Has the Inflation Process Changed?

Stephen Cecchetti and Guy Debelle

Brandeis University, and NBER; and Reserve Bank of Australia

Initial Draft: June 2004
This Revision: May 2005

Has the inflation process changed?

Stephen Cecchetti and Guy Debelle*

1. Introduction

Low inflation has been a fact of economic life in many countries now for at least a decade. While some have had low inflation since the early 1980s (at least), the 1990s was notable for the widespread incidence of low inflation, including amongst a large number of formerly high inflation economies in Latin America and Eastern Europe.[1]

On a number of occasions, the advent of low inflation has coincided with changes in the monetary policy framework, which in many cases has involved the adoption of a form of inflation targeting. But in other countries, there has been no marked change in the policy regime. This raises the important question of whether the nature of the inflation process has changed in a manner which has helped entrench low inflation.

One means by which this could occur would be if the inflation process has become less persistent; that is, a given shock to the price level which boosts inflation now has a smaller and/or less protracted impact on the rate of ongoing inflation. If the inflation process is less persistent, the task of monetary policy is easier. The sacrifice ratio (that is, the cost of reducing inflation in terms of lost output) will be lower. Moreover, with low inflation persistence and a flexible inflation target, inflation may simply fluctuate around the desired target level with little action required from the policy maker.

If inflation persistence has declined, why has it occurred? One possible explanation is that the decline in persistence itself is related to a change in the conduct of monetary policy. Taylor (1998, 2000) makes this general argument while Sargent (1999) provides a detailed account of the interaction between inflation persistence and the monetary framework in the United States. The case centres on the observation that over the past decade or so, monetary policy has been much more focussed on achieving low inflation, and less on exploiting short-run output gains. These policies have been successful, leading to an increase in the credibility of monetary policy. Increased credibility has, in turn, anchored inflation expectations at a low (and constant) rate of inflation– at the inflation target in those countries which have such a formal target. The dramatic consequence is that inflation expectations are unlikely to adjust to temporary increases in inflation. This reduces the persistence of shocks to both the price level and inflation.

---

* Professor of International Economics and Finance Brandeis University, and NBER; and Reserve Bank of Australia. We would like to thank Rebecca Edwards in particular, Ivailo Arsov and Dimitrios Karampatos for their invaluable hard work, the members of the Eurosystem's Inflation Persistence Network for their help, particularly Benoit Mojon, our discussants at the BIS Conference on Understanding Low Inflation and Deflation, Ignazio Angeloni and Jordi Gali, and Alexandra Heath, David Lebow and Karl Whelan for helpful comments. This is a preliminary version of a paper prepared for the 42nd Panel Meeting of ***Economic Policy*** in to be held in London in October 2005. The views expressed are those of the authors and not necessarily those of the RBA.

[1] According to the IMF's World Economic Outlook database, by 2002, Zimbabwe and Angola were the only countries with inflation over 100%. Over the previous two decades, 48 countries experienced triple-digit inflation for at least one year, with 23 countries having episodes in excess of 1000%.

In this paper we address these questions by studying the univariate inflation process in a number of countries. Levin and Piger (2003) and Gadzinski and Orlandi (2004) conduct a similar exercise.[2] However, in contrast to their analysis, we examine not only the consumer price index (CPI) in aggregate, but also its components. This helps us to identify if changes in the inflation process have a common source. For example, if inflation persistence is lower because of increased credibility in the monetary policy framework and the consequent anchoring of inflation expectations, one should see a similar reduction in persistence in all components of the CPI.

While from a monetary policy perspective, the aggregate inflation process is probably the most relevant, the disaggregated data may reveal some useful insights about the nature of the price-setting process. The disaggregated data may indicate that the level and changes in persistence are solely attributable to persistence in inflation in a few segments of the economy. Analysis of the disaggregated data may also reveal whether the observed persistence (or lack thereof) may be influenced by the statistical methodology employed in calculating the consumer price index. More generally, understanding the source of the change in the inflation process may allow one to make an assessment of how long-lasting such a change may be. The use of disaggregated cross-country price data also allows us to test whether various theories of price setting are consistent with what we see in the data.

A similar approach has been undertaken simultaneously by Ernst and Mojon (2004), focusing on euro area countries. We use a wider sample of countries than they do, which encompass a larger number of monetary policy regimes, and changes to those regimes.

The analysis in this paper focuses on identifying changes in the inflation process attributable to changes in both the mean and the persistence of inflation, as well as the interaction between the two. As other authors observe, there is an important relationship between the mean and persistence of an economic time series. Much of the work analyses the inflation process (see Perron (1989) and Levin and Piger (2003)). We show that after allowing for changes in the mean of inflation (normally one mean break is sufficient over our sample), inflation has generally not had a particularly high level of persistence, similar to the results in Levin and Piger (2003) and Gadzinski and Orlandi (2004). Importantly, this result suggests that measures of inflation persistence that have been obtained previously depend crucially on the sample period over which they are calculated. These results are common across all categories of prices that we examine. After allowing for the change in the mean of inflation, we do find some evidence of a decline in persistence in the recent period, but the order of magnitude of this decline is less than that found previously in the literature.

Thus our primary conclusion is that the principal change in the inflation process over the past two decades has (not surprisingly) been the decline in the mean. The decline in the persistence of the process has generally been of second order importance, and in some cases has been trivial. Hence the common view that inflation persistence is high is not supported by our results.

Our results suggest that the focus of analysis of the inflation process should be on the causes of the changes in its mean. They show that the timing of the decline in the mean of the inflation process is generally difficult to link directly to marked changes in

[2] See also Anderton (1997).

monetary policy frameworks. In several countries, there is some evidence of a simultaneous decline in mean across the different components of inflation that coincides with a change in the policy framework, but often these also coincide with a recession or marked slowdown in growth, so it is difficult to disentangle these two influences.

These results are consistent with the proposition that the shifts in inflation expectations are the proximate cause of the changes in the mean of inflation. Reducing inflation expectations from a high level is inherently costly. However, once inflation expectations are reduced, the intrinsically low persistence of inflation should ensure that a low inflation regime is maintained. The focus of policy makers therefore should be on maintaining the credibility of the regime and carefully monitoring inflation expectations for any indication that they are rising.

Our study of the disaggregated data show that while many sectors of the economy have very low levels of persistence, there are a few sectors, most notably housing, where higher levels of persistence can remain even after allowing for shifts in the mean. As Altissimo, Mojon, and Zaffaroni (2004) demonstrate, aggregate series inherit the persistence of their most persistent component. This proposition is demonstrated in our results.

This allows us to comment on the debate over the appropriate monetary policy target. Mankiw and Reis (2002) and Benigno (2004) both point out that welfare consideration suggest weighting regional and sector inflation not by their economic size but by their degree of nominal rigidity – sectors and regions with higher rigidity should receive higher weight.[3] The simple intuition is that the more sluggish nominal adjustment, the larger the real adjustments are. Since policymakers are concerned with minimizing the latter, as they are what create welfare losses, they should stabilize the sectors where these are biggest. The time-series properties of the data naturally weight sectors and regions in the manner the theory suggests they should, regardless of the weighting that is used. The implication is that by stabilizing aggregate inflation, policymakers will stabilize inflation of the persistent components.

Our results also raise questions about the mapping of various theories of price determination to movements in aggregate inflation. The previous literature argues that the finding of high levels of persistence is inconsistent with most of the standard theories. Our finding of low inflation persistence is thus more supportive of these theories. However, when we examine the link between our estimates of persistence for the disaggregated CPI data, and estimates obtained from other studies of the duration of price-setting, the results obtained are still at odds with conventional theories. Finally, our analysis also reveals that in some cases, statistical methodology for calculating the CPI can influence the time series properties of the data.

2. Theoretical motivation

Theories of aggregate inflation persistence have generally been derived from microeconomic models of price setting that can be classified into three broad categories (and which need not be mutually exclusive): time-dependent models, limited information models and menu-cost or state-dependent models.[4] Many of

[3] Political economy arguments militate against this, as such a policy runs the risk of rewarding industries and regions that fail to implement structural adjustments that would reduce such rigidities.

[4] Taylor (1999) provides a comprehensive survey of the literature.

these models imply high persistence in the price level, which then translates into very low or even negative persistence in inflation.

The canonical time-dependent model of price-setting was developed by Taylor (1980). In Taylor's model, prices are set as a markup over marginal cost in a sequence of overlapping wage contracts which last for a fixed number of periods, *n*. Each contract is set to take account of both the wages in existing contracts and the wages expected to be set in future contracts. At any point in time, the aggregate price level is the average of the level of prices over the past n periods. In this set-up, shocks today affect wages, and hence the price level, for the next *n-1* periods, as each of the *n* contracts is renegotiated. The longer the length of the contract, the more persistent will be the affect of shocks on wages and on the price *level*. However, positive persistence in the price level implies negative persistence in the rate of inflation. [5] Fundamentally then, the Taylor model is one of price level persistence rather than inflation persistence.

Similarly the much used model of Calvo (1983) implies positive persistence in the price level, but *no* persistence in inflation. In this model, firms change their prices in response to a signal they receive with a fixed probability each period. When prices change, firms reset them to minimise deviations from the expected optimal price level. Because only a subset of prices is changed each period, changes are staggered. The timing of the changes generates the persistence in the price level in response to a shock. But since the price setters are forward looking, basing their decisions in part on expectations, the result is that inflation has no persistence. Instead inflation moves immediately to its new level in response to a shock.

