Economic Review

Federal Reserve Bank of San Francisco

1993 Number 3

Adapting to Instability in Money Demand: Forecasting Money Growth with a Time-Varying Parameter Model

Timothy Cogley

Senior Economist, Federal Reserve Bank of San Francisco.
I am grateful to Hopi Huh, Mark Levonian, and Bharat Trehan for helpful comments on an earlier draft and to Desiree Schaan for computational assistance.

Conventional money demand models appear to be unstable, and this complicates the problem of conducting monetary policy. One way to deal with parameter instability is to learn how to adapt quickly when parameters shift. This paper applies a time-varying-parameter estimator to conventional money demand models and evaluates its usefulness as a forecasting tool. In relative terms, the time-varying-parameter estimator improves significantly on ordinary least squares. In absolute terms, we continue to have difficulty tracking money demand through turbulent periods.

According to the quantity theory of money, nominal spending depends on the supply of money and on velocity, and velocity is determined by money demand. If money demand is stable, monetary aggregates can be used as indicators of fluctuations in nominal aggregate demand. Furthermore, if money demand is functionally invariant to changes in money supply, then the Federal Reserve may be able to adjust the money supply in order to offset fluctuations in nominal spending that are due to non-monetary disturbances.

Conventional models of money demand appear to be unstable, however, and this greatly complicates the problem of conducting monetary policy. In particular, since money demand models are functionally unstable, it is difficult to interpret the information in monetary aggregates. For example, in recent years, the Federal Reserve System's money demand models have consistently underestimated M2 velocity. As a consequence, the Federal Reserve has overestimated the rate of M2 growth needed to sustain the projected growth in nominal GDP, and therefore actual M2 growth has fallen below its target range. Ordinarily, the unexpected shortfall in M2 growth would be a sign of serious weakness in the economy. However, in this case, it simply reflected the fact that velocity turned out to be higher than expected. Thus instability in money demand models makes it difficult for the Federal Reserve to keep money within its target range while still trying to achieve its goals for the economy.

As a theoretical matter, there is no reason to believe that conventional money demand models should be stable. For example, since conventional representations are subject to the Lucas critique, changes in central bank operating procedures can alter money demand parameters. Similarly, financial innovation may alter the relation between velocity and opportunity costs. Thus it seems appropriate to treat conventional money demand models as time-varying parameter models.

Roughly speaking, there are two ways to deal with time-varying parameters. One is to seek a deeper theoretical structure whose parameters are time invariant. So far, monetary economists have had little success with this approach. Another way to deal with parameter instability is to learn how to *adapt* to functional changes in money demand by allowing estimated parameters to change quickly when the model begins to show signs of instability. This paper takes the latter approach. It explores a time-varying parameter estimator that gives more weight to recent data and less weight to older data, so that estimates can change quickly when parameters change. The goal is to improve the predictive performance of money demand models.

This intuition is formalized in terms of discounted least squares (DLS). The paper applies recursive DLS to a number of conventional money demand models and compares its predictive performance with ordinary least squares (OLS). In relative terms, DLS compares favorably to OLS. For example, in cases where instability is especially important, DLS reduces the mean square error of one-quarter-ahead forecasts by 55 to 60 percent. Thus DLS can provide an important hedge against gross instability.

In absolute terms, however, conventional money demand models still have a great deal of trouble forecasting through turbulent periods. Thus DLS represents only a partial solution to parameter instability. In particular, since we continue to have difficulty tracking M2 demand, it will continue to be difficult to use M2 as an indicator of economic conditions.

I. Time-Varying Parameters

This section interprets conventional money demand functions in order to motivate the empirical approach taken in the paper. In conventional money demand models, demand for real balances depends on a scale variable, such as income, consumption, or wealth, and on opportunity cost variables. For example, Meltzer (1963) studied variations on the following model:

$$ln\,(m_t\,/\,p_t) = \beta_0 + \beta_1\,ln\,(r_t) + \beta_2\,ln\,(w_t) + u_t,$$

where m_t denotes nominal money balances, p_t is the price level, r_t is a nominal interest rate, w_t is either real wealth or income, and β_0, β_1, and β_2 are parameters.

