Inflation Dynamics and Inflation Regimes

David Demery and Nigel W. Duck

February 2003

Discussion Paper No 03/549

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
UK

Inflation Dynamics and Inflation Regimes

David Demery
Nigel W. Duck[1]

February 2003

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
UK

Corresponding Author: David Demery
E-mail: David.Demery@bris.ac.uk
Telephone: +44 (0)117 928 8437
Fax: +44 (0)117 928 8577

Abstract

In this paper we develop and estimate a new-Keynesian model of inflation and use it to investigate the hypothesis that prices in the UK are re-set more frequently during periods of high inflation. In the model, firms are assumed to condition their expectations on an optimally-selected but incomplete information set and we further assume that the probability that they will reset their price in any quarter depends upon the prevailing inflationary regime. The model implies more complex inflation dynamics than conventional new-Keynesian models predict. We find that we cannot reject the formal restrictions implied by the model (using UK quarterly data) and we estimate that the mean time before prices were reset was around eight months during the 'high' inflationary regime and approximately two years when mean inflation was 'low'.

Classification Code: E21
Key Words: Rational expectations, incomplete information, macroeconomic dynamics, regime switching

[1] The authors acknowledge financial assistance from the Economic and Social Research Council: Award Reference R000223926, 'The imperfect use of information and inflation-output dynamics'.

1 Introduction

In this paper we develop and estimate a new-Keynesian model of inflation and use it to investigate the hypothesis that prices in the UK are re-set more frequently during periods of high inflation. The model differs from existing new-Keynesian models in three key respects. First, we allow firms to differ randomly in their technology and hence in their marginal costs. Specifically, we model a firm's marginal costs as the sum of two separate components: the first is specific to the firm itself, the second is common to all firms. Observation of its own marginal costs does not allow a firm to distinguish these two components perfectly, and this gives it an incentive to acquire information about other firms' marginal costs in order to improve the forecasts it makes of its own.

We assume that the acquisition and processing of this information is costly and, in our second extension, we assume that firms optimally select the information set on which to condition their forecasts - enlarging their information set up to the point where the marginal costs and marginal benefits of extending it further are equal. This in turn implies that firms will typically choose to condition their expectations on an *incomplete* information set. They will therefore not form conventional, fully-rational expectations (RE). We call the expectations they do form *optimally rational expectations* (ORE): they are rational because firms fully exploit their chosen information set; they are optimal because this information set itself is the result of a conventional optimising process.[1]

This departure from RE accords with recent work within the new-Keynesian (and other) literature, much of which has found that the assumption of RE is too strong. At the same time ORE is theoretically more satisfactory than the *ad hoc* assumptions, such as adaptive or quasi-adaptive expectations, that others have used.[2] We show that the incorporation of ORE has important implications for the dynamics of inflation: in particular it can, in principle, explain the observed influence of lagged inflation on current inflation - a feature that new-Keynesian models incorporating RE cannot explain.

[1]Elsewhere in the literature such expectations have been called 'economically rational expectations' (see Crettez and Michel (1992)). The concept is also similar to the notions of near and bounded rationality (see Conlisk (1996)) and Galbraith's (1988) 'errors-in-variable' model of expectations.

[2]For examples see Fuhrer and Moore (1995), Galí and Gertler (1999) and Galí *et al.* (2001). For similar models applied to the UK case see Balakrishnan and López-Salido (2000, 2001). Mankiw (2001) incorporates adaptive expectations into a New-Keynesian model to explain the stylised facts of the dynamic response of inflation and unemployment to monetary shocks. Ball (2000) assumes that firms use only lagged inflation when forming expectations, a feature he labels 'near-rationality'. Roberts (1997, 1998) measures expectations from survey data. In models closer to the one presented in this paper, Mankiw and Reis (2001) assume information disperses slowly across the economy and Woodford (2001) considers the effects of noisy information in a Phelps island model context.

Our third extension is to allow the stickiness of prices, as measured by the probability of a firm resetting its price, to be a function of the underlying inflation regime. Most new-Keynesian models treat this probability as fixed despite Taylor's (1999, p.1021) assessment that ' ... prices at small businesses, industrial prices, and even the prices of products like magazines are adjusted more quickly when the rate of inflation is higher. The dependency of price and wage setting on events in the economy is one of the more robust empirical findings in the studies reviewed here'. As part of our investigation of this effect, we identify 'high' and 'low' inflation regimes for the UK by estimating a two-state Markov switching model; we also examine the behaviour of the re-set probability using moving-window regression techniques.

Our principal empirical findings are that a new-Keynesian model which incorporates ORE and the expected effects of the inflation regime on the probability of resetting prices finds support in UK data: firms' use of an imperfect information set can account for observed UK inflationary dynamics; and that the probability of resetting prices is appreciably higher when inflation is high. We estimate that the mean time before a price is reset was around eight months during the period when UK mean inflation was 'high', and approximately two years when inflation was 'low'.

The paper is in three sections. In the first we develop the model. In the second we describe the data, report the results and use them to show how inflation dynamics are altered by the effects different inflation regimes have on the probability of a firm resetting its price. We end with a summary.

2 The Model

We assume a continuum of firms indexed by $j \in [0, 1]$.[3] Each firm is a monopolistic competitor and produces a differentiated good $Y_{j,t}$, which it sells at the nominal price $P_{j,t}$. Each firm faces an iso-elastic demand curve given by $Y_{j,t} = (P_{j,t}/P_t)^{-\chi} Y_t$ where Y_t and P_t are *aggregate* output and the *aggregate* price level respectively. We assume that output is a simple linear function of the single factor, labour. Specifically, the production function for firm j is $Y_{j,t} = A_{j,t} N_{j,t}$, where $N_{j,t}$ is the quantity of labour employed by firm j in period t and $A_{j,t}$ is a technological factor affecting firm j. Note that by indexing the technology factor on j we are departing from the usual assumption (see, for example, Gali *et al.* (2001)) that this factor is common to all firms.

Nominal prices are assumed to be set as suggested in Calvo (1983): each firm resets its price with probability $1 - \theta$ each period, independently of the

[3] Our underlying assumptions are initially the same as those presented in Gali *et al.* (2001) with the simplification that the production function is linear in the single factor, labour.

time elapsed since the last adjustment. So, each period, $1-\theta$ of firms reset their prices. θ is therefore a measure of price rigidity. For the moment we shall make the standard assumption that this probability is fixed.

Those firms which do reset are assumed to do so with the aim of maximising their expected discounted profits given technology, factor prices, and the constraint on price adjustment defined by θ. The resulting optimal price-setting rule is that each firm should set its price as a markup over a discounted stream of expected future nominal marginal costs, where, if the firm faces a low probability of being able to reset its price, (a high value of θ), the firm places more weight on expected future marginal costs. Formally, a logarithmic approximation to the optimising rule is,[4]

$$p^*_{j,t} = \log(\mu) + (1-\beta\theta)\sum_{k=0}^{\infty}(\beta\theta)^k E_{j,t}MC_{j,t+k} \tag{1}$$

which can be re-written more conveniently for our purposes as,

$$p^*_{j,t} = \log(\mu) + MC_{j,t} + \sum_{k=1}^{\infty}(\beta\theta)^k E_{j,t}(\Delta MC_{j,t+k}) \tag{2}$$

where $p^*_{j,t}$ is the log of the newly-set price of firm j; $\mu \equiv \chi/(\chi-1)$ and is the firm's desired gross markup; $MC_{j,t+k}$ is the logarithm of the nominal marginal cost in period $t+k$ of a firm which last reset its price in period t; and β is a subjective discount factor. Notice that the marginal cost terms in equations (1) and (2) are indexed on j because we are allowing the technological factor - and hence the marginal costs of firms who reset their prices - to differ amongst firms. Firms with a greater than average technological factor will have lower than average marginal cost.

Most new-Keynesian models build on equation (2) by assuming that firms form their expectations about current and future marginal costs fully rationally, that is by conditioning them on *all* available information. However, most empirical studies suggest that new-Keynesian models which incorporate RE are empirically unsatisfactory - specifically they fail to predict the dynamics of inflation - and they have been forced to incorporate other *ad hoc* models of expectations formation.

To avoid both empirical rejection and theoretical awkwardness, we incorporate a model of expectations developed in Demery and Duck (2002) and used by them in simpler form in Demery and Duck (2001). This model of expectations formation assumes that, rather than using *all possible* available information, at least some of which might be regarded as costly to acquire, agents will select their information set by weighing up the costs of acquiring more information against the benefits it confers. Having thus decided their information set, agents fully exploit it. Demery and Duck (2002) term these expectations 'optimally rational expectations' because the selection of the

[4]See Galí *et al.* (2001 p. 1244).

information involves a standard optimising process, and, *given that information set*, the agent is assumed to make rational expectations. They suggest that, in general, ORE will involve the use of an incomplete information set.

To apply this idea here we begin by assuming that ΔMC_t, the change in the log of the economy-wide average nominal marginal cost, follows a stationary process and therefore has the moving-average representation,

$$\Delta MC_t = d + \sum_{i=0}^{T} \alpha_i \varepsilon_{t-i} \equiv d + \alpha(L)\varepsilon_t \tag{3}$$

where $\alpha_0 = 1$; ε_t is white noise; and d is a constant.

