

THE DYNAMICS OF EUROPEAN INFLATION EXPECTATIONS

JÖRG DÖPKE, JONAS DOVERN, ULRICH FRITSCHE AND JIRKA SLACALEK

December 18, 2006

ABSTRACT. We investigate the relevance of the Carroll's sticky information model of inflation expectations for four major European economies (France, Germany, Italy and the United Kingdom). Using survey data on household and expert inflation expectations we argue that the model adequately captures the dynamics of household inflation expectations. We estimate two alternative parametrizations of the sticky information model which differ in the stationarity assumptions about the underlying series. Our baseline stationary estimation suggests that the average frequency of information updating for the European households is roughly once in 18 months. The vector error-correction model implies households update information about once a year.

Keywords: Inflation expectations, sticky information, inflation persistence

JEL Classification: D84, E31

Jörg Döpke, Deutsche Bundesbank, joerg.doepke@bundesbank.de; Jonas Dovern, Kiel Institute for World Economics (IfW Kiel), jonas.dovern@ifw-kiel.de; Ulrich Fritsche, University Hamburg and German Institute for Economic Research (DIW Berlin), ufritsche@diw.de; Jirka Slacalek, German Institute for Economic Research (DIW Berlin), jslacalek@diw.de. The second author worked on the paper while staying in the Department of Macro Analysis, German Institute for Economic Research (DIW Berlin). We thank Christopher Carroll, Christina Gerberding, Heinz Herrmann and Christian Schumacher for helpful comments. We are grateful to various seminar audiences. We also thank Cédric Viguiè for providing us with disaggregated survey response data from European Commission's Harmonized Business and Consumer Surveys and Christina Gerberding for her data on inflation and GDP expectations. The views presented in this paper are the authors', and do not necessarily reflect those of the Deutsche Bundesbank, DIW Berlin or IfW Kiel.

1. INTRODUCTION

Several recent papers (including Mankiw and Reis, 2002, 2006) argue that sticky information models, in which agents update their information occasionally rather than instantaneously, resolve some puzzles in the output–inflation dynamics. For example, sticky information models are able to account for considerable inflation persistence and substantial sacrifice ratios (recessionary disinflations) typically observed in the data.

Microeconomic foundations for the sticky information paradigm were elaborated in Carroll’s (2003) work on the “epidemiological model of expectations.” Carroll argues that US survey data on inflation expectations are consistent with a model in which, in each period, only a fraction of households adopts inflation forecasts of rational experts. The remaining households find it costly to update their information and continue using their past expectations rather than form fully rational predictions. In related work Sims (2003), Branch (2004) and others provide alternative justifications for models with agents that do not instantaneously incorporate all available information as implied by most standard macro models.

While the sticky information approach seems to be useful for modelling US data, corresponding evidence for European countries is, to the best of our knowledge, still lacking.¹ This paper attempts to fill this gap by investigating inflation expectation data from four major EU economies (France, Germany, Italy and the United Kingdom). We believe it is particularly interesting to compare these results to each other and to the US since the institutional settings in Europe and the US differ substantially in at least two ways. First, the monetary policy set-up and recent inflation dynamics in various EMU countries, the UK and the US are quite varied. For example, whereas Germany, under the Bundesbank regime, has always had

¹The only work testing sticky information models on international data is Khan and Zhu (2002) and Handjiyska (2004). However, these two papers have to adopt some restrictive assumptions to circumvent data limitations: Khan and Zhu approximate agents’ expectations with forecasts from a VAR model. Handjiyska uses only interpolated data on expert expectations.

moderate and stable inflation rates, Italy faced considerably higher inflation rates in the early 1990s and has witnessed pronounced declines in price level increases over the past decade in the run-up to and after the introduction of the euro. In addition, different communication strategies of central banks might affect how information spreads across households.² Second, both the size and structure of the forecasting industry are dissimilar. In the US it is dominated by private forecasters, while in Europe public forecasters play a more prominent role. These factors may, in principle, affect how much the Carroll's sticky information model is relevant for European countries as well as the implied speed of adjustment of household expectations.

Our findings in general support the usefulness of the Carroll's sticky information model for the description of inflation dynamics in European countries. We find that household inflation expectations adjust sluggishly to the more precise predictions of professional forecasters. The average speed of this adjustment varies little across the four countries we investigate and is somewhat higher than that in the US: a typical household updates its inflation expectations roughly once in eighteen months (compared to once a year previously found in the US). While this result is quite robust across the estimation methods, we find that the frequency of information updating in Europe is somewhat lower for the vector error-correction specification, amounting to about once a year. Similarly to the US, European households are not backward-looking: they tend to update their expectations from experts' rational forecasts rather than actual past inflation rates. Finally, the estimates are stable over time: our data do not make it possible to distinguish any statistically significant differences between various institutional settings (e.g., inflation-targeters and non-targeters).

²See for example Ehrmann and Fratzscher (2005) and the literature cited therein.

2. THE EPIDEMIOLOGY OF HOUSEHOLD INFLATION EXPECTATIONS

Carroll (2003) proposes the following micro-founded model of the transmission of inflation expectations between professional forecasters and households. He argues that the dynamics of aggregate household expectations is adequately captured with a model in which households choose to update their expectations occasionally rather than instantaneously. New information about inflation spreads slowly across households in the following “epidemiological” way. Suppose a number of informed agents, experts, collects relevant information on future inflation in every period and makes rational inflation forecasts. These forecasts are published in newspapers. Households, on the other hand, find it costly to read the newspapers and to stay completely up-to-date (or make informed inflation forecasts). For that reason, in each period only a randomly chosen fraction λ of households follows the latest inflation stories in the newspapers and updates its inflation expectations. The remaining $1 - \lambda$ households stick to their forecasts from the previous period. The evolution of the (average) household (denoted HH) inflation (π) expectation (\mathbf{E}) follows:

$$\mathbf{E}_t^{HH} \pi_{t,t+1} = \lambda \mathbf{E}_t^{EX} \pi_{t,t+1} + (1 - \lambda) \mathbf{E}_{t-1}^{HH} \pi_{t,t+1},$$

where $\mathbf{E}_t^{HH} \pi_{t,t+1}$ and $\mathbf{E}_t^{EX} \pi_{t,t+1}$ denote one-period-ahead inflation expectations of households and experts (EX), respectively.

