–18.
- . 2001. Improving federal-funds rate forecasts in VAR models used for policy analysis. *Journal of Business and Economic Statistics* 19, no. 3:324–30.
- Sims, Christopher A., and Tao Zha. 2006. Were there regime switches in U.S. monetary policy? *American Economic Review* 96, no. 1:54–81.
- Stock, James H., and Mark W. Watson. 2003. Has the business cycle changed? Evidence and explanations. In *Monetary policy and uncertainty: Adapting to a changing economy*. Federal Reserve Bank of Kansas City.
- Woodford, Michael. 2005. Central-bank communication and policy effectiveness. In *The Greenspan era: Lessons for the future*. Federal Reserve Bank of Kansas City.

Merchant Acquirers and Payment Card Processors: A Look inside the Black Box

RAMON P. DEGENNARO

The author is the SunTrust Professor of Finance at the University of Tennessee and a visiting scholar at the Federal Reserve Bank of Atlanta. He thanks Jerry Dwyer, Dick Fraher, Scott Frame, Will Roberds, and Lynn Woosley for useful comments and discussions. He is grateful to Timothy Miller and Mario Beltran of NOVA Information Systems for explaining important institutional details and to Lee Cohen and Victoria L. Messman for research assistance.

Like most consumers, you probably take your credit and debit card transactions for granted. You and others like you carry millions of cards and use them billions of times annually. But unless a transaction goes awry, you rarely think about how your cards work. In fact, a great deal happens after you produce your card to pay for a purchase and before the merchant receives funds and you receive your bill.

What happens during the few seconds between the time you swipe your card and the terminal flashes a result? How does that swipe translate into a line on your bill from the institution that issued the card? When making a purchase using a card online or over the telephone, why are you sometimes asked for the three- or four-digit number printed on the back of the card, the card's expiration date, or arcane information such as your mother's maiden name?

From the merchant's perspective, how is that same card swipe turned into cash to pay for the goods or services provided? Why does a merchant pay a larger fee when it accepts a card in some circumstances than it does in others? And why was the representative from the payment card company so interested in the merchant's personal information before the merchant was even permitted to accept cards?

This article answers such questions. It explains how the card network signs up merchants to accept payment cards and how the sales slips that consumers sign are converted into cash for the merchants. The discussion begins with an explanation of the simplest type of card transaction—one using a private-label card (one that is accepted by only one merchant)—but the focus is primarily on the Visa and MasterCard networks in the United States. The major aspects of payment cards are similar in other countries, although details may differ, especially for cards other than Visa and MasterCard. The key institutions in this transactions process are the merchant acquirer and the payment card processor. The largest of these often perform both functions. Together, merchant acquirers and processors serve as the communications and transactions link between the merchants and the card issuers.

Merchant acquirers and card processors are important for several reasons. First, every card issuer deals with at least one payment processor, and every merchant that accepts cards has a relationship with a merchant acquirer. Without them, the payment system as we know it would not exist. According to Gerdes et al. (2005), U.S. consumers used credit cards for 19 billion transactions and debit cards for another 15.6 billion in 2003. These figures represent a dollar volume of \$1.7 trillion for credit cards and \$600 billion for debit cards. In terms of dollar value, annual growth for credit cards between 2000 and 2003 was 9.9 percent, and for debit cards, 21.9 percent. According to the *Nilson Report* (2005a), in 2004 consumers in the United States

held 795.5 million MasterCard and Visa cards (about three cards for every man, woman, and child in the country).

Second, the industry generates revenues through merchant fees, which merchants must recover either through higher prices or more sales, and the dollar amount

is substantial. Lucas (2004), for example, reports that debit and credit card fees are the fourth-largest expense for gas stations and convenience stores after labor, rent, and utility costs.

Third, the merchant acquiring and processing industry employs many workers. Jeff Johnson, vice president of search and recruitment with CSH Consulting, estimates that the industry employs about 50,000 people.¹ Despite the size of the industry, few people understand the function of merchant acquirers and processors, and almost no academic research on this topic exists.²

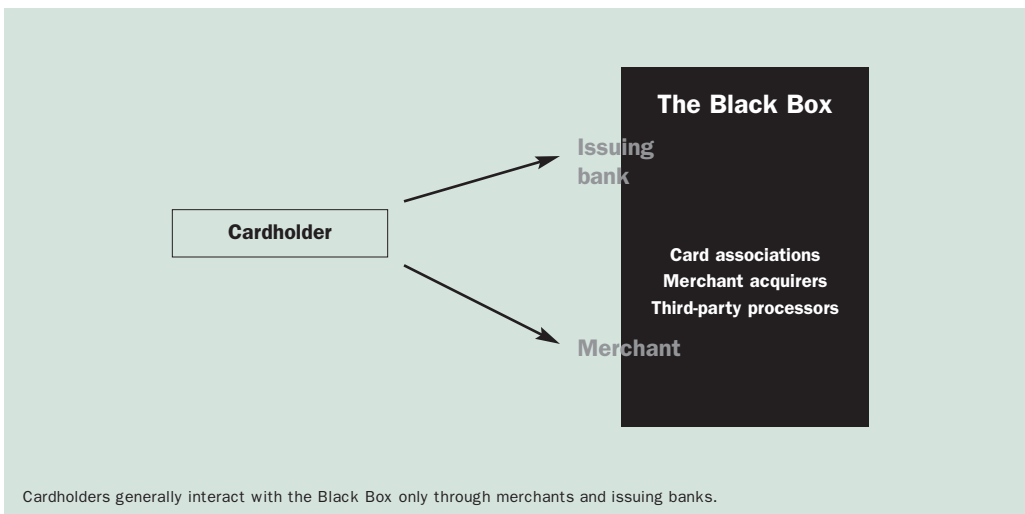
The next section describes how regulations and card association rules set the boundaries of the Black Box. A description of a private-label transaction follows. This is the simplest type of card transaction because the card is accepted by only one merchant. The article then identifies some major types of institutions in the payment card industry and traces the transactions process. The following sections describe how chargebacks and fraud affect a merchant acquirer and identify cross-sectional risk differences among card transactions from the perspective of the merchant acquirer.

The Boundaries of the Black Box

Figure 1 presents a schematic of a credit and debit transaction, in which the cardholder is typically aware only of the issuing bank and the merchant. The cardholder deals with the issuing bank and with the merchant under the protection of Regulations Z and E. The issuing bank and the merchant are liminal figures that deal with the cardholder in the realm of these regulations and with the Black Box through the associations and the merchant acquirers. Unbeknownst to the cardholder, card-based transactions actually travel through the Black Box—a highly evolved group of intermediaries that sign up merchants to accept cards, handle card transactions, manage the dispute-resolution process, and, along with regulatory agencies, set rules that govern card transactions.

Things can and do go wrong with card purchases and billings. Sometimes the culprit is poor quality or bad service. Sometimes the merchant fails to deliver the product. Cardholders and merchants may dispute a refund, and fraud by both cardholders and merchants is a constant challenge. Although these matters can cause serious headaches for cardholders and issuing banks, in most cases the financial impact is relatively minor from the cardholder's perspective. This situation exists because Regulation Z and card association rules limit an innocent cardholder's liability to at most \$50 in

Figure 1
The Boundaries of the Black Box



almost all cases involving credit card fraud, and Regulation E and association rules provide essentially the same protection for debit card users.³ Regulations Z and E thus shift liability for fraud from the (innocent) cardholder to other parties. By means of contracts, the parties within the Black Box and the issuing banks assume and allocate this liability.

In practice, then, Regulations Z and E ensure that most of the losses that result from card-based transactions are allocated among the entities within the boundaries of the Black Box. Aside from initiating a transaction with a merchant at the point of sale, the only time a cardholder interacts with the Box itself is during a dispute. Even then, if an attempt at resolution between the cardholder and the merchant fails, the cardholder typically turns to the issuing bank for relief. For their part, issuing banks usually interact with the Black Box only through the card associations.

Private-Label Cards

This section describes a simplified example of a transaction using a private-label card—a card accepted only by the merchant that issued it. Examples include department stores such as Macy’s and Sears. The transaction begins when the consumer presents the card at the point of sale. The sales clerk enters the purchase amount and, depending on the equipment available, either records the card number and obtains a signature or swipes the card. Depending on the specific merchant, the rest of the transaction cycle is handled either in-house or by a third party such as GE Capital. Sears handled its own processing until 2003, when it sold that part of its business to

1. Jeff Johnson, e-mail messages and telephone conversations with author (November and December 2005).
2. An exception is Rochet and Tirole (2002). Their focus differs from this article’s. They develop a theoretical model of optimal interchange fees and the merchants’ decision to accept payment cards.
3. Section 226.13 of Regulation Z addresses credit card “billing errors.” Section 205.11 of Regulation E contains error-resolution procedures for debit cards, and Section 205.6 of Regulation E covers consumers’ obligation for unauthorized transfers.

