Inflation persistence and the rationality of inflation expectations

Brissimis, Sophocles and Migiakis, Petros

Bank of Greece

December 2010

Inflation persistence and the rationality of inflation expectations

Sophocles N. Brissimis
Bank of Greece
Economic Research Department
and University of Piraeus
Tel.: 0030 210 3202388
Email: sbrissimis@bankofgreece.gr

Petros M. Migiakis
Bank of Greece
Economic Research Department
Tel.: 0030 210 3203587
Email: pmigiakis@bankofgreece.gr

Abstract

The rational expectations hypothesis for survey and model-based inflation forecasts − from the Survey of Professional Forecasters and the Greenbook respectively − is examined by properly taking into account the persistence characteristics of the data. The finding of near-unit-root effects in the inflation and inflation expectations series motivates the use of a local-to-unity specification of the inflation process that enables us to test whether the data are generated by locally non-stationary or stationary processes. Thus, we test, rather than assume, stationarity of near-unit-root processes. In addition, we set out an empirical framework for assessing relationships between locally non-stationary series. In this context, we test the rational expectations hypothesis by allowing the co-existence of a long-run relationship obtained under the rational expectations restrictions with short-run "learning" effects. Our empirical results indicate that the rational expectations hypothesis holds in the long run, while forecasters adjust their expectations slowly in the short run. This finding lends support to the hypothesis that the persistence of inflation comes from the dynamics of expectations.

Keywords: Inflation; rational expectations; high persistence
JEL classification: C50; E31; E52

1. Introduction

Inflation expectations are known to play a key role in the transmission of monetary policy to the economy. Dynamic stochastic general equilibrium (DSGE) models, which are now considered the workhorse of monetary analysis, assume that expectations are formed rationally. This is essentially the assumption that individuals and firms behave so as to avoid systematic mistakes, i.e. they succeed, by some process known to themselves, in eliminating systematic components from their expectational errors (see McCallum, 2010). Theoretical and empirical research, however, has shown that this may not always be the case.

Brissimis and Magginas (2008) in the context of the New-Keynesian Phillips curve (NKPC) found that inflation forecasts, used as measures of inflation expectations, deviate from rationality because when the assumption of rationality is imposed in estimation, this necessitates the inclusion of lagged inflation in the model (hybrid NKPC), which appears to account for these deviations. This study, however, did not explore further the nature of the deviations from rationality. An alternative to the rational expectations hypothesis is the adaptive learning hypothesis that focuses on the process by which agents learn in forming their expectations and this under certain conditions may lead to rationality. Adaptive learning has had so far limited empirical application (see e.g. Chevillon et al., 2010). Finally, behavioral economics, a new field of economics, making use of concepts from other social sciences, suggests that economic agents may not always be the rational, optimizing agents we assume in theoretical models. Incorporating this assumption into macro-models can make a major difference in the way these models work (Akerlof, 2007).

How important are, then, deviations from rationality and are they short-lived or do they have a permanent component? Moreover, how should deviations from rationality be modeled and embedded in macro-models? Answering these questions will give useful insights into the working of monetary policy and its transmission mechanism and might further enhance the realism of models used for policy purposes (Mishkin, 2010).

This paper investigates the issue of rationality of inflation expectations in the US economy in view of the empirical regularity that the inflation rate displays high persistence. We first examine the degree of persistence of both the actual and forecasted inflation series, the latter being used as proxy for expected inflation, so as to properly assess the permanent or temporary nature of the effect of shocks on inflation and thus highlight the "inflation inertia" problem. Departing from the literature on near unit roots (e.g. Elliott et al., 1996; Lanne, 2001), we address the question whether inflation can be treated as a strictly stationary process (as would be expected a priori on theoretical grounds), given recent econometric advances that permit the distinction of highly persistent series into stationary and locally non-stationary ones (Phillips et al., 2001; Lima and Xiao, 2007). Thus we accurately estimate whether the parameter measuring the degree of inflation persistence falls within the stationarity boundaries.

Based on the proper assessment of the persistence properties of inflation, we proceed to test the rationality of inflation expectations. The empirical analysis of the rationality hypothesis allows for a learning process to evolve in the short run that might be converging towards a rational expectations equilibrium. Thus, we specify a framework which separates the formation of expectations in the long run from short-run learning dynamics that involve adjustment in expectations.

Our work contributes to the literature on inflation expectations in three respects. First, it examines the rational expectations hypothesis by using an econometric methodology that takes proper account of the persistence properties of the inflation and inflation forecasts data series. Second, it bridges two distinct strands of the literature by allowing for short-run learning dynamics that involve adjustment in inflation expectations to be combined with a rational expectations relationship in the long run. Third, our empirical analysis indicates that rationality of inflation expectations exists in the long run but the learning process develops slowly in the short run, a finding that may be viewed as a root cause of the high inflation inertia observed in the real world.