We further assume that the deviation of $\Delta MC_{j,t}$ from this average is itself a stationary process such that,

$$\Delta MC_{j,t} = d + \Delta MC_t + \sum_{i=0}^{T} \gamma_i u_{j,t-i} \equiv \Delta MC_t + \gamma(L) u_{j,t} \tag{4}$$

where $\gamma_0 = 1$; and $u_{j,t}$ is also white noise.[5]

Combining equations (3) and (4) gives,

$$\Delta MC_{j,t} = d + \alpha(L)\varepsilon_t + \gamma(L)u_{j,t} = \sum_{i=0}^{T} \alpha_i \varepsilon_{t-i} + \sum_{i=0}^{T} \gamma_i u_{j,t-i} \tag{5}$$

One way of interpreting the conventional RE assumption is that each firm conditions its expectations of $\Delta MC_{j,t+k}$ $(k \geq 0)$ in period t on an information set which contains the *separate* histories of $\{\varepsilon_t\}$ and $\{u_{j,t}\}$, by which we mean $\{\varepsilon_t, \varepsilon_{t-1} .... \varepsilon_{t-\infty}\}$ and $\{u_{j,t}, u_{j,t-1}, ... u_{j,t-\infty}\}$. In this context, this is what we mean by conditioning on all the available information. So, assuming RE, it follows that,

$$\begin{aligned} \text{If } 0 &\leq k < T: \\ E^F_{j,t} \Delta MC_{j,t+k} &= d + \sum_{i=k}^{T} \alpha_i \varepsilon_{t+k-i} + \sum_{i=k}^{T} \gamma_i u_{j,t+k-i} \\ \text{Otherwise} &: \quad E^F_{j,t} \Delta MC_{j,t+k} = d \end{aligned}$$

Observation of the *separate* history of $\{\varepsilon_t\}$ could arise from *direct* observation of the aggregate variable ΔMC_t from some published source. Observation of the history of $\{u_{j,t}\}$ would then occur naturally from observation of $\Delta MC_{j,t}$. An alternative interpretation of RE, which is more useful for what follows, is that a firm which has no *direct* observation of the aggregate variable ΔMC_t, but which observes the *histories* of $\{\Delta MC_{s,t}\}$ for

[5] Although it is important that $\alpha(L)$ and $\gamma(L)$ differ, we assume that the lag lengths of the two components are the same. There is nothing restricting or substantive in this assumption. If the lag lengths were different they could be made to be the same by adding the required number of zero coefficients to the component with the lower lag length.

$s = 1, 2, ..j, ..S$ (where S is the total number of firms), in effect observes the *separate* histories of $\{\varepsilon_t\}$ and $\{u_{j,t}\}$ since the average value of $\{\Delta MC_{s,t}\}$ is ΔMC_t. From this alternative perspective, a conventional rational expectation of $\Delta MC_{j,t+k}$ can be seen as one formed by a firm which has been willing to incur the costs of acquiring the histories of $\{\Delta MC_{s,t}\}$ for $s = 1, 2, ..j, ..S$.

Consider now the case of firm j which has found it optimal *not* to acquire this complete information set, but which has instead acquired the histories of R such series, one of which is its own, where $1 \leq R < S$. Each of the R observed histories, *taken on its own*, amounts to an observation on $\{d + \alpha(L)\varepsilon_t + \gamma(L)u_{s,t}\}$ from $\{\Delta MC_{s,t}\}$, where $\Delta MC_{s,t}$ is one of the R series observed by firm j.[6] However, *taken together*, they also provide an observation on the history of the mean series, $\left\{\overline{\Delta MC}_{R,j,t}\right\}$, where $\overline{\Delta MC}_{R,j,t} \equiv \frac{1}{R}\sum_{s=1}^{R} \Delta MC_{s,t}$. This mean series can be written as,

$$\overline{\Delta MC}_{R,j,t} = d + \alpha(L)\varepsilon_t + \gamma(L)\overline{u}_{R,j,t} \tag{6}$$

where $\overline{u}_{R,j,t} \equiv \frac{1}{R}\sum_{s=1}^{R} u_{s,t}$; and the variance of $\overline{u}_{R,j,t}$ is $\sigma^2_{u_R}\left(=\frac{1}{R}\sigma^2_u\right)$.

The sum of two MA processes can be reparameterised as a single MA process.[7] So equation (6) can be rewritten as,

$$\overline{\Delta MC}_{R,j,t} = d + \rho_R(L)\overline{\eta}_{R,j,t} \tag{7}$$

where $\overline{\eta}_{R,j,t}$ is white noise by construction; $\rho_R(L)$ is a polynomial in the lag operator with $\rho_{R,0} = 1$; and the variance of $\overline{\eta}_{R,j,t}$, $\sigma^2_{\eta_R}$, and the elements of $\rho_R(L)$s are functions of the αs, the γs, and the two variances, σ^2_ε and $\sigma^2_{u_R}\left(=\frac{1}{R}\sigma^2_u\right)$.

From equations (6) and (7) it follows that, for this firm, $\Delta MC_{j,t}$ can be represented as the sum of two *separately observed* white-noise error processes,

$$\begin{aligned}\Delta MC_{j,t} &= d + \rho_R(L)\overline{\eta}_{R,j,t} + \gamma(L)(u_{j,t} - \overline{u}_{R,j,t}) \\ &= d + \sum_{i=0}^{T}\rho_{R,i}\overline{\eta}_{R,j,t-i} + \sum_{i=0}^{T}\gamma_i(u_{j,t-i} - \overline{u}_{R,j,t-i})\end{aligned} \tag{8}$$

The history of $\left\{\overline{\eta}_{R,j,t}\right\}$ is observed from the *joint* observation of R series; and the *separate* history of $\{u_{j,t} - \overline{u}_{R,j,t}\}$ is the *extra* information obtained from observing the *separate* history of the jth series.

It follows from equation (8) that a firm observing the R *separate* histories $\{\Delta MC_{s,t}\}$, and fully exploiting that information, will form expectations of

[6]We assume that σ^2_u and the parameters d, α_i and γ_i are common to all firms.

[7]See, for example, Hamilton (1994, pp. 102-107). Note that the ρ_Rs will be common to all firms provided the γs and the variance of $u_{s,t}$ are common to all firms, which we shall assume they are.

the current and future values of $\Delta MC_{j,t}$ as follows:

$$\begin{aligned}
\text{If } 0 &\leq k < T: \\
E^R_{j,t}\Delta MC_{j,t+k} &= d + \sum_{i=k}^{T} \rho_{R,i}\overline{\eta}_{R,j,t+k-i} + \sum_{i=k}^{T} \gamma_i (u_{j,t+k-i} - \overline{u}_{R,j,t+k-i}) \\
\text{Otherwise} &: E^R_{j,t} MC_{j,t+k} = d \qquad (9)
\end{aligned}$$

where the notation $E^R_{j,t}$ defines firm j's expectations conditioned on an information set consisting of R histories. Once again, 'fully exploits' means that no element in the histories of, in this case, $\left\{\overline{\eta}_{R,j,t}\right\}$ and $\{u_{j,t} - \overline{u}_{R,j,t}\}$ can be used to reduce the firm's forecast error.[8]

If, as appears plausible, the marginal benefits of greater forecast accuracy are decreasing, whereas the extra costs of enlarging the information set are increasing, it is likely to be optimal for firms *not* to acquire the complete information set.[9] Furthermore, RE is an empirically detectable, special case of ORE. Consequently, we shall develop the model on the assumption that firms condition their expectations on an incomplete information set.

With this assumption it follows that we can write $p^*_{j,t}$ as,

$$p^*_{j,t} = \mu^* + MC_{j,t} + \sum_{i=0}^{\infty} k_i \overline{\eta}_{R,j,t-i} + \sum_{i=0}^{\infty} k^\gamma_i (u_{j,t+k-i} - \overline{u}_{R,j,t+k-i}) \qquad (10)$$

where μ^* is a constant, $k_0 = \sum_{i=1}^{\infty}(\beta\theta)^i \rho_{R,i}$; $k_1 = \sum_{i=1}^{\infty}(\beta\theta)^i \rho_{R,i+1}$; .....; $k_n = \sum_{i=1}^{\infty}(\beta\theta)^i \rho_{R,i+n}$; $k^\gamma_0 = \sum_{i=1}^{\infty}(\beta\theta)^i \gamma_i$; $k^\gamma_1 = \sum_{i=1}^{\infty}(\beta\theta)^i \gamma_{i+1}$; .....; $k^\gamma_n = \sum_{i=1}^{\infty}(\beta\theta)^i \gamma_{i+n}$

We shall assume that a sufficiently large number of firms reset their prices each period to justify the assumption that the average value of u_{t-i+k} across those firms will be zero.[10] We also shall assume that any other firm has an equal probability of being included in any particular firm's information set so that the average value of $\overline{u}_{R,j,t+k-i}$ across all firms is also zero.[11]

[8] It is relatively straightforward to show that the forecast error from equation (9) is white noise. See Demery and Duck (2002).

[9] As mentioned earlier, firms could in principle distinguish between the common and idiosyncratic histories by accessing published data on *aggregate* marginal cost. Demery and Duck (2001) justify the assumption that firms do not use such information on two grounds. First, that, whilst, in principle, information about aggregate behaviour is relatively cheaply available in official or non-official sources, in practice there are the usual publication delays and revisions which make current and recent aggregate data unavailable or unreliable. And second, they carry out a series of simulation experiments to gauge the loss of profits a firm would incur by basing its expectations on limited rather than full information. They find that, in general, these profit losses are likely to be very small.

[10] We are also assuming that $u_{j,t}$ is independent across firms and has a finite variance.

[11] In reality, it is more likely that fims will 'network' and that the probability any other series has of being included in firm j's information set will depend upon (say) its geographical or industrial proximity to firm j. In this 'networking' case, groups of firms will share the same information set, each being informed of the others' $\Delta MC_{s,t}$ and their histories.

