

NBER WORKING PAPER SERIES

HAS US MONETARY POLICY CHANGED?
EVIDENCE FROM DRIFTING COEFFICIENTS AND REAL-TIME DATA

Jean Boivin

Working Paper 11314
<http://www.nber.org/papers/w11314>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2005

I thank Ben Bernanke, John Donaldson, Marc Giannoni, Hong Li, Rick Mishkin, Athanasios Orphanides, Mark Watson and an anonymous referee, as well as participants at the Columbia Macro Lunch and NBER Monetary Economics Meeting for useful comments. I thank Athanasios Orphanides for sharing his real-time output gap data and Dean Croushore for assistance with the Greenbook data. Mauro Roca provided helpful research assistance. Financial support from the National Science Foundation (SES-0214104) is acknowledged. All errors are my own. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Jean Boivin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Has US Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data
Jean Boivin
NBER Working Paper No. 11314
May 2005
JEL No. E52, C32, C50

ABSTRACT

Despite the large amount of empirical research on monetary policy rules, there is surprisingly little consensus on the nature or even the existence of changes in the conduct of U.S. monetary policy. Three issues appear central to this disagreement: 1) the specific type of changes in the policy coefficients, 2) the treatment of heteroskedasticity, and 3) the real-time nature of the data used. This paper addresses these issues in the context of forward-looking Taylor rules with drifting coefficients. The estimation is based on real-time data and accounts for the presence of heteroskedasticity in the policy shock. The findings suggest important but gradual changes in the rule coefficients, not adequately captured by the usual split-sample estimation. In contrast to Orphanides (2002, 2003), I find that the Fed's response to the real-time forecast of inflation was weak in the second half of the 1970's, perhaps not satisfying Taylor's principle as suggested by Clarida, Galii and Gertler (2000). However, the response to inflation was strong before 1973 and gradually regained strength from the early 1980's onward. Moreover, as in Orphanides (2003), the Fed's response to real activity fell substantially and lastingly during the 1970's.

Jean Boivin
Columbia Business School
Uris Hall, Room 821
3022 Broadway
New York, NY 10027-6902
and NBER
jb903@columbia.edu

1 Introduction

Since the end of the great inflation of the 1970's, the U.S. has experienced only mild inflation and a considerable reduction in the volatility of real activity.¹ An important question is what, ultimately, brought about these improved economic outcomes.

One possible explanation is that the Federal Reserve changed its behavior in an important way. From a narrative perspective, this seems fairly uncontroversial, although the reasons for the changes might be more so.² The monetarist influence of the late 1970's, culminating in the "experiment" of 1979-1982 under chairman Volcker, is often cited as an example. It has also been argued that the composition of the Federal Open Market Committee (FOMC) and other political pressures might have fundamentally altered monetary policy decisions over time. Further, the Fed might have learned from past experiences, changed its views about the economy and, accordingly, modified its conduct of monetary policy.³

However, despite the large amount of empirical research on monetary policy rules, there is surprisingly little consensus on the nature or even the existence of these hypothesized changes.

Clarida, Galí and Gertler (2000) provide empirical evidence of important changes in the U.S. conduct of monetary policy over the last forty years. In particular they find that while monetary policy accommodated inflation in the 1970's, this drastically and suddenly changed with the ap-

¹See among others, McConnell and Perez-Quiros (2000) and Stock and Watson (2002).

²See Romer and Romer (2002) and Sargent (2002)'s discussion.

³Romer and Romer (1989) argue that the Fed's concerns for inflation and unemployment may have changed significantly with the arrival of Volcker. The importance of changing political pressures and ideologies on the conduct of monetary policy has been advanced by Froyen (1974), Potts and Luckett (1978) and Blinder (1986) among others. Finally, DeLong (1997), Taylor (1998), Sargent (1999) and Romer and Romer (2002) suggest learning or changing beliefs stories – not necessarily continuous according to the latter – that could explain an evolution in the Fed's response to the economy.

pointment of Volcker in 1979. They emphasize that the pre-Volcker conduct of monetary policy did not satisfy the so-called Taylor principle, so that a given increase in inflation was typically associated with a smaller increase in the nominal interest rate, thus resulting in a lower real interest rate. They argue that such behavior on the part of the Fed did not rule out non-fundamental fluctuations. Lubik and Schorfheide (2004) reach similar conclusions while formally testing for indeterminacy in the context of an estimated general equilibrium model.

This evidence has been challenged along three dimensions. First, Cogley and Sargent (2001) argue that the Fed's evolving views about the economy might have been more gradual. Cogley and Sargent (2001) specify a reduced form vector autoregression (VAR) with drifting coefficients which produces results, they argue, broadly consistent with those of Clarida, Galì and Gertler (2002), although the time variation is not clearly of a discrete nature. Second, Sims (2001) and Stock (2001) argue that Cogley and Sargent's (2001) conclusion might be contaminated by the presence of heteroskedasticity. This claim is supported in part by the evidence of Sims (1999) and Sims and Zha (2004), who find that most of the observed changes between the pre- and post-Volcker periods can be attributed to changes in the variance of the shocks. Yet, Cogley and Sargent (2005), extending their earlier model to allow for heteroskedasticity in the reduced form VAR shocks, still find important changes in the implied policy rule parameters. Finally, the evidence of changes in the conduct of monetary policy has also been challenged on the ground that it does not properly account for real-time issues, both in terms of the vintage of data used and the misperceptions about potential output. For instance, Orphanides (2001) argues that estimating monetary policy rules on *ex post* data, which were not available to policymakers in real-time, can lead to a very distorted picture of the historical conduct of monetary policy. Orphanides (2002) concludes that when the Clarida, Galì and Gertler (2000) rule is estimated using real-time data on inflation and

unemployment the conduct of monetary policy in the 1970's is not greatly different from the one thereafter. Orphanides (2003) reaches a similar conclusion, using real-time estimates of the output gap, except that he finds evidence of a reduction in the response to real activity.

The conclusion on whether or not systematic monetary policy changed thus appears to be particularly sensitive to: 1) the way the time variation of the coefficients is modeled, 2) the treatment of heteroskedasticity and 3) whether or not the real-time nature of the data is taken into account.

The goal of this paper is to contribute to this debate by considering an empirical framework that accounts simultaneously for these three potential issues. Given that there is considerable uncertainty about the true nature of the changes in monetary policy, I entertain the possibility that the conflicting evidence reported in the literature might be stemming from the use of empirical framework not sufficiently flexible to allow for a rich evolution of monetary policy. I consider a forward-looking Taylor rule with time-varying coefficients. Following the pioneering work of Cooley and Prescott (1978), the time variation is modeled as driftless random walks, and is estimated using the median-unbiased estimator proposed by Stock and Watson (1998). I argue that this is an appealing and flexible way of uncovering changes in policy. Importantly, the estimation recognizes explicitly the presence of heteroskedasticity, allowing for changes in the variance of the policy shock. Finally, given the possibility that the existing evidence suggesting important changes in monetary policy might be due to an inappropriate account of the real-time nature of monetary policy, I follow Orphanides (2001, 2002, 2003) and model the forward-looking behavior directly using the forecasts produced in real time by the staff economists at the Board of Governors, published in the Greenbook.

The findings suggest substantial changes in the Fed's response to inflation and real activity, not properly accounted for by the typical split-sample approach. In contrast to what is typically

assumed in the literature, a single regime does not appear to properly characterize the pre-Volcker conduct of monetary policy. The response to inflation was strong until around 1974, but then fell dramatically in the second half of the 1970's, perhaps not satisfying the Taylor principle. In parallel, the response to real activity weakened continuously throughout the 1970's. Moreover, the transition under Volcker appears to have happened gradually, with the most important changes occurring between 1980 and 1982 and not in 1979, as is often assumed. From the mid-1980's onward, monetary policy consistently responded strongly to inflation and weakly to real-activity. More than the real-time data issue, the failure to properly account for this rich evolution of monetary policy appears to be the reason why Clarida, Galì and Gertler (2000) and Orphanides (2002, 2003) find conflicting results.

The paper proceeds as follows. Section 2 presents the specification of the forward-looking Taylor rule and describes the Greenbook forecasts. Section 3 specifies a time-varying parameter model and Section 4 describes the estimation. In section Section 5, I present the results and show their robustness to an alternative treatment of heteroskedasticity, various specifications of the forecasting horizons and different measures of real activity. In particular, to make sure that the main conclusions are not due to unknown particular features of the conditioning scenarios embedded in the Greenbook forecasts, I report results based on the current period real-time value of inflation and real activity. Section 6 concludes.

2 A forward-looking Taylor rule

A popular structural characterization of the monetary policy conduct is the forward-looking version of Taylor’s (1993) rule:

$$\begin{aligned} R_t &= \alpha_t + \beta_t \pi_{t+h_\pi|t} + \gamma_t x_{t+h_x|t} + \rho_t(L)R_{t-1} + \varepsilon_t \\ &= \Phi_t' Z_t + \varepsilon_t \end{aligned} \tag{1}$$

where R_t is a short-term interest rate and $\pi_{t+h_\pi|t}$ and $x_{t+h_x|t}$ are the Fed’s expectations of inflation and some proxy of the output gap, respectively at horizons h_π and h_x . Note that all coefficients have a t subscript to emphasize that they are potentially time-varying. The vector Φ_t denotes the collection of these parameters and Z_t the corresponding regressors.

This specification has been widely used both as a successful empirical description of monetary policy and as a normative prescription.⁴ Taylor (1993) argues that a constrained version of this rule – with no lags of R_t , $h_\pi = h_x = 0$, the contemporaneous value of linearly detrended GDP, and time invariant parameters – fits the Greenspan’s conduct of monetary policy between 1987 and 1992 impressively well. This led to numerous applications of the Taylor rule on U.S. data, and for other countries as well.⁵

Various approaches have been employed to estimate such a rule. One is to explicitly model and estimate the expectation process, as in Clarida, Gali and Gertler (2000). They assume that the Fed has rational expectations, which implies a set of moment conditions that can be used in a GMM framework.⁶ Their split-sample estimates suggests important changes in the conduct of monetary

⁴See for instance the papers collected in Taylor (1999) and the references therein.

⁵For instance, Clarida, Gali and Gertler (1998) show the empirical success of (1) for Germany, Japan, UK, France and Italy, over the post-1980 period.

