U.S. Coins: Forecasting Change

BY DEAN CROUSHORE

Although the government annually produces about 70 new coins for every man, woman, and child, the economy's need for coins can vary from year to year. So how do the U.S. Mint, which makes the coins, and the Federal Reserve, which distributes them, decide how many coins the economy needs? Dean Croushore highlights some facts about coins and describes how demand for change is forecast.

Every year, the U.S. government produces about 70 new coins for every man, woman, and child in the country, or about 20 billion coins. In recent years, the program to produce new state quarters, plus the introduction of the golden dollar, increased total demand for new coins. When the demand for the new quarters and dollars became surprisingly strong in 1999 and 2000, shortages of some coins developed in different parts of the country.

To help prevent such shortages, a team of economists and analysts at the Federal Reserve are developing new models to forecast demand for coins. This article describes some of the work the Philadelphia Fed has undertaken since the project began in 2000.

Let's see how the Federal Reserve and the U.S. Mint decide on the number of coins to be produced and distributed and how difficult it is to forecast demand for coins. We'll begin by looking at some basic facts about the institutions involved and how demand for coins is calculated.[1]

THE FACTS ABOUT COINS

The U.S. Mint is in charge of producing coins. The Mint must obtain the raw metals to be used in production, procure the equipment, and hire workers to produce the coins needed for the economy. Once coins are produced, the Mint sells them at their face value to the Federal Reserve. The Mint makes a considerable profit (called seignorage) on the sales of coins. For example, one of the new Sacagawea golden dollars costs about 12 cents to produce and yields one dollar in revenue. The resulting profit of 88 cents goes to the U.S. Treasury. The Mint's profits increase the government's revenue by a billion dollars or more each year.

The Federal Reserve distributes coins from 37 coin offices, mainly Reserve Banks and their Branches, and more than 100 coin terminals operated by armored carriers such as Brinks and Loomis-Fargo. The armored carriers hold inventories of coins for the Fed and for banks and other financial institutions (hereafter just called banks). The carriers move coins between their terminals, banks, and Federal Reserve coin offices.

The U.S. Mint generally produces coins for circulation according to orders from the Federal Reserve. The Federal Reserve, in turn, generally orders an amount of coins based on the expected demand from banks. And, of course, those banks want coins to satisfy the demands of their customers. So, ultimately, the amount of coins produced depends on the demand for coins by people and businesses. Let's take a look at how that demand is measured.

The main concept in analyzing the demand for coins is called *net pay*. Net pay is an unusual economic concept because it represents the change in banks' demand for coins, which can be either positive or negative: It's positive

[1] The author thanks David Griffiths of the U.S. Mint for comments on an earlier draft of this article, and Brian Coulter of the Federal Reserve Cash Product Office, Geoffrey Gerdes of the Board of Governors, and Christopher Sims of Princeton University for consultation on the coin forecasting models. The author further thanks the team of economists and analysts at the Federal Reserve Bank of Philadelphia who worked on this project, including John Chew, Brian DiCiurcio, James Gillard, James Sherma, Keith Sill, and Tom Stark.

Dean Croushore is vice president and economist in the Research Department of the Philadelphia Fed.

when banks ask the Federal Reserve to deliver coins because demand has increased; it's negative when those institutions return coins to the Federal Reserve because demand for coins has declined.[2] Because of the durable nature of coins and because they can be returned to the Federal Reserve, coins (as well as paper money) are different from all other goods and pose unique challenges.

Coins flow between people or businesses and banks (Figure 1). When people and businesses want more coins, their banks order more from their local Federal Reserve offices, which 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 results in positive net pay, since banks' demand for coins is positive.

However, if people begin turning in more coins than banks want to keep on hand, the banks may return the extra to the Federal Reserve, either directly (line *D*) or through armored-carrier terminals (line *E*). So lines *D* and *E* represent negative net pay because demand for coins is negative.

In any given month, some banks may have positive net pay and others may have negative net pay. Net pay 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*:

$$\text{Net pay} = (A+B+C)-(D+E).$$

Net pay is calculated separately for six different denominations of coins: penny, nickel, dime, quarter, half-dollar, and dollar. Total net pay is the sum across all banks.