The *average* price set by those firms who are resetting can therefore be written as,

$$p_t^* = \mu^* + MC_t + \sum_{i=0}^{\infty} k_i \overline{\eta}_{R,t-i} \tag{11}$$

where $\overline{\eta}_{R,t-i}$ is the average value of $\overline{\eta}_{R,j,t-i}$ over the firms resetting in period t.

From equations (5) and (8) and our assumptions that u_{t-i+k} and $\overline{u}_{R,j,t+k-i}$ sum to zero across firms who are resetting their price, we can write,

$$\alpha(L)\varepsilon_t = \rho_R(L)\overline{\eta}_{R,t} \tag{12}$$

We define the current price level as a weighted average of the prices of those firms which are resetting and those which are not. Since all previous prices have the same probability of being reset, the current price level can be seen as a weighted sum of the average prices of those resetting and the average price level in the previous period. Formally,

$$p_t = (1-\theta)p_t^* + \theta p_{t-1} \tag{13}$$

From equations (11) and (13) we can derive the inflation rate ($\pi_t \equiv p_t - p_{t-1}$),

$$\pi_t = \theta\pi_{t-1} + (1-\theta)[\Delta MC_t + k_0\overline{\eta}_{R,t} + \sum_{i=1}^{\infty}((k_i - k_{i-1})\overline{\eta}_{R,t-i})] \tag{14}$$

and hence,

$$\pi_t = \theta\pi_{t-1} + (1-\theta)\sum_{i=0}^{\infty} k_i^* \overline{\eta}_{R,t-i} \tag{15}$$

where $k_0^* = \rho_{R,0} + k_0$; $k_1^* = \rho_{R,1} + k_1 - k_0$; $k_2^* = \rho_{R,2} + k_2 - k_1; ....; k_n^* = \rho_{R,n} + k_n - k_{n-1}$.

A more informative way of representing our model can be derived by first re-writing equation (12) as,

$$\overline{\eta}_{R,t} = \frac{\alpha(L)}{\rho_R(L)}\varepsilon_t \tag{16}$$

Substituting equation (16) into equation (15) we obtain,

$$\pi_t = \theta\pi_{t-1} + (1-\theta)\sum_{i=0}^{\infty} k_i^* \frac{\alpha(L)}{\rho_R(L)}\varepsilon_{t-i} \tag{17}$$

Such informal channels are likely to be important sources of low-cost information and often integral to the firm's participation in economic activity. In the interests of simplicity we assume that each firm obtains information from a random draw of $R-1$ realisations of $\Delta MC_{s,t}$. This means that no two firms will share the same information set other than by chance.

which we can re-write as,

$$\begin{aligned}\pi_t &= [\theta - \rho_{R,1}]\pi_{t-1} + [\theta\rho_{R,1} - \rho_{R,2}]\pi_{t-2} + [\theta\rho_{R,2} - \rho_{R,3}]\pi_{t-3} + ... \\ &\quad +\theta\rho_{R,n}\pi_{t-(n+1)} + (1-\theta)\sum_{i=0}^{n} k_i^* \Delta MC_{t-i} \qquad (18)\end{aligned}$$

If we had made the normal rational expectations assumption - that each firm is fully informed about the separate realisations of ε_t and $u_{j,t}$ - then $\frac{\alpha(L)}{\rho_R(L)}$ would equal 1, $\eta_{R,t}$ would equal ε_t and equation (18) could be written as,

$$\pi_t = \theta\pi_{t-1} + (1-\theta)\sum_{i=0}^{\infty} k_i^* \varepsilon_{t-i} \qquad (19)$$

A comparison of equations (18) with (19) shows that a major implication of ORE is that inflation is likely to be more heavily influenced by lags in inflation than the standard RE model implies.[12] The intuition for this is that with less knowledge of the precise nature of the shocks affecting their marginal costs, firms will be unable to react to them as accurately and promptly as they would if they were fully informed. This, of course, is not the result of any irrationality but of their deliberate choice to trade off some lower forecast accuracy for lower information costs.

In deriving equation (18) we followed the convention of treating the sticky price parameter θ as fixed. Casual theorising suggests that this is at best a convenient approximation. The probability that a firm will re-set its price - a convenient analytic way of representing those factors responsible for price stickiness (e.g. 'menu costs') - is likely to be different in an economy experiencing zero inflation from one experiencing substantially higher inflation.

A more formal theoretical basis for the prediction that θ will be a function of an economy's inflation rate is developed in Ball, Mankiw and Romer (1988). In their model, imperfectly competitive firms face a fixed cost of changing their prices but can select the length of the interval between price changes. Maximisation of expected discounted current and future profits implies that the optimal interval will be a function of the average inflation rate. High inflation causes a firm's profit-maximising nominal price to change rapidly, which raises the benefits from frequent adjustment.[13] In our empirical work in the next section, we build on this work and that of Hamilton (1990) by allowing the economy to experience different inflation regimes

[12] It is relatively straightforward to rewrite the more conventional New-Keynesian model - equation (19) - in its more familiar form, $\pi_t = \beta E_t \pi_{t+1} + \lambda(MC_t - p_t)$ where λ ($= \frac{(1-\theta)(1-\beta\theta)}{\theta}$) is positive.

[13] Ball, *et al.* (1988) also find that it will decrease the greater the variance of aggregate and firm-specific shocks. 'When either variance is large, a firm's future profit-maximising price is highly uncertain, so the firm does not wish to fix its price for long' (p. 25).

and by allowing θ to be different in the different regimes. As is clear from equations (18) different values for θ will imply different inflation dynamics.

In the next section we consider how (18) can be tested and report the results of testing it using UK data.

3 Data and Empirical Results

To test our model we estimate equation (18) by GMM and test the overidentifying restrictions it implies.[14] To see the nature of these restrictions, assume that we have determined the appropriate value of n in equation (18) and hence the order of the lag on π_t, and, by implication, on ΔMC_t. The number of reduced-form (unrestricted) coefficients to be estimated is then $2(n+1)$. Because the k_i^*s in equation (18) are themselves functions of β, θ, and the ρ_Rs, there are $n+2$ *structural* parameters. The model therefore has n overidentifying restrictions.

In equation (18) the parameter θ is treated as a constant. We wish to test for the influence of the economy's inflation regime on θ and therefore need to identify different inflation regimes. To do so we estimated a two-state Markov switching model.[15] The results, presented in Figure 1, suggested two low-inflation regimes, 1963Q2-1970Q4 and 1982Q1-2000Q4 with a mean inflation rate of 1.02% per quarter, and one high-inflation regime, 1971Q1-1981Q4 with a mean inflation rate of 3.28% per quarter. Sample size limitations[16] prevented estimation of our model over the three separate periods and we therefore adopted two estimation strategies. In the first, we estimated and tested the model *separately* over the two periods 1963Q2-1981Q4 and 1982Q1-2000Q4. The first period we view as the 'high-inflation regime' and the second period the 'low-inflation regime'. We then compared the estimates of θ for the two periods. Because the first of these two subsamples covers an initial epsiode of low inflation, in the second approach we estimated the model over a 15-year moving window, starting with the period 1963Q2-1978Q1 and ending with the period 1986Q1-2000Q4. This approach gives us a set of 91 estimates of each parameter, in particular of θ. From these estimates of θ we can derive an implied series for the mean length of time before a price is reset, and we can compare this with the behaviour of inflation.

Our data, full details of which are given in Appendix B, are from the UK and cover the private (non-government) sector. We adopt data definitions similar to those employed by Batini, Jackson and Nickell (2000). The

[14] All versions of the model are estimated using weights based on a consistent estimator of the asymptotic covariance matrix of the *unrestricted* model. This permits tests of the overidentifying restrictions we discuss below.

[15] Maximum likelihood estimates of the parameters were obtained using Hamilton's (1990) EM algorithm.

[16] The initial low-inflation period covers only 29 quarters.

inflation rate (π) is defined as the quarterly change in the log of the overall GDP price deflator.[17] A series for unit labour costs was constructed by taking the ratio of nominal non-government compensation of employees to real non-government GDP. The log of this ratio defines our variable MC. We adjust the published compensation estimates to include a labour income component of the income of the self-employed.[18]

The adjustment we make to employee compensation implies that the average return to labour of the self-employed is equal to the average remuneration of employees in employment. Self-employment income is not separately identified in the UK accounts[19] so we follow the procedure used by Batini *et al.* (2000), who adjust compensation by the ratio of total employment to the number of employees. The imputation of labour income of the self-employed is particularly important given the growing importance of these sectors: the proportion of self-employment to total employment rose in the UK from around 8% in 1960 to 13% in 2000.

We cannot reject the null hypothesis that the two series - π and ΔMC - are stationarity, the ADF test statistics of -2.921 and -4.018 (respectively) being significant at the 5% level (assuming one lagged term in both cases)[20].

A number of econometric issues arise in the estimation of equation (14). The presence of the contemporaneous term ΔMC requires the use of an instrumental variable estimator. The error term added to (14) may also be serially correlated or heteroskedastic. A number of potential sources for this error suggest that the errors may be serially correlated in (14). The markup parameter (μ) may be subject to random variation; there may be random departures from the optimal price (1); equation (1) is also a logarithmic approximation. These terms would appear in difference form in equation (14): and since we do not know their statistical properties in equation (1), they may introduce serial correlation of unkown form in equation (14). For these reasons we estimate this equation by GMM and, in Table 1 we report the estimates of equation (18) for the two assumed inflation regimes (together with our full-sample estimates).[21] The lag length, n, for the three periods

[17] The prices are basic prices. As Batini *et al* point out, the use of basic prices means that value added is measured *net of indirect taxes*, which is theoretically more appropriate than measures in market prices. It was not possible to construct the non-government GDP deflator due to the lack of a constant price government value added series.