⁶A much earlier reference on the direct estimation of the Fed’s expectations and of their inclusion, through IV estimation, in the policy rule, is Abrams, Froyen and Waud (1980).

policy. Orphanides (2001, 2002, 2003) argues, however, that this evidence, based on *ex post* data, does not properly account for the information that was actually available to policy makers at the time decisions were made. If the data available to policy makers and their perception about the output gap were, in real time, considerably different from what we now know, the historical characterization of monetary policy based on *ex post* data could be distorted. In fact, using real-time forecasts of inflation and real activity to estimate equation (1), Orphanides (2002, 2003) finds that the conduct of monetary policy in the 1970's is not greatly different from the one thereafter, thus contradicting the conclusion reached by Clarida, Galì and Gertler (2000).

Since this difference in result has been attributed in part to the treatment of the real-time nature of monetary policy, it is important to acknowledge this issue in the empirical framework. I thus follow Orphanides and use, as a proxy for the Fed's expectations, the forecasts computed by the staff of the Federal Reserve Board of Governors before each FOMC meeting and published in what is known as the Greenbook.⁷ As Orphanides argues, one of their key appeal is that they do not assume knowledge of the subsequent revisions of these data, embedded in the *ex post* data typically used. Moreover, they might contain more information than small-scale empirical models, such as VARs. As illustrated by Romer and Romer (2000), the Fed might have an informational advantage over the private sector that is reflected in the Greenbook. Second, as argued by Bernanke and Boivin (2003), a large set of conditioning information might be required to properly model monetary policy. In that respect, the Greenbook forecasts presumably include information from a wide range of sources, including the forecasters' "judgement", which is not otherwise directly measured.

⁷The Greenbook forecasts have been used originally in the estimation of policy rules by McNees (1986, 1992) and more recently by Orphanides (2001, 2002, 2003) and Romer and Romer (2003). Romer and Romer (2000) document their superior historical performance, relative to private sector forecasts, in forecasting inflation. For a discussion of the properties of the Greenbook forecasting errors, see Romer and Romer (2000, 2002) and Orphanides (2002).

Implicit in the estimation of (1) by OLS, as in Orphanides (2001, 2002, 2003), or with time-varying parameters, as in the present paper, is the assumption that the Greenbook forecasts are contemporaneously uncorrelated with the policy shock, ε_t . In particular, this requires that for any given period, the forecasts are not based on knowledge of the actual deviation of the policy setting from the one prescribed by the rule itself, i.e. the policy surprise. While some casual arguments could be made to support this assumption⁸, its validity can unfortunately not be checked directly, as there is no official detailed documentation of the conditioning scenarios underlying the Greenbook forecasts.⁹ One potential limitation of the direct estimation of Taylor rules using Greenbook data is thus that the endogeneity could introduce some bias in the policy parameter estimates. If this is the case, one should be careful before concluding that the inflation response, say, took a specific range of values at some point in the sample. However, two important observations should be made at this stage. First, as long as this endogeneity was systematic throughout the sample period, it could not be responsible for a low frequency evolution of the policy parameters that our empirical

⁸For instance, that the forecasts are produced by the Staff, before the FOMC meeting, independently of the actual decision subsequently taken.

⁹Reifschneider, Stockton and Wilcox (1997) provide some information on how these forecasts were constructed in last ten years at least. First, the forecasts are based on a set of conditioning assumptions, in particular about the expected path of monetary policy. Typically, this assumption is that the “nominal federal funds rate will remain unchanged over the next sixth to eight quarters”. By construction, they thus do not provide the best forecasting performance possible, but rather a benchmark summarizing the momentum already in place in the economy. This “neutral” assumption about the future path of monetary policy reflects the institutional division between the Board staff and the FOMC, and in particular the fact that the Staff aims to avoid being construed as making policy recommendations. This could suggest that for most of that period, the forecasts were not conditioned on the policy surprise. The second feature is that while econometric models, like the large-scale FRB/US model, are used in constructing the forecasts, there is a large “judgmental” component. As a result, the forecasts cannot be mechanically reproduced from any particular forecasting model.

framework is designed to uncover. Second, if inflation and real-activity respond only with a lag to changes in policy, as is often assumed in monetary VAR analysis using quarterly data¹⁰, the real-time current quarter values of inflation and real activity, i.e. $h_\pi = h_x = 0$, should not be affected contemporaneously by ε_t , and are thus not subject to the potential endogeneity problem. Based on this observation, the strategy adopted in this paper is to document the evolution of monetary policy using different specifications of the forecasting horizons in the Taylor rule. Since, as I show below, the characterization of the *evolution* of monetary policy is essentially the same for all horizons, and in particular for the specification using the forecasts for the current quarter, this suggests that the evidence of changes in monetary policy, the main conclusion of this paper, is not an artifact of this potential endogeneity issue.

3 Modeling the time variation

How did the conduct of monetary policy change and by how much? Answering these questions requires modeling the time variation in the parameters. In doing so, it is important to keep in mind that the conduct of monetary policy might have evolved in various ways. In particular, there is no guarantee that the implied changes in the parameters were discrete. For instance, time variation stemming from an evolving view of the Fed on the economy would suggest gradual and continuous drifts in the policy parameters. Even the Volcker “experiment”, which is most often thought as a one-time discrete change in policy, might have implied an adjustment process taking place over a few years. Furthermore, not all the parameters are necessarily changing exactly at the same date. If this is the case, split-sample estimates could be misleading. Unfortunately, standard

¹⁰See among others, Bernanke and Blinder (1992), Bernanke and Mihov (1998) and Christiano, Eichenbaum and Evans (2000).

stability tests, such as the Quandt (1960) likelihood ratio test, cannot discriminate between the different forms of instability: they all have power against various alternatives and more than one type of change might be present in the policy reaction function.

It thus seems useful to explore an alternative to the often assumed discrete break model. The time-varying parameter (TVP) model considered in this paper assumes instead that the policy parameters follow driftless random walks:

$$\Phi_t = \Phi_{t-1} + \omega_t$$

where $E[\omega_t] = 0$. This is the model proposed by Cooley and Prescott (1976), partly as a way to empirically account for the Lucas (1976) critique on the inappropriateness of stable econometric models for policy evaluations. It has been widely used in forecasting applications,¹¹ and Cogley and Sargent (2001, 2003) use this specification for the parameters of their reduced-form VAR model.

Given the various factors that might have influenced monetary policy, the main appeal of the TVP model is its ability to uncover changes of a general, and potentially permanent, nature, for each parameter separately. Of course, if the true change is a single discrete jump, and is simultaneous for all parameters, the TVP model would be mis-specified. But this is arguably a very specific case and even then, the TVP model can provide, as seems to be the case in the empirical application below, a useful approximation. Intuitively, the TVP estimator of Φ_t applies weights, decreasing the importance of observations farther away from t , whereas OLS gives equal weights to all observation.¹² As such, the TVP estimate would produce a smooth estimate of the discrete change, with the parameters starting to change before the actual break date. It is thus possible to determine whether the uncovered changes are consistent at all with a discrete break

¹¹See Stock and Watson (1996).

¹²Note that this down-weighting scheme implies, for the sample size considered in this paper, that observations at the beginning and at the end of the samples have very little influence on the mid-sample estimates of the parameters.

model.

Moreover, as soon as there is more than one discrete change, or different dates for different parameters, it might in fact be desirable to approximate the changes with the TVP model rather than estimating a multiple break model. One reason is that the multiple break model requires the estimation of many parameters, an estimate on each side of the breaks, for each coefficient. Another reason is that it requires an estimate of the break dates, which in macroeconomic applications often involves considerable uncertainty.¹³ Consequently, although such an estimation approach is feasible in principle, it tends to produce unsatisfactory results when the number of breaks is large relative to the sample size or the break dates are too close to each other. The TVP specification can approximate the multiple discrete changes in a parsimonious way, the number of estimated parameters being independent of the number of break dates.

The TVP model thus appears to be a flexible, parsimonious and tractable way of uncovering the nature of the time variation.

4 Heteroskedasticity-robust TVP estimation

Provided with an estimate of the variance of ω_t , the time series of the parameters, $\{\Phi_t\}$, can be obtained using the Kalman filter or equivalently, by generalized least squares. All the parameters of the model, including the variance of ω_t , can be estimated jointly by maximum likelihood estimation (MLE), using the Kalman filter to construct the likelihood function from its forecasts error decomposition representation. But if the variance of ω_t is small, its maximum likelihood estimate tends to be biased toward 0.¹⁴ This is potentially an important problem in the context of macroeconomic relations because, even though the time variation in the parameters is found to be statistically

¹³See for instance Stock and Watson (2002). This is also a feature of the results presented below.

¹⁴That is, the MLE has a (large) point mass at 0. See Stock and Watson (1998) for a discussion of this point.

significant, and the changes over the sample are economically important, they might still imply a small period to period variation in the parameters. As we show below, this is indeed the case for the current application.

This paper thus considers an alternative way of estimating the variance of ω_t using an estimation approach proposed by Stock and Watson (1998). Their strategy is explicitly designed to account for the deficiency of the MLE when the parameters' variance is small. Another important appeal is that it is computationally easy to implement.

The general idea of the estimation method is to exploit the fact that the distribution of a stability test, under the alternative of a TVP model, depends on the variance of the parameters. Hence, if all other parameters are known or consistently estimable, an estimate of the variance of these parameters can be inferred from a realization of a given stability test.

More specifically, we can rewrite the random walk specification for the time-varying policy parameters as:

$$\Delta\Phi_t = \omega_t = \tau v_t$$

where v_t and the policy shock, ε_t , are assumed to be mutually uncorrelated mean 0 random disturbances.¹⁵ The scalar τ governs the size of the variance of the parameters. To permit approximations based on non-degenerate asymptotic distributions, the case where it is small is formalized by making τ local to 0, that is:

$$\tau = \lambda/T.$$

In order to understand this parameterization, one needs to make the distinction between the small

¹⁵Stock and Watson (1998) and Boivin and Watson (1999) discuss the set of assumptions underlying the asymptotic theory of this estimator.

sample quantity τ , which we wish to estimate, and the asymptotic device that is used to estimate it. In fact, while asymptotically this parameterization implies that the variance of the parameters vanishes, for any given sample T and $\lambda \neq 0$, τ is a fixed non-zero quantity. Besides, since the estimation is based on the asymptotic distribution of a test statistic, this parameterization gives a meaning to τ being “small” asymptotically. More specifically, as it is the case when computing the local power of a test, the standard deviation of the parameters is assumed to shrink at the same rate as the sample size, which is sufficiently slowly to make the asymptotic distributions of the test depend on λ .

