



Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Can aggregation explain the persistence of inflation? ☆

Filippo Altissimo^a, Benoit Mojon^b, Paolo Zaffaroni^{c,*}^a Brevan Howard Asset Management LLP and CEPR, UK^b Banque de France, France^c Imperial College London, UK

ARTICLE INFO

Article history:

Received 5 December 2006

Received in revised form

18 December 2008

Accepted 18 December 2008

Available online 20 January 2009

Keywords:

Inflation

Persistence

Aggregation

Euro area

Price stickiness

ABSTRACT

An aggregation exercise is proposed that aims at investigating whether the fast average adjustments of the disaggregate inflation series of the euro area CPI is coherent with the slow adjustment of euro area aggregate inflation. Estimating a dynamic factor model for 404 inflation sub-indices of the euro area CPI allows to decompose the dynamics of inflation sub-indices into a part due to a common macroeconomic shock and to sector specific idiosyncratic shocks. Although idiosyncratic shocks dominate the variance of sectoral prices, one common factor appears to be the main driver of aggregate dynamics. In addition, the heterogenous propagation of this common shock across sectoral inflation rates, and in particular its slow propagation to inflation rates of services, generates the persistence of aggregate inflation. We conclude that the aggregation mechanism explains a significant amount of aggregate inflation persistence.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Recent contributions show that heterogeneity in price stickiness across sectors implies a slower response of aggregate prices to a monetary policy shock than would be the case if prices of all sectors were identically sticky (see [Imbs et al., 2005](#); [Carvalho, 2006](#); [Boivin et al., 2006](#); [Clark, 2006](#) and references therein). In view of this result, it appears legitimate to analyze the role of the aggregation process across heterogenous sectors for the dynamics of the business cycle.

We make a step in this direction by assessing the role of aggregation across heterogenous sectors for the persistence of aggregate inflation. We empirically investigate whether the prediction that persistence of the aggregate is (much) larger than the cross-sectoral average persistence holds true for our data, which gathers inflation rates from 404 sub-indices of the euro area CPI between 1985 and 2005. It turns out that this prediction is not rejected for our data. In particular, we show that the heterogenous responses of sectoral inflation rates to a common shock imply an aggregate inflation response that is indeed much slower than if we impose that all sectors respond similarly to the median/average sector. We therefore argue that the aggregation process can solve the apparent dilemma between the flexibility of sectoral prices¹ and the persistence of macroeconomic inflation.

☆ We are grateful to Laurent Baudry, Laurent Bilke, Herve Le Bihan and Sylvie Tarrieu (Banque de France), Johannes Hoffman (Bundesbank) Roberto Sabbatini and Giovanni Veronese (Banca d'Italia), Anna-Maria Agresti and Martin Eglersperger (ECB) for making the data available to us. We thank Gonzalo Camba-Méndez, Steven Cecchetti, M. Hashem Pesaran, Jim Stock and the participants to the Eurosystem Inflation Persistence Network. Finally we thank the Editor (Bob King) and two anonymous referees for their valuable comments. Any remaining errors are of course the sole responsibility of the authors. The views expressed in the paper are the authors' own and do not necessarily reflect those of the Banque de France, the ESCB or Brevan Howard Asset Management.

* Corresponding author.

E-mail address: p.zaffaroni@imperial.ac.uk (P. Zaffaroni).

¹ See [Chari et al. \(2000\)](#), [Bils and Klenow \(2004\)](#) and [Dhyne et al. \(2005\)](#).

Table 1
Impulse response function μ_k for the Beta(a, b) density with $a = 4b$.

k	μ_k				0.8^k
	$b = 0.2$	0.7	1	3	
1	0.8	0.8	0.8	0.8	0.8
2	0.72	0.67	0.66	0.65	0.64
5	0.61	0.48	0.44	0.37	0.33
10	0.54	0.33	0.28	0.17	0.11
50	0.39	0.12	0.07	0.01	1.4×10^{-5}
200	0.36	0.09	0.05	0.01	4.1×10^{-20}

Our analysis is based on results on aggregation of time series² and it is built upon an explicit model of the aggregation process that is then estimated and simulated. In the model, each sectoral inflation rate is decomposed into its response to an aggregate shock that affects all sectors and a sector specific shock. Our main findings can be summarized as follows: (i) one common factor accounts for 30 percent of the overall variance of the 404 series. This share is twice as large if one focuses on business cycle and lower frequencies, i.e. on the persistent components of the series; (ii) the propagation mechanism of the common shock is highly heterogeneous across sectors; (iii) the persistence implied by the aggregation exercise mimics remarkably well the persistence observed in the aggregate inflation. In particular, the cross-sectional distribution of the sectoral parameters implies an autocorrelation function of the aggregate CPI inflation which decays hyperbolically toward zero and displays long memory; and (iv) the high volatility and low persistence, observed on average at the level of sectoral inflation, is consistent with aggregate smoothness and high persistence.

Such results are important for two distinct reasons. First, they show the importance of heterogeneity and aggregation in shaping up the dynamic response of macrovariables. This is a clear warning against the naïve use of average-median microparameters in calibrating macromodels. Second, they shed further light on the long standing debate on the persistence of inflation (see Pivetta and Reis, 2007 and reference therein). These results provide support to the view of inflation as a stationary process, albeit with long memory.

The paper is organized as follows. Section 2 illustrates the intuition on the theoretical link between aggregation, heterogeneity and persistence. Section 3 presents the data used in the empirical analysis and addresses the presence of common factors across the price sub-indices. Section 4 introduces the micromodel and the estimation methodology. It then discusses the estimation results and their implications for the aggregate persistence. Section 5 concludes.

2. Aggregation, heterogeneity and persistence: a link

To build the intuition on the theoretical link between heterogeneity, aggregation and persistence, we make use of a simple example which summarizes the results of Zaffaroni (2004) (see also Robinson, 1978; Granger, 1980). Let us consider n units and let the i th unit be described by a first-order autoregressive model

$$y_{it} = \alpha_i y_{it-1} + u_t + \varepsilon_{it}, \quad i = 1, \dots, n, \quad (1)$$

where u_t , the common shock, is an *i.i.d.* sequence $(0, \sigma_u^2)$, and the ε_{it} , the idiosyncratic shock, an *i.i.d.* sequence $(0, \sigma_{\varepsilon,i}^2)$ for any i, t . The propagation of the shocks is heterogeneous across units and the α_i are *i.i.d.* with probability density function $f(\alpha)$ over the support $(-1, 1)$.

