Anchoring Inflation Expectations in Unconventional Times: Micro Evidence for the Euro Area*

Jonas Dovern[a] and Geoff Kenny[b]
[a]FAU Erlangen-Nürnberg and CESifo
[b]European Central Bank

We exploit micro data from professional forecasters to examine the stability of the distribution of long-term inflation expectations in the euro area following the Great Recession. Although mean expectations declined somewhat, we find no evidence that the central tendency of the long-run distribution became unanchored. Also, the degree of co-movement of expectations with other variables did not increase noticeably. In contrast, long-term inflation uncertainty increased and expectations became negatively skewed. Such findings are in line with the predictions of theoretical models emphasizing the impact of the lower bound on policy rates and uncertainty about the transmission of unconventional monetary policies.

JEL Codes: E31, E58.

1. Introduction

A tight anchoring of medium- to long-term inflation expectations around the central bank's target is commonly seen as crucial for

*We would like to thank participants at the May 17–18, 2018 Federal Reserve Bank of Cleveland conference "Inflation: Drivers and Dynamics," and the September 2017 ECB conference "Understanding Inflation: Lessons from the Past, Lessons for the Future" for helpful comments and suggestions, as well as participants at seminars at the ECB, Heidelberg University, the Freie Universität Berlin, LMU Munich, and the 10th International Conference on Computational and Financial Econometrics. We also thank the editor and two anonymous referees for very helpful comments on earlier drafts of the paper. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the ECB or the Eurosystem. Corresponding author (Dovern): Lange Gasse 20, 90403 Nuremberg, Germany. Phone: +49-911-5302-290. E-mail: jonas.dovern@fau.de.

steering inflation toward this target without suffering substantial economic costs. However, during recent years, large macroeconomic and financial shocks associated with the Great Recession and the fact that policy rates reached their effective lower bound (ELB) have led to concerns about a possible de-anchoring of long-term inflation expectations in the major currency areas. In the case of the euro area, concerns have focused on persistently too-low inflation or even deflationary risks and an associated departure of inflation expectations from levels consistent with the objective of the European Central Bank (ECB). For example, Draghi (2014) highlights "the risk that a too prolonged period of low inflation becomes embedded in inflation expectations." Indeed, recent unconventional monetary policies are often motivated as addressing such risks.

In this paper we examine empirically this central policy question: How well anchored did euro-area inflation expectations remain in the wake of the Great Recession and euro-area sovereign debt crisis? By the anchoring of long-term inflation expectations, we refer to their stability over time and their consistency with the objectives of a central bank. In principle, a necessary and sufficient condition for long-term expectations to be perfectly anchored is that they are constant and equal to the central bank's target. However, a direct test of these conditions is complicated by uncertainty about what constitutes "well" in practice and also about the precise target and/or the horizon over which the central bank aims to achieve it.[1] A situation in which expectations become unanchored would imply a significant and substantial movement in the level of long-term inflation expectations away from the level implied by the central bank's objective. Clearly, the concept of anchoring is also a matter of *degree* in the sense that it is insightful to consider how tightly or firmly expectations may be anchored at a given level. For example, even in a situation where the level of inflation expectations remains aligned with the central bank's objective, economic agents may become less confident about this outcome and may therefore attach a lower probability or likelihood that this objective will be

[1]For example, in the case of the euro area, while price stability is clearly defined to be an increase in the Harmonised Index of Consumer Prices that is below, but close to, 2.0 percent, the ECB indicates that it aims to achieve this "over the medium term," which retains some degree of vagueness.

achieved and, at the same time, may attach higher probabilities to more extreme inflation outcomes. Such an analysis of changes in the degree to which expectations are anchored requires information extracted from the probability distribution surrounding long-term inflation expectations.[2]

Much of the recent economic literature attempting to quantify the evidence and risks of such a de-anchoring—both in the euro area and elsewhere—has focused only on the *mean* or first moment of the distribution of long-term inflation expectations (Demertzis, Marcellino, and Viegi 2009, 2012; Gürkaynak, Levin, and Swanson 2010; Beechey, Johannsen, and Levin 2011; van der Cruijsen and Demertzis 2011; Dräger and Lamla 2013; Mehrotra and Yetman 2014).[3] In this paper, we provide new micro evidence about the anchoring of inflation expectations in the euro area considering the *full subjective forecast distribution*. Strong theoretical arguments justify the need to study the properties of the full distribution and not simply focus on mean expectations. In particular, shifts in the variance of this distribution, its skewness, or tail risk can offer additional evidence of change in agents' beliefs about future inflation over the longer term and the factors that may be shaping them. For example, the model of imperfect credibility in Bodenstein, Hebden, and Nunes (2012) suggests that achieving the central bank's inflation objective may become more challenging following a period such as the Great

[2]Mehrotra and Yetman (2014) highlight the importance of the full probability distribution, noting that "there are at least two dimensions to anchoring: both the level at which expectations are anchored ... and how tightly expectations are anchored at that level."

[3]In particular, such studies have focused on possible changes in the mean or in the strength of its co-movement with other economic variables. Most evidence emerging from this literature suggests that long-term inflation expectations were affected by the Great Recession. For the U.S. economy, Galati, Poelhekke, and Zhou (2011), Autrup and Grothe (2014), Ciccarelli and García (2015), and Nautz and Strohsal (2015) all suggest that inflation expectations in the United States started to react more strongly to macroeconomic news. Ehrmann (2015) also reports evidence of a similar increased sensitivity during periods of low inflation. For the euro area, Galati, Poelhekke, and Zhou (2011) identify a structural break in the responsiveness of European inflation expectations to macroeconomic news and Lyziac and Paloviita (2017) conclude that there are "some signs of de-anchoring." However, Autrup and Grothe (2014), Strohsal and Winkelmann (2015), and Speck (2016) have argued that the degree of anchoring did not change around that time.

Recession, as central bank credibility becomes more relevant than in "normal" times. In a similar vein, Beechey, Johannsen, and Levin (2011) demonstrate how imperfect information and potential time variation in the central bank's objective can be associated with a sizable increase in long-term inflation uncertainty as measured by the variance of the distribution for long-term expected inflation.[4] However, it is the incidence of the ELB that makes the strongest case for studying the full distribution. Under the ELB, models of the business cycle exhibit multiple equilibriums, implying that the distribution of long-term inflation expectations may change and attach non-negligible probabilities to quite distinctive outcomes.[5] Clearly, the risks of such bad equilibriums will tend to first show up in increased long-term inflation uncertainty or tail risks, i.e., in the second and fourth moments of the distribution. In addition, models which take into account the ELB also emphasize that limitations in the central bank's ability to respond to deflationary shocks lead to a negatively skewed distribution for long-term expectations (see, for example, Coenen and Warne 2014; Hills, Nakata, and Schmidt 2016).

To address our central question, we proceed in three steps. In a first, we test for possible structural change in the distribution of the long-term subjective forecast distribution, exploiting the methods of Andrews and Ploberger (1994), Bai and Perron (1998, 2003),

---

[4]Regarding possible changes in the central bank's objective, there has been considerable discussion in the wake of the financial crisis about the need for central banks to consider adjusting upward their inflation objectives, and the euro area has not been immune to these discussions. For example, Ball et al. (2016) recently make this recommendation as a means to avoid the incidence, severity, and costs of hitting the ELB constraint. This discussion is conceptually distinct from other recommendations which have emphasized increasing short-run inflation expectations as a demand-management device at the ELB.

[5]For instance, Benhabib, Schmitt-Grohé, and Uribe (2001) highlight the existence of a deflationary equilibrium where the ELB is binding and inflation is stuck below target. More recently, Aruoba and Schorfheide (2015) construct a two-regime stochastic general equilibrium model in which the economy may alternate between a "targeted inflation regime" and a "deflation regime." Busetti et al. (2014) also study the risks of such a regime in a model with learning and show that it may imply considerable risks of a de-anchoring of long-term inflation expectations and may give rise to a period of sustained low real output growth.

and Inoue (2001) to shed light on possible shifts, their magnitude, and timing. This provides direct evidence about whether or not the Great Recession and its aftermath resulted in any significant changes in key features of that distribution.

In a second step, we exploit the available micro data in a panel setting and explain changes in long-term inflation expectations by studying their co-movement with other macroeconomic variables. Our use of micro data contrasts with most other recent studies mentioned above which have focused on average measures of expectations or representative proxies extracted from asset prices. Although we do not claim a causal interpretation, the analysis of such co-movements sheds light on possible changes in the degree to which the distribution is anchored. For example, allowing for uncertainty about the central bank's objective, and learning on the part of private agents about its ability to hit that objective, we can expect some positive co-movement between short-term macroeconomic news and long-term inflation expectations (e.g., as in Beechey, Johannsen, and Levin 2011). In the spirit of Levin, Natalucci, and Piger (2004), we also consider the co-movement of long-term expectations with an ex post measure of central bank performance to assess whether agents partly update their future long-term expectations by taking into account the rate of inflation that the central bank has actually delivered. In line with recent discussions about secular stagnation and deflationary equilibriums with simultaneously weak trend growth and excessively low inflation expectations, e.g., as discussed in Eggertson and Mehrotra (2014) and Summers (2014), we also examine the co-movement between long-term inflation expectations and corresponding long-term expectations about the real economy. Using matched individual-level expectations, our panel data allow us to study these interrelationships while also controlling for other sources of variation that are common to all forecasters, such as observed inflation rates. In addition, we provide direct tests of whether these interrelationships have changed following the Great Recession and during the more recent period when the ELB has been binding and the ECB has employed nonstandard monetary policy.

In a third and final step, using a similar set of appropriately transformed covariates, we extend the above micro-level analysis to shed light on the factors which co-move with changes in long-term inflation uncertainty which we measure by the individual variance of the subjective forecast distributions.[6] Again, such an analysis speaks directly to the question of anchoring because higher uncertainty about long-term inflation prospects implies that the distribution is less tightly anchored even if overall mean expectations remain unchanged.

