Volume Title: Asset Prices and Monetary Policy

Volume Author/Editor: John Y. Campbell, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-09211-9

Volume URL: http://www.nber.org/books/camp06-1

Conference Date: May 5-6, 2006

Publication Date: September 2008

Chapter Title: Noisy Macroeconomic Announcements, Monetary Policy, and Asset Prices

Chapter Author: Roberto Rigobon, Brian Sack

Chapter URL: http://www.nber.org/chapters/c5375

Chapter pages in book: (p. 335 - 370)

8

Noisy Macroeconomic Announcements, Monetary Policy, and Asset Prices

Roberto Rigobon and Brian Sack

8.1 Introduction

The relationship between economic data, on the one hand, and asset prices and monetary policy, on the other, has become a widely studied topic in the academic literature—and for good reason. Macroeconomic conditions are a key factor determining near-term policy expectations, and those expectations reverberate throughout the financial system by influencing the returns expected on all asset classes.

But despite being widely studied, our current knowledge of the interactions between economic news and asset prices has many shortcomings, and the results are puzzling in some dimensions. Perhaps most importantly, the estimated effects of data releases on monetary policy expectations and asset prices are found to be relatively small. This is the case even for those assets that are known to be very sensitive to near-term monetary policy expectations, such as eurodollar futures and short-term Treasury securities.

This finding is surprising. After all, the literature over the past two decades has argued that monetary policy to a large extent responds systematically to economic conditions. Indeed, the literature has made tremendous progress estimating monetary policy rules that account for these systematic responses in terms of low-frequency data (such as quarterly data). If monetary policy is so systematic, one would expect to see evidence of it also in the higher-frequency movements in interest rates and

Roberto Rigobon is an associate professor with tenure at the Sloan School of Management, Massachusetts Institute of Technology, and a research associate of the National Bureau of Economic Research. Brian Sack is vice president of Macroeconomic Advisers, LLC.

We thank John Y. Campbell and Laurence H. Meyer for valuable comments and suggestions, and Mike McMorrow for excellent assistance with the analysis.

asset prices around data releases. That is, the major economic data releases would be expected to explain an extensive amount of the variation in assets sensitive to near-term policy expectations.

In our view, the puzzle of the "detachment" of monetary policy expectations and asset prices from the incoming economic news is partly related to the difficulties associated with measuring the surprise component of that news. Most studies to date compute a "surprise" measure for a given release based on expectations taken from a survey conducted ahead of the release. They then regress changes in an asset price on this surprise measure, which we refer to as the standard "eventstudy" approach. The attempt to isolate the unexpected component of the release was a vast improvement over earlier efforts that could not make such a separation, as only the unexpected component should prompt a market reaction. However, this approach likely falls short of accurately measuring the market effects of the incoming news—perhaps considerably.

A problem with the standard eventstudy approach is that the macroeconomic news is likely to be measured very poorly, for several reasons. First, it is hard to accurately measure what the markets are expecting for a given release at the time it comes out, including the full distribution of risks seen for the release. Second, even if one accurately measured expectations, the actual release may be seen as a noisy indicator of the underlying true fundamental factor that drives market responses. And third, the variable measured is usually only one component of a report. After all, most of these reports are complicated, providing lots of information of varying relevance.

Thus, it is quite likely that the macroeconomic surprise included on the right-hand-side of the eventstudy is only a very rough measure of the true incoming news. This chapter focuses on measuring the reaction of asset prices and monetary policy expectations to the "true" economic news embedded in the major U.S. data releases. Rather than attempting to better measure the data or the expectations, we focus on developing econometric techniques that will adequately deal with the measurement problems associated with the data surprises used in the existing eventstudy literature.

Our efforts take us in two directions. First, we modify the standard eventstudy regression framework to account for the possibility that the measured surprises contain error. The measurement issues considered here lead to a classical error-in-variables problem of a standard regression, one that biases downward the estimated sensitivity of asset prices to the incoming data. We develop a new estimator that allows for measurement error and, hence, eliminates this downward bias. The procedure could be used in other applications to correct for the error-in-variables problem.

Second, we employ a principal components approach that removes the need to even try to measure the data surprises. In effect, the approach uses the observed market reactions to infer what the true data surprises were. Such an approach may have appeal if one regards the incoming data as being complex and having many dimensions that could affect asset prices—conditions that make it difficult to measure the data surprise in the manner of the standard eventstudy exercise.

The results provide us with unbiased estimates of the response of monetary policy expectations and asset prices to the "true" surprise contained in all of the major data releases. They also allow us to recover the importance of those true surprises. An important finding from the chapter is that macroeconomic data releases matter to a much greater extent than found in previous studies—that is, they account for a greater portion of the fluctuations in market interest rates. Moreover, using these estimators, we are able to refine a set of patterns in the responses that should be explained by any model addressing the interactions between economic variables, monetary policy, and asset prices.

8.2 Estimating the Effects of Macroeconomic Announcements: Current Methods

Researchers in both macroeconomic and financial economics are very interested in understanding the linkage between monetary policy and asset prices. To that end, one strand of literature has attempted to measure the response of asset prices to monetary policy "shocks," or the erratic and unpredictable component of monetary policy decisions. But such shocks are limited in size and account for only a very small portion of the variation in asset prices. Instead, most of the movement in short-term interest rates likely represents the systematic response of monetary policy to economic developments. Thus, it may be more relevant to investigate the responses of monetary policy expectations and asset prices to incoming news about the economy.

A sizable literature has taken up this topic and has provided us with some valuable results. The studies to date almost uniformly take an approach that is commonly referred to as "eventstudy."

8.2.1 The Eventstudy Specification

Papers in the eventstudy literature typically proceed in a simple regression framework in which the reaction of a given asset price (or market yield) is regressed on the surprise components of the data release, as in the following specification:

$$\Delta s_t = \gamma \cdot z_t + \varepsilon_t, \tag{1}$$

$$z_t = M_t - E_{t-\tau}[M_t], \tag{1'}$$

where M_t is the released value of the macroeconomic announcement, and $E_{t-\tau}[M_t]$ is a measure of the market's expectation ahead of the release. The specification assumes that the only market-moving information is the surprise component of the release z_t, and the parameter γ is the market sensitivity to that surprise—which is the primary interest of this chapter.

The basic approach implicit in specification (1) has not varied much over time, but the empirical implementation of the equation has changed in two dimensions.

First, the measure of expectations has improved. Early papers in this area had to model the market's expectations either as past realized values of the macroeconomic variables or as the outcome of forecasting models that do not necessarily perform very well. More recently, researchers have increasingly relied on surveys to measure expectations and to better isolate the surprise component of data releases. Hence, the measurement of the variable z_t has likely improved over time.

Second, studies have increasingly used a narrower window to measure the market response to the data release. Whereas earlier papers may have used monthly or quarterly data, the eventstudy literature has moved to using changes at a daily frequency (see, for example, McQueen and Roley [1993] and Gürkaynak, Sack, and Swanson [2005]) or even in some cases on an intraday bases (see, for example, Fleming and Remolona [1997] and Balduzzi, Elton, and Green [2001]).[1] The idea of using a narrower window is to reduce the influence of other events that might be affecting the asset price in addition to the data surprise. In terms of the equation (1), it reduces the variance of the error term ε_t, which should improve the accuracy of the estimate of the parameter γ.

The eventstudy approach has importantly contributed to our understanding of the manner in which monetary policy expectations and asset prices react to incoming economic data. Indeed, as we will show in the following, this approach finds that the market reaction to a number of releases is statistically significant. Nevertheless, in our view, the eventstudy approach has some shortcomings that prevent it from recovering the market response to a "true" macroeconomic data surprise.

8.2.2 The Econometric Problem: Noisy Data Surprises

The potential problem that arises with the eventstudy approach is that the results will only be as good as the measure of data surprises included on the right-hand side of the equation. Indeed, the model (1) implicitly assumes that the measured data surprise z_t truly captures the true macroeco-

1. Several other papers use intraday data but focus primarily on foreign exchange rates, including Andersen et al. (2003) and Faust et al. (2003). As an example of a paper using low-frequency data, Cutler, Poterba, and Summers (1989) attempt to measure the influence of macroeconomic shocks on equity prices using a monthly vector autoregression (VAR). Stock and Watson (2003) provide a review of other papers that examine the relationship between financial variables and macroeconomic conditions at a monthly or quarterly frequency. Note, however, that the primary interest in that paper is measuring the predictive power of financial variables for economic outcomes rather than the effects of economic outcomes on financial variables.

nomic news arising from the releases. If that is not the case, the estimated parameter γ will be biased.

Even with the improvements noted in the preceding, it is a somewhat dubious assumption that the variable z_t is perfectly measured—or that it is even well measured. Instead, it is more plausible that the variable z_t contains considerable measurement error, from a variety of sources.

First, it is unlikely that the survey measures used accurately capture the market expectations at the time of the release. In the results presented in the following, we collect those expectations from two surveys and splice them together to create a full time series. Before September 2004, we use the median response from the Money Market Services (MMS) survey, which is a survey of professional forecasters taken the Friday before each release. Since then, we instead use the median response from the regular survey taken by Bloomberg. This figure is the most commonly discussed measure of consensus expectations in the financial markets.

But there are a number of reasons to believe that the expectations measured from these surveys are not necessarily appropriate for gauging the market response. The survey respondents are not the relevant market participants whose expectations matter. Moreover, the survey covers a variety of respondents with very different backgrounds and skill sets, raising questions about whether certain individual responses could distort the measures. It is not even clear that the respondents have the correct incentive scheme, as we suspect that they may assign greater utility to having an out-of-consensus call that comes in correct than having a consensus call that comes in correct. And, last, we arbitrarily use the median from the panel, though the argument for using this over the mean or some other measure is not clear-cut.

In addition to concerns about the cross-section of panelists, we also have some concerns about the timing of the surveys. Ideally, we would like to know the market expectations the moment before the data release. The MMS survey is instead taken the Friday before the release, making it somewhat stale. For those releases that come out on a Friday (e.g., the employment report), that leaves an entire week (and all the data released that week) for expectations to evolve and move away from the survey response. And the situation for the Bloomberg expectations is even worse. Those responses are submitted at irregular times. Most respondents enter their estimates about a week before the releases, but many, instead, do it two weeks in advance, while others wait until the week of the release.[2]

Another source of mismeasurement of the macroeconomic surprise is

2. To take an example, consider the employment report that was released on November 5, 2004. Of the seventy-eight responses to the survey, thirteen were submitted more than two weeks in advance. Most of the responses, thirty-nine, came in about one week in advance (with two others coming in earlier that week). And twenty-four respondents waited until the week of the release to submit their views.

the data release itself. The released data can be thought of as a noisy version of the "true" economic fundamental to which the market responds. Researchers usually focus on just one aspect of the release, and often that one aspect can appear anomalous. A recent example was the advance gross domestic product (GDP) report for the fourth quarter of 2005, which came in well below the market's expectations. That surprise owed in large part to a puzzling drop in defense spending that quarter, and, hence, Wall Street analysts generally dismissed the implications of the report.[3]

Overall, we believe that the measured data surprises could be quite noisy. Market expectations are probably not measured particularly well, as the survey used is a random variable that at best can be considered to be unbiased but not measured without error. And the actual release is likely to contain some noise relative to the true macroeconomic news that affects markets.

