Economic Review

Federal Reserve Bank of San Francisco

1992 Number 2

Predicting Contemporaneous Output

Bharat Trehan

Senior Economist, Federal Reserve Bank of San Francisco.
I am grateful to John Judd, Brian Motley and Sun Bae Kim for comments. Robert Ingenito provided capable and conscientious research assistance.

This paper presents an update of a simple model for predicting real GDP using contemporaneous monthly data. These forecasts are based on just three variables, all of which are available early in the quarter. The earlier version of this model was used at this bank for more than four years. An analysis of the real-time forecasts made over this period shows that the forecasting errors were reasonable, and that the model's forecasts compare well to the Blue Chip consensus forecasts.

It is easy to appreciate the value of an accurate reading on the current state of the economy during times of great uncertainty. During the first quarter of this year, for example, observers were trying to determine whether an economic recovery would take hold, or whether the economy would slip back into recession. Such a determination is likely to be especially important for monetary policymakers. For instance, information that the economy was contracting over the first quarter could very well have led to a further easing of policy. Yet data on broad measures of economic activity are available only with a lag. In this paper I discuss a method of obtaining estimates of contemporaneous aggregate activity using data that are available with relatively short lags.

Earlier work at this bank showed that reasonable forecasts of real gross national product (GNP) growth in the current quarter could be obtained using a set of only three variables: nonagricultural employment, industrial production, and real retail sales.[1] This paper adds to the earlier analysis in three ways. First, I update the model so that it can be used to predict real gross domestic product (GDP) instead of real GNP. After searching over a list of about a dozen or so variables, I find that the same three variables (with one small modification) still provide reasonably accurate forecasts of real GDP.

Second, I present data on the ex ante forecast accuracy of this model. The model (which I shall refer to as the Monthly Indicators model or MI model below) has been used to predict real GNP at the Federal Reserve Bank of San Francisco for about four years now. The model's (real time) forecasts over this period have been more precise than the Blue Chip consensus forecast (which is the average of the forecasts of roughly fifty leading private sector forecasters).

Last, I look at what the MI model contributes to the accuracy of real GDP forecasts over a time horizon of one to two years. I use a quarterly Bayesian vector autoregression (BVAR) model which forecasts GDP (plus some other variables) to examine this issue. I present evidence which

[1]See Trehan (1989) for a discussion.

suggests that attempts to improve the MI model's forecast of current quarter real GDP growth are unlikely to have large payoffs in terms of forecasting real GDP growth over longer horizons.

The rest of the paper is organized as follows. Section I briefly reviews earlier work on the model and then discusses the process that was used to choose the variables to predict real GDP. Sections II and III present tests of the forecasting accuracy of the model. I present both results for the variables used to predict real GDP (the indicator variables) and the results of predicting real GDP itself. Section III also contains the comparison with the Blue Chip forecasts. Section IV takes up the issue of what the BVAR contributes to real GDP forecasts beyond the one-quarter horizon, mainly to determine the likely benefits of making the current quarter forecast more precise. Section V concludes.

I. Choosing the Indicator Variables

The Original Model

When the model was first specified, three criteria were employed to choose variables that would be used to predict real GNP. The same criteria will be used this time as well. For a variable to be included in our model, the first (and most important) test it must pass is purely statistical: variables will be ranked on the basis of their usefulness in predicting real GDP. Second, in order to limit the costs of collecting and processing data, I also impose the requirement that only a relatively small number of variables be used to predict real GDP. This rules out methods that attempt to predict each (or most) component(s) of GDP in the National Income and Product Accounts (NIPA). Finally, since I am interested in obtaining current quarter real GDP forecasts as early as possible, I impose the requirement that the monthly variables that are to be included be available relatively early.

Based largely on considerations of timeliness, a set of more than a dozen variables was chosen for statistical analysis. These included different measures of interest rates, sales, labor inputs, and so on.[2] I found that reasonable forecasts could be obtained on the basis of three variables: nonagricultural employment, industrial production, and retail sales deflated by the producer price index.

[2]In addition to the three variables included in the model, the list of variables I looked at contains manufacturing shipments and inventories, housing starts, automobile sales, retail sales net of autos, total labor hours, average weekly hours, manufacturing hours, and short-term and long-term interest rates. For a detailed description of the variable selection strategy, see Trehan (1989).