The layout of the remainder of the paper is as follows. In section 2, we briefly describe our data set and the manner in which we estimate key moments of the subjective long-term forecast distributions. Section 3 presents the empirical evidence for changes in the distribution. In section 4, we turn to the analysis of the variables which co-move with changes in long-term mean expectations and uncertainty and analyze whether the strength of this co-movement has changed following the Great Recession. Section 5 concludes by discussing the economic significance of our findings and their implications for monetary policy.

2. Data and Estimation of Moments

We base our analysis on the individual density forecasts provided by the ECB's Survey of Professional Forecasters (SPF), which has been conducted and published by the ECB at a quarterly frequency since the beginning of 1999. Since we are interested in long-term inflation expectations, we primarily focus on those forecasts that have a forecast horizon of roughly five years (i.e., 20 quarters or $h = 20$). Our sample period ranges from 1999:Q1 to 2017:Q1.[7] Online appendix A (available at http://www.ijcb.org) provides detailed information about the SPF histograms and other data sources used in the study. Reflecting its survey origins, the SPF data set is heavily unbalanced because forecasters leave the panel or enter later and because they

[6] In the remainder of this paper we will use the two terms "variance" and "uncertainty" interchangeably. For an overview of other definitions and measures of uncertainty see, for instance, Rossi, Sekhposyan, and Soupre (2016).

[7] Note that during the first two years of the SPF long-term forecasts were only surveyed on an annual basis in 1999:Q1 and 2000:Q1. We make use of these observations whenever possible. However, we have to drop them from our econometric estimation whenever we relate long-term expectations to lagged information from the SPF.

are not required to report their forecasts in every survey round. The number of individual respondents therefore varies from one quarter to the next. We focus on the analysis of this unbalanced panel and, in particular, do not attempt to interpolate any missing observations. The number of respondents providing the histograms depicting the probabilities they assign to a range of future inflation outcomes averages 39.1 in our sample. Although this is slightly lower than the number of respondents who report point forecasts for inflation (44.6 on average), it nonetheless provides a rich cross-sectional basis for econometric analysis. Moreover, when considering co-movement with other variables, we are able to match the estimated moments for inflation at the individual level with corresponding individual moments estimated from equivalent histograms for gross domestic product (GDP) growth and the unemployment rate.

Before attempting to study the properties of the full distribution, it is necessary to estimate the key moments summarizing its location, spread, symmetry, and tail risks. The density forecasts are provided as histograms for which every forecaster reports probability forecasts that reflect their assessment of the likelihood that future inflation will fall within certain intervals. Formally, denote with $p^j_{i,t+h|t}$ the probability that forecaster i $(i = 1, \ldots, N)$ in survey period t attaches to the event that the inflation rate in period $t + h$ falls into a particular interval j $(j = 1, \ldots, J)$.[8] To compute mean long-term expectations, the corresponding inflation uncertainty, and higher moments of the density forecasts, we adopt the most common approach as our baseline estimates and we then consider robustness with respect to alternative approaches. The baseline approach is non-parametric and assumes that all the probability mass in a particular interval j is compressed at the midpoint of this interval, which we denote by μ_j. We assume that the open intervals at both ends of the distribution have the same width as all other intervals and accordingly compute the midpoints of these intervals in the same way. The

[8]Note that J, which represents the total number of surveyed intervals, changes over time, as the survey design was changed at several points in time (see online appendix A). As the probabilities must sum to 100 percent, it can be reasonably assumed that a probability of 0 is assigned by agents to intervals that were not included in a particular survey round.

mean $\pi_{i,t+h|t}$ is, for instance, computed as the probability-weighted sum of each midpoint (equation (1)). The first four moments are computed as follows:

$$\text{Mean:} \qquad \pi_{i,t+h|t} = \sum_{j=1}^{J} p^{j}_{i,t+h|t}\mu_j \tag{1}$$

$$\text{Variance:} \qquad \sigma^2_{i,t+h|t} = \sum_{j=1}^{J} p^{j}_{i,t+h|t}(\mu_j - \pi_{i,t+h|t})^2 \tag{2}$$

$$\text{Skewness:} \qquad S_{i,t+h|t} = \sum_{j=1}^{J} p^{j}_{i,t+h|t}(\mu_j - \pi_{i,t+h|t})^3/\sigma^3_{i,t+h|t} \tag{3}$$

$$\text{Kurtosis:} \qquad k_{i,t+h|t} = \sum_{j=1}^{J} p^{j}_{i,t+h|t}(\mu_j - \pi_{i,t+h|t})^4/\sigma^4_{i,t+h|t} - 3 \tag{4}$$

As highlighted in Engelberg, Manski, and Williams (2009), it must be acknowledged that our moment estimates are subject to measurement and estimation error. In our case, the assumptions we make regarding the allocation of probability mass to the midpoints of the surveyed intervals may be a particularly important source of such error, and it is therefore of interest to test the overall robustness of our conclusions with respect to the above nonparametric approach. Therefore, we present alternative results based on different assumptions and parametric methods below and, in more detail, in online appendix B. These are based on (i) the assumption that the probability mass is distributed uniformly within bins, (ii) the assumption that the true underlying distribution is normal, and (iii) the approach of Engelberg, Manski, and Williams (2009) that assumes a beta distribution for cases with more than two used bins and a triangular distribution for cases with fewer used bins.[9] Furthermore,

[9]For the standard deviation, we find very close correspondence across all methods considered (see online appendix B for further details). For the skewness and the kurtosis, also both discrete approximations are highly correlated, while the correlations with the moment estimates produced by the beta function tend to

rounding on the part of respondents[10] may be an important source of measurement error and, as highlighted in online appendix C, it may be associated with bias in particular affecting the estimation of the second and fourth moments.

To summarize the common pattern across individual SPF respondents, from the individual moments computed according to equations (1) to (4), we can construct the cross-sectional averages by summing each of the individual moments and dividing by the number of density forecast available at each point in time. One reason to focus on the average moments rather than the individual moments when testing for structural breaks in the next session is that measurement error is likely to affect more strongly the estimation of moments at the individual level. When averaging across individuals, idiosyncratic individual-specific approximation errors tend to offset each other, resulting in an overall smaller measurement error (see online appendix C for further details on this).

Note that an estimate of the central tendency of expectations is also available directly from the reported point forecasts that are collected in the SPF. We refrain from using this information in the computation of moments for two main reasons. First, it is not clear whether panelists report their expected mean, mode, or median as their point forecasts and whether density forecasts and point

---

be lower. These lower correlations suggest that a key possible source of measurement error may link to the assumption of a discrete versus a continuous distribution function. We have opted for the discrete approach because it is in line with the discrete nature of the survey design (i.e., the discrete bins used in the survey questionnaire) and does not attribute a specific continuous functional form to the respondents' subjective densities. Nonetheless, the analysis points to a potentially important role for measurement error in influencing the third and fourth moment estimates.

[10] There is empirical evidence that a large fraction of reported interval probabilities are integers or multiples of 5 or even 10 (Manski and Molinari 2010; Glas and Hartmann 2018; online appendix C) which, on the one hand, suggests that forecasters tend to round the raw probabilities derived from a model for the density of the continuous random variable to be forecast. On the other hand, many forecasters state that they use judgment (instead of formal models) when producing the density forecasts (ECB 2014). This could imply that the reported probabilities are the (raw) outcome of a more intuitive heuristic process rather than representing the rounding of an underlying continuous distribution. This latter view suggests that rounding per se may not be a source of measurement error in the reported probabilities.

Figure 1. Different Measures of Central Tendency of Long-Term Inflation Expectations

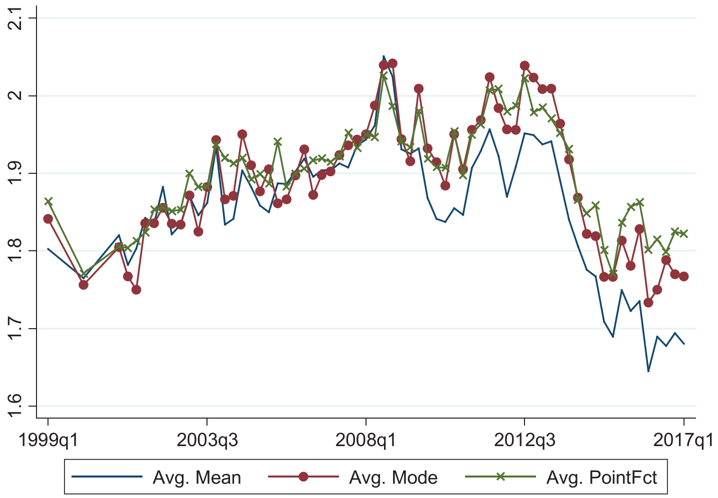

Notes: Average (Avg.) mean expectations and average modal expectations are computed based on the reported density forecasts. The average point forecasts are computed based on the reported individual point forecasts.

forecasts are consistent.[11] Second, a focus on mean expectations from equation (1) is also justified because it draws on all the probabilities collected from respondents. As such, it may contain more information than the long-term point forecasts. Figure 1 plots the mean estimated according to equation (1) together with two other measures of central tendency taken from the survey. The first is the

[11]The SPF survey itself does not offer any information concerning what part of the distribution the point forecasts capture, e.g., whether panelists report a most likely or modal outcome or a conditional expectation (e.g., based on specific assumptions about policy or other exogenous variables) or whether density forecasts and point forecasts are even consistent. Engelberg, Manski, and Williams (2009) study the consistency between density and point forecasts and García and Manzanares (2007) provide evidence that the density forecasts from the SPF are more reliable than the point forecasts. In a separate (unreported) robustness analysis we also compute our moment estimates under the assumption that the point forecasts equate with the mean of the distribution. As expected, we do find that this assumption can change significantly our moment estimates, especially the estimated skewness.

mode computed as the midpoint of the interval which is assigned the maximum probability, again averaged across the responding forecasters in a given round. The second alternative measure is simply the average point forecast. Overall, one observes a very clear comovement and similarity between these three measures of central tendency. In particular, the average point forecast and the estimated mode are very closely related, suggesting that when they give their point forecasts survey respondents may be giving a modal prediction rather than an estimate of their mean expectations. In the period since the financial crisis mean expectations dropped slightly below both the estimated average mode and the reported point forecasts. Given this divergent pattern, we also include the estimated mode of the probability distributions in our empirical analysis of potential shifts in the distribution's location.