Thus, news about inflation can be thought of as a disease that spreads slowly across the population, infecting a fraction λ of all households in each period. The calculation outlined in detail in Carroll (2003) leads to the equation formulated for annual inflation rates, which are typically reported in surveys of inflation expectations:

$$\mathbf{E}_t^{HH} \pi_{t,t+4} = \lambda \mathbf{E}_t^{EX} \pi_{t,t+4} + (1 - \lambda) \mathbf{E}_{t-1}^{HH} \pi_{t-1,t+3}. \quad (1)$$

Equation (1) holds if (i) inflation follows a random walk process (or households believe that inflation is a random walk) or (ii) $\mathbf{E}_{t-1}^{HH} \pi_{t-1,t+3} \approx \mathbf{E}_{t-1}^{HH} \pi_{t,t+4}$.

Both of these assumptions are likely to be satisfied in our dataset. As discussed below, the underlying CPI inflation process in the core European economies has, indeed, been very persistent recently, warranting the random walk approximation. Second, given the high persistence of the inflation process, there is not much difference between household expectations as of time $t - 1$ of inflation rates at $t + 3$ and $t + 4$, which, in turn, implies that condition (ii) is also likely to be met.

3. EXPECTATION DATA

To test the model of information diffusion, two kinds of inflation expectation data are needed: inflation forecasts of households and professional forecasters. The forecasts of households were obtained from the European Commission's (EC) consumer survey and those of professional forecasters from Consensus Economics, a London-based macroeconomic survey firm.³

Household expectations were constructed using the EC survey's question 6, which asks how, by comparison with the last 12 months, the respondents expect that consumer prices will develop in the next 12 months.⁴ Unfortunately, the answers are qualitative rather than quantitative (unlike, for example, the question on expected inflation in the US Michigan Survey of Consumer Sentiment). This means that the respondents are asked about the direction of the expected movement of consumer prices (increase/fall), not about the exact quantitative value of this movement. Consequently, care needs to be taken when transforming these data into quantitative measures of household expectations, required to test equation (1). We follow much of the existing literature (including Gerberding, 2001 and Mankiw et al.,

³ We only investigate Germany, France, Italy and the United Kingdom because expectations of professional forecasters from other European countries, such as the Netherlands and Spain, are only available since 1996.

⁴The exact wording of question 6 of the Consumer Survey of the Joint Harmonised EU Programme of Business and Consumer Surveys is: "By comparison with the past 12 months, how do you expect that consumer prices will develop in the next 12 months?" For more information on the survey, see the Commission's web page, http://europa.eu.int/comm/economy_finance/indicators/.

2003) in adopting the Carlson and Parkin (1975) method, explained in the Appendix.

The procedure requires that specification of a variable that captures the perceived current level of inflation rate to rescale the expectations appropriately. We investigated three alternatives that have been proposed in the literature: recursive Hodrick–Prescott filter, (ii) past inflation and (iii) inflation extracted as fitted values from the regression of inflation on the past balance statistics from the survey (lagged by four quarters). One normalization that works well in terms of low mean squared error and is used in the paper is the recursive Hodrick–Prescott filter, in which the inflation trend was extracted in a quasi-real-time way.⁵

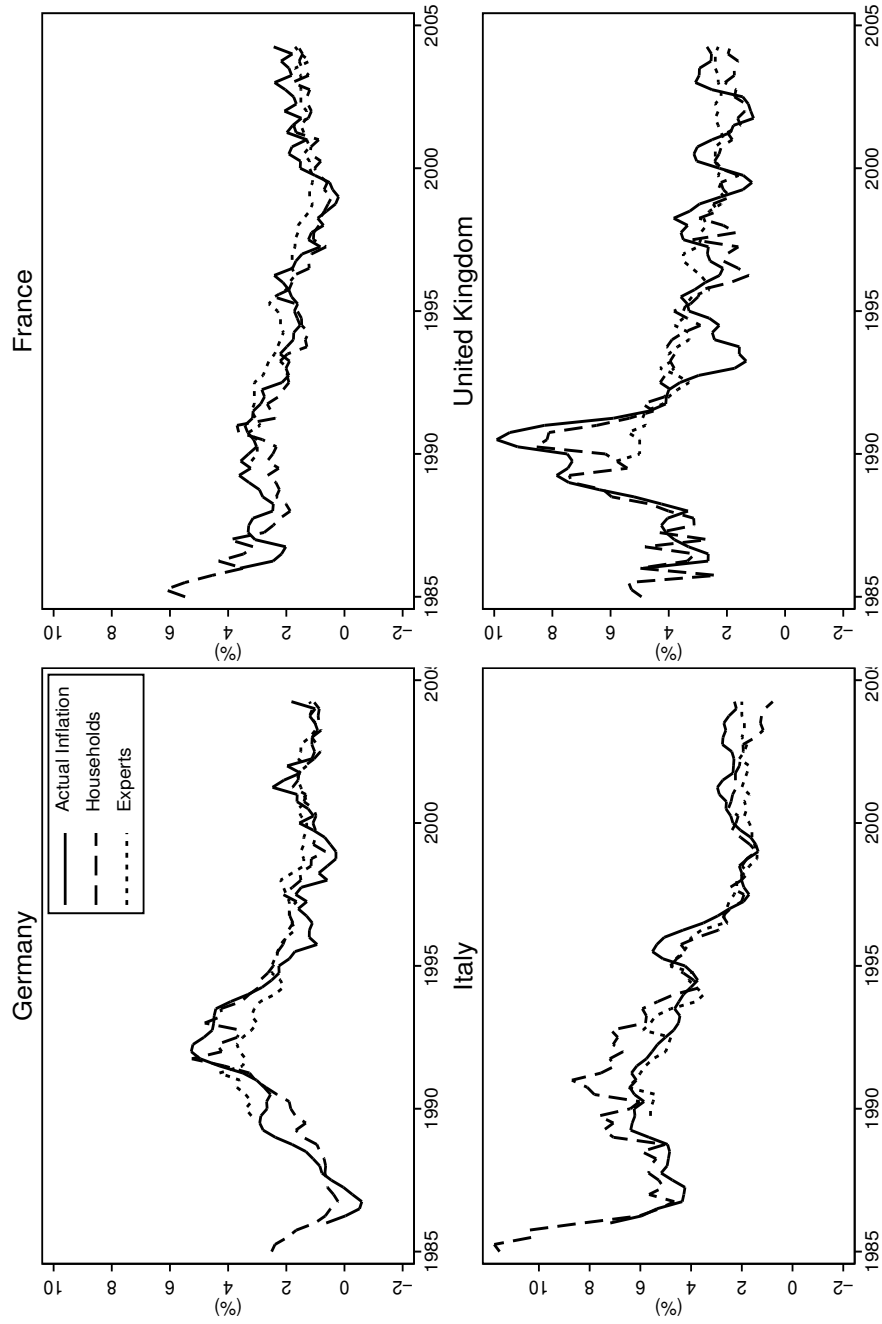
Figure 1 compares expert and household inflation expectations with actual inflation rates. Apparently, both expert and household predictions are roughly in line with actual inflation. However, sometimes there are rather persistent differences between expectations and actual inflation. More importantly, household and expert expectations differ quite considerably in certain time periods. Thus, a closer examination of the dynamic interaction of both variables is warranted.