From Theorem 1 of Stock and Watson (1998), the asymptotic distribution of the Quandt (1960) likelihood-ratio test,¹⁶ QLR_T , under the TVP alternative, is a functional of Brownian motions, indexed by λ :

$$QLR_T \Rightarrow \sup_{s_0 < s < s_1} \left\{ \frac{\mathbf{h}_\lambda^0(s)' \mathbf{h}_\lambda^0(s)}{ks(1-s)} \right\}$$

where $\mathbf{h}_\lambda^0(s) = \mathbf{h}_\lambda(s) - s\mathbf{h}_\lambda(s)$, $\mathbf{h}_\lambda(s) = \mathbf{W}_1(s) + \lambda D \int_0^s \mathbf{W}_2(r) dr$, $D = \Sigma_{zz}^{-1/2} \Sigma_{vv}^{1/2} / \sigma^2$, $\Sigma_{zz} = E[Z_t Z_t']$, $\Sigma_{vv} = E(v_t v_t') = \Sigma_{zz}^{-1} \Omega \Sigma_{zz}^{-1}$, $\Omega = E[Z_t \varepsilon_t \varepsilon_t' Z_t']$ and \mathbf{W}_i are Brownian motions.¹⁷ Given D , an estimate of λ can be inferred from the realization of QLR_T . Stock and Watson (1998) use $D = I$ in their application, since it simplifies the computation of the estimator and makes the time variation comparable across parameters, if the variables are measured in standard deviation units. Since there is no clear alternative, I follow this approach.

¹⁶This is a standard approach to test the null of stability against the alternative of a structural break occurring at an unknown date. Andrews (1993) derived the asymptotic null distribution of this test.

¹⁷Stock and Watson (1998) assume that Z_t does not contain lagged dependent variables. However, the results of Li (2004) imply that the derivations of Stock and Watson (1998) are also valid when Z_t contains lagged dependent variables.

In order to make sure that the time variation found in the policy rule is not contaminated by the presence of heteroskedasticity in the policy shocks, an heteroskedasticity-robust version of this estimation is implemented. The estimation of $var(\Delta\Phi_t)$ proceeds as follows:

1. An estimate of λ is readily obtained by inverting the heteroskedasticity-robust version of the QLR_T test, which can be performed simply by using a lookup table, such as the one in Stock and Watson (1998);
2. Σ_{zz} is estimated with $T^{-1} \sum_{t=1}^T Z_t Z_t'$ and, based on the Stock and Watson (1998) results, under the local-to-zero time-variation in the parameters, Ω can be estimated by $T^{-1} \sum_{t=1}^T \hat{\varepsilon}_t^2 Z_t Z_t'$, which is the White estimator of $E[Z_t \varepsilon_t \varepsilon_t' Z_t']$, based on the OLS residuals, $\hat{\varepsilon}_t$;
3. $\hat{\Sigma}_{vv}$ is constructed from $\hat{\Sigma}_{zz}^{-1} \hat{\Omega} \hat{\Sigma}_{zz}^{-1}$.

Conditional on $var(\widehat{\Delta\Phi}_t) = \left(\hat{\lambda}/T\right)^2 \hat{\Sigma}_{vv}$, the time series $\{\hat{\Phi}_t\}$ is obtained in the same way as in the standard MLE approach, using the Kalman smoother. This effectively means that in the estimation of $\hat{\Phi}_t$, observations away from t , on both sides, are down-weighted. In fact, the larger the estimated variance of the parameters, the more rapidly these weights are declining as we get away from t .

Importantly, and again to make sure that the results are not contaminated by the presence of heteroskedasticity in the policy rule, the implementation of the Kalman filter allows for changes in the variance of the policy shock. The variance of the policy shock is estimated separately over different regimes, from the OLS residuals.¹⁸

¹⁸Note that, given the local-to-zero time variation assumed in this paper, this way of estimating the variance of the shocks is entirely consistent with the TVP specification, asymptotically.

5 Empirical results

5.1 Data and baseline specification

Equation (1) is estimated using monthly data and two lags of R_t . R_t is the annualized monthly average of the federal funds rate and $\pi_{t+h_\pi|t}$ is the Greenbook forecast of the annualized percentage change in the GNP or GDP deflator between t and $t+h_\pi$.¹⁹ Because of repeated changes in the base year, no consistent time series of forecasted real GNP or GDP can be constructed from the Greenbook, over the sample considered. As a result, in the benchmark specification, $x_{t+h_x|t}$ is proxied using the difference between the natural rate of unemployment and the forecasted unemployment rate. Defining the unemployment gap in this way makes its sign consistent with that of the output gap, as conventionally defined.

There is considerable uncertainty, not only about the appropriate measure but also, the relevant concept of real activity.²⁰ To maintain the real-time nature of the estimation, in the baseline specification, I use, as a measure of the natural rate of unemployment at time t , an historical average of the unemployment rate up to that point. This is the same measure that Orphanides (2002) presents and appears broadly consistent with narrative accounts of this period.²¹ I investigate below the robustness of the baseline results to alternative real-time measures of real-activity, including Orphanides' (2003) real-time output gap measure.

The earliest Greenbook forecasts available are for the October 1965 FOMC meeting and are made publicly available with a five year lag.²² The forecasts have been produced for various

¹⁹In October 1991, the Greenbook switched from forecasting the GNP deflator to the GDP deflator. In terms the inflation measure, this change in definition has essentially no effect on our results.

²⁰Note that the time-varying constant accommodates, to some extent, mis-specification in the real-activity measure.

²¹See for instance Orphanides (2003).

²²Note that the current values of inflation and real activity are not known within the period. Besides, the Greenbook forecasts are not available at a standard frequency: usually eight data points a year are available, and they are not

horizons, in some instances ranging up to eight quarters ahead. But only the shorter horizons are consistently available throughout the period. As a result, the forecasting horizons considered in this paper vary from the current quarter up to three quarters ahead.²³ The baseline specification assumes a three quarter (3-Q) ahead forecasting horizon on both inflation and the unemployment gap, for which the available sample is 1969:5 to 1998:6. While this is broadly consistent with the horizons considered in the literature, I also document below the robustness of the results to alternative choices.

Figure 1 displays the time series of the various measures used in the baseline specification.

5.2 Baseline results

As discussed in the Introduction, the existing literature suggests that the conclusion on whether or not systematic monetary policy changed depends on various aspects of the empirical framework, including the type of time variation in the parameters, the treatment of heteroskedasticity and the real time nature of the data. To sort out the influence of each of these factors, in this subsection I first test the stability and perform a split sample estimation of (1). Since a central theme of this paper is that a proper model of the time variation is crucial to the conclusion, I then consider a baseline TVP model that is otherwise identical the split-sample model. While this baseline specification allows to identify the effect of the TVP assumption, it is not necessarily meant evenly spaced. In this paper I assume that in a given month, the Fed only observes the Greenbook forecasts for the FOMC meeting that occurred before the beginning of that month. This presumably assumes slightly less information than the Fed actually had. Sensible perturbations around this assumption, including linear interpolation between the meetings, did not affect the results.

²³For the different horizons, a consistent time series, with at least two observations per quarter, starts as follows: Current quarter (C-Q): 1967:11. One quarter ahead (1-Q): 1968:7. Two quarters ahead (2-Q): 1968:9. Three quarters ahead (3-Q): 1969:5.

to represent the most satisfactory specification. This is why in Section 5.3 I check the robustness of the conclusions to alternative, perhaps more realistic, specifications.

The first step is to provide statistical evidence that (1) indeed displays instability. Table 1 reports the p -values of the heteroskedasticity-robust version of the QLR_T test, applied to different combinations of the parameters characterizing the long run policy responses.²⁴ There is evidence of instability for all these combinations. The test, applied jointly to all the coefficients, suggest at least two breaks: one in February 1974 and the other in May 1979. These results already suggest that the time-variation might be more complex than the one-time discrete change that is often assumed. The TVP estimates I report below will shed further light on the nature of these changes.

Table 1: Hetero-robust QLR_T test

Parameters Tested	p -values
$\{\alpha, \beta, \gamma, \rho_1, \rho_2\}$	0.000
$\{\beta, \gamma, \rho_1, \rho_2\}$	0.001
$\{\beta, \rho_1, \rho_2\}$	0.015
$\{\gamma, \rho_1, \rho_2\}$	0.002

The joint stability test on all coefficients result implies a median unbiased estimate of λ equal to 7.63, which suggests a relatively small period to period variation in the parameters. As I alluded to in the description of the estimation approach, this estimate is within the range of values of λ for which the MLE tends to run into problems.²⁵

²⁴The p -values are obtained using Hansen (1997)'s simulation procedure.

²⁵Stock and Watson (1998) show that for the local level model and for $\lambda = 8$, the probability of getting $\hat{\lambda} = 0$ with MLE is 42%. Comparatively, for the median-unbiased estimation approach, it is only 13% (See Table 1 of their paper).

Given the existing empirical results, it is useful to consider the estimates of the coefficients, assuming a one-time structural break. A natural break date, within the confidence intervals of the break dates just estimated, is October 1979. Romer and Romer (1989) argue that it corresponds to a significant shift in policy, soon after the start of Volcker’s tenure as chairman of the Fed. The results from the policy rule estimated on each side of this date are reported in Table 2. The estimates of β and ρ_i imply a long-run response to inflation, $\frac{\beta}{1-\rho(1)}$, equal to 1.37 and 1.57 in the first and second samples respectively. In terms of the response to inflation, these estimates would suggest that monetary policy in the 1970’s was not too different than under Volcker or Greenspan. This is in line with the point made by Orphanides (2002, 2003).

Table 2: Split sample policy rule estimates with break in October 1979

	α	β	γ	ρ_1	ρ_2
Pre-1979:9	0.159 (0.133)	0.109 (0.032)	0.178 (0.047)	1.335 (0.086)	-0.415 (0.082)
Post-1979:10	0.203 (0.124)	0.195 (0.042)	0.020 (0.041)	1.253 (0.062)	-0.377 (0.060)

(Standard errors in parenthesis)

Turning to the TVP results, the estimation is implemented as described in the previous section, with the variance of the policy shock assumed to experience a structural break at the same date as the parameters in the split sample estimation, i.e. October 1979. This is the same account of heteroskedasticity implicit in the split-sample estimation. Figure 2 reports in the top panels the estimates for the long-run response to inflation and unemployment, together with one standard error confidence bands,²⁶ as well as the split-sample estimates reported in the previous table. The

²⁶These are obtained from the Kalman filter, conditional on the median-unbiased estimate of the variance of the

bottom panels report the estimates of the federal funds rate persistence, $\rho_t(1)$, and the intercept, α_t , together with one standard error confidence bands.