8.2.3 The Bias in Eventstudy Estimates

We start with the assumption that the macroeconomic surprises used in the eventstudy literature are measured with error for the reasons discussed in the preceding. In this case, the estimates obtained in the standard literature are plagued with error-in-variables bias.

To provide some structure for discussing the problem, we assume the asset price change immediately around the release at time t is denoted by Δs_t. This market reaction is driven by the true macroeconomic news contained in the announcement, which we denote z_t^*, according to the following equation:

$$\Delta s_t = \gamma \cdot z_t^* + \varepsilon_t. \tag{2}$$

We are interested in measuring the sensitivity of financial markets to the true economic news, captured by the parameter γ. The residual ε_t captures movements in the asset price in that window that are not driven by the data surprise (or at least not under this linear structure).

To estimate equation (2), most researchers attempt to measure the true macroeconomic news z_t^* as the difference between the released data and the expectation of that data, where the expectation is typically determined from a survey taken in advance of the release. But, as previously discussed, there are two potential problems with that measure—that the release may be seen as a noisy version of the true relevance of the news and that the expectations may be measured poorly. Considering this, we should perhaps

3. For example, David Greenlaw from Morgan Stanley summarized the report as follows: "Much weaker than expected report. Both final sales and inventories came in well below expectations in Q4 [2005]. However, we believe that a significant portion of the downside is likely to be recouped in Q1 . . . Defense [spending] plunged 13% in Q4. We suspect that at least some of this drop reflects a timing quirk that will be unwound in Q1." (David Greenlaw and Ted Weisman, *GDP [Q4 Advance],* Morgan Stanley Economic Data Bulletin, January 27, 2006)

take the measured data surprises to be a noisy representation of the true economic news, as follows:

$$(3) \qquad z_t = z_t^* + \eta_t,$$

where z_t denotes the measured data surprise. In this case, the mismeasurement of the true data surprise is captured in the variable η_t.

Using this proxy for the true macroeconomic news, researchers typically resort to estimating the following equation:

$$(4) \qquad \Delta s_t = \gamma \cdot z_t + \upsilon_t,$$

using an ordinary-least-squares (OLS) regression. However, given the preceding structure, the error term from the estimated equation is

$$(5) \qquad \upsilon_t = \varepsilon_t - \gamma \cdot \eta_t,$$

which is negatively correlated with the right-hand-side variable in the regression. This correlation, of course, results in the bias in the regression estimate of γ.

To quantify the bias, we assume that the true macroeconomic news has a variance of $\sigma^2_{z^*}$ and that the measurement error is mean zero conditional on the true surprise ($E_t[\eta_t|z_t^*]$) and has a variance of σ^2_η. We also assume that the portion of the asset price movement not explained by the macroeconomic surprise (ε_t) is mean zero conditional on the true news and the measurement error ($E_t[\varepsilon_t|\eta_t, z_t^*]$) and has a variance of σ^2_ε. Under these assumptions, the estimate obtained by an OLS regression is:

$$(6) \qquad \hat{\gamma}_{\mathrm{OLS}} = \gamma - \gamma\frac{\sigma^2_\eta}{\sigma^2_{z^*} + \sigma^2_\eta}.$$

This estimate has the standard downward bias (toward zero), which is the standard result in the presence of an error-in-variables problem. Based on this consideration, we argue that the typical eventstudy estimation may understate the influence of macroeconomic news on asset prices.

At this point, it is useful to note that we have considered two forms of mismeasurement of the macroeconomic news—one based on noise in our reading of the market's expectations, and one based on noise in the release itself. Both forms are captured by equation (3), and, hence, the bias in the OLS estimates applies to both of them. Nevertheless, the interpretation of the results is different depending on which of the two sources predominantly accounts for the mismeasurement. If the mismeasurement is in terms of measuring the market's expectations, then the OLS estimates are actually missing part of the market reaction. If instead the noise is contained in the actual data, then the market is reacting by less, as it is doing the signal extraction problem and discounting the value of the released data. In that case, the OLS estimates are an accurate measure of the true (but limited) market reaction to the released data.

We are interested in discovering the market reaction to the true surprise, adjusting for the measurement error from these two sources. There are several potential solutions. One is to find an instrument, something that is correlated with the true macroeconomic news but uncorrelated with the measurement error. But such instruments do not exist, leaving the problem of estimation unresolved. Another solution is to improve the data itself, for example, by better measuring market expectations. In that regard, the emergence of economic derivatives may be useful in that they may provide a more accurate and timely reading of market expectations. Still, given all of the preceding considerations, it is not clear that we will ever have a fully accurately measure the macroeconomic news.

In this chapter, we take an alternative approach in which we attempt to address the issue through econometric technique. We will ultimately develop two methods that help us resolve some of these issues and allow us to better understand the linkages from economic news to asset prices and monetary policy expectations.

8.3 Identification through Censoring

The problem of error-in-variables that we discuss in the preceding is, in fact, a problem of identification. To see that, consider the case of measuring the effect of a single data release on a single asset price. In that situation, we can compute only three statistics: the variance of the asset price, the variance of the macroeconomic news, and the covariance between them. The problem is that these moments are determined by four underlying parameters: γ, $\sigma^2_{z^*}$, σ^2_{η}, and σ^2_{ε}. Thus, the solution is not identified, or there is a continuum of solutions.

In the preceding we noted that an instrumental-variables approach is one way of solving the problem, if one were able to find an appropriate instrument. Note that the availability of such an instrument basically solves the identification problem. For a variable ω_t to be a valid instrument, it must be correlated with the true news but uncorrelated with the measurement error, as follows:

$$z^*_t = \beta \cdot \omega_t + \kappa_t, \tag{7}$$

The availability of this instrument adds three pieces of information (the variance of ω_t, its covariance with the measured news, and its covariance with the asset price response), while only adding two unknown variables (β and the variance of κ). As long as β is different from 0, these additional conditions resolve the identification problem. However, as noted in the preceding, we cannot think of an instrument that is valid in the circumstances studied in this chapter.

In the absence of a valid instrument, the question is whether we can solve the identification problem through some other means. We will do so by developing a new technique that we label "identification through censoring."

8.3.1 The Case of One Macroeconomic Announcement

To demonstrate the methodology, we first assume that there is only one macroeconomic announcement at a given time. One special feature of macroeconomic announcements is that they occur at prespecified days. This is important because it implies that we can find a sample of other days (or times) at which the magnitude of the surprise variable is exactly 0. When the variable is exactly equal to 0, it means that its error-in-variables is 0 as well. This "censoring" of the measurement error will provide the identification.[4]

Formally, this situation can be described by the following equation:

$$\Delta s_t = \begin{cases} \gamma \cdot z_t + \varepsilon_t & t \in D \\ \varepsilon_t & t + 1 \in D \end{cases} \tag{8}$$

where D is the set of days (or times) on which the announcements take place. We are assuming that no announcements take place the day before those included in D. Under the assumption that the disturbance ε_t is homoskedastic, we can use the variance of the asset price observed at time $t - 1$ as additional information in the identification. In that case, the following equations hold:

$$\begin{aligned} var(\Delta s_{t-1}) &= \sigma_\varepsilon^2 \\ var(\Delta s_t) &= \gamma^2\sigma_{z^*}^2 + \sigma_\varepsilon^2 \\ var(z_t) &= \sigma_{z^*}^2 + \sigma_\eta^2 \\ cov(\Delta s_t, z_t) &= \gamma\sigma_{z^*}^2 \end{aligned} \tag{9}$$

This is a system of four equations and four unknowns that can be solved for the parameters. Most importantly, the sensitivity of the asset price to the incoming news can be solved as follows:

$$\gamma = \frac{var(\Delta s_t) - var(\Delta s_{t-1})}{cov(\Delta s_t, z_t)}. \tag{10}$$

This estimator is in the spirit of Rigobon and Sack (2004), in which the estimator depended on the change in the variance relative to the change in the covariance. Here, the change in the covariance is just the covariance itself, as the macroeconomic surprise has no variance when it is censored.

The preceding computations rely on the assumption that the structural

4. This intuition comes from Goldberger (1991), who argues that the variance of the error-in-variables in survey data depends on the size of the announcement. He used the following example: If you ask how many cigarettes a person smokes in a day, a nonsmoker will answer zero—and that reply has no error-in-variables whatsoever. But someone who smokes a pack and a half a day will probably have a sizable error. In other words, the magnitude of the error depends on the magnitude of the reply, with complete censoring of the error at zero.

shocks in the asset price equation (ε_t) are homoskedastic. This is a fairly strong assumption, and one that is not necessary. To derive an estimator like equation (10), all we need is a prediction of what the variance would have been like in the absence of the macroeconomic news. Thus, we can incorporate heteroskedasticity to the degree that it is predictable. In other words, the identifying assumption is that the variance of ε_t is predictable.

For example, suppose we observe a release at 8:30, and as a "control window" we use a thirty-minute interval from the previous afternoon at 2:30. The preceding assumes that the variance of ε_t around 8:30 is the same as that around 2:30 on the previous day. But even on days of no announcements, this does not seem to be the case. Instead, we require a much weaker assumption—that the shift in the variance of ε_t on announcement days is the same as the shift on nonannouncement days, or

$$(11) \qquad \sigma^2_{\varepsilon,t-1,8:30} - \sigma^2_{\varepsilon,t-2,2:30} = \sigma^2_{\varepsilon,t,8:30} - \sigma^2_{\varepsilon,t-1,2:30}.$$

This assumption allows for the data to have heteroskedasticity over our sample, as long as that heteroskedasticity looks the same on announcement and nonannouncement days. In this case, the estimator (10) still works if we replace the variances with the shift in the variances. This is the assumption that we will employ in the empirical results in the following.

This estimator eliminates the bias coming from error-in-variables that affects the typical OLS estimates. However, the estimator is only as good as its identifying assumptions. The two main identification assumptions needed are that the errors-in-variable are classical and that the variance of the asset prices is predictable (so that we can make an accurate judgment of what the variance would have been in the absence of the macroeconomic surprise). Conditional on those identifying assumptions, the coefficients from this procedure are accurate. However, if either of the two main assumptions is violated, the estimates are biased. We will return to these issues in the following.