Table
**The Ten Largest U.S. Merchant Acquirers in 2004,
 Excluding Partnerships and Alliances**

Ranking (transactions)	Ranking (dollar volume)
1. First Data	1. Chase Merchant Services
2. BA Merchant Services	2. BA Merchant Services
3. Chase Merchant Services	3. First Data
4. Paymentech	4. Paymentech
5. Fifth Third Bank	5. Nova Information Systems
6. Global Payments	6. Fifth Third Bank
7. Nova Information Systems	7. Global Payments
8. Wells Fargo	8. Wells Fargo
9. Alliance Data Systems	9. First National Merchant Solutions
10. Heartland Payment Systems	10. Heartland Payment Systems

Merchant acquirers holding at least 1 percent of U.S. market share in 2004 (by dollar volume), including partnerships and alliances
1. First Data (including Chase Merchant Services, Paymentech, Wells Fargo, SunTrust, and PNC)
2. BA Merchant Services
3. Nova Information Systems (including KeyCorp)
4. Fifth Third Bank
5. Global Payments
6. First National Merchant Solutions
7. Heartland Payment Systems
8. TransFirst

Source: *The Nilson Report* (2005b)


Citigroup. In this simplified example, the processor bills the cardholder and remits funds to the merchant.

Private-label transactions are relatively simple because only one merchant and one processing entity are involved. For universal cards such as Visa and MasterCard, the situation is more complex not only because many different merchants could have made the sale but also because many different banks could have issued the card. Specialized institutions have evolved to route transactions to the correct business entities, and others have evolved to manage the relationship between the card networks and the merchants.

Payment Cards: The Industry and Transactions Processing

The industry. The payment card industry comprises many different entities that perform various tasks, and because many of them have formed alliances, the lines between them are often blurred. The card issuer provides the cards to the consumer and, in the case of credit cards, extends credit to the consumer. (See the sidebar on page 32 for information about different types of cards.) The relationship is business-to-consumer. The merchant acquirer signs up merchants to accept payment cards for the network. This relationship is business-to-business. These acquirers also arrange processing services for merchants. Processors handle transaction authorization and route a (usually electronic) transaction from the point of sale to the network (front-end processing). Later, they handle the information and payment flows needed to

Figure 2

Parties Involved in a Card Program: A Four-Party Network

convert the electronic record created at the point of sale into cash for the merchant (back-end processing). Some merchant acquirers perform the processing themselves; others resell the services of a third-party processor. That is, they are merchant acquirers who resell front- and back-end processing services but do not provide those services themselves. Most of the larger merchant acquirers also function as processors, but almost all of the smaller ones are resellers. The table lists the ten largest merchant acquirers by the number of transactions processed and by dollar volume. Because some acquirers have formed partnerships and alliances, the table also reports the eight groups with more than 1 percent of U.S. market share (by dollar volume).

Only a bank may join Visa or MasterCard; as a result, many merchant acquirers and processors form an alliance or partnership with a sponsoring bank. In addition, depending on the needs of the merchant, an acquirer might sell front-end processing from any of several companies and back-end processing from yet another one. These arrangements make the web of relationships messy, complicating the transactions process. The next section clarifies this process.

Transactions processing. Figure 2 illustrates the institutions participating in a transaction involving either of the two major payment card associations, MasterCard and Visa, which are examples of four-party networks.⁴ The network includes the card issuers and the merchant acquirers/processors plus the cardholders and the merchants. The card issuer distributes cards to consumers, bills them, and collects payment from them. The merchant acquirer recruits merchants to accept cards and provides the front-end service of routing the transaction to the network's processing facilities. The processor is responsible for delivering the transaction to the appropriate card issuer so that the customer is billed and the merchant receives funds for the purchase. Acquirers often delegate the actual processing to third-party service providers. The sidebar on page 36 provides a brief explanation of the differences between four-party networks and three-party networks (for

4. The associations, as umbrella organizations, are not counted as a separate group. Neither are service providers because their function is often served by merchant acquirers.

Types of Payment Cards

Consumers today can choose from a wide variety of payment cards, and the universe of cards can be partitioned in several ways. For example, one way to differentiate cards is according to the merchants who accept them. Some retailers issue private-label cards that are accepted only in their stores. Examples include Sears and Macy's. General-purpose cards, by contrast, are accepted by a wide variety of merchants. Visa and MasterCard are the most common examples.

Another way to classify payment cards is by the amount of time consumers have before payment is due. Debit cards enable a direct withdrawal from the user's savings or checking account, and payment is due much sooner than for a credit or charge card. Debit cards can be used in either online or offline mode. When used in online mode, the card is swiped through a terminal equipped to handle a personal identification number (PIN). In this case, the cardholder enters a PIN instead of signing a transaction slip, and funds are deducted from the user's account immediately. In offline mode, the card is swiped through a standard terminal, and no PIN is entered. Instead, the merchant obtains the cardholder's signature. In this case, the customer's

account is debited within two or three days. A debit card user can purchase any amount up to his balance in that account, and some of these cards even come with overdraft protection.

In contrast, credit cards and charge cards allow the purchaser a longer period of time before he must deliver funds to cover the purchase, and the card may or may not have a predetermined spending limit. Charge cards require the cardholder to pay the balance in full each month unless special arrangements have been made while credit cards allow him the option to make only a minimum payment and pay interest on the balance carried from month to month.

Still another way to distinguish payment cards is by the type of issuer. Financial institutions issue bankcards, which may be either charge cards or credit cards. Visa and MasterCard are the most popular examples. Nonfinancial institutions issue non-bankcards. Market participants subdivide these non-bankcards into two subcategories. Nonbank credit cards, such as Discover Card, enable the cardholder to roll over a balance from month to month while some travel and entertainment cards, such as the American Express Rewards Green Card and the Diners Club Charge Card, are charge cards.

example, American Express, Discover Card, and Diners Club), in which the card issuer and merchant acquirer are the same entity.

The transactions process has two major parts. The first is authorization, and the second is clearing and settlement. Authorization is the process of obtaining permission from the bank that issued the card to accept the card for payment. Clearing and settlement is the process of sending transactions through the Visa or MasterCard network so that the merchant can be paid for the sale. Authorization begins when a consumer presents his card to the merchant for a purchase. Usually, this authorization happens at the point of sale, though an increasing number of transactions are being done in "card not present" situations (for example, online). Merchants usually obtain authorization electronically, either by having the consumer swipe the card through a terminal at the point of sale or by entering the card information manually. However, some transactions still rely on voice authorization, which entails the merchant calling an authorization center to obtain permission to accept the card.⁵ The terminal sends the merchant's identification number, the card information, and the transaction amount to the card processor. The processor's system reads the information and sends the authorization request to the specific issuing bank through the card network. The issuing bank conducts a series of checks for fraud and verifies that the

cardholder's available credit line is sufficient to cover the purchase before returning a response, either granting or denying authorization. The merchant acquirer receives the response and relays it to the merchant. Usually, this process takes no more than a few seconds.

After authorization, the second major part of the transactions process—clearing and settlement—begins. When a consumer purchases an item with a payment card, the consumer and the merchant form a contractual obligation. The merchant agrees to deliver the goods or services, and the consumer agrees to pay for them. Settlement is the process by which assets are delivered to discharge that obligation. Clearing comprises the series of transaction activities from the moment the trade or purchase occurs until it is settled. Usually, clearing involves the transfer of information rather than assets. Examples include netting numerous trades to reduce the number of deliveries, meeting reporting requirements, or handling failed trades (say, due to an error in recording). In the payment-card industry, the most common example of clearing is the process of transferring transaction information from the merchant to its bank. Clearing, then, includes activities that facilitate settlement.

In practice, clearing and settlement for payment cards is more complicated because several entities are involved. Recall that payment card networks include four distinct parties (Figure 2). Moreover, each of those parties for a transaction could be one of hundreds or thousands of different acquirers or issuers and one of millions of cardholders or merchants. The process differs somewhat depending on the specific merchant acquirer and the type of network. The following discussion outlines the major steps of a typical clearing and settlement process for payment cards.

Figure 3 illustrates a typical transaction cycle. In the first step, the merchant sends its transactions to its merchant acquirer. The merchant acquirer sends this information to the merchant accounting system (MAS) servicing that particular merchant's account. In some cases, the MAS is a part of the merchant acquirer; in others, it is a different entity. The MAS distributes the transactions to the appropriate network—Visa transactions to the Visa network, MasterCard transactions to the MasterCard network, and so forth.⁶ Next, the MAS deducts the appropriate merchant discount fee (to cover the costs of the merchant acquirer's activities) from the transaction amount and generates instructions to remit the difference to the merchant's bank for deposit into the merchant's account. The MAS sends these instructions to the automated clearinghouse (ACH) network, which is a computer-based system used to process electronic transactions between participating depository institutions.⁷

Specialized institutions have evolved to route card transactions to the correct business entities, and others have evolved to manage the relationship between the card networks and the merchants.

