



WORKING PAPERS

RESEARCH DEPARTMENT

**WORKING PAPER NO. 02-15/R
FORECASTING COIN DEMAND**

Dean Croushore and Tom Stark
Federal Reserve Bank of Philadelphia

March 2003

FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • (215) 574-6428 • www.phil.frb.org

WORKING PAPER NO. 02-15/R

FORECASTING COIN DEMAND

Dean Croushore and Tom Stark

Federal Reserve Bank of Philadelphia

March 2003

The views expressed in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. John Chew, Brian DiCiurcio, Jim Gillard, and Jim Sherma provided outstanding research assistance on this project. Thanks to Geoff Gerdes, David Griffiths, Chris Sims, and Dennis Smock, as well as participants in a seminar at the Federal Reserve Bank of Philadelphia, for helpful suggestions. Please send comments to the authors at Research Department, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106-1574, or e-mail: dean.croushore@phil.frb.org.

FORECASTING COIN DEMAND

ABSTRACT

Shortages of coins in 1999 and 2000 motivated us to develop models for forecasting coin demand. A variety of models were developed, tested, and used in real-time forecasting. This paper describes the models that were developed and examines the forecast errors from the models both in quasi-ex-ante forecasting exercises and in real-time use. Tests for forecast efficiency are run on each model. Real-time forecasts are examined. We conclude with suggestions for further refinements of the models.

FORECASTING COIN DEMAND

In early 1999, the demand for pennies increased sharply. At the same time, the U.S. government introduced a program to produce five new quarter-dollar coins each year, each depicting information about one of the states in the union. Though demand for the first several quarters issued was relatively small and not far from what the government expected, the popularity of these state quarters grew over time. By the second half of 1999 and throughout 2000, the demand for the state quarters grew so much that the U.S. Mint was unable to produce enough coins fast enough to meet demand; as a result, shortages of quarters developed. Once those shortages began, people began using other denominations (mainly dimes and nickels) instead of quarters in change, so demand for those coins went up as well. Adding to the problems of inadequate production was the government's decision to produce a new golden dollar coin; the Mint produced 1.2 billion of the coins by September 2000.

Given the sharp increases in coin demand in 1999 and 2000, and the difficulty in getting enough coins of various denominations into the hands of the public, the Federal Reserve Bank of Philadelphia began in 2000 to investigate how to improve forecasts of demand for coin. The idea was that if the Fed had its own projection of the demand for coins, it could work with the Mint to ensure an adequate supply.

This paper describes the development of coin forecasting models. We begin in section 1 by discussing the historical data on demand for coins. Then, in section 2, we discuss the development of four different models of coin demand: a structural model, a time-series model, a VAR model, and a Bayesian VAR model. In section 3, we illustrate the real-time forecasts produced by the models, although a complete analysis will not be possible until we have a longer

history of forecasts and actuals. Section 4 concludes the paper by discussing potential refinements of our models and methods.

1. Coin Demand: History and Definitions

An important institutional constraint in determining the amount of coins that circulate is that the U.S. Mint controls production of coins, while the Federal Reserve System distributes them around the country. Unless both agencies are in synch, problems with coin production and delivery may arise. The Fed distributes coins from 37 coin offices, which include mostly the Reserve Banks and their branches and through over 100 coin terminals with armored carriers.

Developing a forecasting model involves several steps: choosing among different types of models, testing the different models to see how well they perform, then running forecasts in real time and investigating the quality of the forecasts. Because the demand for each coin denomination seems to behave differently than other coin denominations, the models that we examine will contain a separate forecasting equation for each denomination, rather than modeling overall coin demand in a single equation.

The first difficulty we faced in forecasting coin demand was a lack of data. The Fed's database contained national data from 1989 to 1999 on coin demand, but coin demand has huge monthly seasonal fluctuations, so the data amounted to roughly 11 observations. We could not proceed further until we obtained a data set from the U.S. Mint containing data back to 1957.

We measure the flow demand for coins by net pay, as data on stocks were unavailable. Net pay is an unusual economic concept because it represents the change in the demand for coins, which can be either positive or negative. Net pay equals the payout of coins from the Mint, Fed, and armored-carrier terminals to financial intermediaries minus the inflow of coins to

the Fed and armored-carrier terminals from financial intermediaries (hereafter called banks). Net pay is positive when banks ask the Federal Reserve to deliver coins to them. But if banking customers decide to return many coins to their banks, the banks in turn may send them back to the Federal Reserve, resulting in negative net pay. Because coins are durable and can be returned to the producer, coins are different from all other goods. Note that net pay differs from the change in the demand for coins in the rare instance in which there is a shortage of coins. Because such cases are rare and we do not have reliable data on the amount of shortages, our models of coins ignore such instances, and we proceed by assuming that net pay equals the change in demand.

Figure 1 illustrates the flows to and from banks. Coins flow between people or businesses and banks. When people and businesses want more coins, their banks order more coins from their local Federal Reserve offices. The Federal Reserve offices then ship coins to the banks, either from inventories of coins located at armored carriers (the line labeled *A* in the chart) or directly from the offices' own inventories (line *B*). Occasionally, the U.S. Mint ships coins directly to a bank (line *C*). Each of these shipment methods to banks represents payout of coins, which increases net pay.

However, if people begin turning in more coins to banks than the banks want to keep on hand, the banks may return the extra to the Federal Reserve, either directly (line *D*) or through armored carriers (line *E*). So, lines *D* and *E* represent inflows of coins, which decrease net pay.

In any given month, some banks may have positive net pay and others may have negative net pay, so total net pay could turn out to be positive or negative. Net pay is the payout minus inflow, so it is the sum of the amounts shown by lines *A*, *B*, and *C* minus the sum of the amounts shown by lines *D* and *E*:

$$\begin{aligned} \text{Net pay} &= \text{payout} - \text{inflow} \\ &= (A + B + C) - (D + E). \end{aligned} \tag{1}$$

There are six different denominations of coins (penny, nickel, dime, quarter, half dollar, and dollar) and net pay is calculated for each separately.

The data on net pay for each denomination show significant changes over time. We will look at the net pay for each denomination in units of millions of coins each month (Figures 2a to 2f). The data run from January 1957 to February 2003 for pennies, nickels, dimes, and quarters, from July 1960 to February 2003 for half dollars, and from January 1995 to February 2003 for dollars. For each denomination, the thin line shows the actual monthly value of national net pay (summed across all Federal Reserve coin offices), and the thick line is a 12-month moving average, which is shown to help illustrate the long-term trend in the data.

The charts show that month-to-month seasonal fluctuations in net pay are huge. For some coin denominations, national net pay is even negative in some months. The net pay of different denominations swings dramatically from one month to the next, mostly because of changes in people's spending patterns. People need more change in the summer months for parking meters at the shore and soda machines. They use more change around holidays at the end of the year, as well. But people need much less in the middle of the winter.

Over half of all coins produced recently are pennies. The net pay of pennies has averaged over 800 million coins per month in the last five years, while the sum of all other denominations has been about 700 million coins per month. The long-run trend in net pay for pennies has been fairly constant since 1980 or perhaps declining slightly. For other main denominations (nickels, dimes, and quarters), the trend over time is slightly upward, which suggests that they may be replacing pennies gradually in terms of quantity used for making

change. There are some interesting variations in net pay for those coins, especially in the 1960s when the value of silver, which was a major component of dimes and quarters, increased sharply. Demand for those coins declined sharply when the coins were redesigned with no silver in them. The other two denominations, half-dollars and dollars, had net pay near zero for much of the 1980s and 1990s. The introduction of the Sacagawea golden dollar coin in 2000 caused a sharp increase in the net pay of dollars in that year.