Contrary to the split-sample estimates, the TVP results suggest that the conduct of monetary policy greatly evolved since the early 1970's. The response to inflation varies roughly between 0.7 and 2.4, that to unemployment between 0.4 and 2.6, and the persistence of the federal funds rate between 0.85 and 0.93.²⁷

The point estimates of the policy responses to inflation and unemployment reveal interesting patterns. In the early 1970's, the response to inflation was well above one. It fell substantially in 1973 – 1974, and by 1975 the response was below one, not satisfying the Taylor principle. The response remained weak until the late 1970's and then started to gradually increase until the mid-1980's. It eventually stabilized at a level above 2, with perhaps a mild decrease in the mid-1990's. Interestingly, the response to unemployment displays a quite different pattern. It experienced a sustained decrease throughout the 1970's, evolving from a level above 2 to a level below 1 from 1978 onward. Concerning the persistence of the federal funds, two main observations stand out. First, it was stable throughout the 1980's and 1990's at a level higher than in the 1970's. Second, there was a sharp reduction in the persistence of the federal funds rate in the early 1980's; this is in fact to be expected since this period was characterized by an atypical volatility of the federal funds rate.

Overall, the characterization of monetary policy that emerges appears quite reasonable. In fact, many aspects, including the timing of the changes, are consistent with existing accounts.²⁸

parameters.

²⁷An Okun's law coefficient, of 2.5 say, can be used to translate γ_t in terms of a response to the output gap. The resulting coefficient then varies roughly between 0.1 and 1.

²⁸Concerning the timing of these changes, it is important to note that given that the estimated time series of the parameters are based on a full-sample, two-sided filter, there is no reason to expect the TVP to pick up a change

The behavior of policy during the early 1970's matches Orphanides' (2003) conjecture that the Fed initially intended to respond aggressively to inflation, but failed to detect an adverse shift in trend productivity. The weak response to inflation between 1975 and 1978 is consistent with the account of Romer and Romer (1989), who argued, based on their readings of the FOMC minutes, that "the actual commitment to combat inflation was weak" during that period. Like Clarida, Galì and Gertler (2000) and Cogley and Sargent (2001, 2005), the actual size of the response to inflation during this period is estimated below one. In fact, the whole time series of the response to inflation is surprisingly similar to the one estimated by Cogley and Sargent (2005). The fact that such similar conclusions are obtained from a markedly different empirical framework and different treatment of real time issue is reassuring. Further, the inflation response reached a temporary peak precisely in October 1979, a date singled out by Romer and Romer (1989) as corresponding to an unusual tightening of monetary policy. Finally, the sustained decline in the response to real activity corroborates the recent results of Orphanides (2003).

But the TVP estimates uncovers important features which have not been detected in previous empirical analysis. The changes in both the response to inflation and real activity are sizeable and important to characterize the evolution of monetary policy. Moreover, the uncovered behavior of the parameters is more complex than any of the existing split-sample estimates suggest. First, the changes in the response to inflation and unemployment do not appear to have been synchronized. Second, as Romer and Romer (2002) argue, this evolution does not appear unidirectional and seems, at least in part, gradual. We see, for instance, that while some important and sizeable changes are apparent around the appointment of Volcker, the main increase in the response to inflation occurred during the 1980 – 1982 period. Moreover the uncovered changes are not limited to that

with some delay.

episode. This is consistent with evolving views and their gradual implementation within the Fed.²⁹

Of course, precise statements about the size of the policy response need to be qualified by the statistical uncertainty surrounding these estimates. For instance, we cannot conclude at conventional statistical level that the response to inflation is really below one between 1975 and 1978.³⁰ It is important to stress, however, that the *changes* in the coefficients over time are statistically significant, as implied by the stability test reported above.³¹ That is, the response to inflation is significantly smaller between 1975-1978 than anywhere else in the sample.

In light of the TVP results, it is also instructive to look more closely at the robustness of the split-sample results. At first sight, the split-sample and TVP estimates appear to yield dramatically different conclusions concerning the evolution of monetary policy, particularly about the strength of the response to inflation. However, given the nature of the changes uncovered by the TVP estimation, this is not surprising. The long-run response to inflation varies considerably throughout the 1970's, being high initially and low in the mid- to late 1970's. Other things equal, the constant parameter OLS estimate for the pre-1979:9 sample should thus tend to be downward biased in the earlier part of that period and upward biased in the later part. The results reported above are certainly consistent with this explanation: the OLS estimates is 1.37, which corresponds to the mid-range of the TVP estimates obtained for that period. To investigate this further, I re-estimate

²⁹It is true, as it should be clear from the description provided in the previous Section, that the TVP estimation should tend to make discrete changes look more gradual than they really are. But that is not sufficient to explain the gradual nature of the uncovered changes. In fact, the TVP estimates has no problem picking up a rapid increase in the response to inflation in late 1979 and the changes are not limited to that period, as the behavior of the response to unemployment makes clear. I investigate this further below by considering a specification that nests a TVP with a discrete change in 1979:10.

³⁰The marginal probability that the response to inflation is below one at the trough is about 80%.

³¹Note that the significance of the changes in the coefficients cannot be formally inferred from the confidence intervals.

the constant parameter model over the 1974:2 – 1979:9 period, where the starting date corresponds to the first break date estimate reported above. The results are reported in Table 3. Note that if the true nature of the change in policy was a discrete jump in 1979:10, starting the first sample later or ending it earlier, should not dramatically affect the results. Yet, they appear extremely fragile: the long-run response to inflation in this case drops from 1.37 to 0.83, consistent with the TVP estimates results. Moreover, ending this sample just one month earlier, in 1979:8, makes the estimate drop further to 0.43. This evidence suggests that the way instability is modeled is crucial to uncover properly the nature of the changes in monetary policy during the 1970's. This provides further evidence of the gradual nature of the actual changes and also of the fragility of the split-sample estimates in the present context.

Table 3: Policy rule estimates for alternative sub-samples

	α	β	γ	ρ_1	ρ_2
1974:2 – 1979:9	0.276 (0.354)	0.040 (0.080)	0.107 (0.055)	1.485 (0.115)	-0.533 (0.114)
1974:2 – 1979:8	0.400 (0.375)	0.018 (0.083)	0.118 (0.056)	1.464 (0.117)	-0.507 (0.117)

(Standard errors in parenthesis)

Given the behavior of the TVP estimates around 1979 and the widely held belief of a drastic change in policy soon after Volcker's appointment, it might still be argued that a discrete shift would be more appropriate to describe this period. One way to check more directly if this potential mis-specification of the TVP model is important, is to nest together the TVP and discrete change model. In order to maintain a parsimonious specification, β_t and γ_t are assumed to follow driftless

random walks, with a potential discrete shift in 1979:10.³² Letting D_t be a dummy equal to one for t greater than 1979:10, the estimated policy reaction function is:

$$R_t = \alpha_t + \beta_t \pi_{t+h_\pi|t} + \tilde{\beta} D_t \pi_{t+h_\pi|t} + \gamma_t x_{t+h_x|t} + \tilde{\gamma} D_t x_{t+h_x|t} + \rho(L)R_{t-1} + \varepsilon_t.$$

This can be estimated as before, accounting for the presence of the time invariant parameters, $\tilde{\beta}$ and $\tilde{\gamma}$. The time series of the response to inflation is given by $\beta_t + \tilde{\beta} D_t$, and to the unemployment gap by $\gamma_t + \tilde{\gamma} D_t$.

Figure 3 compares the resulting estimates together with the previous TVP results. There are a few differences in the results when a discrete shift is allowed. First, the estimated persistence of the federal funds rate is everywhere lower. Second the time variation in the response to unemployment might be partly discrete in nature and is estimated to be lower post-1980.³³ Third, the response to inflation is slightly higher in the mid-1970's, implying a lower probability that the Taylor principle has been violated. Importantly, however, the conclusions about the general evolution of the responses to inflation and unemployment remain intact: the changes are important, largely gradual and the timing is essentially the same. Perhaps surprisingly, the response to inflation does not show any important jump in 1979, even though such a discrete change is explicitly allowed. The conclusions drawn on the basis of the TVP estimates thus appear robust to the presence of a discrete change in 1979, if such a change indeed occurred.

³²Note that I am not arguing that the nested model is a more appropriate one. In fact, a key advantage of the TVP estimation is that it does not rely on a choice of a simultaneous break date for all parameters. This is not the case for the nested model.

³³Note, however, that there is no guarantee that 1979:10 is the true date for the response to unemployment.

5.3 Robustness

I now investigate the robustness of the previous results along three dimensions: the date at which the variance of ε_t changed, the choice of the forecasting horizons and alternative measures of real activity.

5.3.1 Heteroskedasticity

So far, the TVP estimation assumed a single change in the variance of the policy shock in 1979:10. This has the virtue of making the TVP results directly comparable with the split-sample results, which implicitly rely on the same assumption. It is not clear, however, that a single break in the variance fully captures the heteroskedasticity of the policy shock. One way to gain further insights is to look at the time series behavior of the full sample OLS residuals of (1). These are displayed in Figure 4. Clearly, the early 1980's period stands out. The volatility of the policy shock was a lot larger, which might be due in part to the Fed following a different operating procedure during that period.³⁴ This could have been wrongly attributed to time variation in the parameters. Outside this period, however, there is no clear pattern.

One way to capture this behavior of the policy shock, and non-parametrically account for the potential change in operating procedure, is to allow for three regimes in the variance of the policy shock, with the middle one corresponding to the early 1980's.³⁵ Based on Figure 4, this should be enough to capture the extent of heteroskedasticity present in these shocks. Figure 5 shows the resulting TVP estimates, defining the 1979:10 to 1982:10 as the second of the three regimes. The

³⁴The federal funds rate provides probably a less adequate measure of monetary policy stance between 1979:10 and 1982:10, as non-borrowed reserves were set to achieve a level of interest rates consistent with money growth target. But Cook (1989) argues that the fed funds rate may still provide a satisfactory indicator during this period.

³⁵Boivin (1999) allows explicitly for changes in the operating procedure using the Bernanke and Mihov (1998) model. This requires a multi-variate framework which is beyond the scope of this paper.

second regimes corresponds to the non-borrowed reserves targeting episode. The key difference with the previous TVP estimates is that the variation in the parameters post-1982 is much more pronounced. This is not surprising, however: since the variance of the policy shock post-1984 is about ten times smaller than previously estimated,³⁶ more of the fluctuation in the federal funds rate is attributed to variations in the parameters. In fact, much of this variation seems to be due to the federal funds rate persistence, which induces similar movements in the long-run response to inflation and unemployment gap. Overall, however, the conclusions about the nature of the changes in monetary policy are unaffected.