The rest of the paper is organized as follows: Section 2 reviews the different approaches that have been followed for testing rationality in relation to the degree of inflation persistence. Section 3 sets out the framework to be used for examining rationality as an equilibrium hypothesis in light of the persistence properties of our inflation and inflation expectations data, the sources of possible permanent deviations from rationality and the learning mechanism that may operate in the short run. Section 4 reports the empirical findings on persistence and rationality, while section 5 concludes.

2. Literature review

2.1 Previous findings

The empirical investigation of the hypothesis that inflation expectations are formed rationally has attracted a lot of interest in the past. However, the conclusions of existing studies on the rationality of inflation forecasts, as measures of these expectations, are divergent due to differences in their scope and primarily in their empirical methodology. In this section we summarize the findings of the main relatively recent studies. We start by reviewing previous work that examines the rationality of expectations by assuming that the inflation and inflation forecast series are stationary. Next, we discuss findings from studies that treat inflation data as non-stationary. Finally, we briefly review some of the literature, first, on the persistence of inflation and, second, on adaptive learning, which can be viewed as an alternative to the rational expectations hypothesis.

The strand of literature that deals with rationality of inflation forecasts by using methodologies appropriate for stationary processes provides mixed results. Thomas (1999) concludes that the Survey of Professional Forecasters (SPF) provides unbiased estimators of inflation, a conclusion which is supported by the findings of Romer and Romer (2000); the latter study also provides evidence of rationality for the Data Resources Inc. (DRI) measure of inflation expectations. However, the rationality hypothesis is not confirmed for the inflation forecasts obtained from the Blue Chip Economic Indicators, while the Fed's forecasts contained in the Greenbook do not fail the rationality test and moreover are found to incorporate systematically superior information than professional measures (see Romer and Romer, 2000). Consumer forecasts have also been analyzed by employing empirical frameworks for stationary data. Thomas (1999), using a measure of consumers' inflation expectations from the University of Michigan Survey, finds that consumers form their expectations rationally, while Carroll (2003), using the same measure, rejects rationality. However, this conflicting evidence may in part be the result of the different sample period, as Thomas's sample starts in 1960 and Carroll's in 1981. Finally, Andolfatto et al. (2008), using simulation analysis of a DSGE model of the New-Keynesian type, argue that conventional econometric tests frequently and incorrectly reject the null hypothesis of stationarity of inflation and the unbiasedness of inflation forecasts, due to the small sample size.

On the other hand, a number of studies place the examination of the rationality of inflation forecasts in a co-integration framework. In this case, consumers' inflation expectations are found to be rational (Grant and Thomas, 1999; Forsells and Kenny, 2004), although the same set of data mentioned above (i.e. from the University of Michigan Survey) is used. In contrast, the use of co-integration techniques for the investigation of the rationality of professional forecasts is rather limited, as only Grant and Thomas (1999) assess the Livingston survey measures of expectations with the use of these techniques. Their results are in line with those reported in papers using the stationarity assumption, giving support to the rationality hypothesis. The full specification of their model also involves a gradual adjustment of both inflation and inflation expectations towards the rational expectations equilibrium. Finally, the work of Berk and Hebbink (2006), examining consumers' inflation expectations for various European countries, also falls in this category; their empirical results show that inflation expectations and future realizations of inflation are co-integrated and that the forecast error is stationary.

The different treatment of the inflation rate may be due to the properties of the data series of inflation, which has been reported to exhibit a high degree of persistence. Several explanations have been put forward for this property of the inflation rate. Specifically, Gali and Gertler (1999) relate inflation persistence to the existence of a number of firms in the economy that set their prices according to a backward-looking rule, rather than in a purely forward-looking fashion. Mankiw and Reis (2001) attribute inflation persistence to the slow diffusion of information on economic conditions, which is due to costs of acquiring or acting on new information. In a different vein, the characteristic of persistence underlies the questioning by Roberts (1998) and Ball (2000) of the rationality of expectations formation. The evidence in Brissimis and Magginas (2008) appears to support this view. A distinction, however, which has not yet been pursued in the inflation persistence literature, is the one between a stationary and a locally non-stationary highly persistent inflation (or inflation expectations) series. A discussion of this distinction will be made in the following section.

A final issue concerns the adaptive learning hypothesis for inflation expectations, which has been treated in the literature as an alternative to the traditional rational expectations hypothesis. With adaptive learning, the private sector is assumed to form its expectations based on the past behavior of inflation and a set of information it finds relevant (besides the lagged value of inflation) and, since it has limited knowledge of the precise working of the economy, it updates its knowledge and the associated forecasting rule when new data becomes available. The updating also involves the parameters of the forecast function, which are adapted recursively.