4. EMPIRICAL RESULTS

The choice of the appropriate empirical strategy to estimate equation (1) depends on the time series properties of the underlying expectations. If the series are stationary, model (1) can be estimated directly using OLS (as in Carroll (2003)). If they are non-stationary (I(1)) and cointegrated, the model should be transformed into the vector error-correction (VEC) form.

⁵The details are available from the authors on request. The results reported in the paper typically hold for alternative normalizations considered. In particular, we find that the alternative rescaling procedures typically imply that λ_1 is positive and significant and the summing-up restriction $\lambda_1 + \lambda_2 = 1$ is met (using the same notation as below).

FIGURE 1. Household and Expert Expectations and Actual Inflation



A number of recent papers investigate the degree of persistence of various measures of inflation in Europe.⁶ Although these studies agree that inflation is a very persistent process, the evidence on the order of integration of inflation (i.e. whether it is stationary or $I(1)$) is less conclusive. Most papers cited in European Central Bank (2005) reject the null hypothesis that inflation in large European countries has recently been non-stationary. In contrast, the recent work of O'Reilly and Whelan (2005) (on inflation in the euro area) as well our preliminary tests (investigating inflation and inflation expectations in the EU-4 countries) in general do not reject the null.⁷ A potential criticism of our results is that the sample is too short to allow reliable inference. The fact that we are unable to reject the null may well result from the notoriously low power of the unit root tests under such circumstances, rather than the existence of the unit root.

Since the main focus of this paper is not on providing the definitive answer on the order of integration of inflation (or inflation expectations), we now move on to estimating Carroll's theoretical model and investigate how sensitive its implications are depending on whether we assume stationary or non-stationary environments. Because the tests do not clearly determine the stationarity properties in the relatively short sample we have, we first estimate the Carroll (2003) model in the stationary environment. We then consider how the results are affected if the nonstationary (VECM) set-up is adopted.

4.1. The Stationary Case: The Carroll Model. Before estimating equation (1), we will examine some preliminary evidence on the relationship between expert and household expectations. Given the interest in the interaction between the expectations of both professional forecasters and households, a natural starting point is to ask, (i) which of the two groups

⁶See European Central Bank (2005), table 3.4, page 21 for the summary of the literature on European countries.

⁷The results are available from the authors on request.

TABLE 1. Tests for Granger Non-causality

$$\mathbf{E}_t^i \pi_{t,t+4} = \beta_0 + \sum_{j=1}^p \beta_j \mathbf{E}_{t-j}^{EX} \pi_{t-j,t+4-j} + \sum_{k=1}^p \gamma_k \mathbf{E}_{t-k}^{HH} \pi_{t-k,t+4-k} + \varepsilon_{t+4}$$

Country	Dep. Var.: Expectations of ...	$\beta_j = 0, \forall j$ p value	$\gamma_k = 0, \forall k$ p value
Germany	Experts	0.000	0.125
	Households	0.000	0.000
France	Experts	0.000	0.076
	Households	0.000	0.000
Italy	Experts	0.000	0.010
	Households	0.620	0.000
United Kingdom	Experts	0.000	0.149
	Households	0.009	0.000

Sample: 1989:4 to 2004:2. The tests were computed with $p = 2$ lags of independent variables.

forecasts is on average better and (ii) what the causality is between the two expectations.

Relationship Between Expert and Household Expectations. First, expert expectations are substantially more precise than household expectations. The root mean squared errors of expert forecasts are between 15% to 35% lower in Germany, Italy and the UK than for household expectations. The two expectations are comparably precise in France.⁸ This does not, of course, come as a surprise since households may know expert forecasts when forming their own expectations. According to the epidemiology model, at least some households update their own expectations by following experts.

