Convergence of prices and rates of inflation

Fabio Busetti[*], Silvia Fabiani[†], Andrew Harvey[‡]

May 18, 2006

Abstract

We consider how unit root and stationarity tests can be used to study the convergence of prices and rates of inflation. We show how the joint use of these tests in levels and first differences allows the researcher to distinguish between series that are converging and series that have already converged, and we set out a strategy to establish whether convergence occurs in relative prices or just in rates of inflation. Special attention is paid to the issue of whether a mean should be extracted in carrying out tests in first differences and whether there is an advantage to adopting a (Dickey-Fuller) unit root test based on deviations from the last observation. The asymptotic distribution of this last test statistic is given and Monte Carlo simulation experiments show that the test yields considerable power gains for highly persistent autoregressive processes with "relatively large" initial conditions. The tests are applied to the monthly series of the Consumer Price Index in the Italian regional capitals over the period 1970-2003.

KEYWORDS: Dickey-Fuller test, Initial condition, Integration of order two, Law of one price, Stationarity test, Unit root test.
JEL classification: C22, C32

[*]Bank of Italy, Research Department. Email: fabio.busetti@bancaditalia.it
[†]Bank of Italy, Research Department. Email: silvia.fabiani@bancaditalia.it
[‡]Cambridge University, Faculty of Economics. Email: Andrew.Harvey@econ.cam.ac.uk.

1. Introduction

The issue of price and inflation convergence between different regions has attracted considerable interest in the recent years. Geographical barriers, local monopoly power and the presence of non-tradable goods are possible explanations of why prices may not converge within regions in the same country. Engel and Rogers (2001), Cecchetti et al. (2002), Chen and Devereux (2003) are recent empirical studies on price convergence among US regions.

We consider how unit root and stationarity tests can be used to study the convergence properties of price levels and inflation rates. We show how the joint use of these tests in levels and first differences allows the researcher to distinguish between series that are converging and series that have already converged, and we set out a strategy to establish whether convergence occurs in relative prices or just in rates of inflation. Special attention is paid to the issue of whether a mean should be extracted in carrying out tests in first differences and whether there is an advantage to adopting a (Dickey-Fuller) unit root test based on deviations from the last observation. The asymptotic distribution of this last test statistic is given and Monte Carlo simulation experiments show that the test yields considerable power gains for highly persistent autoregressive processes with "relatively large" initial conditions, the case of primary interest for analysing convergence. A modified version of the test that draws on the ideas in Elliott, Rothenberg and Stock (1996) is also investigated.

The tests are applied to the monthly series of the Consumer Price Index (CPI) in the Italian regional capitals over the period 1970-2003. As this index is not an absolute price level, we are investigating what might be labelled, following Engel and Rogers (2001), the "proportional law of one price".

2. Stability and convergence

2.1. Stability

If the difference between two nonstationary time series, y_t, is a stationary process with finite non-zero spectrum at the origin, we will say they have a stable relationship. The null hypothesis of stability may be tested by a stationarity test. Such a test will reject for large values of

$$\xi_1(m) = \frac{\sum_{t=1}^{T}\left(\sum_{j=1}^{t} e_j\right)^2}{T^2\widehat{\omega}^2(m)}, \qquad (2.1)$$

where $e_t = y_t - \overline{y}$ are the de-meaned observations and, following Kwiatkowski, Phillips, Schmidt and Shin (1992), hereafter KPSS, $\widehat{\omega}^2(m)$ is a non-parametric estimator of the long run variance of y_t, that is $\widehat{\omega}^2(m) = \hat{\gamma}(0)+2\sum_{\tau=1}^{m} w(\tau, m)\, \hat{\gamma}(\tau)$,with $w(\tau, m)$ being a weight function, such as the Bartlett window, and $\widehat{\gamma}(\tau)$ the sample autocovariance of y_t at lag τ.

If the mean is known to be zero under the null, then y_j rather than e_j is used to construct the test statistic, now denoted[1] by $\xi_0(m)$. Under the null hypothesis of zero-mean stationarity of y_t, the asymptotic distribution of $\xi_0(m)$ is given by the integral of a squared Brownian motion process, rather than a Brownian bridge. The $\xi_0(m)$ test will have power against a stationary process with a non-zero mean as well as against a non-stationary process. As shown in Busetti and Harvey (2006), another effective test can be based on the non-parametrically corrected 't-statistic' on the mean of y_t, that is $t(m) = \sqrt{T}\overline{y}/\widehat{\omega}(m)$. Under the null hypothesis of zero mean stationarity $t(m)$ converges to a standard Gaussian distribution. Busetti and Harvey (2006) show that this t-test is nearly as powerful as $\xi_0(m)$ against non-stationarity but is much more powerful against the alternative of a non-zero mean; they advise it be used when either alternative is of interest. Parametric versions of the tests are also possible.