Most of the empirical literature assumes that the parameters are *time invariant*. In practice, however, estimated money demand models appear to be unstable. For example, there was the famous "case of the missing money" in the mid-1970s (see Goldfeld 1976, or Judd and Scadding 1982). More recently, Feinman and Porter (1992) report evidence that M2 demand models have gone off course.

From a theoretical point of view, the instability of empirical money demand models is not puzzling. On the contrary, it is exactly what monetary theory predicts. Two kinds of arguments generate time-varying parameters, one based on the Lucas critique and another based on financial innovation.

The first argument concerns identification and is due to Cooley and LeRoy (1981). Traditional money demand models describe an equilibrium between money supply and money demand. In order to interpret the parameters solely in terms of money demand, however, money supply must be predetermined or exogenous. This condition seems dubious. If money supply is endogenous, the estimated parameters will depend at least in part on supply factors. Furthermore, if there are changes in the determinants of money supply, such as a change in monetary policy operating procedures, the parameters of conventional money demand models will also change. Thus, when money supply is endogenous, a necessary condition for parameter stability is that monetary policy rules not change during the sample. Since post-war U.S. data probably do not satisfy this condition, parameter instability is to be expected.

A second argument emphasizes financial innovation (e.g., Ireland 1992). For example, in cash-in-advance models, agents can buy some goods only with money and other goods with either money or credit. The cash-in-advance constraint gives rise to a transactions demand for money. A financial innovation expands the set of goods that can be bought on credit and thus (other things equal) reduces demand for real balances. Thus financial innovation alters the relation between money, interest rates, and expenditures. Since conventional money demand models do not fully capture the effects of financial innovation, one should expect parameter instability during periods of financial innovation.

Recently, a number of authors have argued that financial innovation may account for the recent bout of instability in M2 demand. In particular, the increased availability of mutual funds may have altered the relation between M2 velocity and interest rates (e.g., Feinman and Porter 1992 or Duca 1993). For example, banks have begun to market mutual funds to retail customers. As these funds become more accessible, the transactions costs of switching between M2 and various non-M2 securities are reduced. This increases the substitutability between M2 and stock and bond funds and thus increases the interest elasticity of M2 demand. The unusually steep yield curve of the last few years also may have induced some investors to switch into mutual funds.

These arguments suggest that we should treat conventional money demand functions as time-varying parameter models. Broadly speaking, there are two ways to deal with time-varying parameters. One is to seek a "deep structural" model of money demand, i.e., one that is invariant to financial innovation and monetary policy regime changes. In order to achieve invariance, however, a deep structural model would have to incorporate decision rules that govern financial innovation as well as rules that govern monetary policy regime changes. This approach is attractive in principle, since it would enable economists to evaluate the effects of policy changes. But this route does not seem promising at present, since monetary theory has not yet advanced to the point where it can deliver empirically useful representations for these decision rules. Given the state of knowledge, it may be worthwhile to seek an alternative solution.

Another approach is to learn to adapt to functional instability. There are at least two ways to think about adaptation. The most common approach is to respecify the model's functional form when it goes seriously off course. For example, recent efforts to respecify M2 demand models are described in Feinman and Porter (1992) and Duca (1993).[1] Another approach is to apply time-varying parameter estimators to conventional models in order to allow parameter estimates to adapt quickly when shifts occur. The time-varying parameter approach may prove useful for forecasting, even if its role in evaluating proposed changes in policy rules may be limited.

These two approaches should be regarded as complementary. Functional respecification is an ex post activity and therefore is not useful at the onset of a turbulent period. It generally takes many years to recognize model instability and to correct the problem, and the time-varying parameter approach may pay important dividends in the interim. Furthermore, when a model is respecified, it may be worthwhile to re-estimate by a technique that gives greater weight to recent data and less weight to older data, and this is precisely what time-varying parameter estimators do. On the other hand, if new financial instruments are introduced or if there are important omitted variables, time-varying parameter estimators may never fully adapt, and functional respecification may be necessary.