[18] This procedure is adopted by Batini, Jackson and Nickell (2000). It has been used in other contexts when calculating aggregate labour income (see for example Blinder and Deaton (1985)).

[19] The income of the self-employed is now consolidated with other incomes in an 'Other Income' category.

[20] The lag length was determined by truncating at the last significant t-statistic.

[21] Using the full sample and first sub-sample the instruments used were as follows: lags 1 to 6 in π, lags 1 to 5 in ΔMC, and lags 1 to 2 in both the output gap (defined as de-trended log output) and wage inflation. For the second sub-sample the additional lags required the addition of a seven-quarter lag in π and a six-quarter lag in ΔMC.

was determined by the last significant ρ_R.[22] For all three periods the diagnostic statistics are satisfactory though there is a hint of higher-order serial correlation in the high inflation regime. We also computed the Newey-West (1987) 'D' test statistic of the model's overidentifying restrictions, a test which is analogous to the likelihood ratio test.[23] The statistic is distributed as chi–square with degrees of freedom equal to the number of restrictions imposed. For all three periods, the null hypothesis that the restrictions are valid cannot be rejected at conventional significance levels. The estimates of β are, for all periods, close to 1 though the point estimate for the low inflation regime is clearly too high.

The estimates of the ρ_R parameters for the three periods all show the same pattern: they are all positive, highest at the first order and then declining, and in all periods the sum of the estimates of the ρ_R parameters is significantly different from zero. However, in the case of the low inflation regime, both the sum of the coefficients on lagged inflation is notably higher, and there are more significant higher-order lags. Both these characteristics suggest that inflation exhibits greater sluggishness in a low inflation regime, over and above any effect due to changes in θ. One possibility is that this greater sluggishness may be due to a lengthening of the period covered by wage contracts.

The three estimates of θ show the pattern we would predict: the lowest estimate is found for the high inflation regime, and the highest estimate for the low inflation regime. For the high inflation regime the estimate of θ suggest a mean interval of just under a year before prices can be expected to be reset - an interval lower than, but not too dissimilar to that found for the whole period. The low inflation regime suggests a mean interval of more than two years. All our estimates of θ are higher than that suggested in Hall *et al.*'s (1997), survey of 654 UK companies. They found that in the year to September 1995 (in a low-inflation period) the median number of times that prices were changed was twice a year, which suggests a value of θ of 0.5. This difference may be due to the sample used in Hall *et al.* which substantially over-represents large firms, though the direction of bias this might introduce into the estimate of θ is not obvious. Taylor (1999) reviewed the direct evidence on the frequency of price changes in the US and concluded that 'price changes and wage changes have about the same average frequency - about one year' (p.1020), implying a value of θ of around 0.75 which is noticeably closer to our estimate.

[22] Using data over the period 1963Q2-1981Q4, we obtained an estimate of $\rho_{R,4}$ of 0.175 with a standard error of 0.109, so for this period we truncate at $n = 3$. For the period 1982Q1-2000Q4 we obtained an estimate of $\rho_{R,7}$ of 0.132 with a standard error of 0.131 and we truncate the lags at $n = 6$.

[23] See Newey and West (1987) p.780, equation (2.9). The test statistic D requires that the same estimate of the covariance matrix is used in both the restricted and unrestricted models as this ensures that $D > 0$.

The results we report in Table 1 are therefore consistent with the new-Keynesian model represented by equation (18) and with the suggestion that the dynamics of inflation are different in different inflation regimes. To illustrate the effects of the different regimes, we present in Figure 2, the impulse response of inflation to an unexpected and permanent jump in the value of ΔMC implied by the estimates for the high and low inflation regimes presented in the second two columns of Table 1. In the high inflation regime inflation moves relatively quickly and smoothly to its new equilibrium level - within approximately 5 periods the inflation rate is within 20% of it. In the low inflation regime the approach is slower and more erratic - even after 10 periods inflation has not moved to within 25% of its new equilibrium.

In Figures 3 to 6 we present selected results of estimating the model using a 15-year moving window.[24] Figure 3 presents the estimates of θ, together with 5% confidence bands and the implied mean interval between price changes. Figure 4 presents estimates of the sum of the first three ρ_Rs and associated 5% confidence intervals. Figure 5 presents estimates of the sum of the second three ρ_Rs and associated 5% confidence intervals. Figure 3 suggests that from the late 1970s or early 1980s there has been a steady rise in θ and the implied mean interval between price changes. The estimated value of θ from roughly 1983 onwards is well above the higher 95% confidence limit estimated from earlier data. Figures 4 and 5 suggest that there has been a similar rise in the sum of the first and second triplets of ρ_Rs beginning at roughly the same time. Here too the sum of the two sets of coefficients tend, by the end of the period, to be above their respective higher 95% confidence limits estimated from earlier data. In Figure 6 we plot the estimated mean interval between price changes in quarters ('contract length') $T = \left(\frac{1}{1-\widehat{\theta}}\right)$ against the mean quarterly inflation rate (%), $\overline{\pi}$, over observations in the moving window. The negative relationship between T and $\overline{\pi}$ is clear from the figure. A simple regression of the estimated contract length (in quarters) against mean inflation (%) and its square gives (standard errors in parentheses),

$$\begin{aligned} \widehat{T} &= \underset{(0.755)}{12.922} - \underset{(0.877)}{5.169}\,\overline{\pi} + \underset{(0.229)}{0.614}\,\overline{\pi}^2 \qquad (20) \\ R^2 &= 0.899 \end{aligned}$$

This regression suggest that, as the mean inflation rate rises over the relevant range, the contract length shortens, but at a diminishing rate. The shortest estimated contract length we estimate is 2.6 quarters based on data centred on 1973Q4, when the mean inflation rate was 2.67% per quarter (around 11% per annum). The longest average contract length we estimate is 9.7 quarters based on data centred on 1990Q1 with a mean inflation rate of 1.08% per quarter (4.4% per annum). Using equation (20), we can calculate

[24] For the moving window regressions we allow for six ρ_R terms.

the implied contract lengths for the two mean inflation rates estimated using the two-regime Markov switching model: a low-inflation mean of 1.016% and a high-inflation mean of 3.284%. The respective contract lengths are 8.3 and 2.6 quarters (approximately two years and eight months respectively).[25] Evaluating equation (20) for the overall sample mean inflation rate of 1.66% per quarter, implies a contract length of 6.2 quarters. The difference between this estimate and that reported in the first column of Table 1 suggests that important information can be lost when analysing inflationary dynamics if one ignores differences in the underlying inflationary regimes. Our results also suggest that the direct evidence reported by Hall *et al.* (1997) underestimates the mean duration of average price contracts in the UK.

4 Summary and Conclusions

Mankiw (2001) has described the new-Keynesian model of inflation as 'ultimately a failure': it simply cannot explain the dynamics of inflation - in particular the observed influence of lags of inflation on current inflation - without recourse to theoretically awkward, *ad hoc* assumptions such as the existence of a fraction of firms who set prices by some rule of thumb. In this paper we have modified the new-Keynesian model in two main ways to overcome this failure. Both modifications are theoretically well-grounded. In the first we have allowed firms optimally to select the information set upon which they form their expectations. Their expectations are still rational in the sense that they are based on a full exploitation of the firms' information set, and they are optimal in the sense that in deciding the content of that information set firms are assumed to weigh up the non-zero costs and benefits associated with enlarging it.

The second modification appeals to a theoretical literature which suggests that the frequency of price setting will be a function of the economy's inflation rate, specifically that in high inflation economies prices will be reset more frequently. The resultant model suggests that lagged inflation rates will have a greater influence on current inflation than the standard new-Keynesian model suggests. We find that we cannot reject the formal restrictions implied by the model (using UK quarterly data) and we estimate that the mean time before prices were reset was around eight months during the 'high' inflationary regime and approximately two years when mean inflation was 'low'.

[25] We also estimated equation (20) with $\widehat{\theta}$s as the regressor and obtained estimates of the 'low' and 'high' inflation regime contract lengths from the predicted θs. Using this approach we estimate the contract length for the low inflation regime to be 8.0 quarters and that for the high inflation regime to be 2.4 quarters. These are very similar to those based on regressions involving T.

Table 1: Imperfect Information Model
Equation (18)

	1963Q2-2000Q4	1963Q2-1981Q4	1982Q1-2000Q4
	n=5	n=3	n=6
Constant	0.000	-0.001	-0.002
	(0.001)	(0.001)	(0.001)
$\widehat{\theta}$	0.774	0.729	0.903
	(0.028)	(0.037)	(0.027)
$\widehat{\rho}_{R,1}$	0.794	0.748	1.076
	(0.089)	(0.078)	(0.116)
$\widehat{\rho}_{R,2}$	0.508	0.440	0.866
	(0.089)	(0.098)	(0.143)
$\widehat{\rho}_{R,3}$	0.388	0.202	0.730
	(0.111)	(0.062)	(0.146)
$\widehat{\rho}_{R,4}$	0.226	-	0.467
	(0.097)		(0.145)
$\widehat{\rho}_{R,5}$	0.171	-	0.532
	(0.083)		(0.112)
$\widehat{\rho}_{R,6}$	-	-	0.360
			(0.089)
$\widehat{\beta}$	0.926	0.952	1.226
	(0.165)	(0.177)	(0.095)
$\sum \widehat{\rho}_{R,i}$	2.088	1.390	4.032
	(0.395)	(0.187)	(0.678)
$Q(4)$	0.789	0.865	0.115
$Q(8)$	0.550	0.044	0.367
$p(J)$	0.852	0.500	0.155
$p(D)$	0.876	0.123	0.118
T	4.415	3.690	10.279

Notes: Estimated standard errors in (.) with a Newey-West correction.
$Q(n)$ is the p-value of the Ljung-Box test for nth order serial correlation.
$p(J)$ is the p-value of the Hansen test of over-identifying restrictions.
$p(D)$ is the p-value of the Newey-West (1987) test of the model's restrictions.
$T \equiv \frac{1}{1-\widehat{\theta}}$ is the expected duration of prices.

