Federal Reserve Bank of Minneapolis
Research Department

EXPLAINING FORECAST REVISIONS

Richard M. Todd*

Working Paper 355

Revised June 1988

NOT FOR DISTRIBUTION
WITHOUT AUTHOR APPROVAL

ABSTRACT

Forecasts are routinely revised, and these revisions are often the subject of informal analysis and discussion. This paper argues 1) that forecast revisions are analyzed because they help forecasters and forecast users to evaluate forecasts and forecasting procedures, and 2) that these analyses can be sharpened by using the forecasting model to systematically express its forecast revision as the sum of components identified with specific data revisions and forecast errors. An algorithm for this purpose is explained and illustrated.

Keywords: Forecasting, Forecast Revisions, Data Revisions, Innovation Accounting

*Federal Reserve Bank of Minneapolis. I have benefited from conversations with Larry Christiano, Preston Miller, David Runkle, John Rust, Christopher Sims, and Warren Weber. Thanks also to Kathy Rolfe for suggesting formats for my tables and to Pat Haswell for editorial assistance. I remain responsible for any faults.

Forecasters in economics and other disciplines frequently forecast the same event more than once, as time passes and information relevant to the event accumulates. As a result, the forecaster generates a sequence of forecast revisions and, not infrequently, a series of questions arise about exactly why the forecast has changed. Since in general forecasts are revised because of revisions in the data or errors in previous forecasts, these questions can be viewed as attempts to evaluate both the current forecast and the underlying forecasting procedure.

Although it is common and reasonable for forecasters to provide and forecast users to demand explanations of forecast revisions, I am not aware of any systematic treatment of how to provide these explanations. In this paper I propose a comprehensive procedure for explaining revisions in the forecasts of complete (no exogenous variables) linear models. The procedure uses the model to decompose its own forecast revisions into components corresponding to data revisions and errors of the previous forecast. I argue and attempt to illustrate that this procedure can provide useful information to both forecasters and forecast users. Section I discusses why people analyze forecast revisions and why a comprehensive analysis would be useful. Section II then explains and illustrates how the suggested procedures attempt to achieve what I regard as the goals of forecast revision analysis.

I. Why Are Comprehensive Analyses of Forecast Revisions Useful?

It is easy to understand why people care about forecasts. In a dynamic world, our current optimal choice depends in part on the probability distribution of future events, and forecasts are supposed to tell us something (for example, the mean or mode) about that distribution.

It is less easy to understand why people should care about forecast revisions. If we have today's optimal forecast in our hands, why should we care about yesterday's optimal forecast? The current optimal forecast is sufficient for making current decisions; yesterday's should be of concern only to historians.

In fact, however, forecasters, journalists, and forecast users all pay attention to forecast revisions. Apparently in response to the perceived desires of their audience, forecasters often provide, in their presentation of a new forecast, an explanation of the factors that led them to revise it (for example, Litterman 1985; U.S. Congress 1987; Chimerine, Behravesh, and Hagens 1987, p. 2; Hagens and Chimerine 1987, pp. 1.9–1.10; Brownstein 1987; Sprinkel 1987a,1987b). In noteworthy cases, the business press amplifies the discussion (Franklin and Cooper 1987; Cooper and Madigan 1988; Hershey 1987; Wessel 1987).[1]

The most likely reason forecast users care about forecast revisions is that they don't believe that any of the forecasts available to them are fully optimal. There are many competing forecasts, each based on simplified models and possibly inaccurate data. Forecast revision analysis can help forecasters and forecast users to evaluate competing current forecasts and their underlying forecasting models or methods. It also helps forecast users to modify the forecast based on personal information or their own views.

Forecast revision analysis generates useful information. Forecast revision analysis provides information about the structure and efficiency of a forecasting procedure. This information is useful in evaluating and using the procedure and its forecasts.

By singling out events that have had big impacts on the forecast, revision analysis identifies the dynamic and cross–variable linkages in the forecasting procedure that were critical in altering its current prediction. Since the linkages identified in this way also tend to be among the critical linkages in the forecasting procedure itself, revision analysis helps both the forecaster and the forecast user to better understand how the procedure works. In light of this information the forecaster may decide to change the model, and the forecast user may decide to change forecasters (or at least pay less attention to forecasters using models whose properties are revealed to be implausible to the user).

In addition to structural[2] information, revision analysis also provides evidence about the efficiency of the forecasting procedure. Using information that was available when the forecast was made, it should be impossible to systematically predict an efficient forecasting procedure's revision. However, a user may believe that some of the factors that caused the revision should have been foreseen at the time of the previous forecast. Used in this way, forecast revision analysis can be both a red flag and a diagnostic tool for a broad spectrum of forecasting inefficiencies, such as failure to take account of systematic bias in data revisions, structural change in coefficients, or omitted variables.

Understanding why forecasts are revised may also help forecast users to combine forecasts with their own private information or opinions. For example, suppose a forecaster in early October 1987 had cited weakness in September in stock prices and hours worked to explain a downward revision in predicted 1988 real gross national product (GNP). Knowing that the September 1987 hours–worked survey had atypically occurred during the week of the Labor Day holiday, a client of this forecaster might discount that GNP revision because of a presumption that hours worked in September had been underestimated.

Thereafter, the forecaster might also make further adjustments in his or her views based on the data released during the interval leading up to the forecaster's next forecast. When the stock market crashed on October 19, the user's understanding of how September's stock price decline had affected the forecaster's previous revision might help the user to guess the forthcoming post–crash forecast revision. Experience with previous analyses of forecast revisions thus helps a forecast user modify the forecaster's previous predictions during the interval between forecasts.[3]

Forecast revision analysis is not redundant. Given that forecast revision analysis can provide useful information, the question remains whether this information might be obtained more precisely or more easily with other techniques. I will argue that forecast revision analysis is a nonredundant source of information that builds upon and complements other approaches.