The aggregate is simply the average of the individual units and by linearity $Y_{n,t} = n^{-1} \sum_{i=1}^n y_{it} = U_{n,t} + E_{n,t}$, with $E_{n,t} = n^{-1} \sum_{i=1}^n \varepsilon_{it} / (1 - \alpha_i L)$ and $U_{n,t} = u_t + \hat{\mu}_1 u_{t-1} + \hat{\mu}_2 u_{t-2} + \dots + \dots$, setting $\hat{\mu}_k = n^{-1} \sum_{i=1}^n \alpha_i^k$. When n gets large, by the strong law of large numbers, each $\hat{\mu}_k$ will converge *a.s.* to the population moments of the α_i , i.e. $\mu_k = \int_{-1}^1 \alpha^k f(\alpha) d\alpha$. The dynamic pattern of the μ_k represents the *impulse response* of the common shocks u_t on the aggregate. Zaffaroni (2004) shows that the rate of decay of the μ_k (i.e. the persistence of the aggregate shock) only depends on the behavior near unity of $f(\alpha)$.

To illustrate this result, let us parameterize $f(\alpha)$ as a Beta(a, b) distribution, with $a, b > 0$.³ The parameter b governs the mass of the Beta distribution around unity and a is set equal to $(\mu / (1 - \mu))b$ in order to assure that the distribution has mean μ for any choice of b . Table 1 reports the behavior of μ_k for various choices of b when $\mu = 0.8$ across all columns.

The smaller is b , the larger is the mass of the distribution $f(\alpha)$ around unity (i.e. the larger is the share of microunits having very high persistence), and therefore the slower would the impulse response converge to zero. In contrast, the last column shows that if all units had the same AR coefficient (equal to $\mu = 0.8$) the impulse response would decay to zero very fast, as μ^k . Zaffaroni (2004) shows that $\mu_k \sim ck^{-b}$ as $k \rightarrow \infty$ for some $c > 0$. This characterization of the impulse response implies a precise representation of the autocovariance function (acf) and spectral density of the limit aggregate U_t , the latter being the limit (in mean-square) of $U_{n,t}$. In particular, Zaffaroni (2004) shows that when $\frac{1}{2} < b < 1$, then

² See Robinson (1978), Granger (1980) and Zaffaroni (2004).

³ A Beta distribution has density function $f(x|a, b) = B^{-1}(a, b)x^{a-1}(1-x)^{b-1}$ for $0 \leq x \leq 1$ where $B(\cdot, \cdot)$ is the Beta function.

$var(U_t) < \infty$, $cov(U_t, U_{t+k}) \sim ck^{1-2b}$ as $k \rightarrow \infty$. In other words, when $\frac{1}{2} < b < 1$ the $cov(U_t, U_{t+k})$ decays toward zero, in agreement with the notion of stationarity, but too slowly to ensure summability, resembling the classical definition⁴ of long memory for the U_t given by

$$cov(U_t, U_{t+k}) \sim ck^{2d-1}, \quad k \rightarrow \infty, \quad (2)$$

with long memory parameter $0 < d < \frac{1}{2}$. It can be shown that there is a simple mapping between the long memory parameter d and the tail index b of the density $f(\alpha)$, which is

$$d = 1 - b. \quad (3)$$

The above results can be generalized to all distributions whose behavior close to unity can be approximated by $(1 - \alpha)^{b-1}$. Furthermore, the results can also be extended within a general autoregressive moving average (ARMA) framework where the implication for aggregation depend on the shape of the distribution around unity of the largest autoregressive root of the ARMA process. See Zaffaroni (2004) for more details.

To evaluate the following empirical application, we will specifically exploit the theoretical link between the distribution of the micropersistence parameters and the persistence of the aggregate.

3. The data

The data set utilized in our analysis comprises 470 seasonally adjusted and annualized quarter on quarter changes of consumer price sub-indices from France, Germany and Italy. The data runs from 1985 Q1 to 2004:Q2. Long time series of CPI sub-indexes are not readily available in most countries of the euro area. Therefore, we have been obliged to limit our data to France, Germany and Italy. Nevertheless, this has not undermined the representativeness of our exercise given that those three countries together account for roughly 70 percent of the euro area population and consumption.

Unfortunately, the availability and cross-countries comparability of the data has induced us to work with consumer price data, which are a measure of the cost of goods and services as purchased by final consumers. CPI includes also the margins charged by retailers at the various steps of the distribution process. This implies that such consumer prices do not properly capture price dynamics at the production level nor the cost of intermediate inputs paid by firms.

Concerning the choice of the sample period, we focus on the post-1985 data for two reasons. First, the German data are not available before 1985. Second, we want to avoid that our results on inflation persistence are impaired by possible deterministic breaks over the estimation period. Many empirical studies have showed that the mid-eighties marked a significant break in average inflation in most OECD countries.⁵ For all these reasons, we deem the post-1985 sample as appropriate to study the persistence of inflation in the euro area. The three CPI original databases together comprise 470 sub-indices, but 66 of these were not suitable for estimation of ARMA models either because they have too few observations (e.g. some sub-indices are available only since year 2000) or because they correspond to items which prices are set at discrete intervals (e.g. tobacco or postal services). We are left with 404 'well-behaved' series each having 78 quarterly observations.

3.1. Sectoral inflation series: descriptive statistics

We now turn to the properties of the sub-sector inflation rates. Table 2 reports descriptive statistics of aggregate inflation and of the distribution of the 404 sectoral inflation rates. Let us emphasize four points.