-
5. For authorization of card-not-present transactions, merchants must follow procedures designed to minimize error and fraud. For example, merchant acquirers can require use of the Address Verification Service (AVS). AVS offers varying levels of detail, including the cardholder's ZIP code, street, city, or state. AVS can even verify which bank issued the card; if the buyer can provide that information, then he probably has the card in hand. This verification process helps rule out fraud by someone who has stolen the card number and does not have the card itself.
 6. The process for transactions routed to networks other than MasterCard or Visa is somewhat different than the one that follows in the text, particularly regarding the handling of payments.
 7. FedACH is part of the Federal Reserve System; the Electronic Payments Network (EPN) is the most notable example of a private ACH network.

To recover these funds, the MAS sends information about the merchant's transactions to Interchange, which is part of the Visa or MasterCard network. Interchange is the clearing and settlement system that transfers data between the card processor and the issuing bank. Interchange determines the interchange fee and Visa/MasterCard assessments (to cover the cost of the issuing bank's services and the network's costs) and sends the information to the card-issuing bank. In turn, the issuing bank remits the transaction amount, less the interchange fee, to Interchange, which passes it on to the MAS. Finally, the issuing bank bills the cardholder and collects the balance.⁸

Merchant acquirers provide other services to merchants besides the processing described above, including installing card terminal equipment, recording transactions, providing reports, and handling problems with card processing (Chang 2004). Some acquirers also provide related services such as analyzing the purchasing patterns of the merchant's customers.

Chargebacks and Fraud


Chargebacks. A merchant acquirer suffers losses if a merchant is unable to make good on credit transactions disputed by customers, called chargebacks. Chargebacks usually occur when a consumer is dissatisfied with a product or service. Beginning with the later of the date on which a transaction is processed or the delivery of the product or service, cardholders have as much as three months to claim a chargeback—sixty days plus up to another month depending on the purchase date relative to the billing cycle.⁹ The presumption is initially in favor of the customer, and the amount of the chargeback is deducted from the merchant's account pending the result of a review. If the dispute is resolved in the merchant's favor, then the merchant recovers the funds. The merchant acquirer is at risk in the event that the merchant fails between the time of the initial sale and the time his account is debited for the chargeback. In this case, according to the card network's rules, the merchant acquirer is liable and must make restitution to the customer.

Because of this feature, the merchant (and ultimately the merchant acquirer) is at risk of loss for up to several months because the transaction can be reversed. In the language of payments, the transaction is not final. This feature greatly enhances the appeal of credit cards to cardholders, but it also shifts the risk of chargebacks to the merchant acquirer. In essence, the merchant acquirer has insured the issuing bank against an adverse result. The risk of a merchant acquirer's contingent liability is similar to that of a bank's guarantee of a debtor's liabilities or an insurance contract. Merchant acquirers include the cost of this implicit insurance in the price that merchants pay for their services.

Quinn and Roberds (2003) argue that payment-finality rules are essentially loss-allocation rules. The rules determine which party to a transaction absorbs the loss if the transaction is not completed. For example, cash transactions are final when goods or services are exchanged for cash. Absent fraud or a private agreement such as a warranty, neither the buyer nor the seller can cancel the transaction after the exchange. In contrast, because of Visa/MasterCard chargeback provisions, credit card transactions are effectively not final for up to three months after delivery of the good or service. This lack of finality is a key determinant of a merchant acquirer's risk because, until a transaction is final, the merchant acquirer bears the risk that a merchant cannot cover a chargeback.

The industry attempts to quantify this risk through the closely related concept of delayed delivery. Magazine subscriptions are a good example. Subscribers pay for subscriptions in advance, and the term of subscriptions can be as much as a few

Figure 3
The Transactions Process



years. If the magazine ceases publication before the term of the contract, then the subscriber has recourse for undelivered issues according to Visa/MasterCard rules. The delay between the sale and the delivery of the goods or services increases the chances that the merchant will fail and be unable to cover the resulting chargeback. The sidebar on page 38 describes an extreme example.

Fraud. Kahn and Roberds (2005) define fraud risk as the risk that a claim cannot be collected because the identity of the person who incurred the debt cannot be established. They identify three distinct types of fraud. First, existing account fraud is usually traced to stolen account information. For example, a thief who steals a card and orders merchandise commits existing account fraud. The second category is new account fraud, popularly called identity theft. In this case, a thief uses information about a third party to open an account, incurring debts in the name of the victim. Finally, those who commit friendly fraud make legitimate transactions that they later deny having made.

The risk of fraud is especially serious if a merchant takes orders by mail, telephone, or over the Internet. In such card-not-present situations, the Truth in Lending Act frees cardholders from liability—they are not responsible for even the first \$50 (association rules provide essentially the same protection for debit card users). This consumer protection shifts the risk to merchants and, in turn, creates a larger contingent liability for merchant acquirers. One notorious example involves a merchant that defrauded customers by taking orders with no intention to deliver. Had the merchant been a traditional storefront operation, red flags would have been more apparent. First, customers would have been interacting with the merchant face to face,


8. For debit cards, the billing is done automatically. Put differently, the cardholder's account is debited, and the cardholder later receives a statement of transactions rather than a bill.
9. Specific details of chargeback terms are complicated because they are governed by law (for example, the Truth in Lending Act), by regulation (Regulation Z for credit cards and Regulation E for debit cards), and by the rules of the card associations and networks. See Furletti and Smith (2005) for more information. The terms in the text are common in the industry.

Three-Party Networks

The figure illustrates the three-party analog to the four-party diagram in Figure 2. The only major distinction is that, in three-party networks, the card issuer and the merchant acquirer are the same entity; in four-party networks, they are separate. In four-party networks, banks that are members of Visa and MasterCard issue the payment cards and extend credit to consumers for credit cards. Separate entities are responsible

for signing up merchants to accept these cards for payment. In practice, some acquirers are affiliated with or have formed partnerships with card issuers. Most payment cards in three-party networks are nonbank cards, issued by institutions such as American Express, instead of a bank. In almost all cases, the difference between three- and four-party networks is unimportant for cardholders.

Figure
Parties Involved in a Card Program: A Three-Party Network



making it easier to detect suspicious behavior. Second, customers would have been more likely to benefit from the experiences of other customers; they might have met in the store or overheard conversations and complaints. Finally, either the merchant would have had no inventory or business history at that location (fueling suspicion), or he would have had at least some inventory and other collateral after the firm failed. Either way, the merchant acquirer would have been better off. Instead, because this was a card-not-present situation, the fraudulent merchant was able to collect a large amount over a period of several weeks. When consumers were no longer willing to wait for delivery and filed chargebacks, they were entitled to relief. Because the fraudulent merchant could not pay, the merchant acquirer was forced to make restitution.

Cross-Sectional Risk Factors

Clearly, a merchant acquirer must consider the credit standing of the merchants it services. Merchant acquirers do perform credit analysis, but the analysis is different from that of a more familiar bank loan. A merchant acquirer's contingent liability is more similar to an insurance contract than to a bank loan. This description fits in part because the acquirer pays only if another entity cannot, but there are other differences. For example, for a bank loan, the bank delivers funds to a borrower. A merchant acquirer, though, advances no funds. Instead, it indemnifies a third party—the card

issuer (who in turn indemnifies the cardholder)—in the event that a merchant cannot cover a chargeback.

Another major difference between bank loans and a merchant acquirer's contingent liability is the term of the contract. Bank loans can have maturities of several years. In contrast, although consumers can file chargebacks for up to several months after a purchase, the effective term of the contingent liability produced by each transaction is usually measured in a very few days. In addition, merchant acquirers review most accounts at least once a year. Cast in terms of the probability of default times the loss given default, the probability of default is affected in part by the time between account reviews, and the loss given default—again absent delayed delivery—rarely represents more than a few days' worth of total processing volume at any one time.

Taken together, the merchant acquirer's annual review of accounts and the short term of the contingent liability have enormous implications for risk. The annual review makes the risk that merchant acquirers face similar to a short-term bond, whereas a bank loan is (sometimes) more similar to a long-term bond. Investors in short-term bonds need not reinvest in the same company when their bonds mature if, for example, a company's credit quality deteriorates. Long-term investors do not have that option. They can only sell their bonds prior to maturity, likely taking a loss because the credit standing of the bonds has deteriorated. Similarly, if a merchant's credit quality deteriorates, a merchant acquirer need not renew the relationship, whereas a bank probably cannot cancel a loan unless a covenant has been violated.

Because the merchant acquirer is at risk if the merchant cannot cover a chargeback, the acquirer must evaluate the credit quality of merchants seeking to use the acquirer's services and monitor the credit quality of the merchants it currently services. The acquirer considers industry effects, firm-specific effects, and even the nature of individual transactions. In fact, merchant acquirers charge different fees depending on whether or not a merchant has followed certain procedures for a transaction.