To explain a forecast revision, a forecaster must identify the important information (data revisions and forecast errors) that has accumulated since the previous forecast and then estimate how much each important bit of news affected the forecast. This can't be done without an understanding of the dynamic properties of the forecasting procedure. For that reason, forecast revision analysis makes use of familiar characterizations of model dynamics, such as elasticities, dynamic multipliers, impulse responses, and variance decompositions. However, forecast revision analysis disciplines the forecaster to use existing tools more frequently and in more revealing combinations. In that way it achieves an effect greater than the sum of the effects of its parts.

For example, forecast revision analysis tends to pick up or focus on structural properties of the forecasting procedure that other sources of information may miss. Elasticities, dynamic multipliers, impulse responses, and variance decompositions are usually presented in a form that focuses on the average static or dynamic effects of shocks to individual variables. For example, impulse responses are often presented as the separate response of each variable in a model to a given variable's typical shock (generally a one–standard–deviation increase). A forecast revision analysis, by contrast, would show the response of the model to the most recent combination of actual shocks to all variables.

Because they focus on the average effects of individual shocks rather than the current effects of multiple shocks, the other techniques may fail to reveal some of the forecasting procedure's important dynamic properties. For example, even when graphs of a model's impulse responses have plausible shapes, a forecast user may find it difficult to judge whether their absolute and relative amplitudes are plausible. Decomposition of the historical variance of the procedure's forecasting errors might sharpen the user's judgment of this issue, but examining the model's implications in a practical situation about which the user is fairly well informed—such as the analysis of the current forecast's revisions—might do so more effectively. Futhermore, the competing sources of structural information tend to be computed only occasionally, whereas forecast revision analyses are naturally computed every time a forecast is made. Thus, forecast revision analysis may discipline the forecaster to more quickly spot defects in the model's structure. Finally, it is possible that some combination of shocks occurs often, but not often enough to dominate a variance decomposition. Because it encourages researchers to track the effects of shocks as they occur, forecast revision analysis might reveal such subdominant patterns.

In providing information about the efficiency of a forecasting procedure's performance, forecast revision analysis is again likely to supplement other approaches. Direct tests of efficiency (Nordhaus 1987) and general tests and diagnostics for misspecification tend to signal the presence but not the source of inefficiencies in the forecasting procedure. By examining the individual contributions of each variable's data revisions and forecast errors to the forecast revision, and tracking these contributions over time, forecast revision analysis can provide a much richer account of why a forecasting procedure has been wrong or inefficient. I suspect, for example, that regular forecast revision analysis would prompt forecasters to rethink the way they treat economic data.[4]

A comprehensive approach is most useful. The advantages of forecast revision analysis that I have cited are best captured by what I will call a comprehensive analysis. This approach processes all new information exactly as it was processed by the forecasting procedure. Failure to do this has limited the usefulness of current discussions of forecast changes.

A comprehensive forecast revision procedure examines all changes in the information used in the forecasting procedure. This includes both revisions of the old data (data used in the previous forecast) and releases of new data (data on the procedure's variables covering time periods for which no data were available at the time of the previous forecast). Many of the forecast revision analyses in the scholarly and popular literature are selective, focusing on a few parts of the new information judged (often for unspecified reasons) to be most important. However, to realize some of the advantages of forecast revision that I have cited, such as sensitivity to data revisions and ability to identify important dynamic properties, it is important to examine all new information

Perhaps the greatest advantage of a comprehensive procedure, however, is that it processes the new information just as the actual forecasting procedure did. For this reason, a comprehensive analysis is possible only with an explicitly documented, replicable forecasting procedure, such as a complete (no exogenous variables or unmodeled add factors) forecasting model. With such a model the incremental effect on the forecast of each new bit of information can be quantified. This means that competing explanations of the forecast revision can be compared and ranked. The results can also be checked for robustness by varying the assumptions used to identify each variable's share of the new information (see section II). No other technique can do these things so well.

II. Comprehensive Analyses of Forecast Revisions Can Be Computed

I have constructed an algorithm that analyzes revisions in the forecasts of complete linear models. Depending on the model to which it is applied, the algorithm is either exactly or nearly comprehensive, in the sense discussed in section I. That is, it processes all new information either exactly or almost exactly as the forecasting model does. The algorithm then summarizes all this information by presenting a decomposition of the forecast revision. The revision in each variable's forecast at each horizon is broken down into 2N components, where N is the number of variables (all endogenous) in the model. N of the components quantify the effects of each variable's data revision, and the other N quantify the effects of each variable's newly released data. The quantification of these components is somewhat arbitrary, in the sense that it depends on identifying assumptions supplied by the user. However, this is typical of causal reasoning in economics, and the results can at least be checked for sensitivity to the identifying assumptions.

The algorithm. The algorithm can be thought of as proceeding in two stages. In stage one, the model and a set of identifying assumptions are used to partition the set of all new information (received since the previous forecast was computed) into distinct subsets. In stage two, the model and another set of identifying assumptions are used to compute the contribution to the forecast revision of each of these subsets.

The new information that causes forecast revisions includes data revisions and newly released data. To make the discussion concrete, suppose a forecast of x (an $n \times 1$ vector) for periods $t + 1$ to $t + k + j$ ($k \geq 1$, $j \geq 1$) was made at time t on the basis of data on x up through t. A subsequent forecast of periods $t + k + 1$ to $t + k + j$ is made at $t + k$ on the basis of data up through $t + k$. The information accumulated between t and $t + k$ will almost certainly cause the forecasting model to revise its forecasts of x_{t+k+1} to x_{t+k+j}. This new information consists of revised values of x_s, $s \leq t$, and newly released data on x_s, $t + 1 \leq s \leq t + k$.