3. Evidence of Change in the Distribution of Long-Term Inflation Expectations

In this section we report our empirical analysis of possible structural breaks in the distribution of long-term inflation expectations. We first discuss the result of a general test for distributional change due to Inoue (2001). This test provides nonparametric prima facie evidence about possible structural change in the whole distribution. We then present the methods used to estimate the timing and magnitude of possible breaks in the first four moments of the distribution before reporting the main findings. Also, to complement the proposed tests of stability in the estimated moments, we apply these breakpoint tests to selected interval probabilities taken directly from the surveyed histograms. The advantage of this complementary approach is that it offers more direct evidence for change in the distribution that does not require the assumptions needed for the estimation of distributional moments.

3.1 Testing for Distributional Change

Before proceeding to tests for change in the moments, it is of interest to examine the prima facie evidence for changes in the distribution as a whole. To do this, Inoue (2001) suggests a nonparametric test which is particularly useful in situations where the correct specification of the underlying datagenerating process is unknown.[12] Inoue (2001) proposes a test of change in the distribution function of a time series x_t that is observed over a total sample of n periods. The test is based on the difference, $(T(s, \tau))$, between the empirical distribution functions for two subsamples defined by a candidate change date $s < n$:

$$T(s, \tau) = \left| \frac{1}{s} \sum_{t=1}^{s} I(x_t \leq \tau) - \frac{1}{n-s} \sum_{t=s+1}^{n} I(x_t \leq \tau) \right|. \qquad (5)$$

In equation (5), the measured distance between the two distribution functions is based on an observed time series and uses the indicator function, $I(\cdot)$, to estimate the relevant probabilities from the time-series observations.

Our data, however, contain a sequence of reported distribution functions in the form of the SPF histograms. We construct a modified test statistic by replacing the empirical distribution functions with subsample averages of the observed distribution functions. Denoting the probability that inflation will be below τ in the long run as reported in period t by $CDF_t^{SPF}(\tau)$, we obtain the modified expression

$$T(s, \tau) = \left| \frac{1}{s} \sum_{t=1}^{s} CDF_t^{SPF}(\tau) - \frac{1}{n-s} \sum_{t=s+1}^{n} CDF_t^{SPF}(\tau) \right|. \qquad (6)$$

Note that in the case of the application to the SPF, the grid for τ is determined by the bins defined in the survey questionnaire. Using the distance measure in equation (6), a weighted Kolmogorov-Smirnov test statistic can be computed as

$$T_1 = \sup_{1 \leq s < n} \sup_{\tau \in \mathbb{R}} \left| \frac{s}{n} \left(1 - \frac{s}{n}\right) n^{1/2} \times T(s, \tau) \right|. \qquad (7)$$

To determine critical values for this modified test, we use the bootstrap that Inoue (2001) proposes.[13] We obtain a test statistic of

[12] We would like to thank the editor Barbara Rossi for pointing us to this test, which is very useful in our context.

[13] For the bootstrap we set the number of iterations to 1,000 and use a window size of nine, following the rule of thumb suggested in Inoue (2001).

$T_1 = 0.174$, which is marginally larger than the 10 percent critical value of 0.173 (the 5 percent and 1 percent critical values are 0.192 and 0.225, respectively).[14] This means that we can reject the null hypothesis that there is no structural break at the 10 percent level. The test statistic is maximized for $s = 2009{:}Q4$, indicating that the distribution of long-term inflation expectations most likely changed during the Great Recession. Although certainly not definitive, this result motivates a further investigation of how the distribution may have changed following the Great Recession. Motivated by this result, we explore the data for further evidence of distributional change by testing for evidence of structural breaks. Below we first describe the implementation of these structural break tests in general and then apply them to the specific interval probabilities collected in the SPF survey and also to the distributional moments.

3.2 Testing for Structural Breaks

One condition for long-term inflation expectations to be well anchored is that their distribution is relatively stable around the central bank's inflation target. To examine this, we test for evidence of any breaks in the probabilities assigned to certain intervals and, subsequently, in the first four moments. Since a priori we know neither the number of breaks nor their timing, we employ the method of Bai and Perron (1998, 2003), who consider the linear regression model with a finite number of possible regimes defined by unknown breaks in the model parameters. Their method yields estimates of the unknown regression coefficients associated with each regime together with estimates of the unknown breakpoints. To select the number of significant breaks and to date their occurrence, we apply these methods by regressing a given interval probability or moment on its first lag and an intercept term. The inclusion of the former term is intended to account for the strong persistence in such long-horizon moments and probability forecasts which would

[14] All data used in this paper and codes for replicating all results can be downloaded from https://www.dropbox.com/s/ue9pyq0cv03mffq/Replication_Files_Dovern_Kenny_2019.zip?dl=0.

otherwise cause substantial residual autocorrelation in our regressions.[15] Formally, the considered model is

$$\bar{m}_{t+h|t} = \alpha_{m,r} + \beta_{m,r}\bar{m}_{t+h|t} + \varepsilon^{m}_{t+h|t}, \tag{8}$$

where $\alpha_{m,r}$ and $\beta_{m,r}$ are regime-specific parameters (with $r = 1, \ldots, R$), $\varepsilon^{m}_{t+h|t}$ is an iid error term, and $\bar{m}_{t+h|t}$ represents the cross-sectional average of individual interval probabilities or moments. For the case of mean long-term expectations ($m = \pi$), for instance, $\alpha_{\pi,r}/(1 - \beta_{\pi,r})$ yields an estimate of the average expected rate of long-term inflation in a given regime. Under the assumption that forecasters believe that, in general, the central bank is able to achieve its inflation target in the absence of further shocks, such an estimate provides a regime-specific measure of the perceived inflation objective. In the absence of any structural breaks, we have $\alpha_{\pi,r} = \alpha_{\pi}$ and $\beta_{\pi,r} = \beta_{\pi}$, which implies that the long-term expectation for inflation is constant. In a similar way, for each of the three higher-order moments and the interval probabilities, $\alpha_{m,r}/(1 - \beta_{m,r})$ provides an estimate for the average perceived future variance, skewness, tail risk, or interval probability embodied in the distribution in any given regime. Breaks of these moments and probabilities may signal changes in the degree to which the distribution of long-term inflation expectations is anchored.

In implementing the Bai-Perron (BP) procedure, it is necessary to specify a *minimum* required distance between any two potential break dates. We set this minimum distance between breaks at a relatively conservative level of eight quarters in order to avoid overfitting and possibly finding an implausibly large number of spurious breaks for each moment. In Bai and Perron's jargon, given our sample size, this implies a "trimming parameter" of $8/72 \approx 0.11$. We determine the number of structural breaks by looking at a modified Schwarz criterion (LWZ) which is suggested for the case of models

[15] According to the Ljung-Box test, we cannot reject the null hypothesis of no residual autocorrelation at the 5 percent level for both the specification without structural breaks and the one with breaks for all moments with one exception: In the case of inflation uncertainty, the large structural break causes mild negative autocorrelation in the residuals of the specification that does not allow for this break.

that include a lagged dependent variable and the sequential SupF-test which is the overall preferred test according to Bai and Perron (2003).[16] Because it turns out that we find evidence for at most one break in all the moments that we consider, we run robustness checks using the test proposed by Andrews and Ploberger (1994). This tests the null hypothesis of no break against the alternative of exactly one break (at an unknown point in time) and is more powerful as well as robust to heteroskedasticity and residual autocorrelation.

Clearly, an important practical consideration for all of the tests discussed above is whether they are sufficiently powerful to identify actual breaks in the data, particularly in the presence of measurement error. To investigate this issue, online appendix C reports the results of two simulation exercises which analyze how the size and power of the breakpoint tests we employ might be affected by measurement error. Overall, the results of these simulations suggest that the tests have reasonable power in samples of the size we are using even in the presence of measurement error. At the same time, the simulations indicate that the power for detecting breaks in higher moments may be lower the more rounding is applied when forming the interval probabilities. Thus, any results on the number of breaks reported below should be considered a "lower bound."[17]

3.3 Breaks in Selected Interval Probabilities

In this section, we apply the above breakpoint analysis to test for breaks in selected interval probabilities over time. In particular, we consider the average probabilities attributed by forecasters to the following five long-term inflation scenarios: (i) outright deflation (long-term inflation outcome below 0 percent), (ii) relatively low inflation of between 0 percent and 1.5 percent (i.e., low compared with the ECB mandate of below but close to 2.0 percent),

[16]The sequential test is based on the idea of sequentially testing the null hypothesis of no l breaks versus $l + 1$ breaks until the null hypothesis can no longer be rejected. In each step of this sequence and given a set of l breakpoints, Bai and Perron (1998) suggest applying $l + 1$ tests of the null hypothesis of no structural break against the alternative hypothesis of a single structural break to the $l + 1$ segments of a time series defined by the l breaks.

[17]However, given that only a fraction of probability statements in the SPF is rounded while in our simulations we assume all of the probabilities are rounded, the issue should certainly not be overrated.

