P* Revisited: Money-Based Inflation Forecasts with a Changing Equilibrium Velocity*

Athanasios Orphanides and Richard Porter
Board of Governors of the Federal Reserve System

May 1998

Abstract

This paper implements recursive techniques to estimate the equilibrium level of M2 velocity and to forecast inflation using the P^* model. The recursive estimates of equilibrium velocity are obtained by applying regression trees and least squares methods to a standard representation of M2 demand, namely a model in which the velocity of M2 depends on the opportunity cost of holding M2 instruments. Equilibrium velocity is defined as the level of velocity that would be expected to obtain if deposit rates were at their long-run average (equilibrium) value. We simulate the alternative models to obtain real-time forecasts of inflation and evaluate the performance of the forecasts obtained from the alternative models. We find that, while a P^* model assuming a constant equilibrium velocity does not provide accurate inflation forecasts in the 1990s, a model based on our time-varying equilibrium velocity estimates does quite well.

KEYWORDS: Inflation, M2 velocity, quantity equation.

JEL Classification System: E30, E50

Correspondence: Orphanides: Division of Monetary Affairs, Board of Governors of the Federal Reserve System, Washington, D.C. 20551, tel.: (202) 452-2654, email: aorphanides@frb.gov. Porter: Division of Monetary Affairs, Board of Governors of the Federal Reserve System, Washington, D.C. 20551, tel.: (202) 452-2497, email: rporter@frb.gov.

* The opinions expressed are those of the authors and do not necessarily reflect the views of the Board of Governors of the Federal Reserve System. We wish to thank John Carlson, John Duca, Jim Lee, and Robert Rasche for comments on earlier versions of this paper.

1 Introduction

Given an estimate of real potential output, Q^*, and an estimate of equilibrium velocity, V^*, P^* is defined as the equilibrium level of prices supported by the current quantity of money in circulation, M:

$$P^* \equiv \frac{MV^*}{Q^*}$$

As Hallman, Porter and Small (1991) showed, P^* can potentially provide a useful anchor for the price level and as such be utilized as a tool for predicting inflation. The framework for understanding the monetary dynamics of inflation relies on the simple idea that if the current price level, P, deviates from its equilibrium level, P^*, then inflation will tend to move so as to close this gap between the actual and equilibrium price levels—the price gap.

However, implementation of the framework requires a firm understanding of what the level of equilibrium velocity is. Much of the original appeal of the Hallman, Porter and Small study was based on the simplicity of their definition of V^*. Using M2, they showed that assuming a constant equilibrium velocity for their sample (1955 to 1988) was a sufficiently accurate representation despite the waves of financial innovation which had taken place during that time period. As monetary practitioners and theorists have always recognized, however, the continuous innovation in financial markets implies that such presumed observed constancies cannot be taken for granted.[1]

By now it is well known that the presumption of constancy of the equilibrium velocity of M2 is no longer valid. As early as 1991 the stability of the historical statistical relationships involving M2 was already being questioned at the Board of Governors. (See Feinman and Porter (1992) for an early accounting of this breakdown.) And consequently, it was recognized that using the P^* framework to forecast inflation, based on the assumption of a constant equilibrium velocity for M2, was no longer reliable.

In this paper we investigate how the P^* framework could be used in an environment in which V^* may be time varying. Specifically, we provide some guidance regarding how the

[1] Indeed, in recognition of this truism, Hallman, Porter and Small warned that "[i]f permanent shifts to velocity are empirically significant, actual prices would diverge from P^* in the long run."

equilibrium velocity of M2 could be adjusted to enable continuing use of the price gap for forecasting inflation.

An obvious ex post "correction" could been obtained directly by simply computing the value of V^* which would have eliminated the inflation forecast errors resulting from the incorrect assumption of a constant V^*. But such an exercise is circular and would be useless for forecasting inflation in real time. Rather, information from other variables observed contemporaneously with velocity should be brought into consideration. To that end, we examine the co-movements of velocity and the opportunity cost of money suggested from traditional money demand formulations as our alternative source of information regarding potential changes in equilibrium velocity and compute the change in V^* implied in that relationship.[2] We examine several alternative specifications of V^* which could have reasonably been obtained in real time by recursive estimation, as soon as the breakdown in equilibrium velocity was recognized. Eventually, each of our estimates exhibits a noticeable upward shift in equilibrium velocity of about the same amount, although they differ somewhat regarding the date when the shift became evident.

Using these alternative specifications of V^* we then show the corresponding one-year-ahead inflation forecasts for the 1990s. The results suggest that much of the deterioration in the inflation forecasts obtained from the P^* framework using the incorrect assumption of constant equilibrium velocity is reversed once we account for the apparent shift in equilibrium velocity.

2 Equilibrium velocity

Traditional theories for the demand for money hold that velocity fluctuates with the opportunity cost of money. Letting $\widetilde{OC}$ denote *deviations* of the opportunity cost of money, OC from its average norm, a simple way to capture this relationship is as follows:

$$V = V^* + \alpha_1 \widetilde{OC} + e$$

[2]The usefulness of bringing information from such sources to bear on the P^* framework has also been recognized by Koenig (1994) who proposed simultaneous estimation of money demand and P^* models.

where α_1 measures the response of velocity to movements in the opportunity cost of money and e is a stationary zero mean error term. In this setting the stability of V^* can be examined in a straightforward manner. If V^* were a constant, it could be easily estimated from the (population) regression:

$$V = \alpha_0 + \alpha_1 \widetilde{OC} + e$$

where, of course, an estimate of V^* could be obtained as the estimated parameter α_0.

Hallman, Porter and Small implicitly relied on the stability of money demand and the absence of any trend in the movement of $\widetilde{OC}$ over the sample period in which they developed the P^* model. Specifically, if in the above population regression the error term, e, averages about zero and departures of OC from its norm, $\widetilde{OC}$, also average about zero, then it is immediate from this regression that the sample mean of V, $\bar{V}$, will give a good estimate of V^*, which is the estimate that Hallman, Porter, and Small selected. However, if V^* shifts up as it apparently has in the early 1990s, the revised estimate of V^* needs to embody some of the upward drift of the velocity error e which occurred then.

Defining the opportunity cost of $M2$ holdings as the rate on the three-month treasury bill minus the average rate paid on $M2$ balances, the relationship between $M2$ velocity and the opportunity cost is shown in figure 1.[3] As can be seen, between 1960 and until about 1990, M2 velocity and opportunity cost moved together quite closely and could likely be described by the simple regression above. Indeed, the relationship implicit in this diagram has formed the basis for models of M2 demand, including the models in Moore, Porter and Small (1990) which were used at the Board of Governors for the last decade.[4] Since about 1990, however, the gap between M2 velocity and opportunity costs widens in a way that does not appear to fit the previous relationship.

Figure 2 provides a scatter plot of the same data. The solid line shows the regression line of velocity on the opportunity cost estimated with data from 1960:1 to 1988:4. The two dotted lines form a 95 percent confidence band for the fitted values of velocity. As can be

[3]See the data appendix for details on the definitions.