2.2. Convergence

If y_t is stationary with finite non-zero spectrum at the origin, the series have already converged. However, they may be in the process of converging, have just converged or have converged some time ago but with a large part of the series dependent on initial conditions. A suitable model will be asymptotically stationary, satisfying the condition that $\lim_{\tau\to\infty} E(y_{t+\tau}|Y_t) = \alpha,$ where Y_t denotes current and past observations. Convergence is said to be *absolute* if $\alpha = 0$, otherwise it is *relative* (or conditional); see, for example, Durlauf and Quah (1999). The simplest such convergence model is an AR(1) process

$$y_t - \alpha = \phi\,(y_{t-1} - \alpha) + \eta_t, \quad t = 2, ..., T, \tag{2.2}$$

where η_t's are i.i.d. innovations and y_0 is a fixed initial condition. By rewriting (2.2) in error correction form as $\Delta y_t = \gamma + (\phi - 1)y_{t-1} + \eta_t,$ where $\gamma = \alpha(1 - \phi)$, it can be seen that the expected growth rate in the current period is a negative

[1] Unlike the case when the mean is subtracted, the statistic is different when reverse partial sums are used; see Busetti and Harvey (2006). This is not of any practical importance in the present context.

fraction of the gap between the two series after allowing for a permanent difference, α. We can therefore test against convergence, that is $H_0 : \phi = 1$ against $H_1 : \phi < 1$, by a unit root test. The power of the test will depend on the initial conditions, that is how far y_0 is from α. If α is known to be zero, the test based on the Dickey-Fuller (DF) $t-$statistic with no constant, denoted τ_0, is known to perform well, with a high value of $|y_0|$ actually enhancing power; see Müller and Elliott (2003).

What happens when testing for relative convergence? Including a constant in the DF regression and computing the $t-$statistic, denoted as τ_1, reduces power considerably. The test of Elliott, Rothenberg and Stock (1996), hereafter denoted ERS, also performs rather poorly as $|y_0 - \alpha|$ moves away from zero; again see Müller and Elliott (2003) and section 2.3 below. A possible way of enhancing power in this situation is to argue that we should set α equal to y_T and then run the simple DF test (without constant) on the observations $y_t - y_T$, $t = 1, ...T - 1$. We will denote this test statistic as τ^*. When $\phi = 1$, the asymptotic distribution of τ^* is

$$\tau^* \xrightarrow{d} \frac{-(W(1)^2 + 1)}{2\left[\int_0^1 W(r)^2 dr\right]^{1/2}} \tag{2.3}$$

where $W(r)$, is a standard Wiener process; for the proof of this result see the working paper version Busetti, Fabiani and Harvey (2006). The 10%, 5% and 1% lower tail quantiles of the limiting distribution are -2.43, -2.69 and -3.16, respectively. The power properties of the τ^* test are evaluated in the next subsection by Monte Carlo simulation experiments. It turns out that it is considerably more powerful than τ_1 for series that start far apart.

A possible objection to τ^* is that it introduces noise into the proceedings because of the variability in the last observation. This effect might be mitigated by estimating α by a weighted average of the most recent observations. Some rationale for this may be obtained by considering the theory for the ERS test. This involves the estimation of α by

$$\widehat{\alpha}_c = \left[y_1 + (1 - \overline{\phi}) \sum_{t=2}^{T} (y_t - \overline{\phi} y_{t-1})\right] / [1 + (T - 1)(1 - \overline{\phi})^2] \tag{2.4}$$

where $\overline{\phi} = 1 + \overline{c}/T$. The recommended value of $\overline{c}$ is 7, as in Elliott, Rothenberg and Stock (1996). If $\overline{c} = 0$ we end up subtracting the first observation. The asymptotic distribution for the t-statistic formed from $y_t - \widehat{\alpha}_c$ is the standard one for τ_0. The de-meaning is based on GLS estimation, assuming that $\alpha = y_0$. If instead we set $\alpha = y_{T+1}$, then we find

$$\widehat{\alpha}_c^* = \left[\overline{\phi}^2 y_T + (1-\overline{\phi}) \sum_{t=2}^{T} (y_t - \overline{\phi} y_{t-1})\right] / [\overline{\phi}^2 + (T-1)(1-\overline{\phi})^2] \qquad (2.5)$$

As in (2.4) the weights sum to unity. Denote the resulting test statistic as $\tau^*_{GLS,\overline{c}}$. As $\overline{\phi}$ approaches one, all the weight goes on to y_T and we obtain τ^*. More generally, a higher order autoregression is used, that is

$$\Delta y_t = \gamma + (\phi - 1)\, y_{t-1} + \gamma_1 \Delta y_{t-1} + ... + \gamma_{p-1} \Delta y_{t-p+1} + \eta_t, \qquad (2.6)$$

The Augmented Dickey-Fuller (ADF) test is based on such a regression. ERS recommend the use of (2.6), without the constant, having first subtracted $\widehat{\alpha}_c$ from y_{t-1}. An alternative would be to estimate α from (2.6) with ϕ set to $\overline{\phi}$. When $y_{T+1} = \alpha$ this leads to an estimator that places relatively more weight on the last p observations. Another possibility is to work within an unobserved components framework where the model is an AR(1) plus noise. In this case $\widehat{\alpha}_c^*$ is replaced by an estimator close to an exponentially weighted moving average (EWMA). The asymptotic distribution of all these modified ERS statistics under the null hypothesis is the same as τ^*.