Second, we can examine whether expert forecasts Granger-cause household forecasts by testing for significance of the appropriate coefficients in the following equations:

$$\mathbf{E}_t^i \pi_{t,t+4} = \beta_0 + \sum_{j=1}^p \beta_j \mathbf{E}_{t-j}^{EX} \pi_{t-j,t+4-j} + \sum_{k=1}^p \gamma_k \mathbf{E}_{t-k}^{HH} \pi_{t-k,t+4-k} + \varepsilon_{t+4},$$

⁸ This is in line with the findings of Gerberding (2001), who investigates rationality, efficiency and unbiasedness of household and expert inflation in detail. She reports that while expert expectations have lower RMSEs (except for France), neither household nor expert expectations are fully efficient (and can be improved upon with available information on past inflation, unemployment and interest rates).

TABLE 2. Sticky Expectations in Europe I.
Restricted Cross-Country Results

$$\mathbf{E}_t^{HH} \pi_{t,t+4} = \lambda_0 + \lambda_1 \mathbf{E}_t^{EX} \pi_{t,t+4} + \lambda_2 \mathbf{E}_{t-1}^{HH} \pi_{t-1,t+3} + \lambda_3 \pi_{t-5,t-1} + \varepsilon_t$$

Model	λ_0	λ_1	λ_2	λ_3	Test	Cross eqn
					p value	p value
M1		0.17*** (0.04)	0.83*** (0.03)		$\lambda_1 + \lambda_2 = 1$ 0.912	0.04
M2		0.17*** (0.03)	0.83*** (0.03)		$\lambda_1 = 0.25$ 0.016	0.62
M3	-0.22*** (0.07)	0.29*** (0.05)	0.78*** (0.03)		$\lambda_0 = 0$ 0.003	0.15
M4		0.31*** (0.05)	0.77*** (0.05)	0.00 (0.04)	$\sum_{i=1}^3 \lambda_i = 1$ 0.003	0.03
M5	-0.22*** (0.07)	0.29*** (0.05)	0.78*** (0.05)	-0.01 (0.04)	$\lambda_3 = 0$ 0.900	0.13
M6			0.92*** (0.04)	0.05 (0.04)	$\lambda_2 + \lambda_3 = 1$ 0.015	0.34

Notes: Sample: 1989:4 to 2004:2. Seemingly unrelated regressions. Newey–West standard errors, 4 lags in brackets. *, **, *** denotes rejection of the null at the 10%, 5% and 1% significance level, respectively.

where regressions are estimated with both expert and household expectations on the right-hand side, $i \in \{EX, HH\}$. This is done in table 1. Columns 3 and 4 indicate that lags of expert expectations are typically significant predictors of household expectations. Household expectations, on the other hand, tend not to Granger-cause experts. Thus, in all countries, except for Italy we conclude that the direction of causality goes from experts toward households.

What Determines Household Expectations? Having found supportive preliminary evidence for the epidemiological model of expectation formation, let us now turn to direct estimation of and inference about the speed of information updating, λ . Table 2 summarizes the estimation results of the following regressions:

$$\mathbf{E}_t^{HH} \pi_{t,t+4} = \lambda_0 + \lambda_1 \mathbf{E}_t^{EX} \pi_{t,t+4} + \lambda_2 \mathbf{E}_{t-1}^{HH} \pi_{t-1,t+3} + \lambda_3 \pi_{t-5,t-1} + \varepsilon_t \quad (2)$$

in various forms. All regressions are estimated with seemingly unrelated regressions (SUR) with coefficients restricted constant across the four countries.⁹

The format of table 2 follows that of Carroll (2003), table III. The left panel (the first four columns) displays the point estimates of λ s together with standard errors; the right panel shows specification tests (p values of various tests of coefficients). The last column (“Cross eqn p value”) tests whether the coefficients are the same across countries.¹⁰

The models are labelled M1–M6. The first model, M1, estimates the following version of (2):

$$\mathbf{E}_t^{HH} \pi_{t,t+4} = \lambda_1 \mathbf{E}_t^{EX} \pi_{t,t+4} + \lambda_2 \mathbf{E}_{t-1}^{HH} \pi_{t-1,t+3} + \varepsilon_t \quad (3)$$

in which coefficients λ_1 and λ_2 are estimated as unrestricted. The estimates of λ_1 and λ_2 are 0.17 and 0.83, respectively. The summing-up restriction implied by the Carroll model, $\lambda_1 + \lambda_2 = 1$, is clearly satisfied.

Model M2 is estimated for the restricted version with the summing-up restriction imposed. Given how close the restriction is to being met in M1, it does not come as a surprise that the point estimates of λ barely change.

Our baseline $\lambda_1 = 0.17$ suggests that the average European household reads economic updates or consults economic experts roughly once in 18 months. In addition, it implies about 47% of households use information outdated more than one year and about 23% more than two years.

⁹Analogous results hold when the models are estimated with equation-by-equation OLS, however, since the cross-correlation between residuals in our dataset is up to 0.3, SUR improve efficiency of our estimates (upon the simple equation-by-equation estimation). In addition, SUR also make it possible to test cross-equation restrictions and answer questions such as: “Does the speed of information updating vary across countries?”

We report some results unrestricted across countries below in table 3. More detailed results of unrestricted SUR and equation-by-equation estimation are available from the authors on request. The results are robust to these modifications.

¹⁰To conserve space we do not report measures of fit, which of course differ slightly for each country (and model). \bar{R}^2 s vary between 0.75 and 0.96. For more information of how well selected models explain household inflation expectations see table 3 below. Also, similarly to Carroll’s findings, the Durbin–Watson statistics detect virtually no autocorrelation in residuals for models M1–M6.

The speed of adjustment $\lambda_1 = 0.17$ is lower than Carroll's baseline coefficient of 0.27. Because the standard error of λ_1 is small, the difference is statistically significant. However, much of the difference between λ in Europe and the US can be accounted for with different time ranges: Carroll's sample (1981:3–2002:1) differs from ours (1989:4–2004:2). Re-estimating model M2 with the US data and our sample range gives $\lambda = 0.16$ (for the US). This matches Carroll's evidence that updating is faster when inflation is in the news, including the early period in his sample. In contrast, in the 1990s λ has fallen, because the number of newspaper articles (in the US) was substantially lower than before 1985. In addition, the recent monetary policy leading to low and stable inflation presumably reduces uncertainty and, together with smaller exogenous shocks hitting the economy, decreases households' incentive to update.

