

Has U.S. Monetary Policy Changed?

Evidence from Drifting Coefficients and Real-Time Data

Jean Boivin*

Columbia University and NBER

September 2003

*Graduate School of Business, 821 Uris Hall, 3022 Broadway, New York, NY 10027. E-mail: jb903@columbia.edu.

This paper builds on the third chapter of my thesis entitled: “The Fed’s Conduct of Monetary Policy: Has it Changed and Does it Matter?” I thank Ben Bernanke, Marc Giannoni, Rick Mishkin, Athanasios Orphanides, Mark Watson and participants at the Columbia Macro Lunch for useful comments and Dean Croushore for assistance with the Greenbook data. Financial support from the National Science Foundation (SES-0214104) is acknowledged. All errors are my own.

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 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 estimation. This paper treats these issues in the context of a forward-looking Taylor rule 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. As in Clarida, Gali and Gertler (2000), but contrary to Orphanides (2002, 2003), the Fed's response to inflation evolved from a weak response in the mid-1970's, not satisfying Taylor's principle at times, to a stronger response thereafter. Moreover, as in Orphanides (2003), the Fed's response to real activity fell substantially in the 1970's.

JEL classification: E52, C32, C50

Keywords: Forward-looking Taylor rule; Greenbook; Parameter instability; Real-time data; Time-varying parameter model.

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. The Fed's behavior during that period is one hypothesis.

The monetary policy hypothesis relies on the existence of important changes in its conduct. From a narrative perspective, this seems fairly uncontroversial, although the reasons of 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 influenced 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 changes.

Clarida, Gali 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

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 might 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.

policy accommodated inflation in the 1970's, it drastically and suddenly changed with the appointment of Volcker in 1979. They emphasize that the pre-Volcker conduct of monetary policy did not satisfy Taylor's 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.

This evidence has been challenged along three dimensions. First, Boivin (1999) and Cogley and Sargent (2001) argue that the Fed 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 broadly consistent with those of Clarida, Gali 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 (2001)'s 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 (2002), 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 (2003), 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 if 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 very distorted pictures of the historical conduct of monetary policy. Orphanides (2002) concludes that when the Clarida, Gali 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 estimation is taken into account.

The goal of this paper is to shed some light on this debate by considering an empirical framework that accounts for these three potential issues. To do so, I consider a forward-looking Taylor rule with time-varying coefficients, estimated using real-time data. Following the pioneering work of Cooley and Prescott (1978), the time variation is modeled as driftless random walks, as in Boivin (1999) and Cogley and Sargent (2001, 2003), and is estimated using the median-unbiased estimator proposed by Stock and Watson (1998). Given that there is considerable uncertainty about the true nature of the time variation, I argue that this is a robust way of uncovering it. Importantly, the estimation recognizes explicitly the presence of heteroskedasticity, allowing for changes in the variance of the policy shock.

The forward-looking behavior is modeled directly using the forecasts produced in real time by the staff economists at the Board of Governors, published in the Greenbook. Using these forecasts has two important advantages worth noting at the outset. First, as emphasized by Orphanides (2001, 2002), these forecasts are based on the information that was actually available whenever historical policy decisions were made. As I explain below, this presumably makes the description of monetary policy more accurate and realistic. Second, in the presence of parameter instability and under some reasonable assumptions, using them makes the analysis robust to the Lucas (1976) critique. The time variation can thus be modeled in a single equation framework, without having

to take a stand on a model for the rest of the economy.

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. The gradual nature of the uncovered changes seems to be a key explanation behind the conflicting conclusions reported in the literature, based on split-sample estimation. Overall, I find that monetary policy improved over that period. It evolved from a weak, accommodative, response to inflation in the 1970’s, to a more aggressive response in the 1980’s and early 1990’s. In parallel, the response to real activity weakened considerably from the mid-1970’s onward.

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 and other specifications of the policy rule. 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 (1993)’s 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, as well as a normative prescription.⁴ Taylor (1993) argues that a constrained version of this rule – with no lags of R_t , π_t^e , current inflation, x_t^e , 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. For instance, Clarida, Gali and Gertler (1998) show the empirical success of (1) for Germany, Japan, UK, France and Italy, over the post-1980 period.

How should we account for the Fed’s forward-looking behavior empirically? One approach is to explicitly model and estimate the expectation process. This is the approach undertaken by Clarida, Gali and Gertler (2000) who assume that the Fed has rational expectations, which implies a set of moment conditions that can be used in a GMM framework.⁵ The alternative considered in this paper is instead to use historical real-time forecasts as a proxy for the Fed’s expectations. In particular, we use 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. They have been used originally in the estimation of policy rules by McNees (1986, 1992) and more recently by Orphanides (2001, 2003), Boivin (1999) and Romer and Romer (2003). Romer and Romer (2000) document their superior historical performance, relative to private sector forecasts, in forecasting inflation.⁶

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

⁵A 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).

⁶In a recent paper, Faust, Swanson and Wright (2003) confirm, using high frequency data, that the Fed forecasts of inflation outperform private sector forecasts. However, they do not find evidence that the private sector learns something from the policy announcements. For a discussion of the properties of the Greenbook forecasting errors, see Romer and Romer (2000, 2002) and Orphanides (2002).

While there is essentially no official detailed documentation on these forecasts, Reifschneider, Stockton and Wilcox (1997) describe two fundamental features that will be instrumental for the present investigation. 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”.⁷ By construction, they 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, whereby only the FOMC members should take part in monetary policy deliberations. The second important feature is that while econometric models, like the large-scale FRB/US model, are used as input in constructing the forecasts, there is a large “judgmental” component. As a result, the forecasts cannot be mechanically reproduced from a forecasting model like a VAR.