8.3.2 The Case of Multiple Macroeconomic Announcements

The Bureau of Economic Analysis (BEA), the Bureau of Labor Statistics (BLS), and other government agencies would make our lives easier if they released one statistic at a time. Unfortunately, this is not the case. Because different releases follow different schedules, often multiple important releases will randomly coincide in both the date and time.

If this problem were just limited to coincidence, we could deal with it by simply eliminating those days with multiple releases. Unfortunately, some of the data releases always coincide with one another. This is the case for those reports that include multiple statistics that have market influence. For example, the employment report involves the simultaneous release of nonfarm payrolls, the unemployment rate, and average hourly earnings—each of which are found to have an independent effect on markets.

In the OLS framework, we can deal with this simultaneity by simply putting the multiple releases into a single regression. We can also address this issue in the identification-through-censoring approach. To achieve identification in such circumstances, it turns out that we simply have to incorporate more than one asset price. For simplicity, we will show this point for the case of two announcements. Also, for simplicity let us assume that the structural shock ε_t is homoskedastic. In this case, the model has the following structure:

$$(12)\qquad \begin{aligned} \Delta s_t &= \gamma_1 \cdot z^*_{1,t} + \gamma_2 \cdot z^*_{2,t} + \varepsilon_t \\ z_{1,t} &= z^*_{1,t} + \eta_{1,t} \\ z_{2,t} &= z^*_{2,t} + \eta_{2,t}, \end{aligned}$$

where the errors in measuring the true surprises ($\eta_{1,t}$ and $\eta_{2,t}$) are likely to be correlated.

Note first that the identification is lost. The covariance matrix of the asset price and the two measures of macroeconomic surprises provides six equations, and the variance of the asset price when there are no surprises provides a seventh moment. But the model has nine unknown parameters: γ_1, γ_2, σ^2_ε, $\sigma^2_{z_1}$, $\sigma^2_{z_2}$, $\sigma^2_{\eta_1}$, $\sigma^2_{\eta_2}$, the covariance between z^*_1 and z^*_2, and the covariance between η_1 and η_2. The underidentification is even more severe in the case of three simultaneous announcements.

The solution to the problem is to consider additional asset prices. If we consider two asset prices, we have the following system of equations:

$$(13)\qquad \begin{aligned} \Delta s_{1,t} &= \gamma_{1,1} z^*_{1,t} + \gamma_{1,2} z^*_{2,t} + \varepsilon_{1,t} \\ \Delta s_{2,t} &= \gamma_{2,1} z^*_{1,t} + \gamma_{2,2} z^*_{2,t} + \varepsilon_{2,t} \\ z_{1,t} &= z^*_{1,t} + \eta_{1,t} \\ z_{2,t} &= z^*_{2,t} + \eta_{2,t}, \end{aligned}$$

where the structural shocks $\varepsilon_{1,t}$ and $\varepsilon_{2,t}$ are possibly correlated, and the errors in the macroeconomic surprises are, as before, also correlated. We have now achieved identification. The variance-covariance matrix of the asset prices and the macroeconomic surprises on both announcement and nonannouncement days provides thirteen moment conditions. These are sufficient to solve for the thirteen unknown parameters.[5]

What delivers the identification? It comes from the fact that the noise contained in our measures of the macroeconomic announcements has to

5. Adding the second asset price brings six new moment conditions—its variance and its covariance with the other asset price on both announcement days and nonannouncement days, and its covariances with the two measures of surprise on announcement days and its variance on nonannouncement days) while adding only four new parameters ($\gamma_{2,1}$, $\gamma_{2,2}$, $\sigma^2_{\varepsilon 2}$, and the covariance between ε_1 and ε_2).

be the same independent of the asset price we are considering. That restriction allows the incorporation of an additional asset to bring new information for the identification.

8.3.3 Implementation of the Estimator

In the following results, we will include five different asset prices and will allow for as many as three simultaneous releases. (All details are described in the next section.) This set-up implies that our estimator is always overidentified. To estimate the parameter values, we use a generalized method of movements (GMM) estimator that seeks to minimize the squared deviations of the errors for each moment condition.[6] It can be shown that this estimator is consistent and that the estimates are asymptotically normal.

8.4 The Estimated Effects of Macroeconomic Surprises

This section begins by describing the data that we use and some of the specific decisions made in implementing the various approaches. It then provides some results from both the standard eventstudy estimator and the identification-through-censoring approach.

8.4.1 Data

In the results that follow, we measure the reaction of five financial variables to incoming macroeconomic news. The set of financial variables is intended to capture the behavior of near-term policy expectations as well as broader asset prices.

Specifically, we include several near-term interest rates that are very sensitive to monetary policy. Eurodollar futures rates ere probably the most useful, liquid instrument for that purpose. We, therefore, include the rates on the second and fourth eurodollar contracts to expire—which will reflect changes in monetary policy expectations roughly at horizons of six and twelve months ahead.[7] We also include the two-year Treasury yield, which is very sensitive to the expected path of monetary policy beyond the hori-

6. So that the relative importance of the moment conditions is not influenced by the unit of measure, we normalize the movements in each asset price by their standard deviation. The results, however, are expressed in terms of basis points for yields and percentage points for equities.

7. The second contract will have between three and six months to expiration (with an average of 4.5). It is tied to the three-month Libor rate, which will be sensitive to the expected average funds rate over those three months (with an average of 1.5). Adding together these averages yields six months. Similar calculations yield twelve months for the fourth contract. We exclude the first and third contracts because we felt that much of their information would be redundant. In addition, we worried that the variation in the expiration of the first contact from zero to three months might be more problematic (given institutional details such as the spacing of meetings).

zon covered by the eurodollar contracts, and the ten-year Treasury yield. Last, we include the S&P 500 index.[8]

For all of these asset prices, we use intraday data. This feature alone provides a sizable improvement over daily eventstudy exercises. As noted in the preceding, with intraday data, we can look at a narrow window around the time of the release—an interval that includes the influence of data releases at a given time but excludes most other market-moving events. In effect, we are shrinking the size of the error term ε_t relative to the influence of the data.

The intraday data slices we consider are thirty-minutes long, beginning five minutes before the time of an announcement to avoid any complications from variation in the precise timing of the quotes or of the releases. The data releases that we consider all take place at either 8:30 a.m. 9:15 a.m., or 10:00 a.m., giving us slices that run from 8:25 to 8:55 a.m., 9:10 to 9:40 a.m., and 9:55 to 10:25 a.m.

For equities, unfortunately, we only have intraday quotes from when the stock market is open, from 9:30 a.m. to 4:00 p.m. Thus, we have to modify our slices accordingly. For the 8:30 a.m. and 9:15 a.m. releases, we use the change in the S&P index from the previous close to 9:55 a.m. For the 10:00 a.m. release, we can use the same slice that we use for the interest rates.

The control window that we use in each case is a thirty-minute window around 2:30 p.m. on the previous afternoon. We use the variance-covariance matrix in that window to predict what the variance-covariance matrix would have been in the event window in the absence of the data release.

The advantage of using the intraday quotes is shown in figure 8.1, which focuses on the response of the two-year Treasury yield to the nonfarm payrolls statistic from the monthly *Employment Situation* report from the BLS. This is the data release that, in recent years at least, has commanded the most attention in financial markets. As can be seen, there is a clear positive relationship between surprises in the payroll release and the movement in the two-year yield. Moreover, this relationship tightens if we use intraday data instead of daily data.

We investigate the market reactions to thirteen different data releases. Those releases are shown in table 8.1, along with some information about the frequency of the release and the sample over which we have a measure of market expectations. We generally begin our sample in 1994, though the sample for the Chicago Purchasing Manufacturers Index (PMI) has a shorter sample because we do not have a measure of market expectations

8. We had hoped to include exchange rates as well, but our intraday data did not extend back far enough to make it a useful sample.

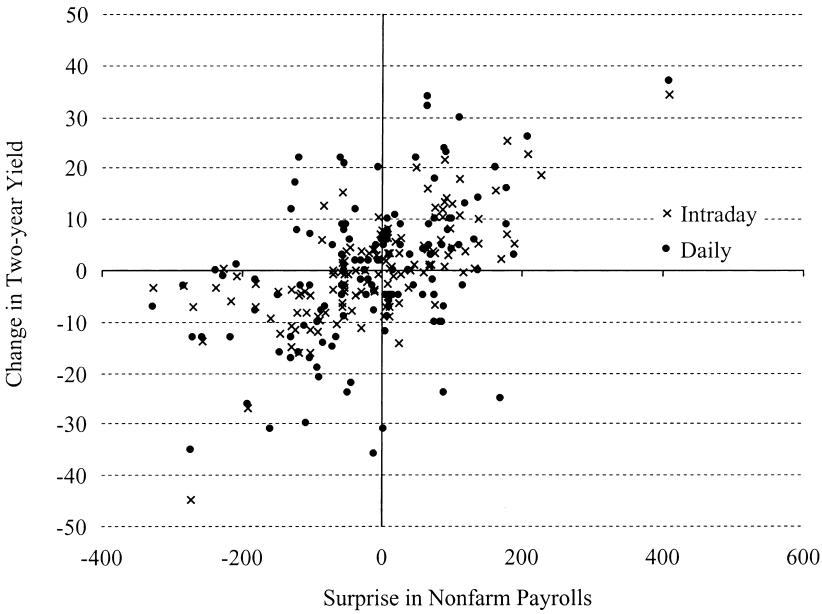

Fig. 8.1 Response of the two-year yield to payroll surprises

Table 8.1 Macroeconomic data announcements

Release	Release time	Frequency	Date of first observation	No. of observations
Nonfarm payrolls	8:30	Monthly	7-Jan-94	137
Hourly earnings	8:30	Monthly	4-Feb-94	134
GDP (advance)	8:30	Quarterly	28-Jan-94	46
Retail sales (excl. autos)	8:30	Monthly	13-Jan-94	137
Core Consumer Price Index	8:30	Monthly	13-Jan-94	137
Core Producer Price Index	8:30	Monthly	12-Jan-94	137
Housing starts	8:30	Monthly	20-Jan-94	135
Durable goods	8:30	Monthly	27-Jan-94	135
Capacity utilization	9:15	Monthly	14-Jan-94	137
Institute for Supply Management Manufacturing Index	10:00	Monthly	3-Jan-94	133
Chicago Purchasing Manufacturers Index	10:00	Monthly	31-Dec-99	58
Consumer confidence	10:00	Monthly	25-Jan-94	137
New home sales	10:00	Monthly	2-Feb-94	136

until December 1999. Our list includes nearly all of the major macroeconomic indicators that are generally seen as significant market movers.[9]

8.4.2 Evenstudy Estimates

Even though it may have the shortcomings discussed in the preceding, we still view the standard eventstudy regression as a very useful exercise, one that can tell us a lot about how asset prices and monetary policy expectations are affected by incoming data. The preceding discussion simply cautions that the resulting coefficients may have some downward bias, thus understating the importance of the data. We implement the eventstudy regression per release, using the data described in the preceding. The results are shown in table 8.2.