Industry effects. Because customers who regret making a purchase have up to three months to act before their credit card purchases are final, businesses that are susceptible to so-called buyer's remorse present higher risk to a merchant acquirer. Consider health clubs, which often sell annual memberships at a discount relative to their monthly fee to encourage customers to commit for a longer period. The problem is that many customers regret their commitment after just a few weeks. Although buyer's remorse alone is not sufficient to win a chargeback dispute, it does give the buyer incentives to try to exploit the process. For example, he might claim that equipment at the club is often broken or that the premises are unsanitary. Because "often" and "unsanitary" are matters of degree, the cardholder has a chance to win the chargeback dispute, putting the acquirer at risk. Merchants that sell items of high and uncertain value—collectibles are an obvious example—are also prone to customer disputes. Customers can be disappointed in artwork, rare coins, or stamps for any of several reasons. Also, fraud is frequently involved in these types of businesses because the goods may not be genuine or their condition might be exaggerated. Mystics, such as fortune tellers, face high chargebacks due to buyer's remorse, and one can easily see how customers of gambling establishments could regret a transaction depending on the outcome of a race or sporting event. For this reason, such businesses usually are not authorized to take credit cards for purchases.

Because the merchant acquirer is at risk if the merchant cannot cover a chargeback, the acquirer must evaluate the credit quality of merchants seeking or using the acquirer's services.

Delayed Delivery in the Extreme

The nature of airline ticket sales and the industry's current financial problems combine to form an extreme example of delayed-delivery risk. Consider a cardholder planning a trip by air. In some cases, the cardholder buys his ticket weeks or even months in advance, and travelers usually pay for their tickets using a credit card. Suppose that the airline fails between the time of purchase and departure. In this case, under credit card association rules, the acquirer must make restitution. How large can the potential losses be?

One merchant acquirer, National City Corporation, reports that as of June 30, 2004, the value of credit card transactions it had acquired for outstanding tickets purchased on United Airlines was \$853 million (National City Corporation 2004a). United Airlines is operating under Chapter 11 protection as of this writing. If United Airlines were unable to honor those tickets, then travelers who purchased their tickets using credit cards would be entitled to refunds under Visa and MasterCard rules, and National City held no significant collateral against this potential liability as of June 30, 2004. The \$853 million worth of unflown tickets, of course, represents the potential liability from exposure to United Airlines alone. National City Corporation (2004a) says that it processed over five times that amount—about \$5 billion worth of delayed-delivery purchases—during the six months ending June 30, 2004. National City Corporation (2004b) reports that as of December 31, 2004, the value of unflown tickets had been reduced to \$547 million.

Of course, the odds are small that National City Corporation would be liable for the full amount of these huge sums. Consider the case of United Airlines. For National City Corporation to be liable for the full amount, three things must happen. First, United Airlines must halt all flights. Second, all ticket holders must file chargebacks within the allotted time limits. Although this is within their rights, many travelers would instead opt to fly on other airlines, which usually honor the stranded travelers' tickets on a standby basis (McCartney 2004).¹ This provision reduces the number of travelers who file chargebacks. Finally, National City Corporation would have to have a recovery rate of zero in liquidation. This outcome is unlikely because, even as a general creditor, the company could probably recover a portion of its losses from the bankrupt carrier. If National City Corporation anticipates problems it can also require a security deposit, a line of credit from a bank, or delay payment to the merchant.

National City Corporation (2004a) puts the problem in perspective. For the first and second quarters of 2003 and 2004, the company processed about \$35 million in chargebacks each quarter, for a total of about \$150 million in the four quarters. Actual losses were about \$1 million each quarter, for a total of about \$4 million. The company had \$5 million worth of chargebacks in the process of resolution as of June 30, 2004. The company believes the chance of a "material loss" because of chargeback rules is "unlikely" (National City Corporation 2004a). Still, losses of this size are not trivial, and "unlikely," of course, does not mean that a material loss is impossible.

1. In November 2005, Congress extended this provision through November 2006. Airlines must honor these tickets but may charge a fee and need only accommodate travelers on a space-available basis.

Instead, customers must get cash advances on their cards and use the cash to make the purchase.

Items that can easily be resold are prone to fraud, so dealers in these products also present higher risks. Consumer electronics and jewelry head the list. Intangible products, particularly downloadable software, tend to attract fraudulent merchants and customers because proof of delivery and the products' performance are difficult to substantiate. Timeshare services have high chargeback rates because customers sometimes place deposits months before developers even begin construction, when

the suitability of the property is difficult to ascertain. Customer dissatisfaction is more common in such cases.¹⁰

Perhaps the best example of an industry effect is the restaurant industry. Most bankers realize that loans to restaurants are very risky. For example, the Cline Group (2003) tracked over 4,000 non-fast-food restaurants in the Dallas area and reported that an average of 23 percent failed during their first year. Yet restaurants are extremely safe customers for merchant acquirers. Why? Consider the nature of a restaurant transaction. The diner finishes the meal, pays using a credit card, and departs. In the vast majority of cases, the consumer is satisfied enough to consider the transaction to be final, and the settling of accounts proceeds normally. Suppose instead that the diner is dissatisfied. Although Visa/MasterCard rules give the diner the right to file a chargeback for several weeks afterward, only in very rare circumstances will the diner pay, leave the premises, and then file a complaint. The diner is more likely to voice his dissatisfaction during the meal, and, almost always, restaurant management accommodates the diner. By the time the consumer uses his credit card, he is satisfied and considers the transaction to be final. The settling of accounts again proceeds normally. Only in very rare circumstances will he still complain after using his credit card. Even then, a complaint does not necessarily imply that the acquirer bears a loss. For the merchant acquirer to incur a loss, the cardholder must win the chargeback dispute (unlikely in such cases), and the merchant must fail between the time of the sale and the chargeback. Otherwise, the merchant itself and not the acquirer is responsible for the chargeback.

Because customers who regret a purchase have up to three months before their credit card purchases are final, businesses that are susceptible to so-called buyer's remorse present higher risk to a merchant acquirer.

Firm-specific risk. Just as insurers and banks evaluate the credit risk of individual companies, so do merchant acquirers. For example, they study standard measures of financial strength, such as financial ratios of individual firms. For unincorporated businesses, financial statements are often unaudited, so acquirers might use business tax returns to supplement the unaudited statements. Especially for small firms, acquirers even proceed beyond the firm level and use information about the owners and managers of companies, especially for unincorporated businesses. Acquirers can use credit scores from the Fair Isaac Corporation, commonly known as FICO scores, at the personal level as well as at the business level. Acquirers also use credit report information and the number of years that a potential customer has been in business to gauge risk. Both traditional lenders and merchant acquirers use information that others have already generated about specific firms—for example, whether or not the merchant has existing banking relationships. Almost surely, an international company will receive greater scrutiny than a domestic one. The processing history of a company that already has a relationship with a merchant acquirer is always important, particularly fraud and chargeback rates.

If the firm's condition is sufficiently weak, a merchant acquirer might require the owner to offer a personal guarantee; such guarantees are common for small business loans. An acquirer might impose conditions similar to restrictive covenants in business loans. For example, the acquirer might impose a processing limit, which corresponds to a commercial bank's lending limits. Like a bank, the acquirer might require

10. For examples of items on restricted lists, see www.internetsecure.com/solutions-faq.htm#2 and www.practicepaysolutions.com/apply/index0007.php.

marginally qualified merchants to provide collateral, usually in the form of a certificate of deposit, cash, or a letter of credit. If the merchant cannot provide collateral, then the acquirer might institute a holdback, or a delayed-payment arrangement. Under such an arrangement, the merchant acquirer withholds payment to the merchant for a predetermined length of time after processing. The duration of the payment delay is usually a function of the delivery delay and, less frequently, the chargeback ratio.

Transaction-related risks. Banerjee (2004) notes that credit cards were originally designed to be physically present at the point of sale. If merchants followed procedures, then nearly all risks except fraud and delayed delivery declined enormously.

This low level of risk is still true for face-to-face transactions. For example, if a merchant swipes a card instead of manually keying the card number, the chance for error drops to near zero. True, the card may have been stolen, but swiping is at least one step toward insuring

These procedures are only partially effective, so merchant acquirers charge higher fees for card-not-present transactions to compensate for the higher risk.

legitimacy: A thief must have stolen the card itself and not just the card number. This consideration goes far toward eliminating theft losses from, say, a dishonest waiter who copies the card number while clearing a diner's tab.