The conditional nature of the Greenbook forecasts justifies estimating (1) as a single equation. In fact, implicit in the estimation of forward-looking Taylor rule is the assumption that inflation and real activity measures respond with at least one lag to changes in the fed funds rate. Although this assumption is common in the literature,⁸ it is not obvious *a priori* that it is appropriate for the Greenbook forecasts. However, since the forecasts were generally produced under the assumption that the nominal fed funds rate does not change over at least the next year and a half, it implies in particular that the forecasts are not responding within the month to changes

⁷Although, as they note, in some occasions the assumption of unchanged fed funds rate is so inconsistent with the committee’s consensus view that the forecasts supplied to the FOMC were conditioned on an alternative policy scenario. Unfortunately, these instances are not documented.

⁸For instance, this is the assumption used for the same equation by Orphanides (2001, 2002). More generally, in the VAR context it corresponds to the recursive assumption used by Bernanke and Blinder (1992), Rotemberg and Woodford (1997), Bernanke and Mihov (1998) and Christiano, Eichenbaum and Evans (1999).

in the policy instrument. Besides, even if in some occasions, the forecasts were conditioned on a different baseline policy scenario, this is not likely to contaminate our results. In fact, given the known lags in the effect of monetary policy – typically estimated from VARs to be more than a year – the policy assumption should not significantly influence the forecasts at the short horizons used in this paper.⁹ At any rate, if these departures from the no change assumption were important, the influence of the policy scenario on the forecasts should grow as the horizon increases. Yet, the qualitative results I report below are essentially unaffected by the choice of horizon.

The Greenbook forecasts are also likely to better account for the Fed’s information set than typical VARs do. This is the case for three reasons. First, 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 lot 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. Third, these forecasts are produced using the information that was actually available to the central bank whenever they were produced. In particular, as emphasized by Orphanides (2001, 2002) and contrary to most empirical applications, the forecasts do not assume that the forecasters or the policymakers knew the *ex post* revisions to the data. As a result, the estimation of the policy rule is entirely “real-time”.

The information set embedded in the Greenbook forecasts should also reflect, at any point in time, the state of knowledge about potential changes in the structure of the economy. This seems reasonable since these are “judgemental” forecasts and if the Board staff observed any such changes, the forecasting paradigm should have been modified accordingly. Under this assumption, the analysis in this paper should be robust to the Lucas (1976) critique. In particular, the impli-

⁹As described below, horizons up to three quarters ahead are considered.

cations of the changes in the policy rule on the rest of the economy does not need to be modeled explicitly. Similarly, if for a while, a structural change went unnoticed by the Fed, the Greenbook should appropriately reflect this ignorance.

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 stability tests, such as the Quandt (1960) likelihood ratio test, do not provide useful guidance on the form 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 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 Lucas (1976) critique on the inappropriateness of stable econometric models

for policy evaluations. It has been widely used in forecasting applications,¹⁰ 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 general form of time variation, for each parameter separately. Of course, if the true change is a single discrete jump, and 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 potentially still provide 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. As such, it cannot reproduce exactly the discrete change. Yet, fitting this model to the data will imply small ω_t for t far from, and larger ω_t for t closer to the break date, potentially uncovering the general movement in the parameters. It is thus possible to determine whether the uncovered changes are consistent at all with a discrete break model.

Moreover, as soon as there are 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 breaks model. The reason is that the multiple breaks model requires the estimation of many parameters, an estimate on each side of the breaks, for each parameter. Moreover, it requires an estimate of the break dates, which in macroeconomic applications involves considerable uncertainty.¹¹ 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.

¹⁰See Stock and Watson (1996).

¹¹See for instance Stock and Watson (2002). This is also a feature of the results presented below.

The TVP model 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 0.¹² 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 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.

¹²That is, the MLE has a (large) point mass at 0. See Stock and Watson (1998) for a discussion of this point.

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. 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 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 makes the asymptotic distributions of the tests dependent on λ .

>From Theorem 1 of Stock and Watson (1998), the asymptotic distribution of the Quandt (1960) likelihood-ratio test,¹³ QLR_T , under the TVP alternative, is a functional of Brownian motions, indexed by λ :

¹³This 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.

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

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. 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.

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 $var(\Delta\Phi_t)$ thus 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. $\hat{\Sigma}_{xx}$ is consistently estimated with $T^{-1} \sum_{t=1}^T Z_t Z_t'$ and, since asymptotically there is no time variation in the parameters, Ω is 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. $\hat{\Sigma}_{vv}$ can be estimated 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.¹⁴

5 Empirical results

5.1 Data

Equation (1) is estimated using monthly data and two lags of R_t . R_t is the annualized monthly average of the fed funds rate and $\pi_{t+h_\pi|t}$ is the Greenbook forecast of the annualized percentage change in the GNP deflator between t and $t + h_\pi$. 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, $x_{t+h_x|t}$ 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 about, not only the appropriate measure but also, the relevant concept of the natural rate of unemployment. To maintain the real-time nature of the estimation, in the baseline specification, I use as a measure of the natural rate at time t an historical average of the unemployment rate up to that point. This is the same measure that Orphanides (2002) presents. I investigate below the robustness of the baseline results to alternative real-time moving

¹⁴Note 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.

average measures of the natural rate of unemployment.¹⁵

The earliest Greenbook forecasts available are for the October 1965 FOMC meeting and are made publicly available with a five year lag.¹⁶ The forecasts have been produced for various 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.¹⁷ Also, the inconsistency arising from the change in the measure of output forecasted – from GNP to GDP – required the elimination of observations after 1991:10. 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 1973:6 to 1991:10. While this is broadly consistent with the horizons considered in the literature,¹⁸ 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.

¹⁵Note that the time-varying constant accommodates, to some extent, mis-specification in the natural rate of unemployment.

¹⁶Note that the current values of these variables are not known within the period. Until 1979, the forecasts are available on a monthly basis. Thereafter, there are usually eight data points a year, typically six weeks apart. 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.

¹⁷The starting dates for the different horizons are as follows. Current quarter (C-Q): 1971:9. One quarter ahead (1-Q): 1972:12. Two quarters ahead (2-Q): 1972:12. Three quarters ahead (3-Q): 1973:6. A consistent time series for the four quarter-ahead (4-Q) horizon only starts in 1977:2.