One of the primary findings from this exercise is that monetary policy expectations react significantly to incoming data. The expected path of the federal funds rate (as captured in eurodollar futures) generally shifts up significantly in response to both strong data on growth (such as retail sales) and high data on inflation (such as core CPI). Overall, we find that twelve of the thirteen macroeconomic variables considered prompt a significant reaction in the eurodollar futures rates.[10]

A second finding is that the effect of the data releases continues to be sizable even as the maturity of the instrument is extended. Indeed, the two-year yield often moves by about the same amount as the eurodollar futures rates, suggesting that any influence on monetary policy is seen as being very persistent.[11] The sensitivity of market yields extends all the way out to the ten-year Treasury note. The magnitude of its reaction is large enough that it suggests that even distant forward rates are reacting to the news, as found by Gürkaynak, Sack, and Swanson (2005).[12]

A final observation from the eventstudy results has to do with the reaction of equity markets. The detachment issue seems particularly problematic for equities, as even the most important data releases (such as nonfarm

9. In all of the results that follow, we discard those days for which we have multiple releases. For the two series from the employment report (nonfarm payrolls and hourly earnings), we always consider their effects together, as discussed in the preceding.

10. Other studies, including Fleming and Remolona (1997), Balduzzi, Elton, and Green (2001), and Gürkaynak, Sack, and Swanson (2005), have found that market interest rates respond significantly to a wide range of macroeconomic data releases.

11. A similar result was found by Kohn and Sack (2003). They noted a similar persistence in the response to Federal Open Market Committee (FOMC) statements and inferred that those statements may be seen as conveying information about the state of the economy in addition to information about the near-term direction of policy.

12. That paper looked explicitly at distant forward rates and found that they often responded to data in the same direction as near-term forward rates. The authors developed the case that this response reflected the fact that long-term inflation expectations in the United States are variable, a case strengthened by the fact that similar sensitivity is not observed in the United Kingdom, perhaps because of its explicit inflation target.

Table 8.2 **Effects of macroeconomic data surprises on asset prices: eventstudy approach**

	ED2	ED4	Y2	Y10	S&P	Percentage of Y2 explained by data surprise
Nonfarm payrolls	5.74**	8.16**	6.67**	5.34**	–0.03	0.46
	(0.57)	(0.78)	(0.66)	(0.57)	(0.07)	
Hourly earnings	1.72**	2.03**	2.01**	1.84**	–0.23**	—
	(0.79)	(1.11)	(0.92)	(0.76)	(0.07)	
GDP (advance)	1.66**	2.28**	1.84**	1.25**	0.15	0.25
	(0.73)	(0.92)	(0.77)	(0.71)	(0.09)	
Retail sales (excl. autos)	2.27**	3.20**	2.05**	1.91**	0.11	0.23
	(0.57)	(0.78)	(0.55)	(0.43)	(0.07)	
Core Consumer Price Index	1.89**	2.62**	1.96**	1.97**	–0.21**	0.28
	(0.37)	(0.50)	(0.41)	(0.44)	(0.09)	
Core Producer Price Index	1.28**	1.78**	1.39**	1.42**	–0.20**	0.17
	(0.33)	(0.46)	(0.40)	(0.40)	(0.08)	
Housing starts	0.33	0.32	0.23	0.28	–0.07	0.01
	(0.21)	(0.27)	(0.24)	(0.19)	(0.07)	
Durable goods	1.02**	1.37**	1.34**	1.16**	–0.04	0.25
	(0.21)	(0.33)	(0.26)	(0.24)	(0.05)	
Capacity utilization	0.91**	1.36**	1.03**	1.01**	0.05	0.19
	(0.20)	(0.26)	(0.19)	(0.17)	(0.06)	
ISM Manufacturing Index	1.78**	2.56**	2.07**	2.04**	–0.05	0.31
	(0.25)	(0.41)	(0.33)	(0.34)	(0.04)	
Chicago Purchasing Manufacturers Index	1.30**	1.99**	1.55**	1.55**	0.08	0.30
	(0.31)	(0.48)	(0.43)	(0.36)	(0.06)	
Consumer confidence	1.80**	2.15**	1.74**	1.70**	0.08	0.26
	(0.33)	(0.44)	(0.35)	(0.29)	(0.04)	
New home sales	1.02**	1.25**	1.01**	0.85**	–0.05	0.16
	(0.25)	(0.31)	(0.25)	(0.24)	(0.04)	

Notes: The table shows the estimated response of the financial variable (in basis points for rates and percentage points for equities) to a 1 standard deviation surprise in the economic release. ED2 = the rate on the second eurodollar futures contract (a proxy for monetary policy expectations about six months ahead); ED4 = the rate on the fourth eurodollar futures contract (a proxy for policy expectations about twelve months ahead); Y2 = the two-year Treasury yield; Y10 = the ten-year Treasury yield; and S&P = the S&P 500 index. The last column reports the R-squared statistic for the Y2 regression. No statistic is reported for hourly earnings because it is estimated in the same regression as nonfarm payrolls.

**Significant at the 5 percent level.

payrolls) do not prompt a significant market reaction.[13] But looking at the response of equities to all of the releases provides us with an important clue about why that may be the case.

The likely explanation for this finding is that a release such as nonfarm payrolls contains offsetting forces on equity prices. On the one hand, a strong report would suggest more strength in the economy and, hence, bet-

13. By contrast, equities do appear to react significantly to monetary policy shocks, as shown by Bernanke and Kuttner (2003).

ter earnings prospects, which should boost equity prices. On the other hand, it also raises long-term interest rates, which should lower equity prices. These two forces offset one another, leaving the net effect on equity prices insignificantly different from 0. A similar story could be told for all of the demand-side indicators, which all have no effect on equities.

If this were in fact the case, then we should more clearly see a negative response of equity prices to data that is directly about inflation. The reason is that there is no offsetting news in that case—higher inflation implies that rates will be higher but not that growth will be higher. Thus, equity prices should fall. Indeed, this is precisely what we find. Indeed, the S&P index reacts negatively and significantly to positive surprises in core CPI, core Producer Price Index (PPI), and hourly earnings—every single inflation measure considered.[14]

Overall, the eventstudy regressions provide an interesting pattern of market responses to different types of incoming news. Nevertheless, the *R*-squared statistics from the regressions are relatively low, generally ranging from 0.15 to 0.50. In other words, the eventstudy regressions typically account for only a small portion of the variance of the market reactions, even if we focus on the movements in the thirty-minute window bracketing the announcement. This last observation is the area in which we will see some improvement under the new estimator.

8.4.3 Identification-though-Censoring Estimates

Table 8.3 shows the estimated responses under the identification-by-censoring (IC) approach. Broadly speaking, the patterns of the responses are the same as in the eventstudy (ES) exercise: stronger-than-expected readings on growth or higher-than-expected readings on inflation tend to boost market interest rates. The stock market response to incoming data on growth is mixed and often insignificant, while it reacts negatively to incoming data on inflation.

The primary difference between the ES and IC approaches is the magnitude of the market responses. The IC coefficients are often two or three times as large as the ES coefficients. This finding suggests that the problem of detachment is, to a large extent at least, associated with the mismeasurement of macroeconomic news.

For example, a 1 standard deviation upward surprise to core CPI (nearly 0.1 percentage point) is estimated to increase yields 6 to 9 basis points, rather than the response of 2 to 2.5 basis points found under ES. It is worth considering again how to interpret this difference. The IC measure is capturing the market response to a true core CPI surprise, one that market par-

14. To our knowledge, this is not an empirical fact that has been emphasized in the literature to date. Fair (2003) finds a positive reaction of equities to inflation news. McQueen and Roley (1993) find a reaction that differs across different states of the business cycle, with negative responses for some variables in some states.

Table 8.3 **Effects of macroeconomic data surprises on asset prices: identification-through-censoring approach**

	ED2	ED4	Y2	Y10	S&P	Percentage of survey-based surprise due to noise
Nonfarm payrolls	8.65	7.65	10.33	9.62	0.09	0.31
	(0.20)	(0.11)	(0.24)	(0.07)	(0.00)	
Hourly earnings	10.52	4.78	7.57	12.71	–1.16	0.70
	(0.06)	(0.11)	(0.10)	(0.29)	(0.00)	
GDP (advance)	5.71	7.39	5.95	5.15	0.02	0.42
	(0.28)	(0.30)	(0.28)	(0.25)	(0.00)	
Retail sales (excl. autos)	6.00	8.19	5.49	4.25	0.02	0.70
	(0.17)	(0.18)	(0.18)	(0.39)	(0.00)	
Core Consumer Price Index	6.43	8.87	6.61	7.59	–0.94	0.77
	(0.72)	(0.64)	(0.66)	(1.16)	(0.05)	
Core Producer Price Index	4.81	6.30	5.33	5.40	–0.68	0.31
	(1.21)	(1.45)	(1.09)	(1.09)	(0.09)	
Housing starts	1.15	1.08	0.95	0.24	–0.01	0.00
	(0.06)	(0.07)	(0.07)	(1.88)	(0.00)	
Durable goods	1.79	2.76	2.35	1.86	–0.01	0.24
	(0.17)	(0.12)	(0.15)	(0.22)	(0.05)	
Capacity utilization	8.63	11.21	7.99	7.06	1.42	0.92
	(1.57)	(0.81)	(0.93)	(0.81)	(0.00)	
ISM Manufacturing Index	10.92	17.13	13.94	14.06	–1.11	0.94
	(0.62)	(0.50)	(0.57)	(0.59)	(0.01)	
Chicago Purchasing	2.42	3.57	3.12	2.94	0.16	0.50
Manufacturers Index	(2.43)	(1.74)	(8.76)	(2.25)	(0.07)	
Consumer confidence	9.69	12.40	9.67	8.58	0.80	0.90
	(0.72)	(0.75)	(0.77)	(0.69)	(0.03)	
New home sales	8.63	8.19	9.12	8.64	–0.88	0.95
	(1.57)	(1.20)	(1.56)	(1.98)	(0.09)	

Notes: The table shows the estimated response of the financial variable (in basis points for rates and percentage points for equities) to a 1 standard deviation surprise in the "true" economic release (that measured without noise). ED2 = the rate on the second eurodollar futures contract (a proxy for monetary policy expectations about six months ahead); ED4 = the rate on the fourth eurodollar futures contract (a proxy for policy expectations about twelve months ahead); Y2 = the two-year Treasury yield; Y10 = the ten-year Treasury yield; and S&P = the S&P 500 index. The last column reports the fraction of the variation in the survey-based surprise measure that is estimated to be noise.

ticipants are convinced has no measurement error in it and one for which the market expectations are measured perfectly. The true CPI release may be discounted if it is seen as containing measurement error (e.g., a higher-than-expected reading driven by a single component, such as the price of lodging away from home), or its estimated effect under the ES may be downward biased if the market's expectations were measured improperly.

