BANK OF ENGLAND

Working Paper No. 434

Evolving UK and US macroeconomic dynamics through the lens of a model of deterministic structural change

George Kapetanios and Tony Yates

July 2011

BANK OF ENGLAND

Working Paper No. 434

Evolving UK and US macroeconomic dynamics through the lens of a model of deterministic structural change

George Kapetanios[1] and Tony Yates[2]

Abstract

Using a model of deterministic structural change, we revisit several topics in inflation dynamics explored previously using stochastic, time-varying parameter models. We document significant reductions in inflation persistence and predictability. We estimate that changes in the volatility of shocks were decisive in accounting for the great moderations of the United States and the United Kingdom. We also show that the magnitude and the persistence of the response of inflation and output to monetary policy shocks has fallen in these two countries. These findings should be of interest in those seeking to resolve theoretical debates about the sources of apparent nominal and real frictions in the macroeconomy, and the causes of the Great Moderation.

JEL classification: C10, C14.

(1) Queen Mary and Westfield College and Bank of England. Email: g.kapetanios@qmul.ac.uk
(2) Bank of England. Email: tony.yates@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. This paper was finalised on 28 April 2011.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/publications/workingpapers/index.htm

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email mapublications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Summary

This paper contributes to a body of work that has sought to describe evolutions in the dynamics of inflation and output in developed economies. That work has been preoccupied with documenting changes in the volatility of inflation and output, changes in the persistence of inflation, and changes in the impact of a monetary policy shock, among other questions.

These facts have been deployed to try to diagnose the causes of the Great Moderation; document evolutions in real and nominal frictions in the economy, and to understand their ultimate causes. The tool of choice for studies of structural change of this kind has been an econometric model that views the parameters that propagate shocks as themselves evolving over time, and behaving as though they were random, but mean-reverting process. This paper applies a very different tool to the same set of questions. We posit that the parameters that propagate shocks evolve smoothly and non-randomly, and may not necessarily be attracted back to the mean.

Why the need for a different tool to the industry standard? First, we provide some suggestive Monte Carlo evidence that models of deterministic structural change do a good job of characterising that change even when in truth that change is random in origin. Second, whether a deterministic or random parameter model is the best choice will depend on the nature of the task in hand. In the macroeconomic dynamics literature that we apply the tool to, there are reasons for at least studying what this deterministic model generates; economic theory is generally silent about the true causes of parameter change, so that we cannot choose on those grounds which econometric tool to use. This theory is however also silent about whether such change should be mean-reverting, so on these grounds it may be desirable to look at evolving macroeconomic dynamics through the lens of the deterministic model which allows structural change to be non mean-reverting.

With these motivations in mind, we take the tools to UK and US data on inflation, GDP and policy interest rates. We document several findings of interest. First, we note significant reductions in inflation persistence (using univariate models) and predictability (using multivariate models). Second, we estimate that changes in the volatility of shocks were decisive in accounting for the Great Moderations in these two countries. Third, the evidence suggests that the magnitude

and persistence of the response of inflation and output to monetary policy shocks has fallen in these two countries.

1 Introduction

This paper offers a complement to work that has sought to document changes in the time-series properties of inflation and other macroeconomic variables. Recent contributions in this field have spanned several strands of thought, including: documenting the increase in the volatility of inflation in the 1970s and 1980s, and its subsequent decline; work describing the rise and fall of inflation persistence; work attempting to account for changes in the moments of inflation in terms of drifting propagation parameters versus drifting volatilities; and work describing the evolving response of the economy to monetary policy shocks. This body of work derives its interest from the light it sheds on questions such as the following: the extent to which changes in macroeconomic performance have been due to good luck or better policy; and the extent to which inflation dynamics derive from intrinsic real or nominal rigidities, or monetary policy itself.

Accounts of the time series for inflation have imported techniques for documenting structural change supplied by the broader time-series literature. One approach has been to assume the existence of infrequent and abrupt changes in the data generating process. Researchers either divide up the time series into portions using prior information about, for example, monetary regimes (eg Benati (2006)); or they test for breaks (eg O'Reilly and Whelan (2004)). If breaks are tested for, the tests can be conducted conditional on some known potential date (work which derives from Chow (1960)) or the tests assume no prior knowledge of the break and allow it to be identified as part of the estimation (see, for example, Brown, Durbin and Evans (1974), Ploberger and Kramer (1992)). This technique does not explicitly model the cause or the incidence of structural change, though in the case of inflation persistence and the Great Moderation there is the implicit association of potential breaks with changes in the monetary policy regime. Recent work by Kapetanios and Tzavalis (2010) provides a possible avenue for modelling structural breaks explicitly.

Another strand of work on the time-series process for inflation has deployed stochastic, time-varying coefficient (hereafter STVC) models. Such models bear resemblance to simple non-linear econometric models such as bilinear models (see Tong (1990)). STVC models have been used by Cogley and Sargent (2001), Cogley and Sargent (2005), Sims and Zha (2006), Benati (2007), Benati and Mumtaz (2007), Stock and Watson (2007) and Cogley and Sargent

(2010). These models articulate an invariant, parametric process for the coefficients and volatilities of the inflation process.

Our main contribution in this paper is to document changes in the inflation process using deterministic time-varying coefficient (DTVC) models. The framework we adopt has a long pedigree in statistics starting with the work of Priestley (1965). That paper suggested that time series may have time-varying spectral densities which change slowly over time, and proposed to describe those changes as the result of a non-parametric process. This work has more recently been followed up by Dahlhaus (1996) and others who refer to such processes as semi-stationary processes. A parametric alternative to this approach has been pursued by Robinson (1989) for linear regression models and Robinson (1991) for non-linear parametric models. Recently, Orbe, Ferreira and Rodriguez-Poo (2005) extended these results to include time-varying seasonal effects.

In common with the 'sample-splitting' literature, our DTVC model is silent about the process that governs structural change from one period to the next. But in contrast to it, we impose that this structural change takes place sufficiently gradually so that our view of the regime at any particular date can be informed by looking at data from adjacent dates.

Like the STVC models, our DTVC model has time-varying parameters, and will potentially uncover as many values for the coefficients that define the inflation process as there are data periods. But this time variation is characterised very differently in the two approaches. To give a concrete example, suppose we model inflation as a time-varying function of a single lag. The STVC describes the time variation in the coefficient on that lag as an invariant, parametric process. Although the coefficient that propagates past inflation into future inflation changes, the inflation process itself is time-invariant. The DTVC model makes no such assumption; the relationship between inflation and its lag would be purely data-determined. Depending on the context, it may or may not be appropriate to impose this time-invariance.

Estimates of STVC models typically impose that coefficients and volatilities follow random walks to avoid estimating autoregressive terms in these processes, and therefore embody the assumption that the shocks to the propagation parameters and volatilities are permanent. Our

DTVC models provide a practical way to avoid imposing this restriction.[1] Many applications combine the random walk assumption with the restriction that the the instantaneous VAR formed by freezing parameters at a point in time is itself stationary. We avoid that assumption here. We explain in an appendix that this assumption can be shown to imply that the time series are stationary. That assumption is desirable because the consistency properties of the maximum likelihood estimators used rely on stationarity. Once again, depending on the context, it may or may not be accurate or desirable to characterise the dependent variable as stationary. Our method at least provides a route to go if a researcher prefers to avoid this assumption.

A final contrast between the approach taken in this paper and the STVC papers is that the latter have typically chosen to use Bayesian methods. We have adopted a frequentist approach here. But nothing about the DTVC methods necessitates us doing this. In particular, there is a considerable literature on non-parametric Bayesian methods, which can accommodate our approach of modelling time variation, as a deterministic function of time. For a review of such methods, see Muller and Quintana (2004).

In some cases, prior knowledge may give a strong steer about which of the DTVC or STVC is the most appropriate tool to use to characterise evolutions in the dynamics of a multivariate time series. However, we offer some food for thought for those researchers who are convinced that the true process has parametric, stochastic time variation in its coefficients. A simple Monte Carlo study shows that DTVC models uncover more accurate descriptions of the evolution of model parameters not only in the case where the true process is a DTVC model, but also in the opposite case where the data are generated by a STVC process. This benefit shrinks as the sample size grows much beyond 1,000 periods, but for the cases considered in recent papers on macro data (40 years of quarterly data or less) the benefit is quite pronounced.

We put a DTVC model to work to describe changes in the inflation process in the United Kingdom and the United States up to 2007 Q1. We document statistically significant changes in both propagation parameters and innovations variances in both univariate and multivariate models of inflation for both countries. These show up as significant reductions in inflation persistence (in our univariate models) and in inflation predictability (in our multivariate models).

[1]It has been suggested to us that in many macroeconomic applications the random walk restriction would be accepted by the data anyway. In which case the option to relax it in our DTVC framework may not be all that valuable.

The numerical contribution of changes in innovations variances is decisive in characterising the rise and fall of inflation volatility. Univariate and multivariate representations of the inflation process tell the same story in this regard for the United States. For the United Kingdom: the multivariate account of the Great Stability assigns less weight to changes in innovations variances. These results serve as a useful footnote to the exchange between between Cogley and Sargent (2001) and Sims (2001). Cogley and Sargent (2001) had identified significant changes in propagation parameters, and Sims commented that these changes were an artifice of not having allowed for evolution over time in the variance of innovations. Cogley and Sargent's subsequent papers documented the extent to which significant changes in propagation parameters survived in the more general model Sims recommended. Following Sim's prescriptions, we allow, like the later Cogley and Sargent pieces, for time variation in both parameters and innovations variances. Both are shown to have changed significantly, though the latter contributes most to the measured Great Moderation in inflation.

Some previous work has remarked on how different changes in inflation dynamics can look when models are estimated on different price indices, so we ran our models on both CPI and GDP deflator data. Unsurprisingly, we find that the story told by our multivariate models is relatively robust to which index is used. There are some differences between the univariate models. Our models of CPI tend to show more evidence of instability than the GDP deflator models.

Finally, we study the evolution of the response of inflation and GDP to identified monetary policy shocks. For both the United States CPI and UK data sets, we find that inflation and GDP respond by more, and stays a little more persistently away from target, at the beginning of the sample, than at the end. In modern DSGE models, the persistence of these responses is a function of the strength of real and nominal frictions (like habits, capital adjustment costs, wage or price indexation).

What interpretation is to be placed on the finding that reduced-form inflation persistence has fallen, and that the size and persistence of the response of inflation and GDP to an identified monetary policy shock has fallen? One interpretation of these results is that, conditional on this class of models being true, the strength of these rigidities was greater before than now. A different interpretation would be that the finding that these moments of inflation and GDP change is evidence that models that have these frictions hard-wired into wages or prices, or in other adjustment costs are simply not adequate models of the data.

One final aspect of the contribution made in this paper is methodological. In order to put the DTVC framework to work to address the questions posed by others of the inflation process we have developed two minor extensions to the DTVC toolkit that were needed to apply it to our exploration of inflation dynamics. We have proposed an estimator for a time-varying error variance-covariance matrix; and implemented new tests for the stability of these error variances.

The rest of the paper is organised as follows. In Sections 2 and 3 we describe the toolkit that we use to model structural change in the inflation process. Section 2 lays out the model of deterministic time-varying coefficients and provides a motivating Monte Carlo study. Section 3 describes a sign-restriction method for identifying monetary policy shocks. Readers familiar with these tools can skip to our results which are set out in the next two sections. Section 4 reports our findings about changing inflation persistence and predictability. Section 5 documents our findings on the evolving responses of inflation and GDP to monetary policy shocks. Section 6 concludes.

2 A deterministic time-varying coefficient model

In this section, we document the DTVC model that we later take to the data. These tools are explained more fully in Kapetanios (2008) and Kapetanios (2007), but we draw out the salient features here.

2.1 Estimating time-varying coefficients

We focus on a VAR-type model of the form

$$y_t = \sum_{i=1}^{p} B_{i_t} y_{t-i} + u_t, \quad t = 1, \ldots, T \qquad \textbf{(1)}$$

where y_t is an m-dimensional vector of variables and the coefficient matrices are functions of time in a sense to be defined below. In practice, m will take either the value 1 when we estimate univariate models for inflation, or 3 when we estimated a VAR with inflation, growth and nominal interest rates.

For the purposes of our empirical analysis we need to provide an assumption on the time varying VAR matrix of coefficients, denoted by B_t and given by

$$B_t = \begin{pmatrix} B_{1,t} & B_{2,t} & \dots & B_{p,t} \\ I & 0 & \dots & 0 \\ 0 & I & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & I & 0 \end{pmatrix}$$

We, therefore, have the following natural assumption.

Assumption 1

$$0 < \lim_{T\to\infty} \lim_{m\to\infty} \frac{\prod_{i=1}^{m} ||B_{T-i}||}{c^m} < \infty$$

for some $0 < c < 1$.

This assumption is similar to that used for the method adopted in the STVC models of Cogley and Sargent (2001) and Cogley and Sargent (2005). We make the assumption because we want to compute 'instantaneous' variances of inflation associated with the values of VAR coefficients and shock variances observed at each point in time. These variances exist only if we make this assumption: they require projecting the hypothetical VAR for some data t infinitely far ahead, hence the restriction on the infinite product of the $B's$. Also this assumption implies mixing which is a weak dependence condition necessary for the theory on which our estimator is based.

We will estimate our model equation by equation. For the purposes of setting out the details of our estimation procedure it is convenient to focus on a given equation of the VAR model, which we will denote by the following:

$$y_{i,t} = \beta(t)' x_t + u_{i,t} \qquad \textbf{(2)}$$

where $y_{i,t}$ and $u_{i,t}$ are the i-th element of y_t and u_t respectively, $x_t = (y'_{t-1}, \dots, y'_{t-p})'$ and $\beta_i(t)$ is the i-th column of $(B_{1,t}, \dots B_{p,t})'$. We define $k = mp$. $u_{i,t}$ is given by $\sigma_i(t) v_{i,t}$.

We make a number of further assumptions, needed to formally justify the estimator, which are stated in the appendix but are informally laid out and discussed below. Our main assumption (Assumption 3) specifies that $\beta_i(t)$ is a smooth deterministic function of time. It is interesting to note that it depends not only on the point in time t but also on the sample size T. This is necessary since in order to estimate consistently a particular parameter one needs the sample size that relates to that parameter to tend to infinity. This is achieved in this context by allowing an increasing number of neighboring observations to be informative about β at time t. In other words we have to assume that as the sample size grows the function β_τ stretches to cover the whole period of the sample. A similar set-up is assumed to hold for the variance of the error term. This set-up has precedents in the statistical literature. For example, the concept of slowly varying processes of Priestley (1965) forms an early instance of similar ideas. Assumptions 4 and 5 are standard temporal dependence and moment conditions for the explanatory variables and the error term. It is important to note that x_t is allowed to be nonstationary. Assumption 4 is easily seen to be satisfied if y_t satisfies Assumption 1. We make a martingale difference assumption for the error term. Note that this assumption is not crucial and is adopted for simplicity. General forms of stationary weak dependence for the error term can be accommodated with minimal changes in the analysis. Finally, Assumption 6 relates to the kernel function, $K(.)$, that will be used for estimation. Note that Assumptions 2 and 3 are much more general than needed for the purposes of our empirical analysis since that involves VAR models. However, we choose to provide a more general setting to illustrate the applicability of the approach.

Following Robinson (1989) and Orbe *et al* (2005), we propose the following estimator for $\beta_{i,\tau}$.

$$\hat{\beta}_{i,\tau} = \left(\sum_{t=1}^{T} K_{t,\tau} x_t x_t' \right)^{-1} \left(\sum_{t=1}^{T} K_{t,\tau} x_t y_{i,t} \right) \qquad \textbf{(3)}$$

where $K_{t,\tau} = (Th)^{-1} K((\tau - t)/Th)$. h is known as the 'bandwidth'. It sets the variance of the kernel function K which determines how fast we reduce the weight placed on data as we move further and further from some t of interest. This estimator bears close resemblance to the standard OLS estimator and it is easy to see that it is the closed form solution of the following optimisation

$$\min_{\beta(t)} \sum_{t=1}^{T} K_{t,\tau} (y_t - \beta(t)' x_t)^2$$

which is obviously related to the OLS problem of minimising a sum of squared residuals.

Some intuition for what we are doing can be got by observing that we can construct 'rolling regressions' as a special case of our estimator: we would do this by specifying a uniform function for K and a bandwidth fixed *a priori*. Our estimator uses adjacent data points to inform the estimate for each period, where instead of a uniform weighting scheme we will use kernels that imply declining weights on data further away. Because we wish adjacent data to be informative about the coefficient that is obtained in a particular period, we have to assume that those coefficients move around 'slowly'. Our assumptions set out above formalise this.

2.2 Estimating time-varying error variances and correlations

In our model, we wish to allow for time variation in the variance of the shocks disturbing our VAR equations. Since our VARs are reduced-form models, we need also to allow for correlation between the shocks disturbing different equations in the VAR. It seems desirable to allow for time variation in these correlations. Here we set out the estimators that we use to do this.

Following estimation of $\beta_{i,\tau}$ we propose the following simple estimator for $\sigma^2_{i,\tau}$.

$$\hat{\sigma}^2_{i,\tau} = \frac{\sum_{t=1}^{T} K^h_{t,\tau}\hat{u}^2_{i,t}}{\sum_{t=1}^{T} K^h_{t,\tau}} = \sum_{t=1}^{T} \tilde{K}^h_{t,\tau}\hat{u}^2_t \tag{4}$$

where

$$\hat{u}_{i,t} = y_{i,t} - \hat{\beta}_i(t)'x_t \tag{5}$$

and $\hat{\beta}_,(t) = \hat{\beta}_{i,t/T}$. Here, we have assumed that the h used in **(3)** is the same as that used in **(4)**. However, they clearly do not need to be the same. Let us denote the parameter h used in **(3)** by h_β whereas the parameter h used in **(4)** is denoted by h_σ. Kapetanios (2007) proves that $\hat{\sigma}^2_{i,\tau}$ estimates consistently $\sigma^2_{i,\tau}$ and that a central limit theorem holds for $\hat{\sigma}^2_{i,\tau}$.

We suggest the following strategy for modeling time-varying correlations that extends the work of Kapetanios (2007). Let

$$\begin{pmatrix} u_{i,t} \\ u_{j,t} \end{pmatrix} = \begin{pmatrix} \sigma^2(i,t) & \rho_{i,j}(t)\sigma(i,t)\sigma(j,t) \\ \rho_{i,j}(t)\sigma(i,t)\sigma(j,t) & \sigma^2(j,t) \end{pmatrix}^{1/2} \begin{pmatrix} v_{i,t} \\ v_{j,t} \end{pmatrix}$$

where $\rho_{i,j}(t) \equiv \rho_{i,j,t/T} \equiv \rho_{i,j,\tau}$ denotes the time-varying correlation of $u_{i,t}$ and $u_{j,t}$. Define $\hat{v}_{j,t} = \hat{u}_{j,t}/\hat{\sigma}(j,t)$. Then, we suggest the following estimator for $\rho_{i,j,\tau}$.

$$\hat{\rho}_{i,j,\tau} = \frac{\sum_{t=1}^{T} K^h_{t,\tau}\hat{v}_{i,t}\hat{v}_{j,t}}{\sum_{t=1}^{T} K^h_{t,\tau}} = \sum_{t=1}^{T} \tilde{K}^h_{t,\tau}\hat{v}_{i,t}\hat{v}_{j,t} \tag{6}$$

We again allow for a different bandwidth for this estimation. We denote this bandwidth by h_ρ. Similarly to β_t we assume that ρ_t is smoothly time-varying (see Assumption 7, in the Appendix).