[4]Figure 1 is in essence an updated and rescaled version of figure 9 in Moore, Porter and Small (1990).

seen, the assumption of a constant V^* for this period is not incompatible with this model. The point estimate for V^* can be read as the value on the fitted line corresponding to the historical norm of the opportunity cost. Starting in about 1991, however, the observed velocity of M2 appears consistently outside the confidence band obtained based on the assumption of a constant V^*. These results suggest that by 1991, the constancy of V^* ought to have been seriously questioned—as was indeed the case.

Unlike Hallman, Porter and Small who posit the V^* is a constant, however, money demand models including Moore, Porter and Small typically allow for the possibility of a trend in V^*. This can be easily captured by specifying

$$V = \alpha_1 + \alpha_1 \widetilde{OC} + \alpha_2 TIME + e$$

where $V^* = \alpha_0 + \alpha_2 TIME$. As it turns out, a small positive trend does appear in the data. Although statistically significant, it is very small compared to the movements of velocity after 1991. Consequently, the possible omission of a trend in computing V^* cannot by itself account for the apparent non-constancy of equilibrium velocity. Nonetheless, for later comparisons, we will use both formulations—including and excluding a trend—as benchmarks.

Another interesting element in figure 2 is that in contrast to the rapidly increasing velocity during the 1992-1994 period, which was not associated with correspondingly large movements in opportunity costs, movements in velocity and opportunity costs since then appear to follow a trajectory roughly parallel to the relationship depicted by the solid line for the early part of the sample. This observation suggests that V^* could now possibly be adequately modeled as roughly constant but at a new higher level. If so, there may be simple alternatives to the constant V^* specification that may well suffice to provide useful estimates of P^*.

2.1 Allowing for a one-time shift in V^*

Perhaps the simplest alternative hypothesis to that of a constant V^* is a specification allowing for a one-time shift in the intercept of the velocity regression described above. In real time, of course, the estimated size of the shift would need to be re-estimated with each additional quarter of data. And since considerable uncertainty might prevail for several quarters regarding the timing of the shift, the same regression could be used to optimally determine this timing. To allow for such a shift in V^* we simply add the variable $D(\tau)$ in the equation and estimate

$$V = \alpha_0 + \alpha_1 \widetilde{OC} + \alpha_2 TIME + \alpha_3 D(\tau) + e$$

where D is a dummy variable that is defined parametrically on an unknown quarter, τ, such that it equals 0 before quarter τ and 1 starting with that quarter. To obtain point estimates of V^* that are relevant for forecasting, starting in 1990 we estimate this regression recursively by adding one additional quarter of data at a time. Further, in each step, we allow the regression to select the quarter in which the shift in the intercept may have occurred, τ, which fits the data best.

Using this simple technique we obtain recursive estimates of V^* which could have been used to construct more accurate equilibrium prices than those obtained using the constant V^* assumption. We carry out this procedure twice, once to obtain a series which allows for a one-time break in V^* from a constant to a new level, and a second time allowing also for the presence of a time trend in the regression.

2.2 Using Regression Trees to Determine V*

Another way of endogenously generating time-varying values of V* involves regression trees as described in Clark and Pregiborn (1991).[5] Application of the technique in our setup involves a binary recursive partitioning of the determinants of velocity, which we specify to be opportunity cost, OC, and time, $TIME$. In particular, if X is the space of OC and $TIME$ over the sample, then the recursive algorithm partitions X into homogeneous rectangular regions x such that the conditional distribution of velocity given x does not depend on the specific values of OC and $TIME$, i.e., the tree is a step function that takes

[5] See also Härdle (1990) for a technical description.

on a constant value of velocity within each region. One way of representing the information from this recursive algorithm is as a tree (see figure 3) with various nodes or branch points and terminal nodes, called leaves, placed at the end of the nodes.

In general, tree-based models can be interpreted as regressions on dummy variables, where the dummy variables are endogenously, recursively determined. They have several advantages: (a) they are invariant to monotone transformations of the independent variables; (b) they are more adapt at capturing non-additive behavior; (c) they allow more general interactions among independent variables. In particular, trees allow us to model relatively abrupt shifts in relationships, perhaps enabling us to track changes in equilibrium velocity non-parametrically. In addition since the trees only depend on the ranks of the underlying data, in constructing our tree-based estimates we do not need to take a stand on whether the relationship between velocity and opportunity costs is linear or logarithimic. To illustrate these concepts, figure 3 displays two panels which represent alternatives ways of viewing the empirical estimates of a regression tree in which the GDP velocity of M2, $V2$, is the dependent variable and OC and $TIME$ are the independent variables, based on data from the entire sample from 1959:2 to 1997:4. The top panel of figure 3 shows the tree structure of the estimates; the full tree is an 11-way partition of the data, that is, it has 11 different terminal values of velocity depending on time and opportunity costs.[6] The top node, called the root, is designated by an oval; the average value of velocity for the whole period (1.748) is listed inside the oval. As we move down the tree the values of OC and $TIME$ are successively split into finer branches, as indicated on the paths between nodes of the tree. For example, the top node splits on the left going down the tree on the basis of $TIME$ for two splits until it splits on the right depending on whether OC is greater or less than 3.46. The eleven terminal nodes or leaves in this tree are represented by rectangular

[6] We used the S-PLUS system to implement the tree-based regression models. The default algorithm implemented there tends to estimate an "overly large tree" with roughly $N/10$ terminal nodes, where N is the sample size. In our application where $N = 155$, the number of nodes in the full model for the full period is 11. Such a size is in accord with the "best current practice" according to Clark and Pregiborn (1991, p. 415). However, as we shall see below, pruning the tree back (even severely) to have a smaller number of terminal nodes, appears to have very little effect on the estimates of V^* relevant for our inflation forecasts.

boxes.

The lower panel in figure 3 represents much of the same information in a way that reveals the time series structure more readily since time is plotted along the x-axis.[7] For example, in the first 6 periods of the sample period, the value of the regression tree is 1.74, that is, the value of the GDP velocity of $M2$, does not depend on the level of opportunity costs. Thereafter, however, the level of the opportunity cost does matter with a three-way split (3 values of velocity depending on the level of opportunity costs (OC) until the late 1980s, and a two-way split thereafter). As estimated, the tree provides a time-varying relationship between opportunity costs and velocity in which abrupt changes may occur.

The estimated regression tree does not necessarily give us a V^* concept directly, except in the first 6 periods, because the summary measure of velocity in most of the terminal nodes depends on the level of opportunity costs. There are two alternative ways to proceed to fashion an estimate of V^* from such a tree. We could prune the tree back to have a smaller number of splits or branches, and hope that the number of time periods in which velocity does not depend on opportunity costs increases. For example, figure 4 shows the splits for ten different possible splits from eleven (the full tree) down to two splits. As the number of splits or partitions becomes small, the regions in which velocity at a given time period or set of time periods does not depend on opportunity costs increase with only one region when there are four partitions, and no regions with either two or three partitions,

[7] These eleven rectangles represent contours of a *step* function that has a constant velocity value within each rectangular region in $OC \bigcap TIME$ space and is zero elsewhere. In a regression tree with a dependent variable that takes on numerical values as in our application using $V2$, the underlying statistical model posits that each observation, say y_i, is distributed as normal with mean μ_i and constant variance where the step function being determined by the regression tree can be thought of as the structural component, $\mu_i = \omega(x_i)$. The regression surface in a regression tree consists of a linear combination of such step functions, $\omega(x_i)$. Specifically, the overall tree surface for a tree with p leaves is $m(x) = \sum_{i=1}^{p} \mu_i I\{x \in N_i\} = \sum_{i=1}^{p} \omega(x_i)$, where I is an indicator function and the N_i are disjoint hyper-rectangles with sides parallel to the coordinate axes. The deviance function is defined to be $D(\mu_i; y_i) = (y_i - \mu_i)^2$. At a given node, the mean parameter μ is constant for all observations. The maximum-likelihood estimate of μ is given by the average of the values at the node, that is within the hyper-rectangle associated with the node. The deviance of a node is defined as the sum of the deviances of all observations in the node and is identically equal to zero when the y's within the hyper-rectangle are the same, and it increases otherwise. Splitting the data into finer partitions involves sorting the observations on a particular variable or set of variables into left (smaller observations) and right groups (larger observations) with the tree algorithm selecting such splits to maximize the change in the deviance.

the top panels in the figure. Alternatively, when the partitions depend on the opportunity cost, in each quarter we could select the value corresponding to the historical average of the opportunity cost as our estimate of V^*.