Given these trends in the demand for different denominations, what can we say about the overall demand for coins? To investigate this issue, we examine the total demand for coins in terms of numbers of coins, adding up the net pay for all six denominations. (From 1957 to mid-1960, we have no data on half-dollars and dollar coins; from mid-1960 through 1994 we have no data on dollar coins. However, we believe that demand is small for these denominations in the periods in which we are missing data, except for a short period in 1979 when the Susan B. Anthony dollar coin was first introduced.) Also, to avoid confusion arising from the seasonal fluctuations, we just look at the total net pay in each calendar year (Figure 3). We look at the total number of coins in billions per year, rather than their dollar value, in part because the Mint's ability to produce enough coins to meet demand depends more closely on the number of coins than on their dollar value.

In the graph, you can see that overall net pay generally increased over time from under 2 billion coins in 1957 to a peak of 23 billion in 1999 and 2000. But the increase was not steady. From one year to the next, net pay sometimes rose and sometimes fell.

We might expect a correlation between net pay and the strength of the economy because people seem likely to use more coins if they buy more goods and services. Looking at Figure 3, you can see that in some years net pay falls when the economy weakens, as in 1990 and 1991.

But in other years, such as 1996 and 1997, net pay falls even when the economy is strengthening. So, there does not appear to be a strong correlation between net pay and economic activity.

Special events caused net pay to be at very high levels in 1999 and 2000. Penny demand was very high in early 1999 (Mullinix, 1999). Also, beginning in 1999, the Mint (directed by laws passed by Congress) rolled out the first quarters in the state commemorative program. The demand for these new quarters turned out to be significantly stronger than was anticipated, thus causing net pay to rise sharply, as you can see in Figures 2d and 3. Then, in 2000, the Sacagawea dollar coin was introduced to much fanfare (Roseman, 2002). Initial demand for the new coin was also strong, and the Mint produced over 1 billion of them (United States Mint, 2001). At the same time, the demand for the new state quarters increased 50 percent from the year before, so again net pay was much higher than expected. With people collecting both the new dollar and the new quarters, fewer of these coins circulated for use as change, so the demand for pennies, nickels, and dimes also rose substantially, as Figures 2a to 2c show.

2. Models of Coin Demand

We initially considered three different types of models: (1) a structural model; (2) a time-series model; and (3) a vector-autoregression model. Descriptions of each model follow, after which we examine tests of forecasting ability for all the models. In determining the exact structure to use with each model, we ran diagnostic tests of various types, then determined the choice of alternative proxy variables by running quasi-ex-ante forecasting exercises and examining simulated forecast performance in the 1990s. Because the models were estimated too soon after the state quarters and golden dollars were introduced, we excluded data from 1999 and

2000 from the sample period. Later, we developed a fourth model, a Bayesian vector-autoregression model, testing it in a similar fashion and adding it to our stable of models.

Structural Model. We began our work on forecasting coin demand using a structural model based on the work of earlier researchers who had modeled coin demand. The payout and inflow of coins in each denomination are modeled separately. To relate the coin data to economic data, we use quarterly averaged data. Economic theory suggests that the demand for money depends mainly on economic activity, interest rates, and the inflation rate. Assuming that the demand for coins by denomination should be similar to overall money demand, we use those variables as the main explanatory variables for money demand.

The equations of the model are:

$$P_t^d = \alpha_1 + \beta_1^d Y_t + \gamma_1^d i_t + \lambda_1^d \pi_t, \quad (2)$$

$$I_t^d = \alpha_2 + \beta_2^d Y_t + \gamma_2^d i_t + \lambda_2^d \pi_t, \quad (3)$$

where P_t^d is payout of coins (in millions of coins) at date t for denomination d , I_t^d is the inflow of coins at date t for denomination d , Y_t is a measure of economic activity at date t , i_t is an interest rate at date t , and π_t is the inflation rate at date t .

We used data from the first quarter of 1958 to the second quarter of 2000 to estimate and test the model. Our baseline version of this model was based on an inherited model from previous analysts, so we used its basic structure of variables. We understand that this model was used to project future coin demand, though details are sketchy because personnel have changed. We modified the model slightly because given the data we had (which included a larger sample than the one that had been used in the past), the original model performed poorly. The economic activity measure was nominal personal consumption expenditures on services, deflated by the

CPI, both seasonally adjusted. The interest rate used was a three-quarter moving average of the federal funds rate, lagged four quarters. The inflation rate was the 12-month percentage change in the CPI. The equations also included a constant term and three dummy variables (for quarters one, two, and three) to handle seasonal effects. We call this structural model A.

We modified the model slightly in several alternative versions. Structural model B used the log of the quarterly average of monthly nominal retail sales as the economic activity variable. Structural model C used the log of the quarterly average of monthly industrial production as the economic activity measure and also included the logarithm of the CPI in the regression. Model D used the log of quarterly nominal personal consumption expenditures as the economic activity variable. Model E used the log of quarterly payroll employment as the economic activity variable and included the log of CPI in the regression. Note that whenever a real variable is used on the right-hand side to represent economic activity, we also include the log of the CPI because the demand for a coin denomination may change as the price level changes.

Diagnostic tests on the results of the regressions for equations (2) and (3) for each denomination suggested fairly clean regressions, with no severe problems. So, we proceeded to generate quasi-ex-ante forecasts with the model. The forecast for the inflow is subtracted from the forecast for the payout to generate a forecast for net pay. Because we knew that we would be comparing the structural models with other models that we thought, a priori, were likely to be superior, we gave the structural models a boost by using realized values for the right-hand-side variables in equations (2) and (3). [By contrast, all the other models that we considered had to generate forecasts for the variables in equations (2) and (3).] For each model, we generated forecasts in a rolling fashion in the following set of steps. First, at the forecast date 1990Q1, given coin data through 1989Q4, we generated one-year-ahead forecasts from the estimated

equations (2) and (3) for pennies, nickels, dimes, and quarters.¹ [Halves and dollars were not forecast because of a lack of data.] Then, we stepped forward one quarter to forecast date 1990Q2, using coin data through 1990Q1, to generate another set of one-year-ahead forecasts. We continued this process until forecast date 1998Q1, at which time we generated one-year-ahead forecasts that extended to the end of 1998Q4. We stopped there so that we would not enter the period in which the state quarter program began.

The results of this exercise were used to generate a number of statistics. Because the Mint's main focus is on the total number of coins produced (adding up the number of coins across all denominations), we chose as a forecast criterion the root-mean-squared-forecast error (RMSFE) for net pay of total coins (adding the number for pennies + nickels + dimes + quarters) during the four quarters of the forecast horizon for the 33 forecast periods from 1990Q1 to 1998Q1. The results for each model are shown in Table 1.

As Table 1 indicates, the worst model in terms of RMSFE is the original one based on services consumption. Much better forecasts are obtained using other variables. The use of the log of payroll employment provides the best out-of-sample quasi-real-time forecasts. We ran the forecasts through a batch of diagnostic tests that are fairly standard in the literature, which will be discussed after the other forecasting models are described.

¹ Throughout this paper, whenever we use the term one-year-ahead forecasts, we always mean the cumulative sum of the forecasts for the variable (usually net pay) over the coming year. For example, for the forecast date 1990Q1, which is based on data through 1989Q4, the one-year-ahead forecast is the sum of the forecasts for 1990Q1, 1990Q2, 1990Q3, and 1990Q4.