Let's take a look at the data on net pay for each denomination to see how each has changed over time. In doing so, we will look at the net pay for each denomination in units of millions of coins each month. The data run from 1957 to 2002, except for half-dollars and dollars. For each denomination, the black line shows the actual monthly amount of national net pay (summed across all 37 Federal Reserve offices). The green line is the average volume over the past 12 months (called a 12-month moving average), which is shown to help illustrate the long-term trend in the data (Figures 2a to 2f).

The charts show some interesting patterns. First, you can see that month-to-month seasonal fluctuations in net pay are huge. For some 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. We need a lot more change in the summer months for parking meters at the shore and for soda machines. We use more change around holidays at the end of the year, as well. But we need much less in the middle of the winter.

About half of all coins produced are pennies. 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 net pay for pennies has been fairly constant since 1980, perhaps trending down slightly (Figure 2a). For other main denominations (nickels, Figure 2b; dimes, Figure 2c; and quarters, Figure 2d), the trend over time has been slightly upward, which suggests that these other denominations may be gradually replacing pennies 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

FIGURE 1

Net Pay

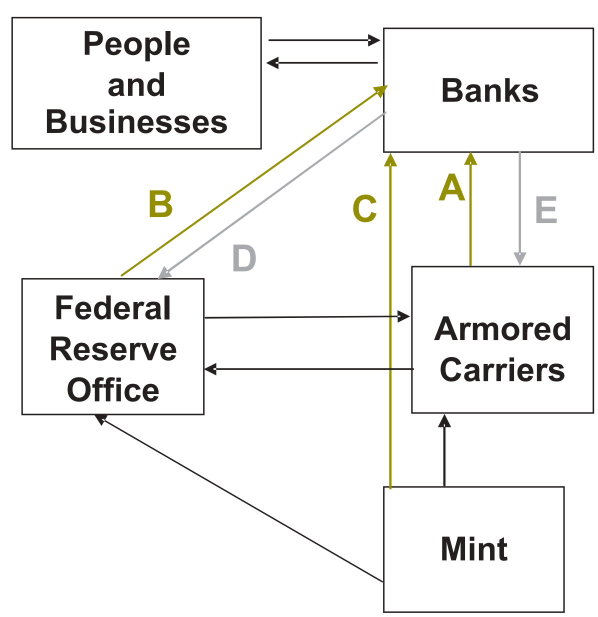

[2] Note that net pay differs from the change in the demand for coins in the rare instance in which there is a shortage of coins. Because such cases are rare and we do not have reliable data on the amount of shortages, our models of coin demand ignore such instances and assume that net pay equals the change in demand.