The contrast between (log) price indices in Florence and Aosta shown in Figure 1 for seasonally unadjusted data provides an illustration. After seasonal adjustment, we use ADF-type regressions to compute the statistics τ_1 and τ^* with number of lags chosen according to the modified AIC criterion (MAIC) of Ng and Perron (2001). We obtain $\tau_1 = -2.53$ and $\tau^* = -2.95$, where α is estimated as the average of the last twelve months. Thus by including a constant term we are unable to reject the null hypothesis, even at 10% level of significance, while with τ^* we reject at 5% level. Notice that in this example the series start quite far apart: the ratio of the initial condition to the residual standard deviation is about 26 in a sample of 408 observations.

2.3. Monte Carlo evidence on the power of τ^* and related tests

Here we report Monte Carlo simulation experiments designed to compare the power of τ^* and the GLS de-meaning test, $\tau^*_{GLS,\overline{c}}$, with the power of the standard DF t-test, τ_1. We consider a near-unit root data generating processes with a range of initial conditions and set $\overline{c} = 10$ for $\tau^*_{GLS,\overline{c}}$; recall that $\overline{c} = 0$ corresponds to τ^*

and that for both τ^* and $\tau^*_{GLS,10}$ the limiting distribution is given by (2.3) with critical values as in table 1. We also examine the $Q^\mu(\overline{\phi}, \infty)$ test of Müller and Elliott (2003); this test belongs to a class of point optimal invariant tests and is designed to give high power for large initial conditions.[2] The test statistic is computed as described in Elliott and Müller (2006, eq. 15), letting $k \to \infty$. Finally, for completeness, we report results for the $DF - GLS$ test of ERS, obtained by GLS de-meaning under the assumption that $y_0 = 0$.

The Monte Carlo experiment considers the AR(1) data generating process, $t = 1, 2, ..., T$,

$$y_t = \alpha + u_t, \quad u_t = (1 - c/T)u_{t-1} + \eta_t, \quad \eta_t \sim NID(0, 1)$$

with c taking on the values $0, 1, 2.5, 5, 10$ and $u_0 = \alpha + K$, with K varying among 0, 5, 10, 15, 20, 25, 30 and 50. The notation $NID(a, b)$ indicates a Gaussian independent and identically distributed process with mean a and variance b. Thus y_t is a highly persistent process for $c > 0$ and a unit root process for $c = 0$. K is the magnitude of the initial condition in units of the disturbance standard deviation. Note that τ_1, τ^* and $Q^\mu(\overline{\phi}, \infty)$ are invariant to α and so this is set equal to zero in the simulations.

Table 1 contains the simulated rejection frequencies of these tests for $T = 100$ and a 5% significance level. This magnitude of the sample size might be most relevant for quarterly data. In this case $c = 5$ is quite plausible as it corresponds to $\phi = 0.95$; a smaller ϕ would mean unusually fast convergence. A value above 0.975 ($c = 2.5$) is quite slow. Table 1 shows that, for $c = 2.5$ and 5, τ^* is considerably more powerful than the standard DF test τ_1 when the initial condition is relatively large. In fact τ_1 is only better when K is 5 or zero and then the power is so low as to render both tests useless. For $c \geq 5$, the use of $\tau^*_{GLS,10}$ would allow further gains over τ^*, however those gains do not seem as large as the losses incurred for $c = 2.5$ and 1. Similarly the $Q^\mu(1 - 10/T, \infty)$ test is generally dominated by τ^* when $c \leq 5$ and $K > 10$. Note that our simulation results are consistent with those reported in table 3 of Elliott and Müller (2005): for instance, their case of $\rho = 0.95$ and $\alpha = 3$ corresponds in our framework to $T = 100$, $c = 5$, $K \approx 10$: their size corrected power of 0.29 is analogous to our reported rejection frequency 0.28. Finally, the power of the DF-GLS test of ERS is much higher than that of the other tests for $K = 0$, but, rather than increasing, it approaches zero rather quickly

[2] The ERS test of Elliott et al. (1996) and the test of Elliott (1999) belong to the family of $Q^\mu(\overline{\phi}, k)$ tests for, respectively, $k = 0$ and $k = 1$. We are grateful to a referee for suggesting that we look at the $Q^\mu(\overline{\phi}, k)$ test.

as the initial condition becomes larger. On balance, the simple test τ^* seems to display the most desirable power properties unless the initial conditions are close to zero. Additional simulation results, for the case $T = 400$ and perhaps more relevant for the case of monthly data, are contained in the working paper version Busetti, Fabiani, Harvey (2006). Note that in the local-to-unity framework (with the autoregressive parameter depending on the sample size and the initial condition fixed), the power of the tests for non-zero initial conditions are lower the larger is the sample. On the other hand, if the autoregressive parameter is kept fixed (e.g. $c = 2.5$ with $T = 100$ versus $c = 10$ with $T = 400$) the power increases with the sample size for given initial condition.