First, the inflation sectoral averages are quite disperse with 50 percent of their distribution ranging from 1.8 to 4.2 percent. The mean of aggregate inflation is 2.6 percent. Second, sectoral inflation is noticeably more volatile than aggregate inflation. On average, the standard deviation of the inflation sub-indices is equal to 3.6 percent, i.e. nearly three times as large the standard deviation of the aggregate. This much higher volatility is a common feature of the inflation rates of sectoral prices as shown in Clark (2006) and Bilke (2005). Third, the persistence of sectoral inflation rates is also clearly smaller than the one of the aggregate inflation series. A first measure of persistence is given by the largest root of an ARMA model fitted on the data. This measure is equal to 0.93 in the case of the aggregate inflation and this is roughly equal to the 75th percentile of the distribution of the same measure on the sector's data.⁶ An alternative measure of persistence is given by the long memory parameter, here denoted d_i for each sub-indices inflation rate π_{it} (see Eq. (2)). If the large majority of the sub-indices inflation rates displays long memory, then the aggregation exercise would be trivial. The last column of the table, which reports the statistics relative to the distribution of estimated long memory parameters d_i based on the parametric Whittle estimator (see Brockwell and Davis, 1987), confirms that only very few sectoral inflation rates exhibit long memory: d_i is not different from or close to $-\frac{1}{2}$ for a vast majority of the cross-section. However, as shown below, our

⁴ When $-\frac{1}{2} < d < 0$ in (2), under further regularity conditions, the U_t is said to display *anti-persistence*, which is a variation of long memory for which the acf becomes summable. We disregard this case because it is not relevant for our empirical application.


⁵ Among the others, Corvoisier and Mojon (2005) suggests that Italian, French and German time series of inflation all admit a break in the mid-eighties.

⁶ The largest root is the one associated to the best ARMA(p, q) as selected by the Akaike (AIC) criteria with $0 \leq p, q \leq 4$, estimated using the ARMAX procedure of MatLab.

Table 2

Descriptive statistics of the 404 sectoral inflation rates (first two columns annualized q–o–q inflation rates).

Aggregate of 404	Mean	Stand. dev.	Larg. ARMA root	Long mem. d
	2.6	1.1	0.93	0.18
<i>Cross-section characteristics</i>				
Weighted mean	2.6	3.6	0.78	–0.33
Unweighted mean	2.4	3.5	0.72	–0.36
Minimum	–11.3	0.7	–0.81	–0.50
25th percentile	1.8	1.7	0.71	–0.50
Median	2.6	2.5	0.83	–0.43
75th percentile	4.2	4.0	0.90	–0.18
Maximum	8.1	25.4	1.02	0.32
<i>Mean for selected sub-sets</i>				
France	1.8	3.0	0.72	–0.34
Germany	1.5	3.2	0.71	–0.34
Italy	3.6	4.3	0.73	–0.36
Processed food	2.6	3.1	0.68	–0.20
Unprocessed food	2.3	3.9	0.52	–0.33
NE indus. goods	2.0	2.2	0.77	–0.38
Energy	1.9	8.4	0.78	–0.37
Services	3.3	3.0	0.74	–0.33

**Fig. 1.** First 10 dynamic principal components: variance accounted.

estimate of d for the aggregate inflation rate is 0.18, well above the 75th percentile of the sectoral d_i parameters (–0.18). Finally, we observe sharper differences in persistence across the main sectoral groupings of the CPI (processed food, unprocessed food, energy, non-energy industrial goods—NEIG—and services) than across countries. The gap between the ARMA largest root of the inflation process of unprocessed food prices (0.52) and one of the energy (0.78) is wider than between the root associated to the ARMA processes fitted on the inflation of German, French and Italian prices. Comparing the main groupings of CPI sub-indices across countries, we also find that the sectoral hierarchy at the euro area level applies within each country.

To conclude, we find clear evidence that the inflation rates of the individual sub-indices are way more volatile and much less persistent than the inflation rate of the aggregate CPI index. Moreover, volatility and persistence are more sector than country specific.

3.2. Behind the aggregation mechanism

This section assesses the two elements that play a crucial role in shaping the effect of cross-section aggregation of time series: the presence of common shocks and of heterogeneity in the propagation mechanism of those shocks. This means that, using the notation of Section 2, we are checking whether there is at least one common shock u_t and whether the α_i are different across i . The analysis here is carried out using non-parametric methods so that the results do not depend on any

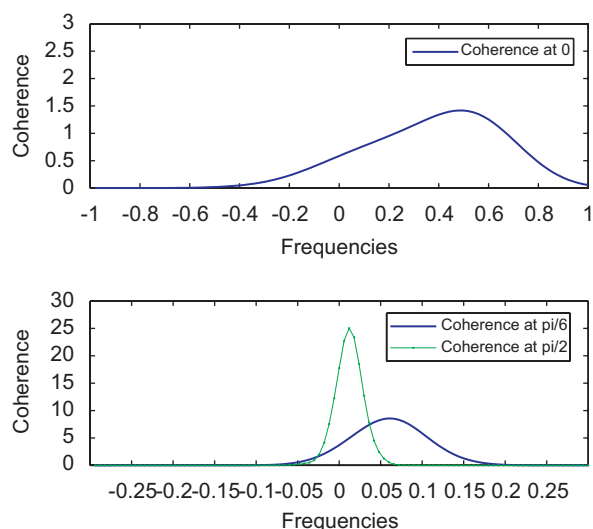


Fig. 2. Coherence between aggregate and sub-indices at various frequencies.

parametric assumptions. Following the recent literature on factor model in large cross-sections⁷ we estimate the first 10 dynamic principal components of the autocovariance structure of the sectoral inflation series. The dynamic principal component analysis provides indications on the number of common shocks explaining the correlation structure in the data (see Forni et al., 2000).

Fig. 1 presents the spectrum of the first 10 dynamic principal components⁸ of the 404 inflation time series.

The variance of sectoral inflation is strikingly dominated by one common factor. Fig. 1 also shows that this first common factor is the only one for which the spectrum is concentrated on low frequencies. Hence this factor drives the persistence observed in sectoral inflation.⁹ The other factors account for a much smaller share of the variance than the first one. They are also more equally relevant at all frequencies, as indicated by their relatively flat patterns in Fig. 1.¹⁰

On the basis of these results, we opt for a factor model of the sectoral inflation series that admits a *single common shock*, modelling the sectoral inflation time series π_{it} as

$$\pi_{it} = \delta_{0i} + \Psi_i(L)u_t + \xi_{it}, \quad i = 1, \dots, n, \quad (4)$$

where u_t is the common shock, ξ_{it} is a stationary idiosyncratic component, orthogonal to the common component and δ_{0i} is the constant mean parameter. $\Psi_i(L)$ is a lag polynomial for the i th unit which represents the propagation of the common shock through the π_{it} process. In the next section we consider parametric specifications of both $\Psi_i(L)$ and ξ_{it} .