Time-Series Model. Whereas a structural model is based on economic theory, a time-series model uses only past data and no theory. Yet research has shown that time-series models often yield superior forecasts, especially in the presence of structural change. So, we developed a univariate time-series model, following standard Box-Jenkins (1976) procedures. We used monthly data from January 1960 to October 2000 to identify several possible models. Because of uncertainty about the optimal procedure to handle the strong seasonal effects, we proceeded with two alternative models: Model A first deseasonalizes the data by regressing it on monthly dummy variables, then bases forecasts on a time-series model on the seasonally adjusted data; Model B does not deseasonalize the data but generates forecasts using a time-series model with seasonal differencing.

The models were run on net pay, rather than payouts and inflow separately. The regression model is:

$$a^d(L)a_s^d(L^{12})(1-L)^i(1-L^{12})^j N_t^d = b^d(L)b_s^d(L^{12})\varepsilon_t^d \quad (4)$$

where d is the denomination (pennies, nickels, dimes, or quarters), N_t^d is net pay of denomination d at date t , L is the lag operator, $a^d(L)$ is a polynomial in the lag operator with p autoregressive lags, $a_s^d(L^{12})$ is a polynomial in the seasonal (12-month) lag operator, i is the order of integration, j is the order of seasonal integration, $b^d(L)$ is a polynomial in the lag operator with q moving-average lags, $b_s^d(L^{12})$ is a polynomial in the seasonal (12-month) lag operator, and ε_t^d is a white-noise error term. Model A sets $j = 0$, $a_s^d(L^{12}) = 1$, and $b_s^d(L^{12}) = 1$ but uses seasonally adjusted data on net pay; Model B uses unadjusted data on net pay and does not restrict the model at all.

Tests of the data for Model A showed that $a^d(L)$ should contain about 14 lags (so $p = 14$), $b^d(L) = 1$ (so $q = 0$), and $i = 1$, so the model is described as an ARIMA(14,1,0). Tests of the data for Model B suggested that $a^d(L)$ should contain about 14 lags (so $p = 14$), $b^d(L) = 1$ (so $q = 0$), $i = 1, j = 1$, $a_s^d(L^{12}) = 1$, and $b_s^d(L^{12}) = 1$, so the model is described as an ARIMA(14,1,0) with one seasonal difference.

A comparison across models with quasi-out-of-sample forecasts for the 1990s suggests that Model A is superior to Model B, as Model A's RMSFE was just 2035 compared with 2465 for Model B. For that reason, we chose Model A to use in subsequent analysis.

Vector-Autoregression (VAR) Model. A vector-autoregression (VAR) model is a natural choice for modeling coin demand as it does not impose too much structure on the model but uses data on economic variables to help forecast coin demand. Our prior was that a VAR model would be superior to other approaches we tried. A VAR model (Sims 1980) is useful because it allows us to conduct “what-if” experiments, such as: “What will happen to the demand for coins if the economy goes into a recession?”

For our VAR model, the net pay of each coin denomination depends on lagged data on economic activity, the interest rate, and the inflation rate, which are the same basic variables used in the structural model. The model-evaluation phase required us to make decisions on the choice of variables, how to seasonally adjust the data, how to choose lag length, and whether to exclude some variables in some equations. We tested a number of different economic variables, including industrial production, personal consumption expenditures, retail sales, and payroll employment. We ran the model using both seasonally adjusted data and not seasonally adjusted data, in the latter case adding monthly dummies to account for seasonality. We used AIC and

SIC testing criteria to choose the lag length. And we examined whether the model was better if we used a full VAR or if we excluded some variables from some equations, a near-VAR. In the latter case, we considered whether the coin demand variables should be excluded from the equations containing the macroeconomic variables. We used monthly data from January 1957 to September 2000.

Preliminary examination of the data and models, including statistical tests of exclusion restrictions, helped to reduce the number of models. Those tests were not able to distinguish between alternative economic variables but suggested that the best model was likely one using seasonally unadjusted data combined with monthly dummies. Several choices of lag lengths were considered after AIC and SIC tests were run. The tests suggested that 13 lags were best. But to avoid the likely overfitting that would result, we also considered a much shorter lag structure (2 lags of each variable) and a mixed lag structure (using lags 1, 2, 3, 6, 9, and 12). Tests indicated that models excluding the coin demand variables from the equations for the macro variables would lead to superior forecasts.

We were unable to clearly determine whether macro variables should be entered in log levels or growth rates, so we ran forecasts using both. The results of testing forecasts for the 1990s are shown in Table 2.

Table 2 shows that nearly all the VAR models, except for those with a small number of lags with growth rate variables, lead to similar root-mean-squared forecast errors. The lowest RMSFE was for model A1, using the log of employment, with 13 lags, so we chose that as the best VAR model. We tested the model's residuals for serial correlation and heteroskedasticity, and found no evidence of any problems. We also tested whether a fixed sample size as we

moved through time would be better than always starting the regression equations with a sample beginning in 1962:1, finding that the latter was superior.

Comparing the Models. We used the testing period of 1990 to 1998 to examine the models and choose the best model within each type. We used similar procedures to compare the models across types and also ran a set of diagnostic tests on the best structural model, the best time-series model, and the best VAR model. Each forecasting model was run in a simulated quasi-ex-ante forecasting exercise from the first quarter of 1990 to the first quarter of 1998, with monthly or quarterly forecasts generated for horizons up to two years. We focused on the one-year-horizon results.

Figure 4 shows a plot of the one-year-ahead forecasts. Each date shown on the horizontal axis is the forecast date, and the value shown on the vertical axis is the forecast amount for net pay over the year from the forecast date to 12 months ahead (for the VAR and time-series models, which are monthly) or 4 quarters ahead (for the structural model, which is quarterly), compared with the actual value of net pay over the same period. (Ignore the line labeled BVAR, which will be discussed later.) As might be expected, the forecasts lag the movements in actual coin demand. Both the VAR and time-series models pick up those movements fairly quickly, so they respond over time, with a lag, to changes in actual coin demand. But the structural model is much slower to respond, even though we gave that model an informational advantage by feeding it actual values of economic variables. The VAR and time-series models move in similar ways, but the time-series model seems to respond a bit more, while the VAR forecasts do not change as much over time. Whether that is a good thing or a bad thing can only be judged with statistical tests, to which we turn next.

Comparisons of root-mean-squared forecast errors and test results for a number of statistical tests on the forecasts (sign test, Wilcoxon signed-rank test, zero-mean test, and test for unbiasedness) are shown in Tables 3 to 7. For additional details on these tests, see Croushore (1998). (The tables each include a row labeled BVAR, which will be discussed later.)

Table 3 compares the root-mean-squared forecast error (RMSFE) across the models in two ways. One column (labeled Quarterly RMSFE) shows the RMSFE for the one-year-ahead forecasts made each quarter from 1990Q1 to 1998Q1. These results are potentially misleading because the observations are overlapping—that is, the forecast errors may be correlated because the forecast horizon is longer than the observation interval. As a check on those results, we also calculated the RMSFE (labeled Annual RMSFE) using only the first forecast made each year, based on data through December of the previous year. The two RMSFE calculations yielded the same result: the time-series model was the best of the three models over the testing period, with the VAR coming in a close second. The structural model was clearly the worst model using the RMSFE criterion.

A useful test of forecasts is a sign test on the forecast errors. If a forecast is optimal, the forecast errors should have a zero median. The sign test examines the null hypothesis of independent errors with a zero median by counting the number of positive observations in the sample, which has a binomial distribution. We ran sign tests for the one-year-ahead forecasts made both quarterly and annually. The results, reported in Table 4, show no rejections of the null hypothesis. So, the forecast errors appear to be independent with a zero median.