One implication of the results is that monetary policy expectations and asset prices may be more systematically related to incoming data than found under the ES approach. This conclusion accords with our understanding of monetary policy from the (lower-frequency) macroeconomic literature, including the view that one way policy has been effective over the past decade is by systematically responding to changes in economic conditions. Our results provide a high-frequency version of that conclusion.

One issue is that the results appear "too good" in some sense. The estimated amount of noise in the data announcements, a statistic that is also identified in the IC procedure, tends to be very high for many of the releases. (This pattern, of course, is directly related to the fact that the IC coefficients are several times larger than the ES coefficients.) For example, the results suggest that 31 percent of the variation in the nonfarm payrolls surprise is due to noise, while 77 percent of the variation in the core CPI release is due to noise.

It is somewhat hard to grasp just how much noise one would expect relative to some actual "truth" that we never observe. However, some of the readings from table 8.2 are clearly implausible. For example, we doubt that 94 percent of the measured surprise associated with the Institute for Supply Management (ISM) index is actually noise.

The extent of the estimated noise may raise some questions about whether the identification assumptions hold. We might be particularly concerned about our efforts to predict what the variance of the asset prices would have been in the absence of the macroeconomic surprise, as needed in the IC procedure. Note that the estimates of both the sensitivity of the market response (γ) and the amount of noise in the surprises (σ_ε) tend to increase in the shift in the variance of the asset price between nonannouncement days and announcement days. Hence, if we are underestimating the variance that would be present in the absence of a macroeconomic announcement, we would be overestimating both of these parameters.

One reason to suspect this pattern is that the macroeconomic surprises measured on the right-hand side of our equations often coincide with the release of other data that might move markets.[15] For example, the employment report not only includes the current month surprise to nonfarm payrolls, but also revisions to payrolls in the two previous months. Thus, even in the absence of a surprise regarding the current month payroll, one might expect more market volatility than on a nonannouncement day because of the possible market reaction to this other information.

If this were the case, the IC estimates presented in table 8.3 may have some upward bias. But note that this upward bias exists because the data release is actually more meaningful than captured by the surprise measure on the right-hand side of the equation. Thus, it still likely reflects that the

15. In addition, the announcement itself (even if it is on expectation) could result in some variance of the asset price because it would presumably reduce uncertainty and cause investors who had different expectations to adjust their positions and views.

macroeconomic news is more important than accounted for by the event-study approach. We might, therefore, want to think of an estimator that can better incorporate that additional information.

8.5 A Principal Components Exercise

This last consideration leads us in the direction of a completely different but complementary approach. The IC estimator was developed out of concern that the macroeconomic surprise variable may be measured poorly, introducing too much variation into that measure. But perhaps the bigger problem is the opposite one—that the right-hand-side variable does not capture *enough* of the surprise in a given data release.

This would be the case if the data release contained market-moving information other than that represented in the surprise measures considered in the preceding. To be sure, most data releases are complicated and convey many pieces of information. It may be difficult to determine a macroeconomic surprise measure that captures all of that information.[16]

An alternative approach that avoids this difficulty is to let the financial market data itself determine the data surprise. Specifically, we again consider the movements of the four interest rates and equity prices in the thirty-minute window around a given release. Our identification assumption is that the primary event driving the markets in those windows is the data release—an assumption that is certainly plausible for the narrow window that we consider around the release. We are not ruling out that other events take place in that window, but if there does appear to be one common event, we will assign its effects to the data release.

The approach that we use to implement this assumption is principal components. For a given release, we stack the market reactions into a matrix with one row per observation and one column for each asset price (the second and fourth eurodollar contracts, the two-year Treasury yield, the ten-year Treasury yield, and the S&P index). The principal components exercise determines a set of orthogonal factors F (same dimensions as X) that are linear combinations of the original data series:

$$F = X \cdot A, \tag{14}$$

where $A'A = I$. As a result, the variance-covariance matrix of the responses of the financial variables is given by $F'F = A' \cdot \mathbf{\Sigma} \cdot A$, where $\mathbf{\Sigma}$ is a diagonal matrix containing the variances of the factors.

The factors are ordered by the magnitude of their variances (with the factor with the highest variance listed first). In this sense, the first factor ex-

16. In the preceding, we have the example of the payrolls release and the relevance of concurrent revisions to past months' payrolls releases. Other examples are quite that retail sales ex-autos coincides with total retail sales (including autos), capacity utilization coincides with industrial production, and so on.

plains as much of the variation across the observable variables as possible, the second factor captures as much additional variation as possible, and so on.[17] The loadings of the financial variables on each of the factors is given by the inverse of the matrix **A**, or A'.

This approach is more general than the IC estimator. It does not require the two identifying assumptions needed in that case, and it can capture a broader set of information than measured by the surprise variables included in the IC and ES approaches. The potential cost, however, is that it could accidentally include some variation not truly associated with the data release. A finding that there is a strong co-movement in the asset prices over the thirty-minute window around the data release would boost our confidence that the procedure is picking up the effects of that release.

As reported in table 8.4, it turns out that a single factor explains the vast majority of the market reaction to each release. This factor typically accounts for 90 percent to 95 percent of the variation in the asset prices in the thirty-minute window.[18] It is this movement that we associate with the data release, as the release is presumably the dominant market event in the window.[19]

In this case, the first principal component provides a measure of the true data surprise, one that incorporates all of the market-sensitive news included in a given release. As we would expect, these data surprises are somewhat correlated with the survey-based surprises used in the preceding. Table 8.4 shows that the survey-based surprises account for as much as 50 percent of the variation in the first principal component. Thus, clearly the surprises used in the ES exercise are an important component of the total news around a data release. However, they are not a complete measure of the market-sensitive news contained in the release, as suggested by the additional (unexplained) variation in the first principal component.

Figure 8.2 presents the example for ex-auto retail sales. On the horizontal axis is the survey-based surprise used in the preceding, and on the vertical axis is the first principal component (normalized in a way to make it most comparable to the retail sales release). Again, the two measures are clearly related, but they are far from identical. The principal component (PC)-based surprise measure has more variation than can be explained by

17. When we apply this technique to the above data set, we normalize each variable by its standard deviation.

18. The table shows the variance of the first factor relative to all of the other factors. But that statistic is nearly identical to the fraction of the variance of the market interest rates explained by the first factor.

19. For comparison, if we conduct the same exercise in the nonevent window considered in the preceding (the thirty-minute window bracketing 2:30), we find that the first factor explains only 80 percent of the variance of the asset prices. Thus, it does appear that the data release window contains an even event that causes a comovement in the asset prices that is larger than that observed at other times.

Table 8.4 **Effects of macroeconomic data surprises on asset prices: principal components approach**

	Factor loadings					Variance explained by first factor	Amount explained by survey-based data surprise
	ED2	ED4	Y2	Y10	S&P		
Employment report	8.4	11.9	9.9	8.1	–0.15	0.98	0.55
GDP (advance)	3.8	5.6	4.8	4.8	–0.14	0.95	0.13
Retail sales (excl. autos)	4.0	5.4	4.4	3.7	–0.04	0.96	0.16
Core Consumer Price Index	3.1	4.2	3.6	3.4	–0.23	0.94	0.21
Core Producer Price Index	3.1	4.2	3.6	3.3	–0.14	0.95	0.14
Housing starts	2.1	2.9	2.2	1.9	–0.03	0.91	0.02
Durable goods	2.6	3.7	2.8	2.6	–0.05	0.94	0.19
Capacity utilization	2.3	3.1	2.2	2.0	0.05	0.91	0.20
ISM Manufacturing Index	3.3	5.1	4.1	4.0	–0.02	0.96	0.42
Chicago Purchasing Manufacturers Index	2.1	3.6	2.6	2.3	0.22	0.92	0.34
Consumer confidence	2.9	3.9	3.1	2.8	0.12	0.96	0.24
New home sales	2.3	3.1	2.6	2.4	0.02	0.94	0.15

Notes: The table shows the responses of the financial variable (in basis points for rates and percentage points for equities) to a 1 standard deviation surprise in the first principal component. ED2 = the rate on the second eurodollar futures contract (a proxy for policy monetary expectations about six months ahead); ED4 = the rate on the fourth eurodollar futures contract (a proxy for policy monetary expectations about twelve months ahead); Y2 = the two-year Treasury yield; Y10 = the ten-year Treasury yield; and S&P = the S&P 500 index. The last column reports the R-squared statistic from a regression of the first factor on the particular survey-based data surprise (two surprises in the case of the employment report).

the survey-based surprise measure, presumably capturing the additional information in the release.

Table 8.4 also reports the loadings of the various asset prices on the PC-based surprise measure. For ease of interpretation, we have normalized each PC measure to have a unitary standard deviation, just as we did with the survey-based surprises used in the ES exercise. The coefficients retain many of the interesting patterns observed in the earlier results. The market interest rates considered have a sizable response to the macroeconomic news, suggesting that the news is affecting the expected path of monetary policy. Those responses are typically also observed at longer-term maturities.

One puzzling aspect of the results is that the equity market no longer appears to have as large of a negative reaction to incoming data on inflation. It is true that the first factor explains a larger fraction of equity price movements for the inflation-related data releases than for other releases, but the response relative to the interest-rate response is smaller than in the preceding results. Instead, the factor analysis essentially finds a separate factor that drives much of the movements in equity prices. We wonder whether

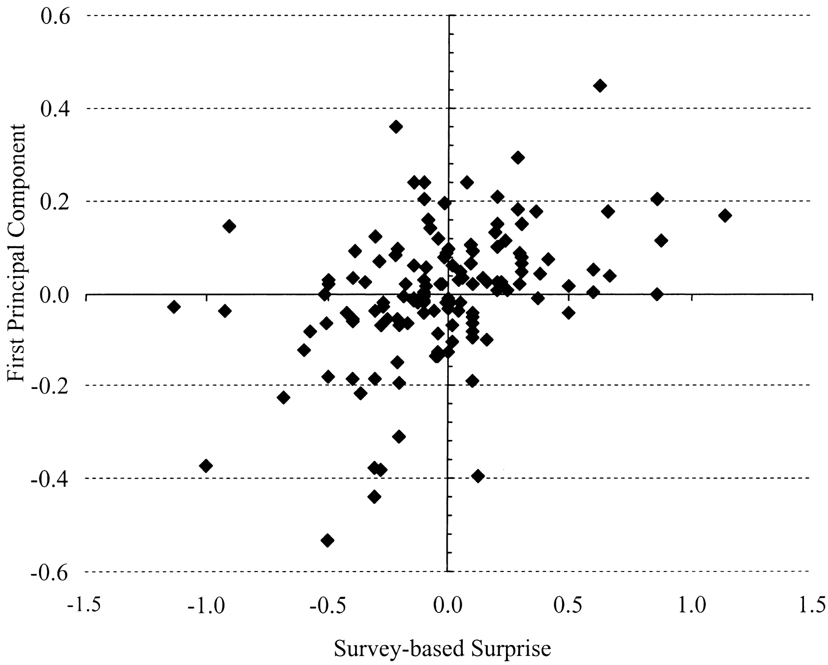

Fig. 8.2 Surprise measures for retail sales: First principal component versus survey-based measure

this finding in part reflects that we are forced to use a wider window for the equity price movements (seventeen and a half hours instead of thirty minutes!), which considerably weakens the identification assumption used in the PC exercise.