Model (4) is here to be intended specifically as a reduced-form model once all the simultaneous, endogenous, relations are resolved. Indeed, the fact that we are working only with consumer price indexes prevents us from spelling out a complete simultaneous model where all the relations between prices at the various stages of the production process are properly accounted for (i.e. the role of the price of intermediate inputs). Furthermore, our model is designed to be in reduced form because we do not pursue any identification of the nature of the shocks, such as a productivity shock or a monetary policy shock, which would instead require a structural model, coupled with a suitable set of identification conditions.

However, one might still wonder whether our reduced-form approach can shed some evidence on which source of shocks is the dominant one, if any. A simple consequence of a purely monetary shock is its neutrality with respect to relative prices long run dynamics, in line with the definition of *pure inflation* given by Reis and Watson (2008). For this reason, we look at whether the propagation mechanism of the common shock in the cross-section of price sub-indices is homogenous across items. Still resorting to spectral analysis, we estimate the coherence¹¹ between the first principal component and each one of the 404 series, in order to obtain ‘correlation’ at different frequencies.

Fig. 2 reports the cross-section distribution of the coherence values at three frequencies: 0 (long run), $\pi/6$ (three years periodicity) and $\pi/2$ (one year periodicity), respectively.

⁷ See Forni et al. (2000), Stock and Watson (2002) and the application by Clark (2006) to US disaggregate consumption deflators.

⁸ Dynamic principal components are calculated as the eigenvalue decomposition of the multivariate spectra of the data at each frequency. The acf up to eight lags has been used in the construction of the multivariate spectral matrix. The data have been standardized to have unit variance before estimating the multivariate spectra.

⁹ The height of the spectrum at frequency zero is a well-known non-parametric measure of the persistence of a time series.

¹⁰ Clark (2006) finds similar results on US inflation data.

¹¹ Coherence measures the degree at which two variables are moving together at a given frequency. It is a quantity always between 0 and 1; see Brockwell and Davis (1987) for details.

The figure supports the presumption that the common shock transmission to sectoral inflation is highly heterogenous. Such evidence plays against the idea that the our common shock can be identified as a monetary policy one. However, we leave the issue of the shock identification to future work.

The following section addresses precisely the issue of how to model and estimate, based on a parametric one-factor linear model, the heterogeneity manifested by Fig. 2.

4. Model: specification and estimation

The quarterly rate of change of each sectoral price sub-index π_{it} is assumed to behave according to the following parametric specification of (4):

$$\pi_{it} = \delta_{0i} + \frac{\Theta_i^u(L)}{\Phi_i^u(L)} u_t + \frac{\Theta_i^\varepsilon(L)}{\Phi_i^\varepsilon(L)} \varepsilon_{it}, \tag{5}$$

meaning that we are considering a particular case of (4) with $\Psi_i(L) = \Theta_i^u(L)/\Phi_i^u(L)$ and $\xi_{it} = \Theta_i^\varepsilon(L)/\Phi_i^\varepsilon(L)\varepsilon_{it}$, where as before $u_t \sim i.i.d.(0, 1)$ is the common shock and now $\varepsilon_{it} \sim i.i.d.(0, \sigma_i^2)$ is the idiosyncratic shock for each $i = 1, \dots, n$, also mutually independent from u_t . No distributional assumptions are made. The above polynomials in the lag operator are assumed to satisfy the usual stationarity-invertibility conditions, that is they have all roots outside the unit circle in the complex plane. Also $\Theta_i^x(L), \Phi_i^x(L)$, respectively, of finite order q_i^x, p_i^x , have no common roots for $x \in \{u, \varepsilon\}$ ensuring identification. Thus, each π_{it} behaves as a stationary ARMA(p_i, q_i) with a possibly non-zero mean δ_{0i} , where standard arguments on sum of ARMA processes yield $q_i \leq \max\{q_i^u + p_i^\varepsilon, q_i^\varepsilon + p_i^u\}$ and $p_i \leq p_i^u + p_i^\varepsilon$, equality holding if the two AR polynomials, $\Phi_i^\varepsilon(L)$ and $\Phi_i^u(L)$, have no roots in common. Note that we allow full heterogeneity of all the parameters characterizing model (5). The aggregation result requires to estimate the cross-sectional distribution of the maximal AR roots of $\Phi_i^u(L)$ and $\Phi_i^\varepsilon(L)$.

We assume that the ε_{it} are mutually orthogonal across units. This is merely a simplification assumption, not required in fact for the aggregation theory, used in this paper, to be applicable. In fact, a factor model, such as (5), represents a parsimonious way to capture cross-sectional dependence. This enters primarily through the common component, representing a form of pervasive dependence. However, dependence is in principle permitted also through cross-correlated ε_{it} (i.e. $cov(\varepsilon_{i,t}, \varepsilon_{j,t}) \neq 0$ for $i \neq j$) as long as this form of cross-correlation is sufficiently weak and not pervasive in the cross-section, ensuring that model (5) belongs to the class of approximate factor structure.¹² For example, one might consider the case of local correlation arising between the prices of items belonging to the same sector as follows: the n units are classified into $m < n$ sectors, with $cov(\varepsilon_{i,t}, \varepsilon_{j,t}) = \rho_l \neq 0$ for all units belonging to the l th sector, with $1 \leq l \leq m$ and $cov(\varepsilon_{i,t}, \varepsilon_{j,t}) = 0$ for all $l \neq r$, any i, j . In this case, the approximate factor structure condition is still satisfied as long as the number of units belonging to each sector do not increase too fast with n . However, we do not endorse these type of parameterizations here. In fact allowing cross-sectional dependence among the ε_{it} is not of primary importance for this paper because, from an econometric perspective, (i) the aggregation mechanism depends solely on the characteristics of the common component (see Zaffaroni, 2004), (ii) our estimation strategy is robust to many forms of cross-sectional dependence among the ε_{it} as, for example, an approximate factor structure, and (iii) we are modelling consumer price indexes, which makes difficult to properly account for input/output relations across sectors.