These models of overlapping contracts have been criticised for their assumption of an exogenous fixed contract length (and fixed probability of receiving a signal). An alternative strategy is to assume state-dependent price setting, such as that implied by menu-cost models, where prices are changed depending on the state of the economy and the gap between the current price and the desired price level. Caplin and Spulber (1987) show that standard menu-cost based models of price adjustment do not generate a straightforward mapping between individual price changes and the behaviour of the aggregate price level. Indeed the relationship can vary considerably over time, depending on the shocks that precipitate the price changes. As a result, the persistence of the price level depends on the size and timing of the shocks, and has no direct implications for inflation persistence.

Limited information models generate some persistence in both the price level and inflation. The Lucas (1972) islands model of price setting is the basis for models of this genre. In the face of increased demand for her product, a price-setter is unsure whether this reflects general upward pressure on prices or an idiosyncratic shock. This creates a signal extraction problem where individuals have to ascertain the extent to which the observed price change is economy wide or firm specific. Price setters will only gradually adjust their prices upwards as the information problem is resolved. The limited information slows down the price adjustment process inducing some persistence in aggregate inflation as price setters learn what the true signal is.

[5] Take the simple case in which $p_t = \rho\, p_{t-1} + \varepsilon_t$ where p_t is the log of the price level, ρ is a coefficient, and ε is a white noise disturbance. If one estimates an AR(1) on the first difference of p_t, the resulting autocorrelation estimate will be $\frac{1}{2}(\rho-1)$. So, for the cases in which the price level has positive persistence, inflation will have negative autocorrelation. A general *n* period model will yield a high-order autoregressive process. In all cases, the persistence in the price level will depend not only on the length of the contracts, but also on things like the relative-price elasticity of demand. Whelan (2004) demonstrates this in a more general context.

Variants of the Calvo and Taylor time-dependent price-setting models are a central part of many New Keynesian models (for example, Woodford (2003)). In these models, persistence in inflation can be attributed to four sources: the price-setting process (along the lines of the models discussed above); the mark-up of prices over marginal cost which in turn is related to persistence in the output gap; inflation expectations; and other shocks to the inflation process itself.

As discussed above, the price-setting process itself is unlikely to be a source of positive persistence – indeed the opposite may be true. With forward-looking behaviour, persistence will be zero.

Persistence in the mark-up is related to persistence in the output gap. However, while monetary policy can contribute to some reduction in persistence from this source, the inherent persistence in economic activity provides a lower bound to this.

In many New Keynesian models, the assumption of forward-looking inflation expectations imply that aggregate inflation has no persistence. This aspect of these models has been criticised by Ball (1994), Fuhrer and Moore (1995) and Rudd and Whelan (2001) inter alia. Ball (1994) pointed out that traditional time-dependent models imply that (credible) disinflations are costless, and, in some cases, can even be associated with a boom in output. This is clearly at odds with the practical experience.

The inflation expectations process is therefore the likely source of inflation persistence. Fuhrer and Moore develop a framework where inflation expectations can be written as a weighted average of backward and forward-looking inflation expectations, where the backward-looking component is simply lagged inflation.[6] (Allowing for a backward-looking component in expectations formation has been a common approach in modelling inflation in applied policy research.)

The influence of the monetary policy framework, and particularly an announced inflation target, on the inflation process can then be modelled in terms of their influence of the degree to which inflation expectations are forward looking. A perfectly credible inflation target would cause all price setters to adopt completely forward-looking inflation expectations, anchored on the inflation target, resulting in a world very similar to that in most New Keynesian models. The announcement of a credible inflation target would thereby lead to a marked decline in inflation persistence from the previous regime where there was a strong backward-looking element to expectations (Taylor 1998).

Erceg and Levin (2003) and Orphanides and Williams (2003) develop models of this sort. They show that inflation persistence can come from the public's limited information about the central bank's policy objectives. The persistence arises because the public only gradually learns about changes in the central bank's policy framework. When there is no change in the framework, or the framework is credible, inflation persistence should be low. Similarly Orphanides and Williams simulate a model with similar features and show that the absence of a long-run inflation objective for the central bank results in markedly higher inflation persistence than a world where the inflation objective is clearly understood by price-setters.[7]

[6] In a more recent paper, Fuhrer (2005) confirms that backward-looking expectations formation is the most likely source for observed inflation persistence in the U.S.

[7] The argument can be illustrate very simply with a basic Phillips curve inflation equation where expectations are a mixture of backward and forward expectations where the forward component is the inflation target, $\bar{\pi}$. Write the standard Phillips

Using an analogous argument, Sargent (1999) links inflation persistence directly to the central bank's understanding of the inflation process and its monetary framework (see also Cogley and Sargent (2001)). He argues that the inflation of the late 1960s and early 1970s resulted from the adoption of an inappropriate monetary framework by the Federal Reserve (a similar argument can be made for other countries). The Federal Reserve misinterpreted the evidence of low inflation persistence in the 1950s and 1960s as implying an exploitable trade-off between inflation and output. The objectives of the central bank did not change, rather their understanding of the economy was incorrect.

Effectively, inflation models failed to take account of the fact that inflation expectations may be sensitive to the level of inflation itself, so that we inflation rose expectations of future inflation rose as well. Over time, following the rise in inflation that resulted from this mistaken approach, the Fed gradually learned the true process for inflation. Policymakers observed the persistence in aggregate inflation as expectations adjusted upwards. Their monetary policy framework changed to take account of the inflation persistence, and policy was once again directed at reinstating low inflation, through the disinflations of Volcker and Greenspan. Sargent expresses the concern that the low persistence being observed today may cause central banks to mistakenly trying to exploit a traditional Phillips Curve trade-off, thinking incorrectly that the inflationary consequences of a rise in inflation induced by an increase in output will be minor.

Most of these theories of price-setting assume a seamless mapping from the firm-level price setting decision to aggregate inflation. However, aggregation issues may be critical. That is, it is possible that there is an important difference between thinking about inflation as though there was only one good in the economy rather than the reality that the consumer price index is an amalgam of many different prices. Hence is it meaningful to talk about persistence in aggregate inflation versus persistence in the components? In terms of Lucas' island model, one can think of all prices rising through time with a common component given by the inflation target and relative price movements around that common mean. There may well be persistence in the individual goods categories, as a result of staggered price changes by individual producers of each good in response to a shock to that particular good, but this would not necessarily translate into inflation persistence in the aggregate inflation process, where the effect of the common stable mean would tend to dominate. Alternatively, as discussed above, Caplin and Spulber (1987) and Caplin and Leahy (1997) show that the relationship between individual behaviour and the aggregate price data can be quite imprecise.

3. Methodology and existing evidence on inflation persistence

Estimates of persistence have been obtained in univariate models of inflation, models of the Phillips curve which also take account of the influence of the output gap, exchange rate changes and oil prices on inflation, as well as in larger macroeconomic models.

---

curve as $\pi_t = \pi_t^e + \beta(y_t - \bar{y}_t) + \varepsilon_t$, and define expectations by $\pi_t^e = \alpha\pi_{t-1} + (1-\alpha)\bar{\pi}$. The persistence of the inflation process is given by the parameter α, the weight on the backward-looking term in the equation for inflation expectations. As the weight on the inflation target increases, α declines and observed inflation persistence declines, until in the limit, with perfect credibility of the inflation target, α=0 and there is no persistence in the univariate inflation process (except to the extent that the output gap term is autocorrelated).

The use of a univariate model means that we are omitting other potential drivers of the inflation process. This omitted variable bias may influence our estimates of persistence. It also does not allow us to identify the source of the observed persistence. For example, as described above, in a new Keynesian model, a portion of the persistence in the inflation process may derive from the process determining the output gap. However, again, the purpose of this exercise is primarily to document the statistical properties of the inflation process. We are interested in the following questions: Is the inflation process persistent? And, has that persistence declined? If the answer to the first question is yes, then it would be useful to identify the source of that persistence from a more fully specified model.[8]

A number of approaches have been used in the literature to measure persistence in a univariate model of inflation. These have included the coefficient on the lagged dependent variable in an inflation equation, the sum of the lagged coefficients in an AR(n) model of inflation, the half-life of a shock to the inflation process and the number of times the inflation process crosses its mean (see Andrews and Chen (1994) and Robalo Marques (2004) for a discussion of these issues). Overall, these measures give broadly similar estimates of inflation persistence (Clark 2003). In this paper we use a measure calculated as the sum of the coefficients of an AR(12) process for monthly price series and an AR(4) process for quarterly price series. [9] We have done a full set of exactly analogous computations based on the AR(1) coefficient and using an alternative measure of persistence, we identify the structural breaks as before, but then calculate the number of times the inflation process crosses its mean, as suggested by Robalo Marques (2004). All three measures yield the same results, so we only report one.

If the estimates of persistence are close to one (that is, inflation is close to a unit root process), Hansen (1999) shows that the point estimates can be biased downwards and provides a bootstrap procedure to calculate the estimates of persistence as well as their confidence intervals. However, as we will show in the next section, many of our estimates are relatively small (in absolute value) and so this is not a major issue.

Using these various approaches, much of the previous literature has tended to find that the inflation process is highly persistent. The AR coefficient is often close to one in a large number of countries when estimated on inflation data over the past twenty years or so (Clark (2003), Gadzinski and Orlandi (2004), Levin and Piger (2003), Batini (2002), Batini and Nelson (2001), O'Reilly and Whelan (2004), Stock (2001)). This is the stylised fact which has motivated much of the subsequent theoretical work.

More recently, some papers have examined whether this estimate of persistence has changed over time. Debelle and Wilkinson (2002), Levin and Piger (2003) and O'Reilly and Whelan (2004) used rolling regressions to examine the evolution of the AR coefficient. Debelle and Wilkinson show that persistence has declined considerably over the past decade in Australia, the United Kingdom, Canada and New Zealand, but there was little evidence of a decline in persistence in the United States. Using more recent data Levin and Piger show that in the United States, persistence has also declined by a similar order of magnitude, but only relatively

[8] Our strategy is supported by the results in Fuhrer (2005) who concludes that, to the extent that inflation is persistent, the forcing process is not the likely source.