Figure 2. Breakpoints for Selected Interval Probabilities

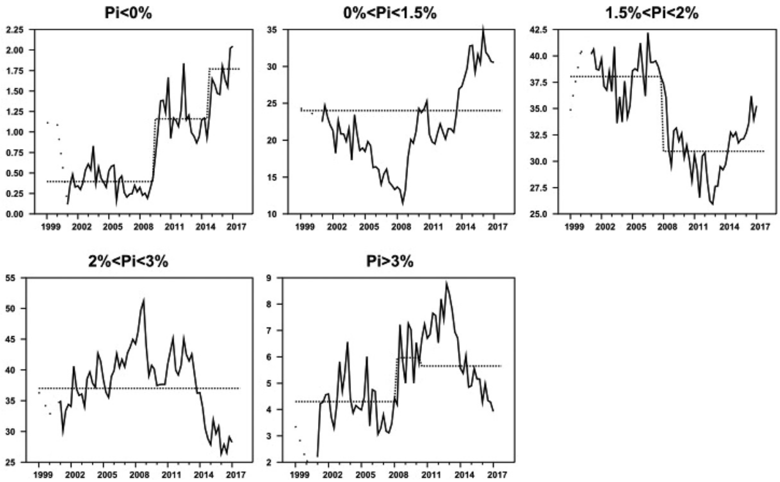

Notes: Selection of breakpoints based on Bai and Perron (2003). The solid black lines refer to the average probability (across forecasters) assigned to the intervals. The dotted lines show the implied unconditional means for different subperiods, with breaks in AR(1) models for the average moments selected using the LWZ statistic. The minimum distance between two breakpoints was set to eight quarters.

(iii) inflation broadly consistent with the ECB's mandate (between 1.5 percent and 2.0 percent), (iv) inflation moderately above target (between 2 percent and 3 percent), and (v) relatively high long-term inflation (above 3 percent). These selected intervals provide an economic narrative relative to the ECB's definition of price stability. For example, a reduced probability mass assigned to the interval between 1.5 percent and 2.0 percent may signal a deterioration in the degree to which long-term expectations are anchored at a level that is below but close to 2.0 percent. In selecting these intervals, we also avoided having to interpolate probabilities within intervals, as this would require the use of auxiliary assumptions that would be hard to justify.

Figure 2 shows the evolution of these probabilities across survey rounds together with their estimated means for which we select significant breaks based on Bai and Perron (1998, 2003) using the LWZ statistic. Overall we find evidence for significant breaks in three of the five intervals examined. On the one hand, we find no significant breaks in the probabilities assigned to the low inflation outcomes (0.0 percent to 1.5 percent) and outcomes moderately above the target (2.0 percent to 3.0 percent). Despite the lack of evidence of a break in these interval probabilities, on the other hand, the sizable movements in the probabilities of these two events are certainly of interest from a policy perspective. In particular, following the Great Recession, there was a large rise in the probability of low inflation, from less than 15 percent to more than 30 percent. Also, the chances of moderately higher inflation dropped from a peak of over 50 percent to below 30 percent toward the end of the sample period. These relative shifts in probabilities are also suggestive of a possible shift toward a more negatively skewed distribution. In contrast, the tests do indicate one downward break for the probability associated with inflation being broadly in line with the target (1.5 percent to 2.0 percent) and two breaks for each of the two outer intervals (deflation and relatively high inflation above 3.0 percent). The direction of the breaks indicate that following the Great Recession forecasters have reduced their assessment of the likelihood that long-term inflation falls between 1.5 percent and 2.0 percent from above 37.5 percent before the Great Recession to below 32.5 percent in the second half of the sample. The lower probability assigned to this outcome range, which is the most consistent with the ECB definition of price stability of below but close to 2.0 percent, can be taken to imply a reduction of the *degree* to which long-term inflation expectations were anchored following the Great Recession. Such a result is suggestive of a more spread-out long-run distribution that has a lower concentration of probability mass in the area surrounding mean expectations and the central bank's target. It may therefore also imply higher long-run inflation uncertainty. There also seem to be (ongoing) shifts toward lower and more negative inflation outcomes, as indicated by the persistent increase in the probabilities associated with negative inflation. In this case, however, it is important to note the considerably smaller shift in the probabilities from below 0.5 percent in the first part of the sample to around 1.75 percent in the more recent period. Overall, therefore, these direct tests of stability in the reported interval probabilities tend to provide empirical support for possible changes in the long-run distribution. In the following sections, we offer further evidence on the nature of these changes based on the estimated moments.

Figure 3. Breakpoints for Moments of Density Forecasts

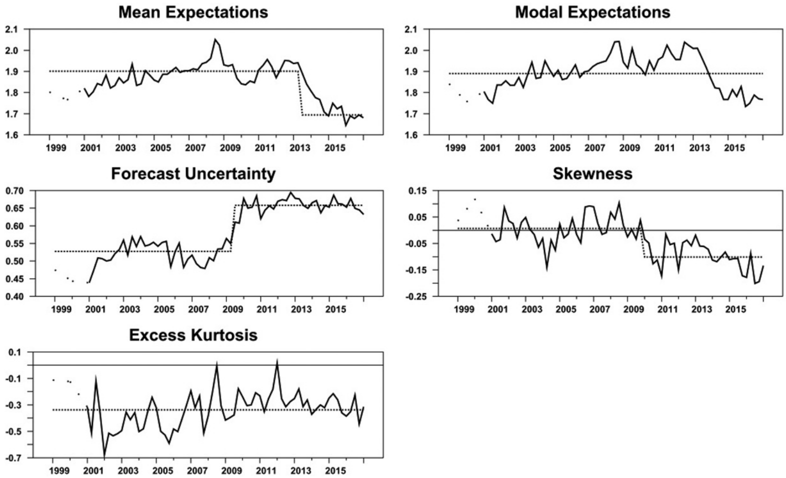

Notes: Selection of breakpoints based on Bai and Perron (2003). The solid black lines refer to the average moments of the density forecasts of the individual SPF participants. The dotted lines show the implied unconditional means for different subperiods, with breaks in AR(1) models for the average moments selected using the LWZ statistic. The minimum distance between two breakpoints was set to eight quarters.

3.4 Breaks in the Mean

We now turn to the analysis of breaks in the estimated moments to provide more formal evidence about which features of the predictive densities may have changed over time. Figure 3 shows the evolution of the modal expectation discussed above together with the first four moments given by equations (1) to (4). In each panel, the regime-specific moment estimates $\hat{\alpha}_{m,r}/(1-\hat{\beta}_{m,r})$ are also plotted (as identified by the LWZ statistic). The breakpoint analysis of density moments is further detailed in table 1. In particular, we report a list of the estimated break dates and the corresponding F-statistics indicating the significance of each break as well as complementary results based on the sequential SupF-test.

For mean expectations, we find one significant break in 2013:Q2 according to both the LWZ statistic and the Andrews and Ploberger (AP) test reported in table 1. The break is downward and occurs in

Table 1. Breakpoints for Average Moments of Density Forecasts

LWZ			Sequential SupF			Andrews-Ploberger		
Period	**F-test**	**Mean**	**Period**	**F-test**	**Mean**	**Period**	**F-test**	**Mean**
A. Mean Expectations								
1999:Q1–2013:Q2 2013:Q3–2017:Q1	7.89***	1.90 1.69	1999:Q1–2017:Q1		1.83	1999:Q1–2013:Q2 2013:Q3–2017:Q1	7.89***	1.90 1.69
B. Model Expectations								
1999:Q1–2017:Q1		1.89	1999:Q1–2017:Q1		1.89	1999:Q1–2017:Q1		1.89
C. Inflation Uncertainty								
1999:Q1–2009:Q2 2009:Q3–2017:Q1	16.56***	0.53 0.66	1999:Q1–2009:Q2 2009:Q3–2017:Q1	16.56***	0.53 0.66	1999:Q1–2009:Q2 2009:Q3–2017:Q1	16.56***	0.53 0.66
D. Skewness								
1999:Q1–2009:Q4 2010:Q1–2017:Q1	8.05***	0.01 −0.10	1999:Q1–2017:Q1		−0.05	1999:Q1–2009:Q4 2010:Q1–2017:Q1	8.05***	0.01 −0.10
E. Excess Kurtosis								
1999:Q1–2017:Q1		−0.34	1999:Q1–2017:Q1		−0.34	1999:Q1–2006:Q2 2006:Q3–2017:Q1	5.95***	−0.45 −0.29

Notes: The number of breaks and their location is selected based on the modified Schwarz criterion (LWZ) and the sequential SupF-test proposed by Bai and Perron (2003). The minimum distance between two breaks is set to eight quarters. The third test is taken from Andrews and Ploberger (1994) and tests against the alternative of only one break at an unknown point in time. Dependent variables are the average moments of the density forecasts of the individual SPF participants. Dates refer to periods in which we observe a significant change in the parameters of an AR(1) model for the respective moment. Implied unconditional means are computed for every break segment based on the estimated coefficients. *** indicates that breaks in the model parameters are jointly different from 0 at a 1 percent significance level.

the wake of the Great Recession and euro-area sovereign debt crisis. Although quantitatively modest, the break is noteworthy, with the regime-specific mean falling to 1.69 percent from 1.90 percent prior to the break, and points to growing beliefs that long-term inflation outcomes could be "below" but "not so close to" 2.0 percent. The drop in mean expectations is found to be statistically significant given the low overall volatility of the time series. Importantly, this break in the mean is detected across three of the four moment estimation methods presented in online appendix B. However, when one fits a beta distribution to the survey histograms, a break in the mean is not detected by both the LWZ-based BP test and the AP test. The case for a significant downward shift of long-term inflation expectations is also weakened by two additional results reported in table 1. First, the sequential SupF-test does not identify any breaks in the dynamics of mean expectations and, as a result, the estimate for the average over the full sample is constant at 1.83 percent. Second, none of the tests identifies any breaks in the case of modal expectations, which average 1.89 percent over the full sample. A finding of a downward break in the mean combined with a stable modal value is consistent with a shift toward a more negatively skewed distribution, and we examine this hypothesis directly below by applying the same test to the estimated third moment.