The estimated equation was

$$\underset{(2.2)}{RGNP_t = 0.8} + \underset{(2.8)}{0.17\,IP_t} + \underset{(3.4)}{0.14\,RSALS_t} + \underset{(5.0)}{1.13\,EMP_t}$$

$$- \underset{(3.0)}{0.21\,RGNP_{t-1}} - \underset{(1.4)}{0.09\,RGNP_{t-2}} - \underset{(4.0)}{0.26\,RGNP_{t-3}},$$

adjusted $R^2 = 0.74$, SEE $= 2.17$

where

RGNP = real GNP

IP = industrial production

RSALS = real retail sales

EMP = nonfarm payroll employment,

(all variables are included as annualized growth rates)

The estimation period was 1968.Q2 to 1988.Q2. The absolute value of the *t*-statistics are shown in parentheses.

Even though purely statistical criteria were employed to select indicator variables, the final selection consists of three of the four key variables that the NBER's Business Cycle Dating Committee used to date the beginning of the current recession.[3] Further, employment and industrial production are two of the four series included in the Commerce Department's Index of Coincident Indicators. That index also includes real personal income and real manufacturing and trade sales.[4] Note that the real retail sales variable included in the MI model is similar to the latter variable and has the advantage of being available roughly one month earlier. This similarity to the coincident indicator index suggests that the model should do reasonably well at turning points. (I will return to this issue below.)

Before going further it is also worth noting that data on the variables included in the model become available relatively quickly. Specifically, data for any month are available by the middle of the following month. For example, data for January are available by mid-February.

Updating the Model

An important reason for updating the model has to do with the benchmark revision of the National Income and Product Accounts (NIPA) that was released in early December 1991. Two of the numerous changes introduced as part of that revision are particularly relevant for the purpose of forecasting GDP. First, the Bureau of Economic Analysis announced that it was shifting from the gross

[3]The fourth variable used by the NBER is real income. See Hall (1991-92) for a discussion of how the NBER dates cycles.

[4]See U.S. Department of Commerce (1984) for a discussion.

national product (GNP) to the gross domestic product (GDP) as the primary measure of production. (GNP includes net receipts of factor income from the rest of the world while GDP excludes it.) Second, the base period of the NIPA was shifted from 1982 to 1987. We need to determine whether the MI model has to be respecified because of these changes.

I did make one small change to the original specification before carrying out this analysis. The first time around, the producer price index was used to deflate retail sales instead of the more obvious consumer price index (CPI), because producer prices typically became available more than one week earlier than consumer prices. However, the gap between the release dates of the two series has narrowed over time, and thus it is now possible to employ the CPI to deflate retail sales and produce the forecasts at around the same time as when the PPI was used.

The search for the best specification was carried out in two parts. Starting with a set of 16 variables, I first isolated variables that were useful in explaining within-sample changes in real GDP. Several alternative statistical criteria were used to help determine the best set of variables.[5] At the end of this procedure I ended up with a set of variables that included the three variables in the original monthly indicators model as well as average weekly hours worked and the 10-year Treasury bond rate. In the second part of this procedure the "out-of-sample" forecasts obtained from this set of variables were compared to the out-of-sample forecasts obtained from the set of variables originally included in the MI model. (This procedure involves estimating the real GDP equation up to a given quarter and using the indicator variables to predict real GDP the following quarter. I used a sample of more than 40 forecasts to carry out this comparison.) It turns out that this larger set of variables does not provide forecasts that are noticeably different from the three variables originally included in the equation. Consequently, I decided not to alter the specification of the original monthly indicators model.

Thus, nonfarm payroll employment, industrial production and real retail sales (which is obtained by deflating nominal retail sales by the consumer price index) turn out to provide reasonably good forecasts of real GDP as well as of real GNP.

[5]These included using the "general-to-specific" strategy recommended by David Hendry (see Hendry and Mizon 1978, for example) as well as the "Final Prediction Error" criterion (see Judge, et al. 1985 for a description) to determine which variables and lag lengths were to be included. Some judgment was also involved; for instance, a variable for which a mechanical procedure included the second lag but not the contemporaneous term was dropped.