For a growing number of transactions, however, the cardholder and the card are not present. As a result, merchants and merchant acquirers face the challenge of developing new procedures for limiting risk. Mail-order and telephone-order (MOTO) transactions—and, more recently, Internet transactions—have presented special problems for payment card associations. The most popular approach has been for merchants to have access to increasingly arcane bits of information during authorization. Some help to confirm that the purchaser has possession of the card itself and not just the card number. For example, card associations have long encoded a verification number into the magnetic stripe on the back of the card. Visa calls this code the Card Verification Value (CVV or CVV1); MasterCard's term is the Card Validation Code (CVC or CVC1). This code, read during the swipe, confirms that the card is actually present at the point of sale. The problem is that this approach cannot help for Internet or MOTO transactions because the card is not present and a swipe is impossible.

Associations have had to devise other ways to confirm that the purchaser is in physical possession of the card at the time of the sale. The result is CVV2 and CVC2. These three-digit numbers (different from the magnetically coded CVV or CVC numbers) are printed on the right side of the signature area on the back of the card. Because this number is not embossed on the card, it does not appear on a paper sales slip, making it harder to steal. The customer must have physical possession of the card—or the printed number stolen by some other means—for the buyer to have access to it. CVV2 and CVC2 are only partially effective, though. First, the network merely flags the transaction if the buyer cannot provide the number; it does not refuse it. Second, some situations make it easy to defeat. For example, a dishonest waiter can steal a CVV2 or CVC2 number while clearing a dinner tab just as easily as he can steal a card number. Because these procedures are only partially effective, merchant acquirers charge higher fees for card-not-present transactions to compensate for the higher risk.¹¹ Still, CVV2 and CVC2 provide one more layer of protection, and Banerjee (2004) reports that they do help discourage fraud.

Another approach is to prearrange a question and answer or series of questions and answers. Card users might be asked to verify their mother's maiden name, for example. By allowing cardholders to select from a list of questions, merchants and

acquirers make it more difficult for a thief to have the necessary information. The Address Verification Service (AVS) is a good example (see footnote 4). This verification process helps rule out fraud by someone who has stolen the card number and does not have the card itself. These procedures are somewhat effective, but as Banerjee (2004) points out, none of the ways to reduce fraud on the Internet seems to be particularly effective. As evidence, he notes that issuers have not lowered the interchange fees they charge for transactions that follow these procedures. More recent innovations are Visa's Verified by Visa and MasterCard's MasterCard SecureCode. Both of these systems use passwords for Internet purchases to insure that only the cardholder can make such purchases.

These examples illustrate that merchant acquirers can help protect merchants (and therefore themselves) from fraud by setting procedures. After all, most merchants are too small to dedicate resources to designing low-cost, effective fraud-protection procedures, so a merchant acquirer can add value by supplying them. Merchant discount rates provide a means for acquirers to give incentives without mandating a specific procedure for each different merchant.

Merchant acquirers provide these incentives by setting qualification levels for the discount fee that merchants pay; the more hurdles the merchant surmounts for a transaction, the higher the qualification rate and the lower the discount fee. A three-tiered system is common, beginning with the nonqualified rate, which is the lowest acceptable category (with the highest fee); moving to the partially qualified rate; and ending with the qualified rate, which is the highest category. For an example of how these tiers are determined, consider the method of entering the card number. Being hand-keyed without AVS might automatically drop a transaction to the non-qualified rate; adding AVS might move the transaction to the partially qualified rate. Swiping the card could move the transaction into the qualified rate.

Different industries sometimes have different qualification criteria. For example, tipping is common in businesses such as restaurants, and the amount of the tip is usually unknown until after the card is swiped or the card number is entered. Therefore, the amount approved is a lower bound on the total amount to be charged. If the final amount including the tip is sufficiently above that lower bound, then the transaction might drop to the partially qualified rate from the qualified rate.

A merchant acquirer's management of individual sales is not limited to the time when the customer places the order. Merchant acquirers often require merchants to follow specific procedures immediately prior to shipping. For example, just before shipping a back-ordered item, a merchant might be required to contact the buyer to verify the customer's telephone number, mailing and shipping address, or e-mail address. For MOTO or Internet purchases, shippers can insist that products be delivered only to the card's billing address (rather than delivering to a destination that a would-be thief designates). This practice helps reduce fraud because the thief is less likely to attempt fraud in the first place if he knows he may not receive the merchandise.

Finally, card associations have set procedures that force acquirers to cooperate to improve network efficiency. One obvious example is the MATCH list (Member Alert to Control High Risk Merchants), maintained by Visa and MasterCard, which comprises problem companies. If a merchant acquirer denies permission to accept cards to a merchant because of adverse processing behavior and fails to add it to the MATCH list, then the merchant acquirer is liable for losses another provider might suffer from that merchant.

11. For example, see AMS's Web site at www.merchant-accounts.com/retail-merchant-account.html.

Summary

Consumer and merchant acceptance of payment cards has been phenomenal. Hundreds of millions of cardholders make billions of transactions worth trillions of dollars each year. Yet few cardholders understand how payment networks operate. Most treat them as a Black Box.

This article demystifies the transactions process for payment cards, emphasizing the roles of the merchant acquirer and card processor. After outlining the regulations and card association rules that set the boundaries of the Black Box, the article describes a transaction with a private-label card. The discussion then considers the complications introduced by general-purpose cards, such as Visa and MasterCard, and introduces a key participant in the payment card market, the merchant acquirer. The description of the risks borne by merchant acquirers demonstrates that they take losses on these transactions only in rare circumstances—usually when a merchant fails to make good on a chargeback. The article also delineates some of the risk factors associated with specific industries, merchant types, and transactions that influence the price merchants pay for these transactions services. Finally, the article discusses some ways that merchant acquirers manage the risks that they face, especially the risk of fraud.

REFERENCES

- Banerjee, Sankarson. 2004. Credit card security on the Net: Where is it today? *Journal of Financial Transformation* 12 (December): 21–23.
- Chang, Howard H. 2004. Payment card industry primer. *Payment Card Economics Review* 2 (Winter): 29–46.
- Cline Group. 2003. *Restaurant start & growth magazine unit start-up and failure study*. Cline Group for Specialized Publications, September.
- Furletti, Mark, and Stephan Smith. 2005. The laws, regulations, and industry practices that protect consumers who use electronic payment systems: Credit and debit cards. Federal Reserve Bank of Philadelphia Discussion Paper No. 05-01, March.
- Gerdes, Geoffrey R., Jack K. Walton II, May X. Liu, and Darrel W. Parke. 2005. Trends in the use of payment instruments in the United States. *Federal Reserve Bulletin* (Spring): 180–201.
- Kahn, Charles M., and William Roberds. 2005. Credit and identity theft. Federal Reserve Bank of Atlanta Working Paper 2005-19, August.
- Lucas, Peter. 2004. Why gasoline retailers are fuming. *Credit Card Management* (August): 20.
- McCartney, Scott. 2004. Bill to protect flyers from shutdowns has a surprising beneficiary. *Wall Street Journal*, October 26.
- National City Corporation. 2004a. *Form 10-Q: Quarterly report pursuant to section 13 or 15(D) of the Securities Exchange Act of 1934—for the quarterly period ended June 30, 2004*, Commission file number 1-10074. Filed August 6, 2004.
- . 2004b. Annual report. *The Nilson Report*. 2005a. Visa & Mastercard—U.S. 2004. No. 828, February.
- . 2005b. Top U.S. acquirers. No. 831, April.
- Quinn, Stephen F., and William Roberds. 2003. Are on-line currencies virtual banknotes? Federal Reserve Bank of Atlanta *Economic Review* 88, no. 2:1–15.
- Rochet, Jean-Charles, and Jean Tirole. 2002. Cooperation among competitors: Some economics of payment card associations. *RAND Journal of Economics* 33, no. 4:549–70.

International Business Cycles: G7 and OECD Countries

MARCELLE CHAUVET AND CHENGXUAN YU

Chauvet is an associate professor of economics at the University of California, Riverside, and a former research economist at the Atlanta Fed. Yu is a research scientist with the New York State Department of Health.

Monitoring economic activity through the use of composite leading and coincident indicators has been a tradition in the United States for over sixty years, since the seminal book by Arthur Burns and Wesley Mitchell (1946). These indicators are some of the most watched series by the press, businesses, policy-makers, and stock market participants. Progressive globalization has sparked a worldwide interest in using economic indicators to analyze cyclical fluctuations. The development of the European Monetary Union and advances in econometric models that explore potential dynamic differences across business cycle phases have given rise to a large recent literature focused on economic indicators and inferences on turning points for European countries.

As markets become more integrated, governments and the private sector seek to conduct their activities in light of both national and international economic conditions. Changes in exchange rates, output, consumption, inflation, and interest rates in different parts of the world can influence the effectiveness of government policies and the competitive position of businesses, even those not directly related to international operations. The benefits of a warning system to detect recessions in major economic partners and in industrialized countries as a whole are considerable. The more reliable the warning system is, the more efficiently economic policy can be implemented as a pre-emptive action against the negative effects of widespread economic weakness and unemployment. Private businesses also benefit from making decisions based on more complete information regarding demand and supply for their services.