Adaptive learning represents a relatively modest deviation from rational expectations resulting from continuous learning on the part of economic agents with imperfect knowledge (Orphanides and Williams, 2003). In this sense, agents can be considered to have bounded rationality in making forecasts. As noted by Evans and Honkapohja (2001), "in many cases learning can provide at least an asymptotic justification for the rational expectations hypothesis" (p.13). Moreover, it appears that adaptive learning models are able to reproduce important features of empirically observed measures of inflation expectations (Gaspar et al., 2009).

2.2 Discussion of the different methodological approaches

The characterization of the inflation data-generating process, in view of its high persistence, is at the root of the differences observed in empirical studies on the rationality of inflation expectations. Thus, a number of different approaches have been followed corresponding to different empirical methodologies. In particular, while in some studies the degree of inflation persistence is assumed to be sufficiently high for the process to be considered non-stationary and this results in the use of co-integration to examine the unbiasedness of survey-based inflation forecasts (see Grant and Thomas, 1999; Forsells and Kenny, 2003 and Berk and Hebbink, 2006, among others), several other papers employ methodologies that presume stationarity (see e.g. Bryan and Palmqvist, 2005 and Andolfatto et al., 2008). Although inflation has been found to be a near-unit-root series, i.e. to have an autoregressive root close to unity (e.g. Lanne, 2001; Gagnon, 2008), it has not been examined whether it is a stationary or a locally non-stationary process.

In order to investigate whether the inflation data contain a unit root, several tests have been proposed, with the standard Dickey-Fuller and Philips-Perron being the most frequently used. Nevertheless, the power of standard unit root tests has been questioned based on the argument that the finding of a unit root in highly persistent series may be attributed to the low power of conventional unit root tests. As a result, modifications of the standard tests have been proposed (Elliott, 1998; Elliott et al., 2002; Ng and Perron, 2001), so as to distinguish unit root processes from stationary ones, when autoregressive roots approach unity.

Also, a number of other tests distinguish between local and global unit root effects. On the one hand, there is the notion of local non-stationarity (Kapetanios et al., 2003) that refers to sub-sample non-stationarity against stationarity in the whole sample (global stationarity). On the other hand, recently popularized local-to-unity models (Phillips et al., 2001; Lima and Xiao, 2007) allow the distinction between local non-stationarity in the sample and stationarity in the population.

Phillips et al. (2001) have suggested a local-to-unity model as a new specification of the first-order autoregressive process that enables the researcher to identify the persistence properties of the time series data more accurately. In this model, the stationarity or explosiveness of the time series process corresponds only to extreme cases. The authors argue that in the case of a sufficiently large autoregressive coefficient the process is better specified as varying between stationarity and non-stationarity. They highlight that a near-unit-root process may not always converge towards stationarity but may exhibit similarities to a unit root process as the sample size tends to infinity. Lima and Xiao (2007) modified the local-to-unity model of Phillips et al. (2001), arguing that any in-sample non-stationarity properties of an economic series (local non-stationarity) should dissipate as the sample tends to infinity. Therefore, the question of whether a process governing a time series is stationary or non-stationary needs careful examination as highly persistent time series may be locally non-stationary.

In this paper, we set out a framework for testing relationships between pairs of locally non-stationary series. In particular, in view of initial findings indicating rejection of stationarity in favor of local non-stationarity in the inflation and inflation forecasts data, we apply the test of Lima and Xiao (2007) to the residuals of the relationship between inflation and inflation expectations in order to examine the rationality hypothesis. This framework permits the separation of the long-run relationship from short-run adjustment (learning) dynamics; it can be seen as the standard Engle-Granger (1987) co-integration framework, applied to locally non-stationary data.

3. Theoretical and methodological issues

The present paper examines the rationality of inflation expectations by properly taking into account the degree of persistence in the underlying series. Our data set includes inflation and two measures of inflation expectations calculated in terms of the GDP deflator for the United States; all variables are expressed in real time and at a quarterly frequency. The first measure of inflation expectations consists of survey responses of professional forecasters about inflation expected one quarter ahead, as collected by the Federal Reserve Bank of Philadelphia and compiled in the Survey of Professional Forecasters; the sample extends from 1966Q3 to 2009Q2. The second measure consists of inflation forecasts of the Federal Reserve Board, as contained in the Greenbook and these span the period 1969Q1 - 2004Q1.[1]

Let $\pi^e_{t+h}|\Omega_t$ denote expectations of the inflation rate (π) h periods ahead, formed at time *t* on the basis of the information set, Ω_t, and π_{t+h} the actual inflation rate at time *t+h*. If expectations are formed rationally, this implies the testable hypothesis that $\pi^e_{t+h}|\Omega_t$ unbiasedly predicts actual inflation at time *t+h* (see Grant and Thomas, 1999; Bryan and Palmqvist; 2005; Andolfatto et al., 2008), i.e.

$$\pi_{t+h} = \pi^e_{t+h}|\Omega_t \tag{1}$$

In order to test for the rationality of inflation expectations, inflation forecasts obtained from the sources indicated above are used. Assuming the forecast horizon is one period (*h*=1), the rationality hypothesis can be tested by the following regression:

$$\pi_{t+1} = \alpha + \beta\pi^e_{t+1} + e_{t+1}, \qquad e_t \sim \mathrm{N}\,(0,\ \sigma^2) \tag{2}$$

For the rationality hypothesis to hold, we require $(\alpha \quad \beta) = (0 \quad 1)$, which constitutes our null hypothesis ($\mathrm{H}_0 : (\alpha \quad \beta) = (0 \quad 1)$). The examination of this hypothesis relies crucially on the time series properties of the data. In view of the high persistence of inflation, the correct specification of the inflation process has significant implications for the robustness and interpretation of the empirical findings.