The estimated equation is

$$\underset{}{RGDP_t} = \underset{(3.9)}{1.1} + \underset{(4.1)}{0.20\,IP_t} + \underset{(5.0)}{0.16\,RSALS_t} + \underset{(5.5)}{0.96\,EMP_t}$$

$$- \underset{(3.4)}{0.20\,RGDP_{t-1}} - \underset{(1.8)}{0.10\,RGDP_{t-2}} - \underset{(4.7)}{0.26\,RGDP_{t-3}},$$

adjusted $R^2 = 0.79$, SEE $= 1.80$.

where
RGDP = real GDP
IP = industrial production
RSALS = real retail sales
EMP = nonfarm payroll employment,
(all variables are included as annualized growth rates)

The absolute value of the t-statistics are shown in parentheses. The equation has been estimated over the period 1968.Q2 to 1991.Q2; as before, the starting date is determined by the availability of the retail sales data. The Lagrange multiplier test for first order serial correlation leads to a test statistic of 0.5, which has a marginal significance level of 50 percent. Thus, it appears that the inclusion of the lagged real GDP terms is sufficient to eliminate serial correlation.

It is worth noting that the new equation is not very different from the original one, despite definitional changes in the dependent variable (specifically, the use of real GDP instead of real GNP, as well as the change in the base year) and in one of the explanatory variables (specifically, the use of the CPI to deflate retail sales instead of the PPI).

It also is tempting to speculate about why the lagged real GDP terms are significant in the estimated equation. One reason that comes to mind is the role played by inventories. This conjecture can be verified by subtracting changes in inventories from real GDP and re-estimating the above equation in terms of final sales:

$$RFSAL_t = \underset{(4.7)}{1.2} + \underset{(0.6)}{0.03\,IP_t} + \underset{(10.9)}{0.30\,RSALS_t} + \underset{(3.5)}{0.55\,EMP_t}$$

$$- \underset{(1.7)}{0.12\,RFSAL_{t-1}} - \underset{(1.0)}{0.06\,RFSAL_{t-2}} - \underset{(0.1)}{0.01\,RFSAL_{t-3}},$$

adjusted $R^2 = 0.76$, SEE $= 1.52$.

where
RFSAL = real final sales.

Note that lags of the dependent variable are significantly less important than in the GDP equation; in fact, the F(3,87) statistic for the null hypothesis that the three lagged RFSAL terms are zero is 1.2, so that the null hypothesis cannot be rejected at any reasonable significance level. This suggests that the lagged RGDP terms in the RGDP equation are capturing the effects of inventory adjustments.

These results suggest that including inventory data may help to make the forecast more precise. Unfortunately, inventory data are released rather late to be useful in this forecast. The lag for data on nominal magnitudes is about two months, while the lag for data on the appropriate deflators is even longer.

II. Predicting the Indicator Variables

Since the forecaster (or policymaker) is likely to be interested in obtaining real GDP forecasts even before three months of information on the indicator variables becomes available, it is necessary to have a method for predicting the monthly values of the indicator variables themselves. I estimate a Bayesian vector autoregression (BVAR) to obtain these forecasts. A vector autoregression (VAR) involves regressing each of a set of variables on lagged values of all variables in the system. Estimating a BVAR implies imposing priors so that the resulting coefficients are a mixture of the coefficients that would be obtained from an unrestricted VAR and the forecaster's prior beliefs. The prior employed here has been termed the "Minnesota prior;" it imposes the belief that most economic time series behave like random walks with drift. For each variable the coefficient on its own first lag is pushed towards one, while the coefficients on all other right-hand-side variables are pushed towards zero. How much should the estimated coefficients be pushed towards this prior? Answering this question involves estimating different versions that vary in how tightly the prior is imposed. The forecasting performance of these different versions is then compared, and the specification that leads to the best forecasts is chosen.[6]

Searching for the best specification to forecast the indicator variables led to a BVAR with five variables: the three indicator variables themselves, plus the interest rate on six-month commercial paper and the average weekly hours of production workers on private, nonagricultural payrolls. Each equation contains 12 lags of each of the variables plus a constant. Since interpreting this many coefficients would be a difficult task, the estimated equations are not presented here. Instead, Table 1 shows cumulative errors from the BVAR over horizons from one

[6]See Todd (1984) for a discussion of Bayesian vector autoregressions.