Perhaps the most important aspect of the PC exercise is its usefulness for assessing the amount of variation in yields that can be attributed to macroeconomic data. The PC exercise indicates that markets are much more sensitive to macroeconomic data releases than suggested by the ES approach. This is a similar finding as the IC estimator used in the preceding. However, in this case, the reason is not only that we are accounting for the measurement error in the survey-based surprise measure, but also because we are accounting for any other relevant information in the release.

One useful aspect of the PC approach is that, unlike the case for the IC estimator, we recover a time series of the true macroeconomic news, as discussed in the preceding. This allows us to cumulate the effects of each release on a particular asset price. Figure 8.3 shows the cumulative effects of each the data releases on the two-year Treasury yield, where each line represents an individual release. (For example, one line represents the effects of all retail sales releases over our sample.) The point of figure 8.3 is not to

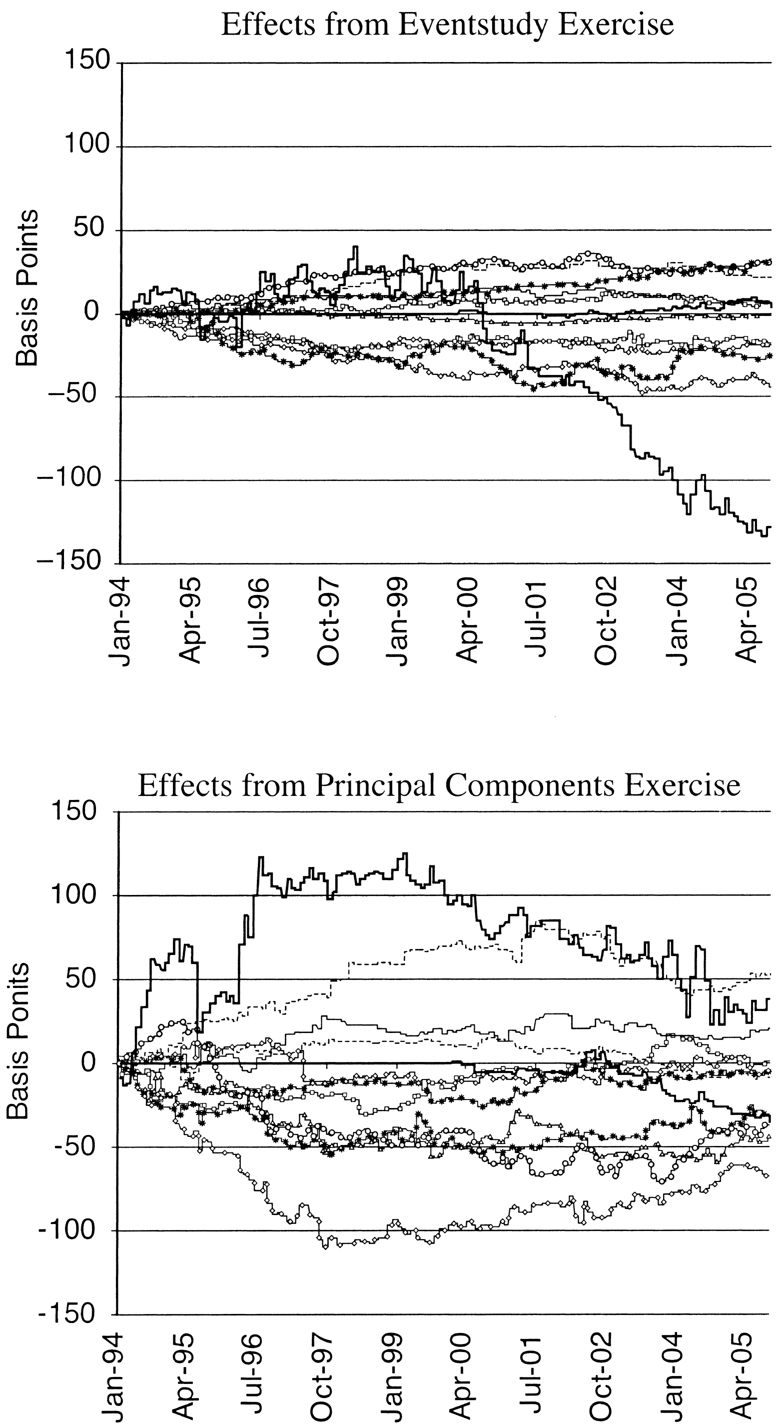

Fig. 8.3 Cumulative effects of individual data releases on two-year yield (one line per release, in basis points)

focus on any particular line, but to get a general sense of the total variation explained under the two approaches. As can be seen, the movements explained by the releases under the ES exercise are much smaller than those under the PC exercise.

Table 8.5 contains some statistics that further quantify the variation explained under the two approaches. It computes the absolute value of the changes attributable to each release, expressed as basis points per year. By this measure, the most influential data release, by far, has been the employment report. Other influential releases include retail sales, the ISM index, the CPI, and the PPI.

More important, the PC measure accounts for much more variation than the standard eventstudy approach. (This, of course, is simply a different way of expressing that the R-squared statistic from the regression increases significantly.) Indeed, this is the case for every single data release considered. We can sum these statistics across all of the releases to obtain a measure of the total variation explained by incoming macroeconomic data (or at least by our releases). By that measure, the PC approach has accounted for nearly twice as much of the variation in the two-year yield than the ES approach. Thus, the new methodology makes an important step toward better understanding the total influence of macroeconomic data on asset prices and monetary policy expectations.

Table 8.5 **Variation in the two-year treasury yield explained by macroeconomic surprises**

	Sum of absolute changes per year	
Release	Eventstudy approach	Principal component approach
Employment report	66	86
GDP (advance)	6	14
Retail sales (excl. autos)	15	36
Core Consumer Price Index	16	30
Core Producer Price Index	10	31
Housing starts	4	19
Durable goods	12	21
Capacity utilization	10	18
ISM Manufacturing Index	24	34
Chicago Purchasing Manufacturers Index	13	25
Consumer confidence	13	23
New home sales	10	20
Total	199	357

Notes: The table reports the sum of the absolute value of changes in the two-year yield attributable to the economic release under the two approaches. These changes are then summed over the sample for each variable and scaled by the number of releases per year divided by the total number of releases in the sample.

8.6 Implications and Conclusions

We have learned a lot from the standard eventstudy literature. This chapter begins with that approach, implementing it with the benefit of using intraday data and looking across a variety of asset prices. The eventstudy exercise clearly establishes a set of facts that macroeconomists should strive to explain when writing down models of the interactions of macroeconomic developments, monetary policy, and asset prices.

There are three broad observations that derive from the eventstudy exercise. First, policy expectations systematically respond to incoming data, with evidence of stronger-than-expected growth or higher-than-expected inflation leading to an upward revision to the expected federal funds rate path (as reflected in eurodollar futures rates and the two-year Treasury yield). Second, the influence of that data on the yield curve extends to very long maturities (the ten-year yield in our exercise). And third, equities show very mixed reactions to incoming data on growth but negative and significant reactions to data on inflation.

Many of these patterns align well with current macroeconomic models. Those models typically assume that monetary policy is systematically related to incoming data—a relationship that should also be apparent in the high-frequency data. The responsiveness of longer-term Treasury yields is somewhat more challenging to explain in current models, in part due to the difficulty associated with understanding the determination of long-horizon expectations, but it, too, has been taken up in the recent literature. Last, as discussed in the preceding, the lack of response of equities to demand-side indicators could reflect that those releases affect both expected dividend growth and interest rates, with offsetting effects on stock prices.

Nevertheless, the eventstudy estimates leave one significant shortcoming in our understanding of market dynamics—that the measured data surprises explain only a small portion of the variation in asset prices and monetary policy expectations. This chapter argues that this shortcoming likely reflects mismeasurement of the macroeconomic news.

We developed two new approaches to better account for the influence of the macroeconomic news under the assumption that the measured surprises are noisy. The first is a new econometric technique for accounting for error-in-variables, one that has the potential to be used in other applications as well. The second is a principal components approach that takes advantage of our ability (using intraday data) to zero in on the asset price movements right around the release.

The new estimators do not significantly change the patterns of the market responses that we obtain from the standard eventstudy approach. That is, the patterns found under the ES approach (including the three observations noted in the preceding) still represent a set of observations that should be explained by macroeconomic models. However, the two new approaches suggest that incoming news generally has a much bigger impact on asset prices than captured by the eventstudy approach.

In the case of the IC estimator, the results suggest that the noise in the measure of the data surprise causes a downward bias in the measured sensitivity of asset prices to that information. The PC estimator also suggests that this may be the case, and it also allows for the possibility that there is other market-sensitive news in the data release beyond the macroeconomic surprise included in the eventstudy regression.

In sum, we argue that the sensitivity of asset prices and monetary policy expectations to high-frequency information on macroeconomic conditions is likely to be greater than captured in previous studies. This finding accords well with the view that monetary policy systematically responds to economic conditions, and that asset prices more broadly are strongly influenced by the evolution of the economy and policy expectations.

References

Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega. 2003. Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review* 93:38–62.

Balduzzi, Pierluigi, Edwin J. Elton, and T. Clifton Green. 2001. Economic news and bond prices: Evidence from the U.S. Treasury market. *Journal of Financial and Quantitative Analysis* 36:523–43.

Bernanke, Ben, and Kenneth N. Kuttner. 2003. What explains the stock market's reaction to Federal Reserve policy? *Journal of Finance* 60:1221–57.

Cutler, David M., Porterba, and Lawrence H. Summers 1989. "What Moves Stock Prices?" *Journal of Portfolio Management* 15:4–12.

Fair, Ray C. 2003. Shock effects on stocks, bonds, and exchange rates. *Journal of International Money and Finance* 22:307–41.

Faust, John, John H. Rogers, Shing-Yi B. Wang, and Jonathan Wright. 2003. The high frequency response of exchange rates and interest rates to macroeconomic announcements. International Finance Discussion Papers no. 2003-784. Washington, DC: Board of Governors of the Federal Reserve System.