Then, it is easy, using the arguments of Theorem 1 of Kapetanios (2007), to show that $\hat{\rho}_{i,j,\tau}$ estimates consistently $\rho_{i,j,\tau}$.

2.3 *Choice of the bandwidth,* h

When a researcher estimates a 'rolling regression', a decision has to be made about the size of the window. The analogue of the window here is the bandwidth h which governs the variance of the kernel. Roughly speaking, the optimal bandwidth involves trading off efficiency versus bias. A large h uses more information from data points further away from the period of interest, and is therefore more efficient. But in an economy with more structural change, information a long way from the current observation of interest is of less value, so a large h generates bias. Thus the ideal bandwidth depends on knowing the evolution of the β's in advance when this is precisely what we are trying to uncover. The procedure for choosing the bandwidth thus involves proposing a value for h, obtaining the β's for this value of h, using this to evaluate an appropriate criterion function and repeating this process until the criterion function is minimised.

The next step, of course, is to decide on a criterion function to use in this optimisation. For h_β we suggest using a 'leave one out' penalised residual sum of squares objective function. We use the 'leave one out' approach because, in the absence of a penalty, if for some observation t of interest we left that observation 'in' we could fit that observation perfectly by putting no weight on neighbouring observations at all. This would not be informative. The formal minimisation used to determine h_β, is given in the Appendix. The penalty term used, is of the generalised cross-validation form as discussed in Orbe *et al* (2005). Penalty terms in the context of choosing h, have a direct analogy to their use for model selection, in the form of information criteria, in the sense that they penalise the use of too many parameters. Usually, when a penalty term is used there is less need for a 'leave one out' approach as the trivial solution of $h = 0$ is not optimal due to the penalty. However, initial empirical work suggests that, in our case, implementing the 'leave one out' approach produces smoother results that are more readily interpretable.[2] Just as for h_β, h_σ and h_ρ can be determined respectively, for some u_t, by similar minimisations that are formally stated in the Appendix.

Our decision to estimate the VAR equation by equation allows us to use non-parametric

[2]Note further that leaving one observation out as in **(C-1)** is asymptotically negligible in the sense that, by **(C-3)**, $Q(\tilde{\beta}) - Q(\hat{\beta}) = o_p(1)$.

estimation with different bandwidths across equations. An alternative would have been to estimate the VAR as a system by imposing a fixed bandwidth for all equations in the VAR. The advantage of our choice is that we allow the data to determine the informativeness of neighbouring observations about the coefficient at a given point in time. The disadvantage, of course, is that we lose the efficiency gains that come from estimation of our VAR model as a system. In principle, it would be feasible to devise a procedure that allowed for system estimation and bandwidths that varied across equations and therefore get the best of both worlds, but we leave that for future work.

2.4 Testing structural stability

We wish to test the hypothesis that $\beta(t)$, $\sigma^2(t)$ and $\rho(t)$ do not change over time. The null hypotheses are

$$H_{0,\beta} : \beta_\tau = \beta, \forall\tau \quad H_{0,\sigma} : \sigma_\tau^2 = \sigma^2, \forall\tau \quad H_{0,\rho} : \rho_\tau = \rho, \forall\tau \tag{7}$$

against the alternative hypotheses that β_τ, $\sigma^2(t)$ and $\rho(t)$ are non-constant and satisfy Assumptions 3 and 7. To this we use the tests proposed by Kapetanios (2008). We focus on β for simplicity. Similar analyses are implied for $\sigma^2(t)$ and $\rho(t)$. Note that Kapetanios (2008) only discusses the test for β. Therefore, our work on the stability testing of variances and correlations extends the results of that paper.

Let us denote the estimate of β, under the null as $\tilde{\beta}$. Depending on the assumptions made about u_t, standard methods can be used to estimate β under the null. One test proposed by Kapetanios (2008) looks at the difference between the estimates of β_t under the two hypotheses. The starting point for developing the test focuses on a fixed τ. So, the test is based on a statistic which takes the form

$$T^\tau = (\hat{\beta}_\tau - \tilde{\beta})'\hat{V}(\hat{\beta}_\tau - \tilde{\beta})^{-1}(\hat{\beta}_\tau - \tilde{\beta}) \tag{8}$$

This statistic focuses on all coefficients of one equation jointly. Of course, statistics that focus on individual coefficients are possible. The above statistic can be used to test H_0. But to do this, we need to jointly consider many points in the interval (0, 1) where τ is defined. To conduct this test we need to use summary statistics for a set of pointwise test statistics. The problem has parallels with the problem of testing when a nuisance parameter is unidentified under the null hypothesis. This problem arises in many areas in econometrics such as linearity testing, tests for structural breaks and others (for more details see Davies (1977) and Andrews and Ploberger (1994)). Let the set of points for which test statistics are available be denoted by $\mathcal{T}_m = \{\tau_1, \tau_2, \ldots, \tau_m\}$, where $\tau_1 < \tau_2 < \ldots < \tau_m$. Three summary statistics are usually considered. These are given by

$$T_{AVE} = \frac{1}{m}\sum_{j=1}^{m} T^{\tau_j} \tag{9}$$

$$T_{SUP} = sup_j \ T^{\tau_j} \tag{10}$$

$$T_{EXP} = \frac{1}{m}\sum_{j=1}^{m} e^{\frac{T^{\tau_j}}{2}} \tag{11}$$

There is no strong evidence suggesting that one test is more useful than another, so we present results based on all three. Kapetanios (2008) provides asymptotic analysis for the properties of tests based on the above test statistics. Given the slow rate of convergence related to non-parametric asymptotics, it is not surprising to note that asymptotic results may not provide good approximations to small sample behaviour. This is the case for these asymptotic tests. Note that the bad performance of tests based on non-parametric asymptotic results is documented in the literature. In particular, Fan (1995, 1998) show that asymptotic tests have rejection probabilities that deviate significantly from the nominal significance level.

A solution for this, suggested by Kapetanios (2008), is the bootstrap. For the bootstrap test there is no need for normalising the test statistic by the variance term $\hat{V}(\hat{\beta}_\tau - \tilde{\beta}_\tau)$. Bootstrap theory suggests that normalising is advantageous, as it results in a pivotal test statistic, but only if the variance can be well estimated. Since this is not, necessarily, the case for the variances of non-parametric estimators we choose not to normalise and we, therefore, use the following modified test statistic.

$$\tilde{T}^\tau = (\hat{\beta}_\tau - \tilde{\beta})'(\hat{\beta}_\tau - \tilde{\beta}) \tag{12}$$

Below we give the bootstrap algorithm suggested in that paper.

Algorithm 1 1. Estimate $\hat{\beta}_\tau$ using **(3)** for all points in $\mathcal{T}_m$. Estimate $\tilde{\beta}$ using OLS if appropriate. Obtain OLS residuals. Denote the set of OLS residuals by $\{\hat{u}_t\}_{t=1}^T$

2. Generate a bootstrap sample for u_t, denoted u_t^*, by resampling with replacement from $\{\hat{u}_t\}_{t=1}^T$ to obtain $\{u_t^*\}_{t=1}^T$

3. Generate a bootstrap sample for y_t, denoted y_t^* by

$$y_t^* = \tilde{\beta}' x_t + u_t^*, \quad t = 1, \ldots, T \tag{13}$$

4. Construct bootstrap version of $\tilde{T}^\tau$ and $\tilde{T}_j$, $j = SUP, AVE, EXP$, denoted $\tilde{T}^{*,\tau}$ and $\tilde{T}_j^*$, $j = SUP, AVE, EXP$.

5. Repeat steps 2-4, B times to obtain the empirical distribution of $\tilde{T}^\tau$ and $\tilde{T}_j$, $j = SUP, AVE, EXP$.

One final point worth noting here is the contrast between the approach to testing for parameter instability here, and that adopted in, for example, Cogley and Sargent (2010). Those authors, having set out a STVC parameter model, compute the posterior probability of the joint event that (for example) inflation persistence was high and has fallen. In our framework, there is no probabilistic dependence of future inflation persistence on today's inflation persistence, hence the need to compute joint densities to assess parameter constancy does not arise.

2.5 A Monte Carlo study to motivate the use of deterministic models

In this section we carry out a Monte Carlo study in order to explore the relative performance of two different estimators of the underlying process of the time-varying coefficient in a dynamic autoregressive model. The model we entertain is a benchmark time-varying AR(1) model given by

$$y_t = \beta_t y_{t-1} + \epsilon_t. \tag{14}$$

This model is the basic building block for much more complicated dynamic models that have been used in the macroeconometric literature to model smooth structural change. These models include time-varying VAR models for the joint modelling of inflation and GDP growth, among other variables, and models that allow for unconditional heteroscedasticity and well as change in the dynamics of the conditional mean. However, focusing on this simple model provides a parsimonious way to determine the relative merits of alternative estimators of β_t. An important question relates to the nature of β_t. The macroeconometric literature assumes uniformly some parametric stochastic model for β_t, which is usually of the form

$$\beta_t = \beta_{t-1} + u_t, \tag{15}$$

ie it is assumed that the process follows a random walk.

An alternative way to model β_t is to assume, as we do in the main body of this paper, that it is a smooth deterministic process and estimate it using the technology discussed in the previous subsections. In our Monte Carlo we wish to examine which approach is better in relative terms. This is of course a difficult undertaking given the radically different philosophies underlying the generation of β_t. We choose to focus on a simple evaluation strategy. We generate β_t with both approaches (stochastic and deterministic). We then estimate $\hat{\beta}_t$ using both estimation procedures. We then examine which approach has a smaller MSE defined as $\frac{1}{T}\sum_{t-1}^{T}(\hat{\beta}_t - \beta_t)^2$, over Monte Carlo replications.

Specifying the process under the stochastic paradigm is easy. We simply adopt the random walk paradigm and generate a sequence $\tilde{\beta}_t$ from **(15)** using $u_t \sim n.i.i.d.(0, 1)$. We impose the usual restriction by setting $\beta_t = \tilde{\beta}_t / \max_t(|\tilde{\beta}_t|)$ which ensures that $-1 < \beta_t < 1$ for all t.

Specifying β_t under the deterministic paradigm is of course more difficult. We have little guidance as to the shape of β_t which of course will be radically different for different empirical settings. We choose to consider three different functions which provide a reasonable coverage of some simple shapes for β_t over time. Letting $\beta_t = \beta_{t/T}$ and $t/T = \tau$ we set $\beta_\tau = f_i(\tau)$, $i = 1, 2, 3$ where

$$f_1(\tau) = -0.5 + \tau, \quad \text{(linear)},$$

$$f_2(\tau) = \frac{0.9}{1 + e^{-10+20\tau}}, \quad \text{(logistic)},$$

$$f_3(\tau) = 0.9\sin(\pi\tau), \quad \text{(sine)}$$

and $\tau \in [0, 1]$. For the non-parametric estimation we use the data dependent method to determine the bandwidth discussed in the previous section. We set $\epsilon_t \sim n.i.i.d.(0, 1)$, $T = 50, 100, 200, 400$ and carry out 1,000 Monte Carlo replications for each experiment.

We report two summary statistics of the mean squared errors (MSEs) over the Monte Carlo replications. The first is the ratio of the average MSE of the non-parametric method over the average MSE of the method based on the state-space model. This is slightly problematic since there is potential for large values of the MSE to skew the average MSE and produce spurious results. So we also use, as a summary statistic, the median of the relative MSE of the two methods (the non-parametric method is again in the numerator) over the Monte Carlo replications.

We report results in Tables A and B. They make for interesting reading. Focusing first on the experiments where the truth is that β_t is a deterministic function of time we see that both summary statistics clearly indicate that the non-parametric method of estimating β_t is preferable. In only one case do we find that the state-space approach is preferred by both summary statistics and that relative superiority is marginal. On the other hand the non-parametric method dominates very clearly in most other cases. It is interesting to note that the relative performance is a function of the true shape of β_t as one would expect. The result though that the non-parametric method is best across all three designs is encouraging and allows us to conclude that our results may have some generality. In one sense, though, this is not surprising since the non-parametric method is designed to handle this kind of true process.

As a result we next examine what happens when the parametric random walk paradigm underlies the true data. Here, we are faced with a very surprising result: the non-parametric method dominates the state-space method comprehensively. This result clearly indicates that even if the data are generated from a model where the parameters follow a stochastic persistent process, using the non-parametric method of estimating β_t is preferable. This benefit tails off the greater the sample size, but for the case studied in this paper (40 years of quarterly data or less) the benefit is quite pronounced. In ongoing work we are studying the theoretical foundations of this result, but for the purposes of this paper, we leave the reader with this small piece of evidence to motivate the use of DTVC models.

Table A: Results on relative average MSE

True Model/T	50	100	200	400
f_1	0.580	0.472	0.799	1.076
f_2	0.211	0.461	0.606	0.755
f_3	1.552	0.539	0.104	0.141
Stochastic	0.525	0.286	0.477	0.612

Table B: Results on median of relative MSEs

True Model/T	50	100	200	400
f_1	0.913	0.731	0.976	1.128
f_2	0.824	0.803	0.911	0.778
f_3	0.548	0.295	0.262	0.352
Stochastic	0.804	0.866	0.898	0.940

2.6 Instantaneous variances, spectra, predictability measures and our counterfactual experiments

We describe changing inflation dynamics in terms of shifts in instantaneous variances, spectra and predictability measures. The discussion focuses on the more complex VAR model but applies obviously to the AR model as well. To calculate the variances, we consider the coefficient and error variance values at each point in time t and obtain the variable variances implied by a VAR(1) with these parameter values. To fix ideas we note that given a VAR(1) coefficient and error variance matrix estimates given by $\hat{B}_t$ and $\hat{\Sigma}_t$ respectively at time t, the vectorisation of the implied variance of a vector variable, y_t, following such a VAR(1) is given by $(I - \hat{B}_t \otimes \hat{B}_t)^{-1} vec(\hat{\Sigma}_t)$.

Similarly, the implied cross-spectrum for y_t at time t and frequency ω is given by $2\pi^{-1} \left(I - \hat{B}_t e^{i\omega}\right)^{-1} \Sigma_t \left(I - \hat{B}_t' e^{i\omega}\right)^{-1}$. The elements of the diagonal of this matrix give the respective spectra for each element of y_t.

Finally, the third means of describing changing inflation dynamics is the use of predictability measures which, as discussed in Cogley and Sargent (2010), can provide a reasonable measure of persistence. We follow Cogley and Sargent (2010) and define predictability for the i-th variable of y_t, at horizon j and time t, as

$$1 - \frac{e_i' \sum_{h=1}^{j-1} \left(\hat{B}_t^h\right) \Sigma_t \left(\hat{B}_t^h\right)' e_i}{e_i' \sum_{h=1}^{\infty} \left(\hat{B}_t^h\right) \Sigma_t \left(\hat{B}_t^h\right)' e_i}$$

where e_i is a vector whose i-th element is 1 and the rest zero.

We consider two scenarios for the calculation of the variances that allow us to measure the relative contribution of variation in the error variances to the variability of the data over time. In the top row of Charts 3, 7 and 11 we report the implied variances allowing both $\hat{B}_t$ and $\hat{\Sigma}_t$ to vary freely. In the second row we report variances where we have fixed $\hat{\Sigma}_t = \hat{\Sigma}$ where

$$\hat{\Sigma} = 1/T \sum_{t=1}^{T} (y_t - B_t y_{t-1})'(y_t - B_t y_{t-1})$$

In words, we fix the disturbances to the VAR equations at their sample average values. This counterfactual experiment allows us to disentangle the effect of changing coefficients and changing error variances on the variance of y_t. Note that similar results for the AR model are reported in the fourth and fifth panel of Charts 1,7 and 13.

3 Structural analysis using sign restrictions

The final objective of our paper is to document changes in the response of inflation and real activity to identified monetary policy shocks. We do this so that we can use any changes we see to infer something about the evolution of real and nominal rigidities in the economy. This section explains our identification strategy. We identify monetary policy shocks in our VAR following, eg, Canova and De Nicolo (2002). We implement the sign-restriction method a little differently than others have done and describe our point of departure below.

The expository analysis is, for simplicity, based on the VAR(1) model given by

$$y_t = A y_{t-1} + u_t$$

The aim is to factorise the covariance matrix of the n-dimensional reduced-form error, denoted by Σ, as

$$\Sigma = BB'$$

where B is usually either given by a Cholesky factorisation or, in our case, an eigenvector-eigenvalue one of the form $B = PDP'$ where P is a matrix whose columns contain the eigenvectors of B and D is a diagonal matrix containing the eigenvalues of B. Of course, there exist multiple such factorisations since for any non-singular orthogonal matrix Q, we have

$$\Sigma = BQQ'B'$$

Traditionally, $n(n-1)/2$ restrictions are sufficient to fully specify a unique Q. A number of schemes deriving from insights from theoretical models have been proposed to specify the $n(n-1)/2$ restrictions. Recently, alternative strategies have been proposed in which Q is specified so as to imply particular signs for the structural impulse responses (IR) given by the sequence. $\{ABQ, A^2BQ, A^3BQ, ...\}$. There is no closed form solution for this problem. Further, it is not known whether there exists a suitable Q such that p sign restrictions can be imposed on the IR sequence when $p \geq n(n-1)/2$. Further, if a suitable Q exists it is not known whether it is unique or indeed how many Q exist such that each of them satisfies the restrictions. Of course if $p \leq n(n-1)/2$ the existence of Q is guaranteed, since there is always a Q that satisfies not just $n(n-1)/2$ sign but $n(n-1)/2$ exact restrictions. The problem is how to determine Q. A popular solution is to parameterise Q in terms of Givens rotations. A Givens rotation is given by

$$I^n_{pq}(\theta) = \begin{pmatrix} 1 & \ldots & 0 & \ldots & 0 & \ldots & 0 \\ \ldots & \ddots & \ldots & \ldots & \ldots & \ldots & \ldots \\ 0 & \ldots & cos(\theta) & \ldots & -sin(\theta) & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & \ldots & sin(\theta) & \ldots & cos(\theta) & \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots & \ldots & \ldots & \ldots \\ 0 & \ldots & 0 & \ldots & 0 & \ldots & 1 \end{pmatrix}, 0 \leq \theta \leq \pi/2$$

where $p < q$ denote the positions on the diagonal taken by $cos(\theta)$. An obvious property of $I^n_{pq}(\theta)$ is that $I^n_{pq}(\theta)I^{n\prime}_{pq}(\theta) = I$. Then, Q is parameterised as

$$Q(\theta) = \prod_{i=1}^{n-1} \prod_{j=i+1}^{n} I^n_{ij}(\theta_{ij}) \quad \textbf{(16)}$$

where $0 \leq \theta_{ij} \leq \pi/2$ and θ is the $n(n-1)/2$ vector containing all the scalar θ_{ij}. It is obvious that

$$Q(\theta)Q(\theta)' = I.$$

Loosely, researchers using sign restrictions follow variants of the following algorithm: 1) choose some θ_{ij}, a rotation of the Q matrix. Check whether the series of shocks it implies generate impulse responses that satisfy the sign restrictions at the specified horizons; 2) if it does, keep it, if not, discard it; 3) go back to step 1. When the search is complete, researchers report statistics about the distribution of impulse responses to the many shocks that they have found, the median, various percentiles, etc. This is exactly what we do. Before we implement this procedure, however, we make a comment.