Table 1

Predicting the Indicator Variables: January 1981–June 1991

Months Ahead	Mean Error	Mean Absolute Error	Root Mean Square Error	Theil's *U*-Statistic[a]
		Nonfarm Payroll Employment		
1	−.06	1.45	2.19	.73
2	−.08	1.15	1.54	.91
3	−.10	1.12	1.43	1.03
		Industrial Production		
1	.79	6.04	8.42	.78
2	.67	4.75	6.17	.82
3	.55	3.10	4.05	.78
		Real Retail Sales		
1	.83	12.54	17.74	.57
2	−.10	7.41	9.80	.58
3	−.23	5.30	7.04	.61

Note: Growth rates are annualized. The errors shown here are cumulative. For instance, the mean error three months ahead is the error in predicting the annualized growth rate between today and three months into the future.

[a]This is the ratio of the RMSE of the model forecast to the RMSE of the naive forecast of no change in growth rates.

to three months; the errors are measured as annualized growth rates.

The sample period covers slightly more than ten years, extending from January 1981 to June 1991, a total of 126 forecasts. For each forecast, the BVAR is estimated up to the prior month and then used to forecast the next three months. For example, for the first forecast the model is estimated through December 1980 and is used to generate forecasts over the January-March period. Next time around the model is estimated through January 1981 and forecasts are generated over the February-April period. Four different measures of forecast accuracy are presented in Table 1: the mean error (ME), the mean absolute error (MAE), the root mean square error (RMSE) and Theil's U-statistic (which compares the RMSE of the model forecast with the RMSE of the naive forecast of no change).

Note that the errors get smaller as the forecast horizon lengthens, a result consistent with the presence of substantial negative serial correlation in the monthly errors. As may be expected, the differences in the size of the errors reflect differences in volatility among the variables; for instance, the standard deviation of the month-to-month growth rates (over the 1981.M1–1991.M6 period) of the employment variable is 2.8 percent, that of industrial production is more than three times as much, and that of real retail sales is roughly seven times as much. The Theil statistics show that the model outperforms the naive forecast by a greater margin when predicting real retail sales than when predicting either industrial production or nonfarm payroll employment.

The errors from the BVAR are smaller than those obtained from univariate autoregressive equations for the same variables, although the differences are not large. Averaging across the three variables, the errors from univariate AR equations are roughly 5 percent larger than those from the BVAR at the one-month horizon and roughly 10 percent larger at the three-month horizon.

III. Predicting Real GDP

Error statistics for the real GDP forecast are shown in Table 2. The full sample period runs from 1981.Q1 to

Table 2

Real GDP Forecast Errors from Monthly Indicators Model

Month of Forecast[a]	Mean Error	Mean Absolute Error	Root Mean Square Error	Theil's U-Statistic[b]
		Full Sample 1981.Q1–1991.Q2 (42 forecasts)		
1	.25	2.08	2.60	.78
2	.27	1.56	1.92	.58
3	.26	1.13	1.59	.48
4	.26	1.11	1.54	.46
		Subsample 1981.Q1–1986.Q1 (21 forecasts)		
1	.52	2.39	2.86	.70
2	.60	1.77	2.17	.53
3	.70	1.45	1.88	.46
4	.75	1.44	1.89	.46
		Subsample 1986.Q2–1991.Q2 (21 forecasts)		
1	−.02	1.77	2.32	1.02
2	−.06	1.34	1.63	.71
3	−.18	0.81	1.22	.53
4	−.23	0.78	1.06	.47

Note: Growth rates are annualized.

[a]These dates refer to the month of the quarter in which the forecast becomes available. The fourth month is the month after the quarter ends. Each forecast is based on complete data for the previous month.

[b]This is the ratio of the RMSE of the model forecast to the RMSE of the naive forecast of no change in growth rates.

1991.Q2, a total of 42 forecasts. In addition, I also show the results for the two halves of the sample period, that is, for the subperiods 1981.Q1–1986.Q1 and 1986.Q2–1991.Q2. For each forecast, the GDP equation is estimated up to the previous quarter, and the resulting coefficients are used, together with the current quarter values of the indicator variables, to predict real GDP growth in that quarter.