The Wilcoxon signed-rank test is related to the sign test, since it has the same null hypothesis, but assumes a symmetric distribution. It accounts for the relative sizes of the forecast errors, not just their sign. The test statistic is calculated by taking the absolute values of

the forecast errors, ranking them by size in increasing order, then finding the sum of the ranks for the positive forecast errors. Table 5 reports the results, which show no rejection of the null hypothesis for any of the models. Again, the forecast errors appear to be independent with a zero median.

A simple test that optimal forecasts should pass is that the mean of the forecast errors be zero. This can be tested with a t-test that depends on the sample mean, standard deviation, and number of observations. Table 6 reports the p-values for the t-tests, and again the null hypothesis is not rejected for any of the models. (Because the quarterly forecasts are overlapping observations of one-year-ahead forecasts, the covariance matrix was adjusted using the Newey-West procedure.)

So, the models pass three nonparametric tests on the forecast errors, which suggests that the models are reasonable.

A more stringent test is a test for unbiasedness. A set of forecasts is unbiased if a regression of the actual values (the dependent variable) on a constant term and the forecasted values (the independent variable) yields coefficients that are not significantly different from 0 for the constant term and 1 for the forecast term. The regression is:

$$\pi_t = \alpha + \beta\pi_t^f + \varepsilon_t, \quad (5)$$

where π_t is the actual inflation rate and π_t^f is the forecast at each date t . The test for unbiasedness requires that over a long sample period, $\hat{\alpha}$ should be close to zero and $\hat{\beta}$ should be close to one, if the forecasts are not biased.

Table 7 reports the results of testing this hypothesis. For quarterly forecasts, the null hypothesis that $\hat{\alpha} = 0$ and $\hat{\beta} = 1$ is rejected, but it is not for annual forecasts. (As with the zero-mean test, because of the overlapping observations in the quarterly forecasts, the covariance

matrix was adjusted using the Newey-West procedure.) For annual forecasts, the p-value on the F-test shows no rejection of the null hypothesis, but with only 9 observations, the sample size is so small that rejection was unlikely.

How much should we worry about the rejection of the unbiasedness test for the quarterly forecasts? Particularly worrisome is the estimated β coefficient in the structural model, which is negative, suggesting that the forecasts move in the opposite direction of actual coin demand. The very low values of the \bar{R}^2 statistic for the structural and VAR models suggest that they fit poorly; the time-series model seems to be more consistent with the data. Possibly, the poor fit and rejection of the null hypothesis of unbiasedness arise because the sample is small. But, overall, the findings suggest that we should continue our efforts to improve the models.

The biggest surprise in our tests was how well the time-series model performed relative to the other models. The other two models had the advantage of using data on economic variables that had been revised, rather than the original data that would have been available to forecasters in real time; on the importance of this issue, see Croushore and Stark (2001). That is, the data on employment used in the VAR and the structural model were taken from a database in October 2000 (when these tests were run). The VAR does not use data on employment beyond the forecasting date; the model generates forecasts for employment for dates after the forecasting date, whereas the structural model does not. To give the structural model the benefit of the doubt, we used actual values instead of forecasts for employment in periods after the date at which the forecast was made. Despite this advantage, the model did worse than any other model. And as we began to use the model to forecast in real time, it began producing unreasonable forecasts that were out of line with the other forecasts. Given that the structural model forecasts also tended to move counter to the underlying trend in coin demand (as shown in Figure 4 and in

the negative slope coefficient in Table 7) we decided to drop the structural model and instead to work on a Bayesian VAR, which will be discussed next.

Because the time-series and VAR models passed the tests we posed, we began using them for forecasting coin demand in real time in February 2001, reporting the results to officials in the Fed responsible for coin distribution, first at the Cash Fiscal Product Office at the Philadelphia Fed, then at the Cash Product Office at the Los Angeles Branch of the San Francisco Fed, to which responsibility for coin distribution was moved in spring 2001.

Bayesian VAR. The research on these coin models took place in several stages. As we proceeded to forecast with the time-series and VAR models in February 2001, we also began development work on a new model, a Bayesian VAR (Litterman 1986). Preliminary work on that model was completed in September 2001.

We had expected the VAR model to be superior to the time-series model, but in our quasi-ex-ante forecasting exercise, that proved not to be the case. Most likely, the cause was overfitting the model. But overfitting probably arose because of the extreme seasonality in the data, which required many lags in the VAR. To reduce the problem of overfitting, we chose a Bayesian approach.

The main differences between the VAR and the Bayesian VAR are in the amount of past data used (24 months in the Bayesian version versus 13 months in the non-Bayesian version) and in the coefficients of the forecasting equation, which are fixed in the VAR but allowed to change over time in the Bayesian VAR. In the Bayesian VAR, a number of hyperparameters are chosen to make the model perform well, the most important being those that concern the seasonal patterns in the data.

The model consists of data on seven variables, just as in the VAR: an economic activity variable (the log of not seasonally adjusted payroll employment), a price variable (the not seasonally adjusted CPI inflation rate), an interest rate (the federal funds rate), and four net pay variables for each main coin denomination (pennies, nickels, dimes, and quarters). Seasonal dummies are included in the model, as described below.

Each variable is modeled with the following state-space representation, comprised of an observation equation and a state equation:

$$y_t = Z_t' \beta_t + \varepsilon_t, \quad (6)$$

$$\beta_t = \beta_{t-1} + \nu_t, \quad (7)$$

where y_t is one of the variables in the system; Z_t is a 180×1 vector composed of a constant, 11 seasonal dummies, and 24 lags of the 7 variables in the system; β_t is a 180×1 coefficient vector, ε_t is a shock distributed as: $\varepsilon_t \sim N(0, \sigma^2)$, and ν_t is a shock distributed as: $\nu_t \sim N(0, Q)$, with covariance matrix Q .² With 180 coefficients per equation and seven equations, the model is large. Equation (6) is a regression equation with a time-varying coefficient vector, β_t . Equation (7) shows that the coefficients move as a random walk, which allows for permanent shocks, which may be important for times such as the mid-1960s when the value of silver in coins rose so high that silver coins were driven out of circulation, the similar events in the early 1980s when the value of copper in a penny exceeded one cent and penny demand rose sharply, and changes in seasonality. If $Q = 0$, then there are no permanent shocks

² When we need to differentiate among the σ^2 terms for each equation, we write σ_j^2 .

to coefficients, and the model reverts to one with fixed coefficients. Following standard practice, we restrict the model by assuming that Q is diagonal.

The Kalman filter is used to generate the conditional expectation and variance of the coefficient vector β_t , which we denote $\beta_{t|t}$ and $P_{t|t}$, where:

$$(\beta_t | \Omega_t) \sim N(\beta_{t|t}, P_{t|t}) \text{ and } \Omega_t = (y_t, \dots, y_1, Z_t, \dots, Z_1).$$

With the aid of equation (7), we can find the moments of $(\beta_{t+1} | \Omega_t)$, denoted $\beta_{t+1|t}$ and $P_{t+1|t}$, where:

$$(\beta_{t+1} | \Omega_t) \sim N(\beta_{t+1|t}, P_{t+1|t}).$$

We use the recursive algorithm described in Hamilton (1994) to compute $\beta_{t|t}$, $P_{t|t}$, $\beta_{t+1|t}$, and $P_{t+1|t}$.

Starting values for the recursion are given by the Bayesian prior:

$$\beta_1 \sim N(b, P_{1|0}),$$

where β_1 is the 180 x 1 coefficient vector described above at time $t = 1$, which we assume is distributed normally with mean b and variance $P_{1|0}$.