The two-way partition gives us a representation of V^* as having a constant value a little over 1.7 until the early 1990s when it shifts up to a value between 1.9 and 2.0. The three-way partition adds an earlier step in 1978. Since the tree regression does not have the capability of directly representing a linear time trend, the presence of such a trend in the velocity data would be captured by splitting velocity in the middle of the full sample period, which is what our estimates do in the sample midpoint of 1978. Most notable, at least given our interest in forecasting inflation in the 1990s, is the fact that the right-most part of the tree, corresponding to this more recent period, remains invariant to the number of partitions that are used so that our V^* estimates do not appear to depend on the number of leaves in the tree, that is the results hold under both parsimonious and profligate tree parameterizations.

2.3 Estimation

In figure 5, we provide estimates of V^* obtained from recursive estimation using the techniques described above. As indicated there, we are interested in determining how the various estimates of V^* evolve in real time as they are fed more observations, so we will update the estimates quarterly as each new observation becomes available.[8] The four alternative time-varying estimates shown with dashed and dotted lines correspond to the cases of a constant with one intercept shift, (Estimate A); a constant, constant trend and one intercept shift (Estimate B); the simplest regression tree specification with just 2 splits (Estimate C) and the full regression tree evaluated at the average opportunity cost, (Estimate D).

Three things are noteworthy here. First, even at the beginning of 1990, (when we first consider the additional flexibility in the specification of V^* than the constant estimate) the point estimate of equilibrium velocity for the quarters shown appears somewhat higher than

[8]These calculations are only exercises in one sense: We do not deal explicitly with the fact that in real time many of our data sources are preliminary.

the fixed constant estimate. This result is not surprising since, as already mentioned, a slight upward trend would appear to provide a better characterization of V^* during the 1960 - 1988 period over which the constant V^* is computed. Second, and most important, compared with the fixed V^* estimate all four alternative recursive estimates show a dramatic upward shift in equilibrium velocity during the 1992-1994 period. The exact time at which this change is most noticeable differs quite significantly in the four estimates, with estimate B indicating a major shift as early as 1992, D in 1993, and estimate C detecting a large change only in late 1994. Third, despite the important timing differences in the four estimates during the 1992-94 period, it appears that by the end of the sample the four alternative estimates are in substantial agreement with one another.

Overall, while it would be difficult to judge the relative merits of these estimates of V^* with such a short sample, the recursive estimates should be noted more for their similarity relative to the constant V^* estimates, instead of their differences. The next step, then, is to use these estimates to construct the implied time series for P^* and associated inflation forecasts.

3 Inflation forecasts

Using the alternative concepts of V^* shown in figure 5, we next turn our attention to the inflation forecasts corresponding to the implicit alternative estimates of P^*. We concentrate on the simplest specification of the inflation equation in Hallman, Porter and Small which can provide one-year-ahead forecasts of inflation (their equation 14):

$$\Delta\pi_t = \alpha(p_{t-1} - p^*_{t-1}) + u_t$$

Here π_t denotes inflation over the four quarters starting with quarter t and p and p^* denote the natural logarithms of P and P^* respectively. Using their data covering the period from 1955 to 1988, Hallman, Porter and Small estimated α to be -0.22 with a standard error of 0.044.[9] Because of changes in the definitions of the variables and sample coverage, however,

[9]The annual model in Hallman, Porter, and Small uses only non-overlapping fourth quarter observations on inflation and the price gap. Since we are interested in detecting changes in V^* in real time, we adopt a

we need to reestimate this equation. Ending the sample in 1988 for comparability, we use overlapping quarterly data from 1960:1 to 1988:4. This yields an estimate of α of -0.16 with a standard error of 0.041 (obtained using the asymptotic correction for the moving average error term).[10]

Figure 6 shows the in-sample inflation forecasts obtained from this equation as well as the out-of-sample forecasts based on the assumption of constancy for equilibrium velocity. As can be seen, while the forecast appears to move closely with realized inflation for most of the sample, since about 1992 forecasted inflation diverges from actual inflation with the forecast error becoming progressively worse until about 1995. The pattern of errors suggests that over this period, P^* was constructed using an estimate of equilibrium velocity which was likely systematically smaller than it should have been. This assessment, is in agreement with the estimated upward shifting V^* implied by the relationship between velocity and opportunity costs.

Staring in 1990, figure 7 provides the alternative forecasts based on the four time-varying concepts of V^* shown in figure 5. As can be seen, allowing for the time variation in the estimate of V^* during this period would have yielded substantially better forecasts of inflation especially after 1992. Of the four alternative estimates, it appears that Estimate B (obtained by estimating a one-time shift in V^* and allowing for a time trend) and Estimate D (obtained from the full regression tree) would have resulted in the smallest forecast errors during this period. As shown in table 1 for these two forecasts, the mean error was one-quarter of a percent or less and the mean absolute error was about half a percent. Indeed, this performance is better than the in-sample performance of the constant V^* specification over the earlier period—as shown in the memo item in the table.

Finally, the simplicity of the formulation we are examining allows us to easily characterize the specification error arising from incorrectly assuming a constant V^* specification

setup that allows us to use all quarterly observations.

[10]That is, since we are using overlapping, year-over-year inflation rates as the dependent variable in the regression, the proper estimation of the asymptotic variance of α needs to take the error structure induced by using overlapping observations into account.

when the correct specification should account for a one-time shift in velocity. If there were a one-time shift in equilibrium velocity from V_1^* to V_2^*, then the associated forecast error in forecasting inflation using the original V^* estimate, V_1^*, would be

$$-\alpha(log(V_2^*) - log(V_1^*))$$

Since the estimated $\alpha = -.16$ and the shift in V^* appears to be from about 1.7 to 2, the implied inflation forecast error is about two and a half percent per four-quarter period. Indeed, as shown in figure 7, this is consistent with the magnitude of the recent forecast errors corrresponding to the forecast obtained from the constant V^* specification.

3.1 Alternative Non-Structural Approaches

The P^* framework depends on only two parameters, V^*, the equilibrium velocity, and α, the coefficient indicating what fraction of the lagged gap between the log of the price level and the log of P^* is expected to close over the next four quarters. The forecasting results suggest that of these two parameters, the behavior of only one has been problematical, namely V^*. By relaxing the assumption that this parameter is constant but leaving the other parameter fixed we were able to restore the accuracy of the model. The specification analysis presented at the end of the previous section suggests that a trial-and-error process could also lead to the eventual discovery that setting $V^* = 2.0$ would have corrected the forecasting problems of P^* framework.