This article constructs an international business cycle indicator using a broad production measure of the G7 countries and the Organisation for Economic Co-operation and Development (OECD) member countries.¹ It also builds national business cycle indicators for each of the G7 countries individually using series that comove with their aggregate economic activity. A dynamic factor model with Markov switching (DFMS) is used to combine these macroeconomic series and to estimate probabilities of current

business cycle phases for each of the G7 countries and for the aggregate G7 and OECD measures, which can be used as a warning system to monitor country-specific and international business cycles.²

The novelty of this approach is that we extend the DFMS model to include a filter that minimizes the occurrence of false turning points as it sorts out minor contractions and estimates only major economic recessions and expansions. This feature is especially

The results of this study indicate that some economic recessions and expansions were common to the majority of OECD countries studied, characterizing an international business cycle.

important in situations in which an economy often slows down but does not enter a recession, occurrences that lead to a high rate of false alarms.

The phases of business cycles are well characterized by the model probabilities, which show a clear dichotomy between expansions and recessions for each of the

G7 countries and for the aggregate OECD and G7 measures. The proposed model detects only probabilities of major recessions compared with the probabilities obtained without the filter, which capture several minor contractions for some of the G7 countries. Discerning between major downturns and minor contractions helps to avoid identifying false turning points. This quality is especially important for monetary policy purposes because central banks may want to act only in the event of major recessions affecting several sectors of the economy at the same time, such as employment, sales, output, and income.

OECD countries differ in their institutions, monetary and fiscal policies, industrial compositions and structures, and average aggregate growth rates. The results of this study indicate, however, that OECD countries share some common business cycle phases despite their idiosyncrasies. Some economic recessions and expansions were common to the majority of countries studied, characterizing an international business cycle. The results from the probabilities also suggest that the business cycle derived from the OECD and G7 output data coincides with the swings in the euro area. The OECD countries altogether have experienced three major recessions in the period analyzed: during the oil crisis in the mid-1970s, in the early 1980s, and in the early 1990s.

Comparing the U.S. business cycle with the international business cycle shows that recessions in the United States are more frequent and of shorter duration than in the aggregate OECD in the sample analyzed. The U.S. economy led the beginning and end of the contractions occurring in the rest of the world in the early 1970s and early 1990s, whereas the 1980s recession started and ended at about the same time in the United States and the OECD countries. Some patterns of lead-lag relationship are also revealed in the business cycle phases among the G7 countries.

The article begins with an intuitive explanation of the model and then presents the empirical results for the aggregate OECD countries and for each of the G7 countries.

Constructing the Model

This analysis uses a multivariate system to model business cycle fluctuations in G7 and OECD countries. The model is an extension of the DFMS model, which has been successfully applied to represent business cycles worldwide. As in the DFMS model, an unobservable variable is computed as a nonlinear weighted average of the observed coincident macroeconomic series, and it represents the common information related to business cycles contained in these series. This latent variable switches regimes following a two-state Markov process, which represents expansion and contraction phases of the business cycle.

We extend the model by including a self-adjusting variable-bandwidth filter, which enhances signal-to-noise ratio cycles. The advantage of this filter is that it minimizes the occurrence of false turning points because it removes minor economic contractions and estimates only major recessions and expansions (Chauvet 2005). This filtering is especially important in situations with low signal-to-noise ratios, where the detection threshold in Markov-switching models can be low to capture recessions and can thus lead to a high rate of false alarms when the economy slows down but does not enter a recession.

We apply the model for each of the G7 countries' macroeconomic variables that display simultaneous movements with national gross domestic product (GDP), such as consumption, production, sales, employment, and income, among others. The resulting dynamic factor model characterizes country-specific business cycles. We also apply the model to an aggregate measure of output of twenty-nine OECD countries and to the GDP of each of the G7 countries to obtain a broad measure of the international business cycle shared by most industrialized and semi-industrialized countries. The proposed method tracks business cycle fluctuations and generates coincident probabilities of business cycle phases, which can be used to predict business cycle turning points.

The Data

We use quarterly data to build the coincident indicators for each of the G7 countries individually. These data were obtained from the International Financial Statistics database, Datastream Systems Inc., and the OECD database, with different sample ranges. For the United States, we use the same four coincident variables used by the National Bureau of Economic Research (NBER): measures of sales, personal income, industrial production, and employment.³ For the other six countries, we select four series that correspond closely to the same measurement variables used to build the coincident indicators of the U.S. economy (see Table 1). In particular, industrial production and employment are common variables used for all countries. Different measures of income (such as personal income or wages and salaries) are used for all countries except Japan. Other variables used are sales (retail or manufacturing), electricity consumption, GDP, consumption, and manufacturing orders.

In order to represent a broad measure of international business cycles, we use the aggregate OECD quarterly industrial production series for its country members combined with the GDP of each G7 country. Table 1 summarizes the information about all the series used.

Empirical Results⁴

Business cycle phases are well characterized by the estimated probabilities, which display a clear dichotomy between expansions and recessions for each G7 country

-
1. G7 members are Canada, France, Germany, Italy, Japan, the United Kingdom, and the United States. OECD members are Australia, Austria, Belgium, Canada, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Japan, Korea, Luxembourg, Mexico, the Netherlands, New Zealand, Norway, Poland, Portugal, Spain, Sweden, Switzerland, Turkey, the United Kingdom, and the United States. Since the Slovak Republic became a member only in December 2000, the aggregate industrial production series we use does not include this country.
 2. See Chauvet and Hamilton (2006) for a detailed explanation of the multivariate DFMS model and the univariate Markov switching model.
 3. The NBER's decisions regarding business cycle dates are considered the official U.S. turning points and are used as the benchmark for model comparison.
 4. The model selected by diagnostic and predictive performance tests in identifying turning points is an autoregressive specification of order two for each country and for the aggregate OECD and G7 series.

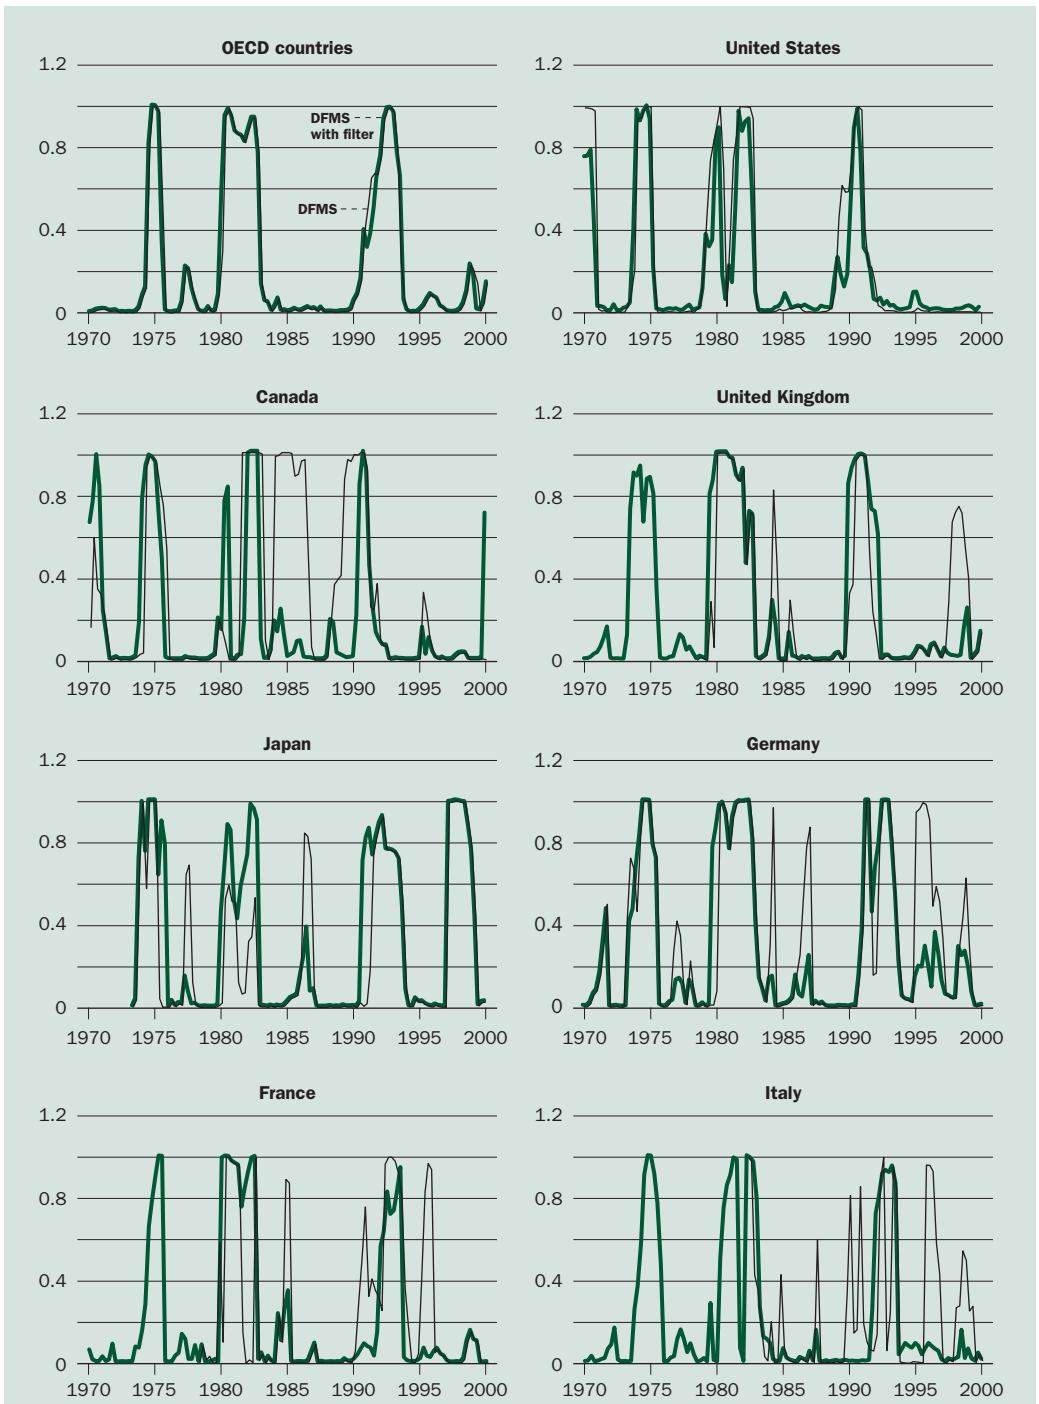
Table 1
Coincident Variables of G7 and OECD Countries