[9] Hereafter, a reference to an AR(12) model should be translated to AR(4) for Australia and New Zealand which only have quarterly CPI data.

recently. O'Reilly and Whelan and Gadzinski and Orlandi (2004) find little evidence of a change in persistence in euro area countries. These rolling regressions can indicate whether persistence has changed, but are not very precise in determining the exact timing of the change in persistence and hence it is difficult to map the change in persistence to factors such as a change in the monetary policy framework.

In obtaining these estimates of high inflation persistence, and measuring the change in persistence over time, very few papers allow for the possibility of a shift in the mean of inflation over the sample period. Perron (1989) shows that failing to account for a shift in the mean of a process will give misleading estimates of persistence. Against this, allowing for too many shifts in the mean can lead to an under-estimate of persistence. In the extreme, controlling for a shift in the mean each period would generate an estimate of zero persistence. In this paper we take an agnostic view on the appropriate number of mean shifts to allow for. Our primary purpose is to demonstrate the significant impact that mean shifts can have on existing estimates of persistence. As will be discussed in the results below, the largest decline in the estimate of persistence tends to occur when only one mean shift is allowed for.

Clark (2003) and Levin and Piger (2003) do allow for an explicit shift in the mean, and demonstrate that, even so, inflation persistence is markedly lower in the more recent period. For example, Clark finds that allowing for a break in the mean of inflation in 1993Q1, inflation persistence in the United States in an AR model is reduced substantially from roughly 0.9 to 0.13. Similarly, the results in Levin and Piger show that once a structural break is allowed for, the null hypothesis of a unit root can be rejected at the 95 percent confidence level for 29 of the 48 inflation series that they examined whereas the null hypothesis could only be rejected for eight series when no structural break was allowed for. Gadzinski and Orlandi (2004) also find that once shifts in the mean are controlled for, persistence has generally been low in euro area countries and the United States. Bilke (2004) and Corvoisier and Mojon (2005) show that persistence in aggregate measures of inflation in France and OECD countries (respectively) has been low for a number of decades once mean breaks are allowed for. Bilke finds evidence of only a single mean break in France in the mid 1980s, while Courvosier and Mojon find a number of significant breaks in OECD countries over the past four decades: in the late 1960s/early 1970s, the early-mid 1980s and the early 1990s.

To test for shifts in the mean, we use the methodology described in Bai (1999). First we conduct a Quandt (1960) test on the AR(1) and AR(12) models of inflation. This finds the maximum value of the Chow test over all possible break points. Once we have identified this point, we re-estimate the model allowing for a structural break in the mean at this date and obtain a second estimate of persistence. We then adopt the same procedure to identify a second and third mean break, in each case obtaining another estimate of persistence. Levin and Piger (2003) find that Bayesian methods of testing for structural breaks generate very similar results to the approach we have adopted here.

We also conduct the same Quandt procedure to identify a structural shift in persistence in the original AR(12) regressions with no mean breaks, as well as the regressions which include one break in the mean of the series.

In addition to testing for persistence in the aggregate CPI, which has tended to be the focus of much of the previous literature, we also examine the persistence properties of disaggregated components of the CPI (Ernst and Mojon (2004) adopt a similar

approach using euro area data). Along with the use of cross-country data, this allows us to investigate where the changes in the properties of the aggregate CPI series are coming from. If the change in the mean or the persistence parameter is associated with a change in the monetary policy framework, then it is likely that the timing of the change will be similar for the disaggregated components. As a result we have obtained CPI data at the first level of disaggregation, which generally includes the following categories of goods and services: food, alcohol & tobacco, clothing & footwear, housing, furniture, health, transport, recreation, communication and education.[10]

We can also investigate whether the statistical methodology used to calculate the CPI from the raw price data may be artificially generating some of the persistence. For example, the calculation of the housing component of the CPI in some countries involves an estimate of owner's equivalent rent that embodies a moving average component. And in some countries, the treatment of some prices has changed as the timing of sales changed.

As mentioned above, Altissimo, Mojon, and Zaffaroni (2004) demonstrate, aggregate series inherit the persistence of their most persistent component. The argument is summarised in Box 1. Hence, using the disaggregated data we can identify which of the components of the CPI are the primary contributors to the persistence in the aggregate CPI, or whether persistence is similar for all components.

Other authors have examined the issue at an even more disaggregated level, using the price data that are the basic inputs to the CPI. Bils and Klenow (2002) examine the properties of 123 price components of the US CPI. They that few of these series exhibit much persistence; and that there is little relationship between the frequency of price change and what persistence (and volatility) there is in the price series. Clark (2003) uses a similar data set and finds that again, after controlling for a structural shift in the mean, there is very little evidence of persistence in either the aggregated or disaggregated data. In Europe, Aucremanne and Dyhne (2005) examine disaggregated data for Belguim, while Álvarez and Hernando (2004) use Spanish data to determine whether time or state dependent models best describe pricing behaviour. The former find that a mix of both models appears to be present in the data, while the latter find evidence in support of the Calvo model. But as the discussion in the previous section shows, this is consistent with zero persistence in inflation.

Finally, one issue which arises is whether the analysis should be conducted using seasonally adjusted (SA) data or data that are not seasonally adjusted (NSA). From a theoretical perspective, the NSA data would be preferable as they relate directly to the actual price decisions taken by the firms. However, for a number of goods and services such as school fees and clothing, there is a large seasonal element to price changes and some price decisions are taken only on an annual basis. This suggests that the SA data will give a more accurate indication of the true underlying inflation process.

[10] The data are seasonally adjusted using the X-12 procedure in all cases except for the US and Germany where we were able to obtain seasonally adjusted data directly. The seasonal adjustment is particularly important for some components where there tends to be a large shift in the price level once a year, such as education. Although, as we note in the next section, all of our conclusions are robust to the use of data that are not seasonally adjusted.

4. Results

The results of our examination of the inflation process fall naturally into three groups. First, there is the set of results based on the aggregate inflation data. Here we have data on 19 countries over various time periods. Second, we present results based on data covering prices in as many as 12 disaggregated categories for each of the countries. Finally, we look at the timing of the changes in the mean of inflation to see if they coincide either within countries or across commodities.

We present the results based on an AR(12) process for inflation estimated with the seasonally adjusted data. Our focus is on the simple sum of the twelve coefficients on the lags in the autoregression. The results obtained with the seasonally *un*adjusted data were not substantively different, as were the results obtained using the measure of mean-crossings suggested by Robalo Marques (2004).

Aggregate Inflation

We begin with the measures of persistence in the aggregate inflation series, conditional on allowing for up to 3 changes in the mean of the series. The results in Table 1 are ordered by the estimate of the persistence that assumes no breaks in the mean of the time series.

The results have the following noteworthy characteristics. First, without allowing for any changes in the mean, persistence is high and close to 1 in many countries. The exceptions tend to be a number of European countries for which we only have data from the 1990s onwards - an issue to which we return below.

The AR(12) estimates of persistence reported in Table 1 are higher than those obtained using an estimate based on an AR(1) process, although those countries which have higher persistence on the AR(12) measure also tend to have higher persistence on the AR(1) measure. The AR(1) measures were however, significantly affected by seasonality in the inflation data.

Second, allowing for the mean of inflation to change just once over the sample significantly reduces the estimate of persistence in most countries (the results are shown in Figure 1). Allowing for more changes in the mean further reduces estimated persistence but often by a much smaller amount.

After allowing for one break, seven of the nineteen estimates of persistence exceed one half; allowing for three breaks, only three countries have inflation with persistence larger than one half, namely Germany, Italy and Japan.

For a number of the European countries reported in Table 1, the results are obtained using data beginning around 1990. In nine of the countries in our sample, however, we have longer time series and are able to examine the importance of the short sample period. For this subset of nine countries, the results comparing the full sample with a shorter sample beginning in 1990 are reported in Table 2. In each there is a marked decline in measured persistence once we use the shorter sample, although for Germany, Italy, Japan and Sweden, a low level of persistence is only obtained once we allow for one mean break.

Table 1: Persistence of Aggregate Inflation AR(12/4) model, seasonally adjusted data				
	No Breaks	One Break	Two Breaks	Three Breaks
Italy	0.97	0.92	0.86	0.80
Brazil	0.95	0.41	0.38	0.38
Portugal	0.93	0.45	0.18	0.04
US	0.89	0.57	0.40	0.45
Japan	0.88	0.80	0.72	0.93
Chile	0.87	0.71	0.57	0.15
Australia	0.86	0.59	0.55	0.29
Germany	0.85	0.74	0.60	0.60
New Zealand	0.82	0.45	0.45	0.32
UK	0.82	0.66	0.61	0.33
Canada	0.80	-0.09	-0.19	-0.30
Sweden	0.79	0.41	0.32	0.13
Finland	0.79	0.26	0.03	-0.33
Austria	0.70	0.31	0.16	0.19
Spain	0.61	0.24	0.18	-0.54
Netherlands	0.56	0.00	0.07	0.04
France	0.45	0.22	-0.11	-0.42
Luxembourg	0.34	-0.41	0.07	0.96
Belgium	0.00	-0.21	-0.73	-2.39
Median	0.82	0.41	0.32	0.19
EU Aggregate	0.75	0.83	0.89	0.55

Computed as the sum of the coefficients in an AR(12) or AR(4) autoregression. Breaks are determined by sequential Quandt (1960) tests on individual series. The countries are sorted from most to least persistent, as measured by the case without breaks.