¹⁸Orphanides (2002) uses the same horizon specification.

5.2 Baseline results

Before modeling the time variation in the parameters of the policy rule, the first step is to provide statistical evidence that (1) indeed displays instability. The p -value of the heteroskedasticity-robust version of the QLR_T test, applied jointly to all coefficients of this equation, is 0.000.¹⁹ The 95% confidence interval for the break date estimate is [1977:9 1983:9].²⁰ This is consistent with a change in policy occurring in the late 1970's or early 1980's, but there is considerable uncertainty about the exact timing.

The stability test result imply a median unbiased estimate of λ equal to 8.5. The size of $\hat{\lambda}$ implies relatively small period to period variation in the parameters. As I alluded to in the description of the estimation approach, this is within the range of values of λ for which the MLE tends to run into problems.²¹

Given the existing empirical results, it is useful to consider the estimates of the parameters, assuming a one-time structural break. A natural break date, within the confidence interval 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 tenure as chairman of the Fed. The results from the policy rule estimated on each side of this date are reported in Table 1. The estimates of β and ρ_i imply a long-run response to inflation, $\frac{\beta}{1-\rho(1)}$, equal to 1.33 and 1.23 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

¹⁹The p -value is obtained using Hansen (1997)'s simulation procedure.

²⁰The confidence interval is computed using the asymptotic distribution of the break date estimate, as derived by Bai and Perron (1998).

²¹Stock 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).

point made by Orphanides (2002, 2003).²²

Table 1: Split sample policy rule estimates with break in October 1979

	α	β	γ	ρ_1	ρ_2
Pre-1979:9	0.138 (0.290)	0.096 (0.058)	0.134 (0.053)	1.338 (0.102)	-0.415 (0.099)
Post-1979:10	0.648 (0.238)	0.227 (0.054)	-0.028 (0.058)	1.228 (0.077)	-0.416 (0.074)

(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 1979:10. 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 their 90% confidence bands, as well as the split-sample estimates reported in the previous table.²³ The bottom panels reports the estimates of the fed funds rate persistence, $\rho_t(1)$, and the intercept, α_t , together with their 90% confidence bands.

Contrary to the split-sample estimates, the TVP results suggest that the conduct of monetary policy greatly evolved since the 1970's. The response to inflation varies roughly between 0.1 and

²²I get a lower estimates than Orphanides (2002, 2003) for the long-run response to inflation in the first sample. Note that Orphanides (2002, 2003) uses quarterly data, assumes that first regime ends in the second quarter of 1979, and his estimation starts in 1966:1.

²³Under the estimation framework described in the previous section, the asymptotic variance of the time varying parameter are obtained from their OLS estimates.

2.3 , that to unemployment between 0.2 and 2.6, and the persistence of the fed funds rate between 0.82 and 0.96.²⁴

The uncertainty about the exact size of the parameters, as suggested by the confidence intervals, is not uncommon for such macroeconomic data. Fortunately, however, the parameters are sufficiently precisely estimated to detect a clear pattern in their behavior. Moreover, the time variation in these parameters is statistically significant, as established from the stability tests above.

In the early part of the sample, i.e. late 1973 and early 1974, the response to inflation was temporarily above one. This short episode could be consistent with Orphanides (2003)'s conjecture that the Fed intended to respond aggressively to inflation, but could not detect an adverse shift in trend productivity. At the end of 1974, however, it fell drastically and persistently to a level well below one. It increased continuously thereafter with a very sharp increase in the early 1980's. It remained more or less constant from that point onward, perhaps with a mild increase in the late 1980's. The salient feature of the response to unemployment is its sharp and persistent decline in 1976. The persistence of the fed funds dropped considerably in the early 1980's; this is in fact to be expected since this period was characterized by an atypical volatility of the fed funds rate.

Concerning the timing of these changes, it is important to note that given that the 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 with some delay. In fact, the timing of the changes in the Fed's response to inflation is surprisingly consistent with the historical account of this period provided by Romer and Romer (1989). The peak in the early part of the sample corresponds precisely to April 1974, one of the dates they single out as an unusual tightening of monetary policy. From 1975 to August 1978 they argue that "the actual commitment to combat inflation was weak," consistent with the

²⁴An 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.

TVP estimates. In 1978 and 1979 the estimated response to inflation increases more rapidly, to a peak in October 1979, another date they single out.

Overall, monetary policy appears to have gradually paid more attention to inflation and less to real activity. On the one hand, the general evolution of the response to inflation tends to confirm the finding of Clarida, Gali and Gertler (2000), which was too weak for most of the period before Volcker. On the other hand, the decline in the response to real activity is consistent with the recent results of Orphanides (2003).

But the TVP estimates differ in important dimensions from the existing results in the literature. First, the changes in both the response to inflation and real activity are sizeable and important to characterize the evolution of monetary policy. Second, the uncovered behavior of the parameters is more complexed than any of the existing split-sample estimates suggest: as Romer and Romer (2002) argue, this evolution does not appear uni-directional and seems, at least in part, gradual. It is true that, based on the description provided in the previous Section, the TVP estimation would tend to make discrete changes look more gradual than they really are. But the TVP estimates seem to have no problem picking up a rapid increase in the response to inflation in late 1979. Moreover the uncovered changes are not limited to that episode and over the whole sample they appear largely gradual. This is consistent with evolving views and their gradual implementation within the Fed.