	Series	Sample
OECD countries	Aggregate industrial production for 29 countries	1960Q1–2000Q1
United States	Industrial production Total civilian employment Personal income less transfer payments Manufacturing and trade sales	1959Q1–2000Q2
Canada	Industrial production Employment, business and personal services Personal income Personal consumption expenditures	1967Q2–2000Q1
United Kingdom	Industrial production Employee jobs Disposable income Retail sales	1980Q1–2000Q2
Japan	Industrial production Employees nonagricultural ind. Department stores sales Electric power consumption	1973Q1–2000Q2
Germany	Industrial production Employed persons Gross wages and salaries Manufacturing orders	1962Q2–2000Q2
France	Industrial production Employment except agriculture Gross disposable income Gross domestic product	1978Q1–2000Q1
Italy	Industrial production Employment Wages and salary earnings Private consumption	1982Q1–1999Q4

Source: International Financial Statistics database, International Monetary Fund, Datastream Systems Inc., and OECD database

and for the aggregate G7 and OECD measures, as shown in Figure 1. The coincident probabilities of recessions increase substantially during recessions and display low values during expansions. Figure 1 also compares the probabilities of recession from the DFMS model with and without the self-adjusting variable-bandwidth filter. For the model with the filter, the probabilities of recessions detect only major recessions, but in the model without the filter the probabilities also capture several other minor contractions in addition to the major recessions for some G7 countries. The fact that the probabilities estimated without the filter capture minor contractions is not a disadvantage per se if the goal is in fact to capture them. However,


Figure 1
Coincident Probabilities of Recessions for G7 and Aggregate OECD Countries



Note: The graphs show the probabilities of recessions using a DFMS model with a self-adjusting variable-bandwidth filter and a DFMS model without a filter.

Source: Estimated probabilities from the proposed DFMS model with filter

Figure 2

Probabilities of Recessions for the Aggregate OECD Countries and the United States

if the aim is to discern between major downturns and minor contractions, then the filter reduces the risk of calling false turning points. This feature is especially important for monetary policy purposes because central banks may want to change the size and direction of changes in interest rates depending on the severity of the economic downturn.

In order to analyze business cycle phases, we define turning points based on whether the probabilities of recessions and expansions are smaller or greater than 50 percent. For example, the beginning of a recession occurs when the probability of a recession moves from below 50 percent to above 50 percent. This rule provides a good definition of turning points because the estimated probabilities clearly distinguish times when an expansion is more likely from those when a recession is more likely.

OECD countries.⁵ Figure 2 shows the full-sample probabilities of recession for the aggregate output of the OECD and GDP of each G7 country. The probabilities of recessions and expansions can be interpreted as a representation of business cycle phases for industrialized and semi-industrialized countries given that the analysis includes twenty-nine member countries.

Table 2 summarizes some features of the probabilities of recession measure. The average duration of a recession shared by OECD countries is eight quarters, and the average probability that the economy will enter a recession is 87 percent. Expansions last twenty quarters on average, and the average probability of entering an expansion is 95 percent. According to the recession probabilities, OECD countries altogether have experienced three major recessions in the period analyzed: during the oil crisis in the mid-seventies, in the early eighties, and in the early nineties. The results from the probabilities suggest that the business cycle obtained from the broad OECD output measure coincides with the euro area's business cycle. The timing of recessions is very close to the euro area's recessions as dated by the Centre for Economic Policy Research (CEPR) Business Cycle Dating Committee (see the shaded area in the second panel of Figure 2), which is a European counterpart to the NBER Business Cycle

Table 2
Estimated Business Cycles of OECD and G7 Countries

	Number of full recessions	Average expansion probability	Average expansion duration (quarters)	Average recession probability	Average recession duration (quarters)
OECD	3	0.95	20	0.87	8
United States	4	0.94	17	0.84	6
Canada	4	0.95	20	0.88	8
United Kingdom	3	0.96	25	0.86	7
Japan	4	0.95	20	0.83	6
Germany	3	0.90	10	0.89	9
France	3	0.96	25	0.87	8
Italy	3	0.95	20	0.87	8

Source: Authors' calculations based on estimated probabilities from the proposed DFMS model with filter

Table 3
Business Cycle Dating for OECD Countries and the Euro Area

CEPR dating for the euro area		Model dating for OECD countries	
Peak	Trough	Peak	Trough
1974Q3	1975Q1	1974Q3	1975Q2
1980Q1	1982Q3	1980Q1	1982Q4
1992Q1	1993Q3	1991Q3	1993Q3

Source: CEPR (2003); authors' calculations based on estimated model probabilities

Dating Committee.⁶ During periods that the CEPR classifies as expansions, the probabilities of recessions are generally close to zero. At CEPR peak dates (the onset of recessions), the probabilities of recession increase substantially above 50 percent and stay high until the trough dates (the end of recessions).

Table 3 compares the CEPR recession dating for Europe and the recession dating obtained from our model of OECD countries. From the six estimated turning points in the period studied (three peaks and three troughs), three match exactly, and the other three are off by only one or two quarters. This dating also coincides with the the euro

5. Since G7 members are also OECD members, the results of the combination of aggregate G7 outputs are subsumed in the OECD results.

6. Although the techniques used differ between the NBER and the CEPR, the dating generated by these institutions is similar in the sense that it is often used as a benchmark. The euro area considered by the CEPR includes Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, and Spain.

Figure 3
Probabilities of Recessions for the United States





Figure 4
Probabilities of Recessions for Canada



area dating by Artis, Marcellino, and Proietti (2003), Dopke (1999), Artis, Krolzig, and Toro (2004), Anas and Ferrara (2004), and Krolzig (2001), among others.