This result raises the question how well a simpler approach that updates V^* based on the recent inflation error without using the additional information we bring to bear, namely the opportunity cost of money, would have performed. As is evident from figures 5 and 6 such updating would have created wildly divergent estimates of V^*; estimates of V^* would have tended to fall in 1991-92 before eventually moving upward later in the period. A less extreme non-structural approach lets us evaluate the importance of opportunity costs to our forecasting results. Suppose as an alternative to Estimates A and B we dropped the term in opportunity costs, OC, from the regressions but computed everything else in

Estimates A and B as above. As would be expected, this method would also eventually pick up the upward movement in the equilibrium velocity. However, the recognition is delayed (by at least one year), resulting in a marked deterioration in the forecast accuracy of this alternative.

4 Conclusions

Our paper lays out a structural strategy constructing real-time estimates of V^*. The estimates we obtain go quite far in restoring the forecasting accuracy of the P^* model. That the stability of the P^* approach could be restored in the face of fairly massive changes in the financial service industry that occured over this period, suggests to us that this simple dynamic version of the quantity equation is still worth having in the monetary practitioners' toolkit.

Data Appendix

In constructing the data for this paper we follow Hallman, Porter and Small (1991) but make allowances for three changes that need to be taken into account to update the results. First, we use the new redefined version of M2. Early in 1996 the Federal Reserve made a small change in the definition of M2, excluding overnight RPs and Eurodollars from this broad measure but including them in M3. (See Whitesell and Collins, 1996.) Although this change does not appear to have changed the statistical properties of the aggregate in major ways it is noticeable in our analysis in two ways. First, it resulted in an overall increase in the velocity of the aggregate. Second, as overnight RPs and Eurodollars were not nearly as negligible late in our sample period as earlier, it makes estimates of a positive overall trend in the equilibrium velocity of M2 more noticeable. In our estimation, we use the quarterly seasonally adjusted data for the redefined M2 series which are available on a consistent basis from 1959. We construct the opportunity cost of M2 as the difference between the yield on the three-month Treasury bill and the average rate paid on M2 balances. On the BEA side of the ledger, there were two important relevant changes regarding income and price concepts used in basic macro discussions of policy. In 1992 the BEA switched first to a GDP income concept instead of GNP, and more recently the use of a chain-type price deflator for GDP. Given the changes adopted by the BEA, we modified the concepts of P and Q commensurately. Finally, for potential output, Q^*, we rely on the estimates made by the Congressional Budget Office. (*The Economic and Budget Outlook*, United States Government Printing Office, 1998.)

References

Clark, Linda A. and Daryl Pregiborn, "Tree-Based Models" in John M. Chambers and Trevor J. Hastie, editors, *Statistical Models in S*, Wadsworth and Brooks, 1991.

Feinman, Joshua, and Richard D. Porter, "The Continuing Weakness in M2," Federal Reserve Board, Finance and Economics Discussion Series, No 209, September 1992.

Hallman, Jeffrey, Richard D. Porter, and David H. Small, "Is the Price Level Tied to the M2 Monetary Aggregate in the Long- Run?," *American Economic Review*, (September 1991), p. 841-858.

Härdle, Wolfgang, *Applied Nonparametric Regression*, Cambridge University Press, 1990.

Koenig, Evan F., "The P* Model of Inflation, Revisited," Federal Reserve Bank of Dallas Working Paper 94-14, September 1994.

Moore, George R., Richard D. Porter, and David H. Small, "Modeling the Disaggregated Demands for M2 and M1: The U.S. Experience in the 1980s" in *Financial Sectors in Open Economies: Empirical Analysis and Policy Issues*, edited by Peter Hooper et al, Board of Governors of the Federal Reserve System, Washington DC, 1990.

Orphanides, Athanasios, Brian Reid and David H. Small, "The Empirical Properties of a Monetary Aggregate That Adds Bond and Stock Funds to M2," *The Federal Reserve Bank of St Louis Review*, 76(6), November/December 1994, p. 31-51.

Whitesell, William C., and Sean Collins, "A Minor Redefinition of M2," Federal Reserve Board, Finance and Economics Discussion Series, 96-7, February 1996.

Table 1: Summary Error Statistics from Inflation Forecasts

Statistic	Constant V^*	Time-Varying V^* Estimates A	B	C	D	Memo Item: Constant V^* 1961:1–1989:4
Mean Error	1.65	0.44	0.22	0.54	0.03	0.15
Standard Deviation of Error	1.14	0.61	0.47	0.84	0.59	1.05
Mean Absolute Error	1.74	0.66	0.44	0.81	0.46	0.78

NOTES: The 28 out-of-sample forecast errors (1991:1 to 1997:4) are defined as inflation over the four quarters ending with quarter t minus the forecast for the corresponding period, in percent. Constant V^* is estimated as the mean of $V2$ over the period from 1960:1 to 1988:4; Estimate A allows for a one-time shift in the intercept; Estimate B allows for a time trend and a one-time shift in the intercept; Estimate C allows for a two-branch regression tree; Estimate D allows for a full regression tree evaluated at the average opportunity cost of holding M2.