FIGURES 2a - 2c

Figure 2a
Pennies Net Pay

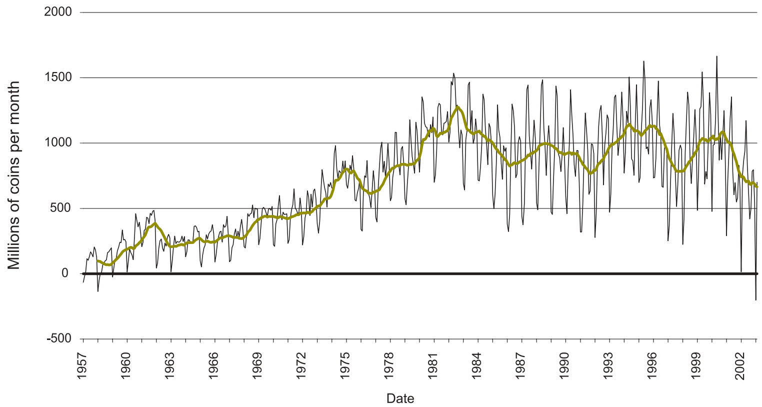

Figure 2b
Nickels Net Pay

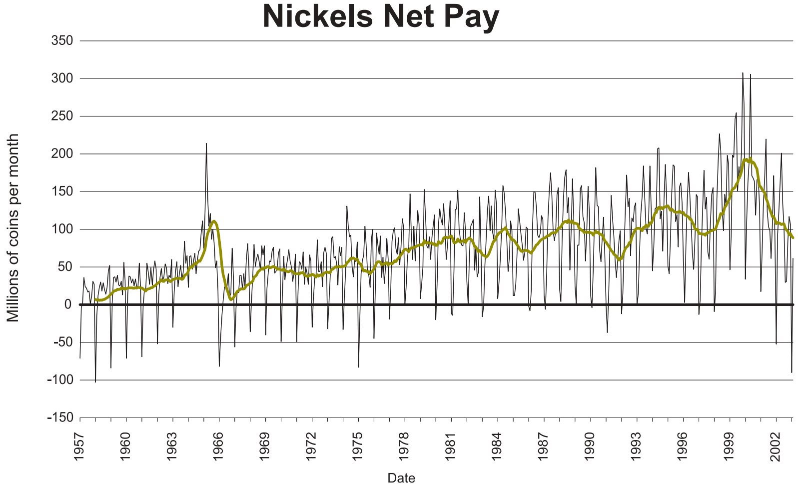

Figure 2c
Dimes Net Pay

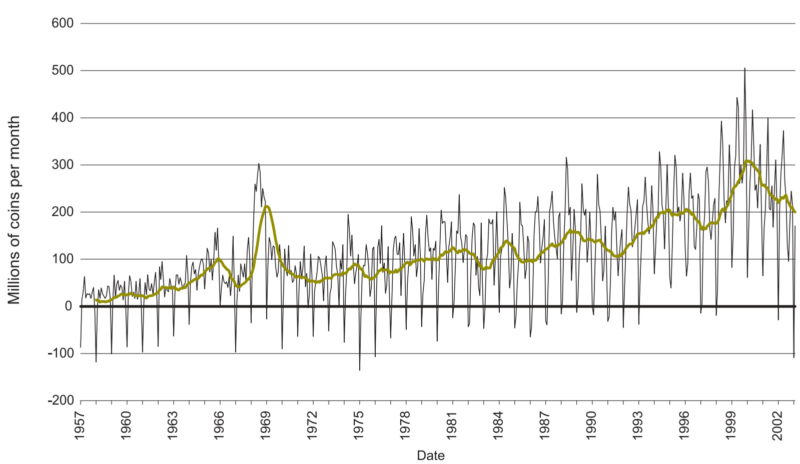

Note: The black line shows the monthly value of net pay; the green line indicates the 12-month moving average.

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 in 2000 caused a sharp increase in the net pay of dollars that year.

Given these trends in demand for different denominations, what can we say about overall demand for coins? To investigate this issue, we'll examine total demand for coins in terms of numbers of coins, adding up the net pay for all six denominations. Also, to avoid confusion arising from seasonal fluctuations, we just look at the total net pay in each calendar year (Figure 3). We look at the total number of coins rather than their dollar value, in part because the Mint's ability to produce enough coins to meet demand depends on the number of coins rather than 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, sometimes net pay rose and sometimes it fell.

We might expect a correlation between net pay and the strength of the economy because it seems likely that people will use more coins if they're buying more goods and services. However, there does not appear to be a strong correlation between net pay and economic activity. For example, while net pay fell when the economy weakened, as in 1990 and 1991, it fell even when the economy was strong, as it was in 1996 and 1997.

Special events raised net pay to very high levels in 1999 and 2000. First, 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 anticipated, thus causing net pay to rise sharply (Figures 2d and 3). Then, in 2000, the Sacagawea dollar was introduced to much fanfare. Initial demand for the new coin was also strong, and the Mint produced over 1 billion of them. 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. At the same time, the demand for nickels and dimes also rose substantially (Figures 2b and c).

FORECASTING COIN DEMAND

Sharp, unexpected increases in net pay during 1999 and 2000 led the Federal Reserve Bank of Philadelphia to investigate ways to improve forecasts of demand for coins. Developing a forecasting model involves several steps: choosing among different types of models, testing the different models to see how well they perform, seeing how they deal with changes, such as the introduction of the new quarter and dollar coins, then running forecasts in real time and investigating the quality of the forecasts. Because demand for each coin denomination seems to behave differently from that of the other coin denominations, the models we examine will contain a separate forecasting equation for each denomination, rather than modeling overall coin demand in a single equation.

