Tail Co-movement in Inflation Expectations as an Indicator of Anchoring*

Filippo Natoli and Laura Sigalotti
Bank of Italy

We analyze the degree of anchoring of inflation expectations in the euro area during the post-crisis period, with a focus on the time span from 2014 onwards when long-term beliefs have substantially drifted away from the policy target. Using a new estimation technique, we look at tail co-movements between short- and long-term distributions of inflation expectations, estimated from daily quotes of inflation derivatives. We find that, during 2014, average correlations between short- and long-term inflation expectations rose sharply; moreover, negative tail events impacting short-term beliefs have been increasingly channeled to long-term views, triggering both downward revisions in expectations and upward changes in uncertainty. Overall, our results signal a risk of downside de-anchoring of long-term inflation expectations.

JEL Codes: C14, C58, E31, E44, G13.

1. Introduction

Headline inflation in the euro area has been falling since 2012 and became negative at the end of 2014. The five-year, five-year forward (5y5y forward) inflation swap rate, a commonly used indicator of medium- to long-term inflation expectations, has fallen well below

*We would like to thank David Altig, Antonio Antunes, Pietro Catte, Sara Cecchetti, Paolo Del Giovane, Andrea Finicelli, Giuseppe Grande, Marco Taboga, Jean-Paul Renne, Jelena Stapf, and the participants of the March 2015 Bank of Italy workshop on "Low Inflation and Its Implications for Monetary Policy" for their helpful comments and suggestions. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Italy. All the remaining errors are ours. Author e-mails: filippo.natoli@bancaditalia.it and laura.sigalotti@bancaditalia.it.

2 percent since September 2014; on September 4, the phrase *inflation expectations for the euro area over the medium to long term continue to be firmly anchored* was abandoned for the first time in the monetary policy statements of the European Central Bank. Inflation expectations persistently below the 2 percent level may suggest a loss of confidence in the ability of the monetary authority to achieve its inflation target; if inflation expectations are de-anchored, then consumption and investment could be postponed, leading to a deflationary spiral.

In this paper we analyze whether there has been a downside de-anchoring of long-term inflation expectations in the euro area after the global financial crisis. We propose a new method of assessing the degree of anchoring of expectations that overcomes weaknesses in standard techniques used to measure the *pass-through* of expectations. Looking at a set of indicators that includes the co-movement between extreme changes in long- and short-term market-based inflation expectations, we find that signs of downside de-anchoring have emerged in 2014.

Ordinary pass-through models of inflation expectations assume that when expectations are firmly anchored to the central bank's target, long-term expectations should be quite stable, not reacting (as short-term expectations do) to inflation or macro data releases; according to this claim, linear correlations of long-term expectations with short-term expectations provide insights on the sensitivity of long-term beliefs and might reveal possible signs of de-anchoring. However, linear correlations may not be sufficient to assess the degree of anchoring for two main reasons: first of all, they measure average correlations while, in times of falling expectations and interest rates close to the zero lower bound, concomitant downswings in short- and long-term expectations could be more informative than concomitant upswings to identify a de-anchoring below the target; secondly, they disregard the pass-through of uncertainty and, in general, variations in the entire distribution of long-term expectations.[1]

[1]Even when long-term expectations are relatively stable, the uncertainty around these expectations might be significantly volatile. In light of that, as pointed out by many authors, inflation targeting should help anchor market perceptions of the *entire* distribution of long-term expectations (e.g., Gürkaynak, Levin, and Swanson 2010).

On the grounds of the aforementioned elements, we investigate possible signs of de-anchoring by estimating the co-movement between changes in breakeven inflation (proxying risk-neutral inflation expectations) as well as between variations in the dispersion around these prices, at different maturities.

The whole investigation relies on inflation swaps and options (caps and floors). Without further identification assumptions, estimates based on their market quotes jointly represent inflation expectations and the inflation risk premium attached to them: while inflation risk premia are also informative, as well as expectations, in assessing the degree of anchoring, we choose not to extract risk premia and to loosely refer to breakeven inflation as inflation expectations throughout the analysis.[2]

The analysis is carried out in two steps: in the first one, we derive the option-implied probability distributions of future inflation one, two, three, five, seven, and ten years ahead on a daily basis in the period between October 2009 and February 2015 using the semi-nonparametric technique of Taboga (2016); secondly, we evaluate the co-movement between short- and long-term expectations and standard deviations using (i) linear correlations and (ii) tail correlations (i.e., correlations between extreme variations in short- and long-term expectations or long- and short-term standard deviations) based on the theory of copulas and on the TailCor estimator of Ricci and Veredas (2013).

We find that, during 2014, the co-movement between the short- and long-term distributions of expected inflation was asymmetric: on the one hand, long-term expectations reacted more to downswings than to upswings in short-term ones; on the other hand, the standard deviation of the distribution of long-term expectations reacted more to increases than to decreases in short-term uncertainty. This joint evidence suggests that signs of de-anchoring have emerged.

A growing literature investigating anchoring in the most recent period is reporting mixed results. Strohsal and Winkelmann (2015) find more firmly anchored inflation expectations in the euro area than in the United States as of 2011; moreover, estimates based on inflation swaps and options suggest only mild reactions of inflation

[2] A discussion of the role of risk premia in market-based estimates is provided in section 2.2.

beliefs to macroeconomic announcements during the crisis (Autrup and Grothe 2014) and post-crisis period (Scharnagl and Stapf 2015 and Speck 2016). Our result is in line with Ehrmann (2014), who studies the stability of long-term beliefs in a panel of countries before and after inflation targeting: under persistently low inflation, he finds that a sign of downside de-anchoring with respect to a target is that inflation expectations get revised down in response to lower-than-expected inflation but do not respond to higher-than-expected outturns. Our result is also coherent with the findings in Lyziak and Paloviita (2016) such that longer-term inflation forecasts have become somewhat more sensitive to shorter-term forecasts and to actual HICP inflation.

The most common method used to assess the degree of anchoring in one economy involves testing the sensitivity of inflation expectations to surprises in macro news (the *news-regression* approach of Gürkaynak, Sack, and Swanson 2005). While news regressions heavily depend on how the surprise component of each announcement is estimated, our approach is totally market based, so it is free from identification issues.[3] In addition to the newly designed technique to evaluating anchoring, we also contribute to the option literature by providing robust estimates of euro-area inflation densities and by applying copula-based and TailCor methods to option-implied data.

This paper is organized as follows. In section 2, we describe the data set of inflation swaps and options on euro-area inflation (section 2.1) and the derivation of option-implied distributions of future inflation (section 2.2). Section 3 presents the measures employed to assess the degree of anchoring (section 3.1) and describes the building blocks of the empirical analysis (section 3.2). In section 4 we present the estimates of the option-implied distributions of inflation and discuss the evolution of deflation probabilities, quantiles, means, standard deviations, and skewness for the available maturities (section 4.1); the measures of anchoring are computed on inflation swaps and standard deviations of the estimated distributions and results are discussed in section 4.2. Section 5 concludes.