References

Balakrishnan, Ravi and López-Salido, David (2000), 'Understanding UK inflation: the role of structural change', *Mimeo, Bank of England and Bank of Spain.*

Balakrishnan, Ravi and López-Salido, David (2001), 'Understanding UK inflation: the role of openness', *Mimeo, Bank of England and Bank of Spain.*

Ball, Laurence(2000), 'Near-Rationality and inflation in two monetary regimes', *NBER Working Paper 7988.*

Ball, Laurence, Mankiw, Gregory N. and Romer, David (1988), The new-Keynesian economics and the output-inflation trade-off' *Brookings Papers on Economic Activity*, 1988:1, pp 1-65; reprinted in *new-Keynesian Economics, Volume 1: Imperfect Competition and Sticky Prices*, N.G. Mankiw and D. Romer, eds., (Cambridge, MA: MIT Press).

Batini, Nicoletta, Jackson, Brian and Nickell, Stephen (2000), 'Inflation and the labour share in the UK', *Bank of England External MPC Discussion Paper No. 2.*

Blinder, Alan S. and Deaton, Angus (1985), 'The time series consumption function revisited.' *Brookings Papers on Economic Activity* 2, 465-521.

Calvo, Guillermo A. (1983), 'Staggered prices in a utility maximizing framework', *Journal of Monetary Economics,* 12, 383-398.

Conlisk, John (1996) 'Why Bounded Rationality?', *Journal of Economic Literature*, 34(2), 669-700.

Crettez, B. and Michel, P. (1992) 'Economically rational expectations equilibrium', *Economics Letters* 40(2), 203-06.

Demery, David and Duck, Nigel W. (2001), 'The new-Keynesian Phillips curve and imperfect information', *University of Bristol Discussion Paper 01/528.*

Demery, David and Duck, Nigel W. (2002), 'Optimally rational expectations and macroeconomics', *University of Bristol Discussion Paper 02/533.*

Fuhrer, Jeffrey C., and Moore, George R. (1995), 'Inflation persistence', *Quarterly Journal of Economics*, 110, 127-159.

Galbraith, John W. (1988) 'Modelling expectations formation with measurement errors' *The Economic Journal*, 98, 412-428.

Galí, Jordi; Gertler, Mark and López-Salido, J. David (2001), 'European inflation dynamics', *European Economic Review* 45, 1237-1270.

Galí, Jordi and Gertler, Mark (1999), 'Inflation dynamics: a structural econometric analysis'. *Journal of Monetary Economics* 44, 195-222.

Hall, Simon, Walsh, Mark, and Yates, Anthony (1997) 'How do UK companies set prices?', *Bank of England Mimeo.*

Hamilton, James D. (1990)**, '**Analysis of time series subject to changes in regime' *Journal of Econometrics* 45, 39-70.

Hamilton, James D. *Time Series Analysis*, (1994), Princeton University Press, Princeton.

Mankiw, N. Gregory (2001), 'The inexorable and mysterious trade-off between inflation and unemployment', *Economic Journal* 111,C45-C61.

Mankiw, N. Gregory and Reis, Ricardo (2001), 'Sticky information versus sticky prices: a proposal to replace the new-Keynesian Phillips Curve', *Harvard University Mimeo.*

Newey, Whitney K. and West, Kenneth D. (1987)**,** 'Hypothesis testing with efficient method of moments estimation' *International Economic Review* 28,777-787.

Roberts, John M. (1997) 'Is inflation sticky?', *Journal of Monetary Economics,* 39 173-196.

Roberts, John M. (1998) 'Inflation expectations and the transmission of monetary policy', *Board of Governors of the Federal Reserve, Finance and Economics Discussion Series Paper No: 1998-43*

Taylor, **John B**. (1999) 'Staggered price and wage setting in macroeconomics', in John B. Taylor and Michael Woodford (eds.) *Handbook of Macroeconomics*, North Holland, Elsevier, 1009-1050.

Woodford, Michael (2001), 'Imperfect common knowledge and the effects of monetary policy' *Mimeo, Princeton University.*

Data Appendix

The raw data used in this paper can be downloaded from the following University of Bristol web site:

http://www.ecn.bris.ac.uk/www/ecdd/newk/newk.htm

The data were obtained from the National Statistics DataBank Online at http://www.data-archive.ac.uk/. The four-digit codes are the relevant National Statistics codes for the series used.

π is the inflation rate defined as the first difference in the logarithm of the GDP deflator: $\pi_t = \log(DEF_t) - \log(DEF_{t-1})$, where $DEF = \frac{\text{ABML}}{\text{ABMM}}$, ABML is Gross Value Added (average) in current basic prices, seasonally adjusted; and ABMM is Gross Value Added in 1995 basic prices, seasonally adjusted.

The logarithm of nominal marginal cost (MC) is defined as:

$$MC = \log\left[\frac{(\text{DTWM–NMXS}^a)\left(\frac{\text{DYZN+BCAJ}}{\text{BCAJ}}\right)}{\text{ABMM-}\left(\frac{\text{NMXV}^a\text{+NMXS}^a}{DEF}\right)}\right]$$

where DTWM is total compensation of employees (£m) seasonally adjusted; NMXS^a is the variable NMXS seasonally-adjusted (X11), where NMXS is compensation of employees in government seasonally unadjusted; similarly NMXV^a is the variable NMXV seasonally-adjusted, where NMXV is general government gross operating surplus; DYZN is the number of self-employed workforce jobs (000, seasonally adjusted); and BCAJ is the number of employee workforce jobs (000, seasonally adjusted). Prior to 1978, the two employment series were available for the second quarter in each year only, so for these years observations for other quarters were derived by linear interpolation. This definition of labour share follows the preferred definition adopted by Batini, Jackson and Nickell (2000). In the absence of a constant price series for government value added, we have assumed that the government value added deflator is the same as that for Gross Value Added. The growth in nominal marginal costs is defined as: $\Delta MC_t \equiv MC_t - MC_{t-1}$.

Real output (y) is ABMM, gross value added in 1995 basic prices, seasonally adjusted. The wage rate is defined as:

$$W = \frac{\text{DTWM}\left(\frac{\text{DYZN+BCAJ}}{\text{BCAJ}}\right)}{\text{DYZN+BCAJ}}$$

and wage inflation is defined as $\Delta w_t = \log(W_t) - \log(W_{t-1})$.

Figure 1
UK: Probability of Low Inflation Regime
[Low = 1.016% High=3.284% (per quarter)]