Figure 1

Velocity and Opportunity Cost of M2

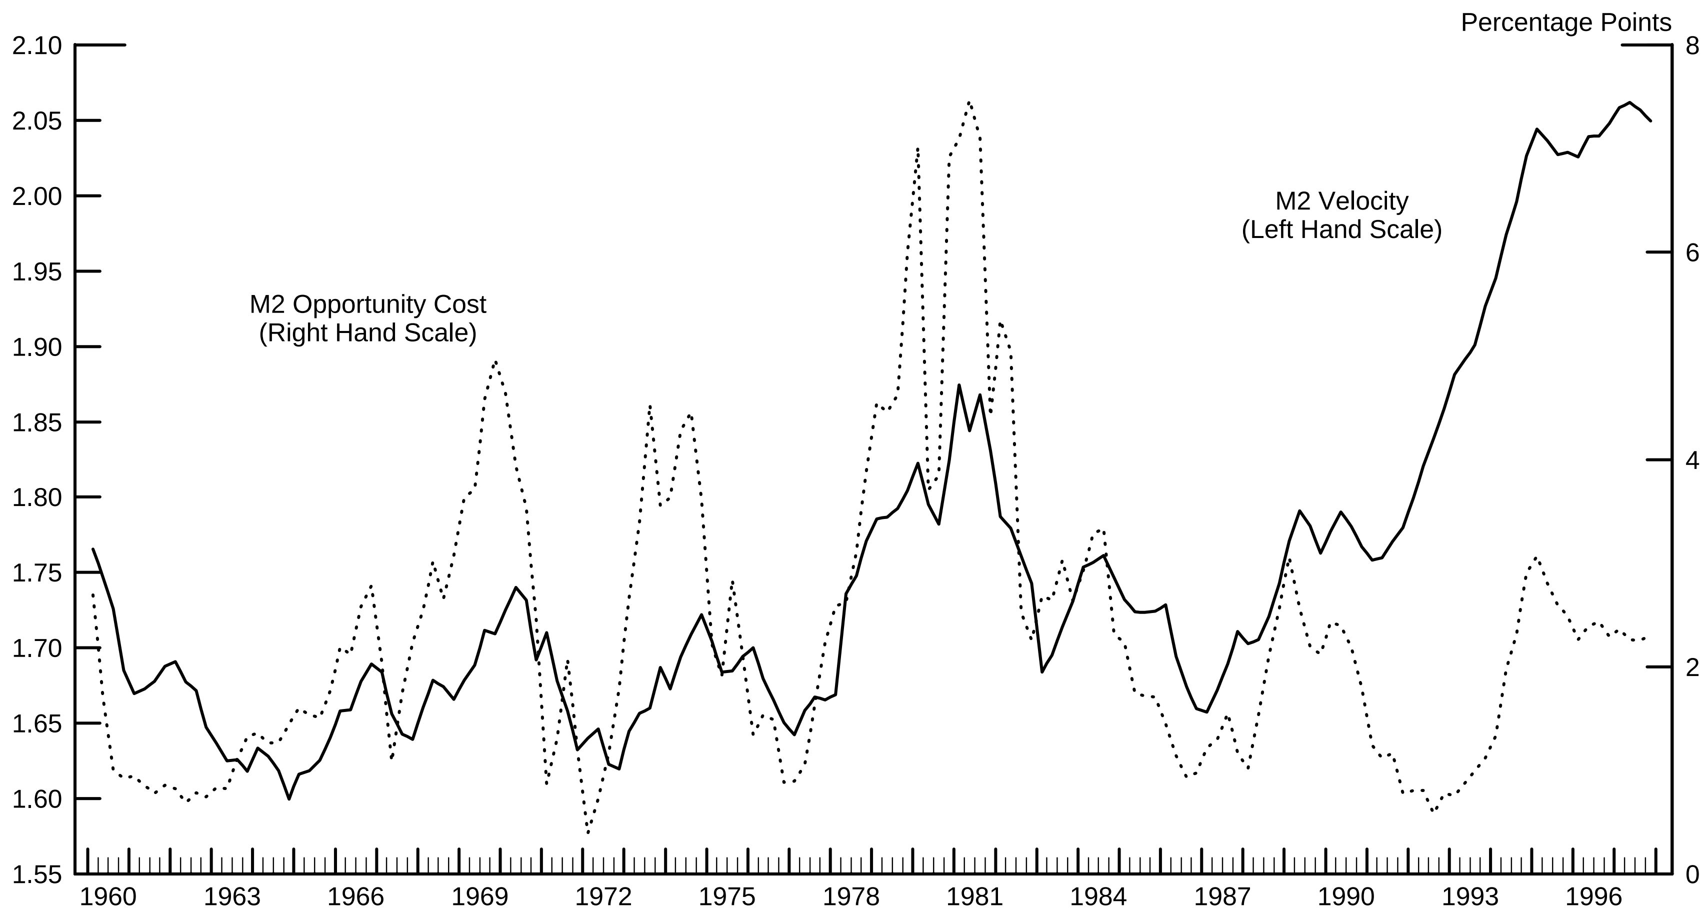

Figure 2

Scatter Plot of the Shift in M2 Velocity in Relation to its Opportunity Cost