The algorithm begins by identifying each variable's data revision with itself. That is, we observe revisions $r_{i,t-h} = {}_{t+k}x_{t-h} - {}_{t}x_{t-h}$, for $i = 1, 2, ..., n$ and $h = 1, 2, ..., H$, where ${}_{q}x_{s}$ denotes the data available at q on the value of x at s and H is the maximum number of periods before t for which data revisions will be analyzed. (H is bounded by the number of observations on x available at t+k.) My method treats the sequence $r_{i,t-1}$ to $r_{i,t-H}$ as an independent source of new information attributable to variable i only. I thus follow the common practice of neither explicitly modeling, nor even simply taking account of the correlations among, data revisions.[5] This gives me N distinct subsets of new information, one for each variable's sequence of data revisions.[6]

Having partitioned the data revisions into N subsets, I turn to the newly released data. I start by converting the new data into a sequence of one–step–ahead forecast error vectors. To do this, I first reestimate the model's coefficients using revised data through t. (Since I have already identified the data revision part of the new information, I now work only with the revised data.) Call these coefficients ${}_{t+k}\beta_t$. I then use these coefficients and revised data through t + h to forecast ${}_{t+k}x_{t+h+1}$, for $h = 0, 1, \ldots, k-1$. Subtracting these forecasts from the actual values of ${}_{t+k}x_{t+h+1}$ gives me a sequence ${}_{t+k}f_{t+h+1}$ of one–step–ahead forecast errors.

I could now use the same simple identification scheme I used for the data revisions. That is, I could simply identify the sequence of i^{th} components of ${}_{t+k}f_{t+h+1}$ as the part of the forecast error since t attributable to variable i. However, I could also use any other scheme for attributing the components of the forecast error vectors to the individual variables. Devising such schemes is what economists usually have in mind when they talk about identifying an econometric model. In my own programs, and in the examples presented below, I use the technique of orthogonal decomposition (Hakkio and Morris 1984; Doan and Litterman 1986, pp. 12–24). This gives me N subsets of information derived from the newly released data. Subset i is a sequence ${}_{t+k}u_{i,t+h+1}$ of error vectors that contain the contribution of variable i to each component of ${}_{t+k}f_{t+h+1}$, $h = 0, 1, \ldots, k-1$. Because the u sequences add up to the f sequence, they can be said to decompose the f sequence.

Since the choice of an identification scheme is controversial among economists, I want to emphasize that my procedure doesn't depend on any particular one. Orthogonal decomposition is simple to program and is convenient for illustrating my technique (see below). In other contexts it might be appropriate to devise an alternative identification based on economic theory and observations. Such alternatives can be substituted for my orthogonal decomposition, provided that the N sequences of forecast error components they generate also add up to the f sequence.

Alternative identification schemes for the forecast errors, including alternative orthogonalizations (based on different orderings of the variables), will change the algorithm's decomposition of the forecast revision. In fact, it is not hard to find cases of fairly substantial change, and one example is presented below.

In the second stage of my algorithm, I compute the components of the forecast revision attributable to each of the 2N subsets of new information. To do this I employ what can be thought of as additional identifying assumptions. However, the results do not seem to be very sensitive to these additional assumptions.

Note that in choosing to decompose the forecast revision into N additive components, I am generally suppressing some information. A more complete decomposition of the forecast revision, for example, might include components for all possible multivariate interaction effects. Unless the true mapping from the subsets of new information to the components of the forecast revision is additive, any N–component decomposition must suppress information about interaction effects between the subsets. Furthermore, the class of models for which the true mapping is additive is fairly restrictive, consisting principally of linear models whose coefficients are not updated between forecasts. (See Appendix A for a demonstration that even linear models generate nonadditive mappings if their coefficients are updated.)

As a practical matter, however, I am not too worried about these interaction effects. They are likely to be small for linear models whose coefficients change only slightly in response to new information. This has at least two implications.

Since the models I generally work with seem to fall into this category, one implication is that I can usually ignore the small degree of ambiguity caused by interaction effects in stage two of my algorithm. If necessary, such as when very surprising new data have caused an unusually large change in the coefficients, I can check the sensitivity of my results to interaction effects by varying my identifying assumptions.

Another implication is that for linear models with slowly changing coefficients, my complete algorithm can usually be closely approximated by a simpler algorithm discussed in Appendix B. The simpler algorithm uses a fixed set of coefficients and thereby ignores interaction effects altogether. Currently I have programmed it only to analyze the effects of newly released data (that is, not data revisions) and use it mainly as a check on my complete algorithm. I have found only a few cases where the simple and complete algorithms differ to an interesting degree on the effects of new information. (The rather mild example presented below is the strongest I have found so far.) However, because of these few cases, the likelihood that they would be more common in models that allow more rapid coefficient change, and the speed with which my complete algorithm can be computed when the Kalman filter is used to calculate the coefficient updates, I will focus my remarks on the complete algorithm.

The identifying assumptions I normally use suppress interaction effects by, in effect, decomposing each of them into shares attributed to the subsets involved in the interaction. First I order the 2N subsets of new information (N for data revisions and N for new data). Starting with the first subset and proceeding sequentially, I then attribute to each subset its incremental contribution to the forecast revision. In abstract terms, let S_i be the sequence of new information vectors that defines subset i, i = 1, 2, ..., 2N. Let Z be a sequence of zero vectors (of conformable dimension). Then define a function F mapping from the space in which the 2N and Z sequences lie to the space of forecast revisions in such a way that $F(Z,Z,...,Z) = 0$ and $F(S_1,...,S_N)$ equals the forecast revision. Then I attribute $F(S_1,Z,Z,...,Z)$ to subset 1 and $F(S_1,S_2,...,S_i,Z,Z,...,Z) - F(S_1,S_2,...,S_{i-1},Z,Z,...,Z)$ to subset i, i = 2, 3, ..., 2N. This completes the algorithm.

Note that as it computes its decomposition of the forecast revision, the algorithm implicitly parcels out interaction effects to each of the variables involved, based on their incremental contribution. The incremental contributions obviously depend on the ordering of the subsets, so the importance of interaction effects can be checked by rerunning the algorithm with different orderings.

The abstract discussion in the previous paragraph highlights how interaction effects are handled but does not provide many details of just how the algorithm operates. Let me correct this by describing how I order the subsets and use the model to implicitly compute the F function.

The only restrictive feature of the way the algorithm is programmed is that I must put the N data revision subsets first. The order of the data revision subsets within this group is not usually important, because my models are linear and normally their coefficients change only slightly. Typically I use the same order of variables here that I used in the orthogonal decomposition of forecast errors in stage one.