The gradual nature of the changes is also reinforced by a comparison of the split-sample and TVP estimates. It is interesting that these two approaches give so different results. Looking at the TVP estimates, the choice of a break date appears particularly problematic: the response to inflation is quite erratic around 1980 and its change do not coincide with those of the response to real activity. To investigate the influence of choice of the break dates, I re-estimate the first sample

for slightly different ending dates. The results are reported in Table 2. If the true nature of the change is really discrete, ending the first sample slightly earlier should not dramatically affect the results. Yet, the results appear extremely fragile to the choice of the break date. For instance, when ending the first sample in 1979:7 – that is just two months earlier than in the benchmark – the long-run response to inflation in this case drops from 1.33 to 0.85! Alternatively, ending the sample two months later, in 1979:11, the estimates increases to 1.65. In light of the TVP results reported, this is not entirely surprising since the estimates in 1979 are quite above average, compare to their values throughout the 1970's. Moreover, imposing a break in the response to unemployment in 1979 should bias its estimate downward, which also biases the response to inflation. This evidence suggests that the way instability is modeled is crucial for a proper characterization of monetary policy in the 1970's. This provides further evidence of the gradual nature of the actual changes in monetary policy and also of the fragility of split-sample estimates in the present context.

Table 2: First sample policy rule estimates for alternative break dates

	α	β	γ	ρ_1	ρ_2
Pre-1979:7	0.350 (0.315)	0.062 (0.061)	0.148 (0.054)	1.312 (0.103)	-0.385 (0.101)
Pre-1979:8	0.230 (0.303)	0.083 (0.060)	0.142 (0.054)	1.321 (0.104)	-0.397 (0.101)
Pre-1979:9	0.138 (0.290)	0.096 (0.058)	0.134 (0.053)	1.338 (0.102)	-0.415 (0.099)
Pre-1979:10	-0.304 (0.334)	0.149 (0.069)	0.082 (0.063)	1.435 (0.121)	-0.508 (0.117)
Pre-1979:11	-0.064 (0.340)	0.145 (0.072)	0.150 (0.061)	1.246 (0.107)	-0.333 (0.106)

(Standard errors in parenthesis)

Given the behavior of the TVP estimates in 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 in 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. The choice of the break date is not innocuous, but based on the previous discussion, 1979:7 appears to be a sensible choice for the break date in β_t .²⁵ Letting D_t

²⁵Note 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 not the case for the nested model.

be a dummy equal to one for t greater than 1979:7, the estimated policy reaction function is:

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

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 estimate together with the previous TVP results. There are two key differences: 1) the response to inflation is larger for most of the sample when a break in β and γ is allowed and 2) the change in the response to unemployment is not as drastic in this case. But importantly, the conclusion about the evolution of the response to inflation remains intact: the changes are important and largely gradual. Notice in particular that for the volatile period of the late 1970's, early 1980's, assuming or not a discrete change yields very similar results. Consequently, the TVP model does provide a good approximation to a discrete change in this period – if there were indeed such a change.

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 the natural rate of unemployment.

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 to make 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.²⁶ 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.²⁷ 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 1984:1 as the second of the three regimes. The key difference with the previous TVP estimates is that the variation in the parameters post-1984 is more pronounced. This is not surprising, however: since the variance of the policy shock post-1984 is about ten times smaller than previously estimated,²⁸ more of the fluctuation in the fed funds rate is attributed to variations in the parameters. In fact, a lot of this variation seems to be due to the fed 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.

²⁶The federal funds rate provides probably a less adequate measure of monetary policy stance in that period, 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.

²⁷Boivin (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.

²⁸In the baseline estimation, the post-1979:10 variance of the policy shock was 0.740, whereas now it is 1.907 for the 1979:10-1984:1 period and 0.080 for the post-1984:1 period. The variance of the shock in the pre-1979:10 period is 0.149.

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 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 $h_\pi = 3, h_x = 0$.²⁹ As is clear from the Figure, no matter the horizon combination selected, there is important variation in the parameters, implying a clear evolution of monetary policy in the 1970's and early 1980's. The only notable difference is for the specification where the Fed responds to the current value of both inflation and unemployment gap. In this case, the response to inflation is smaller, and near one in the mid to late 1980's. 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 even when the Fed is assumed to respond to current inflation and unemployment gap, there is a clear distinction in the conduct of monetary policy before and after the early 1980's.

5.3.3 Natural rate of unemployment

Finally, perhaps the most important source of uncertainty about this specification concerns the appropriate measure of the natural rate of unemployment. 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 measures, corresponds to the 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. For instance, the true natural rate

²⁹Clarida, Gali and Gertler use $h_\pi = h_y = 1$ as their baseline specification.

might be more variable. Bernanke and Boivin (2003) consider on a one-sided five year moving average of the unemployment rate as a proxy for the natural rate. Figure 7 displays the baseline measure together with three alternative measures, displaying varying degree of smoothness. These are based on a one-sided moving average of the unemployment rate: 1) a five year moving average, 2) a three year moving average and 3) a one year moving average. Of course, these univariate measures do not cover the wide set of potential definitions of the natural rate.³⁰ But it is apparent from the Figure that 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 measure of the natural rate differ substantially, the implied time variation is overall 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 twelve month moving average measure, it falls sooner in the late 1970's, and increases more rapidly from the mid-1980's onward. But for all the other measures, the behavior of the response to the unemployment gap is also very similar.

Consequently, the finding of economically important changes in the policy rule coefficients, consistent with a gradual improvement in the conduct of monetary policy, appears to be robust to sensible perturbations along these three dimensions.

³⁰An alternative would be to back out the estimate of the natural rate implicit in the Greenbook forecasts. However, as Romer and Romer (2002) document using a simple Phillips curve, the resulting measure is very volatile and not well defined in the 1970's. One interpretation is that the Fed did not use a concept of the natural rate during that period. An alternative interpretation is that the assumed Phillips curve is inconsistent with the one implicit in the forecasts. In any case, using this measure in the estimation of the forward-looking rule yields the same qualitative conclusions, although the time variation looks even more drastic and arguably unreasonable (e.g. the coefficients on both $\pi_{t+h_\pi|t}$ and $x_{t+h_x|t}$ are negative in the early 1970's). These results are not reported here.