Our priors are taken from the work of Raynauld and Simonato (1993), who considered three alternative modifications of standard Minnesota priors designed to account for variation in the data at seasonal frequencies. As noted above, we use the modification that Raynauld and Simonato call, “the random walk plus dummies specification.” In our implementation, this modification specifies b and $P_{1|0}$ according to the prior that each variable in the VAR follows the process given by:³

³Raynauld and Simonato (1993) also considered priors governed by: $(1-L)(1-L^{12})y_t = \varepsilon_t$ and $(1-L^{12})y_t = \text{constant} + \varepsilon_t$, where L is the lag operator. The first imposes two unit roots at the zero frequency and seasonal unit roots. The second imposes a single unit root at the zero frequency and seasonal unit roots. Our baseline specification assumes a zero-frequency unit root and no seasonal unit roots.

$$y_t = y_{t-1} + \text{constant} + 11 \text{ seasonal dummies.}$$

Thus, b is a 180×1 vector of zeros in all positions except the position corresponding to the own first lag, which is unity. Although the constant and seasonal dummy coefficients are important parts of the specification, given the high degree of seasonality in the coin data, we specify their prior values (somewhat counter-intuitively) as zero, but allow for a large degree of uncertainty about these values by setting large values of the corresponding elements of $P_{1|0}$. This is the same approach used in the standard Minnesota prior for nonseasonal data.

Priors on variances ($P_{1|0}$) are governed by the following. $P_{1|0}$ is diagonal. When a seasonal variable (any variable other than the federal funds rate) appears on the right-hand-side of an equation, the variance prior for the coefficient is:

$$[P_{1|0}(j, m, n)]^{1/2} = TIGHT \times OTHER(j, m) \times SD(a) \times TS(n) \times \frac{\sigma_j}{\sigma_m}$$

where this variance $[P_{1|0}(j, m, n)]$ applies to the equation for variable j , ($j = 1, 2, \dots, 7$), the n th lag of the m th variable, and where $TIGHT$ is a hyperparameter governing overall tightness; $OTHER(j, m)$ is a parameter governing the relative tightness on the m th variable in the j th equation, which we set equal to 1.0 when $j = m$, and 0.2 when $j \neq m$; $SD(a) = SDECAY^{a-1}$, where $SDECAY$ is a hyperparameter, $a = 1$ when $n = 1, 2, \dots, 12$, and $a = 2$ when $n = 13, 14, \dots$

, 24, and so on; $TS(n) = \frac{1}{n^{TOOTH}}$ for $n = 1, 2, \dots, 11$, $TS(n) = \frac{1}{(n-11)^{TOOTH}}$ for $n = 12, 13, \dots$,

23, and so on. These priors allow for lower variances (that is, more confidence in the corresponding elements of b) as the lag, n , varies from $n = 1$ to 11, $n = 12$ to 23, etc. This reflects the influence of the function, $TS(n)$, and captures the belief that we have more confidence that a coefficient is zero when the lag is larger. However, $TS(n)$ also allows for discrete jumps in the variance (less confidence) at the seasonal lags, 12, 24, etc., reflecting the

belief that the coefficients at these lags are often much different than the value of zero assigned in the vector b . When $SDECAY < 1$, the function $SD(a)$ scales down the size of the upward jumps at the seasonal lags occurring in each calendar year. This downward scaling at annual intervals [due to $SD(a)$] parallels that at within-year monthly intervals [due to $TS(n)$].

When a nonseasonal variable (the federal funds rate) appears on the right-hand-side of an equation, the variance prior for the coefficient is:

$$[P_{10}(j, m, n)]^{1/2} = TIGHT \times OTHER(j, m) \times \frac{1}{n^{DECAY}} \times \frac{\sigma_j}{\sigma_m},$$

where $DECAY$ is a hyperparameter.

Two additional hyperparameters complete the specification of the Bayesian prior. First, we specify the elements of P_{10} corresponding to the constant and 11 seasonal dummy variables in each equation, j , as $\sigma_j^2 [TIGHT * CONREL]^2$, where $CONREL$ is set to a large number to reflect a large degree of uncertainty that the mean of these coefficients is zero, as specified in b . In our implementation, $CONREL$ equals 100. Second, the degree of time variation in β , given by the diagonal matrix Q , is determined by the hyperparameter, τ , whose square multiplies the prior on the variance prior, P_{10} . That is, we set $Q = \tau^2 P_{10}$. Normally, we expect τ to be a very small number, indicating that the variance of shocks to the process describing parameter evolution (Q) is a very small fraction of the variance of our prior (P_{10}). For each equation, $j = 1, \dots, 7$, we follow standard practice in estimating σ_j^2 as 0.9 times the variance of the residual from a univariate $ARMA(24, 0)$ including a constant and, for a seasonal variable, eleven seasonal dummy variables.

The hyperparameters *TIGHT*, *DECAY*, *SDECAY*, *TOOTH*, and τ are estimated by a simplex routine, to minimize the log determinant of the one-step-ahead MSE matrix of forecast errors.⁴

Following standard practice, we use the state-space representation (6) – (7) to form one-step ahead forecasts for each variable in the model as $y_{t+1|t} = Z'_{t+1} \beta_{t+1|t}$. Using (7), $\beta_{t+1|t} = \beta_{t|t}$, thus we have a one-step ahead forecasting rule given by: $y_{t+1|t} = Z'_{t+1} \beta_{t|t}$. Forecasts for more than one-step ahead are computed via the chain rule.

Figure 5 plots the square root of the elements of $P_{1|0}$ (that is, standard deviations) corresponding to the federal funds rate and pennies in the equation for pennies. Figure 5a shows the priors for the standard deviations at each lag for a nonseasonal variable (the federal funds rate). As the lag increases, the standard deviations decline, to indicate that we are more certain about the values of b as the lag increases. Figure 5b shows how the standard deviations for our priors about the coefficients on lagged pennies in the equation for pennies change as the lag

⁴ Estimation was subject to the following restrictions on the hyperparameters:

$0 < TIGHT \leq 1$; $0 \leq DECAY \leq 5$; $0 \leq SDECAY \leq 1$; $0 \leq TOOTH \leq 5$; and, $\tau \geq 0$. On the basis of several experiments, the estimates of the hyperparameters do not appear sensitive to starting values for those parameters. Estimation yielded the following point estimates: $TIGHT = 0.2606$, $DECAY = 1.2459$, $SDECAY = 0.7639$, $TOOTH = 0.8514$, and $\tau = 0.3533 \times 10^{-3}$. These estimates are based on a sample of data from 1959:2 to 1989:12, 24 lags (with observations prior to 1959:2 serving as presample values), and one-step-ahead forecasts for the period 1960:2 to 1989:12 for use in constructing the one-step-ahead mean-square-forecast-error matrix. The data were scaled as follows: net pay (in millions of coins) was divided by 100; the federal funds rate (in annualized percentage points) was divided by 400; and the not-seasonally-adjusted CPI inflation rate was constructed as a log first-difference. All computations were performed in RATS 5.02.

increases. The upward spikes at the seasonal lags (12 and 24) suggest an increased level of uncertainty about our prior that the coefficients are zero at those lags, although we are more certain at lag 24 than we are at lag 12. These functions are evaluated at the point estimates described in footnote 4 above.

The result of estimating the Bayesian VAR over the same testing period used in the other models generates an annual RMSFE of 1723, which compares favorably with all the other models, including the time-series model, as Table 3 shows. The forecasts from the Bayesian VAR pass the tests (sign test, Wilcoxon signed-rank test, zero-mean test, and test for unbiasedness) that the other models passed, as shown in Tables 4 to 7, except for the test for unbiasedness with quarterly data. Thus, the Bayesian VAR appears to be our best model, though as Figure 4 shows, its forecasts are very close to those of the time-series model.