Models M3–M6 investigate four alternative structures of household expectations. First, we add a constant to equation (3). This turns out to be significantly different from zero. As advocated by Carroll (2003), it is, however, doubtful a priori that one can have a reasonable structural specification of inflation expectations with non-zero constant term, since this would imply that households' predictions are permanently biased away from experts'. Interestingly, the estimate of constant (λ_0) in M3 is negative. One reason for that may be, as is apparent from figure 1, that over our estimation sample actual inflation rates were actually falling. In such environment some households may have extrapolated this downward trend into the future, which is reflected in the negative values of the constant term.

Models M4–M6 allow for the possibility that consumers are at least in part backward-looking (adaptive) by adding past inflation to the right-hand side of (2). Similarly to the US, there is very little of the backward-looking element in household inflation expectations: the coefficient λ_3 is small both in terms of its size and its level of significance. Thus, households seem to learn from experts rather than naively extrapolate the past inflation rates.

Generally, there appears to be a lot of homogeneity across countries. As indicated in the last column, in four of the six (M1–M6) models considered the null of constant coefficients across countries is satisfied, two models (M1 and M4) yield borderline rejections of homogeneity (at the 5% significance level).

Cross-Country Differences. Having found supportive evidence for the Carroll’s sticky information model in European data, let us now investigate in more detail how the findings vary across countries. Table 3 summarizes estimation results obtained from seemingly unrelated regressions, unrestricted across countries, for models M1–M3.

The findings parallel those in the previous section: First, the speed of updating λ_1 varies between 0.11 and 0.32 (as estimated with models M1 and M2). For all countries, except Italy, it is highly statistically significant. Second, for all countries, except France, the summing-up restriction, $\lambda_1 + \lambda_2 = 1$, implied by the Carroll’s sticky information model, is met. Even for France, the two coefficients sum up to 0.91. Third, the intercept term λ_0 is insignificant for all countries except the UK.

We could now in principle similarly test how stable λ_1 has been over time. However, due to the limited number of observations the tests of structural stability have weak power. We looked for statistically significant differences in the speed on expectation updating by including dummy variables for, e.g., members of the euro area and inflation-targeters. Unfortunately, these investigations did not lead to any statistically clear-cut results about the determinants of λ_1 .

One concern about our results is that the estimates are biased due to the measurement error in household inflation expectations. In presence of such measurement error coefficients λ_2 for model M1 in table 3 are biased toward zero. To get a feel for how serious this bias is we apply four alternative methods to transform the qualitative survey questions to the quantitative responses. For each period we then calculate the cross-sectional variance of

TABLE 3. Sticky Expectations in Europe II.
Country-by-Country Unrestricted Results

$$\mathbf{E}_t^{HH} \pi_{t,t+4} = \lambda_0 + \lambda_1 \mathbf{E}_t^{EX} \pi_{t,t+4} + \lambda_2 \mathbf{E}_{t-1}^{HH} \pi_{t-1,t+3} + \varepsilon_t$$

Model	λ_0	λ_1	λ_2	\bar{R}^2	Test
					p value
Germany					
M1		0.18*** (0.06)	0.82*** (0.06)	0.91	$\lambda_1 + \lambda_2 = 1$ 0.764
M2		0.20*** (0.06)	0.80*** (0.06)	0.91	$\lambda_1 = 0.25$ 0.368
M3	-0.21* (0.12)	0.29*** (0.08)	0.80*** (0.06)	0.91	
France					
M1		0.32*** (0.08)	0.59*** (0.09)	0.85	$\lambda_1 + \lambda_2 = 1$ 0.002
M2		0.18*** (0.07)	0.82*** (0.07)	0.83	$\lambda_1 = 0.25$ 0.322
M3	-0.04 (0.12)	0.33*** (0.10)	0.61*** (0.09)	0.80	
Italy					
M1		0.14 (0.11)	0.86*** (0.09)	0.96	$\lambda_1 + \lambda_2 = 1$ 0.991
M2		0.11* (0.06)	0.89*** (0.06)	0.96	$\lambda_1 = 0.25$ 0.022
M3	-0.18 (0.15)	0.25* (0.13)	0.81*** (0.09)	0.95	
United Kingdom					
M1		0.23*** (0.08)	0.77*** (0.08)	0.89	$\lambda_1 + \lambda_2 = 1$ 0.763
M2		0.23*** (0.08)	0.77*** (0.08)	0.89	$\lambda_1 = 0.25$ 0.781
M3	-0.67** (0.30)	0.53*** (0.16)	0.66*** (0.09)	0.89	

Notes: Sample: 1989:4 to 2004:2. Seemingly unrelated regressions. Newey–West standard errors, 4 lags in brackets. *, **, *** denotes rejection of the null at the 10%, 5% and 1% significance level, respectively.

the four series and average it over time, which gives us an estimate of the variance of the measurement in household expectations. Our calculations imply that the variance of measurement error for Germany, France and Italy is 8.5–9 times and for the UK 5.6 times smaller than the variation in

household expectations. These numbers imply that the measurement error bias causes coefficients λ_2 for model M1 in table 3 to be underestimated by 10-11% for Germany, France and Italy and by 15% for the UK.¹¹

Most of our findings suggest that the epidemiology model of information diffusion performs similarly well, qualitatively as well as quantitatively, for the core European countries as it does for the US. Expert inflation expectations are typically more precise than household expectations. Econometric tests indicate that the Carroll model is adequate along several dimensions.¹² Several models imply that European households update somewhat more slowly than US households, on average once in 18 months compared with once a year. Finally, there is strong evidence that, as suggested by the epidemiology model, European households update information from the professional forecasters rather than the past inflation rate.¹³