This paper concentrates on the potential usefulness of time-varying parameter estimators and does not explore functional respecification. My goal is to provide some insight into the marginal value of time-varying parameter estimators, but I do not claim that this is the only way to proceed.

II. RECURSIVE ESTIMATORS

In real time, Federal Reserve economists reestimate money demand models as new data become available. Since I want to study the reestimation process, it is useful to pose the problem in terms of recursive estimators. Begin by writing the model as

$$y_t = x_t' b_t + u_t,$$

where x_t and b_t are kx1 vectors, y_t and u_t are scalars. The vector b_t denotes the OLS parameter estimate based on data available through date t. If the model is reestimated by OLS each period, then b_t evolves as

$$b_t = b_{t-1} + P_{t-1}x_t(y_t - x_t' b_{t-1}) / (1 + x_t' P_{t-1}x_t),$$

$$P_t = P_{t-1} - P_{t-1}x_t x_t' P_{t-1} / (1 + x_t' P_{t-1}x_t),$$

where $P_t = (X_t'X_t)^{-1}$ and $X_t = (x_1, \ldots, x_t)'$. This is simply the formula for recursive OLS.

Recursive OLS might be appropriate if b were time invariant. However, theory and experience do not support time invariance. We regularly experience parameter shifts in money demand models. After a shift, the model tracks real balances poorly for a while, until the OLS recursions catch up with the parameter shift. The problem with recursive OLS is that it takes too long to catch up. Thus it seems worthwhile to consider alternative estimators that catch up more quickly.

Recursive OLS gives the same weight to all observations in the sample. When the model is subject to parameter shifts, it may be more sensible to give more weight to recent observations and less weight to distant ones. This intuition can be formalized in terms of discounted least squares (DLS) (Harvey 1981). That is, choose the vector b which minimizes the discounted sum of squared errors:

$$DSS = \sum_{t=1}^{T} \delta^{T-t} (y_t - x_t'b)^2.$$

The parameter δ is a discount factor. If $\delta = 1$, each observation is given equal weight, and this simplifies to OLS. If $\delta < 1$, observations close to the end of the sample (i.e., those close to the present) get more weight than those in the distant past. If the model is reestimated period by period by DLS, the parameter vector evolves as follows:

$$\tilde{b}_t = \tilde{b}_{t-1} + \tilde{P}_{t-1}x_t(y_t - x_t'\tilde{b}_{t-1}) / (\delta + x_t'\tilde{P}_{t-1}x_t),$$

$$\tilde{P}_t = \delta^{-1}\tilde{P}_{t-1} - \delta^{-1} \tilde{P}_{t-1}x_t x_t'\tilde{P}_{t-1} / (\delta + x_t' P_{t-1}x_t),$$

where $\tilde{b}_2$ denotes the DLS estimate based on data available through period t and $\tilde{P}_t^{-1} = \Sigma_{j=1}^{t}\delta^{t-j}x_j x_j'$. When $\delta = 1$, these recursions simplify to recursive OLS. When $\delta < 1$, the most recent observation gets more weight in the updating formula than it does under OLS.[2]

1. Feinman and Porter investigate alternative measures of opportunity cost with an emphasis on modeling effects of the steep yield curve. Duca proposes that mutual funds be added to M2 in order to internalize portfolio substitutions.

2. This technique is similar to the random walk parameter model of Cooley and Prescott (1976). One advantage of the Cooley-Prescott

The rationale for using DLS is that it will adapt more quickly to a parameter shift than will recursive OLS. But this comes at the expense of a loss in precision. For example, if the parameters were time invariant, DLS would discount useful information contained in the early observations, and this would increase the variance of the estimates. The parameter δ controls the terms of the trade-off. A value close to 1 favors precision over adaptability. A value far from 1 allows the model to adapt quickly but may produce highly variable estimates even when no shift has occurred. The discount factor must be chosen to balance adaptability against precision.