Using this algorithm invokes an implicit claim made is that $Q(\theta)$ spans the space of n-dimensional orthonormal matrices. In other words for any n-dimensional orthonormal matrix Q, there exists θ such that $Q = Q(\theta)$. Given that n-dimensional orthonormal matrices have $n(n-1)/2$ 'free' parameters and given that θ contains $n(n-1)/2$ elements, it is implicitly conjectured that the above 'spanning' indeed occurs. However, there is no proof of this claim to the best of our knowledge. Since, determining a value for θ such that the signs restrictions hold cannot be done analytically, a grid search over a subset of $\{\{\theta_{11} | 0 \leq \theta_{11} \leq \pi/2\}, ..., \{\theta_{nn-1} | 0 \leq \theta_{nn-1} \leq \pi/2\}\}$ is usually undertaken and grid points for which the sign restrictions hold are retained for further analysis. Our experience with this grid search is that for some IR sequences no solution may be found. We have considered up to 2 million points without finding solutions. This casts further doubt on the possibility that $Q(\theta)$ spans the space of n-dimensional orthonormal matrices, even though an alternative possibility is that there exists no Q such that p sign restrictions hold if $p \geq n(n-1)/2$. Nevertheless, to explore our conjecture that the 'spanning' may not hold we extend the parameterisation in **(16)** to the following

$$Q(\theta^{(1)}, ..., \theta^{(m)}) = \prod_{s=1}^{m} Q(\theta^{(s)}) = \prod_{s=1}^{m} \prod_{i=1}^{n-1} \prod_{j=i+1}^{n} I_{ij}^{n}(\theta_{ij}^{(s)}) \qquad \textbf{(17)}$$

This is obviously a richer parameterisation and indeed considerable experimentation with grid searching this parameterisation suggests two things. Firstly, in a number of cases, if grid searching over $Q(\theta)$ yields a suitable Q, then grid searching over $Q(\theta^{(1)}, ..., \theta^{(m)})$ yields a larger number of suitable Q, for all $m > 1$. Secondly, there are a number of cases when, if $p \geq n(n-1)/2$, grid searching cannot produce a suitable Q, grid searching over $Q(\theta^{(1)}, ..., \theta^{(m)})$ does produce one or more. In our empirical work we set $m = 4$. We found no evidence that higher values of m further improve the outcome of the grid search (in the sense of generating more Q's that satisfy the restrictions. This modification of the algorithm used by others may be useful in other applications that use sign restrictions to identify shocks.

In our data, we find that there are a number of instances where the sign restrictions cannot be imposed for horizons larger than 1. This as we noted before can occur for three reasons: i) the grid search cannot explore effectively the parameterisation $Q(\theta^{(1)}, ..., \theta^{(m)})$. ii) $Q(\theta^{(1)}, ..., \theta^{(m)})$ cannot span the space of n-dimensional orthonormal matrices, iii) no suitable n-dimensional orthonormal matrix satisfying the restrictions exists. We believe that i) is unlikely when the grid search explores 2 million grid points. Therefore, our strategy is to begin by attempting the imposition of the sign restrictions for s^{max} horizons, successively reducing the number of horizons until we can obtain a suitable Q. In a number of instances the number of horizons has to be reduced to 1 before a suitable Q is found. We set $s^{max} = 5$. We start with a higher number of horizons than we finish on the grounds that we can be more confident that the shock has some economic content if it satisfies the restriction for more horizons. We reduce the number of horizons we demand the impulse response to satisfy successively on pragmatic grounds: we would prefer more restrictions rather than less, but responses that satsify few restrictions have some economic meaning, and we take those rather than giving up.

To recap on where our procedure overlaps with and where it departs from that adopted previously in the sign-restrictions literature. Like others, we collect together all shock sequences which generate impulse responses that satisfy the restrictions we invoke from theory. In our case we found that the ways of parameterising the space of shocks that others had used generated very few, sometimes no satisfactory impulse responses. We note that the parameterisations adopted thus far do not span the space of all possible shocks. We adopt a slightly richer parameterisation of that space of shocks and find that it generates more impulse responses that satisfy our restrictions.

4 Documenting and accounting for the evolving volatility and persistence of inflation

In this section we describe our results using the reduced-form time-series models. These reduced-form models are used to document changes in inflation volatility, persistence and predictability; and to determine the extent to which the reduction in inflation volatility over our sample period is due to changes in volatilities or changes in propagation parameters.

4.1 Data and the model

We use the toolkit we have discussed in the previous sections to address this issue. A focal time series for the phenomenon of the 'Great Moderation' is inflation. We therefore carry out both a univariate and multivariate analysis of inflation. The univariate model we consider is an AR(1) with time-varying constant, AR coefficient and error variance. The multivariate model we focus on is a trivariate VAR(1) model of inflation, real activity and interest rates, where each equation has time-varying constants, parameters, and innovation variances.

We estimate on de-meaned data. The time-varying constants in our model echo Cogley and Sargent (2010), who posit a stochastic process for the mean of inflation. Here variations in the mean are deterministic, but this device allows the same decomposition between long and short-run components of inflation. As in many other literatures, those scrutinising inflation dynamics have wanted to be sure not to confuse a rise and fall in the inflation target during the 1970s and 1980s with persistent dynamics around an unchanged mean.

We use three different data sets. For the United States, our first data set comprises the Fed Funds rate, quarterly GDP growth, and quarterly growth of the GDP deflator. Our second data set equals the first but with the CPI as the measure of inflation. Here we are motivated by previous work, for example Benati (2007), which has pointed out that conclusions we draw about both the level and stability of inflation persistence can depend on which US price index is used to measure inflation. The sample period is 1955 Q1-2007 Q1.

The third data set is for the United Kingdom. We use GDP growth, the GDP deflator and 'Bank Rate', the instrument of monetary policy. The sample period is 1975 Q1-2007 Q1. In all models we include only one lag of relevant variables. This was shown to generate residuals free of serial correlation in most of our models.

All our results are set out in the Appendix. The kernel used is the normal density kernel. Our tests for the constancy of the objects in our time-series models are reported in Tables 1, 2 and 3. Charts 1-10 present pictorially the output of our analysis. Charts 1-5 relate to the US data set using GDP deflator as the inflation measure while Charts 6-10 focus on the UK data set.

4.2 Preliminary tests for misspecification

In Tables 1-3 we include the results of preliminary tests for the validity of our model of choice in both the AR and VAR cases. Probability values are reported for all tests. The misspecification tests we use are applied on the standardised residuals: ie the residuals once they have been normalised by the time-varying estimated error standard deviation. We report results on three tests: The LM test for serial correlation using four lags, the LM test for ARCH using again four lags; and the Lee, White and Granger (1993) test for dynamic neglected non-linearity using one lag. The misspecification test results are encouraging. There is little evidence of serial correlation or dynamic neglected non-linearity in any of the data sets we consider, implying that our choice of one lag for the AR and VAR models is appropriate. There is slightly more evidence for variation in the conditional variance. As we do not model this aspect of inflation (but only variation in the unconditional variance) this is possibly to be expected. Further, the evidence is by no means very strong as it is concentrated mainly on the interest rate equation in all three data sets we consider. Overall, we feel that given the above evidence our chosen model appears to be reasonably well specified. We next discuss the rest of the empirical results we have obtained.

4.3 Results: United States

First we look at the results of the AR model, reported in Chart 1. From the top-left and top-centre panels we observe that despite evident and - on casual inspection - significant variations in the constant governing inflation, inflation persistence, measured by the coefficient on lagged inflation, rose sharply from around -0.1 at the start of the sample to about 0.7 in 1970. Since then, persistence seems to have fallen back to around 0.1, although at this point the confidence intervals for the parameters are relatively wide.

The top-right panel shows clear changes in the estimated innovation variance. While inflation persistence is rising through to 1965 or so, the estimated innovation variance is falling: it rises thereafter to peak in the mid-70s. Note how much more pronounced is the subsequent fall in the innovation variance from the peak compared to the fall in inflation persistence from its peak.

The bottom-left and bottom-centre panels of Chart 1 report our counterfactual exercise that attempts to disentangle the effect of changing innovation variances from changing propagation parameters on the instantaneous variance of inflation. The bottom-left panel computes the variance of inflation holding the innovation variances fixed at the sample average. The bottom-centre panel repeats this exercise using our estimated time-varying innovation variances.

Allowing time variation in the variances generates a peak in inflation volatility that is about 30% higher; it generates a more protracted period of higher volatility, beginning in in the early 60s and ending around 1982; finally, with time-varying variances we see that the volatility of inflation is distinctly lower post-1990 than it was pre-1970, in contrast to the fixed-variance case.

The final - bottom-right - panel of Chart 1 plots the instantaneous spectrum implied by our time-varying coefficients model. These spectra reveal the increase in variances at short to medium horizons that took place during the Great Inflation.

We now move on to the VAR model, and compare the story this model tells about the Great Moderation to the story told by the univariate model. The left-hand column of panels in Chart 2 show that the estimated innovations variances have a pronounced peak around 1980, similar to what we found with the univariate model. The multivariate model suggests that these variances were greater at the start of the sample than at the end, in contrast to the univariate model. The right hand panel shows pronounced movements in the correlation of one equation's innovations with another, but we found that these changes are not significant in explaining the changes in the estimated instantaneous variance of inflation reported in Chart 3.

Chart 3 reports the multivariate version of our counterfactual experiments. The multivariate model shows an even more marked contrast between the implied variances of inflation when we allow for time-varying innovations variances compared to when these are fixed. This can be seen from comparing the top-left (variances varying) to the middle-left panel (variances fixed). In the top-left, we trace out two pronounced peaks for the variance of inflation around 1970 and 1980. In the bottom-left panel, we also see two peaks, but these are less than half the size of the peaks when variances are allowed to be time-varying. The numerical contribution of changing innovation variances seems to be even greater for the variance of the growth of GDP. By comparing the central panels, we see that the fixed-variances (bottom-centre) fails to generate the peak in the variance of output we see at around 1980 in the time-varying variance case (top-centre). We do not comment on the bottom panel here as this reports result for instantaneous variances using the structural VAR model discussed in the next section.

In the top panel of Chart 4, we plot a time series of the instantaneous spectrum of inflation associated with the multivariate VAR. This illustrates how the variance of components at many frequencies rose during the Great Inflation.

The bottom panel shows how inflation predictability has changed at all horizons. Predictability is our multivariate counterpart to the univariate concept of inflation persistence. Inflation predictability has clearly fluctuated a good deal up to (roughly) the 21 quarter horizon. It is apparent from the chart that the 1970s was a time when predictability was a little higher at longer horizons. But the striking contrast between the plot of the spectra and the plot of predictability is an informal indication that it may be changes in variances that mark out the Great Inflation from earlier and later periods, rather than changes in parameters: the predictability chart isolates the contribution of changes in parameters. The spectra include the contribution of changes in the variances of shocks. Finally, Chart 5 reports two cross-sections for the predictability measures (1974 Q1 and 2004 Q1) together with 90% confidence bands obtained via a bootstrap procedure. It is clear that the high predictability measure of the mid-70s lies outside the confidence band of the low predictability measure.

Table 1 reports our formal tests for the stability of the objects in our univariate and multivariate models for US GDP deflator inflation. The univariate tests, reported in the bottom panel of the table, show a fair amount of stability in error variances and parameters. By contrast, the multivariate model picks up a lot more instability. We reject stability for error variances in the GDP growth and interest rate equations, and for parameters in the GDP equation.

4.4 Results: United States, inflation measured by CPI

As we noted earlier, Benati (2007) points out that different conclusions can be drawn about the stability or otherwise of the US inflation process depending on which index we use. Is this manifest when we view inflation through our DTVC model? For the sake of brevity, we do not provide figures for CPI-based results, but we offer a short verbal description of our findings and can supply results on request to those who are interested. In short, we find that there are some differences when we use a univariate model to characterise the inflation process, but that by contrast, using a multivariate model leads to very similar results regardless of the inflation index used. The differences between our univariate models for inflation are: first, our CPI results show a more marked fall in inflation persistence from its high period, in the 70s, to the recent period. Second, the CPI results indicate that inflation persistence is clearly lower at the end than at the beginning of the sample period, something that was much less evident in the GDP deflator results. Third, we compute a larger increase from trough to peak in the estimated innovations variance for the CPI. Fourth, if we look at the time series of the spectra, we see that the peak in the variance in inflation in 1980 is contributed to much more evenly across the different frequencies. Fifth, our CPI counterfactual exercise records the contribution of changes in innovations variances to be greater than we saw for the GDP deflator results. Looking at the formal tests for stability of coefficients and variances, we find that the univariate model for CPI shows much more evidence of instability, of both coefficients and variances, compared to the US GDP deflator model. But evidence for instability in the multivariate model is pretty similar for the two data sets.

4.5 Results: United Kingdom

Our final data set is for the United Kingdom, with inflation measured by the GDP deflator. Chart 6 shows results from the univariate model. We observe numerically a large and what looks - from the confidence intervals in the chart - to be a statistically significant fall in inflation persistence (top-centre panel) at the same time as we capture a steady fall in the mean inflation rate (top-left). We find that innovations variances fall markedly throughout our sample period (top-right). These changes translate into a marked fall in the instantaneous variance of inflation (bottom-centre) Our counterfactual exercises suggest that the dominant cause of this is the fall in innovations variances: notice that when we hold the innovations variances flat at the sample average (bottom-left), the implied instantaneous variance is also relatively flat.

Our multivariate account of the UK inflation process is slightly different. Our counterfactual exercises are reported in Chart 8. For GDP growth, we see once again that without time-varying innovation variances, we completely fail to capture the fall in the implied instantaneous variance. However, the picture is not as dramatic for inflation itself. Here we compute a rise and a fall in the instantaneous variance of inflation (top-left panel). And this profile is visible too when we hold innovations variances fixed (bottom-left), albeit that the peak is about 60% of the height of the peak for the time-varying innovations model.

Chart 9 shows our estimated time series for the spectra and for predictability for the United Kingdom. It is interesting to see that just as in the univariate model we record a fall in inflation persistence (Chart 6, top-centre panel) here we see that predictability has fallen at all horizons in recent years compared with the 1970s and 1980s. Finally, Chart 10 reports two cross-sections for the predictability measures (1982 Q3 and 2004 Q1) together with 90% confidence bands obtained via a bootstrap procedure. It is clear that the high predictability measure of the early 80s lies outside the confidence band of the low predictability measure.

We report tests for the stability of the objects in our univariate and multivariate models in Table 3. Both sets of models show widespread evidence of instability in both propagation parameters and innovations variances. (This contrasts slightly with our results for the US data sets where we tended to find that univariate evidence suggested more stability in the inflation process than did the univariate models.)

5 Documenting the evolving response to identified monetary policy shocks

In this section we describe our findings using the structural VAR. We document how the response of inflation and GDP to monetary policy shocks identified using sign restrictions has changed over time.

We report sequences of impulse responses of inflation, GDP growth and interest rates to a positive monetary shock defined as one which implies negative impulse response for inflation, GDP growth over a horizon of s periods and positive impulse response for interest rates for the same horizon. These restrictions are relatively uncontroversial in that they fit a wide class of monetary, sticky price, DSGE models currently in use. The reduced-form estimates for the VAR parameters are taken from our reduced-form analysis that we described in previous sections.

We report the 5% percentile, median and 95% percentile of all the IR satisfying the sign restrictions as described above for the three data sets we consider in Charts 11-19. We choose a subset of time periods to report the results: For the United States these are 1966 Q1, 1976 Q1, 1986 Q1, 1996 Q1 and 2006 Q1; for the United Kingdom we choose: 1980 Q1, 1985 Q1, 1990 Q1, 1995 Q1, 2000 Q1 and 2005 Q1. We report all these results in Charts 11-19.

We start our commentary by looking at the response of US CPI inflation to the structural monetary policy shock: the starting point is the middle panel of Chart 11 which reports changes in the median response. The basic finding is that inflation is estimated to respond less, and less persistently to a monetary policy shock now than in the 1970s. The persistence seems to be at its highest in the 70s with most of the reduction occurring in the ensuing decade. Thereafter, little change is observed. The same holds for the impulse responses of GDP to the monetary policy shock (Chart 12).

The results change somewhat for the GDP deflator data set for both inflation and GDP. For example, the most persistent response of inflation to the monetary shock is recorded for 2005 Q1; the largest response on impact is recorded for 1976 Q1.

Our results for the United Kingdom are a pretty close echo of those for the US CPI inflation data. In Chart 17, we record the median response of inflation to have been the largest on impact in 1990 Q1 and 1980 Q1, and to have been most persistent in 1980 Q1. The impulse response of GDP to the monetary policy shock is recorded to have been largest and most persistent in 1985 Q1, and substantially more so than the response at the end of the sample.

Taking these findings as a whole, there is suggestive but not overwhelming evidence that the response of inflation and GDP to a monetary policy shock in the United Kingdom and the United States is now both less persistent and smaller than in the 1970s and 1980s. In sticky-price DSGE models, these facts would be consistent with there having been some combination of: (i) a reduction in the degree of price stickiness, which ought to reduce the impact of GDP to a monetary policy shock; or some combination of (ii) a reduction in the degree of price or wage indexation, (iii) a reduction in capital or investment adjustment costs (iv) a reduction in the strength of consumer 'habits', all of which would have the effect of reducing the persistence with which inflation and GDP respond to a monetary policy shock.

6 Conclusion

We have deployed a model of deterministic structural change to describe changes in the inflation process in the United Kingdom and the United States. Our intention was that these results complement the growing body of work that has attempted to make concrete accounts of the temporary increase in the volatility and persistence of inflation. We characterised that literature as using 'sample-splitting' models on the one hand, and stochastic, time-varying parameter models on the other.

Our framework is related to the 'sample-splitting' approach in that it deliberately abstracts from attempting to model the process of structural change. But it departs from it by insisting that the structural change occurs gradually, and by therefore allowing adjacent data points to be informative about the regime in any particular period. Our DTVC model mirrors the assumption in stochastic time-varying parameter models that there are as many values for coefficients and volatilities as there are data points. But it relaxes the restriction these models embody that there is an invariant process governing the evolution of these coefficients and volatilities.