Given an order of the data revision subsets, I construct a series of data bases, d_0, d_1, ..., d_N. Data base d_0 is the one actually used in the previous forecast. Data set d_i is obtained from d_{i-1} by adding to each variable in d_{i-1} the effect of data revisions for variable i, i = 1, 2, ..., N. (In the simple identification scheme I use for data revisions, this means that variable i changes by the full amount of its revision and the other variables are unaffected.) The coefficients and forecast based on d_0 are already available, of course. Denote the elements of the latter as ${}_{d_0}x_{t+1}$ to ${}_{d_0}x_{t+k+j}$. I then reestimate the model using d_1. With the reestimated coefficients and data from d_1, I compute a new forecast, ${}_{d_1}x_{t+1}$ to ${}_{d_1}x_{t+k+j}$. The sequence

$$C_{d_1} = \{({}_{d_1}x_{t+k+1} - {}_{d_0}x_{t+k+1}),\ldots,({}_{d_1}x_{t+k+j} - {}_{d_0}x_{t+k+j})\}$$

records the contributions of revisions in the data for variable 1 to each component of the forecast revision.

The next step is to reestimate the coefficients based on d_2 and compute

$$C_{d_2} = \{({}_{d_2}x_{t+k+1} - {}_{d_1}x_{t+k+1}),\ldots,({}_{d_2}x_{t+k+j} - {}_{d_1}x_{t+k+j})\}.$$

This records the (incremental) contributions of revisions in the data for variable 2 to each component of the forecast revision. Repeating these steps for d_3, d_4, ..., d_N, with the obvious modifications, completes the analysis of the effects of data revisions.

To analyze the effects of the newly released data, I parallel the analysis of the data revisions. First I order the sequences of forecast error vectors that were identified in stage one. (The order here need not match either the order used in stage one or the order used in stage two for the data revision sequences, but I usually make them all agree.) I also construct data base e_0 by adding to d_N the forecasts ${}_{d_N}x_{t+1}, \ldots, {}_{d_N}x_{t+k}$ that would have been computed at the time of the previous forecast if the revised data had been available then. Then I construct data bases e_1, e_2, ..., e_N by adding in the sequences of forecast error vectors (that is, ${}_{t+k}u_{1,t+h+1}, \ldots, {}_{t+k}u_{N,t+h+1}$, $h=0,1,\ldots,k-1$) in the specified order. In particular, the value of the j^{th} variable at time $t + h + 1$ in e_i equals its value in e_{i-1} plus ${}_{t+k}u_{i,t+h+1}$, which is the part of its time $t + h + 1$ forecast error attributed to variable i ($i=1,2,\ldots,N$).

With these data bases in hand, I begin a sequence of reestimating coefficients through time $t + k$ and reforecasting periods $t + k + 1$ to $t + k + j$.[7] At each step I take the difference between the forecast based on e_i and the forecast based on e_{i-1} to define the (incremental) contribution of new data on variable i to the forecast revision. That is, paralleling the notation above,

$$C_{e_i} = \{({}_{e_i}x_{t+k+1} - {}_{e_{i-1}}x_{t+k+1}), \ldots, ({}_{e_i}x_{t+k+j} - {}_{e_{i-1}}x_{t+k+j})\}$$

records the contribution of new data on variable i to each component of the forecast revision.

Some examples. I use this algorithm regularly to analyze the forecast revisions of two Bayesian vector autoregressive forecasting models. One model forecasts the monthly values of seven key macroeconomic variables—the Dallas Federal Reserve Bank's index of the foreign exchange value of the U.S. dollar (DOLLAR), Standard and Poor's index of 500 stock prices (SP500), the interest rate on three–month U.S. treasury bills (TBILL), real gross national product (GNP), the GNP deflator (DEFL), the change in business inventories (CBI), and the Federal Reserve Board's measure of the monetary base (MB).[8] The model can stand alone but also acts as the core sector in a large macroeconomic model used regularly to prepare forecasts at the Federal Reserve Bank of Minneapolis. [See Litterman (1984) for a description of both the core model and the larger model.] The other model consists of a 13–variable quarterly macroeconomic sector plus six separate recursive regional sectors. [Amirizadeh and Todd (1984) describe the essential features of an early version of this model.] Its regional economic forecasts are published quarterly in the Minneapolis Fed's District Economic Conditions.

Table 1 shows the variables in the monthly core model and the standard format I use for reporting decompositions of the model's forecast revisions. The initial forecast in this example was made in September 1987, based on data (official, interpolated, and projected) through July 1987 for GNP, DEFL, and CBI and through August 1987 for the other variables. The subsequent forecast was made in March 1988, based on revised and new data through January 1988 for GNP, DEFL, and CBI and through February 1988 for the others. The table analyzes percentage point changes in forecasts of the levels of the variables in December 1988, except for TBILL (basis point units) and CBI ($billion units).

To decompose the revision in a particular variable's forecast (for example, GNP), find the column corresponding to that variable (column 4) and read down. The first row shows the variable's total forecast revision for December 1988 (–1.09 percentage points). The next three rows give a gross decomposition of the total revision into effects attributed to data revisions (+1.70 percentage points), newly released data (–2.77 percentage points), and miscellaneous factors (–0.02 percentage points). (See end note 7 for a discussion of the miscellaneous factors.)

The next block of rows further decomposes the total data revision effect into effects attributed to each variable's data revision. The first row of column 4 in this block, for example, attributes –0.02 percentage points of the total –1.70 percentage point effect of data revisions on the forecast of the December 1988 level of GNP to revisions in the data on the exchange value of the U.S. dollar.

The final block decomposes the total newly released data effect into effects attributed to each variable's newly released data. The third row of column four in this block, for example, says that movements in the GNP deflator that were not predicted in a revised (to incorporate data revisions but not new data) initial forecast contributed +0.29 of the –2.77 percentage point effect of new data on the GNP forecast for December 1988.