[3]The link between tail co-movements and anchoring is first investigated in Antunes (2015), where coefficients of tail dependence between daily revisions of short- and long-term inflation swaps are constructed using different types of bivariate copulas.

2. Option-Implied Distributions of Future Inflation

2.1 Swaps and Options on Euro-Area Inflation

The market for inflation-linked derivatives has witnessed a considerable development in the past few years. The most popular inflation derivatives include inflation swaps and inflation options (caps and floors). An inflation swap is a derivative contract in which two parties agree to exchange a fixed amount of money with a floating amount linked to realized inflation on particular dates in the future. An inflation cap is a derivative contract in which the holder has the right to receive compensation payments at the end of each period in which the inflation rate exceeds an agreed-upon strike rate. The contract involves no obligations when the realized inflation is below the strike. In exchange for the contingent future payment, the holder pays a price (option premium) upfront. A floor is a derivative contract that gives the holder the right to receive payments at the end of each period in which the inflation rate falls below the predetermined strike. Inflation swaps, caps, and floors can be zero coupon or year-on-year. Zero-coupon contracts consist of a single compensation payment at maturity, while year-on-year ones include intermediate payments depending on the level of the inflation rate in each year of the reference period.

The underlying of quoted swaps and options is euro-area HICPxT, lagged by three months in order to be known at the maturity date of the option.[4] Bloomberg provides quotes for both zero-coupon and year-on-year swaps and options on euro-area inflation; for our purpose, we only rely on prices of zero-coupon contracts between October 2009 and February 2015.

The degree of liquidity of caps and floors is not easy to assess; according to Smith (2012), euro-area inflation option markets are more liquid than the U.K. and U.S. ones. Scharnagl and Stapf (2015), whose analysis is also based on zero-coupon options, check for their degree of liquidity by calculating put-call parities and show that no arbitrage violations arise for at-the-money options. In addition, it is worth mentioning that we adopt an estimation methodology (based

[4]Since the HICPxT is not observed daily, the fixed leg of an inflation swap contract over the same horizon, which is traded daily, is taken as a proxy.

on Taboga 2016) that is robust to outliers: pricing errors due to low liquidity should not have a significant impact on the results.

2.2 Extraction of Option-Implied Probability Distributions

The extraction of risk-neutral probability distributions from option quotes is based on the semi-nonparametric method developed in Taboga (2016). In what follows we briefly describe the estimation methodology; see appendix 1 for a more detailed description. The probability distribution of future inflation is assumed to have a discrete support;[5] then, in the absence of arbitrage opportunities there exists a finite set of positive state prices such that the price of any derivative contract on inflation can be expressed as a function of those state prices. Risk-neutral distributions can be simply obtained by rescaling once state prices are estimated. The method assumes that state prices are interpolated by a spline function, which is proved to be equivalent to a set of linear restrictions. The linearity of the problem allows for the derivation of computationally inexpensive estimators. In particular, a least absolute deviations (LAD) estimator can be obtained through a linear programming problem. In addition, this methodology allows for the incorporation of unimodality restrictions on the estimators of state prices. Unimodality of risk-neutral distributions obtained from state prices is a desirable property, and in the previous literature it was not dealt with.[6]

In addition to the computational convenience, the advantages of this methodology include its robustness to outliers, which are known to contaminate data on option prices. Despite the lack of information on the liquidity of option quotes, the robustness of the methodology supports confidence in the estimated state prices.

Throughout the paper, we need to bear in mind that probability distributions extracted from option quotes are risk neutral by assumption, i.e., they are not adjusted for investors' risk preferences.

---

[5]Assuming that the probability distribution of future inflation is discrete does not reduce significantly the scope of the methodology; in fact, continuous distributions can be arbitrarily approximated by discrete ones; moreover, most pricing algorithms require discretization at some stage; finally, market prices are inherently discrete.

[6]As explained in Taboga (2016), multimodality is often an artifact due to estimation procedures rather than an authentic feature of the data.

Risk-neutral distributions incorporate an inflation risk premium in addition to the expectation of future inflation, as well as a liquidity premium. Concerning the inflation risk premium, term structure estimates show that risk premia are significantly volatile, especially on long maturities.

However, inflation risk premia, whose identification gave mixed results in the literature in terms of magnitude and sign, are informative, as well as expectations, in assessing the degree of anchoring: indeed, the investors' willingness to pay large premia in order to protect themselves against a scenario of persistently low inflation would also signal high risks in terms of the central bank's credibility and ability to bring inflation back to target.[7] Bauer and Christensen (2014) point out that risk-neutral probabilities are useful for policy analysis, as policymakers are worried about extreme outcomes just like investors. As stated by Kocherlakota (2013), policy decisionmaking should take into account the evolution of risk-neutral probabilities, since it reflects changes in market participants' views about future possible outcomes.

For these reasons, we choose to not extract risk premia and consider breakeven inflation (and their probability distributions) throughout the analysis, loosely referring to those as inflation expectations.

3. Co-movement between Short- and Long-Term Moments

3.1 Measures of Average and Tail Co-movement

The co-movement between two random variables can be studied in various ways. One standard measure is the Pearson correlation, which estimates the average level of co-dependence. There are some important limits of such a measure: first of all, it does not take into account non-linear relationships between variables; secondly, it does not distinguish between different types of variations (large and

[7]The heterogeneity of the available estimates of inflation risk premia is highlighted also by Pericoli (2012), who provides a comparison of some estimates found in the literature and shows that there indeed are stark differences among them.

small, positive and negative), which are relevant in our assessment of de-anchoring.

In order to detect the co-dependence between extreme variations in short- and long-term expectations, we then turn to two different indicators of tail co-movement: the coefficient of conditional tail dependence, estimated through copulas, and the TailCor index.

3.1.1 Coefficient of Conditional Tail Dependence

The coefficients of conditional upper and lower tail dependence are defined as follows:[8]

DEFINITION 1. *Let X and Y be two random variables with marginal distributions $F_X(x)$ and $F_Y(y)$. Let x_k denote the k-th quantile of variable X and let y_k be the k-th quantile for Y. The conditional upper tail dependence is defined as*

$$\lambda_U = \lim_{k \to 1} Pr\{Y > y_k | X > x_k\};$$

the conditional lower tail dependence is defined as

$$\lambda_L = \lim_{k \to 0} Pr\{Y \leq y_k | X \leq x_k\}.$$

Intuitively, λ_U measures the asymptotic probability of having large outcomes in variable Y, conditional on observed large realizations for variable X. One way to compute the coefficients of conditional tail dependence is through copulas; copula functions are a special class of multivariate cumulative distribution functions that allow for separating the modeling of the marginal distributions from the dependence structure between the variables.[9] The advantage of using this approach is that for some choices of copula functions, the coefficient of conditional tail dependence can be retrieved in closed form.