Figure 2 also compares the probabilities of recession for OECD countries and the U.S. economy, the NBER dating for U.S. recessions, and the CEPR dating of recessions for the euro area. Recessions in the United States are more frequent and of shorter duration than in the aggregate OECD countries during the period studied. The U.S. economy led the beginning and end of contractions occurring in the OECD countries in the early 1970s and early 1990s recessions whereas the 1980s recession

Figure 5
Probabilities of Recessions for the United Kingdom

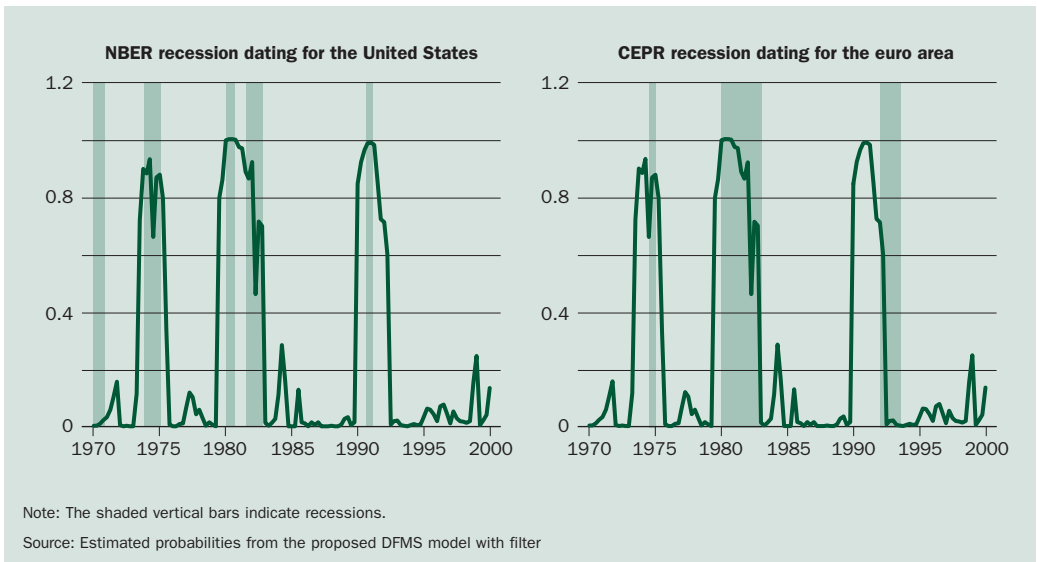



Figure 6
Probabilities of Recessions for Japan



started and ended at about the same time in the United States and the OECD. However, the U.S. economy experienced two recessions between 1980 and 1982 while only one long recession occurred in the OECD countries altogether.

G7 countries. Figures 3–9 plot the probabilities of recession for all G7 countries and contrast these probabilities with the NBER dating for the United States (the first panels of Figures 3–9) and the CEPR dating for the euro area (the second panels of Figures 3–9) for each country. The probabilities of recession show some similarities and differences in the business cycles of the G7 countries. The G7 countries also

Figure 7
Probabilities of Recessions for Germany

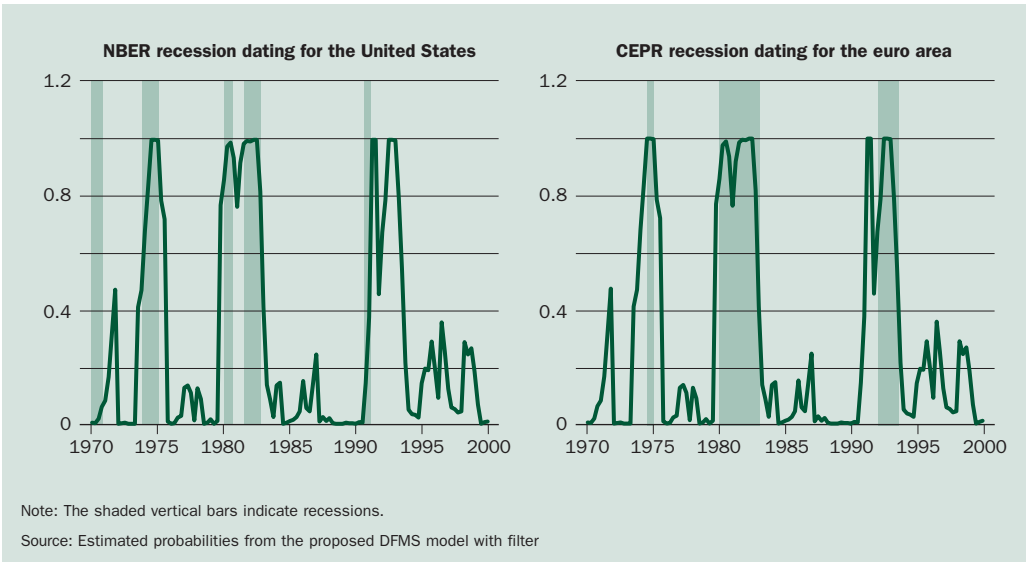




Figure 8
Probabilities of Recessions for France



experienced three or four full recessions in the period studied;⁷ Japan experienced an additional recession in 1997–99.

The most similar recession across the G7 countries is the one that took place in the mid-1970s, which hit all economies at about the same time. The recession in the early 1980s was a long one, lasting three or four years for some countries (France, Germany, the United Kingdom, and Japan) and for the aggregate OECD and G7 measures, whereas for a few countries (Italy, the United States, and Canada), two shorter recessions instead occurred close to each other during the same period.

Figure 9
Probabilities of Recessions for Italy



The main difference in business cycles among these countries concerns the early 1990s recession. This recession started earlier in the United Kingdom, the United States, Canada, and Japan while in Germany and the other G7 countries this recession did not begin until one or two years later. For the aggregate OECD and G7 countries, this recession started and ended at about the same time as the CEPR date for the euro area (Figure 2). The NBER dates the beginning of this recession in the United States in July 1990 while the CEPR dates the start of the recession in the first quarter of 1992.

The closest estimated probabilities of recessions are for the United States and Canada. Recessions began and ended at about the same time in these two countries. Italy and France also have very close recession timing. The one difference between these two countries is in the early 1980s: France experienced a single long recession while Italy had two recessions during this period.

The probabilities of recession for Germany, the United Kingdom, and Japan are also very similar for the first two recessions in the sample. The probabilities suggest that recessions in the United Kingdom occurred slightly ahead of those in Germany and Japan and occurred more closely to recessions in the United States and Canada. In the 1990s recession, the U.K. economy contracted even before the U.S. and Canadian economies. Overall, recessions in the United Kingdom occurred earlier than in other European countries, followed by Germany. Recessions in the United Kingdom also lasted longer than those in the United States, Canada, and Germany.

The Japanese economy displays dynamics similar to the other G7 and OECD countries in the 1970s and 1980s. However, Japan experienced two severe and long recessions in the 1990s: one in 1991–94 and another in 1997–99 (Figure 6). The earlier recession hit Japan at about the same time that it hit the United States but did not end until much later, coinciding with the trough of the recession in the OECD countries.

7. The sample begins in 1970 and therefore does not include the recessions that occurred in the United States and Canada around 1969–70.

The Asian financial crisis that started in 1997 marked the beginning of a second 1990s recession in Japan that was not experienced by any of the other G7 countries studied.

Conclusions

This article constructs business cycle indicators for the G7 countries and for an aggregate measure of output by twenty-nine industrialized and semi-industrialized OECD member countries. We extend the Markov-switching dynamic factor model by adding a self-adjusting variable-bandwidth filter. The model yields output probabilities of the current business cycle phase for each G7 country and for the aggregate OECD and G7 output measures, which can be used as a warning system to monitor country-specific and international business cycles.

As a result of the filter, the probabilities of recession display a clearer distinction between recessions and expansions, reducing the risk of calling false turning points. We find a common business cycle underlying the twenty-nine OECD countries, characterizing an international business cycle. The probabilities of recessions for the aggregate OECD countries indicate that they shared three major recessions in the period analyzed: during the oil crisis in the mid-1970s, in the early 1980s, and in the early 1990s. The most similar recession in terms of timing and duration across countries is the one that took place in mid-1970s, and the most divergent is the one that occurred in the early 1990s.

REFERENCES

- Anas, Jacques, and Laurent Ferrara. 2004. A comparative assessment of parametric and non-parametric turning points detection methods: The case of the Euro-zone economy. In *Papers and proceedings of the third Eurostat colloquium on modern tools for business cycle analysis: Statistical methods and business cycle analysis of the Euro zone*, edited by Gian Luigi Mazzi and Giovanni Savio. Luxembourg: Office for Official Publications of the European Communities.
- Artis, Michael, Hans-Martin Krolzig, and Juan Toro. 2004. The European business cycle. *Oxford Economic Papers* 56, no. 1:1–44.
- Artis, Michael, Massimiliano Marcellino, and Tommaso Proietti. 2003. Dating the euro area business cycle. CEPR Discussion Paper No. 3696, January.
- Burns, Arthur F., and Wesley C. Mitchell. 1946. *Measuring business cycles*. New York: National Bureau of Economic Research.
- Centre for Economic Policy Research. 2003. Euro Area Business Cycle Dating Committee. Press release, September 22. Available online at <www.cepr.org/press/dating.pdf>.
- Chauvet, Marcelle. 2005. Estimating multivariate models with latent variable Markov switching. University of California, Riverside, photocopy.
- Chauvet, Marcelle, and James H. Hamilton. 2006. Dating business cycle turning points in real time. In *Nonlinear time series analysis of business cycles*, edited by Costas Milas, Philip Rothman, and Dick Van Dijk. Vol. 276, *Contributions to economic analysis*. Amsterdam: Elsevier Science and Technology.
- Dopke, Jorg. 1999. Stylised facts of euroland's business cycle. *Jahrbucher fur Nationalokonomie und Statistik* 219, nos. 5–6:591–610.
- Krolzig, Hans-Martin. 2001. Business cycle measurement in the presence of structural change: International evidence. *International Journal of Forecasting* 17, no. 3:349–68.