In many ways the results in Table 1 are typical of my experience with revisions of this monthly model's forecast. Surprises in the new data are usually more important than data revisions, but data revision effects are far from trivial. The omitted information summarized by the "miscellaneous factors" row is trivial. In the lower blocks, diagonal elements generally dominate. This reflects the fact that a variable's own revisions and forecast errors usually cause most of its forecast revisions, at least in this model when its errors are identified by means of orthogonal decompositions.

Table 1 shows some less common results too. Most prominent are some strong cross–variable effects attributed to new data on SP500. Considering that the October 1987 stock market crash intervened between the initial and subsequent forecasts, these strong effects are understandable. The algorithm says that the stock market crash itself cut the forecasts of December 1988 GNP, MB, and TBILL by 2.75 percent, 0.57 percent, and 57.16 basis points, respectively.

Table 1 also illustrates some of the benefits of forecast revision analysis. The large effect attributed to GNP revisions highlights the importance of data revisions in general. In a concerted effort by several economists at the Minneapolis Fed to unravel why the forecast of GNP had not been depressed more by the stock market crash, this factor had not been considered until the algorithm was used to produce Table 1. The table also helped rank competing explanations. For example, the fall in interest rates in late 1987 had been proposed as an explanation for why the model's GNP forecast had changed by only –1.09 percent. However, the table shows that new data on interest rates had a trivial effect (+0.04 percent) on the GNP forecast for December 1988.

Of course, these conclusions could be sensitive to the orderings used to identify the components. To produce Table 1, the algorithm used a single ordering—DOLLAR, SP500, TBILL, MB, DEFL, CBI, GNP—for all decompositions. Table 2 shows an analysis of the same forecast revision based on an alternative ordering—MB, TBILL, DOLLAR, SP500, GNP, DEFL, CBI. The decomposition of the effects of data revisions appears to be completely insensitive to the switch, probably because ordering affects this decomposition only when the model's coefficients are changed a lot by the revisions. The effects of data surprises are sensitive to orderings, but in most of the columns the changes are small. Only for TBILL and, to a lesser extent, MB do the decompositions appear importantly different.[9] This means that the results in these columns must be interpreted with caution, unless the user can supply reasons for preferring a particular ordering (or, more generally, a particular identification scheme).[10]

Table 3 provides an example where the complete algorithm decomposes the effects of new data differently than the simpler approximation discussed in Appendix B. The example is taken from the Wisconsin sector of the quarterly national–regional model, with forecasts of the state's unemployment rate in fourth quarter 1975 based on data through fourth quarter 1973 initially and third quarter 1974 subsequently. This time period was chosen because the coefficients of the Wisconsin unemployment equation shifted relatively sharply then. As a result, SP500's share of the total 1.88 percentage point revision in the forecasted rate rose from 0.79 percentage points in the simple algorithm to 0.91 in the complete algorithm, and the effect ascribed to PR28, an index of 28 commodity prices, fell from 0.46 to 0.29 percentage points. For other variables the differences between the two algorithms are fairly small.

This example of a discrepancy between the simple and the complete algorithms is mild but nonetheless rare, at least in my experience. Even by focusing on periods when changes in the models' coefficients were relatively rapid, I have uncovered only a few other nontrivial examples, none of them even as strong as the one shown in Table 3. That is the basis of my tentative conclusion that the simpler model approximates well when rates of coefficient variation are not extreme.

Concluding Remarks

The technical content of this paper is simple; the algorithms consist of accounting conventions (whose practical application happens to be tedious to describe and program). Nonetheless, the algorithms have already helped me to provide better analyses of the forecasts I publish. I have also benefited from them in discussions with my colleagues of the forecasts used in an actual policymaking setting. I believe that these algorithms, or something like them, should be used routinely to help organize the inevitable but often disjointed discussions of forecasts and forecast revisions.

Footnotes

[1]The reactions of forecast users to forecast revisions are not well documented, so for the most part I must rely on the reader's own observations on this point. However, the records of the Federal Reserve System show that at least one important group of forecast users—the Federal Open Market Committee—receives an account of why its staff's forecasts have been changed and also discusses its members' views of forces that have changed the economic outlook [for example, Federal Reserve Board of Governors (1988, p. 115)].

[2]I use words like "structure" or "structural" here to refer to the characteristics or properties of the forecasting procedure. I do not mean to imply that these properties necessarily reflect the true properties of the economic (or other) system being forecasted.

[3]This function of forecast revision analysis was suggested to me by David Runkle.

[4]Typical current practice is to treat all available data as though they were final and true, and then to replace them as soon as revised data are released. Recognition of the important role data revisions sometimes have in changing forecasts (see the example in section II) should at least cause forecasters to consider widening their forecast confidence bands to account for data uncertainties.

[5]This trivial identification scheme for data revisions could be changed without modifying the rest of the algorithm. For example, I work with a model that includes both real gross national product (GNP) and the change in business inventories (CBI). Since the latter is a component of the former, it is reasonable to suppose that revisions in the CBI data

might prompt, and hence be correlated with, revisions in the GNP data. The part of the GNP revision that is caused by the CBI revision should perhaps be attributed to CBI, not to GNP. This can be done systematically for the entire vector time series of data revisions by supplying a nontrivial identification scheme. One simple possibility is to treat data revisions as correlated contemporaneously but not across time. Then an orthogonal decomposition of each period's vector of data revisions could be computed, using, for example, Choleski factors from the covariance matrix of historical revisions. My algorithm currently doesn't do that, but I am cautious about interpreting any results that seem sensitive to the exact division of data revisions between CBI and GNP.

[6]In principle it would be possible to partition the data revisions as well as the newly released data more finely than I do here. For example, there could be one subset for each revised or newly released bit of data. The steps in my current algorithm would readily generalize to such finer partitions.