[8] The coefficient of tail dependence was first introduced in the finance literature by Embrechts, Lindskog, and McNeil (2003).

[9] See appendix 2 for copula definition and main properties, and Nelsen (2006) for a detailed exposition of the theory and practical aspects of copulas.

Estimating the coefficients of conditional tail dependence through copulas involves (i) choosing the appropriate copula function; (ii) estimating the parameters that maximize the fit of the chosen copula to the data; and (iii) computing the coefficient of tail dependence as a function of the estimated parameters using the closed-form expression.

In particular, we use a Student's t-copula, which belongs to the class of elliptical distributions and displays symmetric tail dependence and potentially very heavy tails. The Student's t-copula is preferred to other copula distributions because of its good fit to inflation swap data in terms of log-likelihood, Akaike information criterion, and Bayesian information criterion (see Antunes 2015). In general, for elliptical distributions, $\lambda_U = \lambda_L$, and in particular, for the Student's t-copula, the coefficients of lower and upper tail dependence are

$$\lambda_U = \lambda_L = 2t_{\nu+1}\left(-\sqrt{\nu+1}\sqrt{\frac{1-\rho}{1+\rho}}\right),$$

where $t_{\nu+1}$ denotes the distribution function of a univariate Student's t-distribution with $\nu+1$ degrees of freedom and ρ is the linear correlation. The stronger the linear correlation ρ and the lower the degrees of freedom ν, the stronger is the tail dependence. However, the coefficient of tail dependence can be positive even if ρ is not.

The copula-based coefficient of conditional tail dependence is an asymptotic and parametric indicator; although its asymptotic nature could be a drawback when the estimation is performed in small samples, this measure is widely used in the literature and its quantitative interpretation is quite straightforward (it is defined as the limit of a conditional probability and hence takes values in the interval [0,1]).

3.1.2 TailCor Dependence Measure

As an alternative to the parametric tail dependence measure implied by copulas, we consider the TailCor index introduced by Ricci and Veredas (2013). This can be implemented under mild assumptions and presents several advantages: (i) it is non-parametric and independent of specific distributional assumptions; (ii) it performs well also in small samples, without relying on asymptotic theory; (iii) it allows to disentangle whether the evidence of tail correlations is

Figure 1. Diagrammatic Representation of TailCor

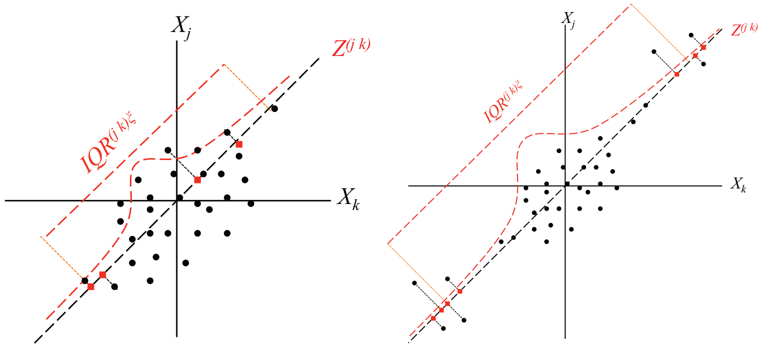

Source: Ricci and Veredas (2013, figure 1 on p. 34).
Notes: The figure shows scatter plots, along with the 45-degree line, where X_j and X_k are positively related (the pairs are depicted with circles). The left plot shows a linear relation, while the right plot shows a non-linear relation. Projecting the observations onto the 45-degree line produces the random variable $Z^{(jk)}$, depicted with squares. For illustrative purposes, the projection is shown only for the observations on the tails, but in the estimation it is done for all the observations.

caused by variables that are linearly correlated and/or non-linearly correlated; (iv) it is exact for any cutoff point of the tail; and (v) it can be computed for tails that are fatter than, equal to, or thinner than those of the Gaussian distribution.

In the following we briefly explain the intuition underlying this dependence measure using a graphical approach. The formal definition is provided in appendix 3, while technical details can be found in Ricci and Veredas (2013). The intuition underlying TailCor is that if two standardized random variables X_j and X_k are positively related (linearly and/or non-linearly), most of the time pairs of observations have the same sign. This means that looking at the scatter plot of the random variables (see figure 1), most of the pairs of observations (depicted with dots) concentrate in the northeast and southwest quadrants. When we project all the pairs on the 45-degree line, we get a new random variable $Z^{(jk)}$ (depicted with squares). The degree of dispersion of the projected dots depends on the strength of the relationship between the tails of the two random variables: if the relation is strong, the cloud is stretched around the 45-degree line and the projected dots are very dispersed. The TailCor measure is equal—up to a normalization—to the difference between upper and lower tail quantiles of $Z^{(jk)}$.

The TailCor index can be decomposed into the sum of two components, which disentangle the degree of dependence between lower tails (DownTailCor) and between upper tails (UpTailCor) of the distributions of the two variables. Using the notation introduced above, DownTailCor is proportional to the difference between the median of the projected variable $Z^{(jk)}$ and its lower tail quantile, while UpTailCor is proportional to the difference between the upper tail quantile of $Z^{(jk)}$ and its median.

Theoretically, the TailCor index takes values between 0 and infinity; however, the actual range of variation in most financial applications is very small. The fatter the tails of the bivariate distribution, the higher the exceedance of the largest attained value over $\sqrt{2}$, which is the largest value under a bivariate Gaussian.

3.2 Data Transformation

In this section we describe the empirical strategy used to measure the co-movement between short- and long-term distributions of future inflation. In particular, we investigate the co-dependence between expectations and between the dispersion around expectations at different horizons; expectations are proxied by spot and forward inflation swap rates, while the dispersion of market beliefs is proxied by the standard deviation of option-implied distributions extracted using the LAD method (see section 2.2).[10]

In order to compute the co-movement measures described in the previous section, we first apply a data transformation procedure to all time series of forward inflation swaps and option-implied standard deviations. We denote by $\{Y_t\}$ the generic time series of data. The procedure involves the following steps:

[10]Proxying expectations with forward inflation swaps allows for tracking of short- and long-term expectations on non-overlapping horizons (e.g., one year ahead after one year versus five years ahead after five years). In our estimate we want to rely only on market data—without making any assumption on the inflation process; hence we cannot estimate forward densities and, in particular, their standard deviations.