Four Models of Coin Demand. We considered four different types of models: (1) a structural model; (2) a time-series model; (3) a vector autoregression (VAR) model; and (4) a Bayesian vector autoregression model. Brief descriptions of each model follow.[3]

Structural Model. Adapting the work of earlier researchers who had modeled coin demand, we first exam-

[3] Additional details about each model can be found in the research paper that I wrote with Tom Stark. The paper is listed in the References section at the end of this article.

FIGURES 2d - 2f

Figure 2d
Quarters Net Pay

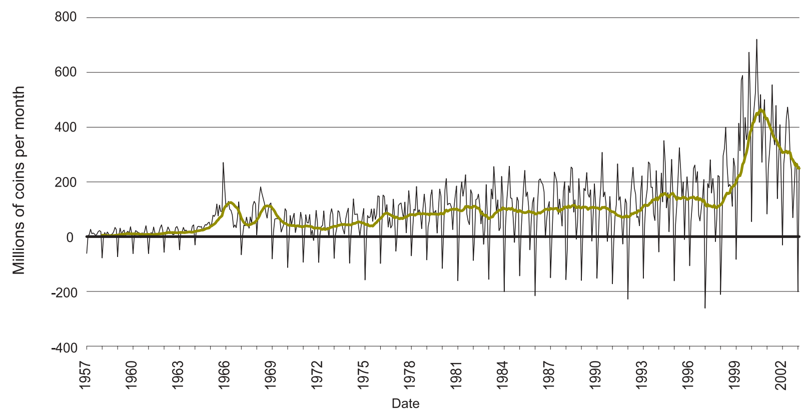

Figure 2e
Halves Net Pay

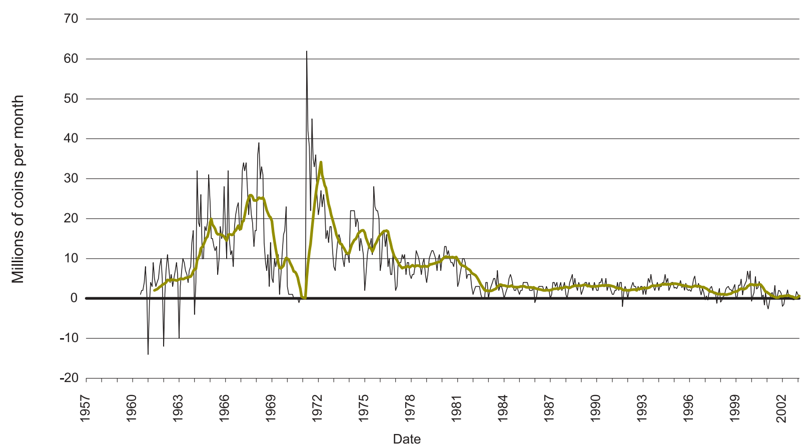

Figure 2f
Dollars Net Pay

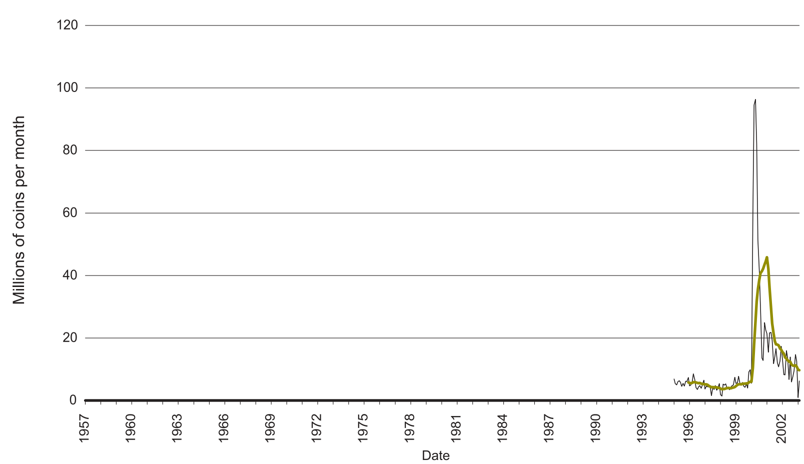

Note: The black line shows the monthly value of net pay; the green line indicates the 12-month moving average.