Fleming, Michael, and Eli Remolona. 1997. What moves the bond market? *Federal Reserve Bank of New York Economic Policy Review* 3:31–50.

Goldberger, Arthur S. 1991. *A course in econometrics.* Cambridge, MA: Harvard University Press.

Gürkaynak, Refet, Brian Sack, and Eric Swanson. 2005. The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American Economic Review* 95:425–36.

Kohn, Donald L., and Brian Sack. 2003. Central bank talk: Does it matter and why? In *Macroeconomics, monetary policy, and financial stability,* 175–206. Ottawa, Ontario: Bank of Canada.

McQueen, Grant, and V. Vance Roley. 1993. Stock prices, news, and business conditions. *The Review of Financial Studies* 6:683–707.

Rigobon, Roberto, and Brian Sack 2004. The impact of monetary policy on asset prices. *Journal of Monetary Economics* 51:1553–75.

Stock, James H., and Mark W. Watson. 2003. Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature* 41:788–829.

Comment Leonardo Bartolini

Rigobon and Sack's nice chapter rightly points its finger to serious problems with our often-too-eager use of survey-based expectations data and offers a solution for some of these problems. Survey data are likely to capture investors' expectations with large errors. Hence, their use to construct measures of macroeconomic news induces a classic attenuation bias that may explain the weak response of asset prices to news documented in recent studies. For instance, Faust et al. (2003) estimate the ten-year U.S. interest rate to rise by 13 basis points in response to a 1 percent unexpected tightening in the Fed Funds rate—a small and just statistically significant effect. Surely, many scholars of financial markets would maintain that most macroeconomic announcements should *not* be expected to have a large impact on asset prices and would not be surprised by the small size of estimates such as those of Faust et al. (2003). Even so, the debate on whether estimated weak responses of prices to news are puzzling would surely benefit by improved measurement of the news content of macroeconomic announcements, which is Rigobon and Sack's goal in this chapter. For this reason, this chapter makes a useful methodological contribution, also offering estimates of the response of asset prices to news that improve on received wisdom.

Despite such improvement, there are reasons suggesting that the response of asset prices to news estimated by Rigobon and Sack—up to ten times stronger than previously estimated—might overstate the true response of asset prices to macroeconomic announcements, possibly by a substantial amount. To assess this view, let me start by agreeing wholeheartedly with the chapter's key premise: survey-based measures of expectations are marred by such deep problems that special care should be taken to develop methods to address their deficiencies. Among such problems, notable is the fact that such surveys are typically conducted well in advance of data announcements, with leads ranging from a few days to a couple of weeks. Such leads imply that by the time of the announcement, much information on the released series has accumulated, making much of the "measured" news just not news any longer. This sampling-lead problem does not, on its own account, invalidate the rationality of the initial forecast. If forecasters form their expectations rationally, "measured news"

Leonardo Bartolini is senior vice president of the International Research Function at the Federal Reserve Bank of New York.

differs from "true news"—that is, news relative to information available one instant before the announcement—by noise only. Such noise compounds with more traditional measurement errors and can be dealt with by the clever variant of standard error-in-variables estimation methods proposed by Rigobon and Sack.

Sampling leads and measurement errors are not the only problem with forecast data, however. Other problems with such data may lead to biased (conditionally and unconditionally) forecasts that are not readily dealt with by the "censored" estimation method suggested by the authors. For instance, individual forecasts may be biased because of the unusual incentives faced by forecasters to maximize their public recognition. Because the names and affiliation of forecasters participating in the Bloomberg survey used by Rigobon and Sack and in other common surveys are publicly listed, forecasters may have an incentive to forecast in the tail of the forecast distribution, as this increases their chances of being the forecaster that comes closest to the actual release (Laster, Bennett, and Geoum 1999; Lamont, 2002). Some forecasters have also been found to distort forecasts toward realizations that benefit their firm, in a way that Ito (1990) dubs "wishful expectations." Depending on how advanced signals on the eventual data release are distributed among forecasters, a bias in individual forecasts may cause a bias in the consensus (that is, median or mean) forecast as well. Finally, forecast data are often found not to reflect efficiently publicly available information. This feature is hardly surprising, as forecasters have weak monetary incentives to offer best forecasts: the forecasters themselves are seldom investment managers, that is, agents with an incentive to put their money where their mouths are. For instance, forecasters have been documented not to learn quickly from mistakes, leading to serially correlated forecast errors (see, for instance, Mankiw Reis, and Wolfers [2003] and Gürkaynak and Wolfers [2005]). None of these problems is attacked by the methodology of Rigobon and Sack, thus leaving scope for overestimating the response of asset prices to news.

To illustrate this point, let's focus on Rigobon and Sack's neatest contribution—"identification through censoring"—which involves the following: first, to view the estimation of the news content of macroeconomic announcements as a classic error-in-variables problem. Second, to cast such problem as a matter of underidentification. And, finally, to identify the model by adding suitable restrictions to the covariance matrix. In fairness to previous research, I should note that viewing error-in-variables as an issue of underidentification has a time-honored tradition, and equally honored is the tradition of adding restrictions to the covariance matrix to achieve identification.[1] The key novelty in Rigobon and Sack's work is to recognize that the variance of the measured news process and of its noise

1. See, for instance, Judge et al. (1982), chapter 19.

component must be zero at all times other than at the time of the announcement. This observation yields a natural restriction to impose on the covariance matrix to achieve full identification.

Take then the authors' basic model of error-in-forecast-measurement:

$$\Delta s_t = \gamma \cdot z_t^* + \varepsilon_t \tag{1}$$

$$z_t = z_t^* + \eta_t, \tag{2}$$

where z_t is "measured news" (that is, data – data expected at $t - 1$) and z_t^* is "true news" (that is, data – data expected at $t - \mathrm{d}t$). We are interested in the coefficient γ, for which ordinary least squares (OLS) provides a biased and inconsistent estimator, as in a classic error-in-variables problem. There is no instrument correlated with true news that can be used to address this problem: if there were one such instrument, z_t^* would not be news anymore.

Let us then follow the authors and obtain the additional condition to identify the model by recognizing that the variance of the asset price on nonannouncement days should equal its "structural" variance σ_ε^2, defined as the variance of the asset when $z_t^* = 0$. (This is a critical assumption, which is further discussed in the following.) However, let us drop the assumption that the measurement error has zero mean conditional on the true surprise. That is, let $E[\eta_t \mid z_t^*] \neq 0$, a relaxation that puts us outside the classic error-in-variables model. Consistent with much previous evidence that news derived from survey forecasts is serially correlated, suppose that measured news follows the partial adjustment model

$$z_t = (1 - \rho)z_{t-1}^* + \rho z_{t-1} + \eta_t. \tag{3}$$

In this case, one obtains

$$\operatorname{var}(\Delta s_{t-1}) = \sigma_\varepsilon^2, \tag{4}$$

$$\operatorname{var}(\Delta s_t) = \gamma^2 \sigma_{z^*}^2 + \sigma_\varepsilon^2, \tag{5}$$

$$\operatorname{cov}(\Delta s_t, z_t) = \gamma(1 - \rho)\sigma_{z^*}^2, \tag{6}$$

where equations (4) and (5) are unchanged relative to Rigobon and Sack's analysis, while equation (6) accounts for the sluggish response of measured news. In this case, the estimator γ becomes

$$\gamma = (1 - \rho)\,\frac{\operatorname{var}(\Delta s_t) - \operatorname{var}(\Delta s_{t-1})}{\operatorname{cov}(\Delta s_t, z_t)}, \tag{7}$$

which differs from Rigobon and Sack's censored estimator $\gamma = [\operatorname{var}(\Delta s_t) - \operatorname{var}(\Delta s_{t-1})]/[\operatorname{cov}(\Delta s_t, z_t)]$ by the term $(1 - \rho)$. Therefore, when news measured from survey data adjusts gradually to true news, Rigobon and Sack's estimator is biased upward.

Is this bias quantitatively significant? Mankiw, Reis, and Wolfers (2003) estimate that about half of the error in inflation expectations from four major surveys remains in the median forecast after one year. So the bias may be significant for some surveys. It may or may not be so for Rigobon and Sack's data, depending on the time series properties of their forecast errors.

Another key feature of forecast data that does not fit neatly into the error-in-variables framework, and which may contribute to an upward bias in Rigobon and Sack's estimated response of asset prices to news, is heterogeneity in forecasters' beliefs. Beliefs heterogeneity, illustrated by the sizable dispersion of individual forecasts in survey data, is a central tenet in the financial analysis of asset trading, from which the macroeconomics literature often abstracts. In the case of Rigobon and Sack's chapter, the cost of such abstraction is that the response of asset prices to news may be overestimated.

To see this point, consider again the censored estimator, $\gamma = [\mathrm{var}(\Delta s_t) - \mathrm{var}(\Delta s_t - 1)]/\mathrm{cov}(\Delta s_t, z_t)$, and note the key identifying assumption that the "structural" (i.e., net-of-news) volatility of asset prices, σ^2_ε, is the same at announcement time (say, 8:30 a.m.) in announcement days, t, and non-announcement days, $t - 1$. This assumption allows attributing the entire rise in volatility at announcement time, $\mathrm{var}(\Delta s_t) - \mathrm{var}(\Delta s_{t-1})$, to the effect of news. The estimator γ then increases linearly with $\mathrm{var}(\Delta s_t) - \mathrm{var}(\Delta s_{t-1})$.

Rigobon and Sack generalize this assumption somewhat, allowing for σ^2_ε to rise at announcement time *in predictable fashion* over its previous afternoon's level (at 2:30 p.m.) as in

(8) $$\sigma^2_{\varepsilon,t-1,8:30} - \sigma^2_{\varepsilon,t-2,2:30} = \sigma^2_{\varepsilon,t,8:30} - \sigma^2_{\varepsilon,t-1,2:30}$$

According to equation (8), the increase in σ^2_ε from its previous afternoon's level in announcement days is the same as the corresponding increase in nonannouncement days. However, rewriting equation (8) as

(9) $$\sigma^2_{\varepsilon,t,8:30} - \sigma^2_{\varepsilon,t-1,8:30} = \sigma^2_{\varepsilon,t-1,2:30} - \sigma^2_{\varepsilon,t-2,2:30}$$

makes it apparent that equation (8) extends the benchmark assumption $\sigma^2_{\varepsilon,t,8:30} = \sigma^2_{\varepsilon,t-1,8:30}$ only if σ^2_ε differs significantly at 2:30 in the two nonannouncement days prior to t. My best guess is that there is *no* systematic difference in 2:30 p.m. volatility between $t - 1$ and $t - 2$ so that equations (8) and (9) effectively reduce to $\sigma^2_{\varepsilon,t,8:30} = \sigma^2_{\varepsilon,t-1,8:30}$.