- We take the first difference in order to get daily variations: $X_t = Y_t - Y_{t-1}$.
- We filter the time series using an AR(1)-GARCH(1,1) model of the following form in order to eliminate persistence and heteroskedasticity, which could induce spurious dependence between variables:[11]

$$\begin{aligned} X_t &= \mu_t + \sigma_t \varepsilon_t, \quad \varepsilon_t \text{ iid} \sim N(0,1) \\ \mu_t &= \lambda X_{t-1} \\ \sigma_t^2 &= a_0 + a(X_{t-1} - \mu_{t-1})^2 + b\sigma_{t-1}^2; \end{aligned}$$

  we denote the filtered daily revisions by $\{x_t\}$.
- We map $\{x_t\}$ into numbers between 0 and 1 through their empirical marginal cumulative distribution function $\tilde{F}_x$.[12] The standardized time series is denoted by $\{u_t\}$.

The resulting filtered and standardized daily changes are used to compute our co-movement indicators.

4. Results

The data set includes daily closing quotes of zero-coupon inflation swaps and options from the first available trading day, i.e., October 5, 2009, until February 18, 2015 (source: Bloomberg). We consider swaps and options with maturity equal to one, two, three, five, seven, and ten years. Concerning options, we use strike rates ranging from 1 to 6 percent for caps and strikes from −2 to 3 percent for floors. Forward inflation swaps such as the 1y1y, the 1y2y, and the 5y5y swaps are computed from quoted spot rates.

4.1 Option-Implied Distributions

For each date in the sample and each maturity horizon, we extract the unimodal LAD estimator of state prices and derive the corresponding risk-neutral distribution, thus getting a time series of

[11] A similar approach is adopted by Christoffersen et al. (2012) to capture dynamic dependence across equity markets.

[12] This transformation, which is equivalent to a ranking, is required for copula estimation.

Figure 2. Option-Implied Risk-Neutral Distributions of Annual Euro-Area Inflation over a Ten-Year Horizon

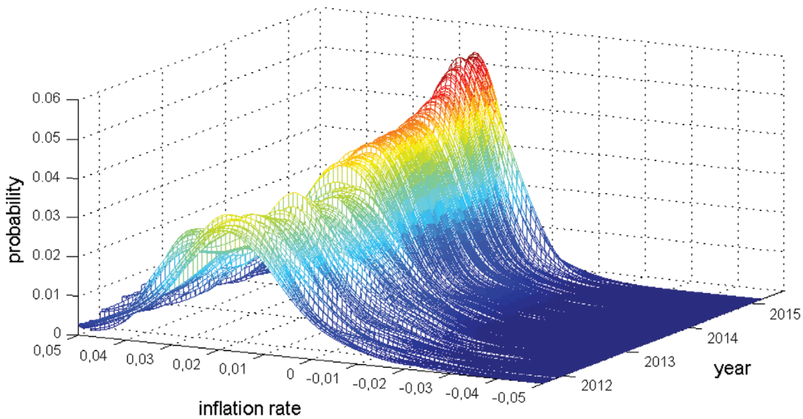

Notes: Option-implied distributions are extracted using the LAD estimator with unimodal restrictions from daily quotes between September 2011 and February 2015. The x-axis corresponds to the annual inflation rate (percentage points); the y-axis indicates the time interval (days).

implied distributions. For instance, figure 2 shows the time evolution of risk-neutral distributions of inflation on a ten-year horizon, as extracted from options data in the period September 2011–February 2015. The plot highlights the tendency of the distributions to become more and more concentrated over time, as well as a shift of the mean towards lower inflation rates.

Figure 3 shows the mean of the option-implied distributions for maturities of one, two, three, five, seven, and ten years. Inflation expectations, as proxied by the expected value of option-implied distributions, have been decreasing since 2012 for all maturities, with sharper falls for shorter horizons. The contraction of investors' beliefs halted around mid-January 2015 for all horizons.

Appendix 1 proves that for a one-year maturity the expected value of the implied distribution coincides with the fixed leg of an inflation swap having the same maturity. Comparing the time series of expected values derived from our estimates with quoted inflation rates, we obtain a very accurate match. For maturities longer than one year, the quoted inflation swap rate must be equal to a non-linear function of the implied distribution: this is confirmed

Figure 3. Means of Option-Implied Risk-Neutral Inflation Distributions (percentage points)

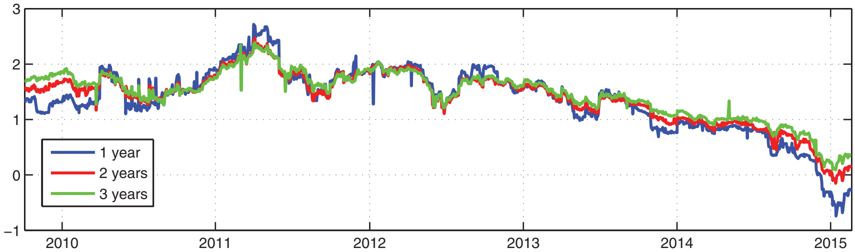

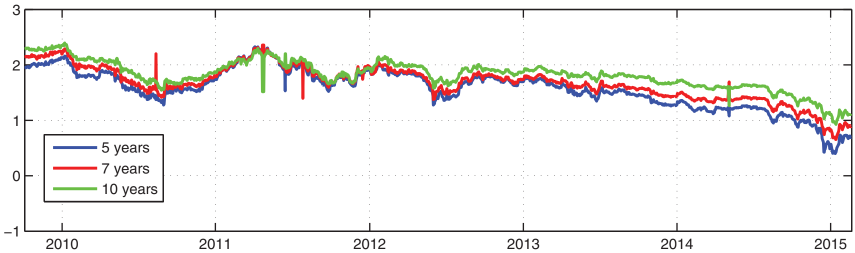

Notes: Option-implied distributions are extracted using the LAD estimator with unimodality restriction. Daily quotes of inflation caps and floors from October 2009 to February 2015 are taken from Bloomberg.

by the results of our estimates. Figure 4 shows that the difference between the quoted inflation swap rate (red line) and the one implied by our probability distributions (blue line) is negligible for all maturities.[13] Although the estimation methodology we adopt does not force this matching through a constraint, we still recover it in quoted prices: this confirms the robustness and reliability of the approach.

Figure 5 shows the evolution of the standard deviation of option-implied distributions over time. This gives insights on the degree of uncertainty in market expectations of future inflation: the higher the standard deviation, the more dispersed are investors' beliefs and/or the more difficult it is to forecast inflation. The figure shows that uncertainty has been decreasing since 2012 for all maturities, like option-implied means. In general, the lowering of the standard

[13]Colors are shown in the online version available at http://www.ijcb.org.