4.1. Estimation strategy

The estimation of model (5) is non-standard. First, the large dimensionality (large n) rules out the recourse to the conventional Kalman filter approach typically used for univariate or for small multivariate ARMA. Second, the recent techniques for estimation of dynamic factor models, all based on the principal component approach such as Forni et al. (2000) and Stock and Watson (2002), would be inappropriate in our model due to the presence of sector specific AR components $\Phi_i^u(L)$ and $\Phi_i^\varepsilon(L)$.

We propose instead the following multi-stage procedure:

- (i) For each unit i , we estimate the following ARMA(p_i, q_i) implied by (5), namely $\Phi_i(L)\pi_{it} = \delta_{0i} + \Xi_i(L)z_{it}$, setting $\delta_{0i} = \Phi_i(1)\delta_{0i}$, $\Phi_i(L) = \Phi_i^u(L)\Phi_i^\varepsilon(L)$ with MA component $\Xi_i(L)z_{it} = \Theta_i^u(L)\Phi_i^\varepsilon(L)u_t + \Theta_i^\varepsilon(L)\Phi_i^u(L)\varepsilon_{it}$, for some white-noise sequence z_{it} and a polynomial $\Xi_i(L)$ of order q_i . Simplification occurs if there are roots in common. Note that here we need not distinguish between the common and the idiosyncratic shock, since this is unnecessary for consistent estimation of the intercept δ_{0i} , and of the AR coefficients in $\Phi_i(L)$, which are the goal of this stage.
- (ii) We average the estimated MA component $\hat{\Xi}_i(L)\hat{z}_{i,t}$ across i yielding

$$\hat{\lambda}_{t,n} = \sum_{i=1}^n w_i \hat{\Xi}_i(L)\hat{z}_{i,t}, \tag{6}$$

¹² An approximate factor structure is qualified by the condition that the maximum eigenvalue of the covariance matrix $E\varepsilon_t\varepsilon_t'$ remains bounded as n increases, where $\varepsilon_t = (\varepsilon_{1t}, \dots, \varepsilon_{nt})'$; see Chamberlain and Rothschild (1983).

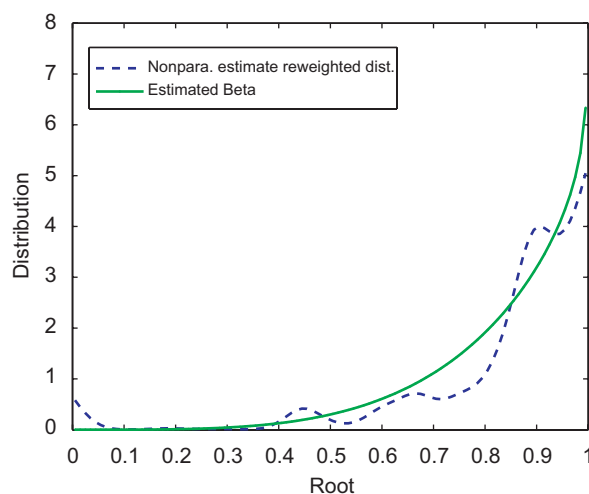


Fig. 3. Distribution of the maximal autoregressive root.

where the w_i are the CPI weights. This yields an estimate of its common part, namely the limit of $(\sum_{i=1}^n w_i \Theta_i^u(L) \Phi_i^c(L)) u_t$, where the approximation improves as n grows since the idiosyncratic part vanishes in mean-square as $n \rightarrow \infty$. We then fit a finite order MA to the $\hat{x}_{t,n}$ in order to get, at this stage, an estimate of the common innovation \hat{u}_t .

- (iii) Using the \hat{u}_t as an (artificial) regressor, we finally fit model (5) to each π_{it} but with the caution of using it as an ARMAX(p_i, q_i, q_i) process (an ARMA process with exogenous regressors). Note that, in doing so, we are not imposing that the AR roots of the common and idiosyncratic components coincide, a feature admitted by the so-called Box–Jenkins model structure. At this stage we get all parameter estimates.

Steps (ii) and (iii) of the procedure can be iterated in order to improve the estimate of the common shock as well as of the model parameters. Such procedure does not require any distributional assumption but its validity requires that both n, T to diverge to infinity. Details of the algorithm can be found in Altissimo and Zaffaroni (2004). The described iterative procedure is similar to a recent modification of the EM algorithm for the estimation of factor models in the presence of idiosyncratic AR components recently proposed by Stock and Watson (2005).

We set $0 \leq p_i^u, p_i^c \leq 2, 0 \leq q_i^u, q_i^c \leq 4$ implying in stage (i) estimation of ARMA(p_i, q_i) with $0 \leq p_i, q_i \leq 4$.¹³ The order of the models (q_i, p_i in step (i), q_i in step (ii) and $p_i^u, p_i^c, q_i^u, q_i^c$ in step (iii)) are selected based on the AIC criteria. At each iteration and for each i this means that we have estimated $3^2 \times 5^2 = 225$ specifications, given the chosen range for $p_i^u, p_i^c, q_i^u, q_i^c$. Our choice on the orders is made solely to contain computational time and to follow the principle of parsimony, since the larger the orders, the larger would be the number of parameters to be estimated. However, in terms of the aggregation mechanism, note that no effect would arise from taking arbitrarily large, yet finite, MA and AR orders.

4.2. Results for sectorial inflation rates

This section briefly describes the estimates of up to 6000 parameters (up to 15 parameters for each of the 404 sub-index time series, namely the AR and MA parameters, the mean and the idiosyncratic innovation variance) of the model. First, the estimated common shock u_t turns out to be white, with a non-significant autocorrelation, corroborating the *i.i.d.* hypothesis. Second, the idiosyncratic volatility σ_i is substantially larger than the common shock volatility, in fact six times larger. The median of the distribution of the absolute value of the estimated first loading is 0.06, whereas we obtain 0.38 for the standard deviation of the idiosyncratic component. This strikingly confirms that most of the variance of sectoral prices is indeed due to sector specific shocks. Third, we turn to $\Phi_i^u(L)$, which dominates the dynamic effects of the common shock on sectorial inflation. In Fig. 3 (blue dashed line), we show the kernel estimate of the distribution of the signed modulus of the maximal AR root of such polynomial.¹⁴

It turns out that this distribution is dense near unity with a median of 0.82 and a long tail to the left. Fourth, we compare the dynamics of sectorial inflation rates and main CPI groupings in Table 3. The table reports the number and the relative

¹³ We use the `armax` and `bj` procedures of `MatLab`.