Figure 1: Changes in Persistence of Aggregate Inflation
AR(12/4) model, seasonally adjusted data

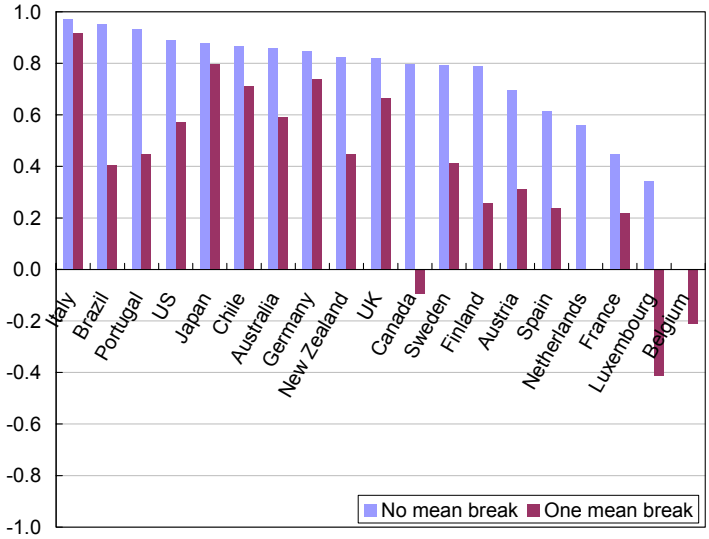

Hence if we were to restrict ourselves to the last decade and a half we would conclude that inflation in these countries showed the same lack of persistence evident in the other euro area countries. Alternatively, if we had a longer run sample for the euro area countries, there may be evidence of higher persistence and/or a higher mean in earlier periods, although Gadzinski and Orlandi (2004) tend not to find this.

These results highlight the importance of controlling for mean breaks when estimating persistence, as well as the critical role played by the choice of sample period. If one does not allow for any break in the mean, then aggregate inflation persistence tends to be higher and close to one, as has been found previously (eg, Fuhrer and Moore (1995), Rudd and Whelan (2001)). However, allowing for even one break in the mean substantially reduces the estimates of persistence, as Levin and Piger (2003) and Gadzinski and Orlandi (2004) have also demonstrated.

Table 2: Impact of Sample Period on Estimated Persistence

AR(12/4) model, seasonally adjusted data

	Sample	No Breaks	One Break	Two Breaks	Three Breaks
Australia	Long	0.86	0.59	0.55	0.29
	Short	0.29	0.13	-0.06	-0.18
Canada	Long	0.80	-0.09	-0.19	-0.30
	Short	-0.08	-0.20	-0.65	-0.93
Germany	Long	0.85	0.74	0.60	0.60
	Short	0.73	0.02	-0.19	-0.26
Italy	Long	0.97	0.92	0.86	0.80
	Short	0.77	0.32	0.35	0.39
Japan	Long	0.88	0.80	0.72	0.93
	Short	0.46	0.11	0.11	-0.39
New Zealand	Long	0.82	0.45	0.45	0.32
	Short	0.36	0.18	0.17	0.10
Sweden	Long	0.79	0.41	0.32	0.13
	Short	0.65	0.34	0.31	0.29
UK	Long	0.82	0.66	0.61	0.33
	Short	0.43	0.14	0.12	-0.04
US	Long	0.89	0.57	0.40	0.45
	Short	0.04	-0.07	-0.35	-0.87

Short samples all begin in 1990. Long samples begin in 1969 in Australia, 1984 in Canada, 1960 in Germany, 1977 in Italy, 1970 in Japan, 1975 in New Zealand, 1980 in Sweden and in the UK, and in 1978 in the USA.

If we test for a break in the estimate of persistence without allowing for a break in the mean, we find a decline in persistence in a number of countries. The two exceptions are the Netherlands and Luxembourg where persistence rises in the second half of the sample from a low negative number.[11] These results (shown in Figure 2) are consistent with those obtained from rolling regressions such as in Debelle and Wilkinson (2002) and Levin and Piger (2003). While inflation persistence declines in nearly every case, the decline is only significant for 6 of the 19 countries. Again,

[11] Estimates of persistence of greater than one were obtained in a few cases prior to the break. This generally was a result of the break in persistence occurring early in the sample when inflation was on a sharp upward trend, such as in Japan in the early 1970s.

account must be taken of the fact that we only have a relatively recent sample of data for the euro area countries.

Figure 2: Changes in the Persistence of Aggregate Inflation
AR(12/4) model, seasonally adjusted data, no break in the mean

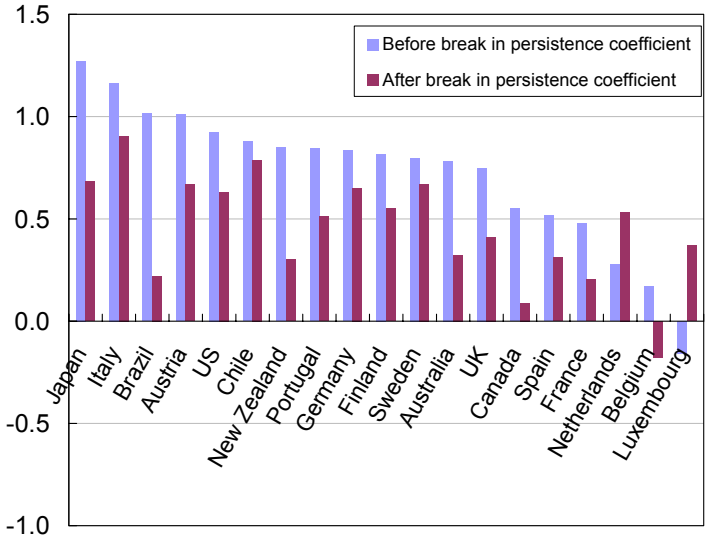

Figure 3: Changes in the Persistence of Aggregate Inflation
AR(12/4) model, seasonally adjusted data, given one break in the mean

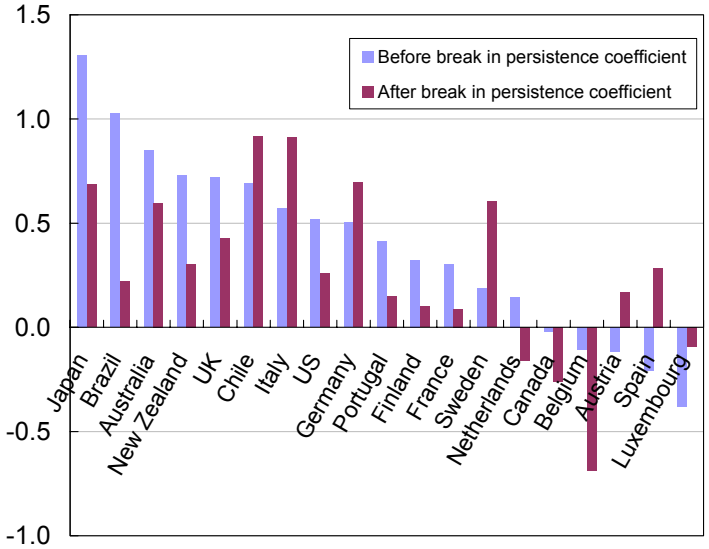

Finally, we look to see whether the estimate of persistence changes, once we allow for a single break in the mean. That is, we first find the most likely date for a break in the mean, and then the most likely date for a break in the slope of the AR(12) regressions. The results are reported in Figure 3. The estimates show declines in persistence in nearly every case. Around half of these changes have conventional t-ratios below -1.6.

The Disaggregated Data

Having identified changes in the mean of the aggregate inflation process, we now examine the component series. We study roughly 12 consumption categories for most of the countries. In total, we have 202 times series that can be used to estimate the AR(12) model. To get a flavour of the results, we begin with a histogram that displays the frequencies of the estimated AR(12) parameters, conditional on the number of assumed breaks in the mean of the individual inflation series. Figure 4 shows both the decline in estimated persistence once mean changes are allowed for, as evidenced by the leftward shift in the distribution, as well as the wide range of estimates across countries and components.

Figure 4: Frequency Distribution of Persistence in Components of Inflation
AR(12/4) model, Seasonally Adjusted Data

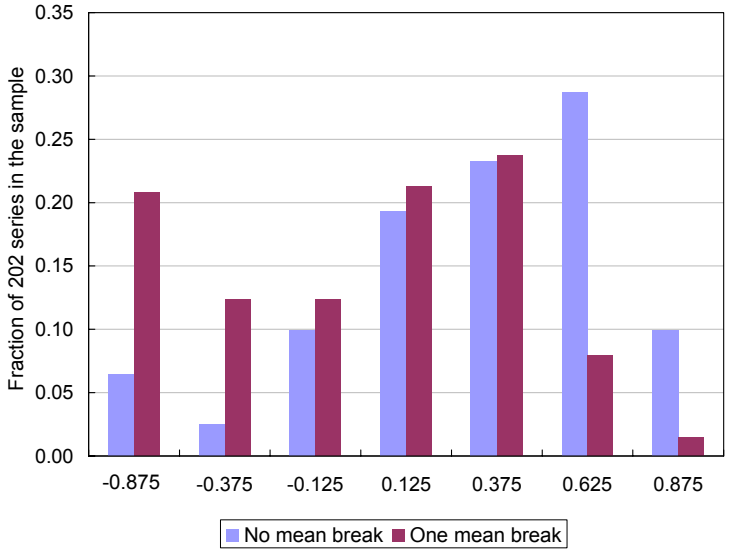

The descriptive statistics in Table 3 bear out the importance of the mean breaks. The fraction of estimates that are positive shrinks from 86%, with no mean breaks, to 59% allowing for one break, and then 48% when there are two breaks. Results for the (mostly European) countries which have the shorter sample show generally less persistence.

Moreover, Table 3 shows that a large fraction of the estimates of persistence are small, even if there is no mean break. Even in the case of the countries for which we have the longer time series, which tend to have higher positive persistence, one half of the series have estimates of persistence below 0.68, while once we allow for one mean break, this declines to 0.38, and to 0.08 once three mean breaks are allowed for.