Overall, our results for mean expectations provide at most only weak evidence for a quantitatively modest decline in long-term inflation expectations toward the end of our sample. Mean inflation expectations also remain in a range that can be considered consistent with price stability as defined by the ECB's price stability objective of "below but close to 2.0 percent." Hence, our analysis provides no grounds to think that the central tendency of euro-area long-term inflation expectations became unanchored. This finding is in line with other recent studies that use market-based inflation expectations, as, for instance, Autrup and Grothe (2014), Strohsal and Winkelmann (2015), and Speck (2016).

3.5 Breaks in the Variance

Figure 3 and table 1 report equivalent results for higher moments. In the case of long-term inflation uncertainty, all three reported tests identify a break in 2009:Q3 that is associated with an increase in uncertainty about long-term inflation. In relative terms, this shift is quantitatively more noticeable than the shift in mean expectations. After the break, i.e., in the immediate wake of the Great Recession, average uncertainty about the long-term inflation at 0.66 percentage points (pp) had increased by about 25 percent compared with its level in the pre-crisis regime (0.53 pp). Following this increase, long-term inflation uncertainty has been quite stable at the new level and has shown no tendency to decline. This pronounced and persistent increase in the variance of the distribution signals that forecasters perceive the long-term inflation outlook to be more uncertain now than before the Great Recession.[18] As a result, the distribution has become less concentrated around levels that are consistent with the definition of price stability. Interestingly, the increase correlates well with an increase of the historical variance of annual inflation rates in the euro area when computed recursively based on an expanding sample starting in 1997; the latter increases from about 0.2 in 2007 to a little below 0.6 in 2014. Thus, it seems as if forecasters anticipate that the increase in inflation volatility which emerged after the Great Recession will be highly persistent—or even permanent—rather than being relevant only to recent years or the short-term outlook for inflation.[19] As a caveat, it should be mentioned that the estimated level of long-term inflation uncertainty could be subject to bias reflecting the rounding behavior of forecasters when they respond to the survey (see online appendix C for further details).

Overall, the identified upward shift in the variance could be taken to imply that the *degree* to which long-term inflation expectations are anchored has diminished. The shift in the variance is consistent with macroeconomic theories highlighting the implications of uncertainty surrounding the central bank's objective (Beechey, Johannsen, and Levin 2011) and, in particular, the potential effects of the lower

[18]It is noteworthy that the upward adjustment came quite soon after the Great Recession and the sovereign debt crisis was not associated with any further rise in long-term inflation uncertainty.

[19]The above-identified upward break in the variance appears to be relatively robust with respect to measurement error. For example, in online appendix B we show how, across four different ways to estimate the variance, we identify the same break in 2009:Q3 using both the LWZ test and the AP test.

bound on nominal interest rates (Benhabib, Schmitt-Grohé, and Uribe 2001). However, the higher variance is also consistent with the view that forecasters believe it may take longer for the central bank to achieve its price stability objective, e.g., as a result of more persistent and volatile shocks in the future or a perceived change in the transmission of monetary policy. It does not necessarily imply that they have reduced their belief in the ECB's ability to ultimately achieve that objective over a longer horizon than the five years to which the survey data relate. Nonetheless, our findings highlight an important challenge for monetary policy and its communication; namely to limit any further the rise in the uncertainty surrounding long-term inflation prospects.

3.6 Breaks in Symmetry and Tail Risk

Figure 3 and table 1 also report the results of the break analysis for the skewness and tail risk in the distribution of long-term expected inflation in the euro area. The time path of the average skewness provides insight on possible changes in the symmetry of the distribution and may thus signal concerns among forecasters about long-term inflation risks either to the upside or the downside. The analysis reveals that, since 2010:Q1, the forecast densities are negatively skewed, on average, whereas prior to this date they were broadly symmetric. This means that since the Great Recession forecasters have assigned a greater share of the overall probability mass to relatively low as opposed to relatively high inflation outcomes. This finding of a negatively skewed distribution is in line with our previous result of a possible downward break in mean expectations with a higher and more stable mode. It is also precisely what is predicted by macroeconomic models which incorporate a lower bound constraint on nominal interest rates (e.g., Coenen and Warne 2014; Hills, Nakata, and Schmidt 2016). Interestingly, according to both the LWZ and AP test results reported in table 1, the break in skewness occurred relatively soon after the Great Recession and prior to the break in mean expectations. The tendency toward a negatively skewed distribution has also been highly persistent and has lasted up to the end of our sample in 2017:Q1. At the same time, these results have to be taken with some caution because the alternative

SupF-test identifies no break in skewness.[20] For the kurtosis, only the AP test identifies a small increase in tail risk. However, the long-term inflation density is platykurtic, which means that forecasters believe that, relative to a normal distribution, less of the inflation uncertainty is associated with infrequent tail events. Overall, therefore, the survey data tend to give relatively small weight to tail risks for inflation.[21] Importantly, as we show in online appendix C, estimates of kurtosis appear to be the most uncertain and the most sensitive to approximation error, in particular linked to rounding behavior on the part of respondents. Using other moment estimation techniques, we find no clear evidence of a break in kurtosis (see online appendix B).

4. Co-movement of Moments with other Variables

The previous section provided evidence that the distribution of long-term inflation expectations in the euro area changed in the period since the Great Recession. In particular, our analysis shows that the distribution experienced a modest downward shift in its mean, a more sizable and significant increase in its dispersion, and a persistent negative skewness. As discussed in the introduction, a second way to shed light on how well-anchored long-term inflation expectations have been is to examine the strength of their co-movement with other economic variables. In a perfect world with no uncertainty about monetary policy effectiveness, we might expect to observe no correlation between moments of the distribution of expectations and other macroeconomic developments. However, in a world with uncertainty and learning on the part of private agents, some co-movement even with short-term macroeconomic developments at business cycle frequencies can be expected. In such circumstances, an increased sensitivity of long-term inflation expectations to such factors would

[20]In online appendix B we provide further information on the robustness of this break in skewness. We find that most moment estimation methods that allow for asymmetry do identify a break in the first half of 2010, in line with our baseline results.

[21]This is a well-known finding in the literature exploiting such data. It is a direct consequence of the fact that many participants in the SPF attach positive probability weights to only a very limited number of the bins in the survey questionnaire (see Kenny, Kostka, and Masera 2015).

be indicative of a change in the degree to which expectations are anchored.

In this section, we use panel regressions to examine the co-movement between the first two moments of the distribution of expected long-term inflation with other economic variables.[22] The proposed panel analysis offers a way to exploit the individual-level replies in order to assess the extent of *average* co-movement of mean long-term expectations and long-term inflation uncertainty with other variables. In particular, the regression coefficients represent the *average* co-movement that is derived from the micro data. Hence, these coefficients, and the co-movements that they identify, match very well the macroeconomic focus of the study. In particular, we exploit the individual expectations in a panel setting to address the following questions: Are there factors that co-move strongly with long-term inflation expectations and uncertainty? Did the role of such factors change after 2007, the year in which the financial crisis began to unfold? Can we identify any coincidences between recent monetary policy events—such as the hitting of the ELB on nominal interest rates or the introduction of nonstandard monetary policies—and changes in the first two moments of the distribution? We first present the results for the mean and then turn to the analysis of the variance.

4.1 Long-Term Mean Inflation Expectations

To examine the co-movement between mean expectations and other macroeconomic developments, we consider the following set of variables: First, we consider the change in forecaster-specific short-term (one-year-ahead) and medium-term (two-year-ahead) inflation expectations ($dE_i[\pi(1y)]_t$ and $dE_i[\pi(2y)]_t$). In addition, we look at the reaction of long-term expectations to past short- and medium-term forecast errors for inflation, again at the individual level ($\pi - E_i[\pi(1y)]_{t-4}$ and $\pi - E_i[\pi(2y)]_{t-8}$). To capture perceived

[22] We focus on the first two moments because these are the ones for which we can readily identify a set of appropriate covariates. However, given the results of the break test analysis in section 3 above, the examination of co-movement with other higher moments or with the interval probabilities discussed in section 3 represents an interesting avenue for future empirical research.

structural changes linked to possible concerns about secular stagnation and deflationary equilibriums, we include the change in long-term expectations for GDP growth ($dE_i[GDP(5y)]_t$) and the unemployment rate ($dE_i[U(5y)]_t$) from the SPF. In the spirit of Levin, Natalucci, and Piger (2004), who show that long-term inflation expectations in the United States and the euro area were highly correlated with a slow moving average of inflation over the period 1994 to 2003, we also consider an inflation "performance gap" as the difference between recent long-term expectations and a (five-year) moving average of past inflation ($\mathrm{MA}(\pi)_{t-1} - E_i[\pi(5y)]_{t-1}$). In addition to these forecaster-specific variables, we consider the possible co-movement with factors that are common across forecasters, such as shocks to the inflation process itself, as reflected in the recently observed change in the inflation rate ($d\pi_{t-1}$). Also, we consider the change in the volume of the ECB's balance sheet (dCBBS_{t-1}) as a monetary policy indicator that is associated with recent unconventional monetary policies and quantitative easing, with the expectation that an expansion of the balance sheet might potentially lead agents to change their expectation of the long-term inflation rate upward. To control for important monetary policy changes, we additionally include a dummy variable for quarters following the announcement of important nonstandard monetary policies (MPA_t)[23] and a dummy capturing the hitting of the effective lower bound (ELB_t).[24]

[23]The events are the introduction of enhanced credit support on May 7, 2009; the introduction of the Security Market Program (SMP) on May 10, 2010; the introduction of the Outright Monetary Transaction (OMT) program on August 2, 2012; the introduction of forward guidance on July 4, 2013; the introduction of the enhanced asset purchase program (APP) on January 22, 2015; and the extension of the APP on March 10, 2016.