[7]Note that ideally the coefficients estimated through t + k from e_0 will be identical to those estimated through t from d_N, because the data in e_0 for periods t + 1 to t + k are simply forecasts based on d_N and hence contain no new information to cause an updating of the coefficients. With no change in the coefficients and with only forecasted data added to the d_N data base, it also follows that the forecast of x_{t+k+1} to x_{t+k+j} based on e_0 is identical to the forecast of these periods that was computed from d_N. Thus, in principle, ${}_{e_0}x_{t+k+i} - {}_{d_N}x_{t+k+i} = 0$, for i = 1, 2, ..., j.

If this were not true, the part of the forecast revision caused by the difference between e_0 and d_0 would escape my method, not being attributed to any variable. As illustrated below, this component is indeed nonzero in the models I typically use, partly because the models use data–based "prior" restrictions on their coefficients (Litterman 1986). In particular, my algorithm does not account for the impact of data revisions and new data on the estimated variances of the equation error terms, which enter as known constants in the prior covariance matrix of coefficients. I also often ignore minor revisions to data early in the sample, such as those caused by recomputing seasonal adjustments. The total impact of the neglected factors is typically small, however, as shown by the size of the "miscellaneous" category in the examples in section II.

[8]Note that some of the variables in the core model, including GNP, DEFL, and CBI, are not officially available as monthly series. Amirizadeh (1985) explains the interpolation and projection procedures used to estimate monthly values for these series. The algorithm treats the official and interpolated data identically. Sometimes this can lead to a counterintuitive decomposition of the revision. For example, new data on GNP can significantly alter the interpolated and projected GNP values for recent months. Such changes, though caused by new data, would be interpreted by the algorithm as pure data revisions. As a result, the algorithm as currently applied to this model can sometimes exaggerate the importance of data revisions.

[9]The reader may wonder whether statements about the significance of the differences between Tables 1 and 2 require that confidence bands be computed for the results in the tables. There is a perspective from which this might be attempted, but it is a somewhat uncommon perspective which I have not tried to exploit. Note that each table analyzes the change in a given model's forecast under given identifying assumptions.

The given model's initial forecast can be thought of as the mode of the distribution of the predicted variables, conditioned upon information available when the initial forecast was made. The given model's subsequent forecast is the mode conditioned on the information available subsequently. Unlike the forecasted variables themselves, these modes are exact functions of known data and thus are not random once the conditioning information and the model are taken as given. The difference between these modes depends only on the difference in the conditioning information, which is obviously not random after the fact. And given the differences in the conditioning information and the modes, the decomposition depends only on the identification scheme.

The only way to compute a confidence band for the decomposition, then, is to find and exploit random elements in the identification scheme. There are two possibilities. For a given identification scheme, the decomposition depends on the estimated coefficients (because they determine the forecast errors) and on the estimated covariance matrix of one–step–ahead forecast errors. To compute the uncertainty in the decomposition for a given identification requires a procedure that samples from the coefficient and covariance matrix distributions while constraining the initial and subsequent forecasts to remain fixed at their actual values. Less conventionally, it would also be possible to posit a probability distribution over all possible identification schemes of the variables. Random samples of identifications from this distribution could then be used to compute the uncertainty in the decomposition that flows from uncertainty over identification schemes. I have not pursued either of these possibilities.

[10]I don't want to imply that the other columns can be interpreted without caution. Tables 1 and 2 are meant to illustrate rather than exhaust the issue of robustness with respect to ordering. A more complete robustness check with respect to all possible orderings could be done, given enough computer time.

Appendix A

The Nonadditivity of Mappings from New Information to Forecast Revisions When Model Coefficients Are Reestimated

Consider a multivariate model that expresses the vector $y_t^T = [y_{1t}, y_{2t}, \ldots, y_{nt}]^T$ as a time–varying linear function of the components of its own first lag plus an error term. Let the function M: $R^n \to R^n \times R^{n^2}$ be given by

$$M(y_t) = \begin{bmatrix} y_{1t} & y_{2t} & \cdots & y_{nt} & 0 & 0 & \ldots & 0 & \ldots & 0 & 0 & \ldots & 0 \\ 0 & 0 & \ldots & 0 & y_{1t} & y_{2t} & \cdots & y_{nt} & \cdots & 0 & 0 & \ldots & 0 \\ \vdots & \vdots & & \vdots & \vdots & \vdots & & \vdots & & \vdots & \vdots & & \vdots \\ 0 & 0 & \ldots & 0 & 0 & 0 & \ldots & 0 & \ldots & y_{1t} & y_{2t} & \cdots & y_{nt} \end{bmatrix}.$$

Then we can write this model at time t as

$$y_t = M(y_{t-1})\alpha_t + S_t\epsilon_t,$$

where

$$\alpha_t = [\alpha_{11t}, \alpha_{12t}, \ldots, \alpha_{1nt}, \alpha_{21t}, \alpha_{22t}, \ldots, \alpha_{2nt}, \ldots, \alpha_{n1t}, \alpha_{n2t}, \ldots, \alpha_{nnt}]^T.$$

The coefficient vector α_t also evolves, according to

$$\alpha_t = T_t\alpha_{t-1} + R_t\eta_t.$$

To apply the Kalman filter, we take α_t as the unobserved state vector; y_{t-1} as the observation vector; and $M(y_{t-1})$, the $(n \times p)$ S_t, the $(n^2 \times n^2)$ T_t, and the $(n^2 \times r)$ R_t as known fixed matrices. We also model $(p \times 1)$ ϵ_t and $(r \times 1)$ η_t as mean zero, serially uncorrelated processes with contemporaneous covariance matrices H_t and Q_t, respectively. Components of ϵ_t are uncorrelated with components of η_t.

The forecast of y_{t+1} at time t, or $y_{t+1/t}$, is given by $M(y_t)T_{t+1}a_t$, where a_t is the current estimate of α_t. The error in this forecast, $y_{t+1} - y_{t+1/t}$, is denoted $u_{t+1/t}$, which is an $n \times 1$ vector. The sequence of estimates, $\{a_t\}$, evolves according to the updating formula

$$a_{t+1} = T_{t+1}a_t + G_{t+1}u_{t+1/t},$$

where G_{t+1} is the Kalman gain matrix.