4.2. The Carroll Model in Vector Error-Correction Form. Having estimated the epidemiology model in a stationary framework, let us now examine how the implications change when we assume that the expectation series are I(1) instead. Suppose we collect expert and household expectations in a vector $x_t = (\mathbf{E}_t^{HH} \pi_{t,t+4}, \mathbf{E}_t^{EX} \pi_{t,t+4})^\top$. If the two series are cointegrated with cointegrating vector $\alpha = (1, -\alpha_1)^\top$, the system has the following vector error correction (VEC) representation:

$$\Delta x_t = \lambda \alpha^\top x_{t-1} + \beta(L) \Delta x_t + \varepsilon_t, \quad (4)$$

where $\lambda = (\lambda_{HH}, \lambda_{EX})^\top$ denotes the vector of loading coefficient and $\beta(L)$ is a matrix lag polynomial. Similarly to the stationary model (1), λ determines

¹¹Measurement error does not affect asymptotic properties of the VEC estimates below, but may lead to a bias in small samples quantitatively similar to that for the stationary specification.

¹²For example, the speed of updating is positive and statistically significant, the summing-up restriction holds fairly well and household inflation expectations are not sensitive with respect to the past inflation.

¹³Consideration might be given to the possibility that households update their expectations by referring directly to other publicly available information, such as foreign prices. However, in the epidemiology framework this information is already captured and processed by professional forecasters, who are assumed to be rational. Moreover, obtaining such information is presumably much more costly than simply referring to the published professional forecasts.

the speed of adjustment toward the (long-run) equilibrium.¹⁴ We are particularly interested in λ_{HH} , which corresponds to the speed of adjustment observed for households. Furthermore, note that the theoretical derivation of the “epidemiology model” predicts a cointegrating vector $\alpha = (1, -1)^\top$. This is due to the fact that in the long-run households completely adapt to the professional forecasts.

Before estimating the VEC representation (4) and its “ α -restricted” counterpart some preliminary specification tests need to be done. First, we test whether there exists a valid cointegrating relationship between the expert and household expectations as shown in table 4. The findings show that, for all four countries, the two series are cointegrated (at the 5% significance level). In addition, we checked whether the theoretical restriction on $\alpha = (1, -1)^\top$ is supported in data. The values for α_1 are close to -1 (the value predicted by the model) and range from -1.21 for the UK to -1.00 for Germany. As illustrated by the likelihood ratio statistics presented in table 5, we find that α is not significantly different from $(1, -1)^\top$ (except in the UK).

The VEC findings are summarized in table 5.¹⁵ All estimates of λ_{HH} are significant (although for the restricted case in France and Italy only at the 10% level) and lie except for the restricted case in Italy in the neighborhood of 0.25, typically somewhat higher than implied by the “stationary” results above. We again find a lot of homogeneity among the four countries with French and Italian households updating presumably somewhat slower than British and German ones. The estimated updating frequencies in table 5 lie between once in three and seven quarters.

¹⁴The adjustment pattern in the partial adjustment version of the model (1), however, differs from the VEC analysis in two ways: First, the adjustment in the VEC is analyzed in an interdependent system and feedback effects are considered and second, the short-run dynamics in the VEC might influence the dynamic adjustment path.

¹⁵The models were estimated for the time frame between 1989:4 and 2004:2, except for Italy, where a valid cointegrating relationship was found between 1992:4 and 2002:4.

TABLE 4. Tests for Cointegration between Household and Expert Expectations

Unrestricted Cointegration Rank Test (Trace)				
Hypothesized		Trace	5%	
No. of CE(s)	Eigenvalue	Statistic	Critical Value	Prob.*
Germany				
None	0.20	16.31	12.32	0.01
At most 1	0.06	3.48	4.13	0.07
France				
None	0.22	15.92	12.32	0.01
At most 1	0.03	1.96	4.13	0.19
Italy				
None	0.36	21.69	12.32	0.00
At most 1	0.09	3.65	4.13	0.07
United Kingdom				
None	0.25	18.90	12.32	0.00
At most 1	0.05	2.60	4.13	0.13

Note: * MacKinnon–Haug–Michelis (1999) p-values. Sample: 1989:4 to 2004:2, Italy: 1992:4–2002:4.