Figure 4. Market Quotes of Inflation Swap Rates (Red Line) and Inflation Swap Rates Implied by the Probability Distributions Embedded in Option Prices (Blue Line)

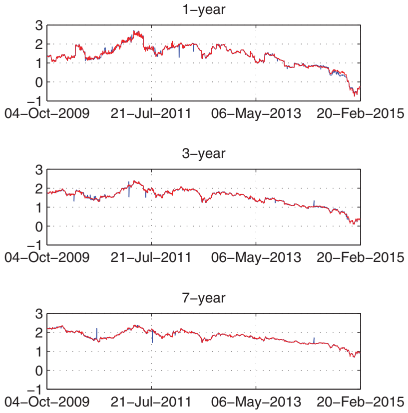

Notes: Market quotes and inflation swap rates are shown at one-, two-, three-, five-, seven-, and ten-year maturities. Option-implied distributions are extracted using the LAD estimator with unimodal restrictions. Daily quotes of inflation swaps and inflation options from October 2009 to February 2015 are taken from Bloomberg.

deviation of inflation distributions does not have a univocal interpretation: if expectations are far from the target, it indicates a higher concentration of beliefs around an undesirable outcome. In the context of long-term inflation expectations falling below the target, the attenuation of uncertainty around those expectations can then be seen as an indicator of diminished credibility of monetary policy.

Figure 6 shows that the skewness of option-implied risk-neutral inflation expectations for short horizons (one, two, and three years) became negative in the past few months after a gradual decline. For unimodal distributions such as the ones we estimated, a negative skewness indicates that the lower tail is fatter or longer than the upper tail; the most recent developments of this indicator for horizons up to three years point towards a predominance of the left tail, suggesting that market views are unbalanced towards negative inflation outcomes. Concerning maturities of five, seven, and ten years, even though the skewness has decreased from mid-2013 until January 2015, it has always remained positive.

Figure 5. Standard Deviations of Option-Implied Risk-Neutral Inflation Distributions (percentage points)

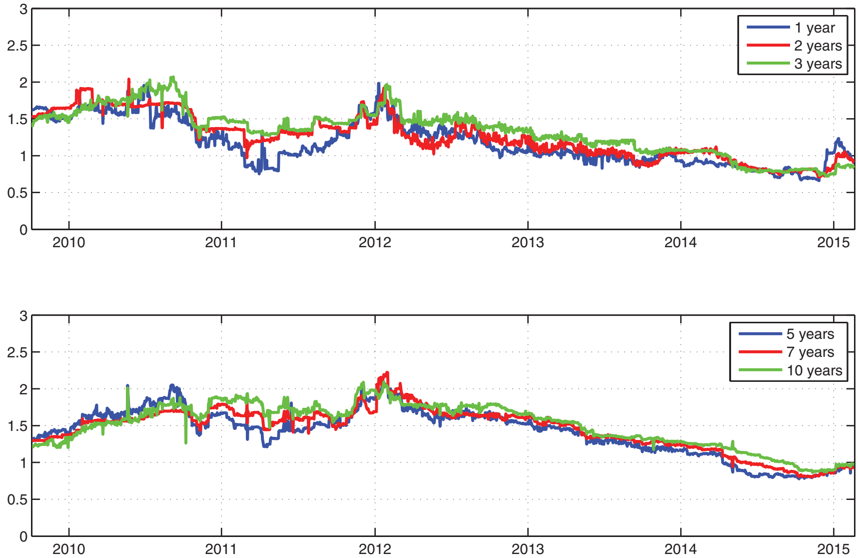

Notes: The upper panel shows maturities of one, two, and three years; the lower panel shows maturities of five, seven, and ten years. Option-implied distributions are extracted using the LAD estimator with unimodality restriction. Daily quotes of inflation caps and floors from October 2009 to February 2015 are taken from Bloomberg.

Figures 3, 5, and 6 highlight that a notable change in market-based inflation expectations has occurred, starting around mid-January 2015: expected values of inflation went up, standard deviations had a small rebound and skewness increased, especially for longer maturities. This can be interpreted as an effect of market agents anticipating the announcement of the quantitative easing program by the European Central Bank (January 22, 2015).

Figure 7 shows the risk-neutral probabilities of the average annual inflation rate over different time horizons falling below zero. Deflation probabilities at all maturities increased sharply in the last few months of 2014; for maturities up to three years, the rise started earlier—in the last quarter of 2013. Around mid-January 2015, the increase halted and deflation probabilities decreased for all time horizons. On the other hand, the probability that the average annual

Figure 6. Skewnesses of Option-Implied Risk-Neutral Inflation Distributions at Different Maturities (percentage points)

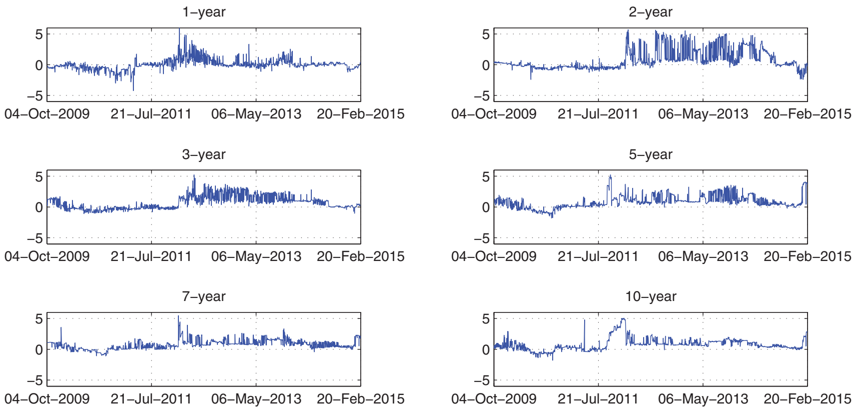

Notes: One-, two-, three-, five-, seven-, and ten-year maturities are shown. Option-implied distributions are extracted using the LAD estimator with unimodality restriction. Daily quotes of inflation caps and floors from October 2009 to February 2015 are taken from Bloomberg.

inflation rate at time different maturities falls between 1.5 and 2 percent shrank during 2014 and rebounded in early 2015 (figure 8).

Having estimated option-implied distributions, we can calculate confidence bands around the mean of expected future inflation. This allows to assess the significance of the decline in short- and long-term inflation expectations observed since 2012. Figure 9 shows the confidence bands for the expected value of option-implied probability distributions of future inflation using the 5th and 95th percentiles of the distributions, over one-, five-, seven-, and ten-year horizons. The upper limit of the confidence band fell below 2 percent for maturities up to seven years; this can be interpreted as the negative gap between expectations and the 2 percent rate being statistically significant at the 10 percent level.[14]

[14]In figure 9 we compare the level of inflation expectations with the 2 percent reference level for illustrative purposes, even though the policy objective of the European Central Bank entails the inflation rate being below, but close to, 2 percent over the medium term.