3.1 Inflation persistence

[1] The Greenbook forecasts are publicly released with a 5-year lag.

To set the stage for a discussion of persistence and of the corresponding empirical framework for testing rationality, consider the general specification of a first-order autoregressive – AR (1) – process:

$$y_t = \rho y_{t-1} + \varepsilon_t \tag{3}$$

where y stands for the series of inflation or inflation expectations, t denotes time ($t = 1,2,...,n$), ρ is the autoregressive coefficient and ε is an error term. When y is stationary, its autoregressive coefficient should lie within the unit circle ($|\rho|<1$) but if $|\rho| \geq 1$, the data generation process of y is governed by a unit or explosive root; in the latter case shocks to ε will have permanent effects.

The standard unit root tests (Dickey and Fuller, 1979 and Philips and Perron, 1988) examine the null $\mathrm{H}_0 : |\rho| = 1$ against the alternative $\mathrm{H}_1 : |\rho| < 1$. Recently, there has been evidence (e.g. Elliott et al., 1996, Lanne, 1999 and Ng and Perron, 2001) that these tests are not powerful enough to distinguish unit root processes from highly persistent stationary ones. Thus, Elliott et al. (1996) modified standard unit root tests by de-trending the series and this allows the distinction between unit root and stationary near-unit root processes. Finally, Kapetanios et al. (2003) proposed a test that examines non-stationarity in subsamples against stationarity in the whole sample. Their test, called the t_{NL} test, involves the following t-type test statistic:

$$t_{NL} = \hat{\delta} \Big/ \sigma(\hat{\delta}) \tag{4}$$

where δ and σ are the coefficient and its standard error respectively estimated from the regression:

$$\Delta y_t = \delta y_{t-1}^3 + u_t \tag{5}$$

The authors tabulated critical values for series with zero mean (case 1), non-zero mean (case 2) and linear trend (case 3).

Still, examining the order of integration of the series (i.e. whether they are I(1) or I(0)) may be misleading, as it does not allow the distinction of highly persistent processes into stationary and locally non-stationary ones (Phillips et al., 2001 and Lima and Xiao, 2007). Instead, the local-to-unity model has been proposed for estimating the degree of persistence. Consider the simple local-to-unity model (Phillips et al., 2001):

$$\rho = 1 + \frac{c}{n} \tag{6}$$

In this model, the distinction between stationary and unit root processes is related to the localizing parameter *c*. According to Phillips et al. (2001), for a finite sample the model accommodates a wide range of values of ρ depending on *c*; the process is stationary if c<0, contains a unit root if *c*=0, while it is explosive if c>0. However, for a given $c \in (-\infty, +\infty)$, as $n \to \infty$, the process will contain a root that converges to unity.

Lima and Xiao (2007), arguing that a shock is likely to affect economic time series for a long period but not for ever, have modified eq. (6) as follows:

$$\rho = 1 + \frac{c}{n^d} \tag{7}$$

This model contains the parameter $d \in D$, with $D \subset (0,1]$, capturing the degree of persistence; if *d*=1 eq. (7) is identical to eq. (6). The fact that eq. (7) provides a more complete categorization of the persistence characteristics of eq. (3) is understood if we consider that for a given *c*<0, with $d \neq 0$, the process approaches stationarity with a rate equal to $n^{d/2}$, as $n \to \infty$ (note that $y_n = O(\sqrt{n^d})$). As in eq. (6), this specification leads to standard stationary or explosive processes only in extreme cases. Also, for a given sample, eq. (7) may lead to local non-stationarity of y in the sense that the confidence interval of the autoregressive coefficient ρ includes a range of values which exceed unity. Thus, the series may be subject to in-sample non-transitory effects that will, however, eventually dissipate (see Kim and Lima, 2010).

From the above it is clear that the hypothesis of local non-stationarity should also be examined, instead of assuming stationarity of near-unit-root processes. To this end, Lima and Xiao (2007) developed a test for examining the null hypothesis that the time series is stationary ($d = 0$) against the alternative hypothesis that it is locally non-stationary ($0 < d < 1$). The test statistic is the following:

$$Q = \underset{1 \le u \le n}{Max} \frac{1}{\sqrt{n}} \frac{1}{\hat{\omega}_y} \left| \sum_{t=1}^{u} y_t - \frac{u}{n} \sum_{t=1}^{n} y_t \right| \tag{8}$$

In (8), u denotes points in time and $\hat{\omega}_y$ is the long-run standard deviation of y, estimated consistently by using non-parametric kernel smoothing. It has been shown (Lima and Xiao, 2007) that the results of standard unit root tests are robust if $\omega_y^2 = \sigma_y^2$, with σ_y^2 denoting the sample variance. Lima and Xiao (2007) computed and tabulated critical values for the above test.