While I have no hard data on hand to document this conjecture, indirect evidence seems compelling enough. Consider, for instance, figure 8C.1, which draws data from Fleming and Remolona (1999), the benchmark study of the minute-by-minute behavior of the Treasury market examined by Rigobon and Sack.[2]

Figure 8C.1 plots data on intraday price volatility in the five-year Treasury note market, distinguishing between announcement and nonannouncement days. To assess my conjecture that $\sigma^2_{\varepsilon,t-1,2:30} = \sigma^2_{\varepsilon,t-2,2:30}$, I would need

2. I thank Michael Fleming for providing me the data needed for figure 8C.1.

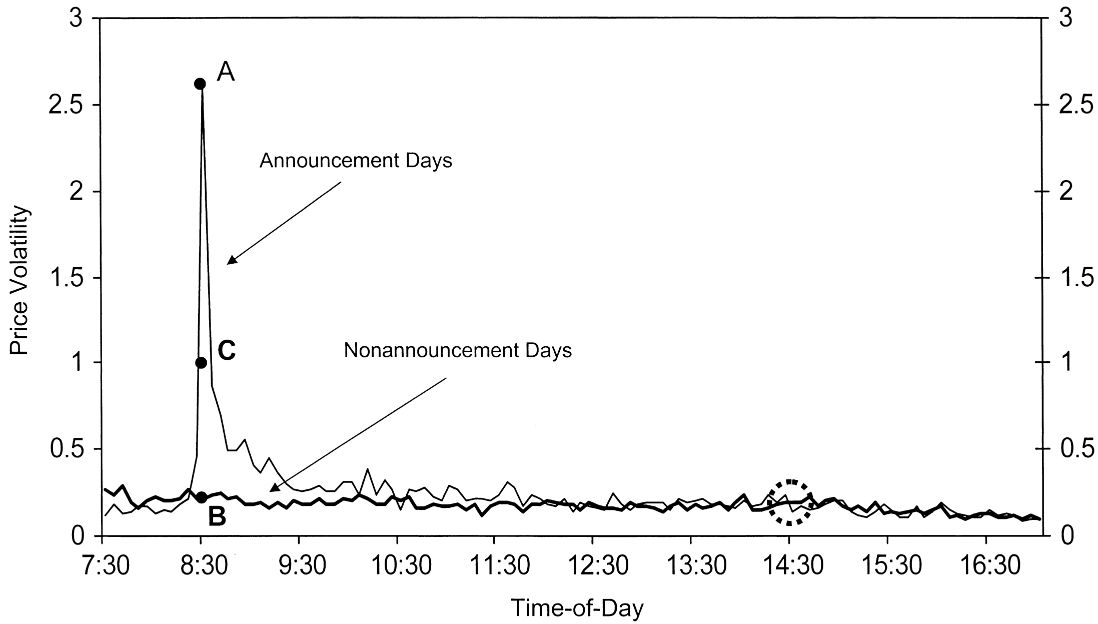

Fig. 8C.1 Intraday price volatility: Five-year treasury notes

Source: Data from figure 1A in Fleming and Remolona (1999). Volatility is measured as standard deviation of log price changes for the five-year note for days with at least one macroeconomic announcements (see Fleming and Remolona [1999] for a complete list of announcements) and days with no announcement, over five-minute intervals starting at the plotted time. Actual values are multiplied by 1,000. The sample period is August 23, 1993 to August 19, 1994.

a breakdown of volatility in the two days prior to announcements. The data I obtained do not offer such breakdown, but figure 8C.1 shows that there is negligible difference in 2:30 p.m. volatility even between announcement and nonannouncement days. I find it hard to believe, then, that there might be a systematic difference in volatility between two (almost random) nonannouncement days. if there is no such gap in volatility, Rigobon and sack's identifying restriction then reduces to assuming that the "structural" price variance at announcement time—that is, the variance when the macroeconomic release comes at its median (or mean) forecast—equals the variance at the same time in nonannouncement days. Rigobon and Sack's censored estimator can then be represented in figure 8C.1 as the distance between point A and point B (scaled by cov[Δs_t, z_t]), where point B captures the price volatility that would have been recorded *with no news*.

There are reasons, however, to believe that the volatility at announcement time will be higher in announcement days than in nonannouncement days even when $z_t^* = 0$. The key reason is that in a world with heterogeneous beliefs, $z_t^* = 0$ *is news to all but the median* (or mean) *forecaster*. Almost all investors will want to trade on the announcement, which has come either higher or lower than they individually expected. Much of this trading may reflect private information about the impact of the announcement, inducing price volatility along channels well studied in the finance literature.

While I am not aware of any direct evidence corroborating the view that even announcements at the median forecast generate price volatility, a considerable amount of indirect evidence comes to its support. Fleming and Remolona (1999), for instance, show that bid-ask spreads widen and trading volumes decline in the Treasury market *in advance* of macroeconomic announcements, with both indicators retracing their way *upon* announcements. This evidence suggests that, in a way, trading volume and price volatility might shift from *just before* to *just after* announcement times, irrespective of the actual data release. Other evidence is provided by studies showing that asset price volatility upon announcement rises with the preannouncement dispersion in beliefs (see, for instance, Green [2004] and Pasquariello and Vega [2006]). Because there is always some dispersion in beliefs about the release in announcement days (while there is none, by definition, in nonannouncement days), structural price volatility is bound to be higher at announcement times even when $z_t^* = 0$. More anecdotally, market participants report opening large positions (long or short, depending on beliefs) in anticipation of releases. These positions get unwound upon announcement, whether the release comes in at its median value or not. As trading volume rises, price volatility is likely to do the same.

In sum, if announcement and nonannouncement days are structurally different in terms of trade dynamics, announcement days are likely associated with higher structural volatility. If properly incorporated, this larger volatility would imply a lower numerator and, hence, a lower value, for Rigobon and Sack's censored estimator γ. Illustratively, in figure 8C.1, news may be responsible for lifting volatility not from point B to point A, but from, say, point C to point A.

Finally, here is a word on the results, which are broadly sensible. The gist of the analysis is that the censored estimates of the impact of news on asset prices are larger than those yielded by event studies that disregard the possibility of attenuation bias. One puzzling aspect of Rigobon and Sack's results is that certain data releases, such as capacity utilization, Institute for Supply Management (ISM), consumer confidence, and new home sales, are estimated to consist almost entirely of noise. These estimates boost the estimated coefficients γ tenfold above the coefficients yielded by earlier event studies. I find the estimated noise component of these indicators to be implausible and view these results as strengthening my belief that the forecast data examined here might not quite fit the classic error-in-variables model.

In sum, this is a nice chapter that offers a considerable improvement in our understanding of the impact of news on asset prices. The methodology offered by the authors may not correct for *all* the shortcomings associated with survey forecast data, but certainly makes a significant advance in correcting for a first-order problem that has been disregarded in previous studies of the impact of news on asset prices.

References

Faust, Jon, John H. Rogers, Shing-Yi B. Wang, and Jonathan H. Wright. 2003. The high-frequency response of exchange rates and interest rates to macroeconomic announcements. Board of Governors of the Federal Reserve System, International Finance Discussion Papers no. 784 October.

Fleming, Michael J., and Eli Remolona. 1999. Price formation and liquidity in the U.S. Treasury market: The response to public information. *Journal of Finance* 54 (5): 1901–15.

Green, T. Clifton. 2004. Economic news and the impact of trading on bond prices. *Journal of Finance* 59 (3): 1201–33.

Gürkaynak, Refet S., and Justin Wolfers. 2005. Macroeconomic derivatives: An initial analysis of market-based macro forecasts, Uncertainty, and Risk. University of Pennsylvania. Unpublished Manuscript.

Ito, Takatoshi. 1990. Foreign exchange rate expectations: Micro survey data. *American Economic Review* 80 (3): 434–49.

Judge, George G., R. Carter Hill, William Griffiths, Helmut Lutkepöhl, and Tsoung-Chao Lee. 1982. *Introduction to the theory and practice of econometrics.* New York: Wiley.

Lamont, Owen A. 2002. Macroeconomic forecasts and microeconomic forecasters. *Journal of Economic Behavior and Organization* 48 (3): 265–80.

Laster, David, Paul Bennett, and In Sun Geoum. 1999. Rational bias in macroeconomic forecasts. *Quarterly Journal of Economics* 114 (1): 293–318.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers. 2003. Disagreement about

inflation expectations. NBER Working Paper no. 9796. Cambridge, MA: National Bureau of Economic Research, June.
Pasquariello, Paolo, and Clara Vega. 2006. Informed and strategic order flow in the bond markets. University of Michigan. Unpublished Manuscript.

Discussion Summary

Richard H. Clarida observed that before announcements of payroll or Consumer Price Index (CPI) numbers, market participants often positioned themselves in volatility trades so that they would do well whether the market moved up or down. In that context, it would be interesting to see what happened when the number came out at consensus.

Drawing a connection with earlier discussions, *Marvin Goodfriend* said that the fact that the authors found a significant announcement effect on long-term interest rates may be a manifestation of classic optimal monetary policy: short rates are expected to be persistent. He also noted that it was interesting that equity prices only tended to respond to inflation news, as optimal monetary policy makes the world behave as much as possible like a flexible price economy.

Thomas Laubach said that news affects asset prices primarily by changing expectations of future policy actions. In that light, it would be interesting to look at time variation, perhaps by splitting into subsamples: for example, in 2001, the CPI was not particularly interesting, but in 2003, when people were worried about inflation, it was probably watched more closely. Sack responded that it was hard to estimate time varying effects but that this could perhaps be achieved using the principal component approach. He also drew attention to the fact that in a separate paper with coauthors, he had found that forward rates ten years out respond strongly to news.

John C. Williams and *Stephen G. Cecchetti* suggested that one could use more information from surveys than just the median forecast. Williams proposed examining the interaction between announcement effects and the dispersion in beliefs in the survey. Cecchetti noted that in principle it should be possible to use an optimal weighting of forecasts.

Clarida commented that Goldman Sachs has economic derivatives contingent on payrolls data. He said that it was surprising how close the first moment was to the Bloomberg consensus. Sack observed that the price history on these assets was still short.

John Y. Campbell said that there was an interesting parallel between the macroeconomic announcements studied in this chapter and the announcements of corporate earnings studied in the finance literature. First, there are private data vendors such as StarMine that reweight individual analysts' earnings forecasts optimally. Second, Andrea Frazzini and Owen

Lamont have found that the average returns on companies making earnings announcements were positive—there is money to be made by buying equities around the times of their earnings announcements. One explanation of this is that risk is concentrated around the earnings announcement. He asked whether the authors had calculated unconditional mean returns around announcements. Sack replied that they had not focused on unconditional means but that the option markets certainly recognized that there was considerable risk surrounding announcements.