Figure 7. Risk-Neutral Probability that the Average Annual Inflation Rate at Different Maturities Is Negative

Notes: One-, two-, three-, five-, seven-, and ten-year maturities are shown. Option-implied distributions are extracted using the LAD estimator with unimodality restriction. Daily quotes of inflation caps and floors from October 2009 to February 2015 are taken from Bloomberg.

4.2 Anchoring of Long-Term Expectations

In order to construct our indicators, we select a set of measures for short- and long-term expectations from forward inflation swaps and option-implied standard deviations. Short-term expectations are proxied by the one-year spot, the 1y1y forward, and the five-year spot rates, while long-term ones are 5y5y forward inflation swaps; concerning dispersions, we compare one-year versus seven-year, one-year versus ten-year, and two-year versus ten-year standard deviations. Summary statistics for levels, daily changes, filtered values, and mapped values of each variable are reported in table 1. Variables in levels, both expectations and the dispersion of expectations, are all strongly persistent: high autocorrelations of variables in first differences are removed by the AR-GARCH filtering.

Figure 8. Risk-Neutral Probability that the Average Annual Inflation Rate at Different Maturities Falls between 1.5 and 2 Percent

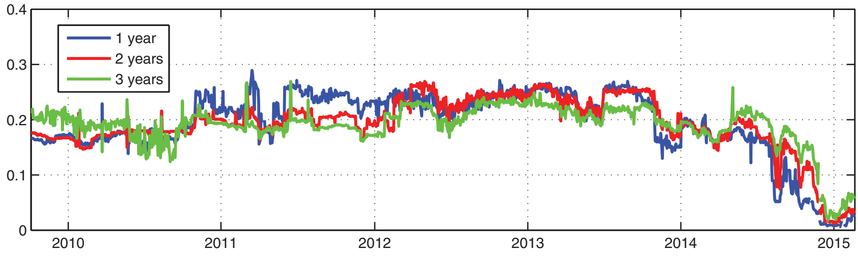

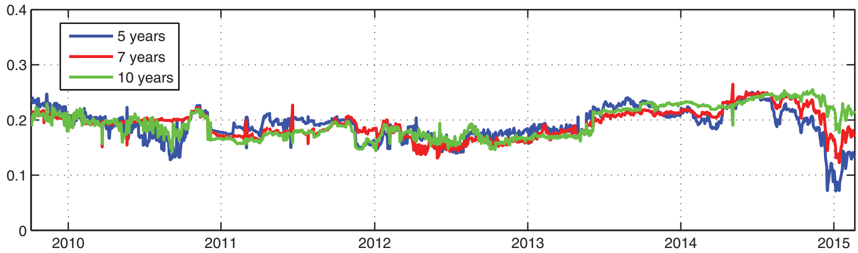

Notes: One-, two-, three-, five-, seven-, and ten-year maturities are shown. Option-implied distributions are extracted using the LAD estimator with unimodality restriction. Daily quotes of inflation caps and floors from October 2009 to February 2015 are taken from Bloomberg.

We report results based on our three measures of co-movement: the Pearson correlation coefficient (average co-movement); the coefficient of tail co-movement estimated with the Student's t static bivariate copula and the TailCor index (two measures of average co-movement in the tails); and the UpTailCor and DownTailCor indexes which track co-movements over time in upper and lower tails. Every statistic is computed using rolling windows of 250 business days of observations; nonetheless, the conclusions we draw are robust to different window lengths.

Results for mean expectations are depicted in figures 10–13. The third panel in figure 10 shows a decline in average correlations between five-year and 5y5y expectations during 2013 and the first half of 2014; a steady increase is then evident from the end of July

Figure 9. Confidence Bands for the Mean of the Option-Implied Probability Distributions of Future Inflation

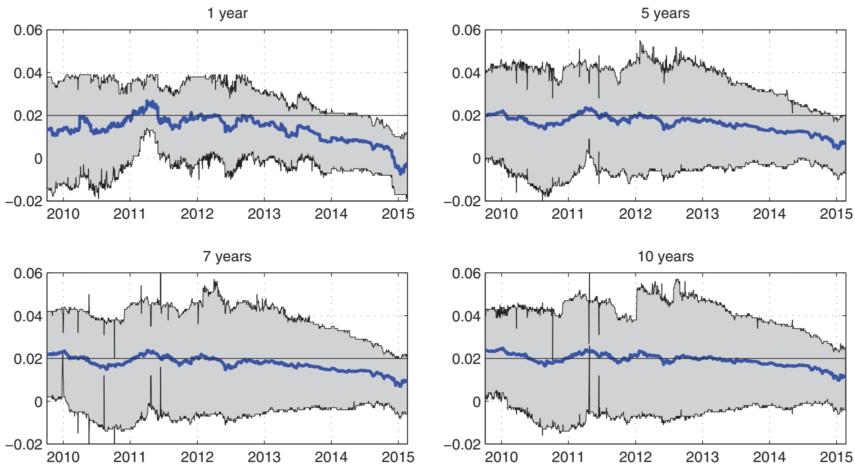

Notes: Confidence bands at the 10 percent level and at one-, five-, seven-, and ten-year maturities are shown.

2014 up to levels close to 60 percent. A similar upward trend for linear correlations is observed in the same period using the other proxies of short-term expectations (first two panels).

These results highlight that an increase in average correlations has happened since the second half of 2014. To investigate further signs of de-anchoring, we look at the path of the copula-based coefficient of tail dependence and at the one of the TailCor index: while the first delivers mixed results (figure 11), the TailCor index suggests that the observed increase in average correlations reflects, at least in part, an increased correlation in the tails (figure 12). To detect possible asymmetries across correlations between left tails with respect to those between right tails, we look separately at each tail using the decomposition of TailCor. Figure 13 depicts the dynamics of the UpTailCor (upper-tail correlation; blue line) and DownTailCor (lower-tail correlation; red line); all panels show that the DownTailCor index increases between early 2014 and early 2015, while the UpTailCor started to rise only towards the end of 2014. This