III. EXPERIMENTAL DESIGN

Robert Lucas warns economists to "beware econometricians bearing free parameters," and the DLS algorithm has a free parameter. Thus it is important to impose some discipline on the choice of δ. In particular, δ must be chosen based on information that is available before the forecast period begins. This section explains how δ is chosen and how the DLS algorithm is evaluated.

I divide the sample into three subperiods. The first covers the period 1954 to 1980 and is used to generate initial parameter estimates. M2 was redefined in 1980, and one of the criteria was that the new aggregate have a stable demand function (see Judd and Trehan 1992). Since parameter instability is not a problem for this subperiod, initial estimates are computed by OLS.

The second subperiod covers 1981 to 1988, and it is used to determine an optimal value for δ. I experiment with values of δ ranging from .8 to .99 and choose the value that minimizes the mean square error of recursive DLS forecasts.[3] M2 demand functions were relatively stable during this period. By choosing δ to optimize goodness of fit over this period, we ensure that the DLS algorithm produces reasonably stable parameter estimates during stable times. This is an important criterion. An algorithm that produced unstable estimates during stable periods would be of no use to anyone.

Some data are saved at the end of the sample to test the DLS algorithm. For the period 1989 to 1992, recursive DLS estimates are computed using the discount factors determined above, and they are compared with recursive OLS estimates. The principal reason for choosing 1989 as the beginning of the test period is a desire to have several years of data available for evaluating the time-varying parameter forecasts, and the results are not sensitive to the precise choice of sample split. Conventional M2 demand functions were unstable over this period. If my intuition is correct, the DLS algorithm should adapt more quickly than OLS, and recursive DLS forecasts should therefore have lower mean square error than recursive OLS forecasts.

IV. RECURSIVE ESTIMATES OF MONEY DEMAND

Basic Specification

This section applies recursive OLS and DLS to a number of standard money demand models. I assume that all the relevant variables are integrated processes and that there is a stable long-run relation between real balances and the scale variable.[4] Given these assumptions, money demand can be expressed as an error correction model. I consider a number of simple specifications which differ according to their scale and opportunity cost variables. The general specification is as follows:

$$ln(m_t / p_t) = a_0 + a_1\, ln(w_t) + z_t,$$

$$b_0(L)\Delta ln(m_t / p_t) = b_1 + b_2(L)s_t + b_3\,(L)\Delta ln\,(w_t) + b_4 z_{t-1} + u_t,$$

where s_t is a vector of interest rate spreads and z_t is the long-run "equilibrium error," in the language of Engle and Granger (1987). The first equation defines the long-run equilibrium relation between real balances and the scale variable.[5] Interest rate spreads are stationary and thus do not belong in the cointegrating relation. The second equation describes the short-term dynamics. The presence of the long-run "equilibrium error" in the second equation ensures that the short-run adjustments in money growth ultimately lead back to the long-run equilibrium level of real balances; hence the name "error correction" model.[6]

approach is that it implicitly allows different discount rates for different parameters. However, this would violate Lucas's dictum to avoid proliferating free parameters. When the Cooley-Prescott model is restricted so that there is only one discount factor, it is basically the same as DLS. I prefer DLS because it is more intuitive.

3. I also experimented with an a priori choice for δ, which was determined by the criterion that the discount function have a half life of 5 years. In general, this produced out-of-sample results that were superior to OLS but inferior to the data-determined value of δ.

4. The data do not contradict these assumptions.

5. I also examined models in which a_1 was set equal to one. This restriction implies that velocity is stationary and is equal to z_t. This had little effect on the result.

6. See Mehra (1991) for a more detailed exposition of error correction models of money demand.

One can write this as a single equation by substituting z_t from the first equation into the second.

I consider various combinations of scale and opportunity cost variables. The opportunity cost of holding M2 depends on the spread between returns on alternative assets and the own rate of return on M2. The latter is calculated by the Federal Reserve Bank of Richmond as a weighted average of the returns on the components of M2 (e.g., Mehra 1991). For alternative rates, I experiment with the six-month commercial paper rate and the 10-year Treasury bond rate.