Four different exercises were performed for each sample period to duplicate the amount of information available over the course of the quarter. The first one tests the forecasting capabilities of the model during the first month of each quarter, when no information is available on the indicator variables. In this case, the BVAR forecasts the values of the indicator variables for all three months of the quarter, and these values are used in the GDP equation to forecast GDP growth. The second assumes that we are in the second month of the quarter, when we have one month of data on the indicator variables, and the BVAR is used to forecast the values of the indicator variables for the remaining two months of the quarter. Similarly, the third set of GDP forecasts is based on two months of data for the quarter, and the BVAR is used to forecast the values of the indicator variables in the third month of the quarter. Finally, the fourth set is based on all three months of actual data for the indicator variables, so that no BVAR forecast is required to predict GDP growth.

Table 2 reveals that the monthly indicators model does not do a very good job of predicting real GDP growth when it has no information about the current quarter. Indeed, for the 1986.Q2–1991.Q2 subsample, the RMSE of the monthly indicators model is slightly larger than the RMSE of the forecast that is based on the simple rule that the rate of real GDP growth this quarter will be the same as it was last quarter (which is why the computed U-statistic is slightly greater than 1).

The model's forecasts become noticeably more precise in the second month of the quarter, that is, once information about the first month of the quarter becomes available. For the full sample, both the MAE and the RMSE fall by around 25 percent. The arrival of the second month of information leads to some further improvement in the forecast.[7]

In comparing the two subsamples, note that while the RMSEs of the first half (that is, the 1981.Q1–1986.Q1 period) are larger than those for the second half (the 1986.Q2–1991.Q2 period), the reverse is true for the U-statistics. This finding suggests that real GDP was more volatile in the first subsample than in the second. Indeed, this conjecture is confirmed by the data, which show that the standard deviation of quarterly real GDP growth fell by more than 50 percent, from 4.1 percent over the 1981.Q1–1986.Q1 period to 1.9 percent over the 1986.Q2-1991.Q2 period.

Real-time versus Final Data

It is possible that the results presented in Table 2 exaggerate the precision of the BVAR forecast, since they are based upon better data than would be available for use in forecasts made in real time. While it is not possible to overcome this problem completely, some information on the model's performance can be obtained from the real time forecasts of real *GNP* that have been made over the past four years. Specifically, we have compiled data on the original model's forecasts since the model began forecasting in 1987.Q3, which gives us a total of 16 forecasts to analyze.

The results of this analysis are shown in Table 3. To provide some sense of the model's relative performance, the table also includes data on the forecasting performance of the consensus real GNP forecast from the Blue Chip Survey. These data are taken from a newsletter titled *Blue Chip Economic Indicators* published by Capitol Publications. This well-known consensus forecast is the average of the individual forecasts of about 50 major forecasters in the private sector.

It needs to be pointed out that it is difficult to line up the two forecasts so that the two are based upon the same amount of information. The Blue Chip forecasts have been dated on the basis of the month in which they are released. For instance, the second quarter Blue Chip forecast released on the June 10 is compared to the model forecast available on June 15. Thus, the Blue Chip forecast will be based on less information than the MI forecast. Further, while the official release date of the Blue Chip survey is the 10th of the month, the survey itself is conducted over the first week of the month. Of the three indicator variables used in the real GDP equation, the only variable likely to be available at that time is payroll employment.[8]

One way to overcome this problem is to compare the model forecast in a given row with the Blue Chip forecast in the following row. Note that such a comparison will tend to overcompensate in those months when employment data for the previous month are released before the survey is

[7]In the original version of the model the arrival of the second month of information did not lead to a reduction in the model's forecast errors. This is reflected in the results of real time forecasting shown in Table 3. The reasons behind this change are not obvious, although experimentation suggests that the change in results has to do with the change in base years and not the change from GNP to GDP.

[8]Forecasters will also know interest rates and labor hours.