TABLE 5. Sticky Expectations in the VEC Form

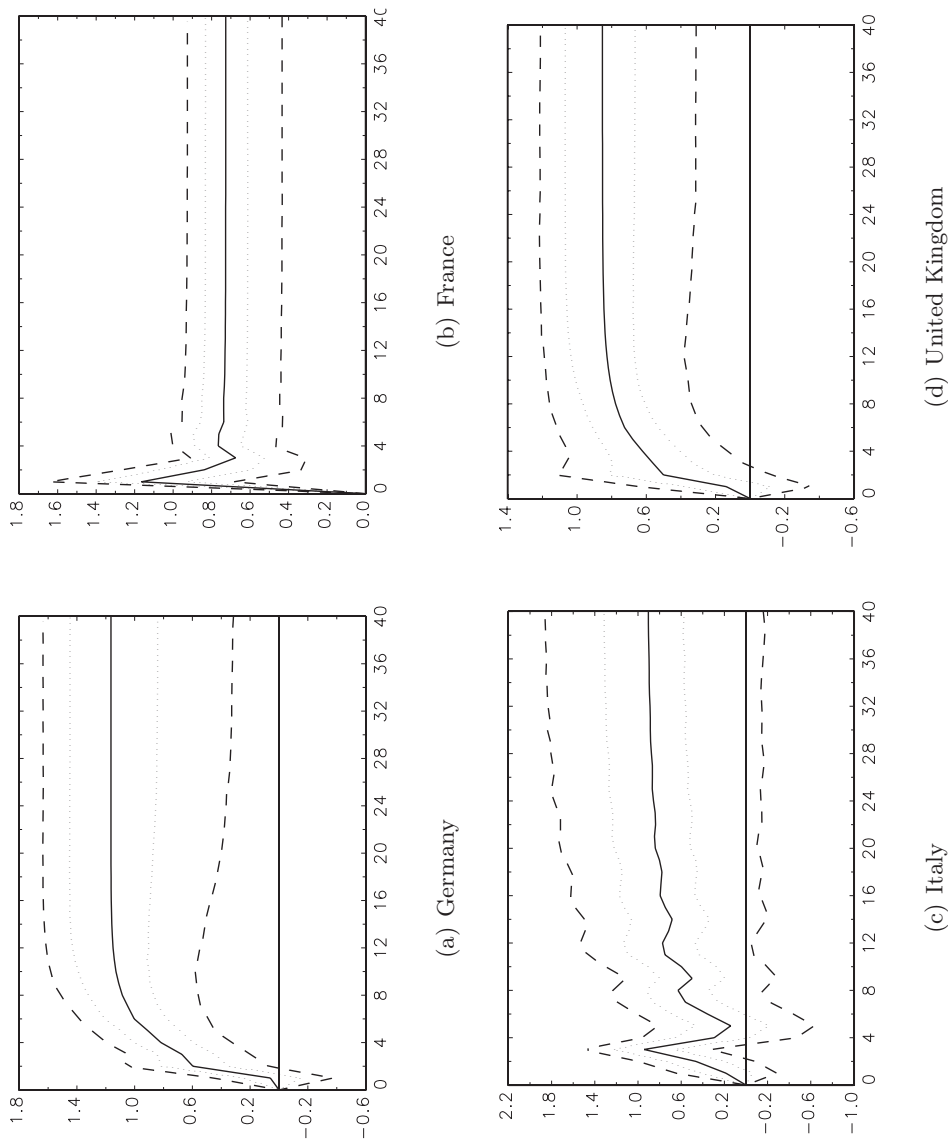
Model		Germany	France	Italy	UK
Unrestricted	$\hat{\lambda}_{HH}$	-0.30***	-0.26***	-0.18**	-0.20***
	std. error	(0.09)	(0.09)	(0.07)	(0.05)
Restricted	$\hat{\lambda}_{HH}$	-0.30***	-0.14*	-0.22*	-0.27***
	std. error	(0.09)	(0.08)	(0.12)	(0.08)
Test for restriction (1, -1) on α	LR stat.	0.00	2.29*	2.97*	3.86**
	p value	0.988	0.070	0.085	0.049

Notes: Sample: 1989:4 to 2004:2, Italy: 1992:4–2002:4. “Unrestricted” refers to the unrestricted VEC model. “Restricted” refers to the VEC estimation results under the restriction $\alpha = (1, -1)^\top$. *, **, *** denotes rejection of the null at the 10%, 5% and 1% significance level, respectively.

The Carroll’s sticky information model is also supported by how the deviations from the long-term equilibrium are corrected.¹⁶ The error-correction process is primarily driven by the adjustment in household rather than expert expectations. This is implied by the estimates of $\hat{\lambda}_{EX}$, which, except for France, are not significantly different from zero. Finally, this confirms the earlier finding that expert forecasts Granger-cause household expectations, whereas household forecasts do not tend to Granger-cause the forecasts of experts.

¹⁶Detailed results are available from the authors on request.

FIGURE 2. Impulse Response of Household Expectations to a Shock to Expert Expectations



Notes: Solid line: impulse responses, dotted line: 68% confidence bound, dashed line: 95% confidence bound. Impulse responses and confidence bounds were calculated using the program JMulTi and the embedded bootstrap procedure with 3000 repetitions. See Lütkepohl and Krätzig (2004) for a detailed description.

TABLE 6. Forecast Error Variance Decomposition of Household Expectations. Proportions of Forecast Error at Different Horizons Accounted for by ...

Quarters	$\mathbf{E}^{EX} \pi$	$\mathbf{E}^{HH} \pi$
Germany		
1	0.02	0.98
4	0.11	0.89
8	0.41	0.59
20	0.78	0.22
40	0.89	0.11
France		
1	0.09	0.91
4	0.45	0.55
8	0.54	0.46
20	0.64	0.36
40	0.70	0.30
Italy		
1	0.39	0.61
4	0.67	0.33
8	0.63	0.37
20	0.78	0.22
40	0.86	0.14
United Kingdom		
1	0.00	1.00
4	0.08	0.92
8	0.25	0.75
20	0.58	0.42
40	0.74	0.26

Figure 2 shows the impulse response functions of household expectations following a one-unit shock in expert expectations together with the 68% and 95% confidence bounds). The dynamic adjustment paths clearly differ among countries. While the results for Germany and United Kingdom point to a smooth and stable adjustment process, those in Italy and France, in contrast, suggest some initial “overshooting”. In all cases, considering the confidence bounds, however, household expectations eventually stabilize around the expected value of 1.

Table 6 displays the forecast error variance decomposition of household expectations into the parts driven by expert predictions and household expectations. In all cases expert forecasts dominate household expectations

clearly in the long-run but not in the short run. Only after five years, the bulk of the forecast error variance in household expectations—58% to 78%—is dominated by expert expectation shocks. This again confirms our general finding: households are relatively sticky in information updating.