¹⁴ We did sign such modulus so that we could distinguish between the effect of a negative root from a positive one and also consider the effect of complex roots.

Table 3

Summary statistics of estimated parameters for sectorial inflation rates.

	#	CPI weight	Largest root > 0.875			Largest root > 0.925		
			#	Freq.	CPI weight	#	Freq.	CPI weight
EA3	404	1	142	0.35	0.48	54	0.27	0.26
Germany	87	0.42	31	0.34	0.23	11	0.12	0.14
France	147	0.30	55	0.38	0.15	23	0.17	0.06
Italy	170	0.28	56	0.35	0.11	20	0.12	0.06
Processed food	65	0.14	9	0.14	0.03	2	0.03	0.01
Unprocessed food	42	0.07	7	0.17	0.01	3	0.07	0.00
Non-energy ind. goods	167	0.33	68	0.42	0.14	25	0.15	0.08
Energy goods	18	0.07	8	0.50	0.05	5	0.31	0.03
Services	114	0.39	50	0.44	0.25	19	0.17	0.14

frequency of series having roots above given thresholds in specific sub-category as well as the total weight of such series in the overall CPI.

Looking at the cross-country pattern, highly persistent sub-sectors in Germany account for a much larger share of the overall euro area CPI; this effect is mainly due the high persistence (AR root of 0.96) of German housing expenditure inflation,¹⁵ which accounts for around 8 percent of the overall CPI. Furthermore, the service sector turns out to be the most relevant for the dynamics of aggregate inflation either because it has the largest share of highly persistent series and the highest weight in the consumption basket.

4.3. Results for aggregate inflation rate

Given the estimates of the microparameters, we are now in a position to infer the dynamic properties of the aggregate induced by the behavior of the microtime series, which represent the core results of the paper. We consider three different types of aggregation schemes. First, we reconstruct the aggregate as an exact weighted average of the individual microtime series and in this way we exactly recover the contribution of the common shocks to the aggregate inflation and to aggregate persistence. This *exact* aggregation result will corroborates the plausibility of one dominant common shock specification. Second, we exploit the theoretical link between the distribution of the largest AR root of sectorial inflation rates and the autocovariance structure of the aggregate, as presented in Section 2, to infer the dynamic properties of the latter. We shall call this *asymptotic* aggregation result, which is the main finding of the paper. Third, we consider a so-called *naïve* aggregation scheme based on the (wrong) presumption that the aggregate model has the same functional form as the individual models, in our case an ARMA. This exercise will warn against not properly accounting for the aggregation effects.

Summarizing the results described below in detail, we claim that the analysis of the microdeterminants of the aggregate inflation supports the view that aggregate inflation in our sample period can be well described by a stationary long memory process. We will show that starting from very simple ARMA process at microlevel we have been able to properly reconstruct the dynamic properties of the aggregate. We will also show that the microvolatility and low persistence can be squared with the aggregate smoothness and persistence.

4.3.1. Exact aggregation

The aggregate inflation data are defined as the weighted average of the sectorial inflation rates.¹⁶ Therefore using the estimates of the model in (5) it follows:

$$\Pi_{n,t} = \sum_{i=1}^n w_i \pi_{it} = \sum_{i=1}^n w_i \hat{\delta}_{0i} + \hat{u}_t \sum_{i=1}^n w_i \frac{\hat{\Theta}_i^u(L)}{\hat{\Phi}_i^u(L)} + \sum_{i=1}^n w_i \frac{\hat{\Theta}_i^v(L)}{\hat{\Phi}_i^v(L)} \hat{\varepsilon}_{it} = \hat{\delta}_0 + \hat{\Psi}(L) \hat{u}_t + \hat{\xi}_t,$$

where the w_i are the CPI weights. Hence, the aggregate inflation is decomposed into two components, one associated with the common shocks, \hat{u}_t , and its propagation mechanism, $\hat{\Psi}(L)$, and a second associated with the idiosyncratic component $\hat{\xi}_t$, where $\hat{\Psi}(L) = \sum_{i=1}^n w_i \hat{\Psi}_i(L)$, viz. a weighted average of the propagation mechanisms at microlevel. Fig. 4 shows the reconstructed aggregate, $\Pi_{n,t}$, versus its common component, $\hat{\Psi}(L) \hat{u}_t$. There is a high correlation between the two components, around 0.76, but the former is clearly more volatile. This suggests that the idiosyncratic component $\hat{\xi}_t$ has not been completely washed out at the aggregate level. However, the effect of the idiosyncratic component has been drastically

¹⁵ The sub-index is rent (including imputed rents from owner-occupied houses), which account for around 20 percent of the German CPI, while it is only 3 percent of the Italian and French.

¹⁶ Statistical offices do not aggregate inflation rates but first they aggregate price indices and then compute the aggregate inflation rate. Here we ignore the possible effect induced by such non-linear transformation since not relevant for the dynamic properties of the aggregate.

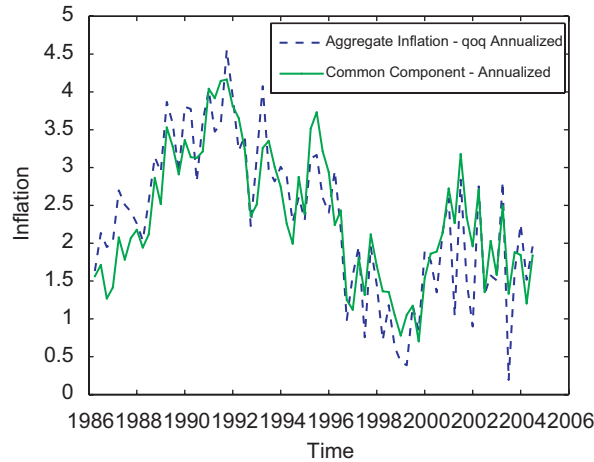


Fig. 4. Aggregate CPI inflation versus estimated common component.

reduced: the average of the estimated variances of the idiosyncratic components, namely $\sum_{i=1}^n w_i \widehat{\text{var}}(\hat{\xi}_{it})$, is equal to 0.2, which is four times larger than the estimated variance of the average of the idiosyncratic components, namely $\widehat{\text{var}}(\hat{\xi}_t) = \widehat{\text{var}}(\sum_{i=1}^n w_i \hat{\xi}_{it})$. Since the idiosyncratic component has little persistence, $\Psi(L)\hat{u}_t$ can be interpreted as a measure of ‘core inflation’, possibly relevant for monitoring and forecasting inflation.