6 Conclusion

Determining whether or not the conduct of monetary policy changed is central to our understanding of the improved U.S. economic performance of the last twenty years. Although there has been 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 importantly over the last 40 years. An important dimension of this evolution is the size of the response to inflation. Monetary policy appears to have been too accommodative in the 1970's and gradually became more responsive in the 1980's. Moreover, changes in the response to the unemployment gap are also sizeable; for instance, monetary policy responded three times more in the mid-1970's to real activity than it did in the early 1980's.
- The observed changes in the response to inflation and unemployment do not appear to be well captured by a discrete shift, occurring close to 1980. In fact, the results suggest that assuming a structural break can be misleading in this case.

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 S. and Jean Boivin (2003), “Monetary Policy in a Data-Rich Environment,” *Journal of Monetary Economics* 50:3, 525-546.
- [6] Bernanke, Ben and Ilian Mihov (1998), “Measuring Monetary Policy,” *Quarterly Journal of Economics*, 869-902.
- [7] Blinder, Alan (1985), “Comment on Federal Deficits, Interest Rates, and Monetary Policy,” *Journal of Money, Credit and Banking*, 17, 685-689.
- [8] Boivin, Jean (1999), “The Fed’s Conduct of Monetary Policy: Has it Changed and Does it Matter?” Ph.D. Thesis, Princeton University.

- [9] Boivin, Jean and Marc Giannoni (2003), “Has Monetary Policy Become More Effective?,” NBER working paper #9459.
- [10] Cogley, Timothy and Thomas J. Sargent (2001), “Evolving Post World War II U.S. Inflation Dynamics,” *NBER Macroeconomics Annual* 16, pp. 331-373.
- [11] Cogley, Timothy and Thomas J. Sargent (2003), “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.,” Manuscript, UC Davis.
- [12] Cook, Timothy (1989), “Determinants of the Federal-Funds Rate,” FRB Richmond Economic Review 75(1), Jan-Feb, 3-19.
- [13] Cooley, Thomas and E. Prescott (1976), “Estimation in the Presence of Stochastic Parameter Variation,” *Econometrica*, 44:1, 167-184.
- [14] Clarida, Richard, Jordi Galí and Mark Gertler (1998), “Monetary Policy Rules in Practice: Some International Evidence,” *European Economics Review*, 1033-1068.
- [15] Clarida, Richard, Jordi Galí and Mark Gertler (2000), “Monetary Policy Rule and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*.
- [16] DeLong, J. Bradford (1997), “America’s Only Peacetime Inflation: the 1970’s,” *Reducing Inflation*, NBER Studies in Business Cycles, Vol. 30.
- [17] Faust, Jon, Eric Swanson, and Jonathan Wright (2003) “Do Federal Reserve Policy Surprises Reveal Inside Information About the Economy?,” Federal Reserve Board.
- [18] Froyen, Richard T. (1974), “A Test of Endogeneity of Monetary Policy,” *Journal of Econometrics*, 2, 175-188.

- [19] Hansen, Bruce E. (1997). "Approximate Asymptotic P Values for Structural Change Tests." *Journal of Business Economic and Statistics* 15(1): 60-67.
- [20] Lucas, Robert E. Jr. (1976), "Econometric Policy Evaluation: A critique." *Carnegie Rochester Conference Series on Public Policy*, 1, 19-46.
- [21] 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.
- [22] McNees, Stephen K. (1986), "Modeling the Fed: A Forward-Looking Monetary Policy Reaction Function." *New England Economic Review*, 3-8.
- [23] McNees, Stephen K. (1992), "A Forward-Looking Monetary Policy Reaction Function: Continuity and Change," *New England Economic Review*, 3-13.
- [24] Orphanides, Athanasios (2001), "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, 91:4, 964-985.
- [25] Orphanides, Athanasios (2002), "Monetary Policy Rules and the Great Inflation," *American Economic Review*, 92:2; 115-121.
- [26] Orphanides, Athanasios (2003), "Monetary Policy Rules, Macroeconomic Stability and Inflation: A View from the Trenches" *Journal of Money, Credit and Banking*, forthcoming.
- [27] Potts, Glenn T. and Dudley G. Lockett (1978), "Policy Objectives of the Federal Reserve System," *Quarterly Journal of Economics*, 17, 525-534.
- [28] 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.

- [29] 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.
- [30] 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).
- [31] Romer, Christina and David Romer (2000), “Federal Reserve Information and the Behavior of Interest Rates,” *American Economic Review*, 90, 429-457.
- [32] 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.
- [33] Romer, Christina and David Romer (2003), “A New Measure of Monetary Shocks: Derivation and Implications,” NBER working paper #9866.
- [34] Sargent, Thomas (1999), “*The Conquest of American Inflation*”, Princeton University Press, Princeton, NJ.
- [35] Sims, Christopher (1999), “Drifts and Breaks in Monetary Policy,” Manuscript, Princeton University.
- [36] Sims, Christopher (2001), “Discussion of Cogley and Sargent ‘Evolving Post World War II U.S. Inflation Dynamics,’ *NBER Macroeconomics Annual* 16, pp. 373-379.
- [37] Sims, Christopher and Tao Zha (2002), “Macroeconomic Switching,” Manuscript, Princeton University.

- [38] 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.
- [39] 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.
- [40] Stock, James H. and Mark W. Watson (2002), "Has the Business Cycle Changed and Why?" *NBER Macroeconomics Annual* 17, pp. 159-218.
- [41] Taylor, John B. (1993), "Discretion versus Policy Rules in Practice," *Carnegie Rochester Conference Series on Public Policy*, 39, 195-214.
- [42] 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.
- [43] Taylor, John B. (1999), editor, *Monetary Policy Rules*, NBER, Chicago University Press.