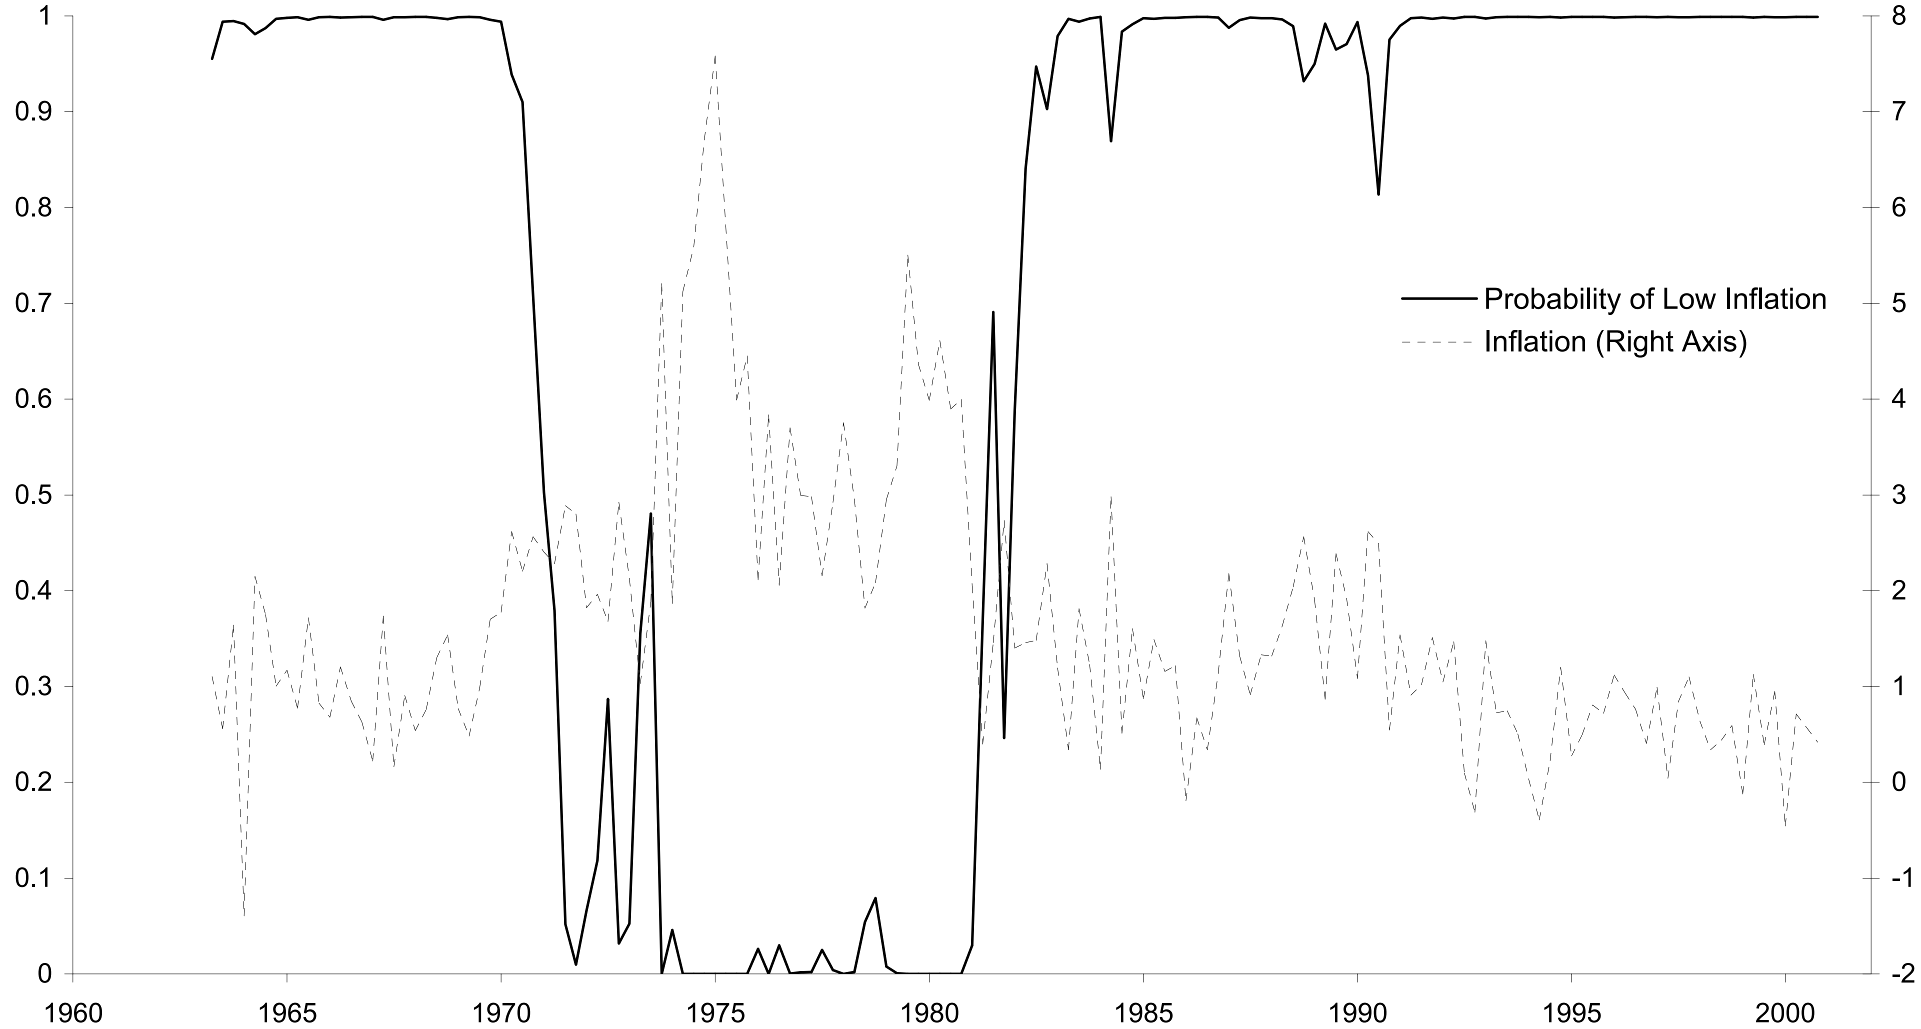

Figure 2
Impulse Response of inflation
[Permanent unit shock to the growth of nominal marginal cost]