4.3.2. Asymptotic aggregation

In this section, we provide a more formal link between microheterogeneity and aggregate persistence based on the arguments introduced in Section 2. However, differently from the simple example discussed there, we first need to take into account that CPI sub-sectors have different weights in the aggregate. To overcome this problem, we implement a relative re-weighting of the 404 maximal AR roots in function of the relative weights of the respective sectors. Precisely, we bootstrap a sample of 10,000 data out of the 404 roots with relative frequency equal to the weighting scheme; hence roots associated to sectors will a larger weight w_i will be re-sampled more often.

For the bootstrapped sample, the distribution of the roots is estimated in two ways, parametrically and non-parametrically. The non-parametric estimate resorts on the use of a bounded kernel in order to account for the bounded support of the roots; such non-parametric estimate of the density function associated to this simulated sample is plotted in Fig. 3 (blue dashed line). Instead, for the parametric case, mirroring the example of Section 2, we fitted the Beta distribution to the data:

$$f(\alpha|\hat{a}, \hat{b}) = B^{-1}(\hat{a}, \hat{b})\alpha^{\hat{a}-1}(1 - \alpha)^{\hat{b}-1}, \tag{7}$$

where the Beta parameter estimates, obtained by maximum likelihood estimation, are respectively, $\hat{a} = 4.7, \hat{b} = 0.87$ (standard deviations equal to 0.17 and 0.02, respectively). The fitted Beta density is also plotted in Fig. 3 (green solid line). The behavior near unity of the two estimated densities, based on the parametric and non-parametric estimator, respectively, is remarkably similar and in the following we will focus on the estimated Beta.¹⁷

We now have all the ingredients to apply the arguments of Section 2. By (3) we immediately get $\hat{d} = 1 - \hat{b} = 0.13$ and plugging this value into (2) yields the (estimated) asymptotic behavior for the acf of the common component of the aggregate, $\Psi(L)u_t$, equal to

$$\text{cov}(\Psi(L)u_t, \Psi(L)u_{t+k}) \sim ck^{-0.84} \text{ as } k \rightarrow \infty. \tag{8}$$

In other words, it follows that the common component of aggregate inflation has long memory parameter equal to $\hat{d} = 0.13$. Therefore, the acf of the common component of aggregate inflation decays toward zero with an hyperbolic decay, and thus is markedly different from the behavior of the sectoral inflation processes.

So far we have backed out the aggregate memory from the microdata. One might wonder at what would be the estimate of the memory based directly on the aggregate data. It turns out that the estimate of the memory parameter¹⁸ d by using only the aggregate $\Pi_{n,t}$ is equal to 0.18, with standard error of the estimate equals to 0.20, which, given the distribution reported in the last column of Table 2, is reasonably close to 0.13, as recovered from the microstructure. Therefore, the aggregate data presents a long memory behavior that is not present in the microtime series; this long memory characteristic appears to be fully accounted for by aggregation.

¹⁷ Using simple regression method, we can show that the non-parametric estimate of the root distribution behaves near unity as $\hat{f}(\alpha) \sim c(1 - \alpha)^{-0.16}$ as $\alpha \rightarrow 1^-$, which is remarkably close to the tail behavior implied by (7) namely $f(\alpha|\hat{a}, \hat{b}) \sim c(1 - \alpha)^{-0.13}$ as $\alpha \rightarrow 1^-$.

¹⁸ We have used the log periodogram regression estimator of Robinson (1995).

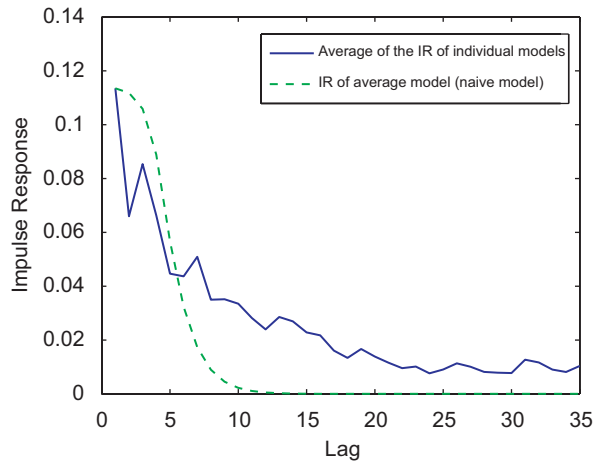


Fig. 5. Impulse responses.

4.3.3. Naïve aggregation

The above results are framed in term of acf and they show that the difference in persistence between the micro and macrodynamic is not necessarily inconsistent. The key element there is the aggregate impulse response function $\Psi(L)$, equal to the limit (in mean square) of $\sum_{i=1}^n w_i \Psi_i(L)$ (cf. Eq. (4)).

Another way to highlight the effects of aggregation on persistence is to consider the following naïve exercise. We construct a hypothetical average ARMA process (here called in fact naïve), whose roots are the average of the individual roots of the 404 estimated ARMAs. This means that we are considering the following impulse response function:

$$\Psi^{naive}(L) = \frac{(1 + \theta_1^u L) \dots (1 + \theta_q^u L)}{(1 - \phi_1^u L) \dots (1 - \phi_p^u L)}. \tag{9}$$

Here θ_h^u , for $1 \leq h \leq q$, is the (cross-sectional weighted) average of the h th root¹⁹ θ_{ih}^u of the MA polynomial $\Theta_i^u(L)$ corresponding to the common shock u_t (cf. model (5)). Similarly, ϕ_k^u , for $1 \leq k \leq p$, is the average of the k th root ϕ_{ik}^u of the AR polynomial $\Phi_i^u(L)$ corresponding to the common shock u_t .