3.2 Rational expectations: the long-run perspective

The local-to-unity model outlined above appears to be able to reconcile conflicting views about the methodological treatment of inflation persistence (see Section 2) and deal with problems regarding the testing of the rational expectations hypothesis when the data series are highly persistent. As we discussed earlier, the rational expectations hypothesis can be tested using eq. (2) considered as a long-run relationship; this means that, if rational expectations hold, any forecasting errors will be short-lived and will not persist in the long run.

If evidence is found of local non-stationarity in the inflation and inflation forecast series, this would imply that, to test for rationality, methodologies suitable for stationary series should not be used since in that case the estimated coefficients would be inconsistent. Conventional co-integration analysis would similarly not be appropriate since it relies on exact unit root inference (see Hjalmarsson and Osterholm, 2010). When our data are locally non-stationary, the residuals of eq. (2) calculated under the rational expectations restrictions, i.e.

$$e_{t+1} = \pi_{t+1} - \pi_{t+1}^{e} \tag{9}$$

should be tested for stationarity against local non-stationarity by applying the Q test. Acceptance of stationarity of these residuals would confirm the validity of the rational expectations hypothesis for the long run. In contrast, rejection of stationarity would indicate a non-transitory bias in the formation of inflation expectations.

In particular, if $\alpha \neq 0$ and/or $\beta \neq 1$ in equation (2), forecasters make systematic errors and deviations of forecasted from actual inflation are permanent. Hence, the estimation of the parameters α and β enables the assessment of the source of the long-run deviations from rationality. Thus, if $\alpha \neq 0$, there exists a constant deviation between actual and forecasted inflation, which could be related to false perceptions of the monetary authorities' (implicit) inflation target. According to Orphanides and Williams (2005), imperfect knowledge of the inflation target may introduce a bias in private sector's forecasts of inflation. Also, deviations from rationality may be related to shifts in inflation trends, which would result in $\beta \neq 1$. For example, in periods of a declining inflation trend, forecasters who base their estimates on past inflation rates, have been reported to overestimate future inflation (see Clarida et al., 2000).

3.3 *'Learning' effects: the short-run dynamics*

Finally, the existence of 'learning' effects in the short-run formation of expectations may be considered. Instantaneous perfect knowledge of future inflation rates is a very stringent assumption; in reality agents do not have perfect knowledge, but rather rely on the formulation of models for producing inflation forecasts. Their forecasts are adjusted as new data become available and forecasters take note of past deviations of their forecasts from actual inflation and gradually learn to utilize the information contained in the errors. In other words, agents engage in learning processes about inflation as they attempt to improve their knowledge of inflation.

The empirical importance of learning effects could be assessed by the following error-correction type of adjustment equation.

$$\Delta\pi_{t+1}^{e} = b_0 + b_1 e_t + \sum_{i=0}^{l} b_{2i}\Delta\pi_{t-i}^{e} + \sum_{i=0}^{l} b_{3i}\Delta\pi_{t-i} + u_{t+1}\ , \tag{10}$$

where e are the residuals from eq. (2). In the above equation, the learning process is understood as an adjustment toward the long-run relationship (2). Thus, the coefficient b_1 ($b_1 < 0$) can be interpreted in terms of the speed with which forecasters adjust their expectations from one period to the next. If stationarity of the residuals in eq. (9) is accepted, this implies that inflation forecasts exhibit rationality in the long run and any deviations of forecasted from actual inflation are temporary, reflecting a learning process that takes place in the short run. If, however, the restrictions imposed in eq. (9) by the rational expectations hypothesis are rejected, forecasters may still adjust in the short run their expectations, which in this case will not converge to rational expectations equilibrium but rather toward a long-run relationship, eq. (2), representing a permanent deviation from rationality.

The findings regarding the speed of the learning process are related to the efficiency of the tools used to forecast future inflation. The higher the adjustment coefficient b_1, in absolute terms, the faster the correction of previous errors and as a result the learning process would be more efficient. On the other hand, if expectations are revised slowly, this may provide an explanation for the high inertia characterizing inflation.

4. Empirical results

In this section we examine the rational expectations hypothesis for the Survey of Professional Forecasters and the Greenbook inflation forecasts. The section is divided in two subsections: the first presents the findings of our analysis on the inflation and inflation forecasts data properties, while the second attempts to provide an answer to our central question, namely whether inflation expectations are formed rationally, and also assess the efficiency of the forecasters in correcting their past errors.