3. Real-Time Forecasts

Each month, after the Federal Reserve coin offices calculate data on net pay, we generate new forecasts for net pay for pennies, nickels, dimes, and quarters at the national level using the time-series model, the VAR model, and the Bayesian VAR. We also use the time-series model to generate forecasts for the very small demand for halves and dollar coins.

The key question put to any forecasting method is: how well does it work? Unfortunately, we have only been forecasting the demand for coin for about 2 years, so we cannot answer that question very well. Table 8 shows the forecasts from each model made every three months from February 2001 to February 2003, along with the actual values in 2001 and 2002. Table 9 shows the same information for averages of the models (the VAR and time-series models for most of 2001; the VAR, time-series, and Bayesian VAR models beginning in

November 2001). As you can see in the tables, the initial forecasts were a bit high for 2001. Given what had happened in 1999 and 2000, with coin demand rising, the forecasting models predicted continued strong demand in 2001 that did not materialize. Instead, coin demand began declining substantially, and it took the forecasting models several months to adjust fully. Coin demand was 17 billion coins for the year, somewhat less than the 21 billion coins suggested by the models early in the year.

Forecasts for 2002 also moderated over time but later increased. Forecasts made in early 2001 pegged coin demand for 2002 at over 20 billion coins. The forecasts were gradually reduced through February 2002, as actual coin demand in 2001 was less than expected. Then, in spring 2002, coin demand grew a bit more than expected, so the forecasts began to rise somewhat. But demand declined in fall 2002, so the forecasts declined as well. The forecasts for 2003 reflect a similar pattern.

So far, the time-series model has done a better job of forecasting than the VAR because it was the quickest to drop the forecasts for 2001 and 2002 as net pay fell. But the period is much too short to favor the use of that model over the others. In a few years, we will have much more data on the forecasts and the errors made by each model, and we will be able to undertake a more complete examination.

4. Conclusions

This paper explained the basic models that we use for forecasting coin demand. We developed a structural model that we shelved because of its poor performance. We developed and currently use a VAR model, a time-series model, and a Bayesian VAR model, all of which

appear to provide sensible forecasts. As time passes, we will be able to evaluate the real-time forecasts of these models.

We are currently working on small modifications to each model. We developed the models in the midst of a major structural change in coin demand because of the state quarter program and the introduction of the golden dollar. As a result, we do not have enough data from that structural change to incorporate any type of break into the models. This could be done when a few years have passed and we are better able to ascertain the implications of this structural break. Also, each model treats all the coins symmetrically; we will now consider alternative structures for the different coins. We may even find that it would be optimal to mix models—the time-series model might be best for some coins and a VAR best for other coins, so the optimal model may be a combination of the two.

We have begun forecasting coin demand at the local office level using the time-series model with data on coin demand at each of the 37 coin offices. This project is designed to provide better analysis of coin demand and inventory control at each Fed cash office.

We are also considering alternative approaches to forecasting. First, once we have enough forecast history, we will run forecast encompassing tests to see if certain models should receive more weight than others. Currently, we simply take the average of the forecasts across the three models and use that as the main forecast for planning purposes. Second, there exist some estimates of the stock of coins, rather than the flow, so we could develop a stock model of coin demand instead of a flow model.

Table 1
Structural Model Results

Model	Economic Activity Variable	RMSFE
A	Services Consumption	3179
B	Nominal Retail Sales	2717
C	Industrial Production	2705
D	Nominal PCE	2737
E	Employment	2647

Notes: Quasi-ex-ante forecast dates 1990Q1 to 1998Q1; one-year-ahead forecasts made quarterly; RMSFE = root-mean-squared forecast error

Table 2
VAR Model Results
RMSFE for 1990 to 1998

Model	Economic Activity Variable	Lag Structure		
		1 to 2	1, 2, 3, 6, 9, 12	1 to 13
A1	Log of Employment	2485	2395	2279
B1	Log of Retail Sales	2495	2423	2299
C1	Log of Industrial Production	2454	2442	2337
D1	Log of Nominal PCE	2469	2456	2355
A2	Growth Rate of Employment	4024	2397	2358
B2	Growth Rate of Retail Sales	4195	2435	2339
C2	Growth Rate of Industrial Production	4081	2443	2412
D2	Growth Rate of Nominal PCE	4207	2470	2407

Notes: Quasi-ex-ante forecast dates 1990Q1 to 1998Q1; one-year-ahead forecasts made quarterly; RMSFE = root-mean-squared forecast error

Table 3
RMSFEs for Different Coin Models

Testing Period 1990 to 1998

Model	Quarterly RMSFE	Annual RMSFE
Structural Model	2647	2606
Time-Series Model	2035	1746
VAR	2279	2014
BVAR	2051	1723
Sample size	33	9

Notes: Quasi-ex-ante forecast dates 1990Q1 to 1998Q1; one-year-ahead forecasts made quarterly and annually; RMSFE = root-mean-squared forecast error

Table 4
Sign Test

Model	Quarterly t-statistic	Reject null?	Annual t-statistic	Reject null?
Structural	0.52	no	0.33	no
Time-Series	1.22	no	1.00	no
VAR	0.17	no	0.17	no
BVAR	0.87	no	0.33	no

Null hypothesis: forecast errors are independent with a zero median

Table 5
Wilcoxon Signed-Rank Test

Null hypothesis:

Model	Quarterly t-statistic	Reject null?	Annual t-statistic	Reject null?
Structural	0.53	no	0.41	no
Time-Series	0.37	no	1.13	no
VAR	0.28	no	0.04	no
BVAR	0.15	no	0.42	no

Null hypothesis: forecast errors are independent with a zero median

Table 6
Zero-Mean Test

Model	Quarterly p-value	Reject null?	Annual p-value	Reject null?
Structural	0.74	no	0.72	no
Time-Series	0.74	no	0.49	no
VAR	0.86	no	0.91	no
BVAR	0.88	no	0.72	no

Null hypothesis: Mean of forecast errors is zero

Note: Quarterly standard errors adjusted for overlapping-observations problem using Newey-West method.

Table 7
Test for Unbiasedness

	$\hat{\alpha}$	$\hat{\beta}$	\bar{R}^2	D.W.	p-value
Model					
Quarterly Forecasts					
Structural	23883 (4.60)	-0.46 (1.47)	0.09	0.11	0.00
Time-series	7047 (3.61)	0.57 (5.05)	0.40	0.41	0.00
VAR	9122 (3.00)	0.44 (2.46)	0.11	0.23	0.01
BVAR	6972 (3.27)	0.57 (4.66)	0.35	0.39	0.00
Annual Forecasts					
Structural	25841 (1.86)	-0.58 (0.69)	-0.07	0.98	0.22
Time-series	5641 (1.48)	0.67 (2.81)	0.46	1.77	0.35
VAR	7707 (1.88)	0.52 (2.08)	0.09	0.22	0.18
BVAR	5120 (1.57)	0.70 (3.44)	0.40	1.58	0.24

Numbers in parentheses are absolute values of t-statistics testing whether the coefficient is zero. The p-value reports the significance level of the test of the null hypothesis that $\alpha = 0$ and $\beta = 1$. For quarterly forecasts, the covariance matrix was adjusted for overlapping observations using the Newey-West procedure.