Table 3: Descriptive Statistics for Persistence in Inflation Subaggregates				
AR(12/4) Model, seasonally adjusted data				
	No Breaks	One Break	Two Breaks	Three Breaks
Full Sample (19 countries + EU, 202 Series)				
25th percentile	0.21	-0.29	-0.58	-0.95
Median	0.53	0.20	-0.04	-0.23
75th percentile	0.77	0.43	0.30	0.19
Percent > 0	86%	59%	48%	38%
Longer sample available (sample starts pre 1987 - 9 countries)				
25th percentile	0.41	-0.04	-0.15	-0.29
Median	0.68	0.38	0.22	0.08
75th percentile	0.81	0.56	0.44	0.31
Percent > 0	93%	72%	66%	56%
Shorter sample only (sample starts in 1987 or later - 10 countries + EU)				
25th percentile	0.09	-0.66	-1.28	-1.48
Median	0.39	-0.03	-0.38	-0.64
75th percentile	0.61	0.34	0.17	-0.05
Percent > 0	80%	48%	34%	23%
Computed using the estimates AR(12/4) coefficients from regressions of the 202 components of inflation. Estimates use the full sample available in each country.				

It is interesting to take the estimates and compute averages by product categories across countries. This provides some indication of whether the persistence in the aggregate inflation is a general feature of the price data or is attributable to any of a small number of the component series.

Tables 4a and 4b report the results of this exercise. As with the aggregate data, there is a clear decline in estimated persistence as we control for breaks in the mean. Once we allow for breaks, very few components show any sizeable amount of positive persistence with the possible exception of housing. The large negative estimate for clothing is likely a consequence of the volatility resulting from changing seasonal patterns of sales in some European countries. (Standard seasonal adjustment techniques are unable to cope adequately with this.)

Table 4a: Average Persistence Across Countries within Components				
AR(12/4) regression, sa data				
Component	No breaks	One Break	Two Breaks	Three Breaks
Food	0.55	0.23	0.07	-0.09
Alcohol	0.28	-0.33	-0.53	-0.77
Clothing	-0.35	-1.07	-1.34	-1.62
Housing	0.63	0.39	0.18	-0.15
Furniture	0.59	0.09	-0.06	-0.17
Health	0.34	-0.18	-0.46	-0.60
Transport	0.49	0.23	0.11	-0.02
Communication	0.16	-0.48	-0.77	-0.97
Recreation	0.35	-0.17	-0.54	-0.66
Education	0.49	-0.07	-0.33	-0.84
Restaurants	0.58	-0.06	-0.31	-0.48
Miscellaneous	0.56	0.20	-0.02	-0.29

Table 4b: Average Persistence Across Countries within Components				
Countries with longer sample only, AR(12/4) regression, sa data				
	No breaks	One Break	Two Breaks	Three Breaks
Food	0.67	0.37	0.19	0.05
Alcohol	0.40	0.03	-0.13	-0.34
Clothing	0.68	0.21	0.10	-0.09
Housing	0.78	0.56	0.45	0.37
Furniture	0.77	0.41	0.35	0.24
Health	0.36	0.04	-0.12	-0.21
Transport	0.56	0.17	0.06	-0.07
Communication	0.37	-0.06	-0.28	-0.43
Recreation	0.63	0.24	0.05	-0.07
Education	0.80	0.58	0.41	0.24
Restaurants	0.84	0.56	0.49	0.48
Miscellaneous	0.54	0.30	0.24	0.03

If we focus on the results for the longer sample (in Table 4b), we can classify the estimates into three categories. Two tend to have high estimates of persistence even when breaks in the mean are allowed for: housing and restaurants. Housing inflation is highly persistent for nearly all countries in the sample, and may well be related to the methodology of measuring prices employed by the statistical bureau. Certainly this is the case for the two countries with which we are most familiar, namely Australia and the United States.[12]

Indeed it is the case that housing inflation is the most common source of high persistence, once we allow for mean breaks, in each of the nine countries for which we have the longer sample of data. [13]

Inflation in the prices of clothing, food, recreation, transport and furniture tend to have high persistence when there are no mean breaks, but this declines markedly once we allow for breaks in the mean. Note that transport prices generally include petrol. The estimates of persistence for health, communication and alcohol are generally low.

Focusing on the countries with high aggregate persistence, even after allowing for mean breaks we find that persistence in only a few components is responsible for the aggregate persistence. For example, in Italy, after allowing for one mean break, aggregate persistence is 0.92, which is a result of persistence of 0.93 in furniture price inflation, 0.72 in food inflation and 0.90 in restaurant food inflation. Persistence in the other price categories is below 0.5. Similarly, in the case of Japan, high persistence in housing, furniture and education is responsible for the high aggregate persistence while persistence in the other components is relatively low.

Overall, these results show that while flexible price-setting behaviour can be pervasive in many sectors of the economy, high persistence in only a few sectors can result in high aggregate inflation persistence. Thus, the evidence of low persistence in disaggregated price data presented by Bils and Klenow (2002) and others is interesting in terms of assessing the various theories of price-setting, but is less relevant for the conduct of monetary policy, which is concerned about the aggregate properties of inflation.

[12] In Australia's case, recent changes in methodology mean this is no longer true.

[13] It is interesting that the shorter sample European data do not have a housing component, which may be a possible cause of the low persistence in the data.

This evidence argues against the proposition raised by Maniw and Reis (2002) and Benigno (2004) that policy makers should focus on those sectors of the economy where prices tend to be more rigid. It supports the rebuttal provided by Otmar Issing (2005): “by assigning higher importance to more rigid units, monetary policy would accommodate structural inefficiencies, ultimately creating perverse incentives. These difficulties would introduce some arbitrariness in the conduct of monetary policy and negatively affect the transparency and accountability of the central bank.”

The Timing of the Breaks

We next look at the timing of the estimated breaks in the inflation process to address several interesting questions:

1. Do the breaks in the mean occur at the same time across commodity groups within a country? Or, alternatively, are they at the same time within commodity groups across countries?
2. Do the estimated changes in persistence (conditional on one mean break) occur at the same time either across commodity groups within a country, or within a commodity group across countries?
3. Finally, we can look to see if the mean and persistence breaks occur at the same time.

Figure 5 summarises the information on the timing of the breaks for the countries for which we have a longer run of data. In each case, the graph shows the date of the first and second mean breaks in aggregate inflation and its various subcomponents, as well as the timing of the break in persistence. For example, the panel for Canada shows that 7 of the first mean breaks in inflation occurred in 1991, while 3 of the breaks in persistence occurred at about the same time.

Overall, there is little evidence that the timing of the breaks in the mean or the breaks in persistence have any similarity *across* countries. The breaks are dispersed throughout the sample period. The same statement is true if we look within product categories across countries. The breaks in mean inflation are neither a worldwide nor a sector-specific phenomenon.

However, *within* countries there is some evidence of a clustering of breaks in mean and persistence, particularly when we look at the longer sample. For New Zealand, there is commonality in break points around 1987, when the country embarked on a major program of reform, including to its monetary policy framework.

In Australia, there is a clustering of mean breaks, and persistence breaks in the early 1990s, when the country went into recession, and there was a marked drop in the level of aggregate inflation. This is also true of Sweden and the UK in the early 1990s, and Germany in the early 1980s. Similarly, there is a clustering of break points in Canada in 1991, when again there was a recession. In Canada, there was also an announced shift in the framework for monetary policy around this time, but it is not possible to disentangle the effect of this on price-setting behaviour from the effect of the recession. Again, in the United States, there are breaks in the mean of inflation of a number of components around the time of the Volcker disinflation. Here again there is a change in monetary policy framework coincident with a contraction in economic activity.