With no clear theoretical grounds for choosing one framework over the other, our DTVC model is therefore offered as a useful complementary tool in documenting changes in the inflation process. We noted, however, that Monte Carlo analysis suggested that it may well pay to use DTVC models even when the true data was generated by an STVC model.

We applied the DTVC model to a VAR model for inflation, and report frequentist estimates. However, the DTVC does not entail these choices. Future work could apply DTVC to a structural model, performing an exercise analogous to that of Fernandez-Villaverde and Rubio-Ramirez (2007). Equally, our estimates could be used as inputs to Bayesian posteriors.

A summary of our findings is as follows: for the United States, we find that there have been significant changes in propagation parameters in the inflation process; that inflation persistence/predictability rose and then fell over our 1955-2007 sample period. This finding accords with Cogley and Sargent (2010). This said, we also find a preponderance of evidence that the dominant cause of the change in instantaneous inflation volatility is changes in the variance of innovations. This echoes what Benati and Mumtaz (2007) found. Like Cogley and Sargent (2010), we find that multivariate models of inflation offer sharper evidence of structural change than do univariate models.

For the United Kingodm, we also find clear evidence of changes in the inflation process. Inflation was less volatile and less persistent at the end of our sample period (2007 Q1) than in 1974, when our sample starts. Multivariate models do not seem to sharpen this conclusion relative to the univariate model, as they do for the United States. However, the univariate and multivariate accounts of the causes of this instability in inflation volatility are different. The univariate model suggests the dominant cause of the decline in inflation volatility is changes in innovations variances. The multivariate model suggests that both changes in innovations variances and propagation parameters contributed. Benati (2007) also finds that changes in innovations variances are the dominant cause of the Great Moderation in the United Kingdom, using an STVC model.

We also document changes in the response of inflation and GDP to monetary policy shocks identified from the reduced-form VAR using sign restrictions. The broad thrust of our evidence on this score is that the responses of inflation and GDP to a monetary policy shock at the end of the sample are smaller and less persistent than those at the beginning of the sample. This is consistent with some combination of: more flexible prices; less price or wage indexation; a reduction in the costs of adjusting capital and/or investment; a weakening of consumer 'habits'. It is also consistent with these models being inadequate characterisations of the data; not that 'habits' are changing, but that there were no habits in the first place, and that some other deeper model that develops evolving persistence is required.

References

Andrews, D W K and Ploberger, W (1994), 'Optimal tests when a nuisance parameter is present only under the alternative', *Econometrica*, Vol. 62, pages 1,383–414.

Benati, L (2006), 'UK monetary regimes and macroeconomic stylised facts', *Bank of England Working Paper No 290*.

Benati, L (2007), 'The great moderation in the United Kingdom', *ECB Working Paper No 769*.

Benati, L and Mumtaz, H (2007), 'US evolving macroeconomic dynamics: a structural investigation', *ECB Working Paper No 746*.

Brown, R L, Durbin, J and Evans, J M (1974), 'Techniques for testing the constancy of regression relationships over time', *Journal of the Royal Statical Association, Series A*, Vol. 138, pages 149–63.

Canova, F and De Nicolo, G (2002), 'Monetary disturbances matter for business fluctuations in the G7', *Journal of Monetary Economics*, Vol. 49, pages 1,131–59.

Chow, A (1960), 'Tests of equality between sets of coefficients in two linear regressions', *Econometrica*, Vol. 28, pages 591–605.

Cogley, T and Sargent, T (2001), 'Evolving post-WWII US inflation dynamics', in Bernanke, B and Rogoff, K (eds), *NBER Macroeconomics Annuals 2001*.

Cogley, T and Sargent, T (2005), 'Drifts and volatilities: monetary policies and outcomes in the post WWII US', *Review of Economic Dynamics*, Vol. 8, pages 262–302.

Cogley, T and Sargent, T (2010), 'Inflation-gap persistence in the US', *American Economic Journal: Macroeconomics*.

Dahlhaus, R (1996), 'Fitting time series model to nonstationary processes', *Annals of Statistics*, Vol. 25, pages 1–37.

Davies, R B (1977), 'Hypothesis testing when a nuisance parameter is present only under the alternative', *Biometrika*, Vol. 64, No. 2, pages 247–54.

Fan, Y (1995), 'Bootstrapping a consistent nonparametric goodness-of-fit test', *Econometric Reviews*, Vol. 14, pages 367–82.

Fan, Y (1998), 'Goodness-of-fit tests based on kernel density estimators with fixed smoothing parameters', *Econometric Theory*, Vol. 14, pages 604–21.

Fernandez-Villaverde, J and Rubio-Ramirez, J F (2007), 'How structural are structural parameters?', *University of Pennsylvania, mimeo*.

Kapetanios, G (2007), 'Estimating deterministically time-varying variances in regression models', *Economics Letters*, Vol. 97, No. 2, pages 97–104.

Kapetanios, G (2008), 'Tests for deterministic parametric structural change in regression models', *Computational Statistics and Data Analysis*, Vol. 53, pages 534–45.

Kapetanios, G and Tzavalis, E (2010), 'Modeling structural breaks in economic relationships using large shocks', *Journal of Economic Dynamics and Control*, Vol. 34, pages 417–36.

Lee, T W, White, H and Granger, C W J (1993), 'Testing for neglected nonlinearity in time series models : A comparison of neural network methods and alternative tests', *Journal of Econometrics*, Vol. 56, pages 269–90.

Muller, P and Quintana, F A (2004), 'Nonparametric Bayesian data analysis', *Statistical Science*, Vol. 19, No. 1, pages 95–110.

Orbe, S, Ferreira, E and Rodriguez-Poo, J (2005), 'Nonparametric estimation of time varying parameters under shape restrictions', *Journal of Econometrics*, Vol. 126, pages 53–77.

O'Reilly, G and Whelan, K (2004), 'Has euro-area inflation persistence changed over time?', *ECB Working Paper No 335*.

Pagan, A and Ullah, A (2000), *Nonparametric Econometrics*, Cambridge University Press.

Ploberger, W and Kramer, W (1992), 'The CUSUM test with OLS residuals', *Econometrica*, Vol. 60, pages 271–85.

Priestley, M (1965), 'Evolutionary spectra and nonstationary processes', *Journal of Royal Statistical Society, Series B*, Vol. 27, pages 204–37.

Robinson, P M (1989), 'Nonparametric estimation of time varying parameters', in Hackl, P (ed), *Statistics, Analysis and Forecasting of Economic Structural Change*, Springer Berlin.

Robinson, P M (1991), 'Time varying nonlinear regression', in Hackl, P (ed), *Statistics, Analysis and Forecasting of Economic Structural Change*, Springer Berlin.

Sims, C (2001), 'Comments on Sargent and Cogley's 'Evolving US postwar inflation dynamics'', Mimeo, Princeton University.

Sims, C and Zha, T (2006), 'Were there regime switches in US monetary policy?', *American Economic Review*, Vol. 96, No. 1, pages 54–81.

Stock, J and Watson, M (2007), 'Why has inflation become harder to forecast?', *Journal of Money, Credit and Banking*, Vol. 39, pages 3–34.

Tong, H (1990), *Nonlinear time series: A dynamical system approach.*, Oxford University Press.

Tweedie, R L (1975), 'Sufficient conditions for ergodicity and recurrence of Markov chains on a general state space', *Stochastic Processes Appl*, Vol. 3, pages 385–403.

Appendix A: The stationarity of time series impied by STVC models

In this appendix, we explain that the common assumption in the STVC literature that instantaneous VAR coefficients at each point in time are stationary can imply that the dependent variables themselves are stationary.

We consider the model

$$y_t = \beta_t y_{t-1} + \epsilon_t. \quad \textbf{(A-1)}$$

$$\beta_t = \beta_{t-1} + u_t, \quad \textbf{(A-2)}$$

discussed in the main text. However, we wish to impose the restriction that $|\beta_t| < 1$. We can do this in a variety of ways and without loss of generality we assume that this is done by imposing the following model on β_t.

$$\beta_t = \beta_{t-1} + I(|\beta_{t-1} + u_t| < 1)u_t. \quad \textbf{(A-3)}$$

and $|\beta_0| < 1$. Geometric ergodicity for y_t can be easily obtained using the drift condition of Tweedie (1975). This condition states that a process is geometrically ergodic under regularity conditions satisfied by assuming disturbances with positive density everywhere, if the process tends towards the centre of its state space at each point in time. More specifically, an irreducible aperiodic Markov chain, y_t, is geometrically ergodic if there exist constants $\delta < 1$, $B, L < \infty$, and a small set C such that

$$E\left[\|y_t\| \mid y_{t-1} = \vartheta\right] < \delta \|\vartheta\| + L, \ \ \forall \vartheta \notin C, \quad \textbf{(A-4)}$$

$$E\left[\|y_t\| \mid y_{t-1} = \vartheta\right] \leq B, \ \ \forall \vartheta \in C, \quad \textbf{(A-5)}$$

where $\|\cdot\|$ is the Euclidean norm. The concept of the small set is the equivalent of a discrete Markov chain state in a continuous context. Small sets are compact.

To prove geometric ergodicity of y_t in this particular instance, we make use of the drift condition in **(A-4)**-**(A-5)**. By **(A-3)** $|\beta_t| < 1$. Then,

$$E\left[|y_t| \mid y_{t-1} = \vartheta\right] = |\beta_t y_{t-1}| \leq |\beta_t| \, |y_{t-1}| \leq \delta |\vartheta|$$

for some $\delta < 1, \forall \vartheta$. Thus, the drift condition holds with a small set follows immediately by geometric ergodicity.

Appendix B: Assumptions underlying the econometric estimator

In this appendix we formally state the assumptions underlying the estimator. These are informally discussed in the main paper.

Assumption 3 $\beta_i(t) = \beta_{i,t/T}$ where each element of $\beta_{i,\tau}$, $\beta_{i,j,\tau}$, $j = 1, \ldots, k$, $\tau \in (0, 1)$, is continuous and twice differentiable on (0, 1). $\sigma_i(t)^2 = \sigma^2_{i,t/T}$ where $\sigma^2_{i,\tau}$, $\tau \in (0, 1)$, is continuous and twice differentiable on (0, 1).

Assumption 4 x_t is an α-mixing sequence with size $-4/3$ and finite 8-th moments. $E(x_{is}x_{jt}) = m_{ij,s,t} = m_{ij}(s/T, t/T) + O(T^{-1})$ where $m_{ij}(.,.)$ is a twice differentiable function of both its arguments.

Assumption 5 $\upsilon_{i,t}$ is a stationary martingale difference sequence with finite 4-th moments which is independent of x_t at all leads and lags.

Assumption 6 The function $K(.)$ is a second-order kernel with compact support $[-1, 1]$ and absolutely integrable Fourier transform.

Assumption 7 $\rho(t) = \rho_{t/T}$ where ρ_τ, $\tau \in (0, 1)$, is continuous and twice differentiable on (0, 1).

Appendix C: Objective functions for the bandwidths

h_β is determined by minimising numerically

$$Q(\tilde{\beta}) = T^{-1} \sum_{t=1}^{T} \left(y_t - \tilde{\beta}(t)' x_t \right)^2 p_\beta(h_\beta) \tag{C-1}$$

where

$$p_\beta(h_\beta) = \left(1 - \frac{1}{T\sqrt{2\pi}} \sum_{t=1}^{T} x_t' \left(\sum_{i=1}^{T} K_{i,t/T}^{h_\beta} x_i x_i' \right)^{-1} x_t \right)^{-1} \tag{C-2}$$

and

$$\tilde{\beta}(t) = \tilde{\beta}_{t/T} = \left(\sum_{i=1, i \neq t}^{T} K_{i,t/T}^{h_\beta} x_i x_i' \right)^{-1} \left(\sum_{i=1, i \neq t}^{T} K_{i,t/T}^{h_\beta} x_i y_i \right) \tag{C-3}$$

h_β, h_σ and h_ρ can be determined respectively, for some u_t, by minimising numerically

$$T^{-1} \sum_{t=1}^{T} \left(\hat{u}_t^2 - \tilde{\sigma}_{t/T}^2 \right)^2 p(h_\sigma) \tag{C-4}$$

and

$$T^{-1} \sum_{t=1}^{T} \left(\hat{\upsilon}_{i,t} \hat{\upsilon}_{j,t} - \tilde{\rho}_{i,j,t/T} \right)^2 p(h_\rho) \tag{C-5}$$

where, using the Rice criterion (see, eg, Pagan and Ullah (2000)),

$$p(h) = \left(1 - \frac{2}{Th\sqrt{2\pi}} \right)^{-1} \tag{C-6}$$

$$\tilde{\sigma}_{t/T}^2 = \frac{\sum_{i=1, i \neq t}^{T} K_{i,i/T}^{h_\sigma} \hat{u}_i^2}{\sum_{i=1, i \neq t}^{T} K_{i,i/T}^{h_\sigma}} \tag{C-7}$$

and

$$\tilde{\rho}_{i,j,t/T}^2 = \frac{\sum_{s=1, s \neq t}^{T} K_{s,t/T}^{h_\rho} \hat{\upsilon}_{i,s} \hat{\upsilon}_{j,s}}{\sum_{s=1, s \neq t}^{T} K_{s,t/T}^{h_\rho}} \tag{C-8}$$

Appendix D: Stability and misspecification tests

Table 1: Structural Stability and Misspecification Tests for US (GDP Defl.)				
Structural Stability Test Probability Values				
Results for VAR model				
Equation	Parameter	Test Statistics		
		Average	Supremum	Exp. Average
Inflation	Constant	0.141	0.045	0.060
	Inflation	0.075	0.025	0.035
	GDP	0.568	0.492	0.492
	Int. Rate	0.181	0.156	0.161
	Err. Variance	0.136	0.116	0.116
GDP	Constant	0.000	0.000	0.000
	Inflation	0.427	0.236	0.241
	GDP	0.236	0.201	0.206
	Int. Rate	0.000	0.000	0.000
	Err. Variance	0.035	0.000	0.000
Int. Rate	Constant	0.286	0.307	0.307
	Inflation	0.563	0.513	0.513
	GDP	0.322	0.362	0.362
	Int. Rate	0.462	0.347	0.347
	Err. Variance	0.000	0.000	0.000
Correlations	Inflation/GDP	0.000	0.000	0.000
	Inflation/Rate	0.000	0.000	0.000
	GDP/Rate	0.005	0.000	0.000
Results for AR model				
	Constant	0.286	0.085	0.111
	AR Coeff.	0.302	0.156	0.241
	Err. Variance	0.643	0.543	0.543
Misspecification Test Probability Values				
	Residual	Ser. Correlation	ARCH	Non-linearity
	Univariate	0.037	0.083	0.737
	Inflation	0.018	0.218	0.736
	GDP	1.000	0.052	0.411
	Int. Rate	0.461	0.000	0.000

Table 2: Structural Stability and Misspecification Tests for US (CPI)				
Structural Stability Test Probability Values				
Results for VAR model				
Equation	Parameter	Test Statistics		
		Average	Supremum	Exp. Average
Inflation	Constant	0.613	0.553	0.573
	Inflation	0.497	0.322	0.362
	GDP	0.447	0.598	0.598
	Int. Rate	0.442	0.417	0.427
	Err. Variance	0.558	0.497	0.497
GDP	Constant	0.000	0.000	0.000
	Inflation	0.176	0.106	0.106
	GDP	0.231	0.151	0.161
	Int. Rate	0.000	0.000	0.000
	Err. Variance	0.060	0.015	0.015
Int. Rate	Constant	0.608	0.628	0.628
	Inflation	0.754	0.784	0.784
	GDP	0.442	0.508	0.508
	Int. Rate	0.688	0.724	0.724
	Err. Variance	0.000	0.000	0.000
Correlations	Inflation/GDP	0.000	0.000	0.000
	Inflation/Rate	0.000	0.000	0.000
	GDP/Rate	0.000	0.000	0.000
Results for AR model				
	Constant	0.010	0.005	0.005
	AR Coeff.	0.005	0.000	0.005
	Err. Variance	0.090	0.040	0.040
Misspecification Test Probability Values				
	Residual	Ser. Correlation	ARCH	Non-linearity
	Univariate	0.998	0.079	0.475
	Inflation	0.905	0.366	0.167
	GDP	0.995	0.068	0.156
	Int. Rate	0.390	0.000	0.000

Table 3: Structural Stability and Misspecification Tests for UK (GDP Defl.)				
Structural Stability Test Probability Values				
Results for VAR model				
Equation	Parameter	Test Statistics		
		Average	Supremum	Exp. Average
Inflation	Constant	0.161	0.171	0.171
	Inflation	0.151	0.065	0.101
	GDP	0.246	0.171	0.176
	Int. Rate	0.201	0.236	0.216
	Err. Variance	0.261	0.231	0.231
GDP	Constant	0.000	0.000	0.000
	Inflation	0.101	0.085	0.085
	GDP	0.020	0.000	0.005
	Int. Rate	0.005	0.000	0.000
	Err. Variance	0.000	0.000	0.000
Int. Rate	Constant	0.317	0.367	0.367
	Inflation	0.246	0.261	0.261
	GDP	0.357	0.362	0.362
	Int. Rate	0.362	0.392	0.392
	Err. Variance	0.000	0.000	0.000
Correlations	Inflation/GDP	0.000	0.000	0.000
	Inflation/Rate	0.020	0.005	0.005
	GDP/Rate	0.095	0.095	0.095
Results for AR model				
	Constant	0.035	0.005	0.010
	AR Coeff.	0.075	0.045	0.050
	Err. Variance	0.201	0.090	0.095
Misspecification Test Probability Values				
	Residual	Ser. Correlation	ARCH	Non-linearity
	Univariate	1.000	0.899	0.901
	Inflation	1.000	0.000	0.852
	GDP	1.000	0.187	0.167
	Int. Rate	0.975	0.000	0.013

Appendix E: Results for the United States

Chart 1: Results on AR inflation analysis: The panels report the constant term, the AR coefficient, the error variance, the implied variance keeping the error variance fixed, the implied variance allowing the error variance to vary and the implied spectrum respectively.