[24]There is some uncertainty concerning the precise date on which the effective lower bound became binding in the euro area. Our empirical analysis reflects this uncertainty, with the ELB dummy taking a value of one for the two survey rounds after July 11, 2012 and June 5, 2014, and zero otherwise. These dates correspond, respectively, to the policy meetings when the ECB deposit rate was cut to zero and the rate on the ECB's main financing operations was cut by 10 basis points to 0.15 percent. Our focus on the June meeting reflects the ECB communication at that time. In particular, when asked at the monetary policy press conference about future interest rate reductions, the president of the ECB stated, "I would say that for all the practical purposes, we have reached the lower bound. However, this doesn't exclude

We use an unbalanced panel regression with fixed effects to look at the co-movement discussed above and to test whether the correlation structure has changed since the Great Recession. Let $\Delta\pi_{i,t+20|t} = \pi_{i,t+20|t} - \pi_{i,t-1+20|t-1}$ denote the change in the long-term inflation expectation of an individual forecaster. We regress this change on forecaster-specific fixed effects and the set of proposed covariates. Collecting the forecaster-specific variables in a vector $X_{i,t}$ and the common covariates in Y_t, the linear panel regression is given by

$$\Delta\pi_{i,t+20|t} = \alpha_i + MPA_t\beta_{MPA} + ELB_t\beta_{ELB} + X_{i,t}\beta_X + Y_t\beta_Y + \varepsilon_{i,t}, \tag{9}$$

where $\varepsilon_{i,t}$ is an error term that we allow to exhibit both spatial and temporal correlation (Driscoll and Kraay 1998).[25] The parameters β_{MPA} and β_{ELB} and parameter vectors β_X and β_Y measure the extent to which long-term expectations co-move with the other variables. For perfectly anchored inflation expectations, we would expect the effect of the two dummy variables to be insignificant and also both $\beta_X = 0$ and $\beta_Y = 0$.

Table 2 lists the results of estimations of equation (9). The first column reports the coefficient estimates over the full sample and without allowing for any break in the parameter values. The second specification allows all parameters in the equation (except for the fixed effects and the constant) to break after 2007:Q3, which we chose because it represents the start of the Great Recession and the

some little technical adjustments and which could lead to some lower interest rates in one or the other or both parts of the corridor. But from all practical purposes, I would consider having reached the lower bound today." See the "Introductory Statement to the Press Conference (with Q&A)" available at https://www.ecb.europa.eu/press/pressconf/2014/html/is140605.en.html.

[25]Driscoll and Kraay (1998) suggest a nonparametric way for estimating standard errors that are heteroskedasticity- and autocorrelation-consistent and robust to general forms of temporal and cross-sectional dependence. Their approach uses cross-section averages of regressors and residuals to compute a heteroskedasticity- and autocorrelation-consistent (HAC) estimator. Given that forecasters form their predictions simultaneously and that professional forecasts are usually found to be subject to information rigidities (Coibion and Gorodnichenko 2012; Dovern et al. 2015), which cause forecast revisions to be autocorrelated, both features are important. Ignoring them and assuming independently distributed error terms is likely to results in an underestimation of the true long-term covariance.

Table 2. Co-movement of Long-Term Inflation Expectations with Other Variables

	Without Break	With Break	
	Full-Sample Coefficients	**Pre-2007:Q4 Coefficients**	**Change after 2007:Q4**
$dE_i[\pi(1y)]_t$	0.011	0.006	0.010
	(0.02)	(0.04)	(0.04)
$dE_i[\pi(2y)]_t$	0.177***	0.205***	−0.052
	(0.02)	(0.05)	(0.06)
$(\text{MA}(\pi) - E_i[\pi(5y)])_{t-1}$	0.155***	0.187***	−0.026
	(0.02)	(0.04)	(0.05)
$\pi - E_i[\pi(1y)]_{t-4}$	−0.006	0.003	−0.010
	(0.01)	(0.02)	(0.03)
$\pi - E_i[\pi(2y)]_{t-8}$	0.007	0.017	−0.013
	(0.01)	(0.03)	(0.03)
$dE_i[GDP(5y)]_t$	0.007	0.019	−0.017
	(0.03)	(0.03)	(0.06)
$dE_i[U(5y)]_t$	−0.013	0.000	−0.019
	(0.01)	(0.02)	(0.02)
$d\pi_{t-1}$	0.031**	0.032	0.001
	(0.01)	(0.02)	(0.03)
dCBBS_{t-1}	−0.160**	−0.882***	0.764**
	(0.06)	(0.34)	(0.34)
MPA Dummy	0.018	—	0.016
	(0.02)	—	(0.02)
ELB Dummy	0.011	—	0.012
	(0.01)	—	(0.01)
Constant	−0.000	−0.004	—
	(0.01)	(0.01)	—
Observations	1,180	1,180	
R^2	0.195	0.204	

Notes: Dependent variable is the change in long-term inflation expectations. Both models include fixed effects for each forecaster. The constant is identified by restricting the average of the fixed effects to equal 0. We report the within R^2. Standard errors are computed using the method of Driscoll and Kraay (1998) and are robust against general forms of spatial and temporal dependence. *, **, and *** denote significance at the 10 percent, 5 percent, and 1 percent level, respectively.

launch of monetary policy measures aimed at mitigating turbulences in financial markets. The table reports the coefficients for the sample period prior to the Great Recession (second column) as well as the estimated change in the coefficients after the Great Recession, i.e., from 2007:Q4 onward (third column).

Looking first at the full-sample results, we observe that only a few of the marginal correlations are significantly different from zero. In particular, those are the correlations with the change in two-year inflation expectations and with the inflation performance gap. For example, long-term inflation expectations are revised by roughly 0.18 pp, on average, when two-year-ahead expectations move by 1 pp. Thus, the co-movement is relatively strong, supporting evidence from U.S. household expectations provided by Dräger and Lambla (2013). Also, according to the estimated parameter values on the inflation performance gap indicator, a 1 pp deviation of the past inflation trend above the announced price stability objective is associated with an upward revision in long-term inflation expectations of just above 0.15 pp. This co-movement potentially highlights the importance of the central bank actually hitting the inflation target in the medium run if long-term inflation expectations are to be stabilized. The observed co-movement is in line with macroeconomic theories which allow for uncertainty about the central bank objective. For example, if private agents are uncertain about the true inflation objective held by the central bank, they may adjust their long-term inflation expectations upward in response to past inflation trends (see, for example, Beechey, Johannsen, and Levin 2011).

In our panel regression results, we also find some quantitatively less important but significant co-movement with the change in actual inflation.[26] The only other significant parameter estimate corresponds to the change in the size of the ECB's balance sheet. The negative sign appears counterintuitive at first and demonstrates that the estimates should not be interpreted causally. A valid interpretation is not that expanding the balance sheet causes inflation expectations to decline. Instead, reverse causality is the most plausible explanation and, also, possible omitted-variable bias. For example, a plausible explanation of the negative co-movement is that the ECB anticipated a decline in inflation expectations and responded with expansionary measures that were associated with an expansion of the balance sheet. On the other hand, both the ECB and the forecasters

[26]This is in contrast to Beechey, Johannsen, and Levin (2011), who find a response of long-term inflation expectations to inflation news for the United States but not for the euro area. However, the sample period used in their study did not cover the period associated with the aftermath of the Great Recession.

included in our panel could simply respond to other macroeconomic news that we do not capture with our variable selection.

Concerning the two monetary policy dummies, we find no evidence for any co-movement with our measure of mean long-term expectations. This implies that the hitting of the ELB and the announcement of nonstandard measures were not immediately followed by a change in inflation expectations. This result is not so surprising given the other controls and co-movement that we have captured in the regressions. For example, it is entirely plausible that the effects of the ELB and the announcement of nonstandard measures are captured via the co-movement with two-year-ahead expectations. The result may also be due to the low frequency of the SPF data, which makes it harder to connect changes in expectations to specific events. Based on higher-frequency data, Karadi (2017) finds that nonstandard measures helped to prevent long-term inflation expectations in the euro area from becoming unanchored after 2013.

The reported changes in the coefficients after 2007:Q4 suggest that the estimated co-movement did not change much after the Great Recession. An F-test of the joint hypothesis that none of the model's coefficients exhibits a structural break in 2007 does not reject this hypothesis. The only coefficient that changes significantly is the one corresponding to the ECB's balance sheet variable.[27] The pre-2007:Q4 estimate is strongly negative and highly significant, suggesting that this period also drives the full-sample results. In contrast, the overall estimate for the sample since 2007:Q4 is much smaller in absolute value and only significant at the 10 percent level.

In summary, the panel regressions suggest that the process governing the means of the subjective forecast distributions is far away from a simple stylized case where the inflation objective is a universal constant and where there is "blind faith" in the ability of the central bank to achieve this objective. Instead, this process is more in line with theories emphasizing uncertainty about the monetary policy transmission mechanism in which agents update their beliefs about long-term inflation in response to relevant shocks. However, there is

[27]These results are clearly a reflection of the major structural break in the balance sheet data after the onset of the financial crisis. Changes in the balance sheet size were very small and regular before 2007 and larger and also more volatile afterward.

no evidence of a substantially higher sensitivity to the main covariates following the Great Recession. This result is also in line with previous results in section 3, where we identified at most only weak evidence that mean inflation expectations had declined following the Great Recession. At the same time, the co-movement that we have identified highlights that the anchoring of the mean of this distribution can in no way be taken for granted. Online appendix B provides more detail on our empirical results and highlights how, although the values of some of the estimates can change slightly, they are generally very robust across different moment estimation strategies. In particular, the strong co-movement of both the medium-term inflation expectations and the central bank performance gap measure with long-term mean expectations as well as the relative stability of coefficients across both subsamples are confirmed for several different moment estimation techniques.