5.3.2 Forecasting horizons

There is also some uncertainty about the forecasting horizons that properly characterize the Fed's behavior. Figure 6 reports the TVP estimates, assuming three regimes in the variance of the policy shock, for various horizon combinations: $h_\pi = h_x = 0$, $h_\pi = h_x = 1$, $h_\pi = h_x = 2$, $h_\pi = h_x = 3$ and $h_\pi = 3$, $h_x = 0$.³⁷ In all cases, there is important time variation in the parameters. The pattern of responses to real activity are similar over the whole period and the to inflation in the 1970's are similar. For short horizons, $h_\pi = h_x = 0$ and $h_\pi = h_x = 1$, the response to inflation is smaller throughout the sample, but especially in the post-1980 period. This smaller response to current inflation is consistent with the Fed being forward-looking: loosely speaking, the estimated coefficient on current inflation should be equal to the actual response to inflation at horizon h_π , multiplied by a coefficient, less than one, related to the degree of persistence in inflation between t and $t + h_\pi$. But, apart perhaps for the very short horizons, the conclusions about the importance

³⁶In the baseline estimation, the post-1979:10 variance of the policy shock was 0.502, whereas now it is 2.271 for the 1979:10-1982:10 period and 0.060 for the post-1982:10 period. The variance of the shock in the pre-1979:10 period is 0.154.

³⁷Clarida, Gali and Gertler use $h_\pi = h_y = 1$ as their baseline specification.

and the pattern of the changes remain essentially the same.

5.3.3 Real activity measures

Finally, perhaps the most important source of uncertainty about this specification concerns the appropriate real-time measure of real activity. In fact, one important conclusion from Orphanides (2003) is that to properly characterize the historical conduct of monetary policy, a proper account of the historically perceived real activity level is essential. While the recursive average of the unemployment rate considered so far is arguably a reasonable measure, consistent with existing empirical applications, there is no guarantee that this, or any other measure, corresponds to the natural rate historically perceived by policymakers. It is thus important to make sure that the results are at least robust to other sensible alternatives.

I first consider alternative measures of the natural rate of unemployment displaying varying degree of smoothness. Figure 7 displays the baseline measure together with three alternatives: 1) a five year moving average as in Bernanke and Boivin (2003), 2) a three year moving average and 4) an exponential smoother obtained as:

$$u_t^* = u_{t-1}^* + 0.075 (u_t - u_{t-1}^*),$$

where u_t^* denote the natural rate of unemployment. This last measure, which can be interpreted as a constant gain update of the mean of unemployment, might be more appropriate than the moving average measure to track the natural rate of unemployment if it is itself moving over time.³⁸ As is apparent from the Figure, they do considerably differ and thus constitute a useful robustness check. Figure 8 compares the resulting TVP estimates together with the baseline results. Even though the underlying measures of the natural rate differ substantially, the implied time variation is overall

³⁸Thanks to an anonymous referee for suggesting this measure.

very similar. In particular, the time path of the long-run response to inflation is essentially the same, no matter the measure used. There might be somewhat more uncertainty about the exact nature of the change in the long-run response to unemployment: with the exponential smoother, it increases more rapidly and to a higher level from the mid-1980's onward. But for all measures, the results suggest that the changes in the conduct of monetary policy were sizeable.

All the real-activity measures so far considered are based on the unemployment rate. As mentioned above, an advantage is that they are constructed from a time series that essentially has a consistent definition throughout the period. This is not the case for measures based on real GNP or GDP. As a robustness check, however, it is interesting to see if a real-time measure based on the output gap yields similar results. I thus re-estimate the version of (1) considered by Orphanides (2003). It is based on quarterly data, the three-quarter ahead Greenbook inflation forecast, one lag of the federal funds rate and the real-time output gap measure he constructed.³⁹ In a given quarter, this alternative output gap measure is based on the headline concept of real output that was in use during that quarter. Orphanides (2003) argues that this measure is as close as is possible to the policymakers' perceptions reported in the available historical record. The results from this specification are displayed in Figure 9. Using quarterly data and this output gap measure yields surprisingly similar results. In fact, applying a simple Okun law calculation to the estimated response to the output gap yields estimates of the same magnitude as those previously obtained for the unemployment gap.

The characterization of the evolution of monetary policy that stems from the benchmark specification is thus robust to alternative measures of the natural rate of unemployment and Orphanides' (2003) account of policymakers' perceptions.

³⁹In this case the sample ends in 1995:4.

6 Conclusion

Determining whether or not the conduct of monetary policy changed is central to our understanding of the effect of monetary policy and of the root causes of the improved U.S. economic performance of the last twenty years. Although there has been a considerable amount of research on the topic, there are still important disagreements, even about the mere existence of these changes.

Accounting for the various issues at the center of this disagreement, the empirical investigation in this paper led to the following conclusions:

- The conduct of monetary policy changed greatly over the last 30 years. The response to inflation was very weak in the second half of the 1970's, but strong elsewhere in the sample. The response to real activity fell substantially and lastingly throughout the 1970's. From the mid-1980's onward, the conduct of monetary policy appears to have been stable.
- The rich pattern of changes in the response to inflation and unemployment are not well captured by a discrete shift, occurring in 1979. The changes are not unidirectional nor synchronized across parameters. Moreover the transition under Volcker appears to have taken place over a few years, with the most important changes occurring between 1980 and 1982. More than real-time data issues, the complex nature of the changes might explain why conflicting results have been reported in the literature.

The timing of these changes is certainly consistent with monetary factors playing a role in the improved economic outcomes of the last two decades. This confirms the view stemming from the narrative evidence cited in the Introduction. Of course, the next step, which requires a fully specified model of the economy, is to measure this evidence against other non-monetary sources of improved economic performance, such as the private sector behavior or simply luck. But, beyond the timing,

the magnitude of the changes identified in this paper is substantial, suggesting a non-trivial role of monetary policy.

References

- [1] Abrams, Richard K., Richard Froyen and Roger N. Waud (1980), "Monetary Policy Reaction Functions, Consistent Expectations, and The Burns Era," *Journal of Money, Credit and Banking*, 12, 30-42.
- [2] Andrews, Donald W.K. (1993), "Tests of Parameter Instability and Structural Change with Unknown Change Point," *Econometrica*, 61, 821-856.
- [3] Bai, Jushan and Pierre Perron (1998), "Estimating and Testing Linear Models with Multiple Structural Changes", *Econometrica*, 66, 47-78.
- [4] Barth, James, Robin Sickles, and Philip Wiest (1982), "Assessing the Impact of Varying Economic Conditions on the Federal Reserve Behavior," *Journal of Macroeconomics*, 4, 47-70.
- [5] Bernanke, Ben and Alan Blinder (1992), "The Federal Funds Rate and the Channels of Monetary Transmission," *American Economic Review*, LXXXII (1992), 901-21.
- [6] Bernanke, Ben S. and Jean Boivin (2003), "Monetary Policy in a Data-Rich Environment," *Journal of Monetary Economics* 50:3, 525-546.
- [7] Bernanke, Ben and Ilian Mihov (1998), "Measuring Monetary Policy," *Quarterly Journal of Economics*, 869-902.
- [8] Blinder, Alan (1985), "Comment on Federal Deficits, Interest Rates, and Monetary Policy," *Journal of Money, Credit and Banking*, 17, 685-689.

- [9] Boivin, Jean (1999), “The Fed’s Conduct of Monetary Policy: Has it Changed and Does it Matter?” Ph.D. Thesis, Princeton University.
- [10] Boivin, Jean and Marc Giannoni (2003), “Has Monetary Policy Become More Effective?,” NBER working paper #9459.
- [11] Christiano, Lawrence, Martin Eichenbaum and Charles Evans (2000), “Monetary Policy Shocks: What Have We Learned and to What End?,” in J. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, (Amsterdam: North-Holland).
- [12] Cogley, Timothy and Thomas J. Sargent (2001), “Evolving Post World War II U.S. Inflation Dynamics,” *NBER Macroeconomics Annual* 16, pp. 331-373.
- [13] Cogley, Timothy and Thomas J. Sargent (2005), “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.,” *Review of Economic Dynamics* 8.
- [14] Cook, Timothy (1989), “Determinants of the Federal-Funds Rate,” *FRB Richmond Economic Review* 75(1), Jan-Feb, 3-19.
- [15] Cooley, Thomas and E. Prescott (1976), “Estimation in the Presence of Stochastic Parameter Variation,” *Econometrica*, 44:1, 167-184.
- [16] Clarida, Richard, Jordi Galí and Mark Gertler (1998), “Monetary Policy Rules in Practice: Some International Evidence,” *European Economics Review*, 1033-1068.
- [17] Clarida, Richard, Jordi Galí and Mark Gertler (2000), “Monetary Policy Rule and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*.
- [18] DeLong, J. Bradford (1997), “America’s Only Peacetime Inflation: the 1970’s,” *Reducing Inflation*, NBER Studies in Business Cycles, Vol. 30.

- [19] Froyen, Richard T. (1974), "A Test of Endogeneity of Monetary Policy," *Journal of Econometrics*, 2, 175-188.
- [20] Hansen, Bruce E. (1997), "Approximate Asymptotic P Values for Structural Change Tests," *Journal of Business Economic and Statistics* 15(1): 60-67.
- [21] Li, Hong (2004), "The Econometrics of the Lucas Critique: Estimation and Testing of Euler Equation Models with Time-varying Reduced-form Coefficients," unpublished manuscript, Princeton University.
- [22] Lubik, Thomas and Frank Schorfheide (2004), "Testing for Indeterminacy: An Application to U.S. Monetary Policy," *American Economic Review* 94, 190-217.
- [23] Lucas, Robert E. Jr. (1976), "Econometric Policy Evaluation: A critique." *Carnegie Rochester Conference Series on Public Policy*, 1, 19-46.
- [24] McConnell, Margaret and Gabriel Perez-Quiros (2000), "Output Fluctuations in the United-States: What Has Changed Since the Early 1980's," *American Economic Review*, 90:5, 1464-1476.
- [25] McNees, Stephen K. (1986), "Modeling the Fed: A Forward-Looking Monetary Policy Reaction Function." *New England Economic Review*, 3-8.
- [26] McNees, Stephen K. (1992), "A Forward-Looking Monetary Policy Reaction Function: Continuity and Change," *New England Economic Review*, 3-13.
- [27] Orphanides, Athanasios (2001), "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, 91:4, 964-985.