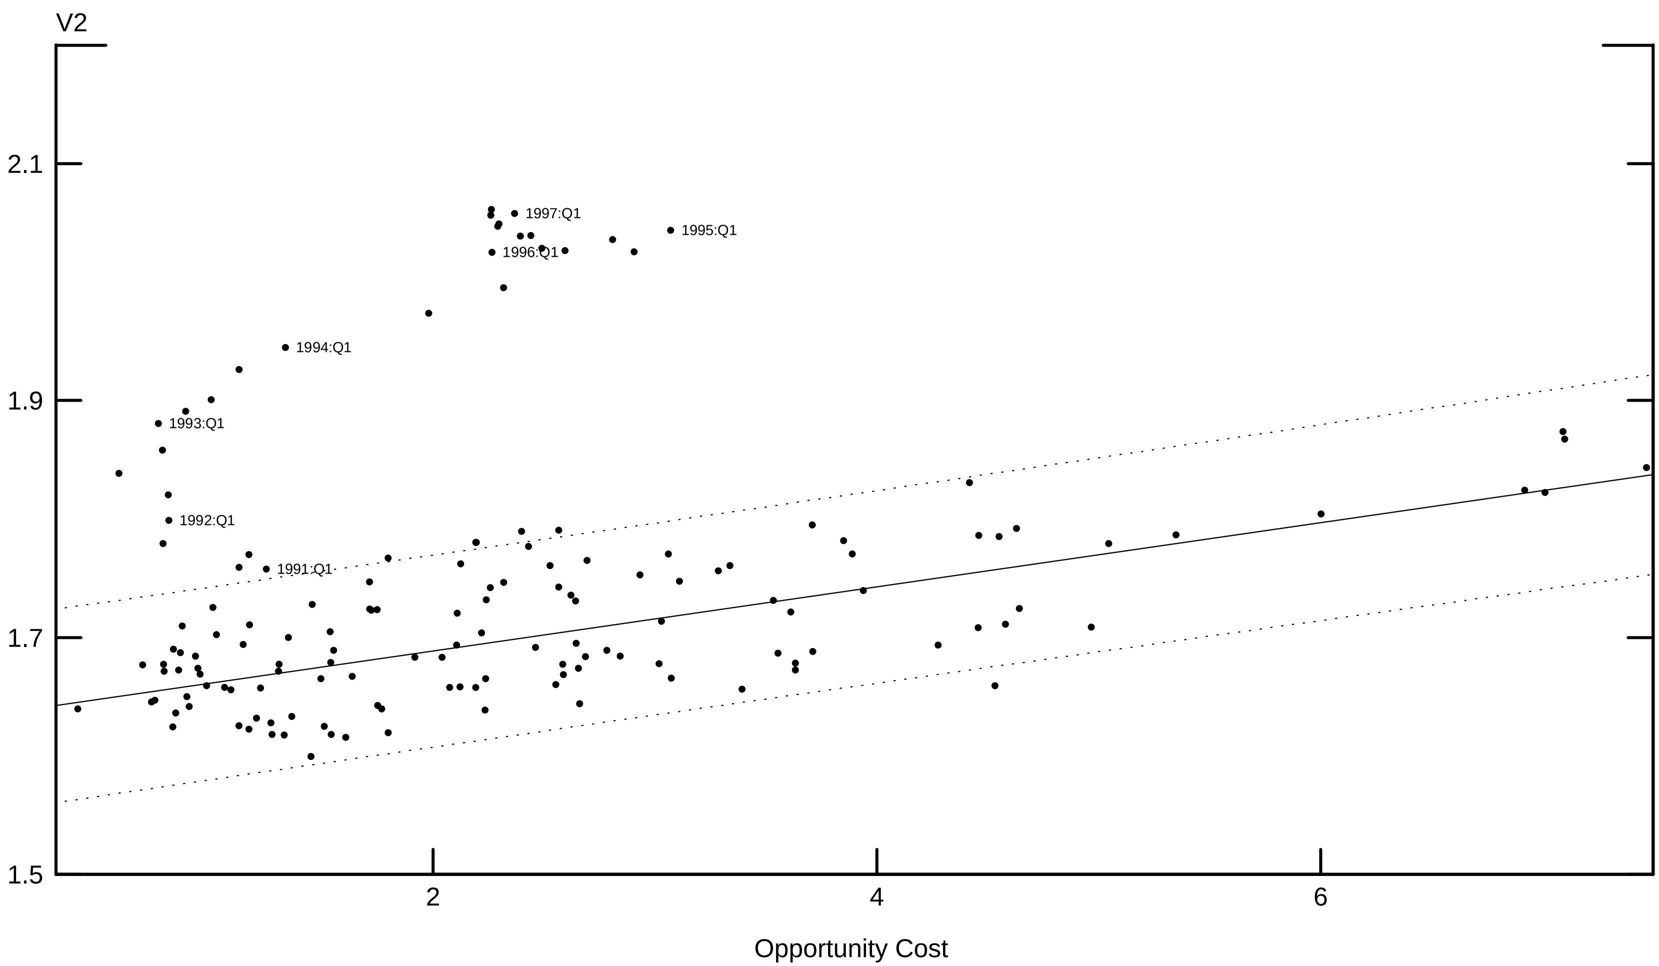

Figure 3

Regression Tree for Full Data Sample

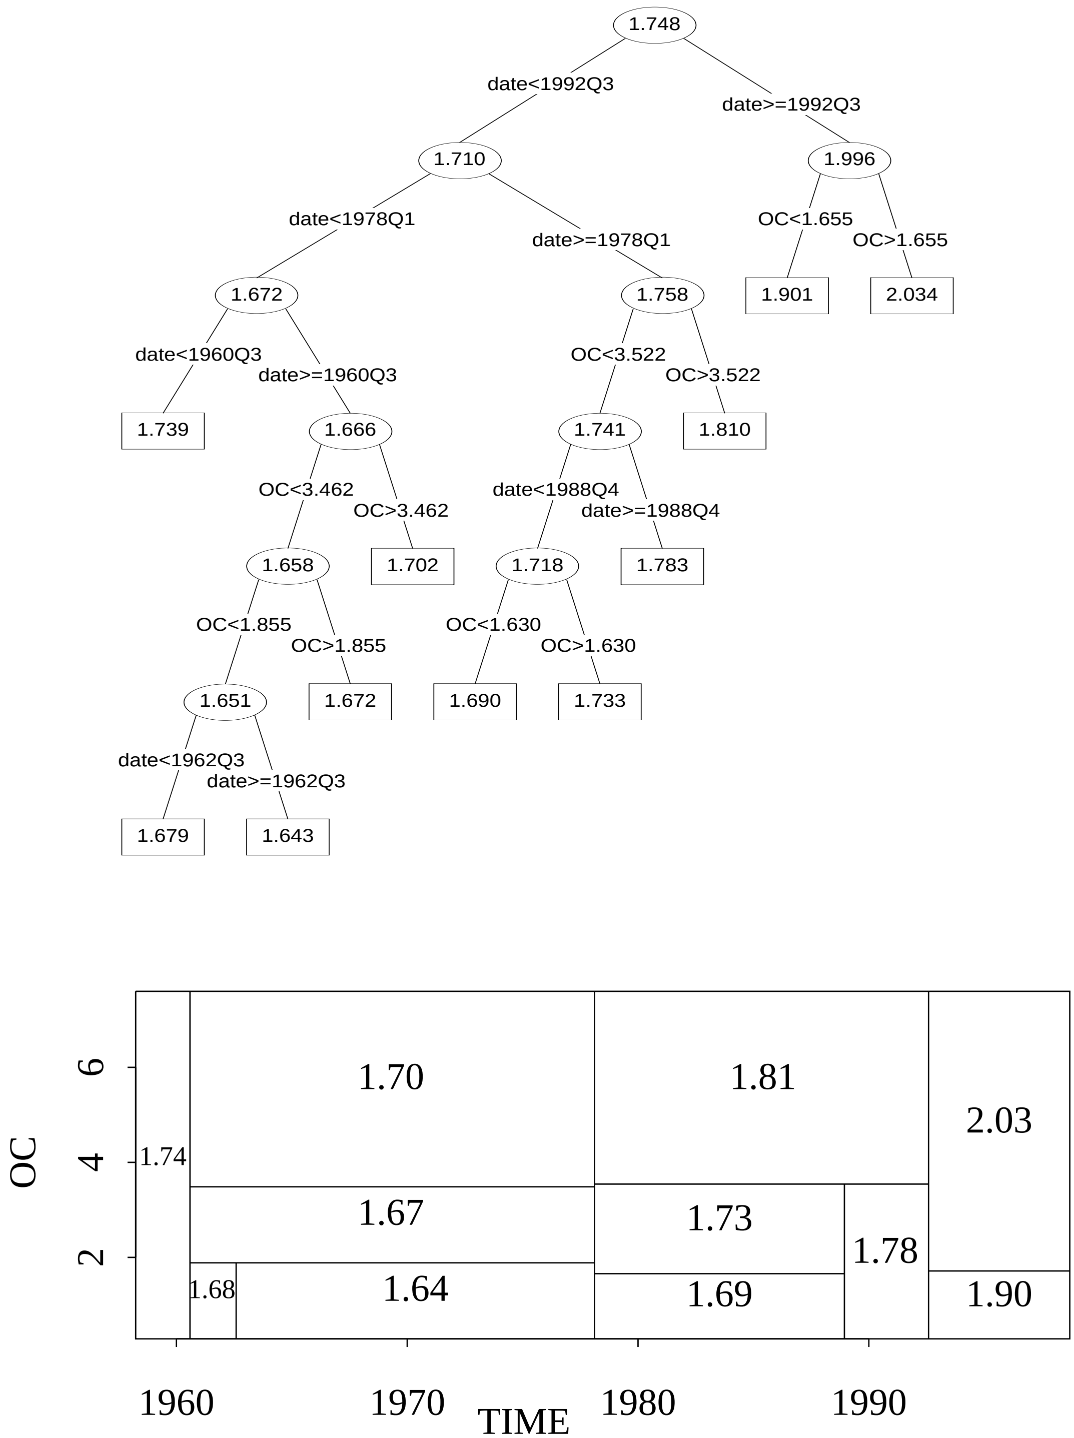

Figure 4

Tree Partitions