To analyze a forecast revision, consider both $y_{t+1/t}$ and $y_{t+1/t-1}$. The latter is given by

$$\text{(A.1)} \quad y_{t+1/t-1} = M[y_{t/t-1}]T_{t+1}a_{t/t-1}$$

$$= M[M(y_{t-1})T_t a_{t-1}]T_{t+1}T_t a_{t-1}.$$

Since M is a linear function, $y_{t+1/t}$ can be rewritten as

$$\text{(A.2)} \quad y_{t+1/t} = M(y_t)T_{t+1}a_t$$

$$= M[M(y_{t-1})T_t a_{t-1}+u_{t/t-1}]T_{t+1}(T_t a_{t-1}+G_t u_{t/t-1})$$

$$= M[M(y_{t-1})T_t a_{t-1}]T_{t+1}T_t a_{t-1} + M(u_{t/t-1})T_{t+1}T_t a_{t-1}$$

$$+ M[M(y_{t-1})T_t a_{t-1}]T_{t+1}G_t u_{t/t-1}$$

$$+ M(u_{t/t-1})T_{t+1}G_t u_{t/t-1}.$$

Subtracting (A.1) from (A.2) gives the forecast revision as

$$\text{(A.3)}\quad y_{t+1/t} - y_{t+1/t-1} = M(u_{t/t-1})T_{t+1}T_t a_{t-1}$$

$$+ M[M(y_{t-1})T_t a_{t-1}]T_{t+1}G_t u_{t/t-1}$$

$$+ M(u_{t/t-1})T_{t+1}G_t u_{t/t-1}.$$

The final term on the right side of equation (A.3) shows that the forecast revision is a quadratic, nonadditive function of the forecast errors. The situation doesn't change if the forecast errors are replaced by orthogonalized innovations or structural equation error terms. These substitutions would simply replace $u_{t/t-1}$ in (A.3) by $Dv_{t/t-1}$, for some matrix D.

Appendix B

A Simpler Algorithm

New data affect a model's initial forecast in two ways. Most obviously, they update the variables in the model's equations. Unless the new data conform exactly to the previous forecast of their values, this will revise the forecasts covering the balance of the initial forecast horizon.

When the model's coefficient estimates are also updated with new data, changes in the coefficients also contribute to the forecast revision. (See Appendix A.) This second effect is usually much smaller than the first, however.

The simpler algorithm ignores the second, smaller effect. It proceeds exactly as the complete algorithm does, except that the model's coefficients are held fixed at the values they assumed in either the initial or the subsequent forecast.

For analyzing the effects of just new data (not data revisions), the simpler algorithm is very simple to program in the RATS statistical package (Doan and Litterman 1986). RATS contains a command, HISTORY, which uses a fixed set of coefficients to decompose the differences between a baseline forecast and the subsequently revealed actual data. To use HISTORY for forecast revision analysis, make the initial forecast (made at time t) be the base forecast, covering periods $t + 1$ to $t + k + j$. For the series that HISTORY regards as actual data, let time periods $t + 1$ to $t + k$ equal the actual data and time periods $t + k + 1$ to $t + k + j$ equal the values predicted in the subsequent forecast (made at time $t+k$). Applying HISTORY over the period $t + 1$ to $t + k + j$ then produces a fixed–coefficient decomposition of a) the initial forecast's errors in periods $t + 1$ to $t + k$, and b) the revisions of the initial forecast for periods $t + k + 1$ to $t + k + j$.

References

Amirizadeh, Hossain. 1985. "Technical Appendix for the Bayesian Vector Autoregressive Model of the U.S. Economy," manuscript. Federal Reserve Bank of Minneapolis, July.

Amirizadeh, Hossain, and Todd, Richard M. 1984. "More Growth Ahead for Ninth District States," Federal Reserve Bank of Minneapolis Quarterly Review Fall, pp. 8–17.

Brownstein, Vivian. 1987. "The Economy: Down, Not Out," Fortune November 23, pp. 61, 64.

Chimerine, Lawrence, Behravesh, Nariman, and Hagens, John. 1987. "Executive Summary — U.S. Outlook," October 30, pp. 1–4. In The WEFA Group (publisher), "U.S. Economic Outlook 1987–90," November 1987.

Cooper, James C., and Madigan, Kathleen. 1988. "Business Outlook," Business Week February 29, pp. 23–24.

Doan, Thomas A., and Litterman, Robert B. 1986. User's Manual: RATS Version 2.00. VAR Econometrics.

Federal Reserve Board of Governors. 1988. "Record of Policy Actions of the Federal Open Market Committee: Meeting Held on November 3, 1987," Federal Reserve Bulletin February, pp. 112–119.

Franklin, William B., and Cooper, James C. 1987. "Business Outlook," Business Week November 9, pp. 27–28.

Hagens, John, and Chimerine, Lawrence. 1987. "Forecast Overview," pp. 1.1–1.15. In The WEFA Group (publisher), "U.S. Economic Outlook 1987–90," December 1987.

Hakkio, Craig S., and Morris, Charles S. 1984. "Vector Autoregressions: A User's Guide." Federal Reserve Bank of Kansas City, Research Working Paper 84–10, November.

Hershey, Robert D. Jr. 1987. "'88 Growth Forecast Cut to 2.4%," The New York Times December 24, pp. 21, 23.

Litterman, Robert B. 1984. "Above–Average National Growth in 1985 and 1986," Federal Reserve Bank of Minneapolis Quarterly Review Fall, pp. 3–7.

__________________. 1985. "How Monetary Policy in 1985 Affects the Outlook," Federal Reserve Bank of Minneapolis Quarterly Review Fall, pp. 2–13.

__________________. 1986. "Specifying Vector Autoregressions for Macroeconomic Forecasting." In Goel, P. and A. Zellner (1986), Bayesian Inference and Decision Techniques, pp. 79–94. Elsevier Science Publishers B.V.

Nordhaus, William D. 1987. "Forecast Efficiency: Concepts and Applications," The Review of Economics and Statistics v. LXIX, #4, November, pp. 667–74.