Multivariate tests of stability and convergence can be constructed by applying (multivariate) stationarity and unit root tests to the vector of contrasts between each unit and a benchmark. Further details are contained in the working paper version Busetti, Fabiani, Harvey (2006).

3. Testing stability and convergence in levels and first differences

For data on prices it is of interest to test the hypotheses of stability and convergence in both levels and first differences, that is to analyze the dynamics of both relative prices and inflation differentials. Let $P_{i,t}$ denote some weighted average of prices in region i at time t. If information is available only for a price index, the observations are $p_{i,t} = P_{i,t}/P_{i,b},\ i = 1, ..., n,\ t = 1, ..., T$, where $b \in \{1, ..., T\}$ is the base year. The difference - or contrast - between (the log of) this price index and one in another region, say region j, denoted $y_t^{i,j}$, is

$$y_t^{i,j} = \log p_{i,t} - \log p_{j,t}, \quad t = 1, ..., T, \qquad i, j = 1, 2, ..., n \tag{3.1}$$

where $y_b^{i,j} = 0$ by definition. This is the logarithm of the relative price between the two regions. The base can always be changed to a different point in time, τ, by subtracting y_τ from all the observations. It is not possible to discriminate between absolute and relative convergence with price indices; all that can be investigated is convergence to the proportional law of one price. The appropriate test for stability is $\xi_1(m)$. Not subtracting the mean gives a test statistic, $\xi_0(m)$, that is not invariant to the base and does not give the usual asymptotic distribution under the null hypothesis of a zero mean stationary process since treating the $y_t's$ as independent is incorrect. A test of convergence, on the other hand, can be based on a DF statistic,τ^*, formed by taking the base to be the last period.

The contrasts in the rate of inflation, that is the inflation differentials,

$$\Delta y_t^{i,j} = \Delta \log p_{i,t} - \Delta \log p_{j,t}, \quad t = 1, ..., T \tag{3.2}$$

are invariant to the base year since this cancels out yielding $\Delta y_t^{i,j} = \Delta \log P_{i,t} - \Delta \log P_{j,t}$. A test of the null hypothesis that there are no permanent, or persistent, influences on an inflation rate contrast amounts to testing that Δy_t is stationary with a mean of zero. The appropriate tests are therefore $\xi_0(m)$ and $t(m)$. Similarly the null hypothesis of no convergence in an inflation rate contrast against the alternative of absolute convergence can be tested using τ_0, the t-statistic obtained from an ADF regression without a constant.[3]

3.1. A testing strategy

Taking account of the results of unit roots and stationarity tests allows the researcher to distinguish between regions that *have already converged* (characterized by rejection of unit root and non-rejection of stationarity test) and regions that *are in the process of converging* (rejection by both tests[4]). However, since both levels and first differences are of interest, the order of testing is also important: do we start the testing procedures with levels or first differences?

As regards convergence tests, Dickey and Pantula (1987), argue that it is best to test for a unit root in first differences and if this is rejected, to move on to test for a unit root in the levels.[5] On the other hand, stationarity of the levels implies that the spectrum of first differences is zero at the origin, thereby invalidating a (nonparametric) stationarity test on first differences. This suggests that the sequence of stability tests should be one in which the stationarity of Δy_t is tested only if stationarity of y_t has been rejected; see also Choi and Yu (1997).

[3] Note that if the time series contrast can be described by a stationary process around a nonzero mean, the τ_0 test will tend not to reject while $\xi_0(m)$ and $t(m)$ will tend to reject; see Busetti and Harvey (2006). Thus if the series are drifting apart because of a non-zero deterministic difference in growth rates, the contrast will tend to be identified as one where there is no convergence, which, of course, is the right outcome. If the τ^* test and $\xi_1(m)$ tests are applied in levels, the latter will tend to reject while the former will not; it is not appropriate to make allowance for a time trend.

[4] As shown in Muller (2005), a stationarity test will tend to reject the null hypothesis for highly persistent time series. In other words, it is difficult to control the size of stationarity tests in the presence of strong autocorrelation; see also KPSS.

[5] The results in Pantula (1989) indicate that the test of a unit root in inflation will tend to reject if the price level is stationary.

Taking those arguments into account we end up with the strategy described in the chart in figure 2, with five possible outcomes A,B,C,D,E. The starting point is the unit root test on inflation differentials. If this doesn't reject we have the case of non-convergence (E), while a rejection will lead to testing the unit root hypothesis in relative prices. The result of the latter test will lead to a stationarity test in either levels or first differences. The final outcomes are as follows.

(A) Relative prices are converging: rejection of unit root in first differences and levels, rejection of levels stationarity test.

(B) Relative prices have converged: rejection of unit root in first differences and levels, non rejection of levels stationarity test.