Table 3

Comparison of Real GNP Forecast Errors, 1987.Q3–1991.Q2

Month of Forecast[a]	Monthly Indicators Model Forecasts: Mean Error	Mean Abs. Error	Root Mean Sq. Error	Theil's *U*-Statistic	Blue Chip Forecasts: Mean Error	Mean Abs. Error	Root Mean Sq. Error	Theil's *U*-Statistic
				Using Real-Time GNP				
2	0.11	0.71	0.90	0.60	0.42	1.13	1.49	0.99
3	0.21	0.86	1.09	0.72	0.36	0.99	1.30	0.86
4	0.14	0.79	1.02	0.68	0.34	0.90	1.14	0.76
				Using Revised GNP				
2	0.14	1.01	1.34	0.86	0.46	1.37	1.99	1.28
3	0.24	1.06	1.48	0.95	0.39	1.26	1.83	1.17
4	0.18	1.05	1.38	0.88	0.38	1.19	1.67	1.07

Note: Growth rates are annualized.

[a]These dates refer to the month of the quarter in which the forecast becomes available. The fourth month is the month after the quarter ends. This dating convention implies that the model forecast may be based on as much as one month of additional information compared to the Blue Chip forecast. See text for details.

conducted. (Employment data for a particular month are usually released on the first Friday of the following month.)

The top half of the table compares both sets of forecasts with "early" GNP data. These early GNP data have been obtained from the Commerce Department's *Survey of Current Business* four months after the end of the quarter. The idea is to reproduce, as closely as possible, the GNP data as it existed when the forecasts were made. The results for the monthly indicators model show that the MAEs average around 0.8 percent, regardless of whether we have one, two, or three months of data on hand. Similarly, the RMSEs are around 1.0 percent.

The results for the Blue Chip consensus forecast show that the MAE varies around 1 percent depending upon the amount of information available, while the RMSE falls from around 1.5 percent for the forecast made in month 2 of the quarter to approximately 1.1 percent for the forecast made in the month after the quarter has ended. While these errors are not that much larger than those of the monthly indicators model, it is worth pointing out that the MI forecasts made in the second month of the quarter (that is, forecasts that are based on one month of information) are more accurate than the Blue Chip consensus forecast made after the quarter has ended (month 4). The Theil statistics show that both sets of forecasts do better than the naive forecast of no change in growth rates.

The second half of the table compares the two forecasts to revised real GNP data. Specifically, the two forecasts are compared to real GNP data as of the fourth quarter of 1991. Note that this increases the forecast errors of both models; the deterioration is especially noticeable in the case of the Blue Chip forecast since it does worse than the simple prediction that real GNP growth this quarter will be the same as it was last quarter.

Chart 1 plots the MI and Blue Chip forecasts as well as early GNP data over this period. The top panel of the chart shows forecasts based on one month of information, while the lower panel shows forecasts based on three months of information. Note that the MI forecast tracks the recession quite well, a result that is not surprising since the forecasts are based on information about the current quarter. Recall also that the set of indicator variables is close to the set of variables included in the Index of Coincident Indicators. Finally, as the results in Table 3 would suggest, while the MI forecasts are more accurate on average than the Blue Chip forecasts, this is not always the case.

Before going further it needs to be pointed out that the Blue Chip consensus forecast has been used only as a benchmark (since it is widely available), and not because it is taken to be the most accurate forecast of real activity in the current quarter. In fact, it is not unreasonable to believe that the forecasters included in the panel were trying to minimize their forecast errors over a time span of a year or so instead of a quarter. In that context, it is useful to ask