Table 8
Quarterly Real-Time Forecasts
(billions of coins)

Forecast Date	Period	Forecasts for 2001				2001
		2001:Q1	2001:Q2	2001:Q3	2001:Q4	
February 2001 (data through January 2001)						
	VAR	3.1	6.9	5.8	6.2	22.0
	Time-series	3.2	6.7	5.3	5.5	20.7
May 2001 (data through April 2001)						
	VAR		6.7	5.4	5.8	21.2
	Time-series		6.6	5.1	5.4	20.3
August 2001 (data through July 2001)						
	VAR			4.0	5.1	18.4
	Time-series			3.8	4.8	17.9
November 2001 (data through October 2001)						
	VAR				3.8	16.8
	Time-series				3.7	16.7
	BVAR				3.7	16.7
Actual for 4 Denominations						
		3.2	6.1	3.7	4.1	17.1

Table 8 (continued)
Forecasts for 2002

	2002:Q1	2002:Q2	2002:Q3	2002:Q4	2002
February 2001 (data through January 2001)					
VAR	3.8	7.0	6.0	6.5	23.3
Time-series	3.2	6.4	5.1	5.5	20.1
May 2001 (data through April 2001)					
VAR	3.6	6.7	5.5	5.9	21.7
Time-series	3.1	6.1	4.8	5.2	19.3
August 2001 (data through July 2001)					
VAR	3.2	6.2	4.7	5.5	19.6
Time-series	2.7	5.5	3.8	4.6	16.7
November 2001 (data through October 2001)					
VAR	2.3	5.5	4.3	4.8	16.8
Time-series	2.1	4.9	3.2	3.6	13.9
BVAR	2.2	5.3	3.6	4.0	15.0
February 2002 (data through January 2002)					
VAR	1.5	5.1	3.8	4.8	15.3
Time-series	1.5	4.7	2.9	3.6	12.7
BVAR	1.6	5.1	3.3	3.9	13.9
Actual for 4 Denominations	2.8	5.7	2.8	3.8	15.1
Forecasts for All 6 Denominations Begin May 2002					
May 2002 (data through April 2002)					
VAR		5.6	4.3	5.0	17.7
Time-series		5.5	3.9	4.1	16.3
BVAR		5.8	4.1	4.5	17.2
Aug 2002 (data through July 2002)					
VAR			4.0	4.9	17.5
Time-series			3.9	4.2	16.6
BVAR			3.9	4.4	16.8
November 2002 (data through October 2002)					
VAR				4.1	15.5
Time-series				3.8	15.2
BVAR				3.8	15.2
Actuals for 6 Denominations	2.8	5.8	2.8	3.8	15.2

Table 8 (continued)
Forecasts for 2003

	2003:Q1	2003:Q2	2003:Q3	2003:Q4	2003
February 2002 (data through January 2002)					
VAR	2.5	5.4	4.2	5.1	17.2
Time-series	1.4	4.4	2.8	3.5	12.2
BVAR	1.8	5.0	3.3	3.8	13.9
Actual for 4 Denominations					
Forecasts for All 6 Denominations Begin May 2002					
May 2002 (data through April 2002)					
VAR	3.3	5.8	4.7	5.2	19.1
Time-series	2.9	5.4	3.9	4.2	16.4
BVAR	2.9	5.8	4.1	4.3	17.1
Aug 2002 (data through July 2002)					
VAR	3.3	5.9	4.5	5.2	18.9
Time-series	2.9	5.6	3.9	4.3	16.7
BVAR	2.8	5.7	3.9	4.2	16.5
November 2002 (data through October 2002)					
VAR	2.7	5.7	3.7	4.8	16.9
Time-series	2.3	5.2	2.9	3.8	14.2
BVAR	2.2	5.2	2.9	3.6	14.0
February 2003 (data through January 2003)					
VAR	0.9	4.9	3.2	4.4	13.4
Time-series	0.9	4.5	2.3	3.2	11.0
BVAR	1.0	4.6	2.3	3.1	11.0

Table 9
Annual Real-Time Forecasts

Fed Average Forecasts					
Forecast Date	Forecasts	(billions of coins per calendar year)			2004
		2001	2002	2003	
Forecasts for Pennies, Nickels, Quarters, and Dimes; Average of Time Series and VAR Forecasts					
Feb 2001		21.4	21.7		
May 2001		20.8	20.5		
Aug 2001		18.1	18.1	18.7	
Forecasts for Pennies, Nickels, Quarters, and Dimes; Average of Time Series, VAR, and BVAR Forecasts					
Nov 2001		16.7	15.2	15.6	
Feb 2002			14.0	14.4	
Actual		17.1	15.1		
Forecasts for all 6 denominations; Average of Time Series, VAR, and BVAR Forecasts					
May 2002			17.1	17.5	
Aug 2002			17.0	17.4	
Nov 2002			15.3	15.1	15.3
Feb 2003				11.8	12.3
Actual			15.2		

REFERENCES

- Box, G.E.P., and G.M. Jenkins. *Time-Series Analysis, Forecasting, and Control*. San Francisco: Holden-Day, 1976.
- Croushore, Dean. "Evaluating Inflation Forecasts," Federal Reserve Bank of Philadelphia Working Paper 98-14, June 1998.
- Croushore, Dean, and Tom Stark. "A Real-Time Data Set for Macroeconomists," *Journal of Econometrics* 105 (November 2001), pp. 111-130.
- Hamilton, James D. *Times-Series Analysis*. Princeton, N.J.: Princeton University Press, 1994.
- Litterman, Robert B. "Forecasting with Bayesian Vector Autoregressions—Five Years of Experience," *Journal of Business and Economic Statistics* 4 (January 1986), pp. 25–38.
- Mullinix, Mark L. "Cash Services: Recent Increases in Coin Demand," Federal Reserve Bank of San Francisco *Announcement Letters: Cash Services*, August 24, 1999.
- Newey, Whitney K., and Kenneth D. West. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica* 55 (May 1987), pp. 703-8.
- Raynauld, Jacques, and Jean-Guy Simonato. "Seasonal BVAR Models: A Search Along Some Time Domain Priors." *Journal of Econometrics* 55 (1993), pp. 203–229.
- Roseman, Louise L. "The Golden Dollar Coin: Testimony before the Subcommittee on Treasury and General Government of the Committee on Appropriations, U.S. Senate," May 17, 2002.

Sims, Christopher. "Macroeconomics and Reality," *Econometrica* 48 (January 1980), pp. 1–49.

United States Mint, *2001 Annual Report*.