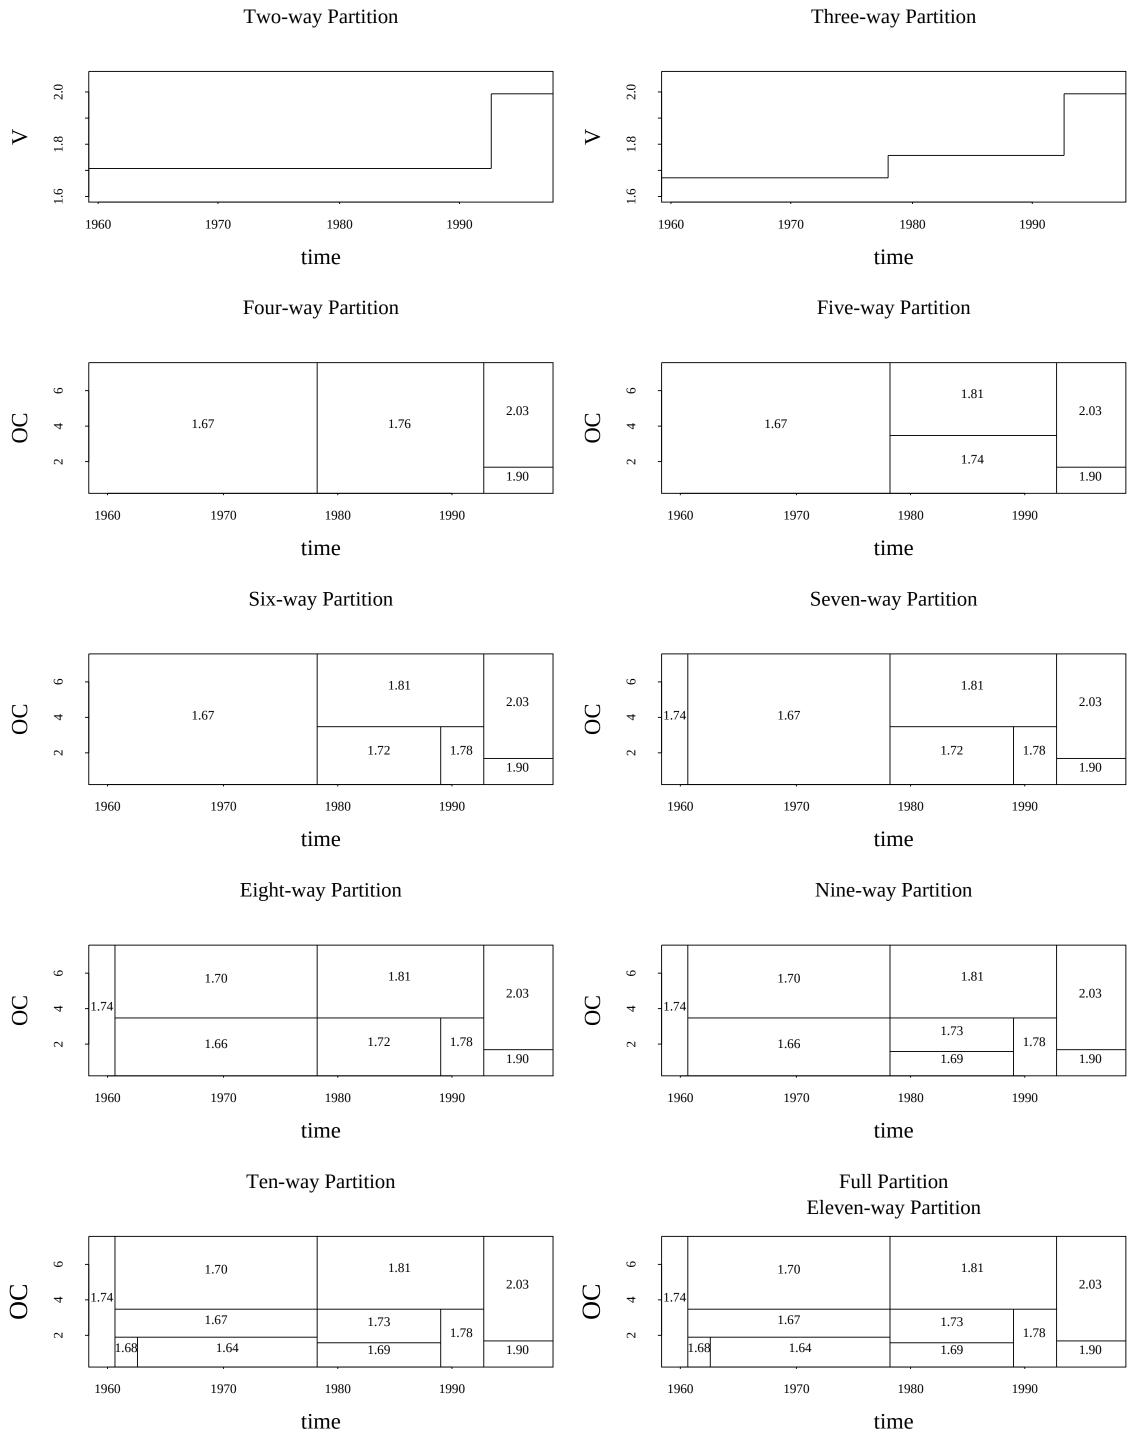

Note: The values inside the rectangles represent the estimated average value of V2 for the region corresponding to the rectangle.