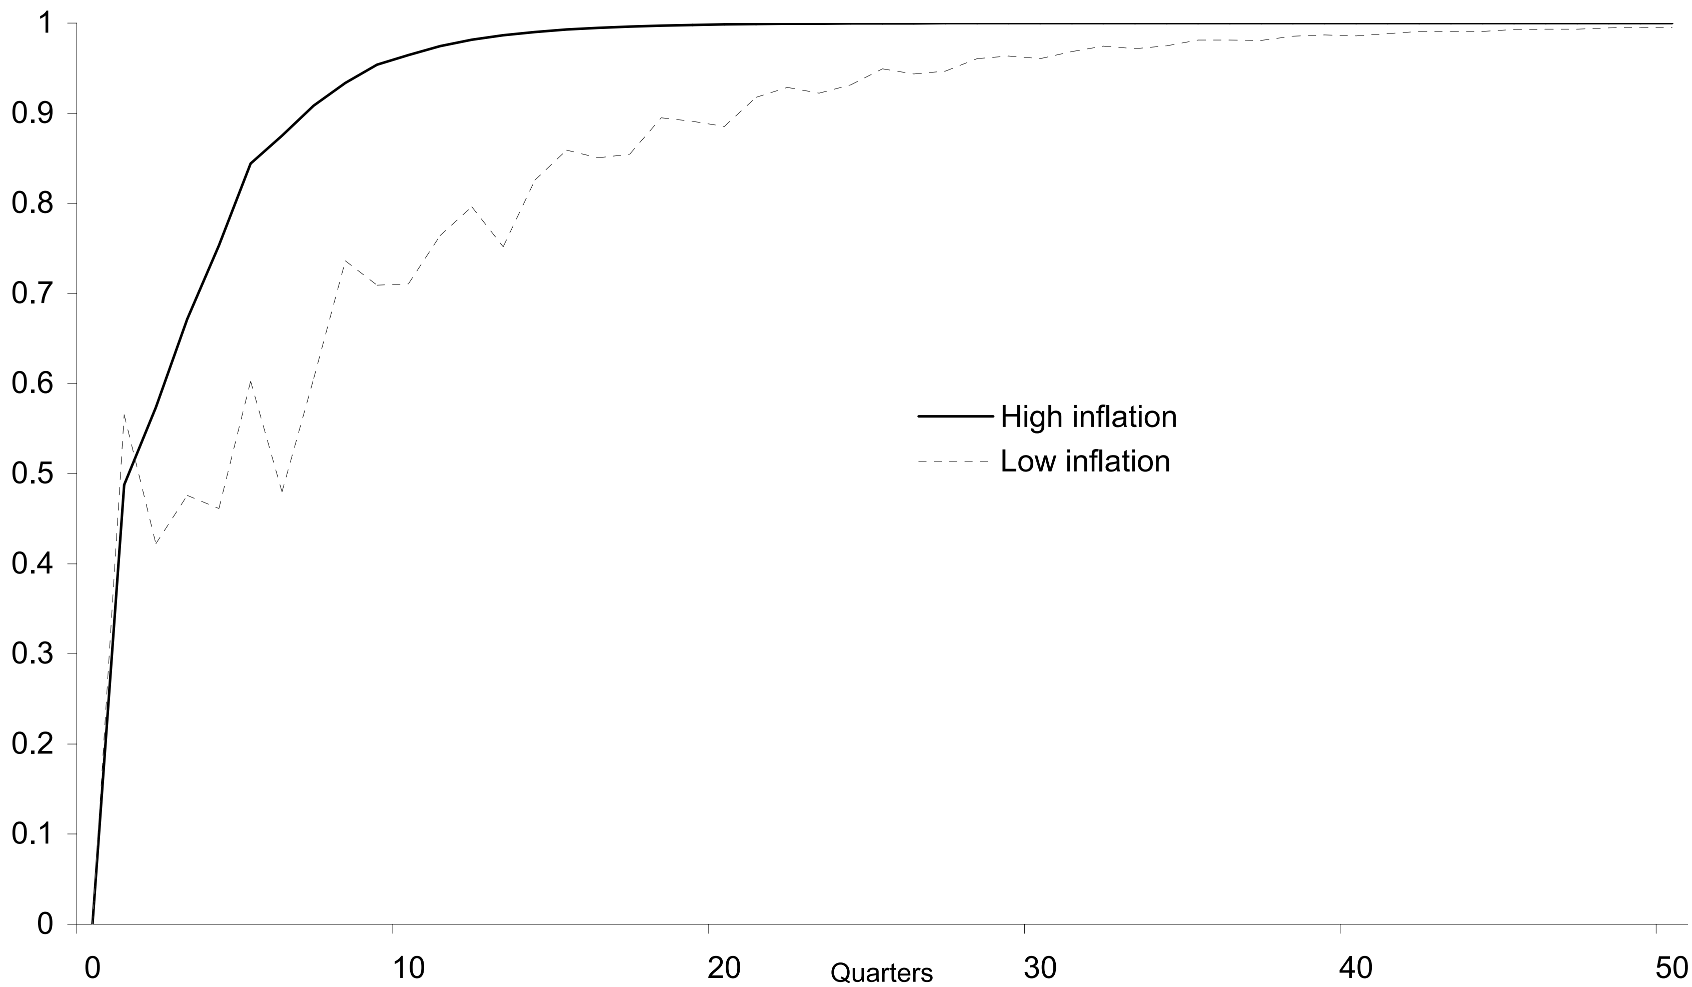

Figure 3
Moving Window Estimates of θ and implied Contract Length

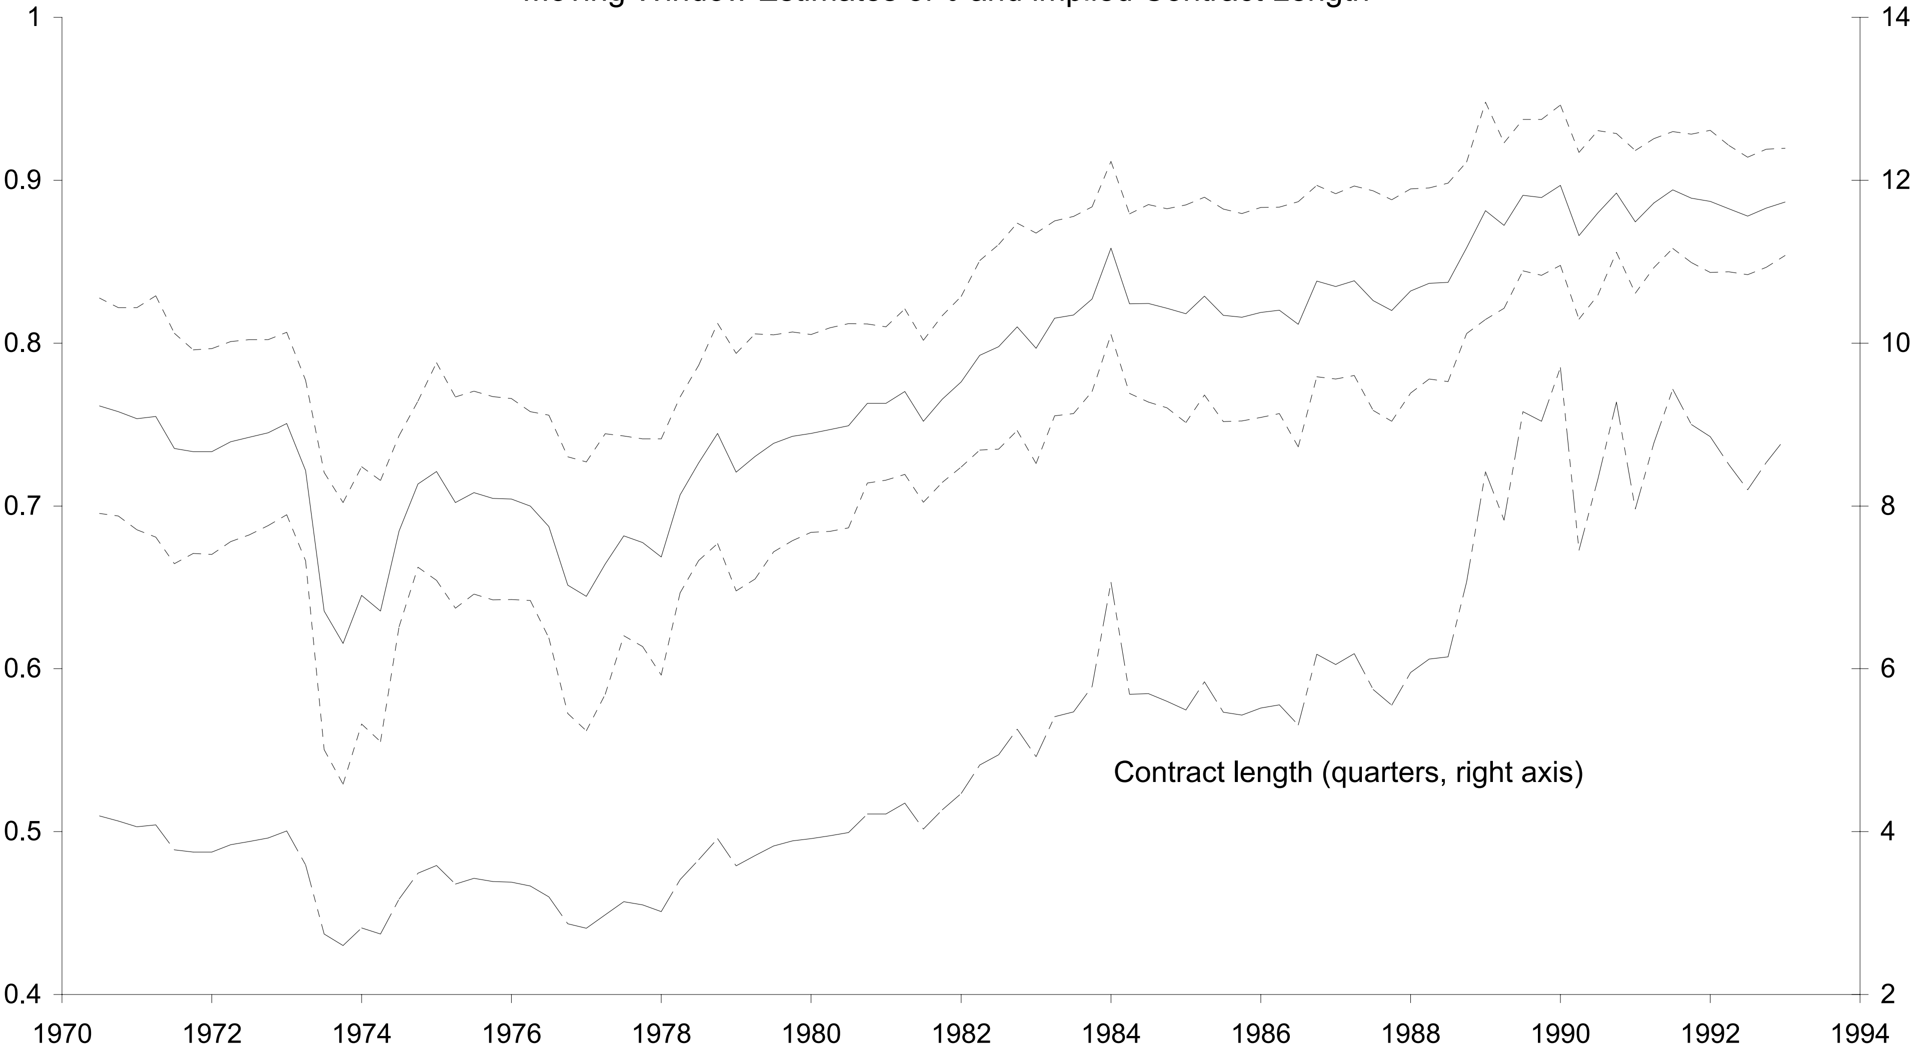

Figure 4
Sum of first three ρs and confidence intervals