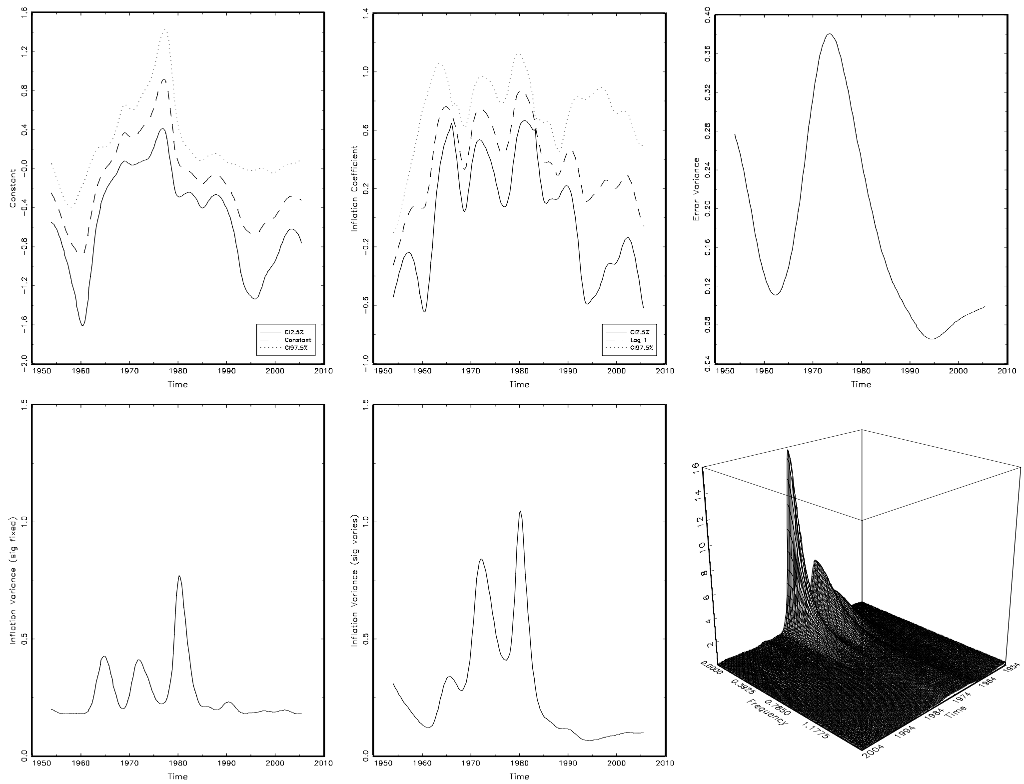

Chart 2: VAR results: estimated error variances

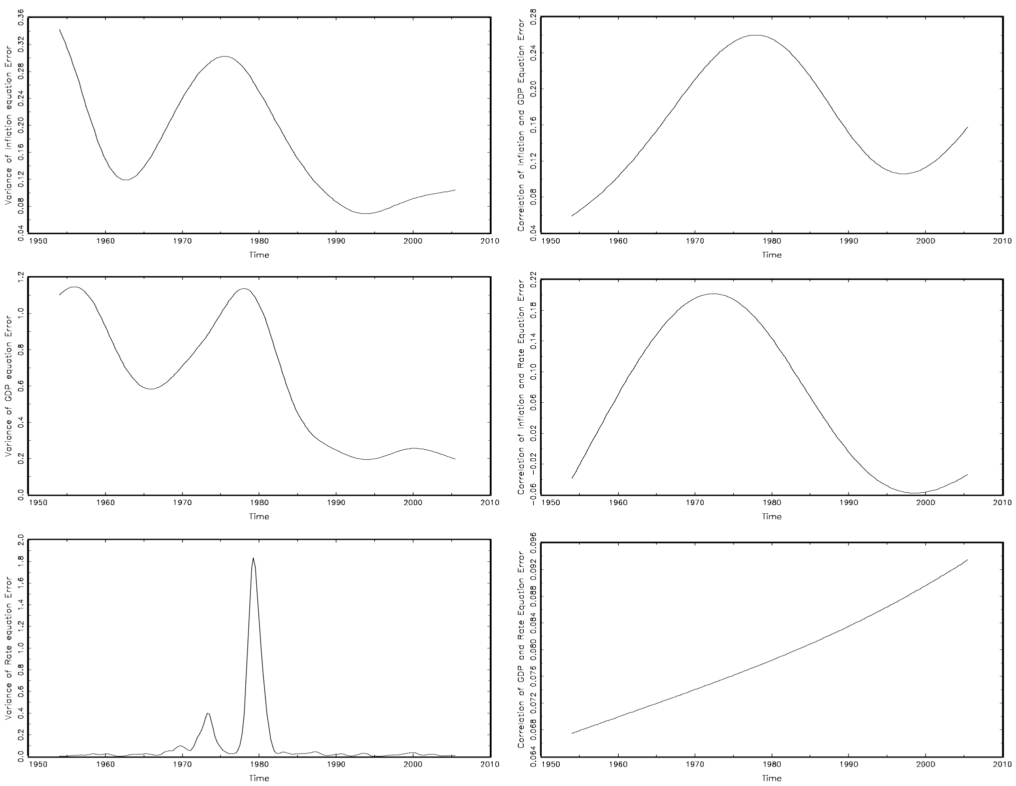

Chart 3: VAR results: implied variances

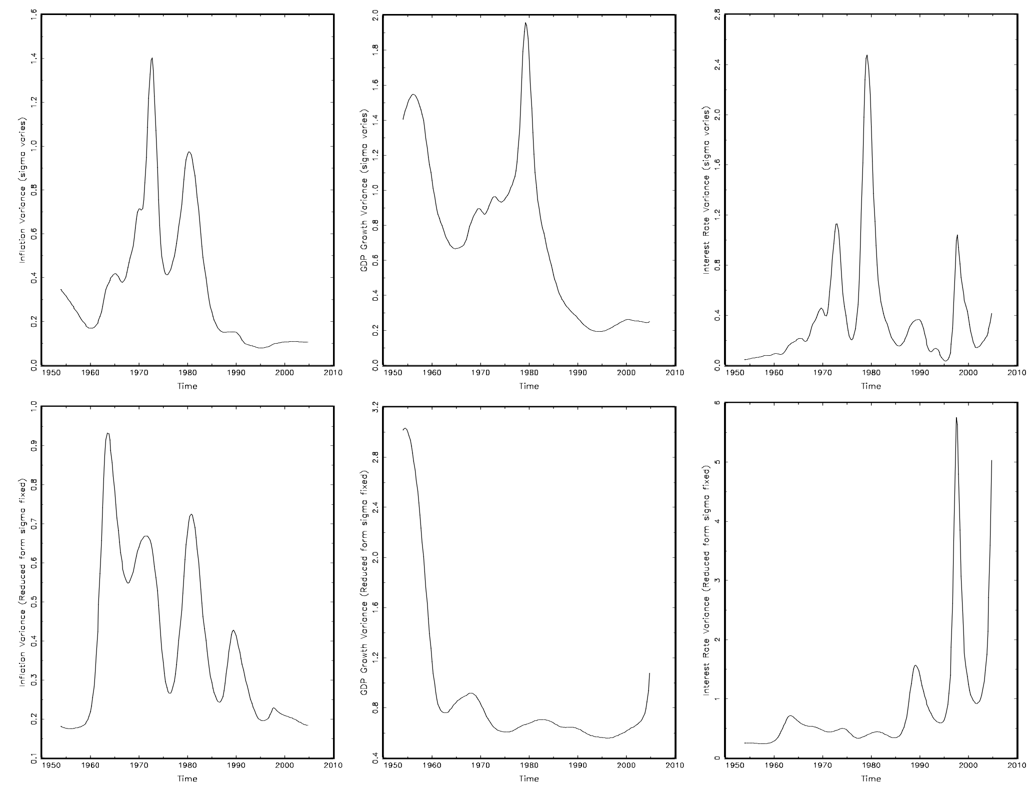

Chart 4: VAR results: implied spectrum and predictability measure for inflation

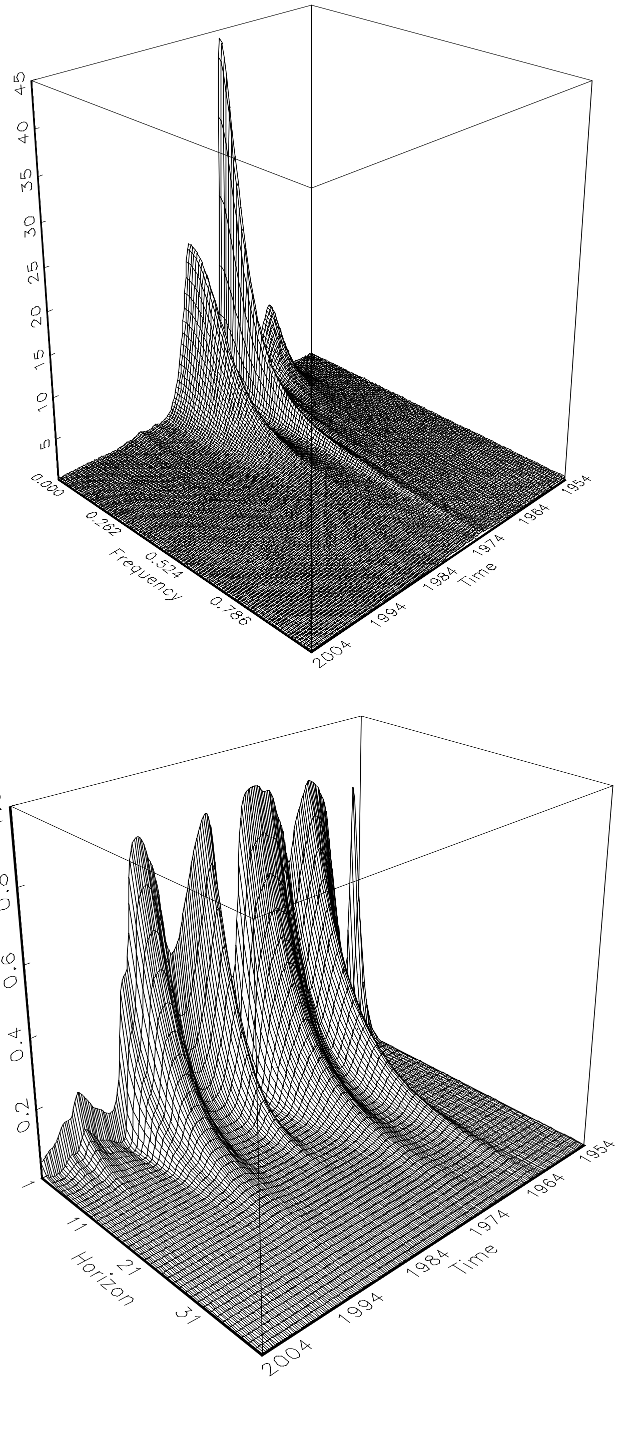

Chart 5: VAR results: predictability measure in 1974 Q1 and 2004 Q1

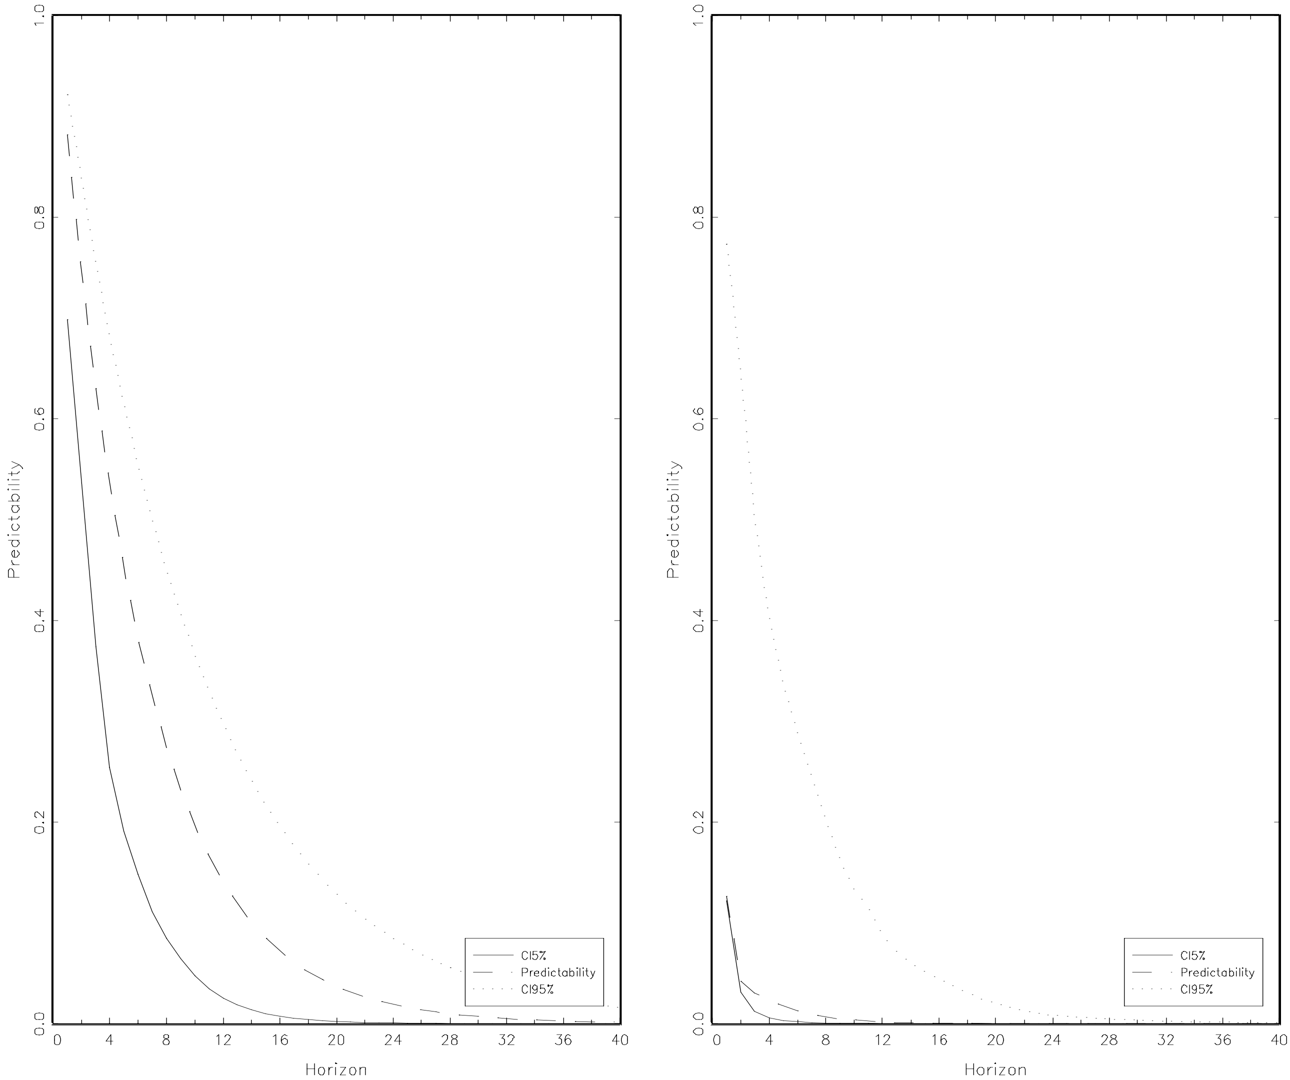

Appendix F: Results for the United Kingdom

Chart 6: Results on AR inflation analysis: the panels report the constant term, the AR coefficient, the error variance, the implied variance keeping the error variance fixed and the implied variance allowing the error variance to vary and the implied spectrum respectively.