FIGURE 3

Annual Net Pay of Coins

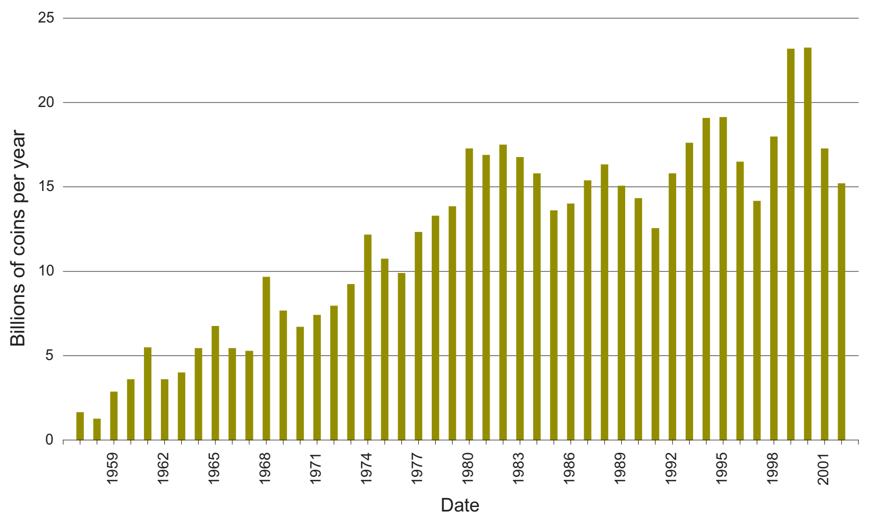

ined a structural model of net pay for each denomination. In a structural model, economic theory dictates which variables should affect demand for each denomination. We then develop a forecasting equation based on that economic theory.

Because many economic data are quarterly, we took quarterly averages of the monthly coin data. Economic theory suggests that demand for coins depends on economic activity, interest rates, and the inflation rate. We experimented with a number of measures of economic activity, including consumer spending on services (used in older models), retail sales, industrial production, personal consumption expenditures, and payroll employment. Payroll employment gave the best results in our tests, so we used that variable to represent economic activity. For an interest rate, we used the federal funds rate, which is the interest rate that banks charge each other on overnight loans. Since the fed funds rate is also the main variable the Federal Reserve targets with its monetary policy, it is a good indicator of the overall level of short-term interest rates. For inflation, we chose the inflation rate as measured in the Consumer Price Index. Again, based on an older model, the inflow of coins from banks to the Federal Reserve is modeled separately from the payout of coins from the Federal Reserve to banks. The forecast for the inflow is then subtracted from the forecast for the payout to generate a forecast for net pay. The forecasting model for each coin denomination also includes a seasonal variable for each quarter of the year to account for the seasonal pattern in coin demand.

Time-Series Model. The second forecasting method we tried was a time-series model, which uses data only from the past and data only on the variable being forecast. For example, the model for net pay of pennies assumes that net pay of pennies in the future depends only on past movements of net pay for pennies; it does not depend on the net pay of any other coin denomination or on any macroeconomic variable.

Though such a model is very simple, we still had to make choices about the forecasting equation of the model: how far back to go in determining the forecast, whether to model the level of net pay or the change in the level of net pay from one month to the next, and how to deal with the seasonal fluctuations — seasonally adjust the data before they go into the model or account for seasonal fluctuations within the model. Experimentation suggested that the best model was arrived at by using 14 months of lagged data for each forecast, modeling the change in net pay from one month to the next, and adjusting the data for monthly fluctuations before running the forecasting equation.[4]

Vector Autoregression (VAR) Model. In the past 20 years, many economists have stopped using structural models for forecasting because these models require more precise economic theory than we usually know; economists have also moved away from time-series models because such models use no economic theory at all. A vector autoregression (VAR) model is a mixture of a structural and a time-series model. The VAR uses economic theory to tell the researcher which variables should be included in the model, but it also incorporates time-series techniques by including past data on each variable in the model. A VAR is useful because it allows us to conduct "what-if" experiments, such as: "What will happen to demand for coins if the economy goes into a recession?"