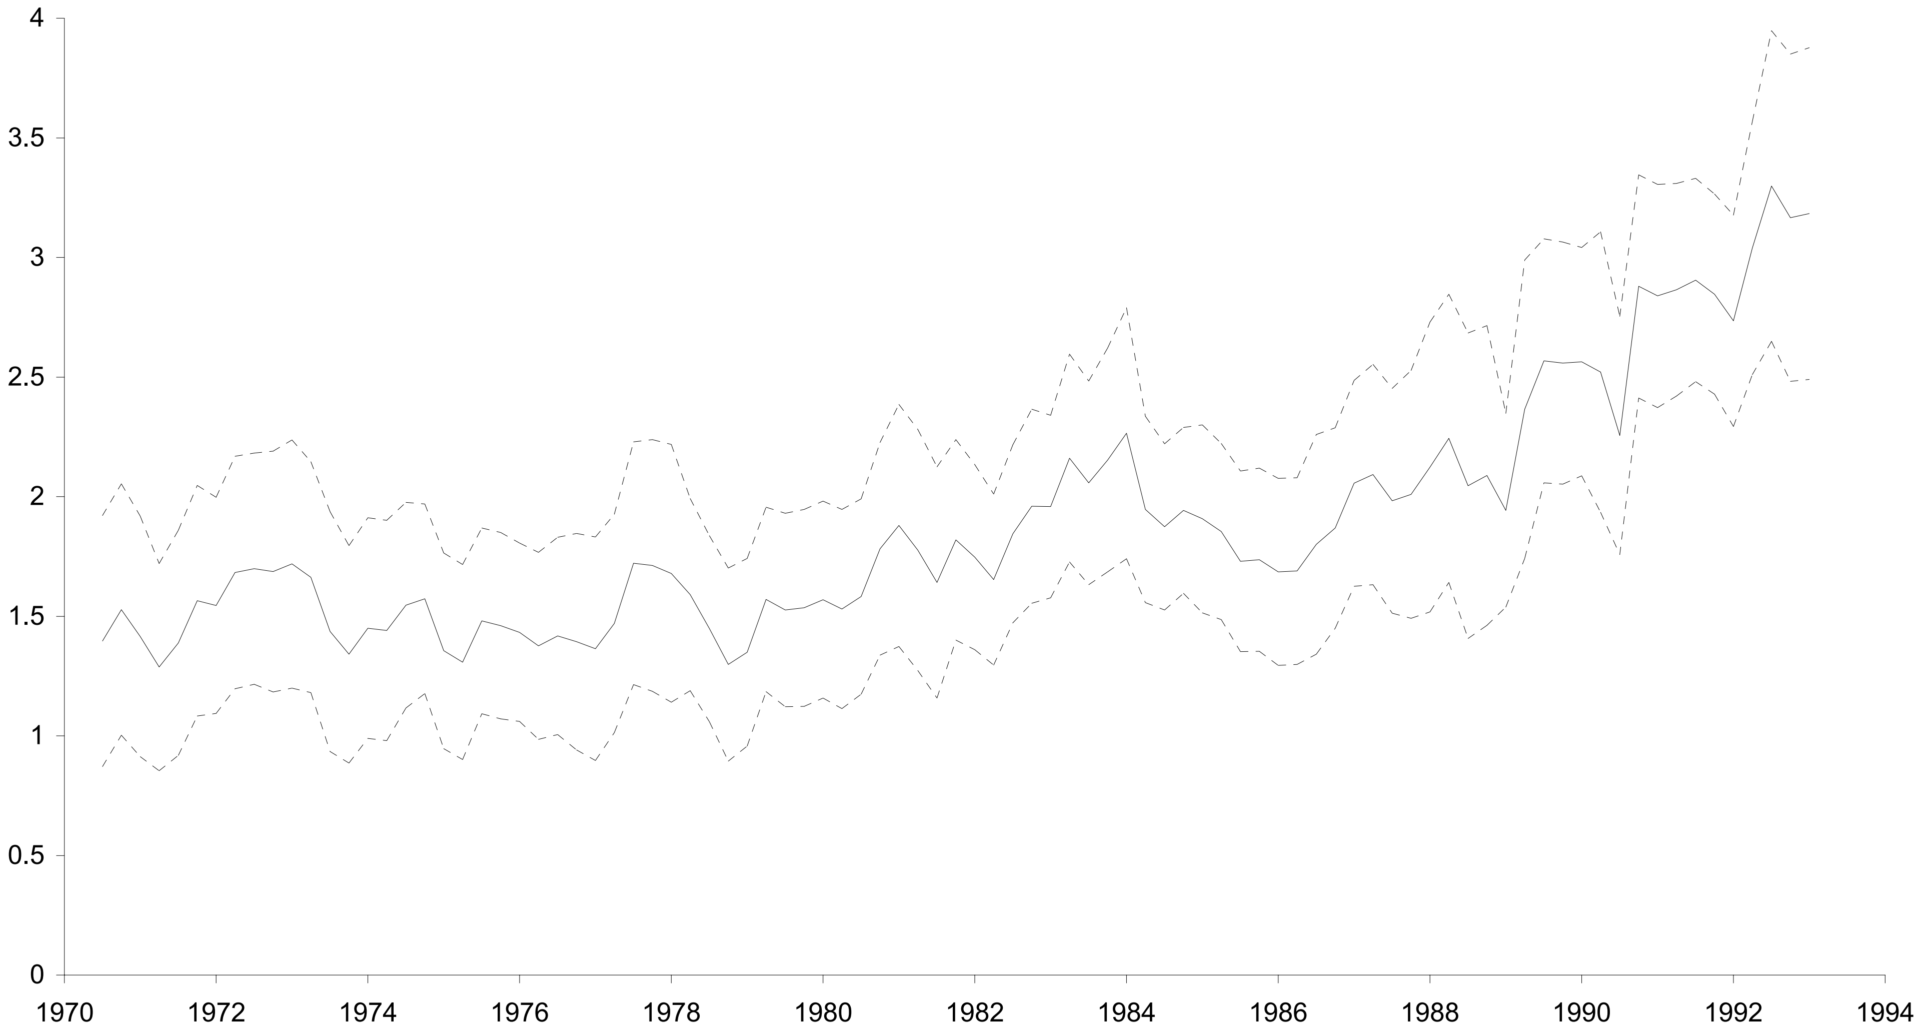

Figure 5

Sum of second three ρs and confidence intervals

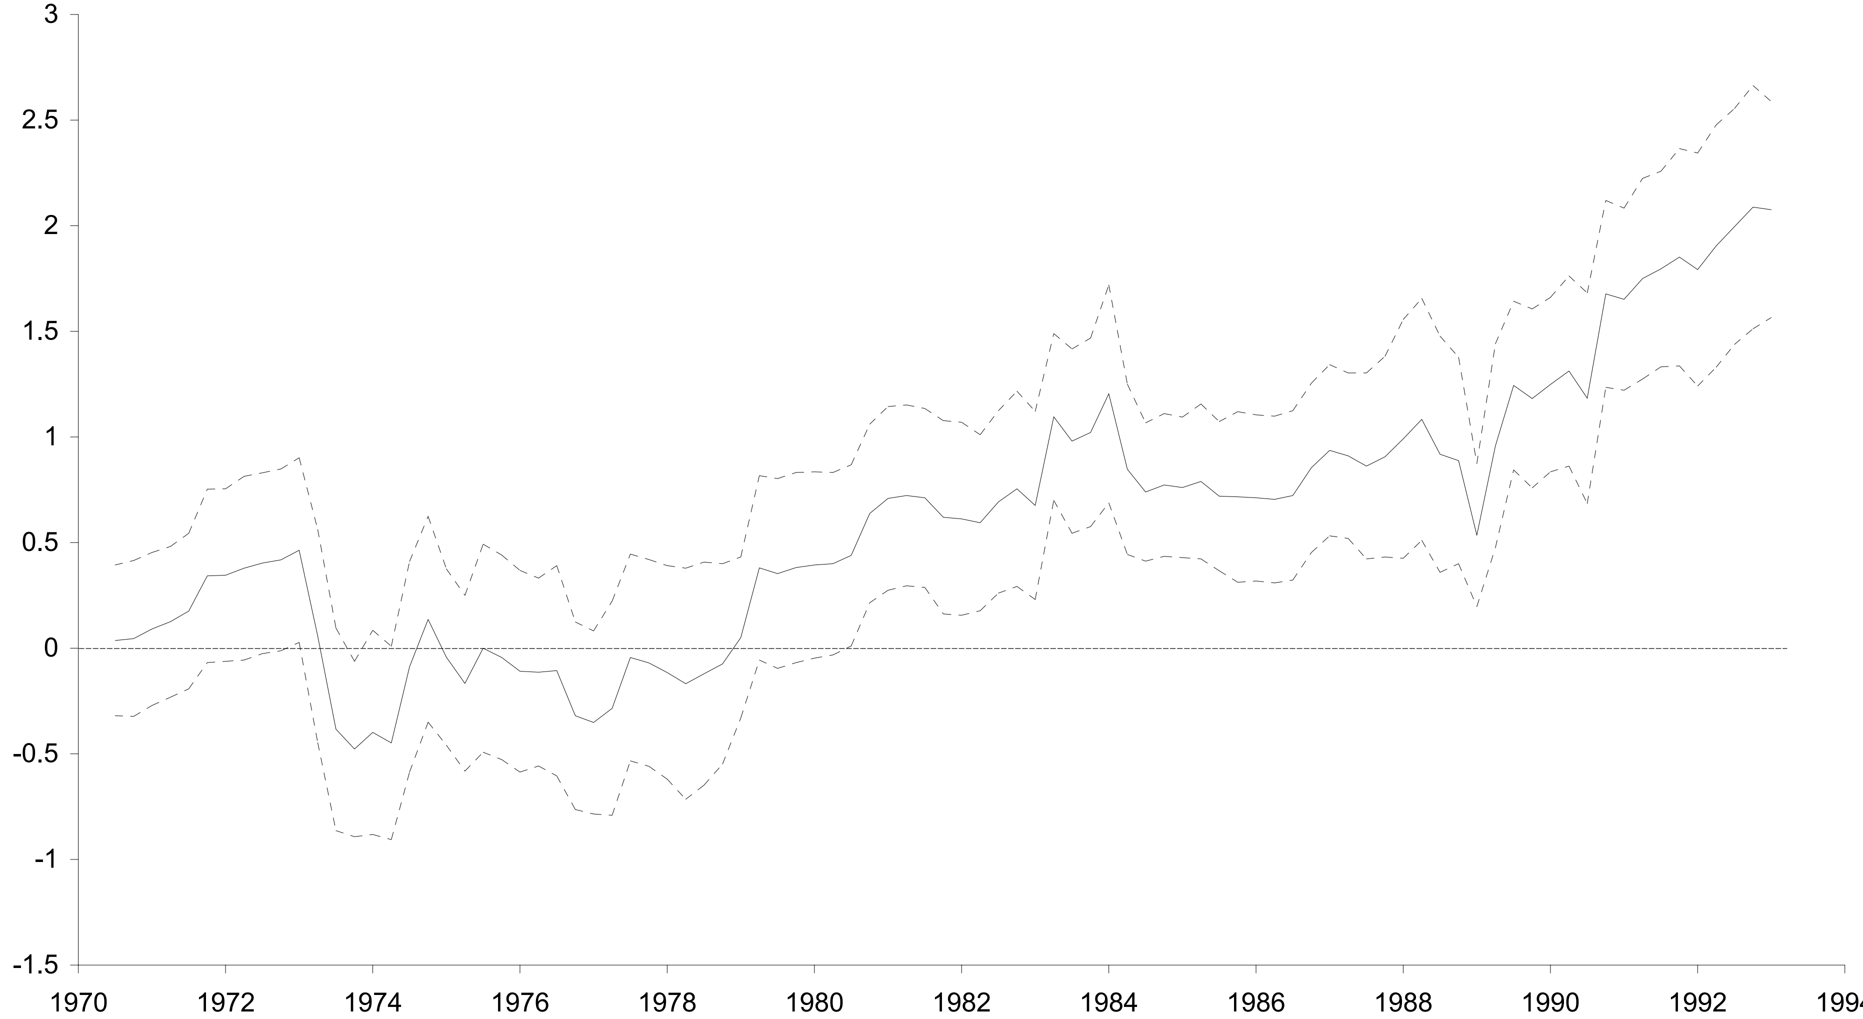

Figure 6
Contract Length and Mean Inflation
[Based on Moving Window Regressions]