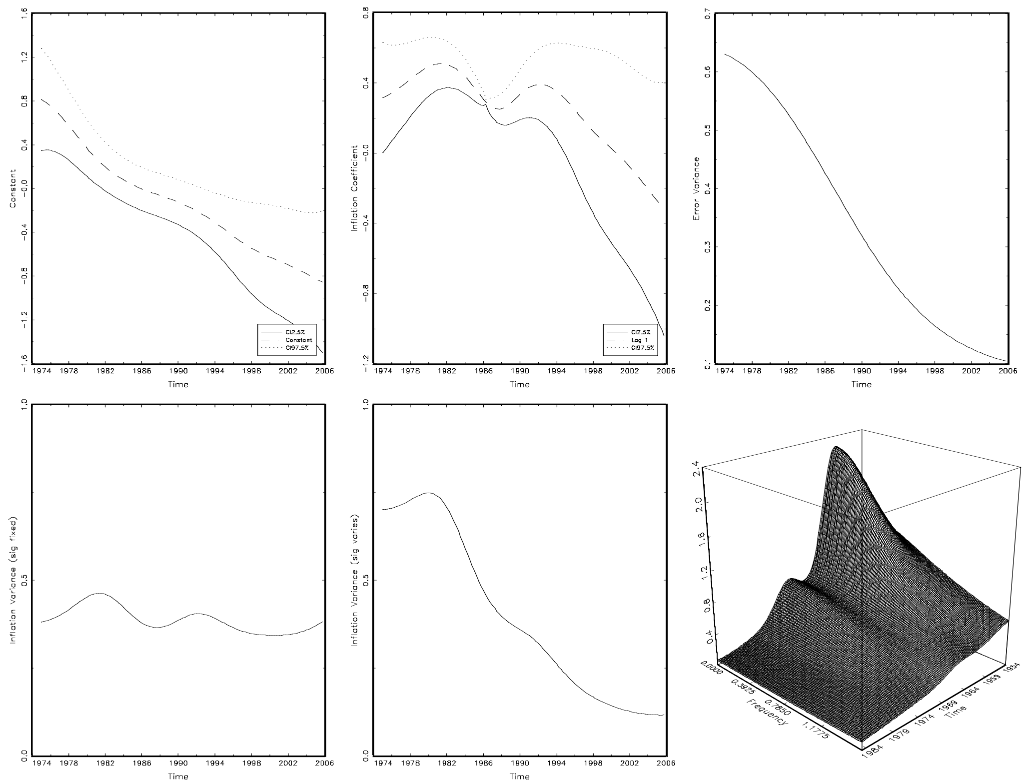

Chart 7: VAR results: estimated error variances

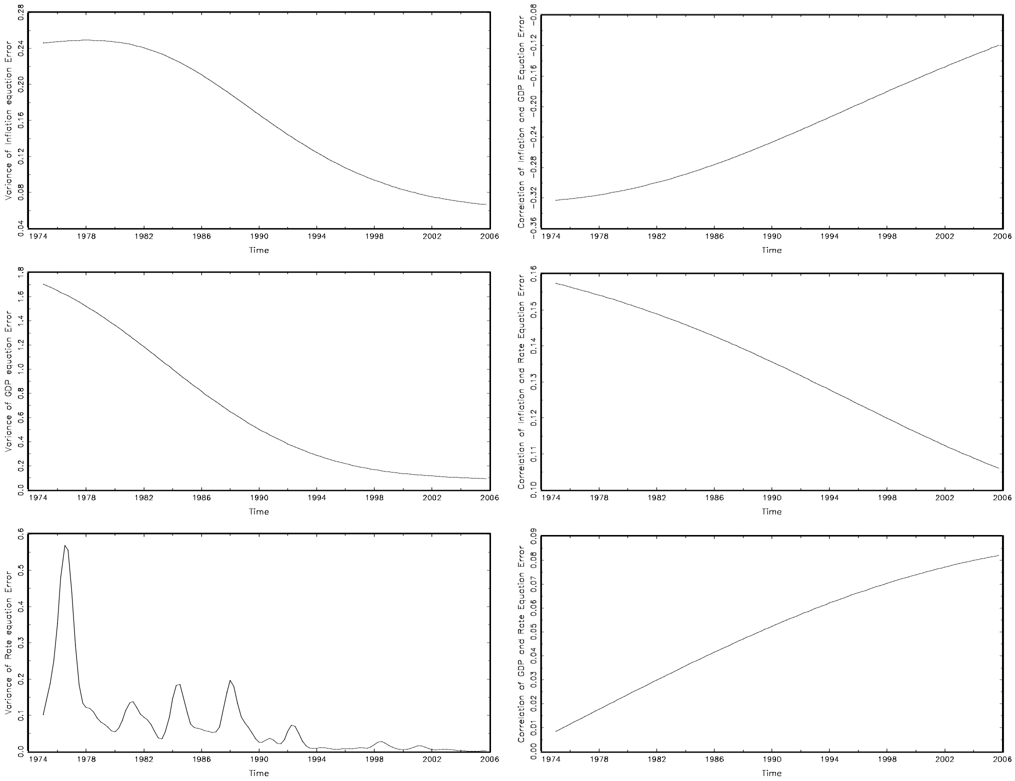

Chart 8: VAR results: implied variances

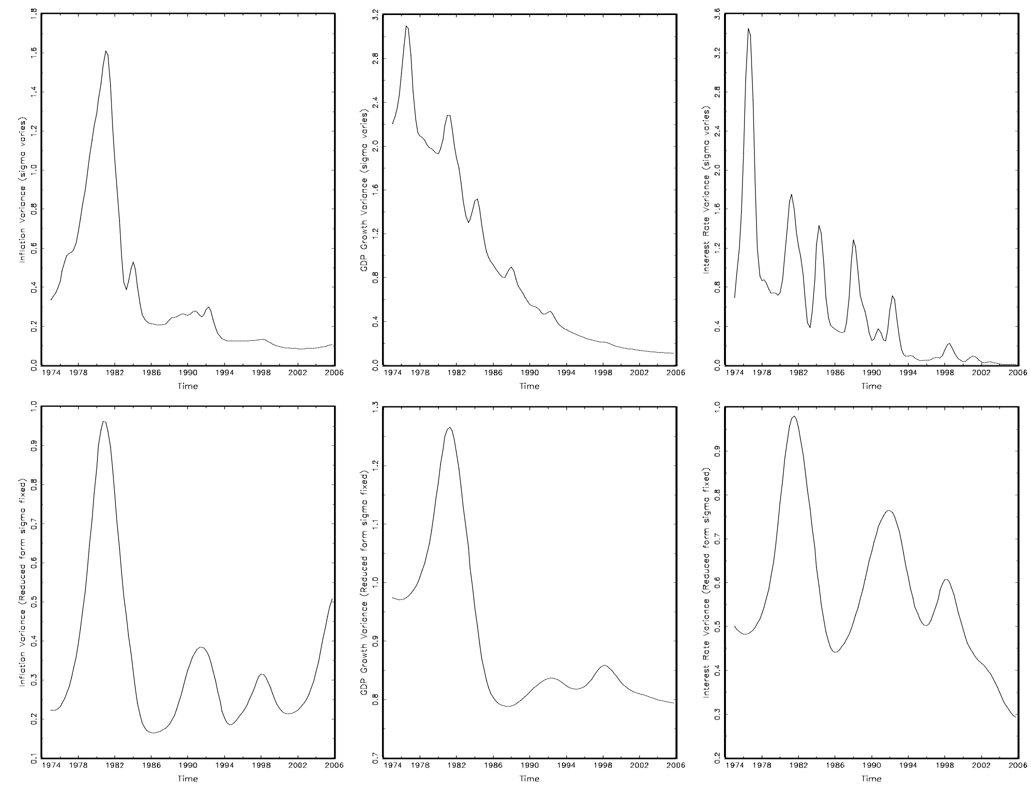

Chart 9: VAR results: implied spectrum and predictability measure for inflation

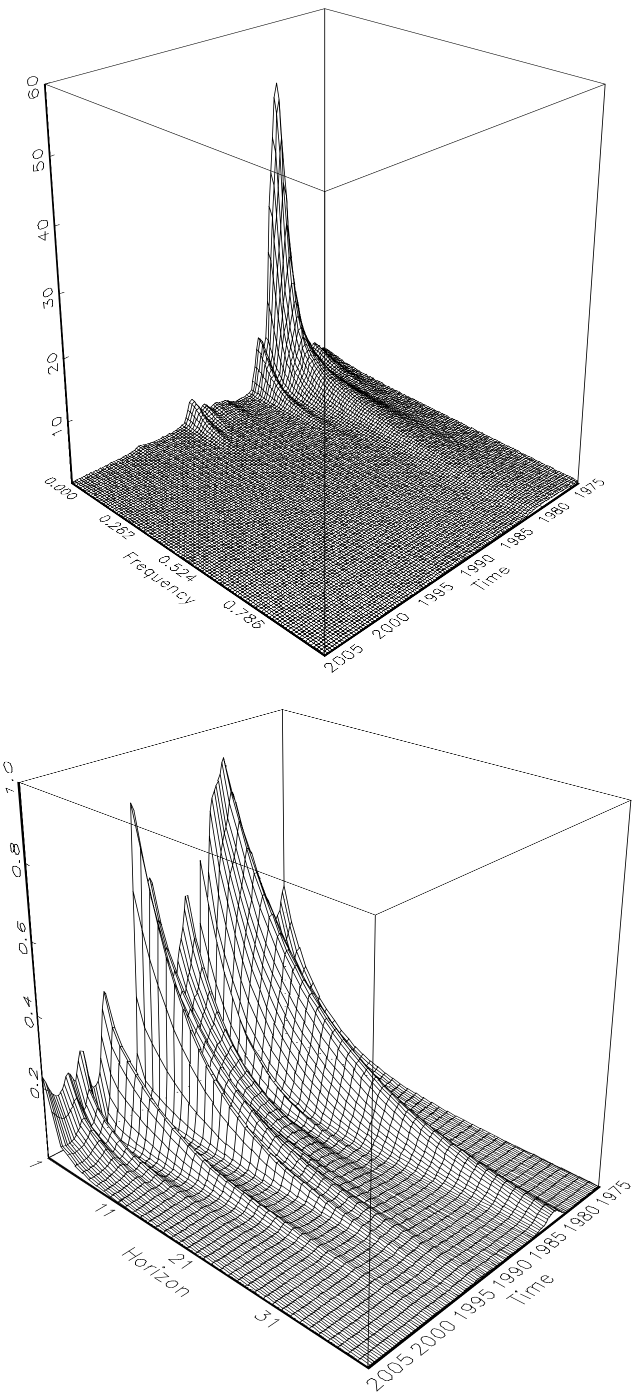

Chart 10: VAR results: predictability measure in 1982 Q3 and 2004 Q1

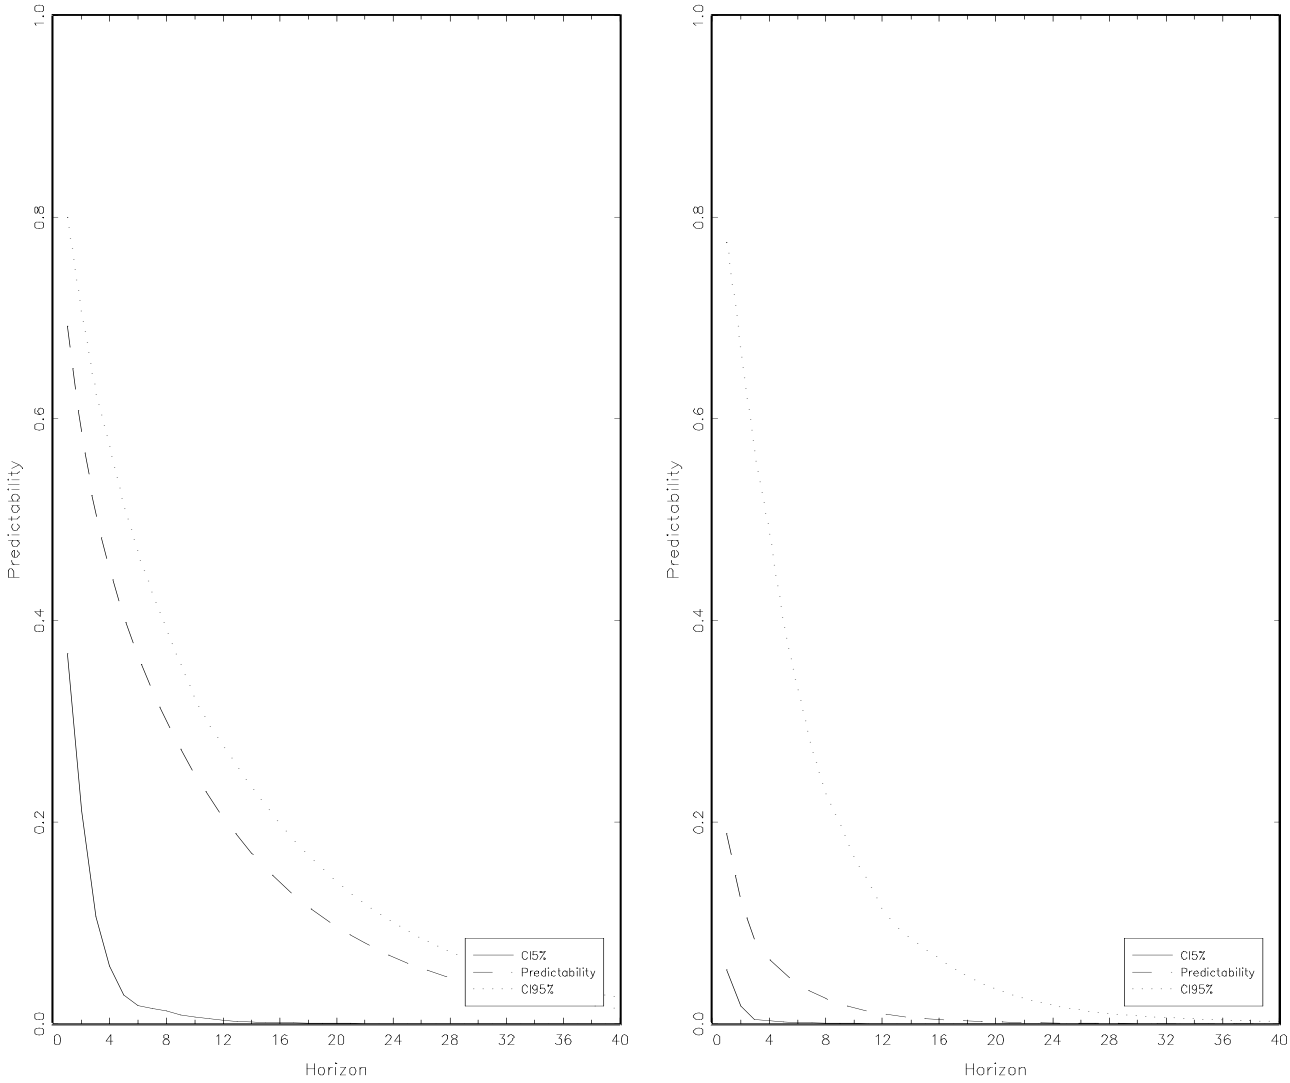