For our VAR model, the net pay of each coin denomination depends on past data on economic activity, the interest rate, and the inflation rate, just as in the structural model. But the equation for each coin denomination depends on that denomination's own history, just as in the time-series model. In the VAR, the economic data

[4] The best model is determined on the basis of the root-mean-squared forecast error, which is described in detail later.

influence forecasts for coin demand, but coin demand is not allowed to influence forecasts of the economic data. Experimentation showed that the best model came from using data on the logarithm of payroll employment as the variable related to economic activity; that variable proved better in our tests than the growth rate of employment and was also superior to other variables, including industrial production, retail sales, and personal consumption expenditures. We also found it was best to use employment data that were not seasonally adjusted but to account for seasonal fluctuations within the forecasting model. Using 13 months of past data also provided the best results for this model.

Bayesian VAR. Bayesian techniques basically involve a researcher's beliefs (for example, concerning seasonality) about the outcome of an empirical investigation: The researcher examines the data in light of those beliefs, then sees if his beliefs change after he has examined the data. Essentially, Bayesian techniques help us rule out certain outcomes that differ so much from economic theory that we do not believe them. The techniques help to keep the estimated forecasting model within certain bounds. Economists use Bayesian methods because research has found that such methods often generate superior forecasts and may handle monthly fluctuations in the data better. Because of the large seasonal fluctuations in the coin data, Bayesian techniques may be very fruitful.

To implement Bayesian methods for forecasting net pay for coins, we applied them to the VAR model described above. 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, some key coefficients are chosen to make the model perform well, the most important being those that concern the seasonal patterns in the data. The Bayesian VAR is also not as restrictive as the VAR because it allows the data on coins to affect the macroeconomic variables (perhaps because people's spending habits are reflected in coin demand, which then helps predict the macroeconomic variables) and it allows data on one coin denomination to affect the forecasts for other coin denominations.

Because of the large seasonal fluctuations in the coin data, Bayesian techniques may be very fruitful.

Comparing the Models. After we built the models, we tested their performance over several periods. To avoid being unduly influenced by the introduction of the new state quarters and by the new dollar coin, we chose 1990 to 1998 as our main testing period. To see how the models would have performed over that time, we generated forecasts at the start of each three months, as if we were at that date and did not know what was to come. That is, we first used the coin data from January 1957 through December 1989, which would have been known to a forecaster making a forecast in January 1990, and generated a forecast for the next 12 months for each coin denomination. Then we stepped forward three months, as if we were in April 1990. Then, using the coin data through March 1990, we generated a forecast for the next 12 months. We continued this process until January 1998, at which point we generated forecasts through the end of 1998 (just before the start of the new state quarters program).

With these forecasts in hand, we calculated summary statistics on how well each forecast did. The most appropriate statistic is the root-mean-squared forecast error (RMSFE), which quantifies deviations — either positive or negative — of the actual from the forecast, where larger errors are penalized more. The RMSFE is calculated by taking a forecast for each year, calculating the forecast error (actual minus forecast) for each month, squaring each error, adding up the squared errors, dividing by the number of forecasted values, then taking the square root. Researchers have found that the RMSFE has a number of desirable properties and gives them a general guide to using forecasts: The best forecasts are those with the lowest RMSFEs. By squaring the forecast errors in calculating the RMSFE, forecasts that are far from actual are penalized more heavily than if we just calculated the average error (Table 1).

As you can see from the table, the time-series model, which was originally proposed as a benchmark model, proved very difficult to beat. In fact, only the Bayesian VAR did better, and its improvement was only slight.[5]

Using these models, we began to generate forecasts periodically, as requested, first by the Fed's Cash-Fiscal Product Office (located at the Federal

[5] Of course, all the forecasting models did much worse in forecasting coin demand in 1999 and 2000 because nothing in the models accounted for the introduction of new coins. But a researcher could have used these models to forecast net pay for purposes of determining how many coins were needed for circulation, then added a projection for demand for new coins that would not circulate because people would keep them as collector's items.

Reserve Bank of Philadelphia), and then the Fed's Cash Product Office (located at the Los Angeles Branch of the Federal Reserve Bank of San Francisco) when that office took over the responsibility for coin issues in spring 2001. Because the structural model performed so poorly in our tests, we stopped generating forecasts with it in early 2001. Instead, we added the Bayesian VAR to our process in September 2001.