(C) Inflation rates are converging: rejection of unit root in first differences but not in levels, rejection of first differences stationarity test.

(D) Inflation rates have converged: rejection of unit root in first differences but not in levels, non rejection of first differences stationarity test.

(E) Non convergence: non rejection of unit root in first differences.

The price and inflation contrasts between Florence and Aosta provide again an illustration. The null hypothesis of non convergence is rejected at the 1% level by the ADF test on the inflation differential: the modified AIC lag selection criterion of Ng and Perron (2001) suggests 19 lags and resulting τ_0 statistic is -3.21. The unit root in levels is also rejected, as was seen in sub-section 2.2, and a rejection also occurs for the levels stationarity tests. Thus, the sequential testing procedure leads to the conclusion that relative prices are converging, that is case A. Further details are provided in table 4 of Busetti, Fabiani, Harvey (2006). In particular, it is interesting to notice that the $\xi_0(m)$ stationarity test applied to the inflation differential would also reject the null hypothesis: the next sub-section explains why this happens.

3.2. First differences stationarity tests for highly persistent process in levels

The properties of first differences stationarity tests when the DGP is a highly persistent process in levels depend on whether the initial condition is small or large. In the former case the test is undersized, in the latter it is oversized with the degree of oversizing increasing with the magnitude of the initial condition. We present a small Monte-Carlo simulation experiment that illustrates the point.

We consider the AR(1) data generating process, $t = 1, 2, ..., T$,

$$y_t = (1 - c/T)y_{t-1} + \eta_t, \qquad \eta_t \sim NID(0, 1) \tag{3.3}$$

for some given initial condition y_0. Thus, as in section 2.3, y_t is a highly persistent process for $c > 0$ and a unit root process for $c = 0$. Notice that a relatively small c and a large initial condition are associated with y_t converging to its long run value of zero.

The validity of stationarity tests in first differences requires that $c = 0$ in (3.3). If this is not the case then the properties of the test depend on the magnitude of the initial condition y_0 relatively to the standard deviation of η_t. In particular, the test is undersized if y_0 is small and (often dramatically) oversized if y_0 is large. We take $\sigma^2_\eta = 1$, $c = 0, 1, 2.5, 5, 10$ and $y_0 = 0, 5, 10, 15, 20, 25, 30, 50$. Table 2 reports rejection frequencies for the stationarity tests $\xi_0(m)$, $\xi_1(m)$ computed on the first differenced data Δy_t, for $T = 100$, where the bandwidth parameter for spectral estimation is equal to $int(m(T/100)^{0.25})$ and $m = 0, 4, 8$.

For $c = 0$ the stationarity tests in first differenced have (approximately) the correct size, while they are undersized when $c > 0$ and the initial condition is small. Oversizing occurs for a large initial condition, at least as large as 15 when $T = 100$. Notice that oversizing can be huge, with the probability of rejecting the null equal or close to 1 in many cases.

Intuitively, this oversizing problem can be explained if we think of a converging path in levels: the first difference is the slope of the series which keeps changing mostly in the same direction in order to bring the level to its long run value. Large initial conditions are not unusual for converging series, as can be seen in the Florence-Aosta example.

4. Convergence properties of the CPI among Italian regions

We apply our testing strategy to inflation and price differentials among Italian regions. The data used are the monthly Istat series of the Consumers' Price Index in nineteen regional capitals for the period 1970M1-2003M12. The series have been rebased, taking 2003 as the base year, and they have been seasonally adjusted by removing a stochastic seasonal component using the STAMP package of Koopman *et al.* (2000). Figure 3 shows the time pattern of the log of relative price levels, computed as the difference between each (log) regional price index and the average national one. As we have set 2003 as the base year the contrasts are constrained by construction to tend to zero near the end of the sample period. The picture seems consistent with high persistence in price differentials, either a unit root or a converging process.

The summary results of the battery of convergence and stability tests on inflation and price differentials on the 171 regional contrasts are reported in table 3. The first panel of the table reports the number of rejections (at 1%, 5% and 10% significance levels) for the ADF test τ_0 on the pairwise inflation contrasts (computed without fitting a mean). The number of lags in the ADF regression, not shown in the table, is chosen according to the modified Akaike information criterion of Ng and Perron (2001). For all inflation differentials the test easily rejects the null hypothesis, thus excluding case E of non-convergence. The results of the unit root tests on price contrasts are reported in the second panel of the table, both for the ADF test with a constant term, τ_1, and for the modified ADF test, τ^* (where the data are transformed by subtracting the average of the observations in the final year). The contrasts are split into two groups, depending on the size K of the initial condition (smaller or larger, in absolute value, than 10 times the residuals standard deviation). This is the second step of the testing strategy summarised in figure 2. According to the results of the new test τ^*, we have 41 rejections (at least at the 10% level of significance) of a unit root in relative prices; notice that, as predicted by the simulation results of table 1, τ^* rejects the null hypothesis much more frequently than τ_1 does for cases where the initial condition is at least 10 times larger than the residuals standard deviation. The third panel of table 3 contains the number of rejections of the KPSS stationarity test, $\xi_1(m)$, carried out on those 41 contrasts identified in the previous step: as the hypothesis of stationarity turns out to be always rejected (at least at the 10% level of significance), we classify all these contrasts as cases A, i.e. pairs of cities where relative prices are in the process of converging. Finally, the last panel of the table presents the results of the stationarity test $\xi_0(m)$ computed on the 130 inflation differentials among pairs of cities for which τ^* could not reject the hypothesis of a unit root in the price contrasts. Here the stationarity test rejects the null hypothesis in 41 cases, which are labelled as cases C of converging inflation rates. The remaining 89 pairs of cities are then attributed to group D: inflation rates have already converged.