Sprinkel, Beryl W. 1987a. Statement before the White House Press Corps, Thursday, August 6.

__________________. 1987b. Press briefing, December 23.

U.S. Congress, Congressional Budget Office. 1987. The Economic and Budget Outlook: An Update August.

Wessel, David. 1987. "Reagan Administration Now Expects Economy to Grow Only 2.4% Next Year," Wall Street Journal December 24, p. 26.

Table 1

One Explanation of the

Monthly Core Model's Revised Forecast for December 1988*

	Forecast Variables						
	DOLLAR (pctg. pts.)	SP500 (pctg. pts.)	DEFL (pctg. pts.)	GNP (pctg. pts.)	MB (pctg. pts.)	TBILL (basis pts.)	CBI ($ bils.)
Change in Forecast	-4.57	-30.04	-2.01	-1.09	-1.24	-101.88	-9.61
Breakdown by Reason for Change							
Revisions in Old Data	-0.55	0.27	-0.49	1.70	-0.59	2.86	-1.97
Surprises in New Data	-3.81	-29.98	-1.50	-2.77	-0.58	-105.72	-7.33
Miscellaneous Factors	-0.21	-0.33	-0.02	-0.02	-0.07	0.98	-0.31
Breakdown by Variable Changed for Each Reason							
Revisions in Old Data							
DOLLAR	-0.65	-0.09	0.01	-0.02	0.03	4.19	0.06
SP500	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DEFL	0.03	0.27	-0.37	0.05	-0.01	-1.58	0.11
GNP	-0.02	-0.13	-0.01	1.71	0.08	3.69	0.49
MB	0.09	0.25	-0.12	-0.04	-0.68	-3.64	-0.88
TBILL	0.00	0.00	0.00	0.00	0.00	0.02	0.00
CBI	0.00	-0.02	0.00	0.00	-0.01	0.18	-1.76
Surprises in New Data							
DOLLAR	-3.61	0.87	-0.04	0.01	0.12	3.16	0.45
SP500	-0.30	-32.16	0.12	-2.75	-0.57	-57.16	-5.68
DEFL	0.07	0.80	-1.50	0.29	0.00	-6.85	0.67
GNP	0.00	0.02	0.00	-0.15	-0.01	-0.42	-0.04
MB	0.00	0.00	0.00	-0.01	-0.06	-0.01	-0.06
TBILL	0.02	0.48	-0.08	0.04	-0.05	-44.08	-0.24
CBI	0.00	0.00	0.00	-0.19	0.00	-0.37	-2.43

*The following order of variables was used to identify the model: DOLLAR, SP500, TBILL, MB, DEFL, CBI, and GNP.

Table 2

An Alternative Explanation of the

Monthly Core Model's Revised Forecast for December 1988*

	Forecast Variables						
	DOLLAR (pctg. pts.)	SP500 (pctg. pts.)	DEFL (pctg. pts.)	GNP (pctg. pts.)	MB (pctg. pts.)	TBILL (basis pts.)	CBI ($ bils.)
Change in Forecast	-4.57	-30.04	-2.01	-1.09	-1.24	-101.88	-9.61
Breakdown by Reason for Change							
Revisions in Old Data	-0.55	0.27	-0.49	1.70	-0.59	2.86	-1.97
Surprises in New Data	-3.81	-29.98	-1.50	-2.77	-0.58	-105.72	-7.33
Miscellaneous Factors	-0.21	-0.33	-0.02	-0.02	-0.07	0.98	-0.31
Breakdown by Variable Changed for Each Reason							
Revisions in Old Data							
DOLLAR	-0.65	-0.09	0.01	-0.02	0.03	4.19	0.06
SP500	0.00	0.00	0.00	0.00	0.00	0.00	0.00
DEFL	0.03	0.27	-0.37	0.05	-0.01	-1.58	0.11
GNP	-0.02	-0.13	-0.01	1.71	0.08	3.69	0.49
MB	0.09	0.25	-0.12	-0.04	-0.68	-3.64	-0.88
TBILL	0.00	0.00	0.00	0.00	0.00	0.02	0.00
CBI	0.00	-0.02	0.00	0.00	-0.01	0.18	-1.76
Surprises in New Data							
DOLLAR	-3.44	0.39	0.02	-0.02	0.14	22.44	0.50
SP500	-0.31	-31.66	0.09	-2.67	-0.31	-101.64	-5.45
DEFL	0.08	0.82	-1.51	0.16	-0.01	-7.18	0.61
GNP	0.00	0.01	0.01	-0.22	-0.01	-0.43	-0.23
MB	0.00	0.06	-0.08	-0.06	-0.40	-3.93	-0.59
TBILL	-0.14	0.39	-0.03	0.04	0.01	-14.95	0.01
CBI	0.00	0.00	0.00	0.00	0.00	-0.03	-2.18

*The following order of variables was used to identify the model: MB, TBILL, DOLLAR, SP500, GNP, DEFL, and CBI.

Table 3

Two Explanations of the Effects of New Data on the Wisconsin Model's Revised Unemployment Forecast for 75:4

	Complete Algorithm	Simple Algorithm
Change in Forecast (hundredths of a percentage point)	192.38	192.38
Breakdown by Variable		
U.S. Variables		
GNP Deflator	18.96	18.54
Business Fixed Investment	7.15	5.55
GNP	143.56	140.06
Monetary Base	-2.40	-4.85
3-Month T-Bill Rate	-1.93	-3.11
Unemployment	-6.41	-4.70
Exchange Value of the Dollar	3.09	5.29
Standard & Poors 500	91.08	79.31
PR28	28.54	46.42
Nonfarm Employment	6.33	5.53
Nonfarm Earned Income	0.25	0.50
Other Personal Income	-4.69	-5.08
Retail Sales	-0.15	-0.09
Wisconsin Variables		
Nonfarm Employment	-25.34	-25.19
Retail Sales	-1.47	-1.13
Unemployment	-67.60	-66.48
Nonfarm Earned Income	-0.38	0.16
Other Personal Income	-0.45	-2.61
Miscellaneous Factors	4.26	4.26