Now, after the Federal Reserve coin offices calculate data on net pay at the end of each month, we generate new forecasts for net pay at the national level using the time-series, the VAR, and the Bayesian VAR models. Because no one knows how long or how large the increased demand for state quarters or the demand for the new dollar is likely to be, the best forecast is likely to be one that simply tracks the overall trend but does not generate forecasts that make strong assumptions about that future demand. All the models we use have that feature. For example, in 2001, demand for coins slowed substantially. Although the models did not predict the slowdown, the forecasts adjusted fairly quickly after the slowdown began.

How Have the Forecasts Performed So Far? The key question for any forecasting method is: How well does it work? Unfortunately, we have been forecasting demand for coins only for about two years, so we cannot answer that question very well. Table 2 shows the forecasts made every three months from February 2001 to November 2002, along with the actual values in 2001 and 2002. As you can see in the table, the initial forecasts for 2001 and 2002 were fairly high. 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.

So far, it appears that the time-series model has done the best job of forecasting because it was the quickest to lower 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.

USING THE FORECASTS

How can the Federal Reserve use these forecasts for coin demand? First, the national coin forecasts can help the Mint in planning its production. The Mint needs to schedule workers and to purchase enough equipment to produce the right amount of coins. Improved forecasts will allow the Mint to reduce production costs by getting a better idea of how many coins it will need to produce. In addition, the Mint will be able to order the appropriate amount of raw materials needed for production.

The forecasts can also help the Federal Reserve in ordering coins. Each Federal Reserve office must order coins every month, and improved forecasts could give them an additional tool for deciding how much to order. Those offices maintain inventories of coins in case of sudden changes in demand, so improved forecasts can help them maintain appropriate levels of inventories. Improved forecasts can also help these offices reduce costs by keeping inventories from becoming too large or too small, since shipping coins between offices is costly.

Because the time-series models performed so well at the national level, we began forecasting net pay for each office based on such models in spring 2002. Because there are 37 offices and six coin denominations, we generated 222 forecasts (37 x 6), each running monthly for the next 30 months. These forecasts are distributed to each office for its use in ordering coins and for planning. Coin offices must also take into account changes in local and national economic conditions that may not be captured in the time-series model that forecasts net pay.

SUMMARY

Forecasting the demand for coins is difficult because of seasonal

TABLE 1

RMSFE for Different Coin Models

Model	RMSFE
Structural Model	2.61
Time-Series Model	1.75
VAR	2.01
Bayesian VAR	1.72

Note: Figures shown are the root-mean-squared forecast error (RMSFE) for each model over the testing period from 1990 to 1998, in billions of coins. A forecast is more accurate if it has a smaller RMSFE. Models for half-dollars and dollars were not run because we have insufficient data.

TABLE 2

Annual Real-Time Net Pay Forecasts

Forecast Date	Actual Data Through	Calendar Year Forecasts 2001	2002
Feb 2001	Jan 2001	21.4	21.7
May 2001	Apr 2001	20.8	20.5
Aug 2001	July 2001	18.1	18.1
Nov 2001	Oct 2001	16.7	15.2
Feb 2002	Jan 2002		14.0
May 2002	Apr 2002		17.1
Aug 2002	July 2002		17.0
Nov 2002	Oct 2002		
Actual		**17.1**	**15.2**

Note: Amounts in billions of coins per calendar year.
Numbers shown for forecast dates from February 2001 to August 2001 are the average forecasts from the time-series model and VAR; numbers shown from November 2001 on are the average forecasts from the time-series model, VAR, and Bayesian VAR. Each forecast is a projection for the calendar year shown in the column header for the last two columns.

fluctuations in net pay and the introduction of new coins. By using some standard types of forecasting models, we have attempted to improve on existing forecasts of net pay. Whether these forecasting models will perform well in practice will require several years of observations. The models we use depend on the stability of historical relationships. As such, changes in how people use coins could cause the models to make large forecast errors in the future. If the models do not do well, we may be able to modify them so that they forecast better in real time.

Our hope is that we will be able to forecast coin demand well enough to prevent any shortages of coins in the future, without the expense of piling up large inventories of unused coins. BR

REFERENCES

Croushore, Dean, and Tom Stark. "Forecasting Coin Demand," Federal Reserve Bank of Philadelphia Working Paper 02-15/R, September 2002.

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.

United States Mint, *2001 Annual Report.*