Figure 5: Frequency Histogram of Location of Breaks in Inflation Series

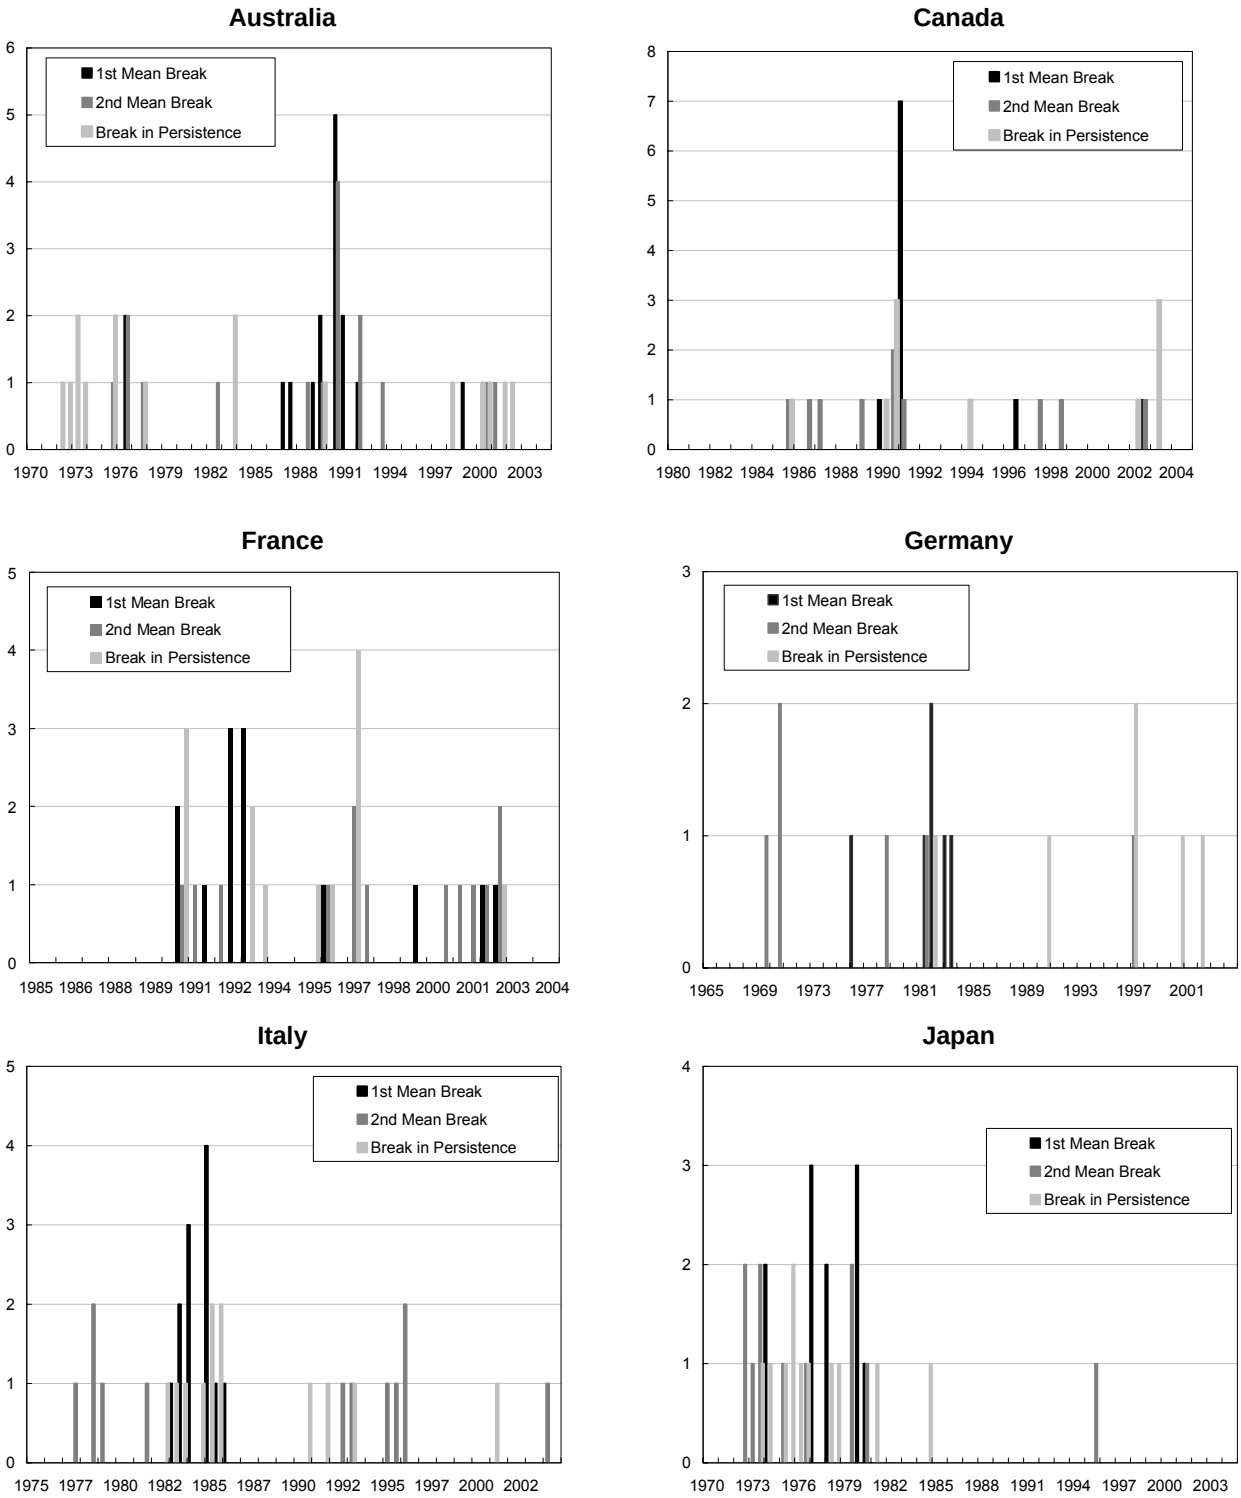

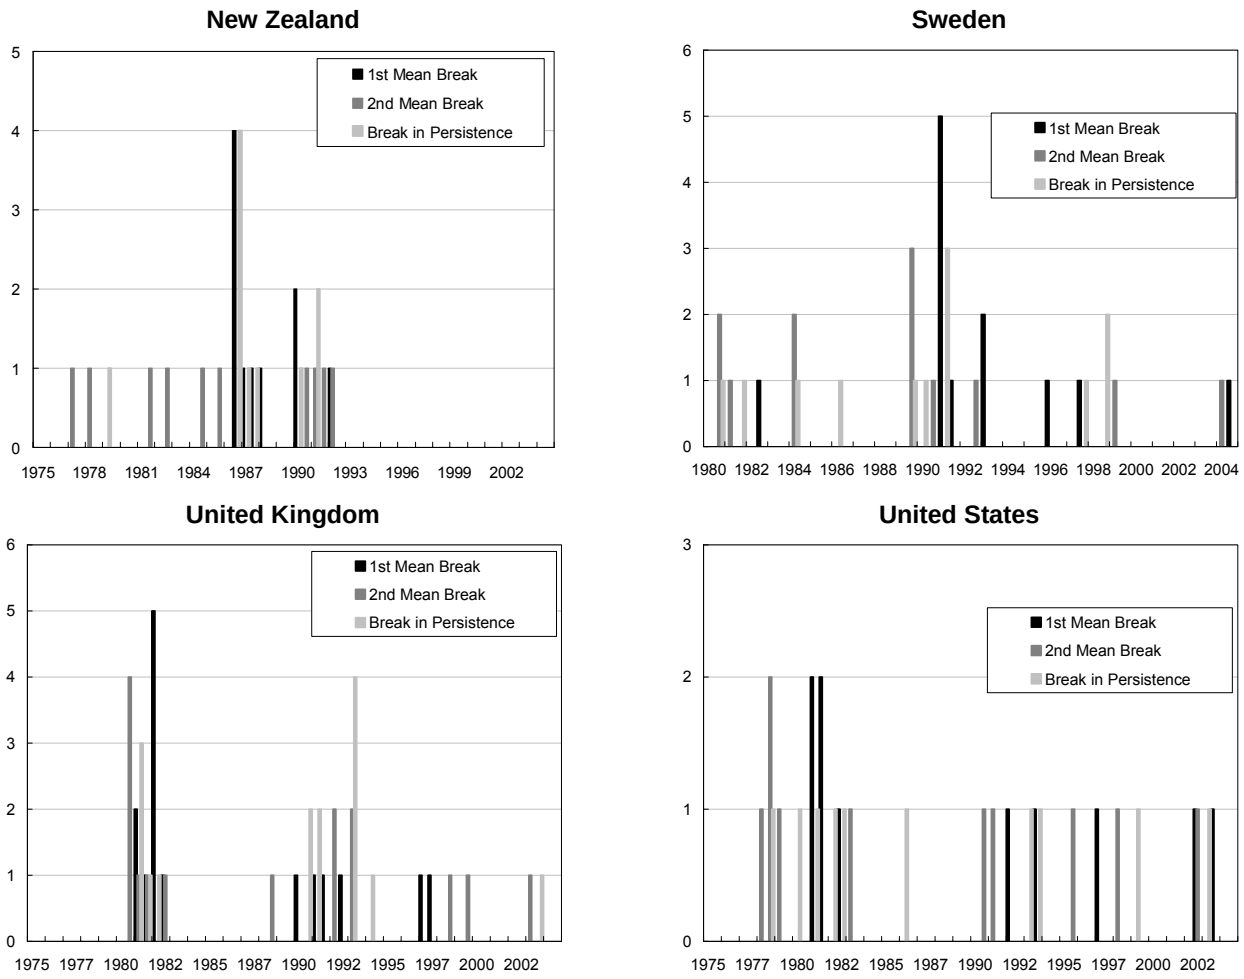

In Japan, the breaks tend to occur in 1974, when the country suffered from the oil price shock. The oil price appears to have been translated into a shift in the mean of a number of components of inflation, as well as in their persistence.

Results for the subset of European countries (not presented here) do not display much commonality in the timing of breaks either across or within countries. There is little indication in these results that the changes in the operation of monetary policy associated with the operation of the ERM or the commencement of the European Monetary Union in 1999 had a material influence on the inflation process. Furthermore, for the sample of EU countries there is also not much evidence that the recessions of 1992 or 2001 had an impact on the inflation process.

The relationship between duration, persistence and variability.

The theories discussed in section 2 have implications for the relationship between the duration, persistence and the variability of innovations to inflation. They have the following important implications:[14]

[14] For a detailed discussion of the theoretical basis for these statements see the summary in Section 2 of Angeloni, Aucremanne, Ehrmann, Galí, Levin and Smets (2004).

1. The longer the time between price changes, the more persistence there will be in the price level, and the more negative the persistence in price changes.
2. The longer the time between changes, the higher the variability of the innovations to inflation (in a reduced-form autoregression).
3. The more persistent are innovations to the markup over marginal cost (or the output gap) the more persistent is inflation.
4. The more backward-looking (or indexed) are pricing decisions, the more positive persistence there is in price changes.
5. The more aggressive monetary policy is in keeping inflation close to its target the higher the variability of innovations to inflation.

Our regressions yield information on persistence and the standard deviation of the innovations to inflation. In addition to this we include data on the duration of price changes for the categories of goods that broadly correspond to the ones used in our analysis for eight European countries and the United States (from Bils and Klenow 2002).[15]

We examine three regressions. In the first we regress the standard error of the inflation regressions on duration (controlling for country-specific fixed effects which may be a proxy for the monetary regime that is present). In the second we look for a correlation between persistence and duration (again, both with and without). Finally, we regress our measure of persistence on the estimated standard deviation of the inflation innovations in the autoregressions.