Figure 1: Real time data used in the estimation

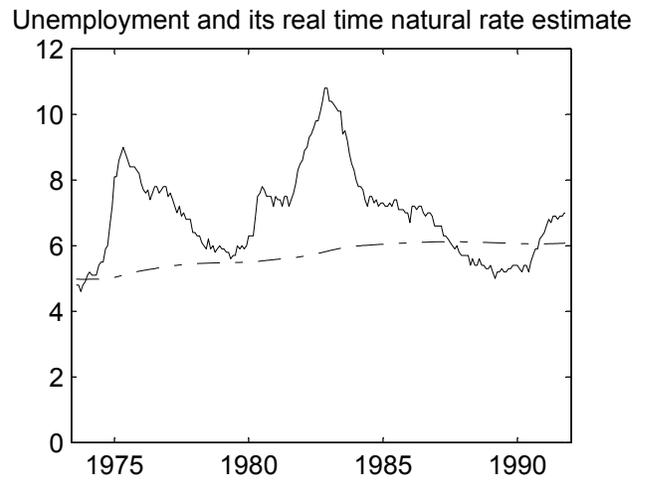
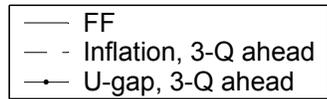
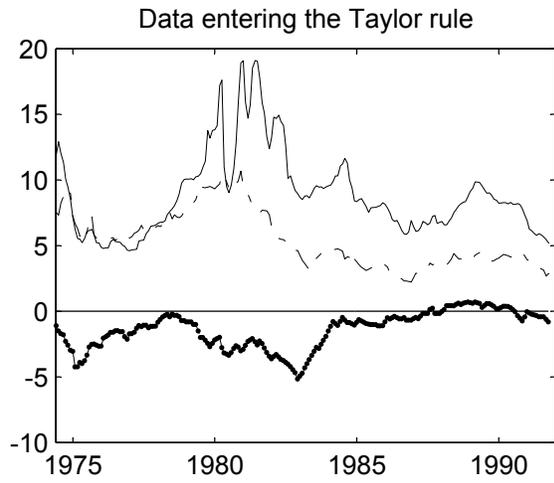


Figure 2: Time Varying Policy Parameters
(Break in the variance in 1979:10)

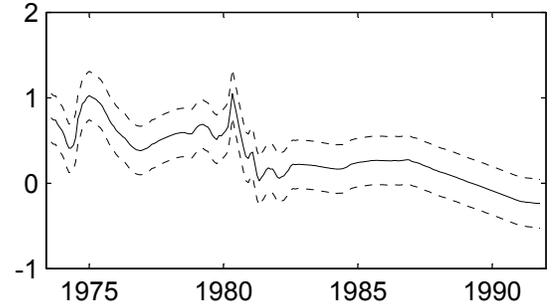
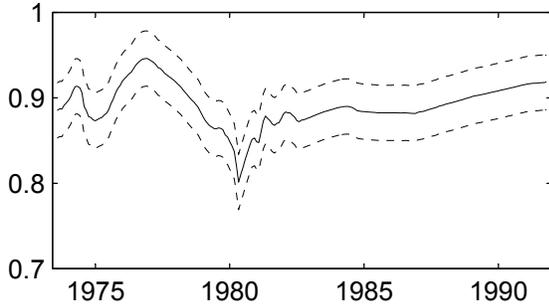
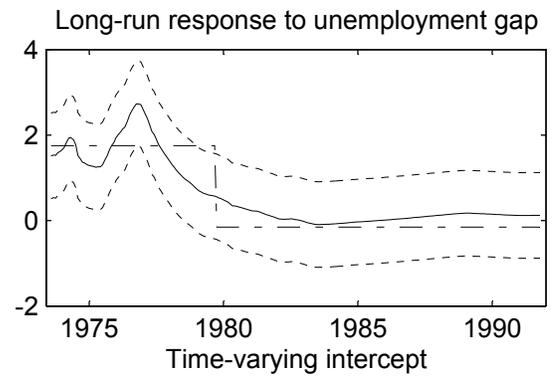
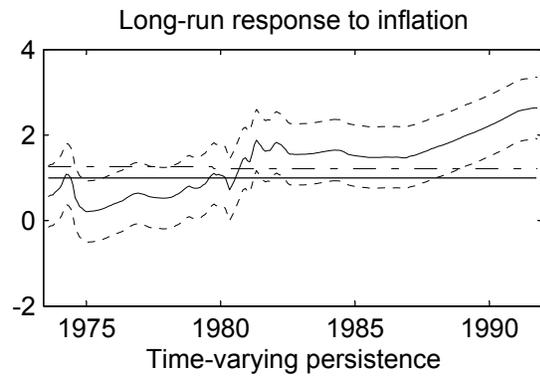


Figure 3: Time-varying parameter with discrete shift in β and γ
(Break in the variance in 1979:10)