**Figure 1
Net Pay**

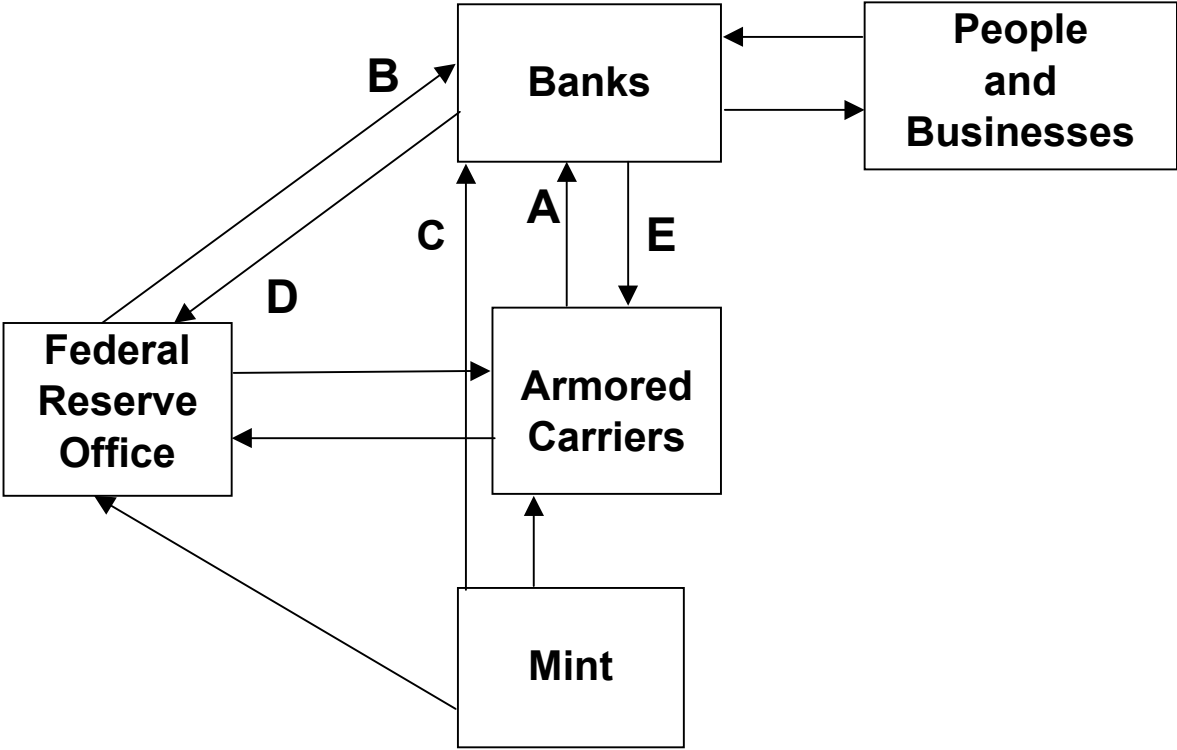



Figure 2a
Pennies Net Pay



Note: The thin line shows the monthly value of net pay; the thick line indicates the twelve-month moving average.

Figure 2b
Nickels Net Pay




Figure 2c
Dimes Net Pay




Figure 2d
Quarters Net Pay




Figure 2e
Halves Net Pay




Figure 2f
Dollars Net Pay




Figure 3
Annual Net Pay of Coins

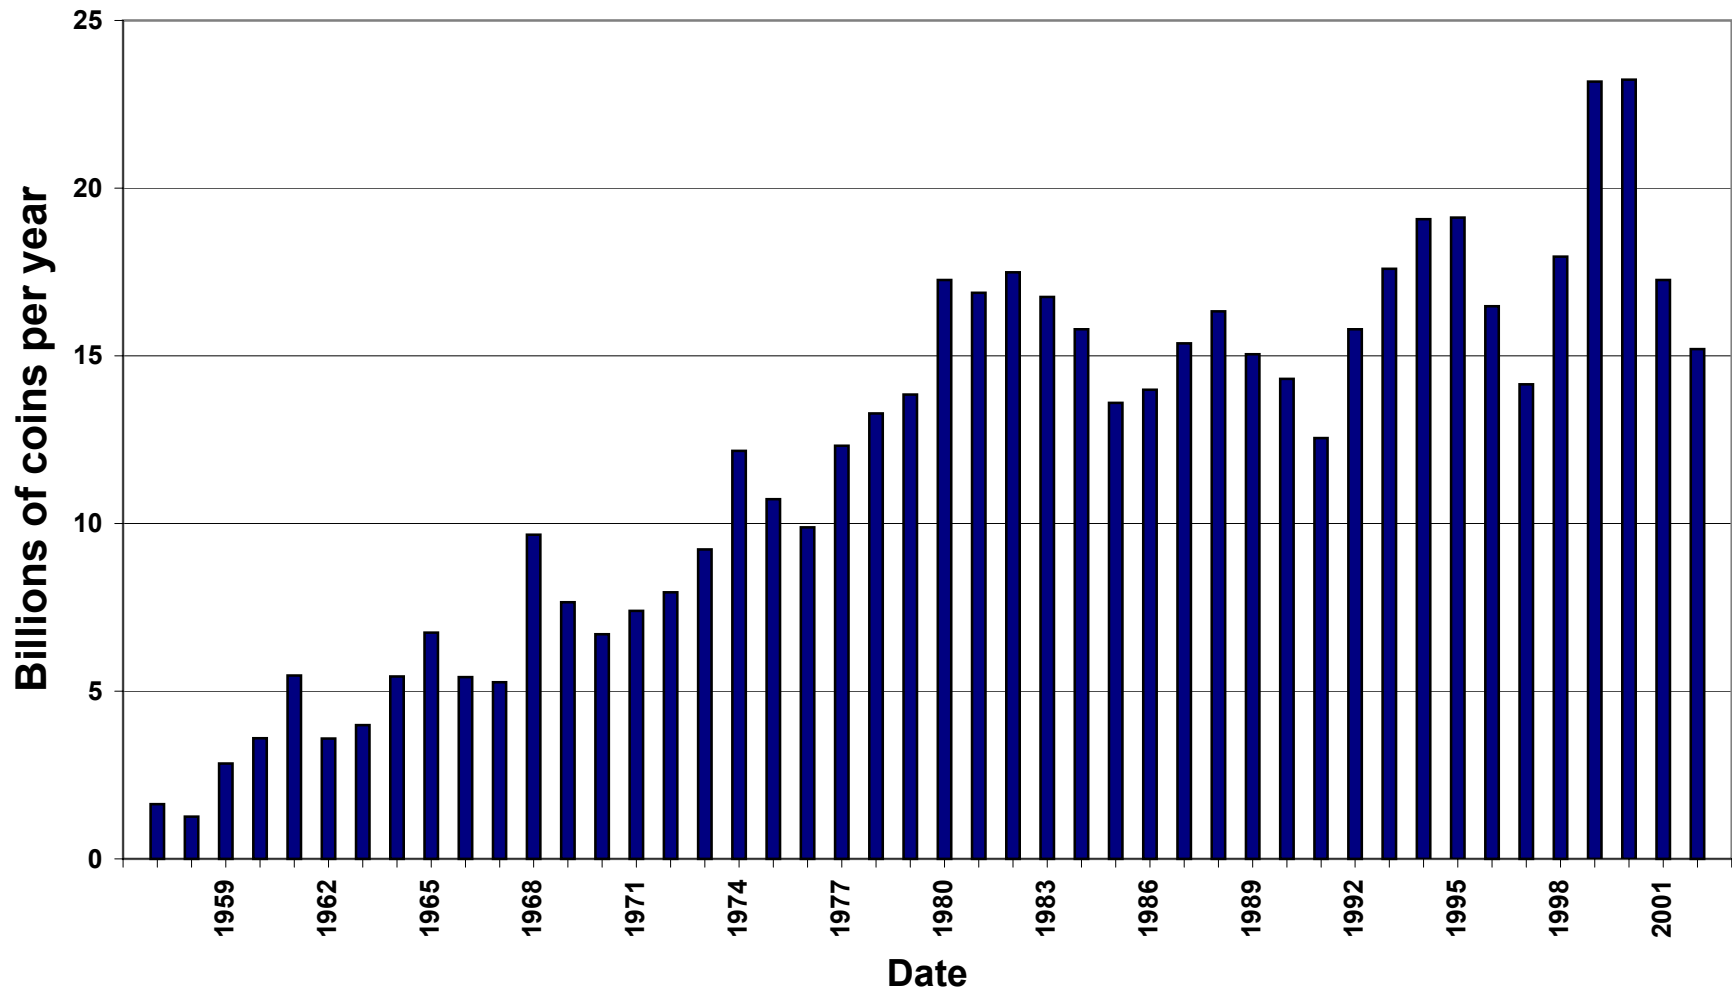


Figure 4
One-Year Ahead Forecasts




Figure 5a
Priors for Standard Deviations in the Equation for Pennies
Coefficients on the Lagged Federal Funds Rate




Figure 5b
Priors for Standard Deviations in the Equation for Pennies
Coefficients on Lagged Pennies