Chart 1
Real Time Forecasts of Real GNP Growth

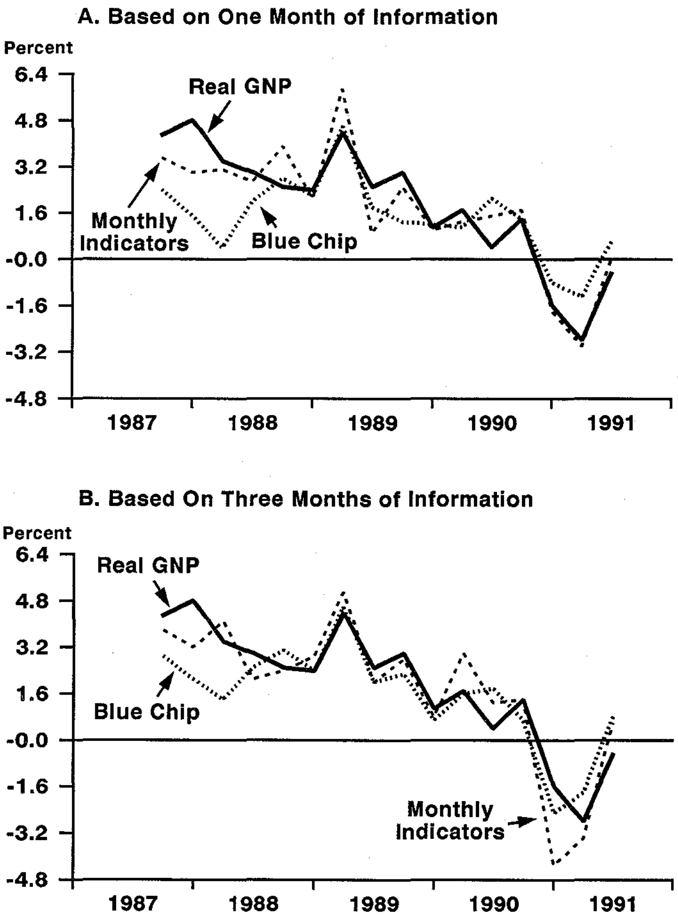

Note: Real GNP data have been taken from the Commerce Department's *Survey of Current Business* four months following the end of each quarter.

what the monthly indicators model contributes to the accuracy of real GDP forecasts beyond the current quarter. We examine this question in the next section.

IV. Evaluating the Usefulness of the Current Quarter GDP Forecast

Usually the forecaster (or policymaker) is interested not just in the forecast of real GDP growth this quarter, but in growth over some longer time period, such as a year or two. It is, therefore, natural to ask what the monthly indicator model's forecast contributes to predicting real GDP over somewhat longer horizons. Perhaps a more important issue for the project at hand concerns the payoff to making the MI forecast more precise. As discussed above, the MI model is a simple one; adding greater detail could improve its accuracy somewhat, especially late in the quarter when more information becomes available. However, greater detail also implies greater cost. Thus, we need to compare the benefits to greater accuracy with the costs of putting together and maintaining a more detailed model.

In the present context (where we are interested in looking at contributions to forecast accuracy over horizons of one to two years), a measure of the benefits can be obtained by examining how the accuracy of real GDP forecasts over one to two years is affected as we increase the accuracy of the current quarter forecast. Here I will make an extreme assumption about how much more accurate the current quarter forecast can be: I will assume that real GDP this quarter is known with certainty.

Forecasts over a two-year horizon will be generated using a BVAR model that is similar to one used for forecasting at the Federal Reserve Bank of San Francisco. This model is estimated on quarterly data (and I will refer to it as the quarterly BVAR). It contains a total of ten variables, including real GDP, consumption, unemployment, the dollar, a measure of money, measures of short and long-term interest rates, and inflation.

Chart 2 plots the percentage reduction in the RMSE of the GDP forecast from the quarterly BVAR when the MI forecast for the first period is included or when the actual value of GDP for the first quarter is included.[9] I show forecasts for an eight-quarter horizon over the 1981.Q1–1991.Q2 period. The errors are cumulative; that is, the RMSE of the four-quarter ahead forecast measures the errors in predicting the level of real GDP four quarters in the future.

Including the MI forecast reduces the RMSE of the one-quarter ahead forecast by about 50 percent and the two-quarter ahead forecast by 35 percent (compared to the case when the MI forecast is not included). The degree of improvement becomes smaller as the forecasting horizon lengthens, falling to less than 15 percent after four quarters and to less than 5 percent in the seventh and eighth quarters.

The degree of improvement we obtain is, of course, dependent upon the model that is being used to forecast real output over the next two years. However, the question of whether the returns to making the MI forecast more precise are worth the effort can be answered in a way that is less model-dependent. We begin by looking at how much the forecast from the quarterly BVAR can be improved

[9]The first quarter here is actually the quarter for which we already have data for the indicator variables. This was termed the contemporaneous quarter in Sections I-III. The change in terminology is necessitated by the introduction of the quarterly BVAR, which contains no contemporaneous information. Note also that the MI forecasts used here are based on three months of information.