For scale variables, I experiment with GDP and personal consumption expenditures. GDP is the standard scale variable in the money demand literature. Consumption can be motivated in two ways. First, consumption is the appropriate scale variable in cash-in-advance models (e.g., Lucas 1988). Second, as an empirical matter, various authors have emphasized that permanent income performs better than current income (e.g., Meltzer 1963), and consumption is a natural, observable proxy for permanent income.

To complete the specification, each version of the model also includes dummy variables for the second and third quarters of 1980, during which credit controls were binding, as well as a dummy variable for the first quarter of 1983, when MMDA accounts were introduced. Finally, the lag polynomials in the second equation are assumed to be of order 1. This is sufficient to capture the dynamics of real M2 growth during the initial estimation period.[7]

The Experimental Period

Each of these models was estimated by recursive OLS and DLS, using quarterly data, and the results are reported in Table 1. The first two columns report the mean square error of recursive one-quarter-ahead forecasts for the various models and time periods. Mean square error is standardized by dividing by the variance of the dependent variable; thus R^2 statistics are equal to 1 minus the mean square error. R^2 statistics are useful for evaluating absolute performance, and mean square error is useful for evaluating relative performance.

The first column of Table 1 reports results for the experimental period 1981–1988. During this period, the various models accounted for roughly 35 to 65 percent of the variation in real balance growth. GDP appears to be superior to consumption, reducing the mean square forecast error by roughly 20 percent. Furthermore, the short-term interest rate appears to perform better than the long-term rate. For example, when GDP is the scale variable, the mean square error for the short-rate model is approximately 35 percent lower than the mean square error for the long-rate model. Including long rates as well as the short rate has no effect on forecast performance. Thus, during the experimental period, the two best models were the ones based on GDP and short-term interest rates.

Even during this period, when M2 money demand was relatively stable, recursive DLS often worked significantly better than recursive OLS. In particular, for the models that include short rates, DLS reduces mean square error by an average of approximately 15 percent. In the long-rate models, the optimal discount factor turns out to be .99, so there is essentially no difference between DLS and OLS. Recall that DLS trades precision for adaptability. Even during the relatively stable subperiod, the gain from adaptability often outweighed the loss of precision.

TABLE 1

PREDICTIVE POWER AND BIAS

	MSE(81-88)	MSE(89-92)	BIAS(89-92)
1. Six-Month Commercial Paper Rate			
GDP			
$\delta = 1.00$	0.471	3.655	−5.42 (.000)
$\delta = 0.82$	0.358	1.430	
Consumption			
$\delta = 1.00$	0.601	1.564	−3.14 (.002)
$\delta = 0.80$	0.510	1.316	
2. Ten-Year Treasury Bond Yield			
GDP			
$\delta = 1.00$	0.637	1.267	−1.35 (.176)
$\delta = 0.99$	0.638	1.270	
Consumption			
$\delta = 1.00$	0.679	1.290	−1.42 (.154)
$\delta = 0.99$	0.681	1.269	
3. Six-Month Commercial Paper and Ten-Year Treasury Bond Rates			
GDP			
$\delta = 1.00$	0.482	3.768	−5.47 (.000)
$\delta = 0.80$	0.361	1.704	
Consumption			
$\delta = 1.00$	0.648	1.149	−2.71 (.002)
$\delta = 0.93$	0.613	1.109	

NOTE: The first two columns report the mean square error of recursive prediction errors scaled by the variance of real M2 growth. The third column reports the statistic $\sqrt{k}E(v_t)/\sigma$, with normal probability values shown in parentheses. When $\delta = 1$, this corresponds to recursive OLS.