In summary, out of 171 regional contrasts we obtained 89 cases D of stability (around zero) of inflation differentials, 41 C's of converging inflation rates, and 41 A's of converging relative prices. Among the largest cities, it turns out that inflation rates have been stable between Milan, Naples and Turin, while relative prices are converging between Rome and Milan and Rome and Naples. Detailed results are provided in the working paper version Busetti, Fabiani, Harvey (2006).

5. Concluding remarks

In examining the behaviour of relative price time series between different regions it is important to distinguish between stability and convergence. Stability is assessed by stationarity tests, while convergence is determined by unit root tests. For pairwise contrasts of inflation rates, these tests are best carried out without removing a constant term. As an alternative to the stationarity test, a 't-test' on the sample mean may be used. For price level contrasts, a Dickey-Fuller unit root test run on data with the base year at the end, and no constant removed, displays good power in testing for relative convergence. We derive the asymptotic distribution of this test statistic, provide critical values and compare its performance with other tests.

We set out a sequential testing strategy to establish whether convergence occurs in relative prices or just in rates of inflation. This strategy is applied to the monthly series of the Consumer Price Index in the Italian regional capitals over the period 1970-2003. It is found that all 171 pairwise contrasts of inflation rates have converged or are in the process of converging. However only about one quarter of price level contrasts appear to be converging.

Acknowledgements

We would like to thank Roberto Sabbatini, Rob Taylor and participants in the *Frontiers in Time Series Analysis* conference, Olbia 2005, and the Bank of Italy's *Workshop on Inflation Convergence in Italy and the Euro Area*, Roma 2004, for comments on earlier versions of the paper and Angela Gattulli for her valuable support with the data. We are also grateful to the editor, Anindya Banerjee, and to two referees for interesting remarks and suggestions. The views expressed here are those of the authors, not the Bank of Italy. Andrew Harvey gratefully acknowledges the hospitality and financial support of the Research Department of the Bank of Italy.

References

[1] Busetti, F., Fabiani, S. and Harvey, A.C. (2006). Convergence of prices and rates of inflation, *Bank of Italy Discussion Papers*, 575, downloadable from http://www.bancaditalia.it/ricerca/consultazioni/temidi

[2] Busetti, F. and Harvey, A.C. (2006). Testing for trend, mimeo.

[3] Cecchetti, S.G., Nelson, C.M. and Sonora, R. (2002). Price Index Convergence among United States Cities, *International Economic Review*, 43(4), 1081-99.

[4] Chen, L.L and Devereux, J. (2003), What can US city price data tell us about purchasing power parity?, *Journal of International Money and Finance*, 22(2), 213-22.

[5] Choi, I. and Yu, B.C. (1997). A General Framework for Testing I(m) against I(m+k), *Journal of Economic Theory and Econometrics*, 3, 103-38.

[6] Dickey, D.A. and S.G. Pantula (1987). Determining the Ordering of Differencing in Autoregressive Processes, *Journal of Business and Economic Statistics*, 5(4), 455–61.

[7] Durlauf, S. and Quah, D. (1999). The new empirics of economic growth. In J.B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1, Ch. 4, 235-308, Amsterdam: Elsevier Science.

[8] Elliott, G. and Müller, U.K. (2006). Minimizing the impact of initial condition on testing for unit roots, *Journal of Econometrics*, forthcoming.

[9] Elliott, G., Rothenberg, T.J. and Stock, J.H. (1996). Efficient Tests for an Autoregressive Unit Root, *Econometrica*, 64, 813-36.

[10] Engel, C. and Rogers, J.H. (2001). Violating the Law of One Price: Should We Make a Federal Case Out of It?, *Journal of Money, Credit and Banking*, 33(1), 1-15.

[11] Koopman, S.J., Harvey, A.C., Doornik, J.A. and Shephard, N. (2000). *STAMP: Structural Time Series Analyser Modeller and Predictor*, Timberlake Consultants.

[12] Kwiatkowski, D., Phillips, P.C.B., Schmidt, P. and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: how sure are we that economic time series have a unit root?, *Journal of Econometrics,* 44, 159-78.

[13] Müller, U.K. (2005). Size and power of tests of stationarity in highly autocorrelated time series, *Journal of Econometrics*, 128, 195–213.