The results, reported in Table 5, show that duration is uncorrelated with both persistence and variability of inflation innovations. This is evidence against implications 1 and 2. However, persistence and variability are highly negatively correlated. These results are strengthened by the inclusion of country-specific fixed effects. The relative importance of country fixed effects in all of our regressions (half of them have p-values less than 0.05) suggests a clear role for monetary policy supporting implication 5.

We interpret these results as consistent with Fuhrer's (2005) conclusion that persistence in the markup is not the source of persistence in inflation. Instead, it is backward-looking expectations, or indexation, are responsible. Our argument has two parts. First, there is no reason to believe that what we are finding is a consequence of omitting other possible drivers from the persistence regressions in panel C of Table 5. Such variable should not have any particular relationship with variance of the innovations in the inflation regressions. Second, as Galí (2004) outlines, more price rigidity generates more backward-looking behaviour and hence more persistence. Intuitively, firms facing bigger shocks – either to demand or costs – will have be less backward looking. The more volatile the environment, the more flexible the firm's pricing policy will have to be and the less inertia it can exhibit.

[15] The data on duration were obtained from a number of sources. The U.S. data are from Bils and Klenow (2002) Table 3. Data for Austria, Belgium, Finland, France, Italy, Netherlands, Portugal and Spain are from the papers summarized in Dhyne et.al (2005).

Table 5: The Relationship among Persistence, Duration, and Variability				
AR(12/4) regression				
	No Breaks	One Break	Two Breaks	Three Breaks
A. Variability and Duration				
Regression without Fixed Effects				
Coefficient Estimate	0.01	0.02	0.02	0.02
t-statistic	1.01	1.38	1.44	1.18
Regression with Fixed Effects				
Coefficient Estimate	0.01	0.02	0.02	0.02
t-statistic	0.97	1.14	1.18	1.00
B. Persistence and Duration				
Regression without Fixed Effects				
Coefficient Estimate	0.00	0.00	0.00	0.00
t-statistic	-0.28	-0.25	-0.23	-0.20
Regression with Fixed Effects				
Coefficient Estimate	-0.01	-0.01	-0.01	-0.01
t-statistic	-1.00	-0.96	-0.96	-0.95
C. Persistence and Variability				
Regression without Fixed Effects				
Coefficient Estimate	-0.15	-0.35	-0.35	-0.32
t-statistic	-2.18	-3.12	-2.73	-2.07
Regression with Fixed Effects				
Coefficient Estimate	-0.69	-0.83	-0.98	-1.03
t-statistic	-8.30	-7.00	-7.39	-6.49

Source: Regressions of the estimated sum of the coefficients in AR12 (or AR4) autoregressions on the standard error of the regression across countries and components. Fixed effects regressions allow the constant to differ for each country. Regressions in panels A and B, with duration, are based on data from 9 countries with 90 time series. Regressions in panel C utilize the entire data set of 19 countries and 202 time-series. Reported t-statistics are robust to heterskedasticity.

5. Conclusion

The main result of this paper is that the most significant change in the inflation process has, not surprisingly, been in its mean. We show that the conventional wisdom, that inflation has a high level of persistence, is not robust. Controlling for a break in the mean of inflation, measured persistence is considerably lower. This is true for both aggregate consumer prices and for disaggregated components. We do find some support for a decline in inflation persistence over the past few decades, but this is generally from a level that is often relatively low.

From our study of disaggregated data we conclude that persistence in one component of inflation is often enough to generate persistence in the aggregate. Thus while micro evidence may indicate low persistence in many categories, this is relatively unimportant for macroeconomic policymakers whose primary concern is the aggregate inflation process.

This is a consequence of the fact that aggregate inflation inherits the persistence of its most persistent disaggregated component. This means that central bankers can achieve their objective of stabilizing inflation in the sector with nominal rigidities by targeting aggregate inflation. Policymakers need not go to the trouble of distinguishing the flexible from the sticky-price sectors. By basing decisions on the aggregate, they will automatically focus on the latter.

While one can argue that the changes in monetary policy frameworks have contributed to a reduction the mean of inflation, it is much harder to make the case that they have had a meaningful impact on persistence. Examining the timing of mean breaks across commodities within individual countries suggests that the relationship between changes in the monetary framework and changes in the mean of inflation is far from perfect (although the data for the euro area may post date the significant changes in monetary frameworks in those countries). There are some exceptions including the disinflations in the US in the early 1980s and in New Zealand in the late 1980s. Both episodes arguably entailed a change in the monetary framework. Recessions also seem to have some influence and are sometimes simultaneous with changes in monetary frameworks.

The finding of little change in persistence also implies that the slope of the short-run aggregate supply curve – that is, the elasticity of inflation with respect to transitory changes in output – has been stable, once account is taken of change in the mean of inflation. Put another way, the sacrifice ratio has not changed much, even in the face of changes in monetary policy regime. But the important role played by changes in the mean of inflation in our results highlights the need to take account of changes in inflation expectations when considering sacrifice ratios. It is likely that the change in the mean is directly related to changes in inflation expectations, and hence that examining the persistence of inflation expectations may be a more fruitful area of future research. If inflation expectations have less persistence because of increased credibility in the monetary framework, then there will not be a need for the central bank to engineer a decline in output below potential to return inflation to its desired level.

We find occasional large shifts in the mean, which one might think of as extremely persistent (i.e. permanent) responses to shocks. But the rest of the time, persistence is modest. So occasionally there are shocks coupled with monetary policy reactions that allow or engender shifts in the mean of inflation, but the rest of the time, the monetary policy response, or the credibility of the framework, ensures that the shocks do not translate into mean shifts. This conclusion is further buttressed by our conclusion that, to the extent that sectoral inflation is persistence, it is backward-looking expectations that are the source. In the end, there is no substitute for credible monetary policy that focuses on keeping inflation low and stable.

Box 1: Autocorrelation and Aggregation

An aggregate time series inherits its persistence properties from its most disaggregated component. Thinking about averaging a random walk and a white noise series, we can see that this makes sense. Such an average would be a random walk, regardless of the weights.

This example is surely extreme, so it is useful to example a simple experiment within the range of data that we study here. Consider a case in which a time series y_t is the equally weighted average of two time series x_{1t} and x_{2t}. That is

$y_t = 0.7x_{1t} + 0.3x_{2t}$

Next, assume that the x's are first-order autoregressive processes with different parameters. So

$x_{1t} = \rho_1 x_{1t-1} + e_{1t}$

and

$x_{2t} = \rho_2 x_{2t-1} + e_{2t}$,
where e_{1t} and e_{2t} are i.i.d. standard normal. Then estimate the AR(1) for y_t

$y_t = \gamma y_{t-1} + u_t$

for various values of ρ_1 and ρ_2 and compare the estimated value $\hat{\gamma}$ to the weighted average $\bar{\rho} = 0.7\rho_1 + 0.3\rho_2$.

Figure A: Autocorrelation of the Average and Average of the Autocorrelations

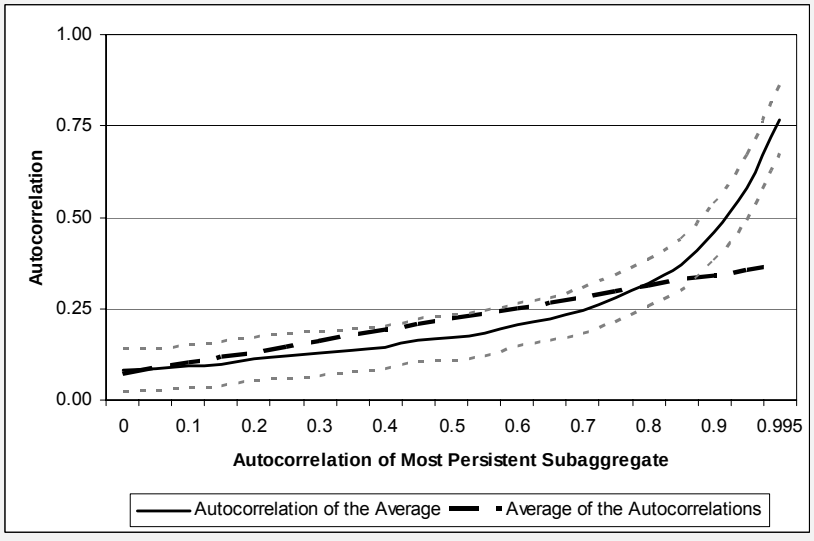

For the purposes of the experiment, we fix ρ_1=0.1 and vary ρ_2 from 0 to 0.995. For each parameter setting, we take 1000 draws of time series of length 180 (that's 15 years of monthly data). In every case, the initial conditions are set such that x_1=x_{20}=0 for t=-20. That is, the first 20 observations are dropped.

Figure A plots the median of the estimated value of $\hat{\gamma}$ obtained from the 1000 draws (the solid line), together with a band equal to plus and minus the mean absolute

deviation of the estimates (the light dashed lines), as well as the weighted average $\bar{\rho}$ (the dark dashed line).

The results are quite striking. For ρ_2 between 0 and 0.8, the autocorrelation of the average is roughly the average of the autocorrelations. But as ρ_2 rises above 0.98, the difference grows rapidly. For ρ_2=0.9, the autocorrelation of the aggregate is 0.46, well above the average autocorrelation of 0.34. As ρ_2 rises to 0.95, the difference becomes even starker: The autocorrelation of the aggregate is 0.580 versus the average of 0.355.

References

Álvarez, L J and I Hernando (2004), 'Price setting behaviour in Spain: stylised facts using consumer price micro data', ECB Working Paper no. 416.