4.2 Long-Term Inflation Uncertainty

We now turn to the co-movement of long-term inflation uncertainty with other variables. We follow the same analytic approach but adjust the set of covariates to correspond to the change in the dependent variable. In particular, we now look into the co-movement with the change in short- and medium-term inflation uncertainty (instead of expectations), and with the *absolute values* of lagged forecast errors and the difference between recently held long-term expectations and the five-year moving average of past inflation. We also consider the co-movement with long-term uncertainty about growth and the unemployment rate (instead of expectations), and with the *absolute values* of changes in the inflation rate and the size of the ECB balance sheet size.

Table 3 reports the results of the panel estimation for long-term inflation uncertainty. Compared with mean expectations, the results suggest that a higher share of variation of long-term inflation uncertainty can be explained by movements in the covariates considered (R^2 of 0.39 for the full sample). The main findings are as follows: First, long-term inflation uncertainty co-moves strongly with inflation uncertainty at a two-year horizon. Second, absolute changes in the current inflation rate are positively correlated with long-term

Table 3. Co-movement of Long-Term Inflation Uncertainty with Other Variables

	Without Break	With Break	
	Full-Sample Coefficients	**Pre-2007:Q4 Coefficients**	**Change after 2007:Q4**
$dV_i[\pi(1y)]_t$	0.067	−0.012	0.134
	(0.05)	(0.08)	(0.09)
$dV_i[\pi(2y)]_t$	0.265***	0.417***	−0.231**
	(0.05)	(0.08)	(0.09)
$\|(\mathrm{MA}(\pi) - E_i[\pi(5y)])_{t-1}\|$	0.005	−0.014	0.026
	(0.01)	(0.02)	(0.02)
$\|\pi - E_i[\pi(1y)])_{t-4}\|$	0.003	0.004	0.001
	(0.01)	(0.02)	(0.02)
$\|\pi - E_i[\pi(2y)]_{t-8}\|$	−0.003	−0.005	0.004
	(0.01)	(0.01)	(0.02)
$dV_i[GDP(5y)]_t$	0.250***	0.282***	−0.077
	(0.04)	(0.05)	(0.07)
$dV_i[U(5y)]_t$	0.101***	0.074**	0.048
	(0.02)	(0.03)	(0.04)
$\|d\pi_{t-1}\|$	0.021***	0.032**	−0.015
	(0.01)	(0.01)	(0.01)
$\|\mathrm{dCBBS}_{t-1}\|$	−0.082*	−0.323**	0.296*
	(0.04)	(0.16)	(0.17)
MPA Dummy	0.020***	—	0.021***
	(0.01)	—	(0.01)
ELB Dummy	−0.004	—	−0.001
	(0.01)	—	(0.01)
Constant	−0.000	−0.003	—
	(0.00)	(0.00)	—
Observations	1,180	1,180	
R^2	0.388	0.402	

Notes: Dependent variable is the change in long-term inflation uncertainty. Both models include fixed effects for each forecaster. The constant is identified by restricting the average of the fixed effects to equal 0. We report the within R^2. Standard errors are computed using the method of Driscoll and Kraay (1998) and are robust against general forms of spatial and temporal dependence. *, **, and *** denote significance at the 10 percent, 5 percent, and 1 percent level, respectively.

inflation uncertainty. Thirdly, we find a strong and highly significant positive co-movement with the perceived uncertainty about long-term growth rates and the long-term unemployment outlook.

The panel regressions for uncertainty also reveal some important co-movement with the indicators linked to monetary policy. In the full-sample estimation, a higher absolute change in the volume of the assets held by the ECB tends to be associated with, on average, a reduction in long-term inflation uncertainty. However, when one considers the persistent rise in long-term uncertainty highlighted in section 3, the changes in ECB balance sheet in absolute terms did not fully insulate long-term inflation uncertainty from the other factors discussed above. Looking at the coefficients corresponding to the monetary policy dummies, we find that the announcement dates for nonstandard measures were generally associated with an increase in inflation uncertainty in the subsequent survey round, although this correlation is quantitatively less important than the above-mentioned downward effects.[28] On the one hand, this might indicate that the announcement of these monetary policy measures, as a side effect, led to a slight increase in long-term inflation uncertainty because forecasters had no historical experience on which to assess their transmission and the long-term implications for inflation. However, the co-movement may also reflect factors that the regression fails to control for and which simultaneously led to an increase in uncertainty and to the monetary policy announcements.

Looking at the second and third columns of table 3, we observe that many of the estimated coefficients do not change significantly after 2007:Q4. Apart from the coefficient on nonstandard monetary policy announcements, which is zero by definition in the pre-2007:Q4 sample, there are only two other individual parameters that exhibit a statistically significant change. First, the correlation with medium-term inflation uncertainty drops significantly after 2007. Second, the negative correlation with the absolute change in the balance sheet volume increases and is no longer significantly different from zero at a 5 percent significance level after 2007. Thus, similar to our results for mean expectations, we do not find evidence of a de-anchoring of the distribution that would be associated with an increase in the co-movement of long-term inflation uncertainty with other variables after the Great Recession. Nonetheless, the co-movement that we

[28]Coenen et al. (2017) have also recently studied the effectiveness of ECB communication during the period of unconventional policies. They show in particular how stock market uncertainty declined following several unconventional policy announcements but also that it rose following other announcements. These authors report evidence that the clarity and detail of the communication was an important factor for those occasions when it reduced uncertainty.

do identify highlights a number of potentially important forces that may be behind the rise in long-term inflation uncertainty that was highlighted in section 3.

Online appendix B analyzes in more detail how the chosen moment estimation technique may affect the inflation uncertainty regressions. It shows again that, with one or two exceptions, much of the significant co-movement is not dependent on a particular measure of the variance, although the value of estimates changes considerably in some cases (perhaps not surprisingly, given that many of the insignificant coefficients exhibit very large standard error). Across all moment estimation strategies considered, we observe significant co-movement with short-term inflation uncertainty and volatility in actual inflation, as well as with long-term growth uncertainty. Also, for all measures, we observe the positive effect on long-term inflation uncertainty associated with nonstandard policy announcements.[29]

5. Discussion and Conclusions

In this paper we have studied the key properties of the distribution of long-term inflation expectations in the euro area and the co-movement of key moments of this distribution with other variables. Our primary purpose has been to assess the extent to which the Great Recession and its aftermath, including the onset of a period in which the effective lower bound on nominal interest rates started to bind, led to any perceptible changes in this distribution. A word of caution is warranted, however, due to the fact that (i) our post–Great Recession sample is small and (ii) the reported probabilistic forecasts are often rounded, which raises the possibility of distortions linked to measurement error. Both factors make it difficult to detect structural breaks and to determine, particularly toward the end of the sample, whether breaks are temporary or permanent. However, our results appear robust across a number of different testing procedures and moment estimation methods. Our main findings add to the recent evidence provided in Autrup and Grothe (2014), Strohsal

[29]There are, however, some differences linked to the changes in the coefficients in the second half of the sample that are discussed further in online appendix B.

and Winkelmann (2015), and Speck (2016) and fall into three broad categories, which we discuss below.

First, and in contrast to most existing studies which have focused only on mean expectations or representative indicators extracted from financial markets, our analysis targets the entire subjective forecast distributions for long-term inflation by looking for breaks in the entire distribution, in interval probabilities, and in the first four moments. Hence, we can provide additional information about how long-term inflation uncertainty, the balance of long-term inflation risks, and the risk of extreme inflation events may have changed since the Great Recession. We find evidence for a change in the distribution around the time of the Great Recession as well as complementary evidence for significant breaks in each of the first three moments of the distribution. All of the latter evidence points toward a heightened risk of lower inflation outcomes. Importantly, we document a small downward shift in the mean long-term inflation expectations around 2013 soon after the intensification of the euro-area sovereign debt crisis. However, the most likely inflation outcome as represented by the mode of the distribution has been much more stable. Also, both the mean and the mode remain broadly aligned with the ECB's definition of its price stability objective of "below but close to 2.0 percent." In addition, however, our analysis of higher moments of the distribution points to a reduction in how tightly expectations are anchored at this level. For example, we document a substantial increase in uncertainty about long-term inflation prospects compared with the period prior to the Great Recession and also a tendency toward a negatively skewed long-term distribution with a higher probability mass attached to relatively low inflation outcomes. The finding of a negatively skewed distribution is precisely what is predicted by macroeconomic models which incorporate a lower bound constraint on nominal interest rates (see, for example, Coenen and Warne 2014 or Hills, Nakata, and Schmitt 2016).

Second, our study has uncovered substantial co-movement between the first two moments of the distribution of long-term inflation expectations and various macroeconomic indicators, including other expectations and indicators capturing the effects of monetary policy. Such co-movement implies that the process governing the distribution is far away from a simple stylized case where the inflation objective is a universal constant and where there is "blind faith" in the ability of the central bank to achieve this objective. Instead, the co-movement that we identify is in line with theories in which agents update their beliefs about long-term inflation in response to certain shocks (Orphanides and Williams 2004, 2007). For example, we find that persistent periods of lower-than-expected inflation are associated with a downward revision in long-term inflation expectations. In this sense, our results suggest that long-term inflation expectations are not completely forward looking. Ultimately, they tend to be influenced by the ex post historical track record of the central bank relative to its announced objective. Such results provide strong support for recent concerns about inflation remaining "too low for too long" (e.g., Draghi 2014) and the motivation behind unconventional monetary policies aimed at avoiding a persistent undershooting of the price stability objective. Regarding our central question, however, these empirical relationships existed also prior to the Great Recession and, overall, they do not appear to have strengthened in its aftermath.