[14] Müller, U.K. and Elliott, G. (2003). Tests for Unit Roots and the Initial Condition, *Econometrica*, 71, 1269-86.

[15] Ng, S. and Perron, P. (2001). Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power, *Econometrica*, 69, 1519-54.

[16] Pantula, S.A. (1989). Testing for Unit Roots in Time Series Data, *Econometric Theory*, 5, 256–71.

Table 1. Power comparison of convergence tests - T=100

		Initial Condition							
		0	5	10	15	20	25	30	50
c=10	τ_1	0.28	0.37	0.64	0.92	1.00	1.00	1.00	1.00
	τ^*	0.14	0.22	0.54	0.89	0.98	1.00	1.00	1.00
	$\tau^*_{GLS,10}$	0.24	0.35	0.70	0.95	1.00	1.00	1.00	1.00
	$Q^m(1\text{-}10/T,8)$	0.12	0.37	0.78	0.96	0.99	1.00	1.00	1.00
	$ERS_{GLS,7}$	0.79	0.09	0.00	0.00	0.00	0.00	0.00	0.00
c=5	τ_1	0.10	0.12	0.19	0.35	0.59	0.82	0.95	1.00
	τ^*	0.05	0.07	0.19	0.53	0.89	0.99	1.00	1.00
	$\tau^*_{GLS,10}$	0.09	0.12	0.26	0.56	0.88	0.99	1.00	1.00
	$Q^m(1\text{-}10/T,8)$	0.05	0.12	0.28	0.40	0.48	0.54	0.59	0.72
	$ERS_{GLS,7}$	0.41	0.08	0.00	0.00	0.00	0.00	0.00	0.00
c=2.5	τ_1	0.06	0.06	0.07	0.09	0.13	0.19	0.27	0.74
	τ^*	0.03	0.04	0.07	0.18	0.43	0.75	0.94	1.00
	$\tau^*_{GLS,10}$	0.06	0.07	0.10	0.19	0.36	0.60	0.82	1.00
	$Q^m(1\text{-}10/T,8)$	0.04	0.06	0.10	0.12	0.12	0.10	0.08	0.02
	$ERS_{GLS,7}$	0.21	0.11	0.01	0.00	0.00	0.00	0.00	0.00
c=1	τ_1	0.05	0.05	0.05	0.05	0.04	0.04	0.04	0.06
	τ^*	0.03	0.03	0.04	0.07	0.11	0.17	0.27	0.84
	$\tau^*_{GLS,10}$	0.05	0.05	0.06	0.07	0.09	0.12	0.17	0.51
	$Q^m(1\text{-}10/T,8)$	0.04	0.04	0.05	0.05	0.05	0.05	0.04	0.01
	$ERS_{GLS,7}$	0.12	0.11	0.07	0.03	0.01	0.00	0.00	0.00
c=0	τ_1	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
	τ^*	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
	$\tau^*_{GLS,10}$	0.06	0.06	0.06	0.05	0.05	0.04	0.04	0.02
	$Q^m(1\text{-}10/T,8)$	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
	$ERS_{GLS,7}$	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08

Table 2. Rejection frequencies of first differences stationarity tests for a highly persistent process in the levels - T=100

			Initial Condition							
		m	0	5	10	15	20	25	30	50
c=10										
	$\xi_0(m)$	0	0.00	0.00	0.00	0.30	0.97	1.00	1.00	1.00
		4	0.00	0.00	0.01	0.26	0.83	0.99	1.00	1.00
		8	0.00	0.00	0.02	0.24	0.64	0.91	0.98	1.00
	$\xi_1(m)$	0	0.00	0.00	0.03	0.28	0.80	0.99	1.00	1.00
		4	0.00	0.00	0.03	0.23	0.62	0.90	0.99	1.00
		8	0.00	0.00	0.04	0.18	0.43	0.70	0.87	1.00
c=5										
	$\xi_0(m)$	0	0.00	0.00	0.01	0.36	0.93	1.00	1.00	1.00
		4	0.00	0.00	0.03	0.37	0.88	0.99	1.00	1.00
		8	0.00	0.00	0.05	0.39	0.82	0.98	1.00	1.00
	$\xi_1(m)$	0	0.00	0.01	0.05	0.20	0.46	0.76	0.93	1.00
		4	0.00	0.01	0.05	0.18	0.39	0.67	0.87	1.00
		8	0.00	0.01	0.05	0.15	0.33	0.56	0.77	1.00
c=2.5										
	$\xi_0(m)$	0	0.00	0.00	0.03	0.22	0.61	0.91	0.99	1.00
		4	0.00	0.00	0.04	0.25	0.63	0.90	0.99	1.00
		8	0.00	0.01	0.06	0.29	0.64	0.90	0.98	1.00
	$\xi_1(m)$	0	0.02	0.02	0.04	0.08	0.15	0.25	0.36	0.84
		4	0.01	0.02	0.04	0.07	0.13	0.22	0.31	0.79
		8	0.01	0.02	0.03	0.07	0.12	0.18	0.27	0.73
c=1										
	$\xi_0(m)$	0	0.01	0.01	0.03	0.09	0.18	0.33	0.49	0.96
		4	0.01	0.02	0.04	0.10	0.21	0.35	0.52	0.96
		8	0.01	0.02	0.05	0.12	0.24	0.38	0.55	0.96
	$\xi_1(m)$	0	0.04	0.04	0.04	0.04	0.05	0.05	0.06	0.11
		4	0.03	0.03	0.03	0.04	0.04	0.05	0.05	0.10
		8	0.03	0.03	0.03	0.03	0.04	0.04	0.05	0.08
c=0										
	$\xi_0(m)$	0	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
		4	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07
		8	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08
	$\xi_1(m)$	0	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
		4	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04
		8	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03