Appendix G: Structural results

Chart 11: Results for US CPI data set: inflation

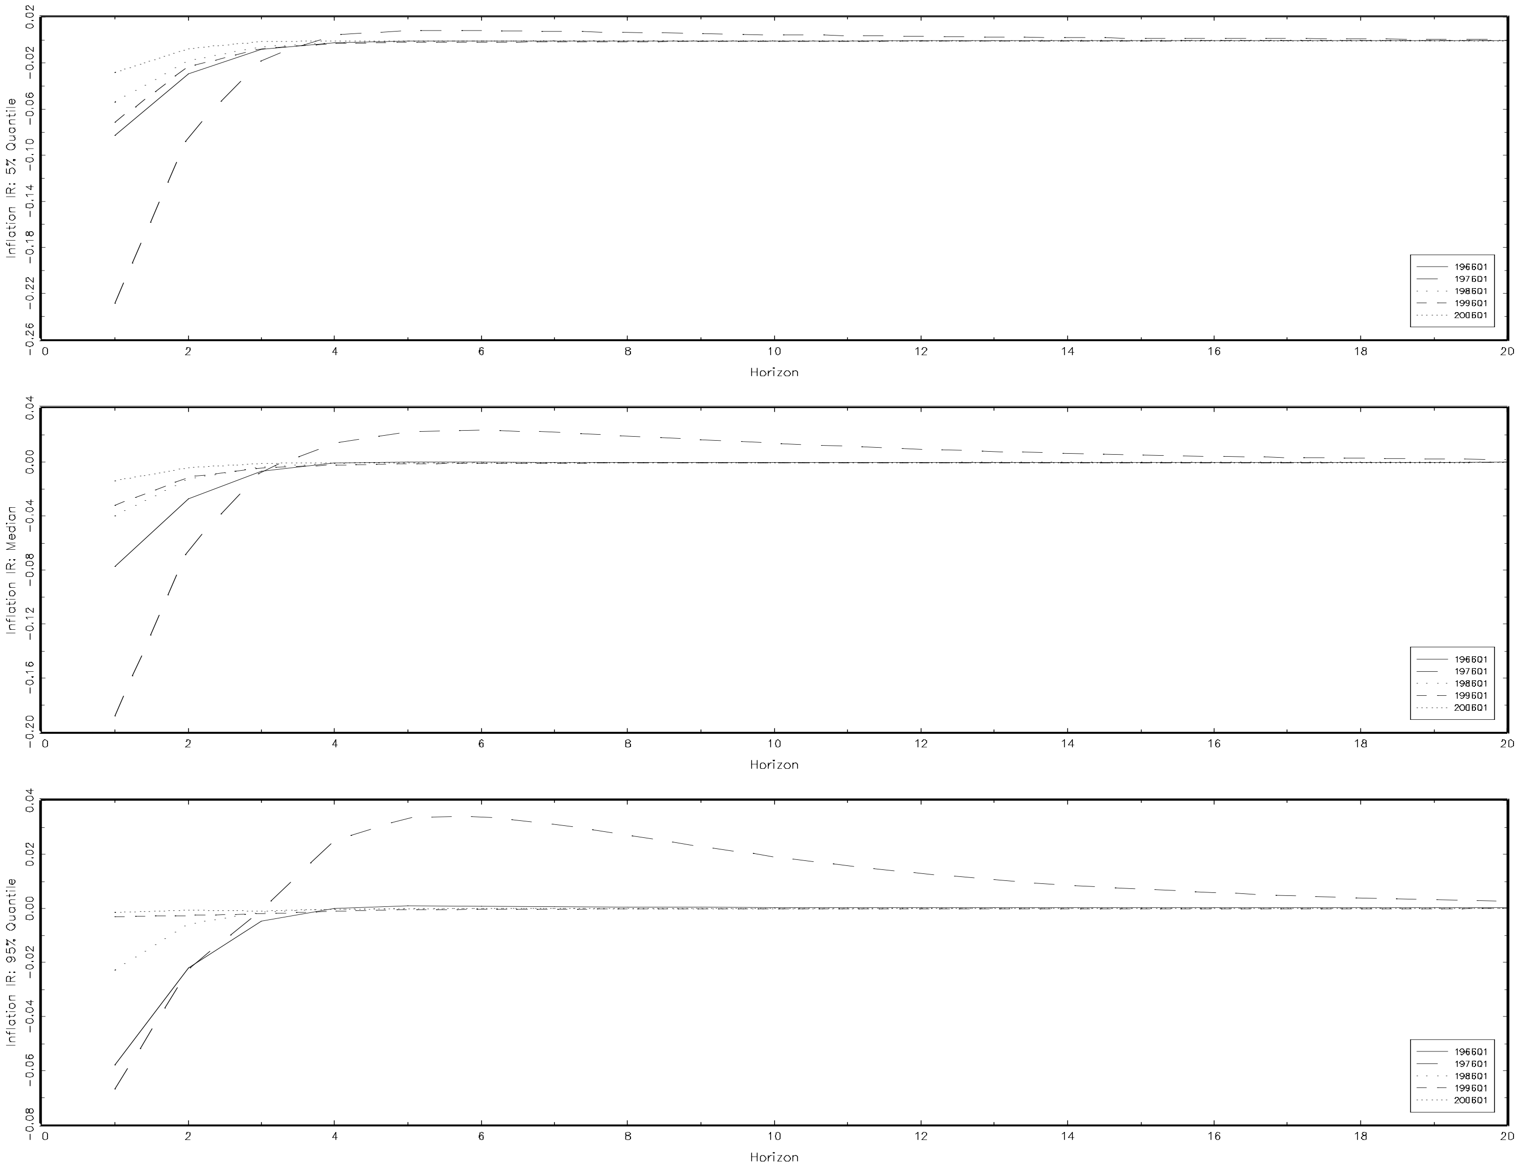

Chart 12: Results for US CPI data set: GDP

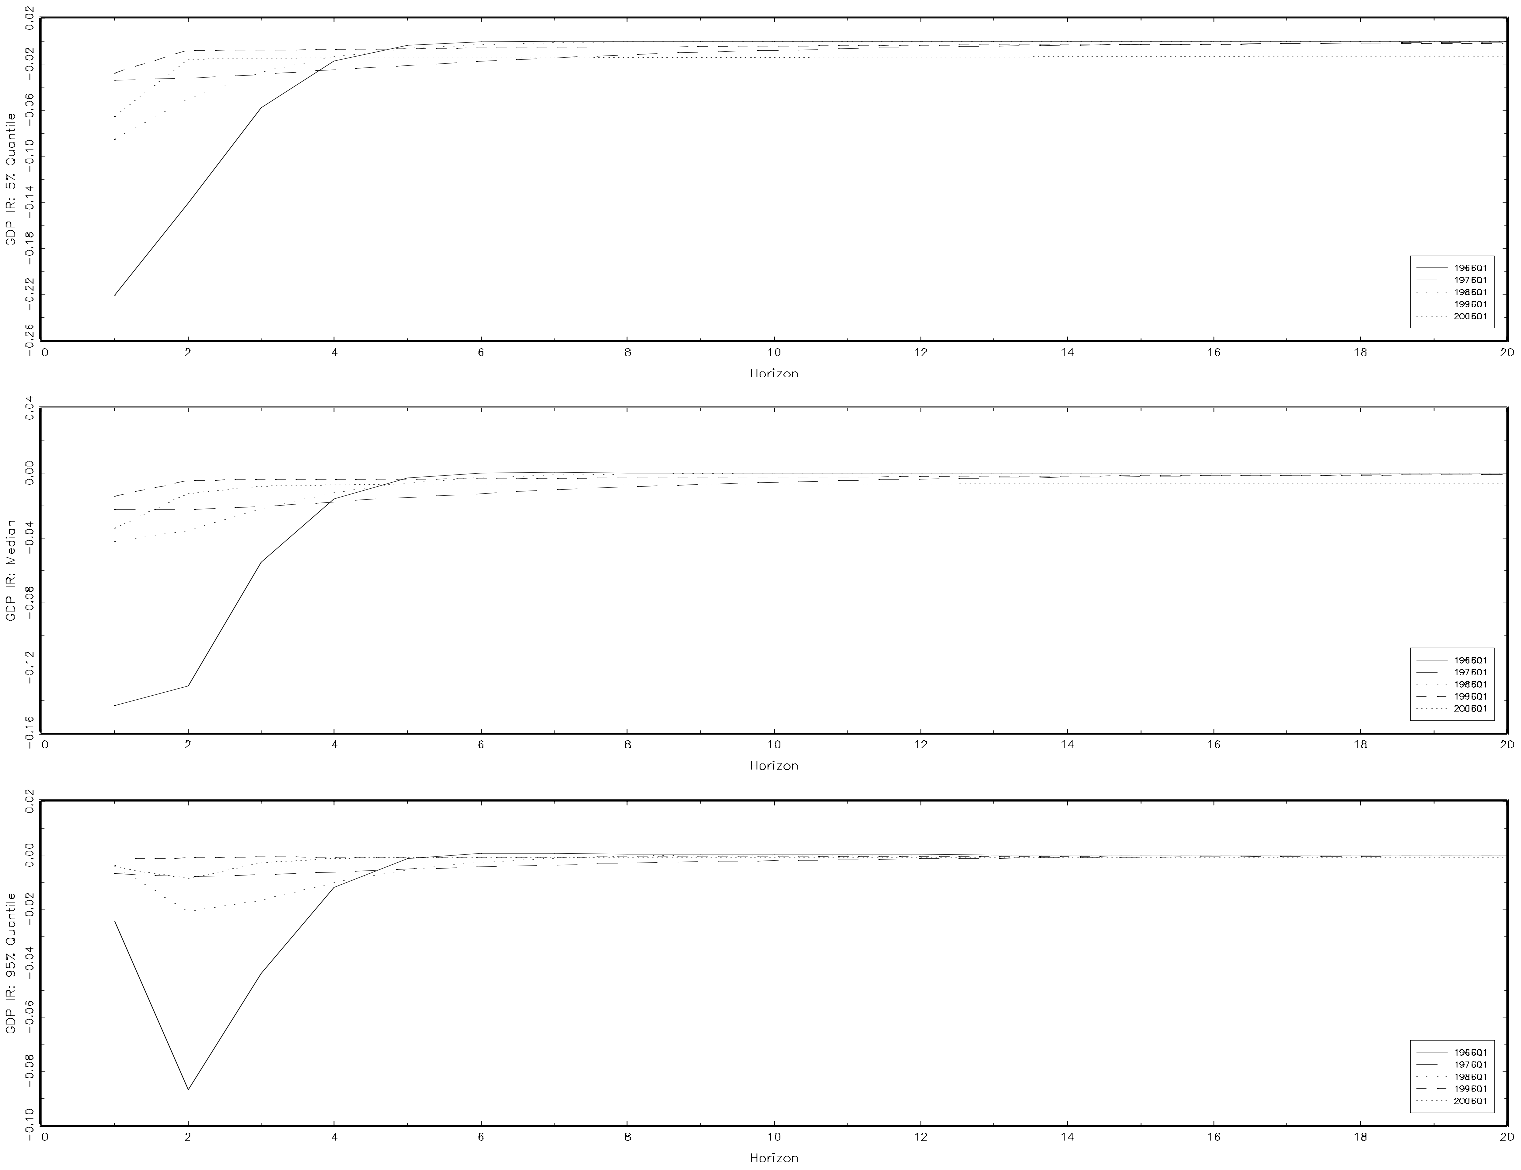

Chart 13: Results for US CPI data set: interest rate

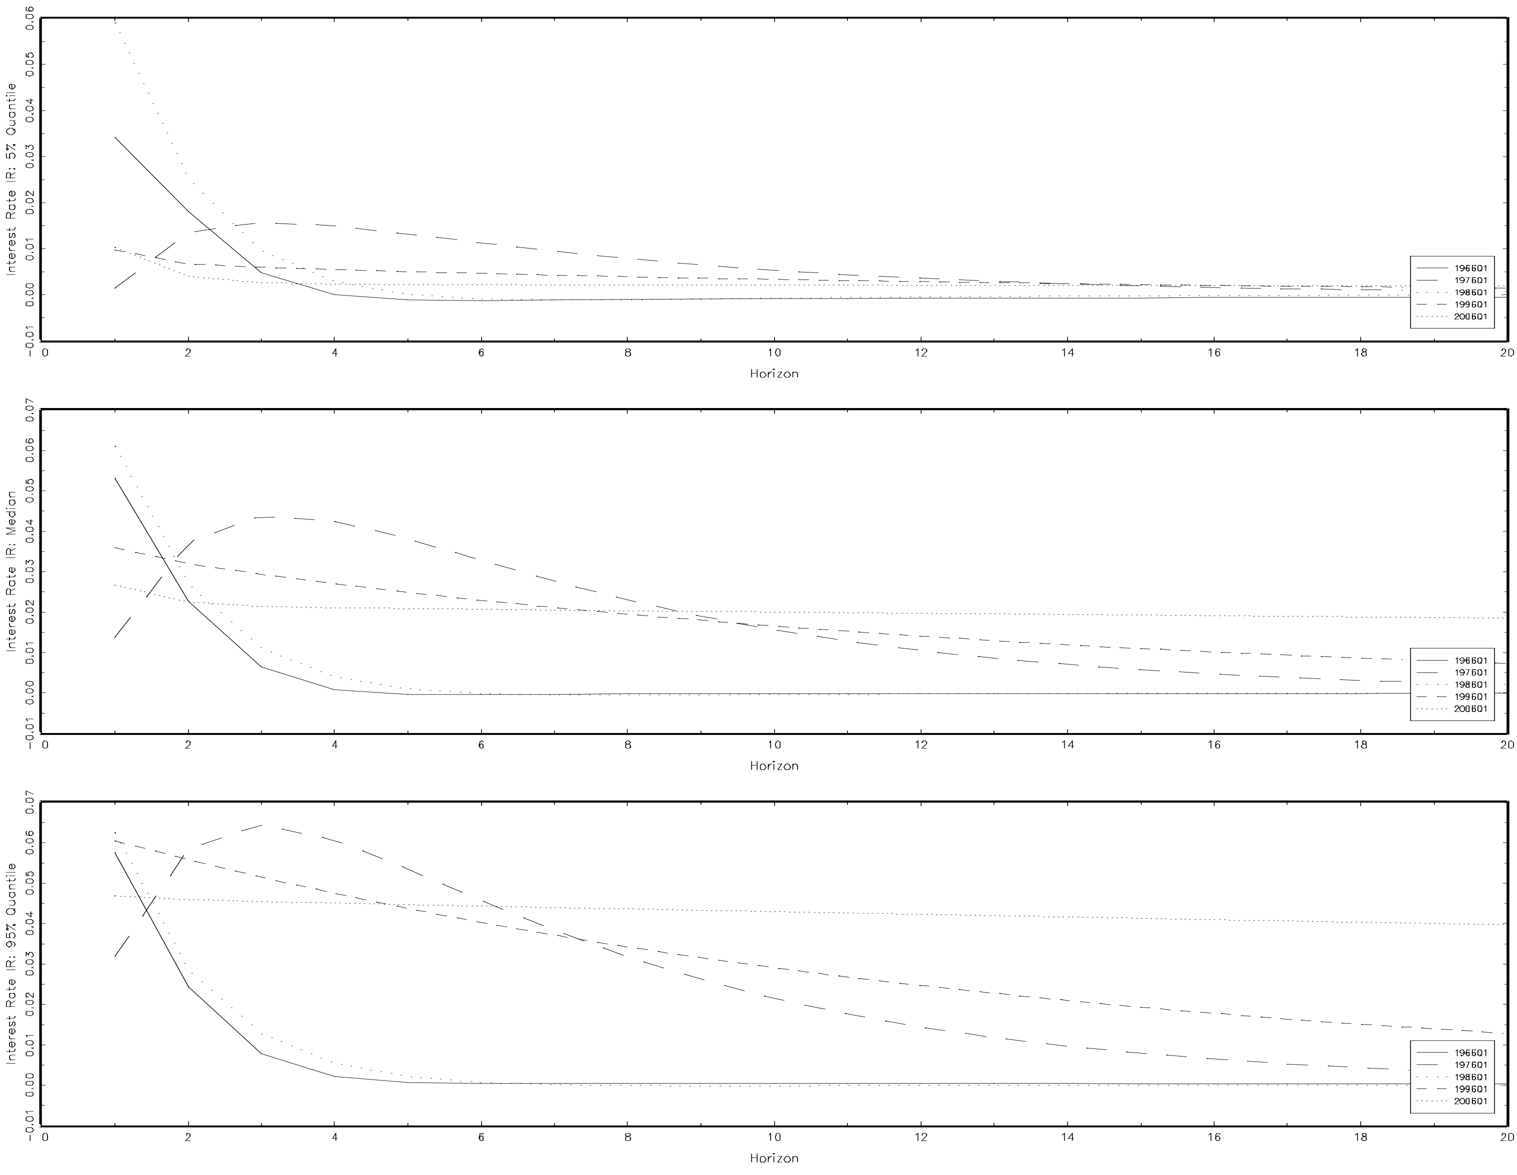

Chart 14: Results for US GDP deflator data set: inflation

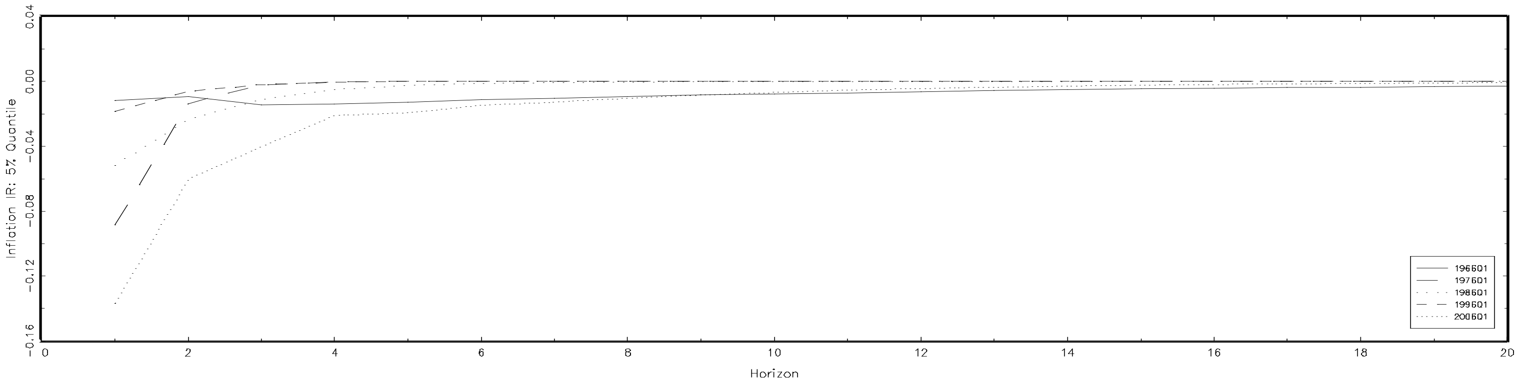

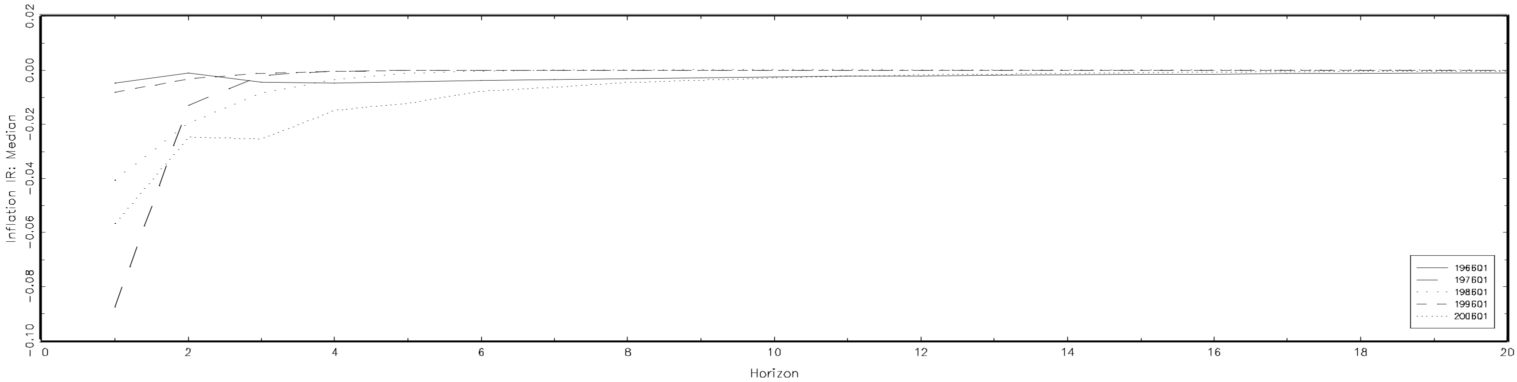