7. Formally, this is sufficient to generate white noise residuals during this period.

The Test Period

The second and third columns of Table 1 report results for the test period 1989–1992. These columns reveal four results. First, when estimated by recursive OLS, the performance of all the models deteriorates badly. The mean square error of recursive OLS forecasts increases dramatically, and "out-of-sample" R^2 statistics are negative in every case (see the second column).[8]

Second, during the test period, the recursive OLS models consistently overestimated real M2 growth. One can test for bias in recursive OLS by computing the mean of normalized OLS prediction errors:

$$v_t = (y_t - x_t' b_{t-1}) / f_t^{1/2},$$

where $f_t = (1 + x_t' P_{t-1} x_t)$. If the model is stable, v_t has mean zero, is serially uncorrelated, and has the same variance as the regression disturbance. Further, if the regression disturbance is normally distributed, then v_t is also normally distributed (see Harvey 1981). Let $E(v_t)$ denote the mean of v_t over the test period: $E(v_t) = (1/k)\Sigma_{t=T+1}^{T+k} v_t$. Given these assumptions, $E(v_t)$ is normally distributed with mean zero and variance equal to σ^2/k. Thus $\sqrt{k}E(v_t)/\sigma$ is distributed as a standard normal random variable.[9]

Bias statistics are reported in the third column of Table 1. The mean recursive residual is negative in all models, and the means are statistically significant in four of the six cases. Since stable models have unbiased recursive residuals, this result confirms our belief that conventional money demand models should be treated as time-varying parameter models.[10]

Third, compared with OLS, DLS performs relatively well, and the percent improvement appears to be positively related to the degree of model instability. For example, the two models that had the lowest mean square error during the period 1981–1988 (i.e., the GDP-short-rate models) turn out to have the highest mean square error during the test period. For these models, DLS reduces the mean square error of one-quarter ahead forecasts by 55 to 60 percent. Thus DLS can be an important hedge against gross instability.

Figure 1 illustrates one of these cases. This is derived from the model that uses GDP as the scale variable and the six-month commercial paper rate as the opportunity cost variable (i.e., the first model in Table 1). The solid line shows real M2 growth, and the dotted lines show forecasts generated by recursive OLS and DLS, respectively. In the second half of 1989, both models systematically begin to overestimate real M2 growth. Recursive OLS is slow to catch on to the apparent structural shift, and it continues to overestimate real M2 growth until the end of 1992. Recursive DLS is more adaptable. It begins to recognize a structural shift around the second quarter of 1990, and its forecasts begin to edge downward. By the second half of 1991, this model seems to be back on track. Whether it stays on track remains to be seen. In this model, recursive DLS reduces mean square error by 60 percent relative to recursive OLS.

DLS is least useful in models that are relatively stable. For example, in models that omit the short-term interest rate, DLS and OLS produce basically the same results. It is worth noting DLS does not significantly *hurt* forecast performance when applied to relatively stable models, so discounting appears to be essentially costless.

Fourth, despite the relative improvement due to DLS, the absolute performance of the DLS models also deteriorated badly during the test period. These models also consistently overestimate real M2 growth and also have negative out-of-sample R^2 statistics. Thus, while DLS is

FIGURE 1

RECURSIVE FORECASTS

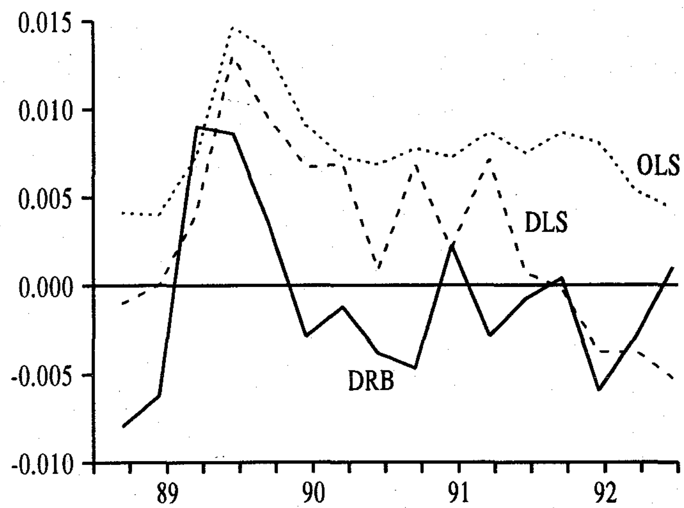

8. Recall that R^2 equals 1 minus mean square error. Recursive prediction errors do not necessarily have mean zero, so their mean square error can be larger than the variance of the dependent variable.