- [28] Orphanides, Athanasios (2002), "Monetary Policy Rules and the Great Inflation," *American Economic Review*, 92:2; 115-121.
- [29] Orphanides, Athanasios (2003), "Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches" *Journal of Money, Credit and Banking*, forthcoming.
- [30] Potts, Glenn T. and Dudley G. Lockett (1978), "Policy Objectives of the Federal Reserve System," *Quarterly Journal of Economics*, 17, 525-534.
- [31] Quandt, Richard E. (1960), "Tests of the Hypothesis that a Linear Regression System Obeys Two Separate Regimes," *Journal of the American Statistical Association*, 55(290), 324-330.
- [32] Reifschneider, David L., David J. Stockton and David Wilcox (1997), "Econometric models and the monetary policy process," *Carnegie Rochester Conference Series on Public Policy*, 47, 1-37.
- [33] Romer, Christina and David Romer (1989), "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," in Olivier Blanchard and Stanley Fisher, eds, *NBER Macroeconomics Annual* (MIT press).
- [34] Romer, Christina and David Romer (2000), "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review*, 90, 429-457.
- [35] Romer, Christina and David Romer (2002), "The Evolution of Economic Understanding and Postwar Stabilization Policy," in *Rethinking Stabilization Policy*, Federal Reserve Bank of Kansas City.
- [36] Romer, Christina and David Romer (2003), "A New Measure of Monetary Shocks: Derivation and Implications," NBER working paper #9866.

- [37] Sargent, Thomas (1999), *“The Conquest of American Inflation”*, Princeton University Press, Princeton, NJ.
- [38] Sims, Christopher (1999), “Drifts and Breaks in Monetary Policy,” Manuscript, Princeton University.
- [39] Sims, Christopher (2001), “Discussion of Cogley and Sargent ‘Evolving Post World War II U.S. Inflation Dynamics,’ *NBER Macroeconomics Annual* 16, pp. 373-379.
- [40] Sims, Christopher and Tao Zha (2004), “Where there Regime Switches in US Monetary Policy?” Manuscript, Princeton University.
- [41] Stock, James H. (2001), “Discussion of Cogley and Sargent ‘Evolving Post World War II U.S. Inflation Dynamics,’ *NBER Macroeconomics Annual* 16, pp. 379-387.
- [42] Stock, James H. and Mark W. Watson (1996), “Evidence on Structural Instability in Macroeconomic Time Series Relations,” *Journal of Business and Economic Statistics*, 14, 11-30.
- [43] Stock, James H. and Mark W. Watson (2002), “Has the Business Cycle Changed and Why?” *NBER Macroeconomics Annual* 17, pp. 159-218.
- [44] Taylor, John B. (1993), “Discretion versus Policy Rules in Practice,” *Carnegie Rochester Conference Series on Public Policy*, 39, 195-214.
- [45] Taylor, John B. (1998), “Monetary policy guidelines for unemployment and inflation stability,” *Inflation, Unemployment, and Monetary Policy*, R.M. Solow and J.B. Taylor (eds.). Cambridge, MIT Press.