Figure 5

Alternative Equilibrium Velocity Estimates

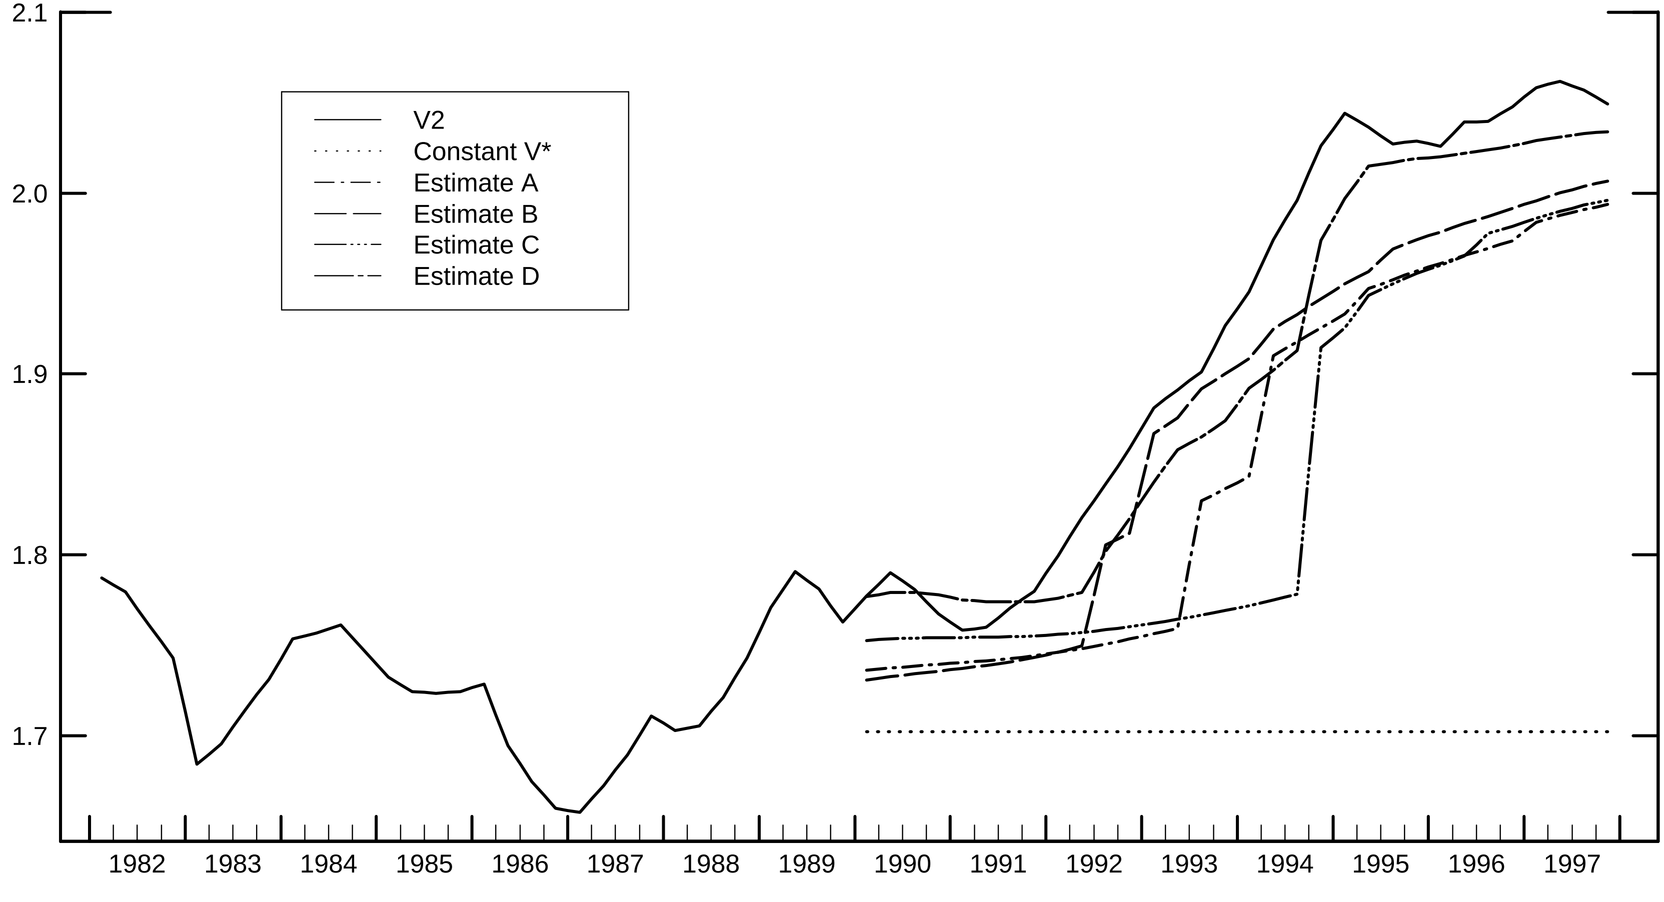

Figure 6

Inflation Forecast Assuming Fixed Equilibrium Velocity

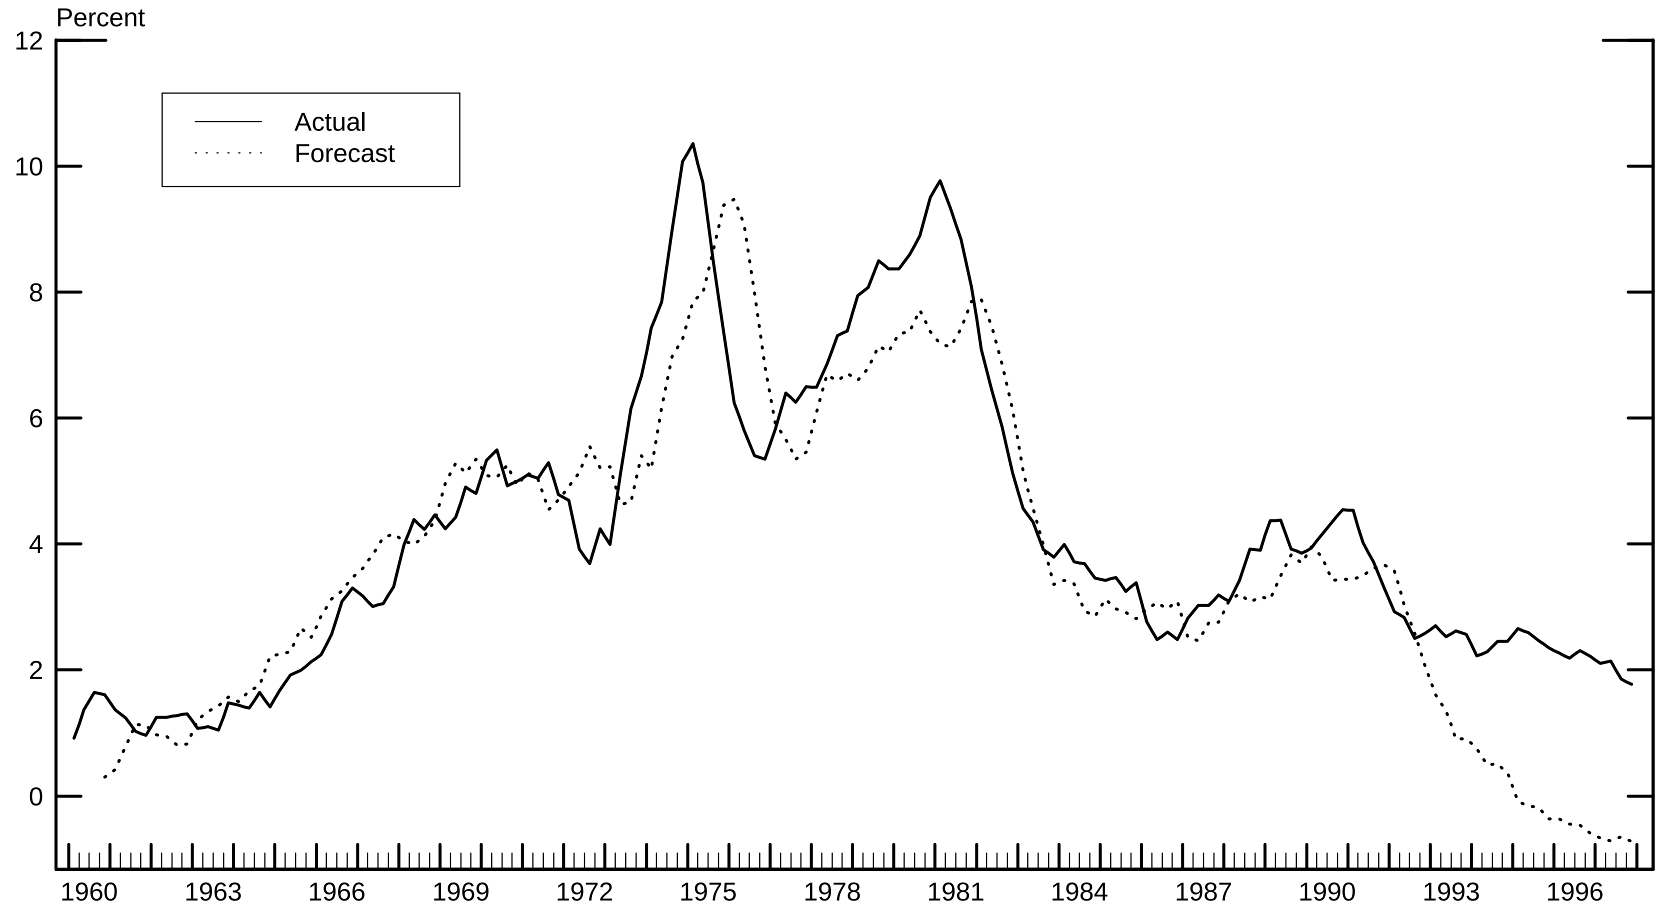

Figure 7

Alternative Inflation Forecasts