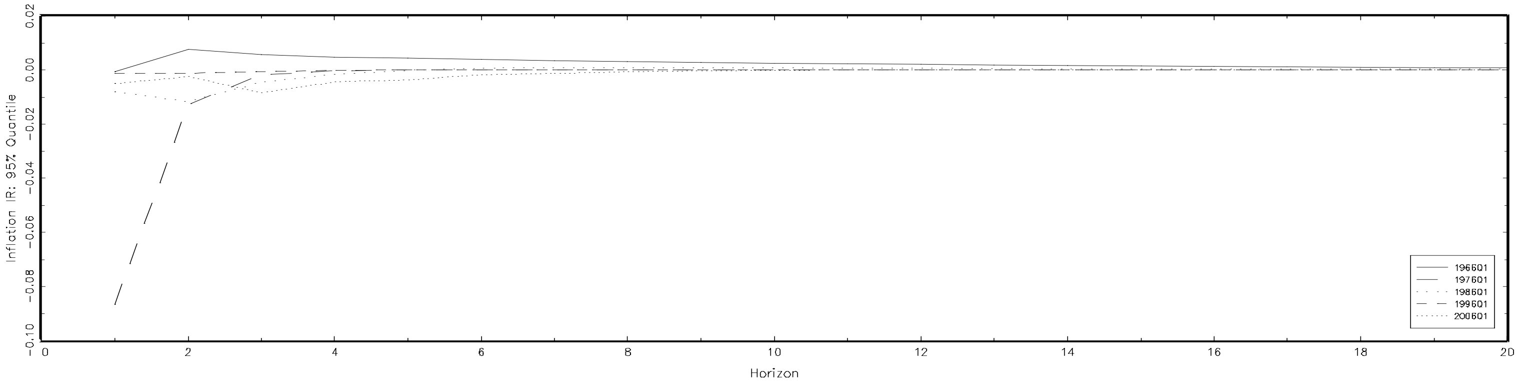

Chart 15: Results for US GDP deflator data set: GDP

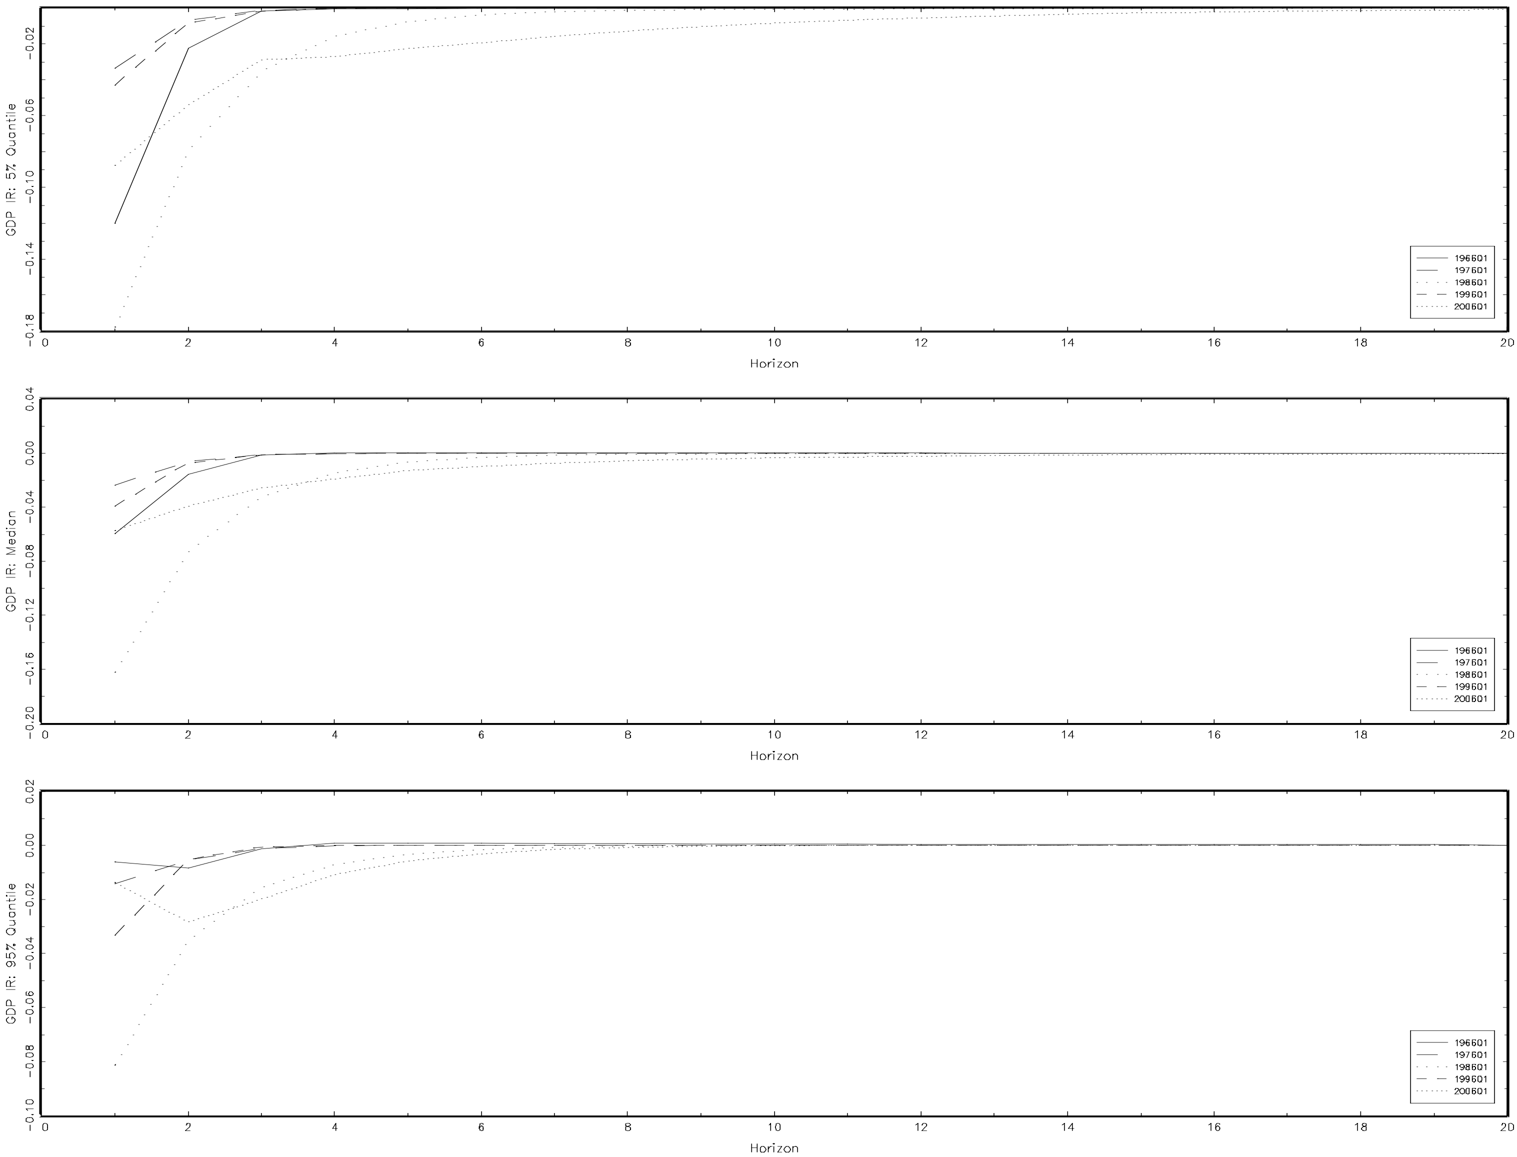

Chart 16: Results for US GDP deflator data set: interest rate

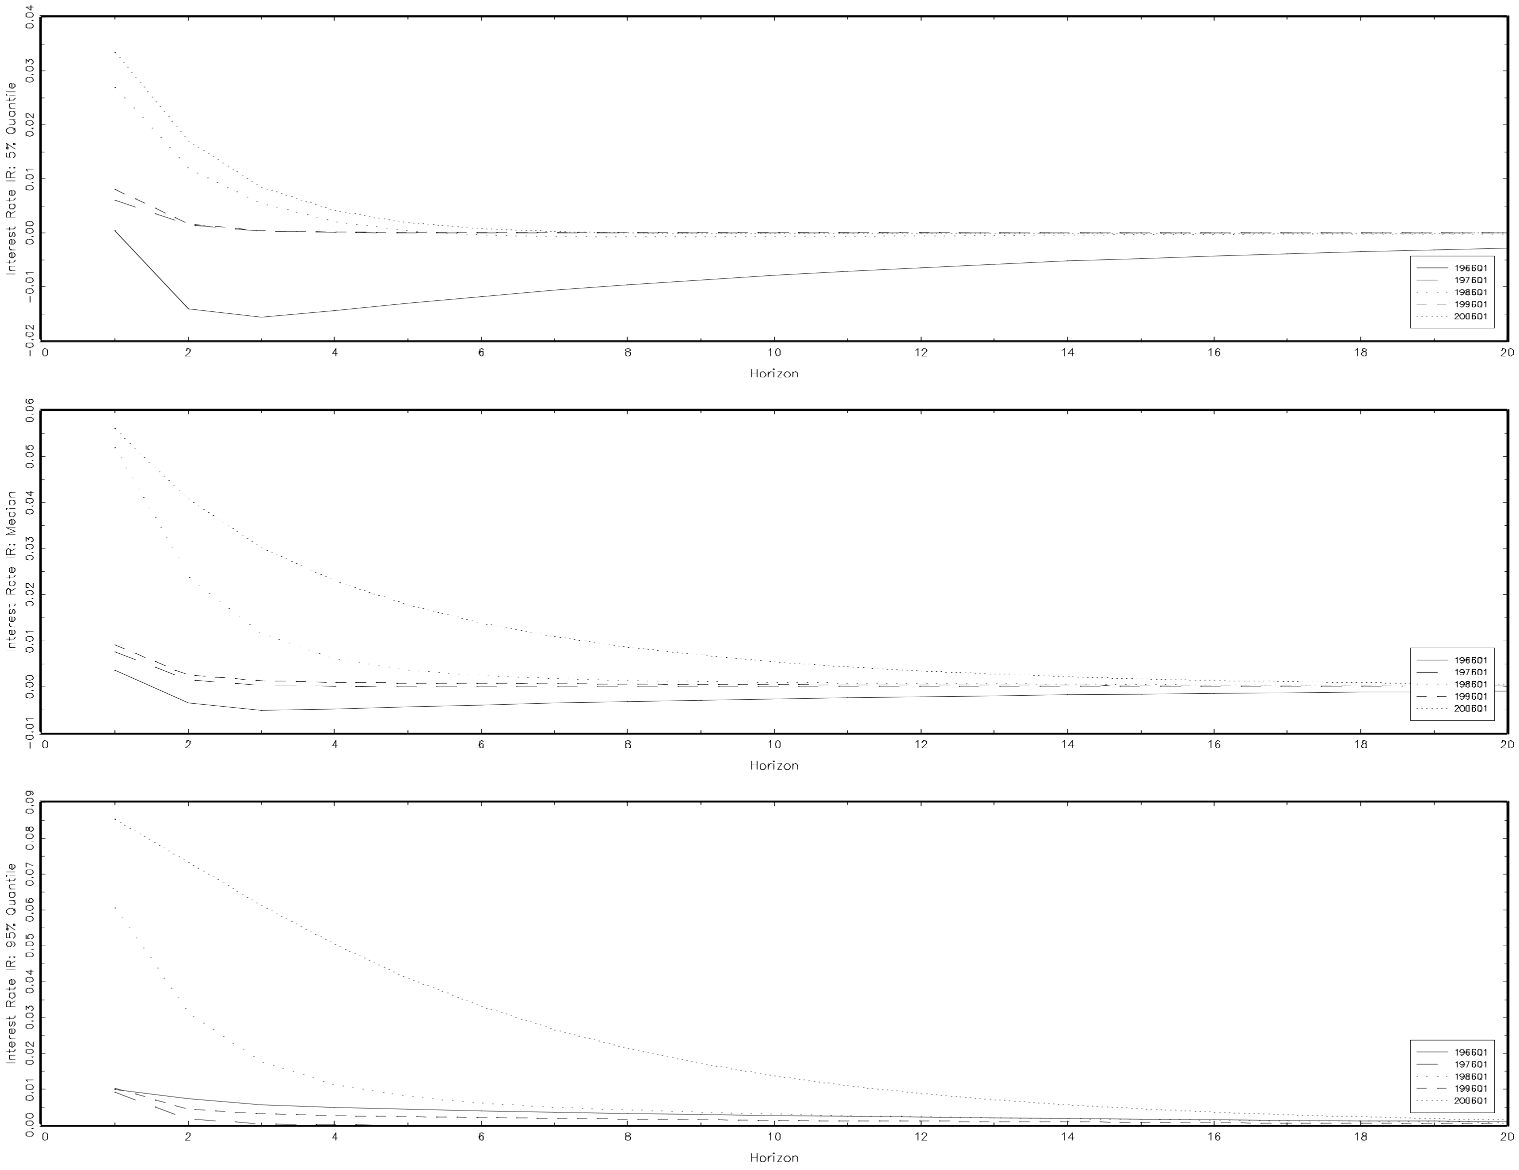

Chart 17: Results for UK data set: inflation

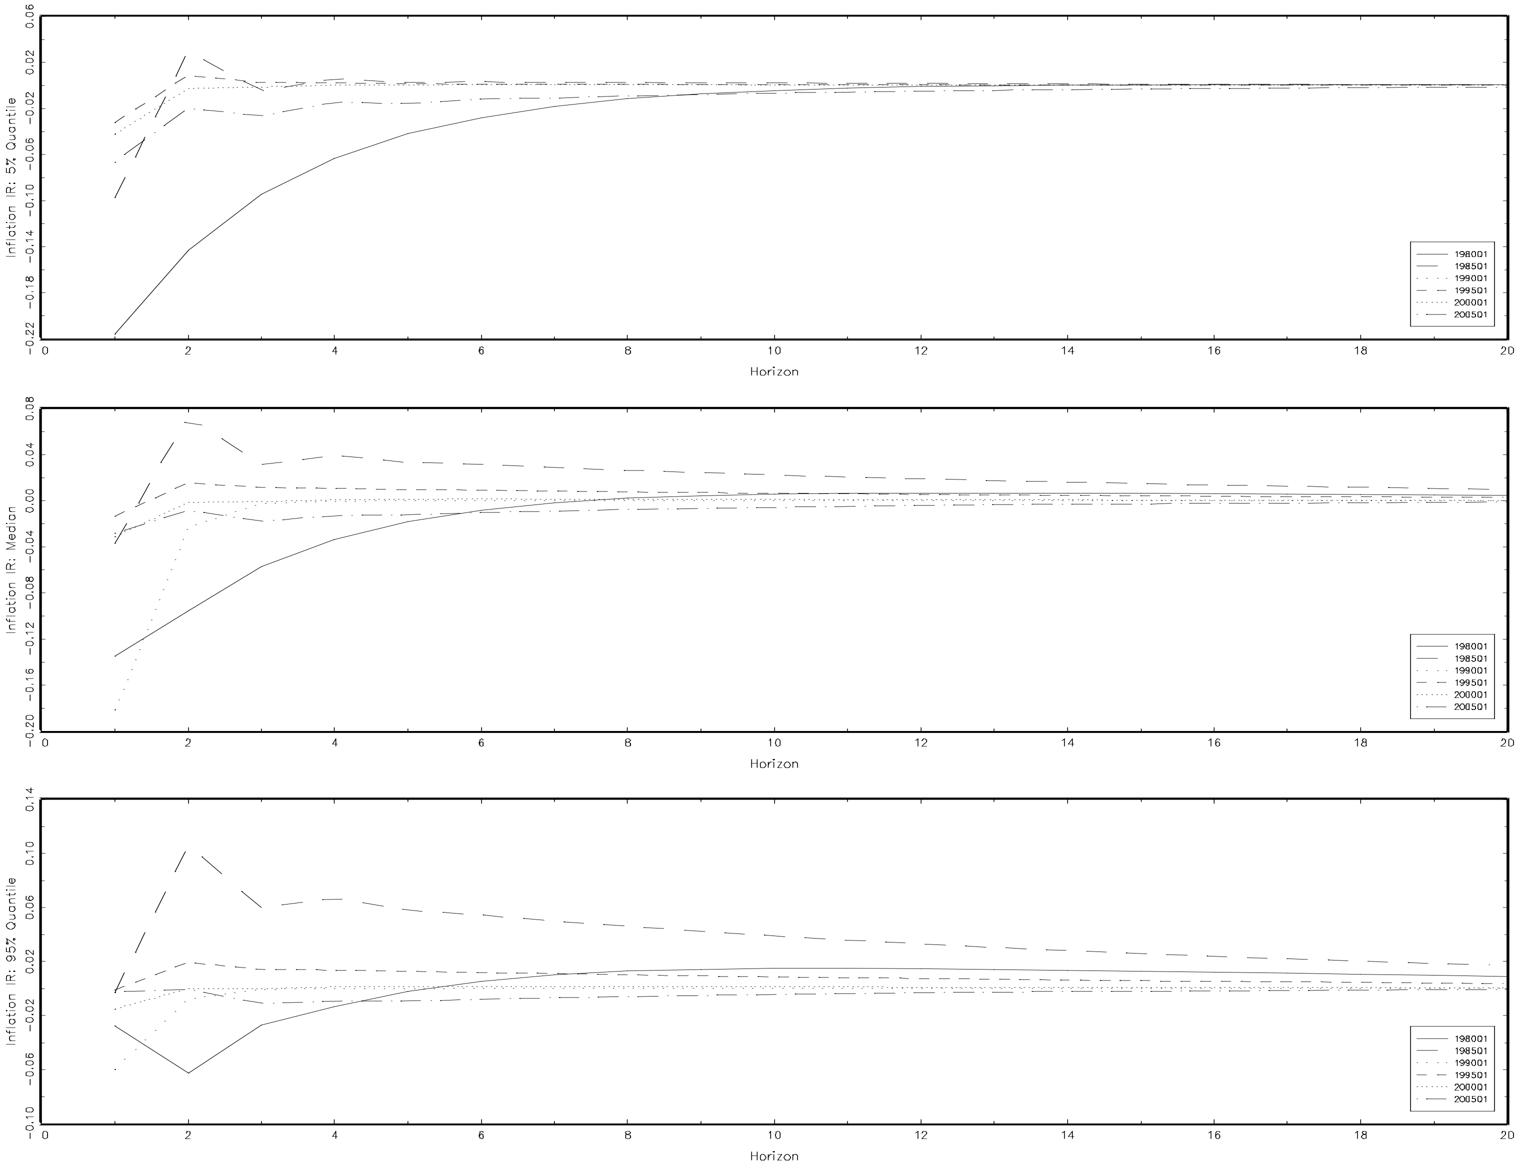

Chart 18: Results for UK data set: GDP

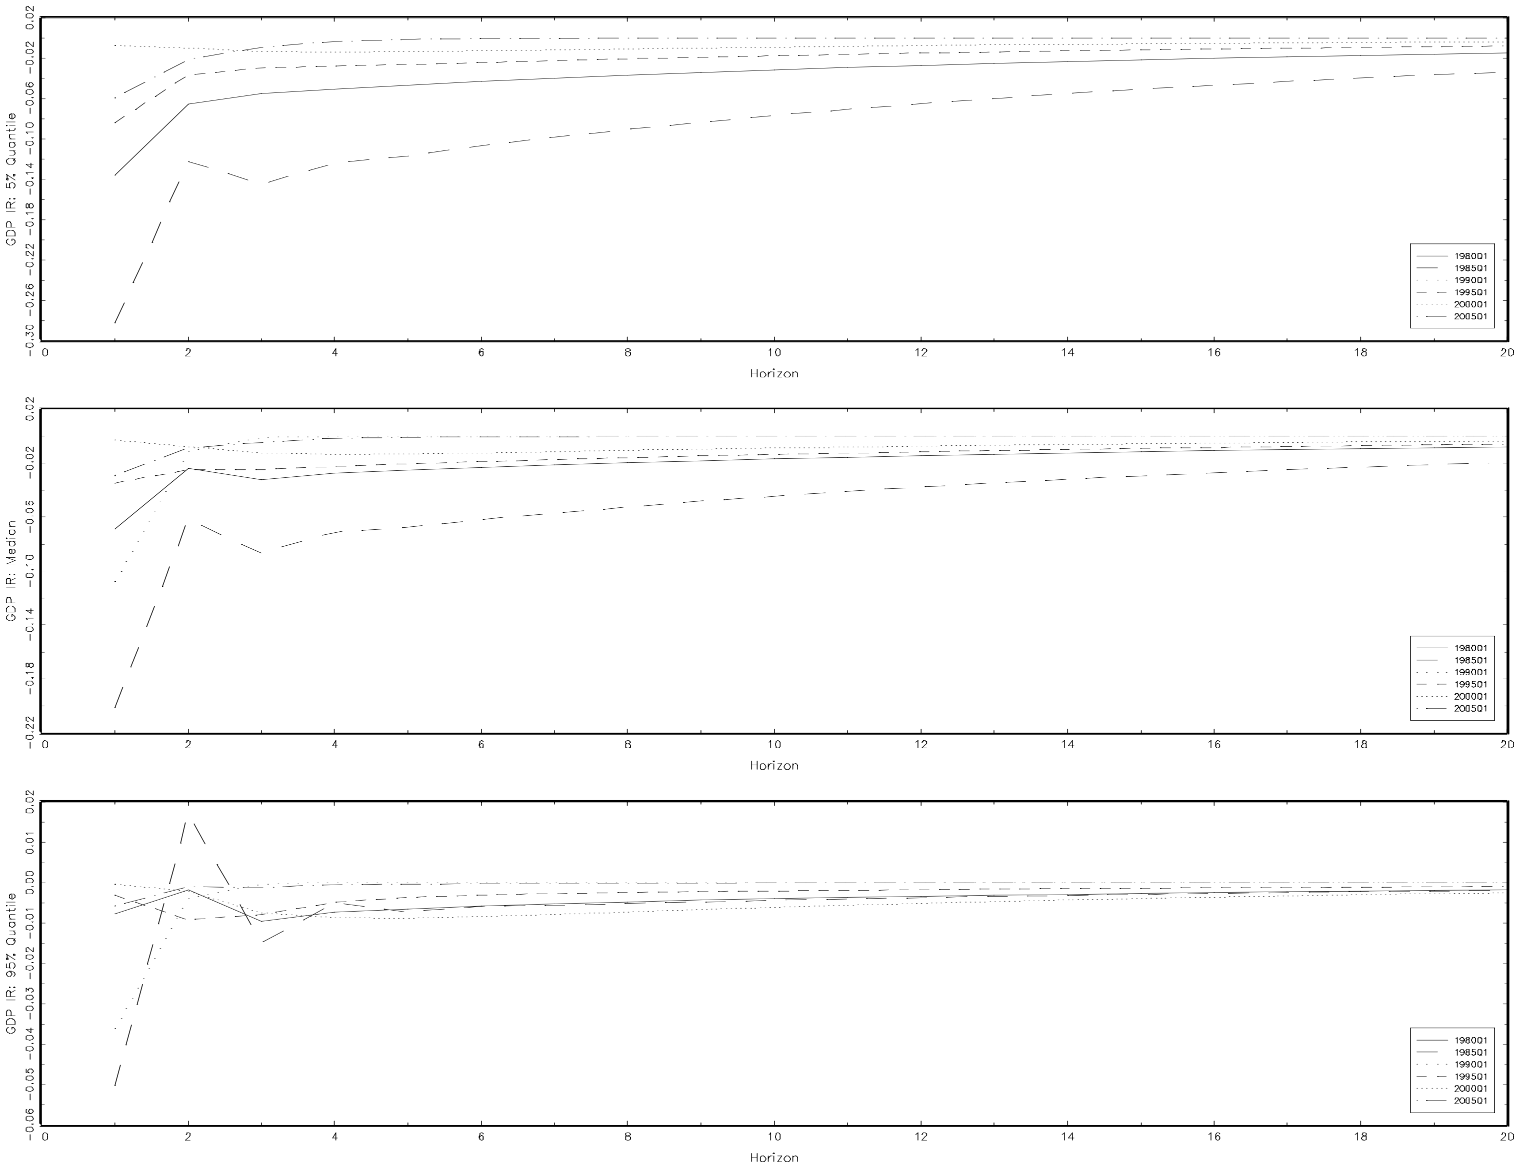

Chart 19: Results for UK data set: interest rate