— TVP
- - TVP with Break in 1979:7

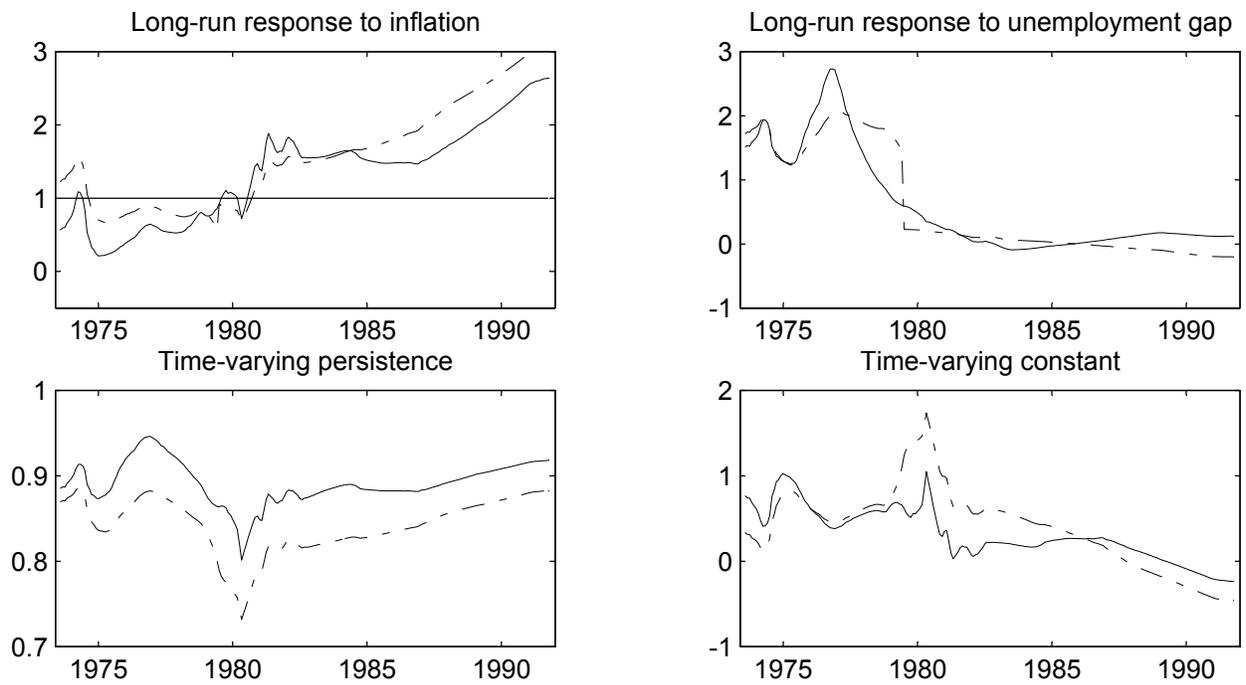


Figure 4: Estimated policy shocks: Full-sample, constant parameters

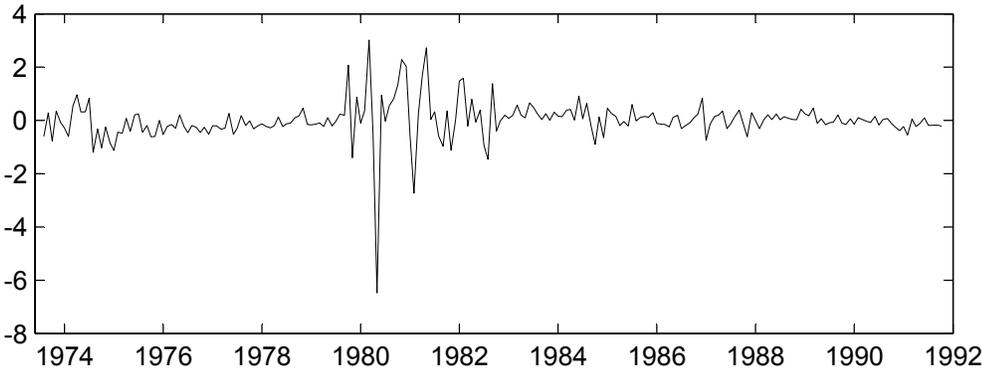


Figure 5: Time Varying Policy Parameters
(Break in the variance in 1979:10 and 1984:1)

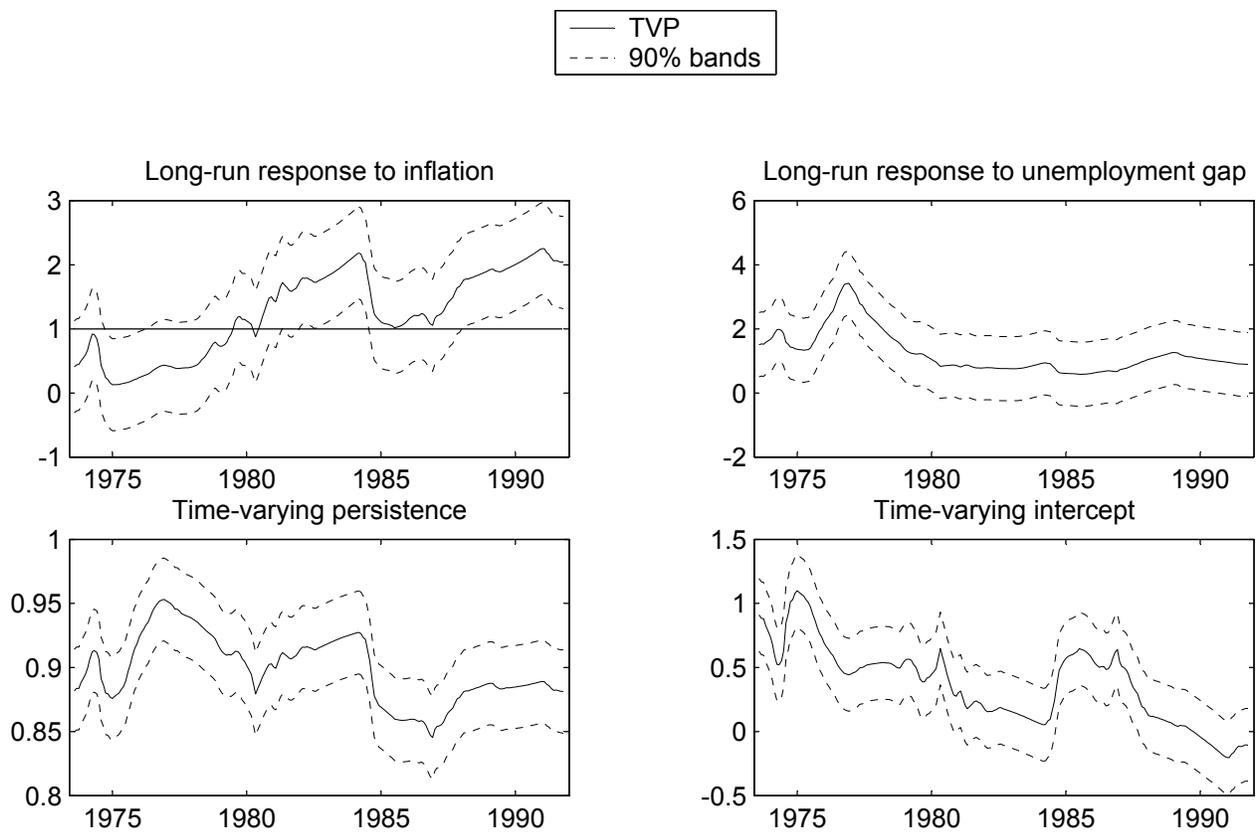


Figure 6: Time Varying Policy Parameters
 Different Forecasting Horizons: [Inflation, Gap]
 (Break in the variance in 1979:10)