Table 1. Summary Statistics

	Obs.	Mean	Std. Dev.	Autocorr.
A. Mean Expectations				
π_{1y}	1,404	1.36	0.60	1.00
$\Delta\pi_{1y}$	1,403	−0.00	0.06	−0.12
x_{1y}	1,403	−0.00	1.00	−0.01
u_{1y}	1,403	0.50	0.29	−0.02
π_{1y1y}	1,404	1.47	0.41	0.99
$\Delta\pi_{1y1y}$	1,403	−0.00	0.06	−0.33
x_{1y1y}	1,403	−0.00	1.00	−0.05
u_{1y}	1,403	0.50	0.29	0.00
π_{5y}	1,404	1.62	0.37	1.00
$\Delta\pi_{5y}$	1,403	0.00	0.03	0.01
x_{5y}	1,403	−0.01	1.00	0.00
u_{5y}	1,403	0.50	0.29	0.02
π_{5y5y}	1,404	2.24	0.22	0.99
$\Delta\pi_{5y5y}$	1,403	−0.00	0.03	−0.16
x_{5y5y}	1,403	−0.00	1.00	0.01
u_{5y5y}	1,403	0.50	0.29	0.05
B. Option-Implied Standard Deviations				
σ_{1y}	1404	1.26	1.01	0.61
$\Delta\sigma_{1y}$	1403	−0.00	0.89	−0.20
x^{sd}_{1y}	1403	0.00	1.01	−0.03
u^{sd}_{1y}	1403	0.50	0.29	0.01
σ_{2y}	1404	1.30	0.74	0.74
$\Delta\sigma_{2y}$	1403	−0.00	0.54	−0.26
x^{sd}_{2y}	1403	−0.00	1.00	−0.03
u^{sd}_{2y}	1403	0.50	0.29	−0.01
σ_{7y}	1404	1.46	0.39	0.80
$\Delta\sigma_{7y}$	1403	−0.00	0.24	−0.01
x^{sd}_{7y}	1403	0.01	1.01	0.02
u^{sd}_{7y}	1403	0.50	0.29	0.10
σ_{10y}	1404	1.51	0.30	0.98
$\Delta\sigma_{10y}$	1403	−0.00	0.06	−0.21
x^{sd}_{10y}	1403	0.04	1.00	−0.03
u^{sd}_{10y}	1403	0.50	0.29	0.06

Figure 10. Pearson Correlation Coefficient on Short-Term versus Medium- to Long-Term Market-Based Inflation Expectations

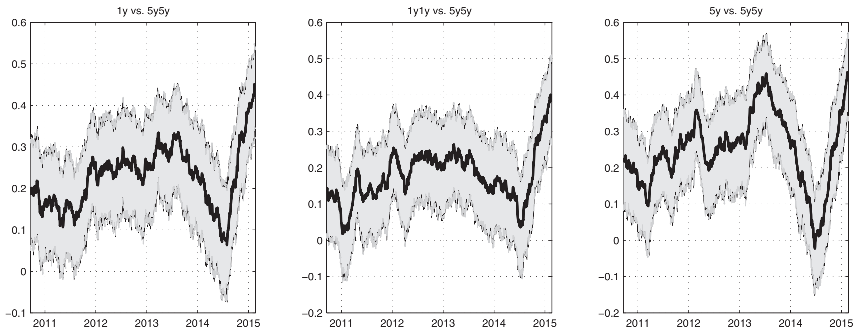

Notes: Short-term mean expectations are one year ahead, one year ahead after one year (1y1y forward inflation swap) and five years ahead, while medium- to long-term expectations are five years ahead after five years (5y5y forward). The coefficient is computed using 250 business days rolling windows. Bootstrapped confidence intervals are obtained with 1,000 replications. Sample: October 5, 2009 to February 19, 2015.

Figure 11. Index of Tail Co-movement Using the Student's *t*-Copula on Short-Term versus Medium- to Long-Term Mean Inflation Expectations

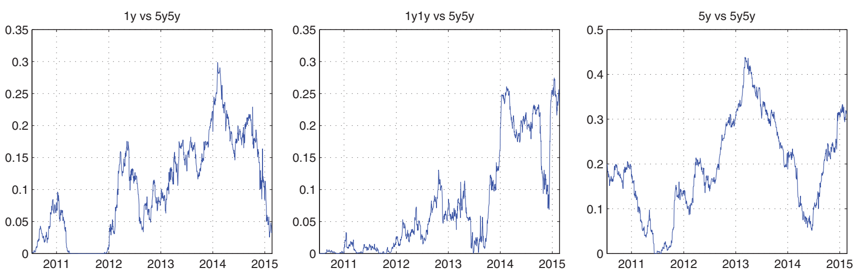

Notes: The index ranges from 0 (no tail dependence) to 1. This index indicates the average co-movement on both upper and lower tails. Short-term mean expectations are one year ahead, one year ahead after one year (1y1y forward inflation swap) and five years ahead, while medium- to long-term expectations are five years ahead after five years (5y5y forward). Values are computed using 200 business days rolling windows. Sample: October 5, 2009 to February 19, 2015.

Figure 12. TailCor Index Computed on Short-Term versus Medium- to Long-Term Mean Inflation Expectations

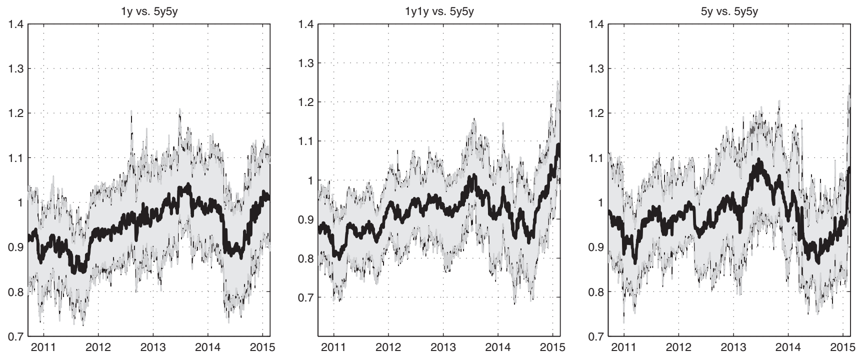

Notes: It takes values between 0 and $+\infty$; under Gaussianity and uncorrelation, the index takes the value 1. This measure indicates the average co-movement in both upper and lower tails. Short-term mean expectations are one year ahead, one year ahead after one year (1y1y forward inflation swap) and five years ahead, while medium- to long-term expectations are five years ahead after five years (5y5y forward). Values are computed using 250 business days rolling windows. Bootstrapped confidence intervals are obtained with 1,000 replications. Sample: October 5, 2009 to February 19, 2015.

evidence suggests that the correlation in the lower tail has increased earlier than the one in the upper tail: according to our interpretation, in 2014 negative events affecting short-term views have been transmitted to long-term expectations more than positive surprises. This stylized fact suggests that some de-anchoring may be occurring, and further investigation is needed.

Figures 14–17 trace the co-movement in option-implied standard deviations. Based on the rolling estimates of the Pearson coefficient, there is no clear evidence of increased correlation between short- and long-run standard deviations during the last part of the sample (figure 14). Interpreting the standard deviation as the uncertainty of market agents around their mean expectations, this points to mixed evidence on the transmission of average uncertainty towards longer maturities.