The difference between $\Psi(L)$ and $\Psi^{naive}(L)$ arises whenever $p > 0$, that is when there is an autoregressive component, since the impulse response is a nonlinear function of the autoregressive parameters. Therefore the coefficients embedded within the impulse response $\Psi(L)$ decay hyperbolically, a symptom of long memory, whereas the coefficients of the naïve impulse response $\Psi^{naive}(L)$ decay exponentially. The idea of the exercise is to see the aggregate response to the common shock in case the propagation mechanism is equal across agents versus the case of different propagation mechanisms, i.e. to quantify the effect of heterogeneity and aggregation.²⁰

Fig. 5 compares the estimated impulse responses²¹ implied by $\hat{\Psi}^{naive}(L)$ and $\hat{\Psi}(L)$, respectively.

The exercise is quite instructive. In the hypothetical case of homogeneity of the micropropagation mechanism, after four years a common shock u_t would be completely absorbed (green dashed line), while in reality, due to the presence of heterogeneity and of some very persistent microunits, around 20 percent of the shock would not have been absorbed (blue solid line). In other words, focusing on the naïve model, that uses (9), leads to understate the degree of persistence in the data in a quantitatively important way; once and again this results warns against the practice of using average microbehavior in order to calibrate macrophenomena.

5. Conclusions

In this paper, exploiting the heterogeneity in the inflation dynamics across CPI sub-indices within the euro area, we investigate the role played by cross-sectional aggregation in explaining some of the difference observed between micro and macroinflation dynamics. We focus in particular on the link between CPI sub-indices and the aggregate CPI. Our results are able to square the volatility and low persistence observed, on average, at the level of sectoral inflation with the smoothness and persistence of the aggregate. The persistence that we obtain through this aggregation exercise mimics remarkably well the persistence observed in the aggregate inflation. In particular, aggregate inflation turns out to be well described by a stationary long memory process. The persistence of the aggregate inflation is mainly due to the high persistence of some

¹⁹ Hereafter, we follow the convention of referring to the inverse of the roots as the roots so that, for example, the root of the polynomial $1 - \phi L$ will be indicated by ϕ rather than by ϕ^{-1} .

²⁰ See also the presentation of similar exercises in the context of microfounded model, though on hypothetical distribution by Carvalho (2006).

²¹ Here $\hat{\Psi}^{naive}(L)$ is obtained by plugging the sample weighted means $\hat{\theta}_h^u = \sum_{i=1}^n w_i \hat{\theta}_{ih}^u$, $1 \leq h \leq q$, and $\hat{\phi}_k^u = \sum_{i=1}^n w_i \hat{\phi}_{ik}^u$, $1 \leq k \leq p$, into (9).

sub-indices mainly concentrated in the service sector, such as housing costs in Germany. Altogether, we show the importance of heterogeneity and aggregation for understanding the persistence of inflation at the macroeconomic level. We leave the design and estimation of stylized models of the business cycle that can be consistent with both heterogeneity at the microlevel and the implied persistence at the macrolevel for future research.

References

- Altissimo, F., Zaffaroni, P., 2004. Towards understanding the relationship between aggregate fluctuations and individual heterogeneity, Imperial College London, preprint.
- Bilke, L., 2005. Break in the mean and persistence of inflation: a sectoral analysis of French CPI. ECB Working Paper 413.
- Bils, M., Klenow, P., 2004. Some evidence on the importance of sticky price. *Journal of Political Economy* 112, 947–985.
- Boivin, J., Giannoni, M., Mihov, I., 2006. Sticky prices and monetary policy: evidence from disaggregated U.S. data. *The American Economic Review*, Forthcoming.
- Brockwell, J., Davis, R., 1987. *Time Series: Theory and Methods*. Springer, Berlin.
- Carvalho, C., 2006. Heterogeneity in price setting and the real effects of monetary shocks, Princeton University, preprint.
- Chamberlain, G., Rothschild, M., 1983. Arbitrage, factor structure and mean-variance analysis on large asset markets. *Econometrica* 51, 1281–1304.
- Chari, V., Kehoe, P., McGrattan, E., 2000. Sticky price models of the business cycle: can the contract multiplier solve the persistence problem? *Econometrica* 68, 1151–1179.
- Clark, T., 2006. Disaggregate evidence on the persistence of US consumer price inflation. *Journal of Applied Econometrics* 21, 563–587.
- Corvoisier, S., Mojon, B., 2005. Breaks in the mean of inflation: how they happen and what to do with them. ECB Working Paper 451.
- Dhyne, E., Alvarez, L., Le Bihan, H., Veronese, G., Dias, D., Hoffman, J., Jonker, N., Lünemann, P., Rumler, F., Vilmunen, J., 2005. Price setting in the euro area: some stylised facts from individual consumer price data. ECB Working Paper 523.
- Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2000. The generalized factor model: identification and estimation. *The Review of Economics and Statistics* 82, 540–554.
- Granger, C., 1980. Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* 14, 227–238.
- Imbs, J., Mumtaz, H., Ravn, M., Rey, H., 2005. PPP strikes back: aggregation and the real exchange rate. *Quarterly Journal of Economics* CXX, 1–43.
- Pivetta, F., Reis, R., 2007. The persistence of inflation in the United States. *Journal of Economic Dynamics and Control* 31, 1326–1358.
- Reis, R., Watson, M., 2008. Relative goods' prices, pure inflation, and the Phillips correlation, Princeton University, preprint.
- Robinson, P.M., 1978. Statistical inference for a random coefficient autoregressive model. *Scandinavian Journal of Statistics* 5, 163–168.
- Robinson, P.M., 1995. Log-periodogram of time series with long range dependence. *Annals of Statistics* 23, 1048–1072.
- Stock, J., Watson, M., 2002. Macroeconomic forecasting using diffusion indices. *Journal of Business and Economic Statistics* 20, 147–162.
- Stock, J., Watson, M., 2005. An empirical comparison of methods for forecasting using many predictors, Harvard, preprint.
- Zaffaroni, P., 2004. Contemporaneous aggregation of linear dynamic models in large economies. *Journal of Econometrics* 120, 75–102.