4.1 Persistence of inflation and inflation expectations

In order to examine the rational expectations hypothesis, we first need to have a clear view on the properties of the data as far as persistence is concerned. This is an essential step because the choice of the methodology to be employed for our main investigation depends on whether the data is stationary or not.

We first examine the data for unit roots by applying both the standard tests (DF and PP) and the modified ones (DF-GLS and PP-GLS), so as to admit cases of near unit roots. Table I below reports the relevant findings. Standard unit root tests fail to reject the unit root null hypothesis for each of the inflation and forecasted inflation series.

However, the estimation of the autoregressive coefficients shows that ρ takes values very close to, but lower than, unity. These results motivate further investigation of the properties of the data. For this purpose, we apply the modified unit root tests (DF-GLS and PP-GLS) that lead to the rejection of the null hypothesis in favor of stationarity. Thus, the data appear to contain near unit roots (for a similar result for inflation see e.g. Lanne, 2001 and Gagnon, 2008).

[Table I]

However, the finding that the autoregressive coefficients lie in a space close to unity ($\rho \in [1/2, 1]$) does not necessarily imply, according to Phillips et al. (2001), that the series are stationary; in this case the stationarity property should be examined in the context of a local-to-unity model. Here we use Lima and Xiao's (2007) version of the local-to-unity model and test for stationarity against local non-stationarity.

[Table II]

Lima and Xiao's test is based on the consistent estimation of the long-run variance of the data series so as to approximate well the properties of the population. The long-run variance of the series and the autoregressive coefficients are estimated consistently by using Epanechnikov kernel smoothing with bandwidth given by Silverman's rule.[2] The results are presented in Table II.

For all the series, the difference between the long-run and the sample variance indicates that $d \in (0,1]$, which signifies the existence of local non-stationarity in the data. Thus, both standard and modified unit root tests do not capture the population properties accurately (i.e. they are not efficient for $n \to \infty$). A consistent estimation of the autoregressive coefficient for all three series reveals that if the specification of eq. (7) is adopted the size of these estimates is smaller compared to the OLS estimates presented in Table I.[3] This indicates that the persistence of actual and forecasted inflation is smaller in the population than in the sample.

---

[2] Recall from section 3.1 that, if the sample variance is equal to the long-run variance, conventional unit root tests are robust.

[3] Note that for the inflation rate, these findings are very close to the ones reported in Lima and Xiao (2007).

[Table III]

We formally tested for local non-stationarity by employing the Q test of Lima and Xiao. The results are reported in Table III above. The Q test statistic indicates rejection of the null of stationarity in favor of local non-stationarity at conventional levels of statistical significance. Also, from the results of Table III one can see that the SPF and Greenbook inflation forecasts have similar properties to inflation regarding persistence. Local non-stationarity in the sense of Kapetanios et al. (2003), i.e. non-stationarity in part(s) of the sample, was further tested by employing the t_{NL} test. The results displayed in Table III show that global stationarity is rejected, while the existence of local non-stationarity is accepted. Having established the local non-stationarity property of the inflation and inflation expectations series, the next step was to investigate the rational expectations hypothesis.

4.2 Tests of the rational expectations hypothesis

Based on the finding that the inflation and inflation expectations series are individually locally non-stationary, we evaluated the rationality of inflation expectations by focusing on the residuals of eq. (2), which were calculated under the rational expectations restrictions (see eq. 9). These residuals were tested for stationarity against the alternative hypothesis of local non-stationarity with the Q test. The results are presented in Table IV for the SPF and Greenbook forecasts. In both cases the null hypothesis that the residuals are stationary is accepted suggesting that forecasters (professional forecasters and the Fed) do not make systematic errors over a long horizon in making their forecasts, i.e. they form their expectations rationally.[4]

[Table IV]

This finding is in line with the results reported by Romer and Romer (2000) and Thomas (1999), although these studies rely on a treatment of inflation appropriate for stationary processes and do not examine the short-run inflation and inflation expectations dynamics.

[4] Technically, these findings imply that the common (locally) non-stationary component of the inflation and forecasted inflation series is eliminated when we take the difference between the two.

[Table V]

Next, we turn to the estimation of the short-run adjustment dynamics. Table V presents the results concerning the estimated from eq. (10) speed of adjustment parameter for inflation expectations. These results confirm the existence of a learning process by uncovering significant adjustment dynamics toward the long-run rational expectations relationship. By modeling expectations using a learning process, we obtain that expectations seem to be moving slowly in response to forecasting errors toward rational expectations. The size of the estimated speed of adjustment parameter indicates that for both the SPF and Greenbook forecasts it takes approximately 7 quarters for the full adjustment to take place. This is a notable result showing that inflation is not inherently sticky but instead its inertia derives from the high persistence of the forecasting errors.