9. $E(v_t)$ is asymptotically normal even if the regression error is not normally distributed, provided that the regression error satisfies a mixing condition (e.g., White 1984). Thus, $\sqrt{k}E(v_t)/\sigma$ is approximately normal for reasonably large k.

10. This result contrasts with Mehra (1992), who fails to reject parameter stability in a similar model.

better than OLS, it does not appear to generate enough improvement to revive the use of M2 as an indicator of short-term fluctuations in nominal aggregate demand.

In retrospect, it is clear that completely naive, atheoretical forecasts would have worked about as well as any of these money demand models over the period 1989–1992. For example, forecasts based on a random walk model of real balance growth would have produced a mean square error of 1.387 over this period, which is comparable to the performance of these models. This simply highlights the difficulty of using historical relationships to forecast during turbulent periods.

V. CONCLUSION

Conventional money demand models often exhibit parameter instability, and this complicates the implementation of monetary policy. Applied macroeconomists might respond to this in two ways. They might seek time-invariant, deep structural representations, or they might apply time-varying parameter estimators to conventional representation in order to allow parameters to adapt quickly when shifts occur. This paper takes the latter approach, exploring the properties of recursive discounted least squares. This technique is designed to give greater weight to more recent data and less weight to older data, and this makes it more adaptable than recursive OLS.

The results suggest that DLS may have a useful but limited role to play in policy modeling. During unstable subperiods, DLS works better than OLS, and the gains can be substantial. For example, in a standard money demand model, in which the scale variable is GDP and the opportunity cost variable is the spread between commercial paper rates and the own return on M2, DLS reduces the mean square error of one-step-ahead forecasts by 60 percent. On the other hand, the absolute performance of DLS estimators also deteriorates badly over the last few years, and the models do not deliver reliable forecasts of M2 money demand. Thus it will continue to be difficult to interpret fluctuations in money growth.

REFERENCES

Cooley, T.F., and S.F. LeRoy. 1981. "Identification and Estimation of Money Demand." *American Economic Review* 71, pp. 825-844.

__________, and E.C. Prescott. 1976. "Estimation in the Presence of Stochastic Parameter Variation." *Econometrica* 44, pp. 167-184.

Duca, J.V. 1993. "Should Bond Funds Be Included in M2?" Working Paper No. 9321. Federal Reserve Bank of Dallas.

Engle, R.F., and C.W.J. Granger. 1987. "Cointegration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55, pp. 251-276.

Feinman, J.N., and R.D. Porter. 1992. "The Continuing Weakness in M2." Working Paper No. 209. Board of Governors of the Federal Reserve System.

Goldfeld, S.M. 1976. "The Case of the Missing Money." *Brookings Papers on Economic Activity* 3, pp. 683-730.

Harvey, A.H. 1981. *Time Series Models*. London: Philip Allan.

Ireland, P.N. 1992. "Endogenous Financial Innovation and the Demand for Money." Working Paper 92-3. Federal Reserve Bank of Richmond.

Judd, J., and J. Scadding. 1982. "The Search for a Stable Money Demand Function: A Survey of the Post-1973 Literature." *Journal of Economic Literature* 20, pp. 993-1023.

__________, and B. Trehan. 1992. "Money, Credit, and M2." Federal Reserve Bank of San Francisco *Weekly Letter* (September 4).

Lucas, R.E. 1988. "Money Demand in the United States: A Quantitative Review." *Carnegie-Rochester Series on Public Policy* 29, pp. 137-168.

Mehra, Y.P. 1991. "An Error Correction Model of U.S. M2 Demand." Federal Reserve Bank of Richmond *Economic Review* 77, pp. 3-12.

__________. 1992. "Has M2 Demand Become Unstable?" Federal Reserve Bank of Richmond *Economic Review* 78, pp. 27-35.

Meltzer, A.H. 1963. "The Demand for Money: The Evidence from the Time Series." *Journal of Political Economy* 71, pp. 219-246.

White, H. 1984. *Asymptotic Theory for Econometricians*. New York: Academic Press.