Chart 2
Percentage Reduction in the RMSE of the GDP Forecast

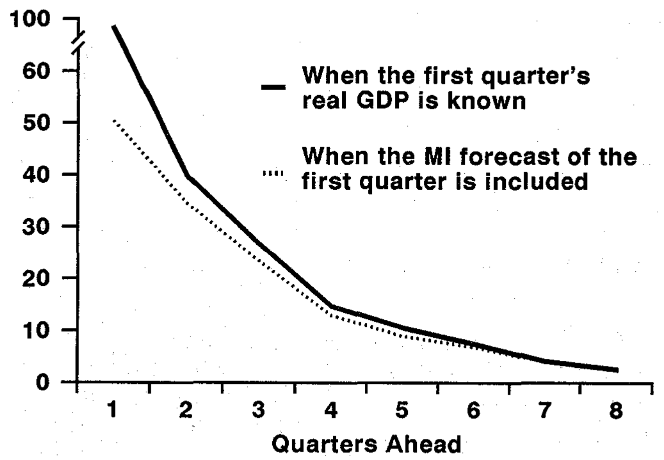

when next quarter's real GDP is assumed to be known, that is, assuming perfect information.

Perfect information implies that the first quarter value of this number is 100 by assumption. More interestingly, knowledge of the first quarter's real GDP reduces the RMSE of the two-quarter ahead forecast by 40 percent and the RMSE of the four-quarter ahead forecast by about 15 percent.

As before, the precise effects of including information about next quarter on the one-year ahead forecast are likely to depend upon the model that is being used, since models differ in their ability to process information about the next—or any other—quarter's GDP. Nevertheless, it is possible to compare the marginal benefit of moving from the no information case to the case where the MI forecast is known to the marginal benefit of moving from knowledge of the MI forecast to knowledge of next quarter's real GDP. (Recall that this is a theoretical upper bound to further improvements in the MI forecast.) Chart 2 provides a simple way of making the comparison. At each point in time, the marginal benefit of moving from the no information case to the case where MI is known is given by the vertical distance between the horizontal axis and the MI line; the marginal benefit of moving from knowledge of MI to perfect information is measured by the vertical distance between the two curves. The greater the difference between the two curves relative to the height of the MI curve, the greater the advantage to improving upon the MI forecast.

The chart indicates that at a two-quarter horizon, the relative improvement in going from the no information case to including the MI model forecast is substantially greater than the relative improvement in going from the MI model to the perfect information case. This continues to be the case at all forecast horizons; in fact, the difference between the two curves is essentially zero from the fourth quarter on. Of course, both curves are close to zero towards the end of the forecast horizon and the difference between them at that point is not very significant.

This exercise suggests that further attempts at improving the current quarter forecast of real GDP are not likely to have substantial rewards in terms of improving our ability to forecast real GDP over somewhat longer horizons. In other words, if the objective is to forecast real GDP beyond the first two quarters, then the simple MI model reaps a large proportion of the gains that would accrue in going from the case of no information about the first quarter's real GDP to the case where the first quarter's real GDP is known with certainty, and does so at relatively little cost.

V. *Summary and Conclusions*

This paper has reviewed a simple method of predicting real GDP. This method requires relatively few resources; the forecasts are cheap to produce and update. The evidence presented above demonstrates that these forecasts compare well to those obtained from major private sector forecasters.

It is possible that the forecast of current quarter real GDP growth could be made more precise by devoting additional resources to the task. However, the evidence presented above also suggests that, if the objective is to forecast real GDP beyond the current quarter, then such an endeavor is likely to lead to relatively limited returns.

References

Hall, Robert E. 1991-1992. "The Business Cycle Dating Process." *NBER Reporter* (Winter) pp. 2-3.

Hendry, D.F., and Mizon, G.E. 1978. "Serial Correlation as a Convenient Simplification, Not a Nuisance: A Comment on a Study by the Bank of England." *Economic Journal* pp. 549-563.

Judge, George G., W.E. Griffiths, R. Carter Hill, Helmut Lutkepohl, and Tsoung-Chao Lee. 1985. *The Theory and Practice of Econometrics*. New York: John Wiley & Sons.

Todd, Richard M. 1984. "Improving Economic Forecasting with Bayesian Vector Autoregressions." Federal Reserve Bank of Minneapolis *Quarterly Review* (Fall).

Trehan, Bharat. 1989. "Forecasting Growth in Current Quarter Real GNP." Federal Reserve Bank of San Francisco *Economic Review* (Winter) pp. 39-52.

U.S. Department of Commerce. 1984. *Handbook of Cyclical Indicators*.