Anderton, R (1997), 'Did the underlying behaviour of inflation change in the 1980s? A study of 22 countries,' Weltwirtshaftliches Archiv, 133(1), pp 22-38.

Andrews, D and W Chen (1994), 'Approximately median-unbiased estimation of auto-regressive models', *Journal of Business and Economic Statistics*, 12, pp 187-204.

Altimissimo, F., B. Mojon and P. Zaffaroni (2004), 'Fast micro and slow macro: can aggregation explain the persistence of inflation?', unpublished manuscript, European Central Bank, November.

Angeloni, I, L Aucremanne, M Ehrmann, J Galí, A Levin and F Smets (2004), 'Inflation persistence in the Euro Area: Preliminary summary of findings', unpublished manuscript, European Central Bank, November.

Aucremanne, L and E Dhyne, (2004), 'How frequently do prices change? Empirical evidence based on the micro data underlying the Belgian CPI', ECB Working Paper no. 331.

Bai, J (1999), 'Estimation of a change point in multiple regression models', *Review of Economics and Statistics*, 79, pp 551-563.

Ball, L, (1994), 'Credible disinflation with staggered price-setting', *American Economic Review*, 84(1), pp 282–89.

Batini, N (2002), 'Euro area inflation persistence', ECB Working paper no. 201.

Batini, N and E. Nelson (2001), 'The lag from monetary policy actions to inflation: Friedman revisited', *International Finance*, 4(3), pp 381–400.

Benigno, P (2004), 'Optimal Monetary Policy in a Currency Area', *Journal of International Economics*, 63(2), pp 293-320.

Bilke, L (2004), 'Break in the mean and persistence of inflation: a sectoral analysis of French CPI', ECB Working paper no. 463.

Bils, M and P Klenow (2002), 'Some evidence on the importance of sticky prices', NBER Working paper no. 9069.

Calvo, G (1983), 'Staggered prices in a utility-maximizing framework', *Journal of Monetary Economics*, 12(3), pp 383–98.

Caplin, A and J Leahy, (1997), 'Aggregation and optimization with state-dependent pricing', *Econometrica*, 65, pp 601-23.

Caplin A. and D. Spulber (1987), 'Menu-costs and the neutrality of money', *Quarterly Journal of Economics*, 102, pp 703-26.

Clark, T (2003), 'Disaggregated evidence on the persistence of consumer price inflation', Federal Reserve Bank of Kansas City Working Paper, 03-11.

Cogley T and T Sargent (2001), 'Evolving post-World War II inflation dynamics', *NBER Macroeconomics Annual*, vol 16, MIT Press.

Corvoisier, S and B Mojon (2005), 'Breaks in the mean of inflation: how they happen and what to do with them', ECB Working Paper no. 451.

Debelle, G and J Wilkinson (2002), 'Inflation targeting and the inflation process: Some lessons from an open economy', Reserve Bank of Australia Research Discussion Paper 2002-01.

Dhyne, E, L Álvarez, H LeBihan, G Veronese, D Dias, J Hoffman, N Jonker, P Lünnemann, F Rumler, and J Vimunen (2005), 'Price setting in the Euro Area: Some stylised facts from individiual consumer price data,' ECB Working Paper, forthcoming.

Erceg, C and A Levin (2003), 'Imperfect credibility and inflation persistence', *Journal of Monetary Economics*, 50(4), pp 915-44.

Ernst, E and B Mojon (2004), 'The determinants of price persistence, part 2. The case of sectoral prices in the euro area', unpublished paper, ECB.

Fuhrer, J and G Moore (1995), 'Inflation persistence', *Quarterly Journal of Economics*, 110, pp 127–59.

Fuhrer, J (2005), 'Intrinsic and inherited inflation persistence,' unpublished manuscript, Federal Reserve Bank of Boston, May.

Gadzinski, G and F Orlandi (2004), 'Inflation persistence in the European Union, the euro area, and the United States', ECB Working paper no. 414.

Galí, J (2004), 'Has the Inflation Process Changed? A Comment' presented at the Third BIS Annual Conference, Brunnen, Switzerland, 18-19 June.

Hansen, B (1999), 'The grid bootstrap and the autoregressive model', *Review of Economics and Statistics*, 81, pp 594-607.

Issing, O (2005), 'One size fits all! A single monetary policy for the euro area', speech given to the Third International Research Forum on Monetary Policy. http://www.ecb.int/press/key/date/2005/html/sp050520.en.html

Levin, A and J Piger (2003), 'Is inflation persistence intrinsic in industrial economies?', Federal Reserve Bank of St. Louis working paper 2002-023, updated version.

Lucas, R (1972), 'Expectations and the neutrality of money', *Journal of Economic Theory*, 4, pp 103-24.

Mankiw, NG and R Reis (2002), 'What Measure Should a Central Bank Target?', NBER Working Paper No. 9375.

O'Reilly, G and K Whelan (2004), 'Has Euro-area inflation persistence changed over time?', ECB Working paper no. 335.

Orphanides, A and J Williams (2003), 'Imperfect knowledge, inflation expectations, and monetary policy', in M. Woodford (ed), *Inflation Targeting*, Chicago: University of Chicago Press.

Perron, P (1989), 'The great crash, the oil-price shock, and the unit-root hypothesis', *Econometrica*, 57, pp 1361-1401

Quandt, R (1960), 'Tests of the hypothesis that a linear regression obeys two separate regimes', *Journal of the American Statistical Association*, 55, 324-330.

Robalo Marques, C (2004), 'Inflation persistence: facts or artefacts?', ECB Working paper no. 371.

Rudd, J and K Whelan (2001), 'New tests of the New-Keynesian Phillips curve,' Federal Reserve Board, Finance and Economics Discussion Series no. 2001-30.

Sargent, T, (1999), *The Conquest of American Inflation,* Princeton: Princeton University Press.

Stock, J. (2001), Comment on Cogley and Sargent, *NBER Macroeconomics Annual* 2001, Vol 16, MIT Press.

Taylor, J (1980), 'Aggregate dynamics and staggered contracts', *Journal of Political Economy*, 88, pp 1-24.

Taylor, J (1998), 'Monetary policy guidelines for unemployment and inflation stability,' in J Taylor and R Solow (eds.), *Inflation, Unemployment, and Monetary Policy*, Cambridge: MIT Press.

Taylor, J (1999), 'Staggered price and wage setting in macroeconomics', Chapter 15 in J Taylor and M Woodford (eds), *Handbook of Macroeconomics*, Elsevier.

Taylor, J. (2000), 'Low inflation, pass-through, and the pricing power of firms', *European Economic Review*, 44, pp 1389–1408.

Whelan, K (2004), 'Staggered price contracts and inflation persistence: some general results', ECB Working paper no. 417.

Woodford, M (2003), *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Appendix: Data Sources

Price data:

All data were seasonally adjusted using X-12, except for the US where the data were seasonally adjusted by the Bureau of Labour Statistics and Germany where the data were seasonally adjusted by the Bundesbank.

Australia: *Quarterly* CPI data published by the Australian Bureau of Statistics data, provided by the RBA. Sample from September 1969 to September 2003 for CPI, food, alcohol, clothing, housing, furniture (household goods), and transport; from September 1980 for communication; from March 1982 for education, recreation; from September 1989 for health, miscellaneous.

Austria: HICP data published by Eurostat, provided by the ECB. Sample from January 1987 to January 2003.

Belguim: HICP data published by Eurostat, provided by the ECB. Sample from January 1991 to February 2003 for all components excluding education from January 2000 to February 2003.

Canada: CPI data published by Statistics Canada, provided by the Bank of Canada and the BIS. Sample from January 1984 to January October 2003.

Chile: CPI data published by the National Bureau of Statistics, provided by Banco Central de Chile. Sample from April 1989 to December 2003 for the CPI, January 1997 to December 2003 for food, housing, furniture (home appliances), clothing, transportation, health, education and leisure and miscellaneous.

European Union: HICP data published by Eurostat, provided by the ECB. Sample from January 1990 to January 2003 for CPI, food, alcohol; from January 1995 to January 2003 for other components.

Finland: HICP data published by Eurostat, provided by the ECB. Sample from January 1987 to January 2003 for CPI, food, alcohol; from January 1995 to January 2003 for other components.

France: HICP data published by Eurostat, provided by the ECB. Sample from January 1990 to January 2003.

Germany: CPI data provided by the Bundesbank. Sample from January 1962 to December 2004, with data for West Germany used prior to 1995 and pan-German data used post 1995.

Italy: HICP data published by Eurostat, provided by the ECB. Sample from January 1977 to February 2003.

Luxembourg: HICP data published by Eurostat, provided by the ECB. Sample from January 1995 to February 2003.

Netherlands: HICP data published by Eurostat, provided by the ECB. Sample from January 1995 to February 2003, except the aggregate HICP from October 1987 to February 2003.

New Zealand: *Quarterly* CPI data published by Statistics New Zealand, provided by the RBNZ. Sample from March 1975 to September 2003 for CPI, food, clothing

(apparel), transportation, alcohol, household operation, credit services; from December 1988 to September 2003 for housing, recreation and education, health.

Portugal: HICP data published by Eurostat, provided by the ECB. Sample from January 1987 to January 2003; except housing, furniture, recreation, restaurants and miscellaneous: January 1988 to January 2003; education January 1995 to January 2003.

Spain: HICP data published by Eurostat, provided by the ECB. Sample from January 1992 to February 2003.

Sweden: CPI data from the Riksbank website. Sample from January 1980 to February 2005.

United Kingdom: CPI data published by the Office for National Statistics, provided by the Bank of England. Sample from January 1980 to February 2004, except personal and leisure goods from January 1987 to February 2004.

United States: CPI data published by the Bureau of Labour Statistics, provided by the Cleveland Fed. Sample from January 1978 to March 2004.