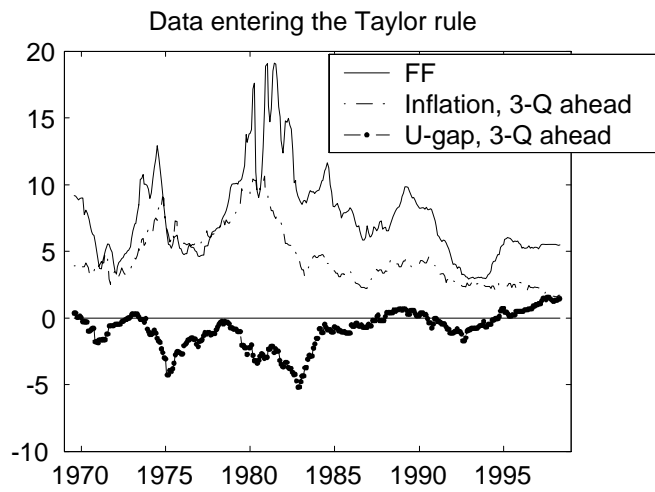


Figure 1: Real-time data used in the estimation

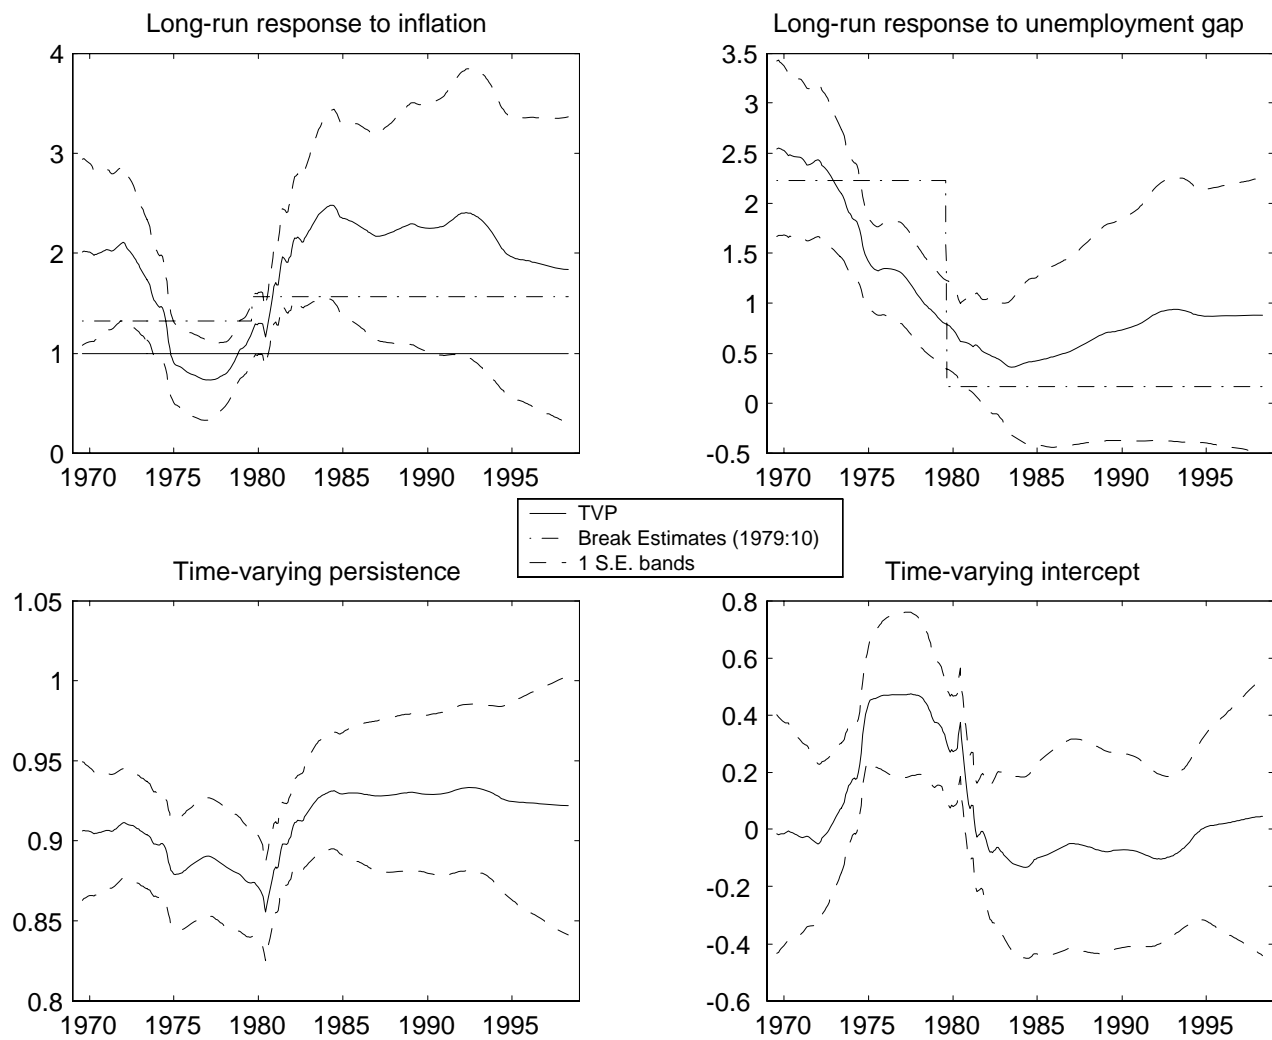


Figure 2: TVP estimates: Benchmark specification

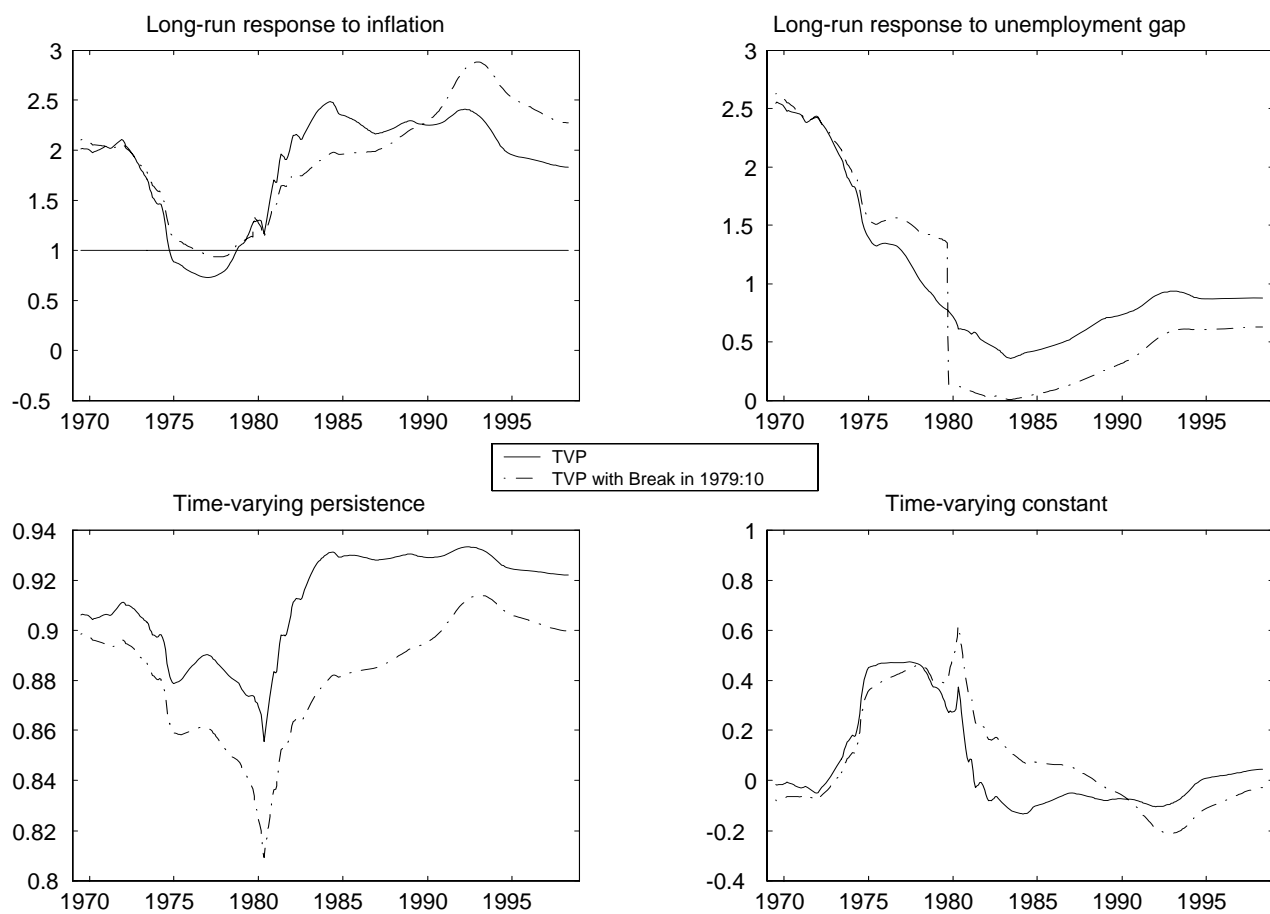


Figure 3: TVP estimates with discrete shifts in β and γ

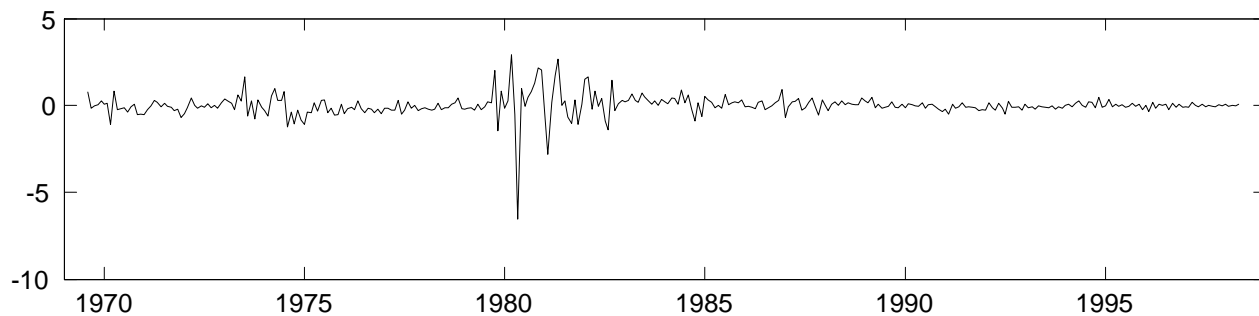


Figure 4: Estimated policy shocks: Full sample, constant parameters

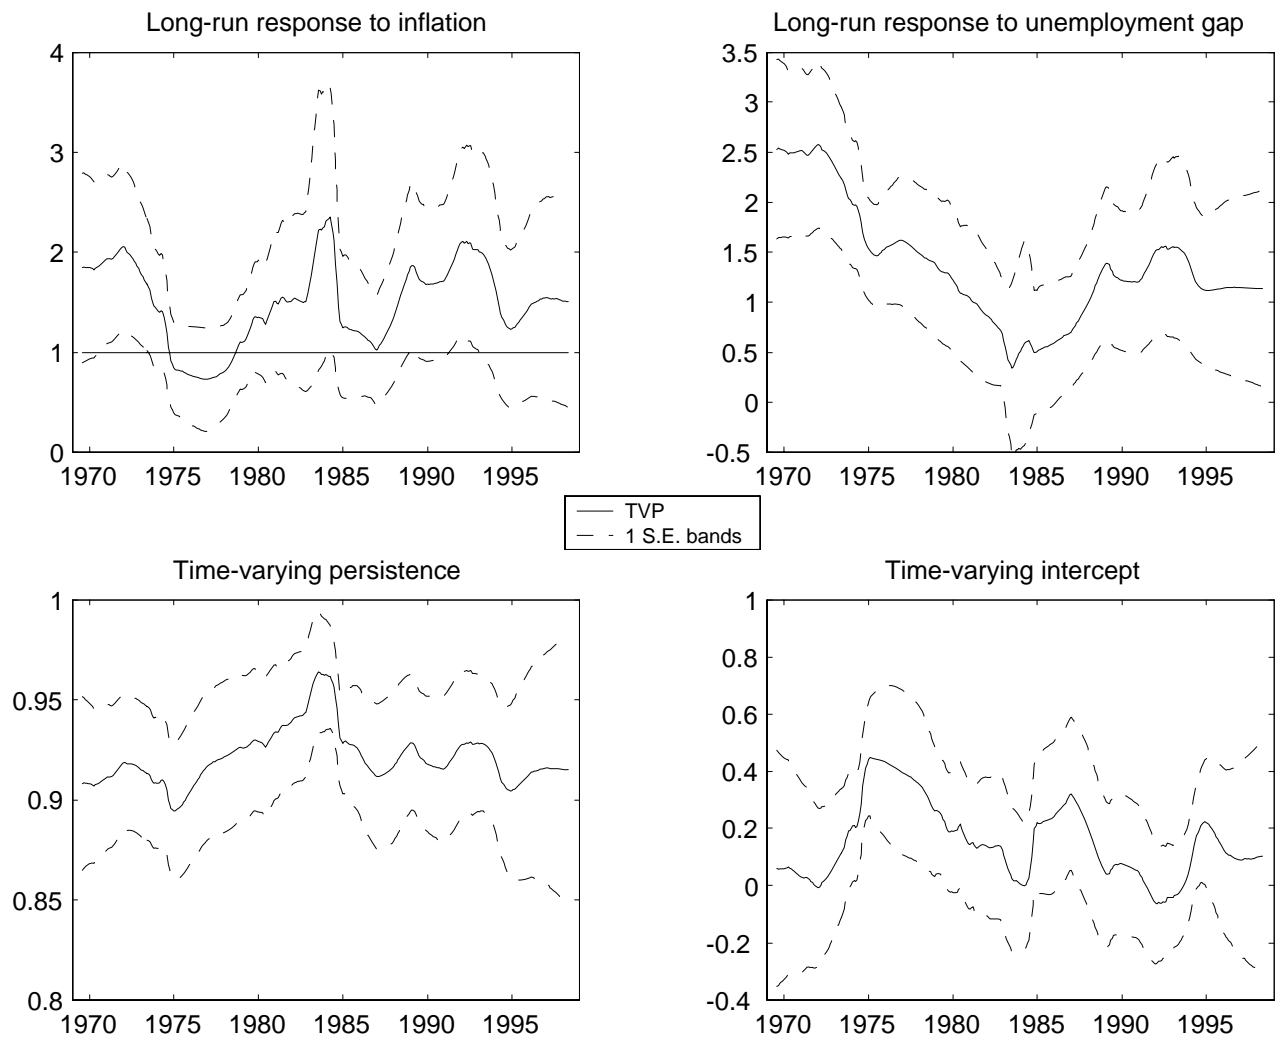


Figure 5: TVP estimates: Alternative treatment of heteroskedasticity (Breaks in the variance in 1979:10 and 1982:10)

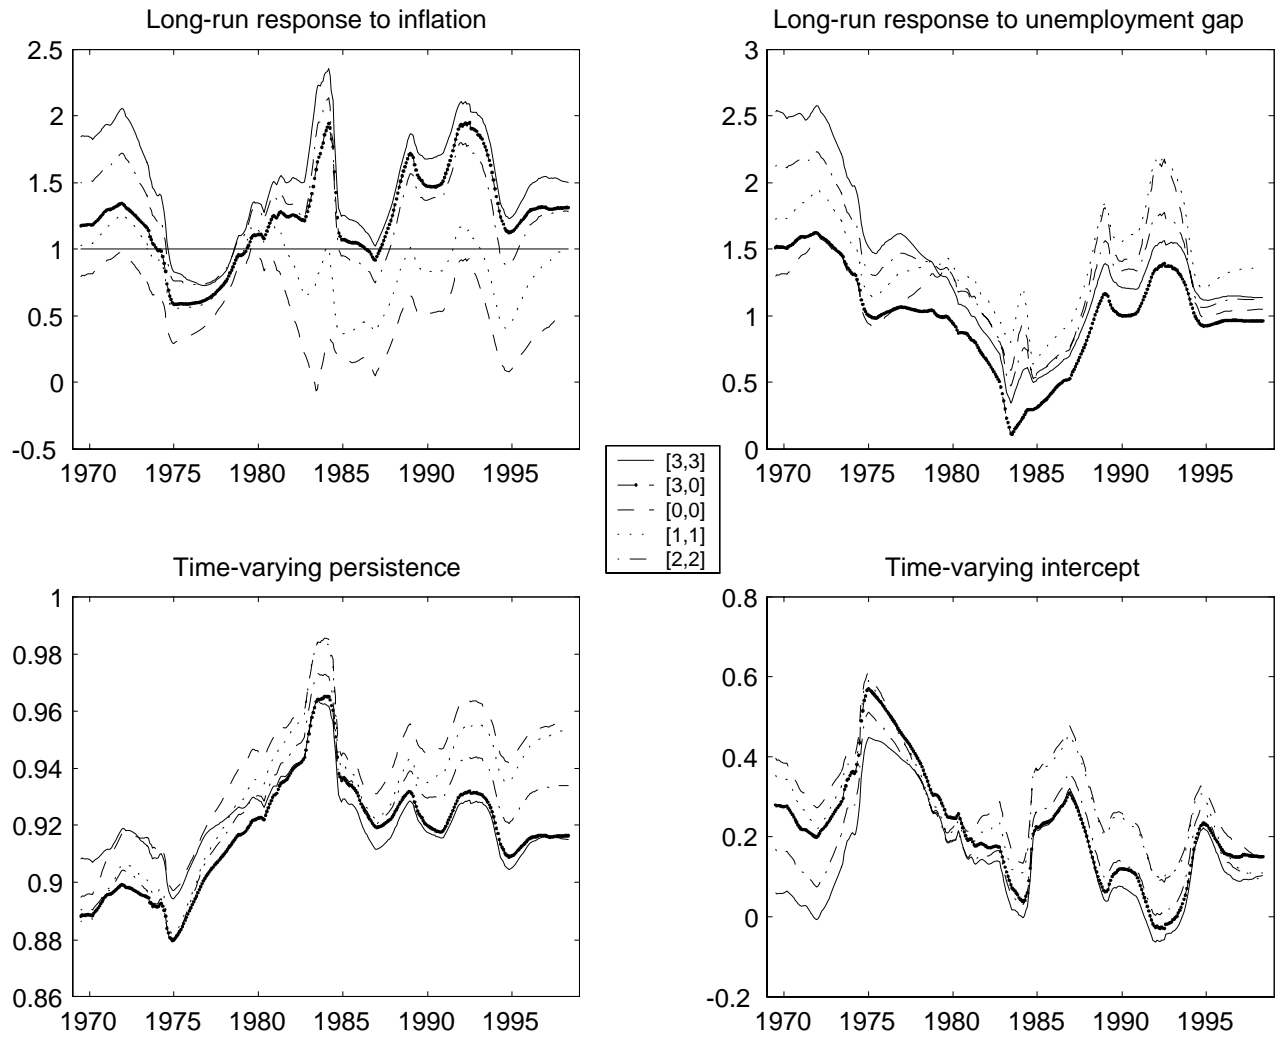


Figure 6: TVP estimates: Alternative forecasting horizons [inflation, gap]