Our findings thus imply that the epidemiology model of Carroll (2003) can be easily extended to the “non-stationary world.” The derived VEC epidemiology model of information diffusion performs similarly well to the stationary model. This result is especially useful for the analysis of European countries, which plausibly have highly persistent inflation rates (see O’Reilly and Whelan (2005) and references in European Central Bank (2005)). Thus, even though it seems to be difficult to draw clear conclusions about the stationarity properties of the series with the small sample size at hand, the VEC representation might be preferable once more data are available.¹⁷

5. CONCLUSION

Inflation expectations are crucial determinants of future inflation dynamics. The model estimated here attempts to analyze how these expectations are formed and how information is transmitted from professional forecasters to households. Our estimates of the speed of information updating have important implications for the persistence of inflation and inflation expectations. We document that the qualitative and quantitative findings previously reported for the US generalize to major European countries. Most European households adjust rather sluggishly to new information; they update their information on average once in twelve to eighteen months. Interestingly, it turns out that households are forward-looking in that they use information processed by experts rather than just past information. These findings are robust to a number of estimation methods (suited for data with various stochastic properties) we consider.

¹⁷This indeterminacy is *ex-post* justified by the similarities between the results from the Carroll model and the results of the VEC models.

The research in this paper can be extended through a number of avenues. Survey data could be used to directly estimate the sticky-information Phillips curve in addition to its epidemiological micro-foundations. Alternatively, it would be possible, in the spirit of Mankiw et al. (2003), to analyze the evolution of cross-section distribution of inflation expectations in Europe rather than just their mean values. Finally, the epidemiology model could, in principle, be estimated for additional countries, using cross-sectional dependence among countries to alleviate problems related to short samples.

REFERENCES

- Batchelor, Ray A., and Adrian B. Orr (1988), "Inflation Expectations Revisited," *Economica*, 55(219), 317–331.
- Branch, William A. (2004), "The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations," *Economic Journal*, 114(497), 592–621.
- Carlson, John A., and Michael Parkin (1975), "Inflation Expectations," *Economica*, 42(166), 123–137.
- Carroll, Christopher D. (2003), "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 118(1), 269–298.
- Ehrmann, Michael, and Marcel Fratzscher (2005), "Communication and Decision-making by Central Bank Committees: Different Strategies, same Effectiveness?" European Central Bank Working Paper 488, European Central Bank.
- European Central Bank (2005), "Inflation Persistence and Price Setting Behaviour in the Euro Area," report, European Central Bank, available at <http://www.ecb.int/home/pdf/research/inflationpersistencepricesettingreport.pdf>.
- Gerberding, Christina (2001), "The Information Content of Survey Data on Expected Price Developments for Monetary Policy," Deutsche Bundesbank discussion paper 9/2001, Deutsche Bundesbank.
- Handjiyska, Boriana (2004), "Adjustment of Household Inflation Expectations in Five OECD Countries," mimeo, Johns Hopkins University.
- Khan, Hashmat, and Zhenhua Zhu (2002), "Estimates of the Sticky Information Phillips Curve for the United States, Canada, and the United Kingdom," Bank of Canada working paper 19, Bank of Canada.
- Lütkepohl, Helmut, and Markus Krätzig, editors (2004), *Applied Time Series Econometrics*, Cambridge University Press, Cambridge.
- Mankiw, N. Gregory, and Ricardo Reis (2002), "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117, 1295–1328.
- Mankiw, N. Gregory, and Ricardo Reis (2006), "Pervasive Stickiness," *American Economic Review Papers and Proceedings*, forthcoming.
- Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers (2003), "Disagreement on Inflation Expectations," *NBER Macroeconomics Annual*, 209–248.

- Nielsen, Hannah (2003), *Essays on Expectations*, Shaker, Aachen.
- O'Reilly, Gerard, and Karl Whelan (2005), "Has Euro-Area Inflation Persistence Changed over Time?" *Review of Economics and Statistics*, 87(4), 709–720.
- Sims, Christopher (2003), "Implications of Rational Inattention," *Journal of Monetary Economics*, 50(3), 665–690.

APPENDIX

Expert Forecasts. The professional forecasts were obtained from Consensus Economics, a London-based macroeconomic survey firm. The survey of experts of private and public institutions in major industrial countries has been collected monthly since 1989. Once every quarter the questionnaire contains a question on forecasts over the next six quarters. The consensus forecast, used in the paper as a measure of expert expectations, is the mean of about 20 to 30 forecasts of local experts from major banks or research institutes in each country.

Household Forecasts. Our measures of household inflation forecasts are based on disaggregated answers to question 6 from European Commission's Harmonised Business and Consumer Surveys. The sample size of the survey is about 2,000 households in Germany, Italy and the UK, and roughly 3,300 households in France. The data are available monthly since 1985.

Extracting Household Inflation Expectations. We first have to re-scale the balance statistics to obtain a measure of inflation expectations. The standard method follows Carlson and Parkin (1975) and its extensions (see, for example, Gerberding, 2001, Mankiw et al., 2003 and Nielsen, 2003). The observed data are from the pentachotomous survey. Consequently, they classify the responses into five subgroups:

Consumer prices will:

- Increase more rapidly,
- Increase at the same rate,
- Increase at a slower rate,
- Stay about the same,
- Fall.

Batchelor and Orr (1988) derive how responses from a pentachotomous survey can be transformed into a measure of inflation expectations ${}_t\mu_{t+1} = \tilde{\mu}_t \times f({}_tA_{t+1}, \dots, {}_tE_{t+1})$, where ${}_tA_{t+1}, \dots, {}_tE_{t+1}$ are the fractions of respondents answering each option and f is a known function (see Batchelor and Orr, 1988, p. 322, formula (11)).

We experimented with three choices of $\tilde{\mu}_t$: (i) recursive Hodrick–Prescott filter, (ii) past inflation, (iii) inflation extracted as fitted values from the regression of inflation on the past balance statistics from the survey. Most results reported in the paper hold for all normalizations considered. In particular, we find that the alternative rescaling procedures typically imply that λ_1 is positive and significant and the summing-up restriction $\lambda_1 + \lambda_2 = 1$ is met (using the same notation as below). One result that does not hold for normalizations (ii) and (iii) is that

household expectations are insensitive to past inflation. This is not surprising because for these two normalizations past inflation rate is a direct input into $\tilde{\mu}_t$.