Third, our analysis sheds light on how forecasters update their assessment of long-term inflation uncertainty in response to macroeconomic developments. Factors which influence this assessment include the volatility in recent inflation rates and perceptions of increased inflation uncertainty at shorter horizons. Also, calling to mind the correlated long-term risks associated with the prospect of secular stagnation discussed in Eggertson and Mehrotra (2014) and Summers (2014), our results suggest that longer-term uncertainty about growth and unemployment can spill over into increased uncertainty about long-term inflation. This empirical relationship may help explain the large upward shift in long-term inflation uncertainty in the euro area following the Great Recession. Concerning the role of recent nonstandard monetary policies, our sample is such that we must limit ourselves to an assessment of how inflation uncertainty changed after key monetary policy announcements. Once we control for other factors, we find that the announcement dates for nonstandard measures were generally followed by a modest increase in inflation uncertainty. Although this result must be interpreted with caution, it may highlight how such measures led to a slight increase in long-term inflation uncertainty because forecasters had no historical experience on which to assess their transmission and the long-term implications for inflation.

In conclusion, by focusing on the full distribution and by exploiting individual-level data, we have been able to make an innovative contribution to the empirical literature trying to understand the process governing the formation of long-term inflation expectations. Many of our findings are well predicted by macroeconomic theory stressing the effects of uncertainty about monetary policy and the implications of constraints such as the lower bound on nominal interest rates. However, a number of open questions remain for future research. These include an analysis of whether expectations of other agents such as private households or financial market participants show similar tendencies to the ones that we have identified. Also, further work is needed to jointly model monetary policy, the business cycle, and the formation of inflation expectations in ways which can help identify the causal mechanisms that may be behind our empirical results.

References

Andrews, D. W. K., and W. Ploberger. 1994. "Optimal Tests When a Nuisance Parameter Is Present Only under the Alternative." *Econometrica* 62 (6): 1383–1414.

Aruoba, S. B., and F. Schorfheide. 2015. "Inflation During and After the Zero Lower Bound." Paper prepared for the 2015 Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 27–29.

Autrup, S. L., and M. Grothe. 2014. "Economic Surprises and Inflation Expectations: Has Anchoring of Expectations Survived the Crisis?" ECB Working Paper No. 1671.

Bai, J., and P. Perron. 1998. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica* 66 (1): 47–78.

———. 2003. "Computation and Analysis of Multiple Structural Change Models." *Journal of Applied Econometrics* 18 (1): 1–22.

Ball, L., P. Gagnon, P. Honohan, and S. Krog. 2016. "What Else Can Central Banks Do?" Geneva Reports on the World Economy, No. 18, ICMB and CEPR.

Beechey, M. J., B. K. Johannsen, and A. T. Levin. 2011. "Are Long-Run Inflation Expectations Anchored More Firmly in the Euro Area Than in the United States?" *American Economic Journal: Macroeconomics* 3 (2): 104–29.

Benhabib, J., S. Schmitt-Grohé, and M. Uribe. 2001. “Monetary Policy and Multiple Equilibria.” *American Economic Review* 91 (1): 167–86.

Bodenstein, M., J. Hebden, and R. Nunes. 2012. “Imperfect Credibility and the Effective Lower Bound.” *Journal of Monetary Economics* 59 (2): 135–49.

Busetti, F., G. Ferrero, A. Gerali, and A. Locarno. 2014. “Deflationary Shocks and De-Anchoring of Inflation Expectations.” Occasional Paper No. 252 (November), Bank of Italy.

Ciccarelli, M. and J. A. García. 2015. “International Spillovers in Inflation Expectations.” ECB Working Paper No. 1857.

Coenen, G., M. Ehrmann, G. Gaballo, P. Hoffmann, A. Nakov, S. Nardelli, E. Persson, and G. Strasser. 2017. “Communication of Monetary Policy in Unconventional Times.” ECB Working Paper No. 2080 (June).

Coenen, G., and A. Warne. 2014. “Risks to Price Stability, the Effective Lower Bound, and Forward Guidance: A Real-Time Assessment.” *International Journal of Central Banking* 10 (2, June): 7–54.

Coibion, O., and Y. Gorodnichenko. 2012. “What Can Survey Forecasts Tell Us about Informational Rigidities?” *Journal of Political Economy* 120 (1): 116–59.

Demertzis, M., M. Marcellino, and N. Viegi. 2009. “Anchors for Inflation Expectations.” DNB Working Paper No. 229.

———. 2012. “A Credibility Proxy: Tracking US Monetary Developments.” *B.E. Journal of Macroeconomics* 12 (1): 1–36 (Article 12).

Dovern, J., U. Fritsche, P. Loungani, and N. Tamirisa. 2015. “Information Rigidities: Comparing Average and Individual Forecasts for a Large International Panel.” *International Journal of Forecasting* 31 (1): 144–54.

Dräger, L., and M. J. Lamla. 2013. “Anchoring of Consumers’ Inflation Expectations: Evidence from Microdata.” KOF Working Paper No. 13-339.

Draghi, M. 2014. “Monetary Policy in the Euro Area.” Speech at the Frankfurt European Banking Congress, Frankfurt am Main, November 21.

Driscoll, J., and A. Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *Review of Economics and Statistics* 80 (4): 549–60.

Eggertsson, G. B., and N. R. Mehrotra. 2014. "A Model of Secular Stagnation." NBER Working Paper No. 20574 (October).

Ehrmann, M. 2015. "Targeting Inflation from Below: How Do Inflation Expectations Behave?" *International Journal of Central Banking* 11 (Supplement 1): 213–49.

Engelberg, J., C. F. Manski, and J. Williams. 2009. "Comparing the Point Predictions and Subjective Probability Distributions of Professional Forecasters." *Journal of Business and Economic Statistics* 27 (1): 30–41.

European Central Bank. 2014. "Results of the Second Special Questionnaire for Participants in the ECB Survey of Professional Forecasters." https://www.ecb.europa.eu/stats/ecb_surveys/survey_of_professional_forecasters/pdf/ecb.spf201401_specialsurvey~357c66c297.en.pdf.

Galati, G., S. Poelhekke, and C. Zhou. 2011. "Did the Crisis Affect Inflation Expectations?" *International Journal of Central Banking* 7 (1, March): 167–207.

García, J. A., and A. Manzanares. 2007. "Reporting Biases and Survey Results: Evidence from European Professional Forecasters." ECB Working Paper No. 836.

Glas, A., and M. Hartmann. 2018. "Overconfidence versus Rounding in Survey-based Density Forecasts." Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3202810.

Gürkaynak, R. S., A. T. Levin, and E. Swanson. 2010. "Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., UK, and Sweden." *Journal of the European Economic Association* 8 (6): 1208–42.

Hills, T. S., T. Nakata, and S. Schmidt. 2016. "The Risky Steady State and the Interest Rate Lower Bound." ECB Working Paper No. 1913.

Inoue, A. 2001. "Testing for Distributional Change in Time Series." *Econometric Theory* 17 (1): 156–87.

Karadi, P. 2017. "The ECB's Announcements of Non-standard Measures and Longer-Term Inflation Expectations." Research Bulletin No. 33, European Central Bank.

Kenny, G., T. Kostka, and F. Masera. 2015. “Density Characteristics and Density Forecast Performance: A Panel Analysis.” *Empirical Economics* 48 (3): 1203–31.

Levin, A. T., F. M. Natalucci, and J. M. Piger. 2004. “The Macroeconomic Effects of Inflation Targeting.” *Review* (Federal Reserve Bank of St. Louis) 86 (4): 51–80.

Lyziak, T., and M. Paloviita. 2017. “Anchoring of Inflation Expectations in the Euro Area: Recent Evidence Based on Survey Data.” *European Journal of Political Economy* 46 (January): 52–73.

Manski, C. F., and F. Molinari. 2010. “Rounding Probabilistic Expectations in Surveys.” *Journal of Business and Economic Statistics* 28 (2): 219–31.

Mehrotra, A., and J. Yetman. 2014. “Decaying Expectations: What Inflation Forecasts Tell Us about the Anchoring of Inflation Expectations.” BIS Working Paper No. 464.

Nautz, D., and T. Strohsal. 2015. “Are US Inflation Expectations Re-Anchored?” *Economics Letters* 127 (February): 6–9.

Orphanides, A., and J. C. Williams. 2004. “Imperfect Knowledge, Inflation Expectations, and Monetary Policy.” In *The Inflation Targeting Debate,* ed. B. S. Bernanke and M. Woodford, 201–48 (chapter 6). Chicago: University of Chicago Press.

———. 2007. “Inflation Targeting under Imperfect Knowledge.” *Economic Review* (Federal Reserve Bank of San Francisco) 1–23.

Rossi, B., T. Sekhposyan, and M. Soupre. 2016. “Understanding the Sources of Macroeconomic Uncertainty.” CEPR Discussion Paper No. 11415.

Speck, C. 2016. “Inflation Anchoring in the Euro Area.” Discussion Paper No. 04/2016, Deutsche Bundesbank.

Strohsal, T., and L. Winkelmann 2015. “Assessing the Anchoring of Inflation Expectations.” *Journal of International Money and Finance* 50 (February): 33–48.

Summers, L. 2014. “U.S. Economic Prospects, Secular Stagnation, Hysteresis, and the Effective Lower Bound.” *Business Economics* 49 (2): 65–73.

van der Cruijsen, C., and M. Demertzis. 2011. “How Anchored Are Inflation Expectations in EMU Countries?” *Economic Modelling* 28 (1–2): 281–98.