Overall, our findings indicate that forecasters form their expectations rationally in the long run, a result that is robust to criticism on the need to properly specify the data's degree of persistence. Furthermore, the proposed framework allows the formation of rational expectations in the long run to co-exist with learning effects in the short run. Note, however, that, unlike other learning mechanisms, our specification leads to adjustment toward the rational expectations equilibrium relationship.

5. Concluding remarks

In the present paper we examined the rational expectations hypothesis for inflation forecasts, having estimated accurately the persistence characteristics of the data series with the use of local-to-unity models. In particular, after running a battery of tests in order to specify the degree of persistence of the data, we found that the series of inflation and inflation forecasts have large autoregressive roots, which exceed unity locally. Thus, applying a local-to-unity framework that distinguishes in-sample effects from those valid for the population, we have found that shocks to the inflation process will produce locally non stationary effects, which however are not permanent.

In light of the above, we have laid out a framework that assesses the rational expectations hypothesis as a long-run relationship. Our findings indicate that expectations of US inflation, contained in the SPF and the Greenbook, are not systematically biased, confirming the rational expectations hypothesis. Furthermore, this framework, in the context of an error correction mechanism, allows forecasters to learn from past errors so that short-run deviations of forecasted from actual inflation will eventually dissipate. Quantifying this error correction process reveals that forecasters correct their errors, relatively slowly, with a mean lag of approximately seven quarters, a result that might well provide an explanation for the stylized fact of high inflation inertia.

References

Akerlof, G.A., 2007. The missing motivation in macroeconomics. American Economic Review 97, 5-36.

Andolfatto, D., Hendry, S., Moran, K., 2008. Are inflation expectations rational? Journal of Monetary Economics 55, 406-422.

Ball, L., 2000. Near-rationality and inflation in two monetary regimes. Working Paper Series 7988, National Bureau of Economic Research.

Berk, J.M., Hebbink, G., 2006. The anchoring of European inflation expectations. Working Papers Series 116, Netherlands Central Bank.

Brissimis, S.N., Magginas, N.S., 2008. Inflation forecasts and the New Keynesian Phillips curve. International Journal of Central Banking 4, 1-22.

Bryan, M.F., Palmqvist, S., 2005. Testing near-rationality using detailed survey data. Working Paper Series 05-02, Federal Reserve Bank of Cleveland.

Carroll C.D., 2003. Macroeconomic expectations of households and professional forecasters. Quarterly Journal of Economics 118, 269-298.

Chevillon, G., Massmann, M., and Mavroeidis, S., 2010. Inference in models with adaptive learning. Journal of Monetary Economics 57, 341-351.

Clarida, R., Gali, J., Gertler, M., 1999. The science of monetary policy: a new Keynesian perspective. Journal of Economic Literature 37, 1661–1707.

Dickey, D.A., Fuller W.A., 1979. Distribution of the estimators for autoregressive time series with a unit root. Journal of the American Statistical Association 74, 427-431.

Elliott, G., Rothenberg, T.J., Stock, J.H., 1996. Efficient tests for an autoregressive unit root. Econometrica 64, 813-836.

Elliott, G., 1998. On the robustness of cointegration methods when regressors almost have unit roots. Econometrica 66, 149-158.

Engle, R.F., Granger C.W.J., 1987. Co-integration and error correction: representation, estimation and testing. Econometrica 55, 251-276.

Evans, G.W., Honkapohja S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press: Princeton, New Jersey.

Evans, G.W., Honkapohja S., 2003. Adaptive learning and monetary policy design. Journal of Money, Credit, and Banking 35, 1045-1072.

Forsells, M., Kenny, G., 2002. The rationality of consumers' inflation expectations: survey-based evidence for the euro area. Working Paper Series 163, European Central Bank.

Gagnon, J.E., 2008. Inflation regimes and inflation expectations. Federal Reserve Bank of St. Louis Review 90, 229-243.

Gali, J., Gertler, M., 1999. Inflation dynamics: a structural econometric analysis. Journal of Monetary Economics 44, 195-222.

Grant, A.P., Thomas, L.B., 1999. Inflationary expectations and rationality revisited. Economics Letters 62, 331-338.

Hjalmarsson, E., Österholm, P., 2010. Testing for cointegration using the Johansen methodology when variables are near-integrated: size distortions and partial remedies. Empirical Economics 39, 51-76.

Kapetanios, G., Shin, G.Y., Snell, A., 2003. Testing for a unit root in the non-linear STAR framework. Journal of Econometrics 112, 359-379.

Kim S., Lima L.R., 2010. Local persistence and the PPP hypothesis. Journal of International Money and Finance 29, 555-569.

Lanne, M., 2001. Near unit root and the relationship between inflation and interest rates: a reexamination of the Fisher effect. Empirical Economics 26, 357-366.