$\xi_0(m)$ is a stationarity test without constant, with bandwidth equal to int(m(T/100)^.25).
$\xi_1(m)$ is a stationarity test with constant, with bandwidth equal to int(m(T/100)^.25).
The initial condition is in units of the error standard deviation.

Table 3. Summary Results for the tests on the CPI in the Italian Regional Capitals

		Number of contrasts	*Number of rejections* 1%	*5%*	*10%*
Inflation contrasts: unit root test					
τ_0		171	167	171	171
Price contrasts: unit root tests					
$\|K\|<10$	τ_1	43	-	1	6
	τ^*	43	-	1	3
$\|K\|\geq 10$	τ_1	128	6	11	21
	τ^*	128	11	26	38
Price contrasts: stationarity test					
$\xi_1(m)$		41	37	40	41
Inflation contrasts: stationarity test					
$\xi_0(m)$		130	8	26	41

Figure 1 – Relative prices and inflation rates in Florence and Aosta

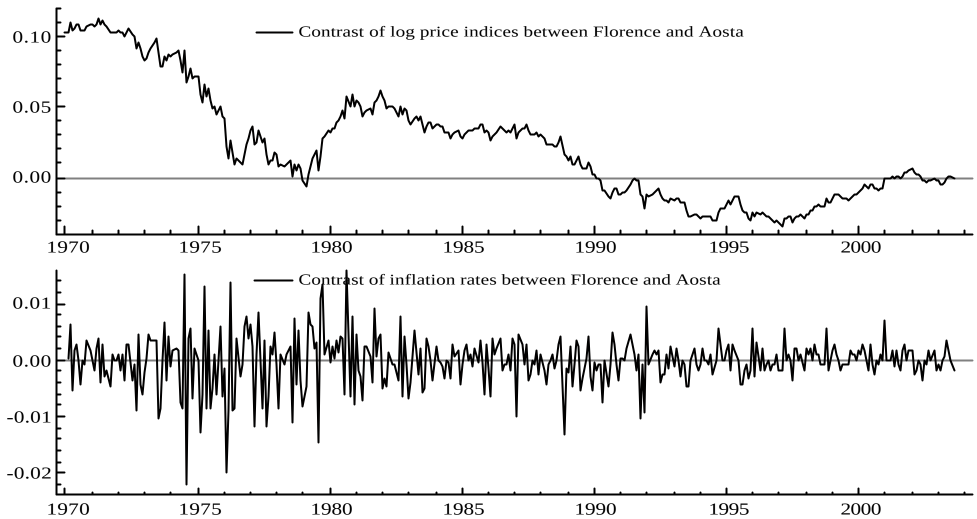

Figure 2 – Testing convergence in levels and first differences

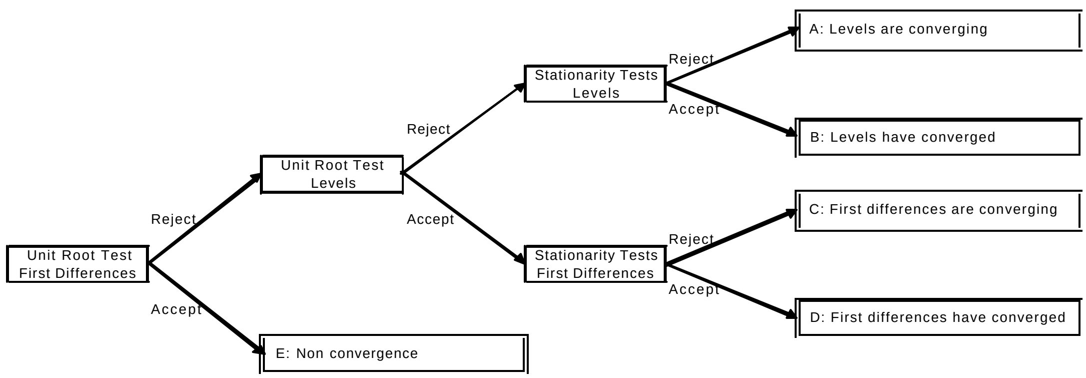

Figure 3 – Regional relative prices, base year=2003

(computed as differences with respect to the Italian average cost of living index)