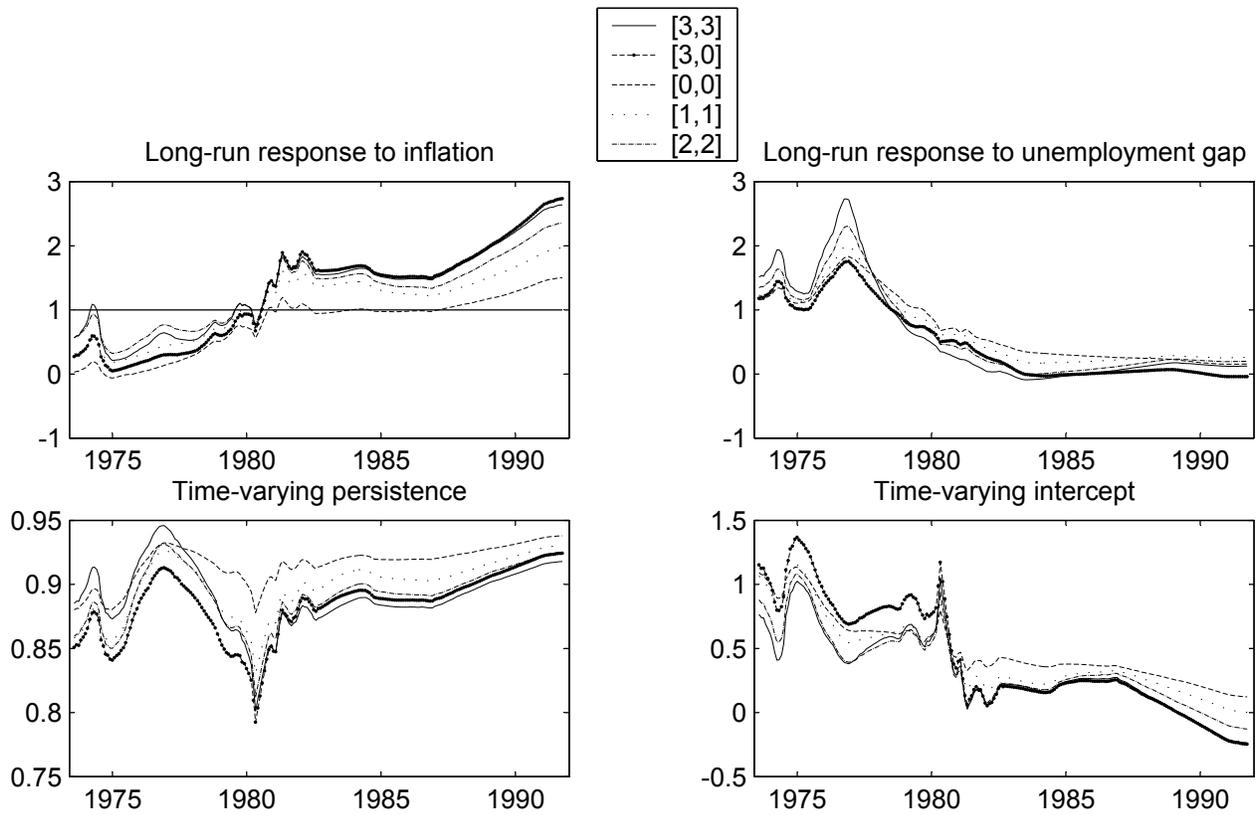


Figure 7: Alternative real-time measures of the natural rate of unemployment

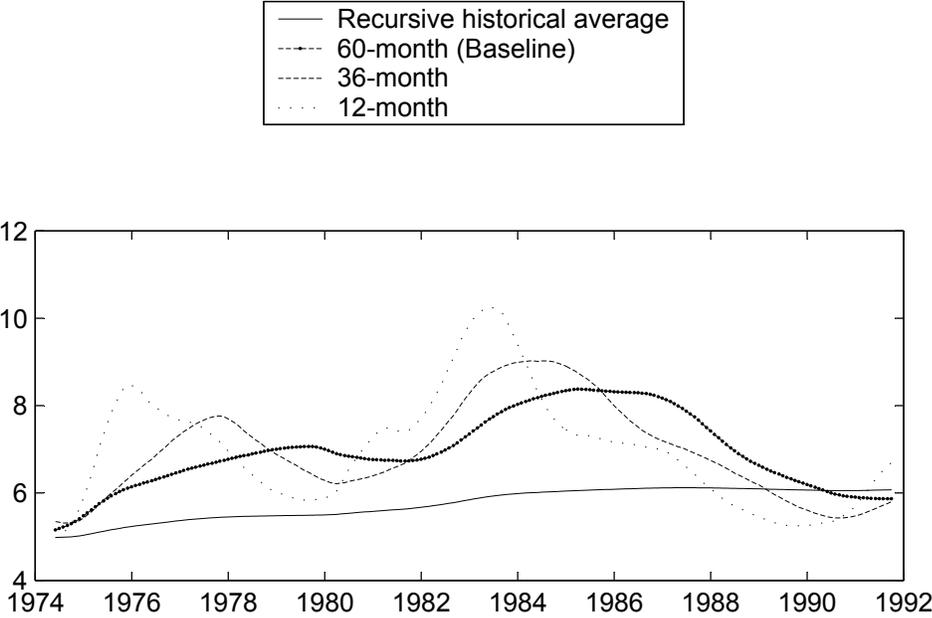


Figure 8: Time Varying Policy Parameters
 Alternative real-time measures of the natural rate of unemployment
 (Break in the variance in 1979:10)