Lima, L.R., Xiao Z., 2007. Do shocks last forever? Local persistency in economic time series. Journal of Macroeconomics 29, 103-122.
Mankiw, N.G., Reis, R., 2002. Sticky information versus sticky prices: a proposal to replace the New-Keynesian Phillips curve. Quarterly Journal of Economics 117, 1295-1328.
McCallum, B.T., 2010. Michael Woodford's contributions to monetary economics. In: Wieland, V. (Ed.), The Science and Practice of Monetary Policy Today. Springer-Verlag, Berlin Heidelberg, pp. 3-8.
Mishkin, F.S., 2010. Will monetary policy become more of a science? In: Wieland, V. (Ed.), The Science and Practice of Monetary Policy Today. Springer-Verlag, Berlin Heidelberg, pp. 81-103.
Ng, S., Perron, P., 2001. Lag length selection and the construction of unit root tests with good size and power. Econometrica 69, 1519-1554.
Orphanides, A., Williams, J.C., 2002. Robust monetary policy rules with unknown natural rates. Brookings Papers on Economic Activity 33, 63-118.
Orphanides, A., Williams, J.C., 2005. Imperfect knowledge, inflation expectations, and monetary policy. In: Bernanke B.S., and Woodford M. (Eds.), The Inflation Targeting Debate. University of Chicago Press, Chicago, pp. 201-234.
Phillips, P.C.B., Perron, P., 1988. Testing for a unit root in time series regression. Biometrika 75, 335-346.
Phillips, P.C.B., Moon H.R., Xiao, Z., 2001. How to estimate autoregressive roots near unity. Econometric Theory 17, 29–69.
Roberts, J.M., 1998. Inflation expectations and the transmission of monetary policy. Finance and Economics Discussion Series 98-43, Board of Governors of the Federal Reserve System.
Romer C., Romer D., 2000. Federal Reserve information and the behavior of interest rates. American Economic Review 90, 429-457.
Thomas, L.B., 1999. Survey measures of expected U.S. inflation. Journal of Economic Perspectives 13, 125-144.

Table I
Inflation and inflation expectations as AR(1) processes

	DF	DF-GLS	PP	PP-GLS	$\hat{\rho}_{OLS}$
π_{t+1}^{SPF}	-1.739	-1.696*	-1.637	-1.665*	0.985 (0.012)
π_{t+1}^{GB}	-2.140	-2.127**	-1.917	-2.625**	0.983 (0.015)
π_{t+1}	-2.279	-2.244**	-2.159	-2.617**	0.991 (0.010)

Note: π_{t+1}^{SPF} is the SPF inflation forecast, π_{t+1}^{GB} stands for the Greenbook inflation forecast and π_{t+1} is the inflation rate in terms of the GDP deflator. Figures in parentheses are standard deviations. DF stands for the Dickey-Fuller and PP for the Phillips-Perron test, while DF-GLS and PP-GLS denote the modified tests. Asterisks (*,**) indicate rejection of the null hypothesis of a unit root (for the 10% and 5% confidence interval respectively).

Table II
Consistent estimation of the autoregressive process parameters

	$\hat{\sigma}^2$	$\hat{\omega}^2$	$\hat{\lambda}$	$\hat{\rho}$	$\hat{d}$
π_{t+1}^{SPF}	5.099	12.974	3.938	0.795	0.312
π_{t+1}^{GB}	5.180	15.101	4.958	0.759	0.287
π_{t+1}	5.291	13.329	4.005	0.796	0.309

Note: As in Lima and Xiao (2007), $\hat{\sigma}^2$ and $\hat{\omega}^2$ are consistent estimators of the sample and the long-run variance of the series, $\hat{\lambda} = \frac{1}{2}(\hat{\omega}^2 - \hat{\sigma}^2)$, $\hat{\rho} = \hat{\rho}_{OLS} - n\hat{\lambda} \Big/ \sum y_{i-1}^2$ and $\hat{d} = -\frac{\ln(1-\hat{\rho})}{\ln(n)}$ is the parameter determining the local non-stationarity/stationarity properties of the series.

Table III
Local non-stationarity tests

Variable	Q test	t_{NL} test
π_{t+1}^{SPF}	3.124	-2.312
π_{t+1}^{GB}	2.509	-2.296
π_{t+1}	2.484	-1.989

Critical values:
Q test, 1%: 1.63, 5%: 1.36, 10%: 1.22 (see Lima and Xiao, 2007)
t_{NL} test, 1%: -3.48, 5%: -2.93, 10%: -2.66 (see case 2 in Kapetanios et al., 2003).

Table IV

Test of the rationality hypothesis

Variable (eq. 9)	*Q* test statistic
e_{t+1}^{SPF}	0.881
e_{t+1}^{GB}	1.147

Critical values: as in Table III

Table V
Estimates of the speed of adjustment parameter

Variable (eq. 10)	$\hat{b}_1$
$\Delta\pi_{t+1}^{SPF}$	-0.149 (0.062)
$\Delta\pi_{t+!}^{GB}$	-0.145 (0.071)

Note: Figures in parentheses are standard errors