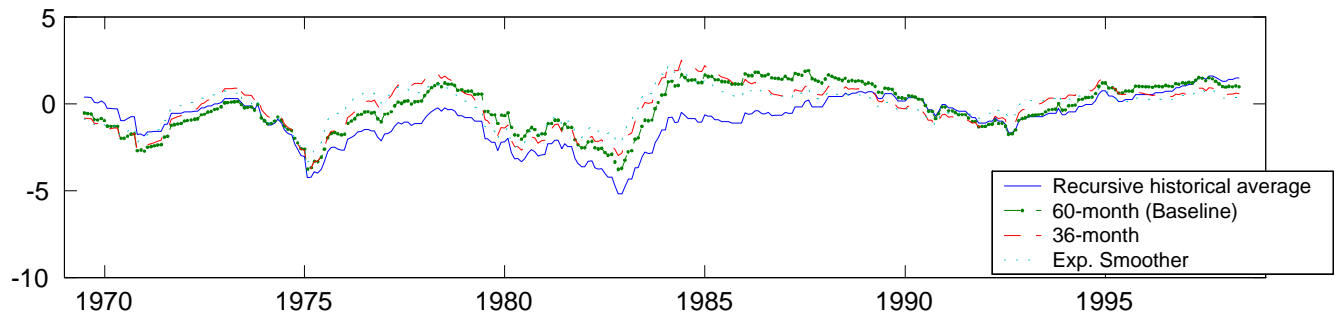


Figure 7: Alternative moving-average measures of the natural rate of unemployment

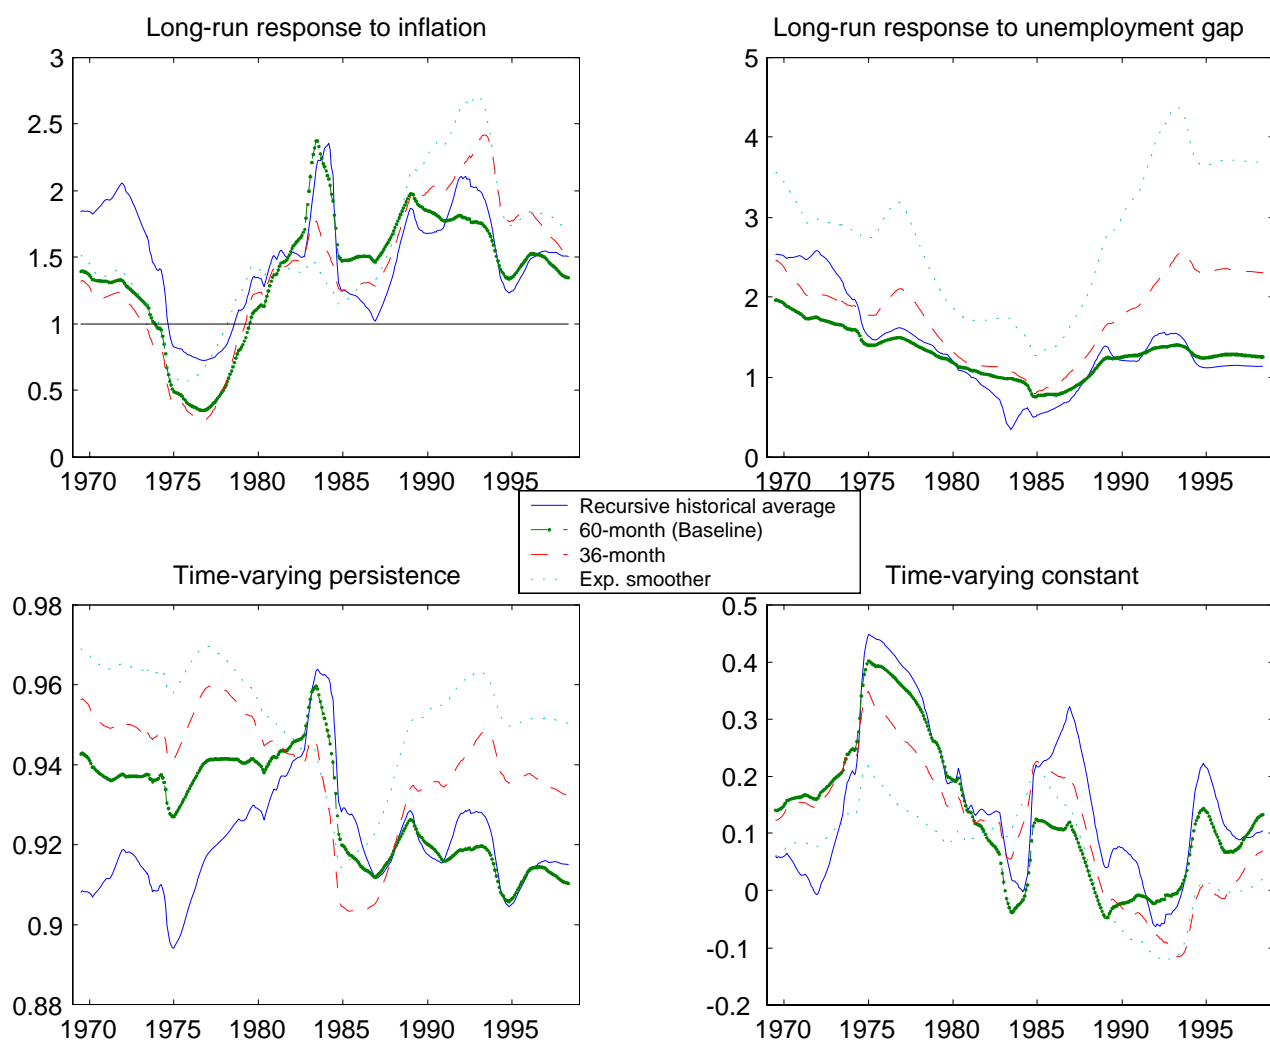


Figure 8: TVP estimates: Alternative measures of the natural rate of unemployment

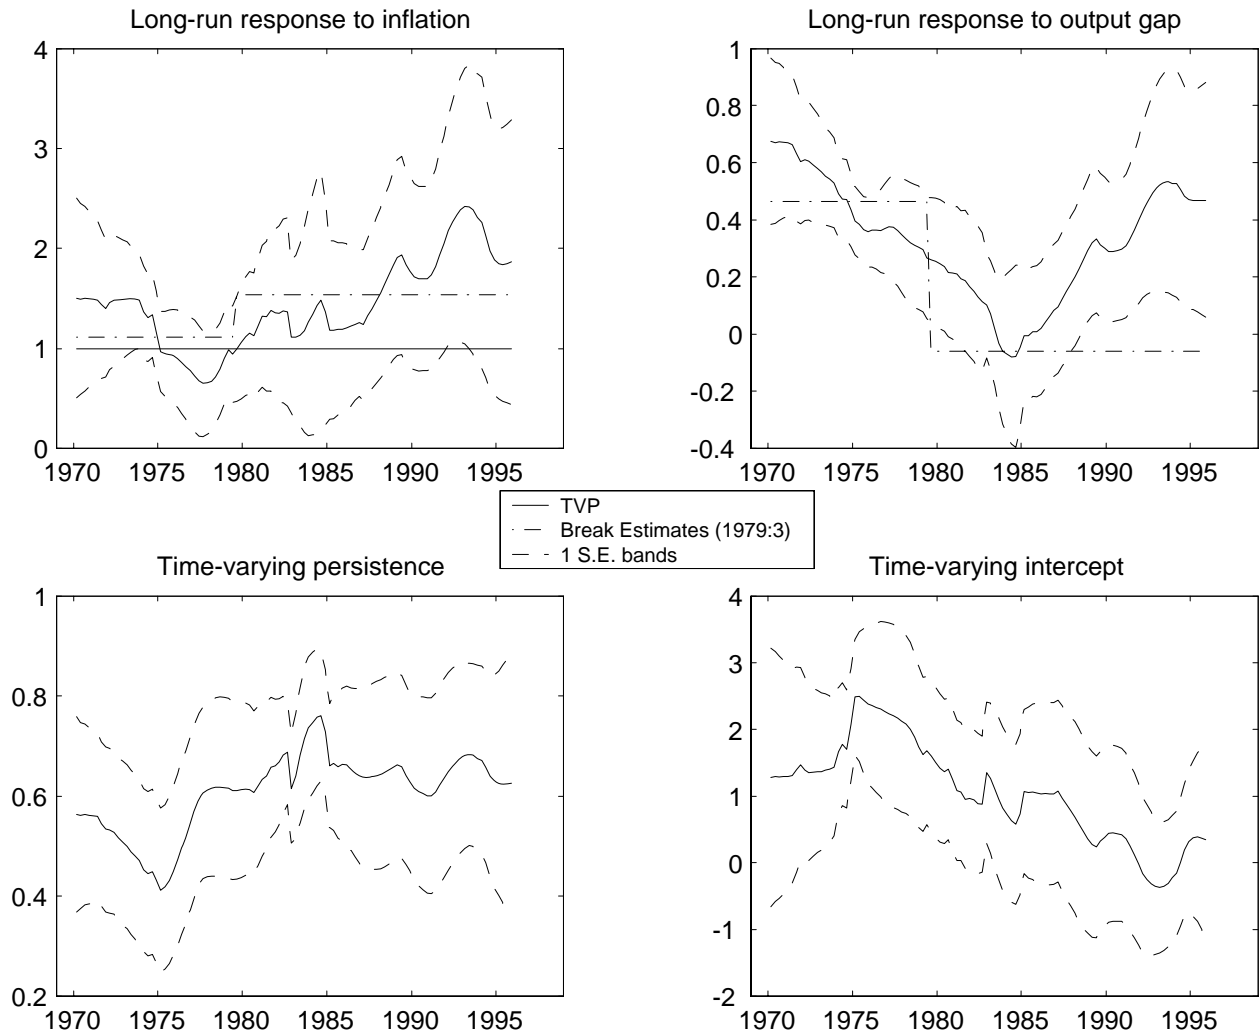


Figure 9: TVP estimates: Quarterly data and real-time output gap measure (Orphanides (2003))