However, the analysis of tail co-movements gives interesting results. While the dynamics of the TailCor index for short- versus long-term standard deviations in the last period is not robust

Figure 13. UpTailCor (Blue Line) and DownTailCor (Red Line) Computed between Short-Term and Medium- to Long-Term Mean Expectations

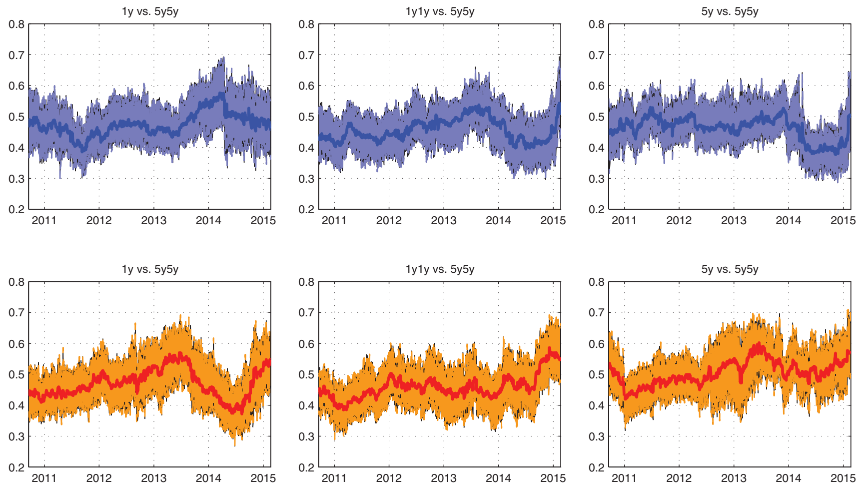

Notes: Short-term mean expectations are one year ahead, one year ahead after one year (1y1y forward inflation swap) and five years ahead, while medium- to long-term expectations are five years ahead after five years (5y5y forward). Values are computed using 250 business days rolling windows; $\xi = 0.85, \tau = 0.75$. Bootstrapped confidence intervals are obtained with 1,000 replications. Sample: October 5, 2009 to February 19, 2015.

to the choice of the employed proxies (see figure 16), the copula-based coefficient of tail dependence suggests a strong recent increase in tail co-movements (e.g., 1y-10y and 2y-10y couples—second and third panels of figure 15). Moreover, the inspection of UpTailCor and DownTailCor leads to opposite results with respect to the one obtained about mean expectations: the UpTailCor increases more than the DownTailCor during most of 2014. This suggests a stronger transmission from upper tail variations of short-term standard deviation to upper tail variations in long-term ones, implying that positive shocks to uncertainty have been longer lasting than shocks reducing uncertainty.

To sum up, the joint reading of the results for average expectations and option-implied standard deviations gives some useful

Figure 14. Pearson Correlation Coefficient on Standard Deviations of Short-Term versus Long-Term Option-Implied Distributions

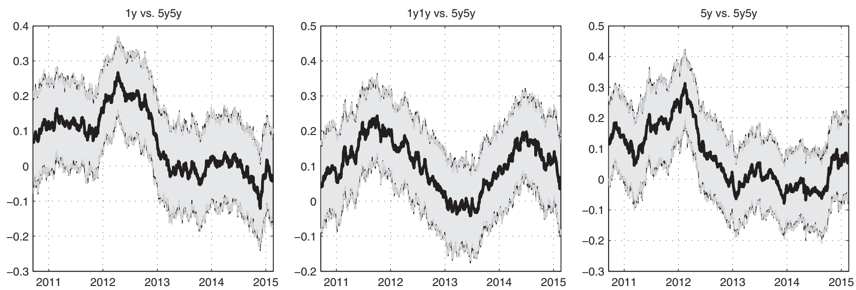

Notes: Pairs of short-term versus medium- to long-term inflation expectations are one year ahead versus seven years ahead, one year ahead versus ten years ahead, and two years ahead versus ten years ahead. The coefficient is computed using 250 business days rolling windows. Bootstrapped confidence intervals are obtained with 1,000 replications. Sample: October 5, 2009 to February 19, 2015.

Figure 15. Index of Tail Co-movement Using the Student's t-Copula on Standard Deviations of Short-Term versus Long-Term Option-Implied Distributions

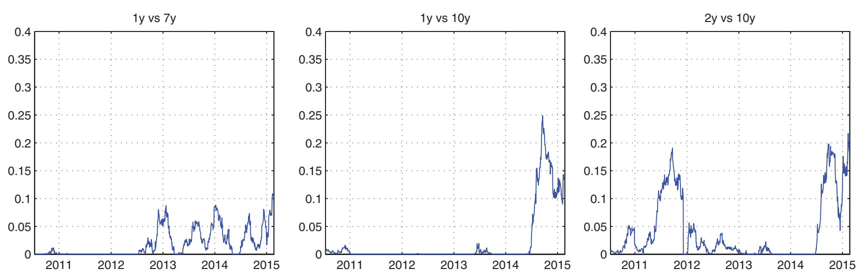

Notes: The index ranges from 0 (no tail dependence) to 1. This index indicates the average co-movement on both upper and lower tails. Pairs of short-term versus medium- to long-term inflation expectations are one year ahead versus seven years ahead, one year ahead versus ten years ahead, and two years ahead versus ten years ahead. Sample: October 5, 2009 to February 19, 2015.

insights. During 2014, average correlations between short- and long-term inflation expectations have increased sharply; moreover, downswings in short-term expectations, possibly reflecting bad macro news or worse-than-expected data readings, have been increasingly

Figure 16. TailCor Index Computed on Short-Term versus Medium- to Long-Term Mean Inflation Expectations

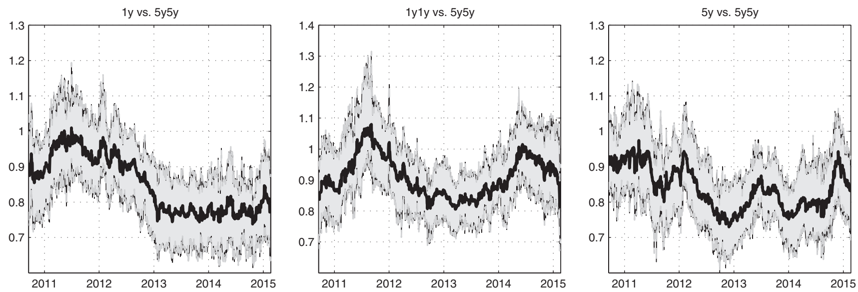

Notes: This measure indicates the average co-movement on both upper and lower tails. Pairs of short-term versus medium- to long-term inflation expectations are one year ahead versus seven years ahead, one year ahead versus ten years ahead, and two years ahead versus ten years ahead. Values are computed using 250 business days rolling windows. Bootstrapped confidence intervals are obtained with 1,000 replications. Sample: October 5, 2009 to February 19, 2015.

Figure 17. UpTailCor (Blue Line) and DownTailCor (Red Line) Computed between Short-Term and Medium- to Long-Term Mean Inflation Expectations

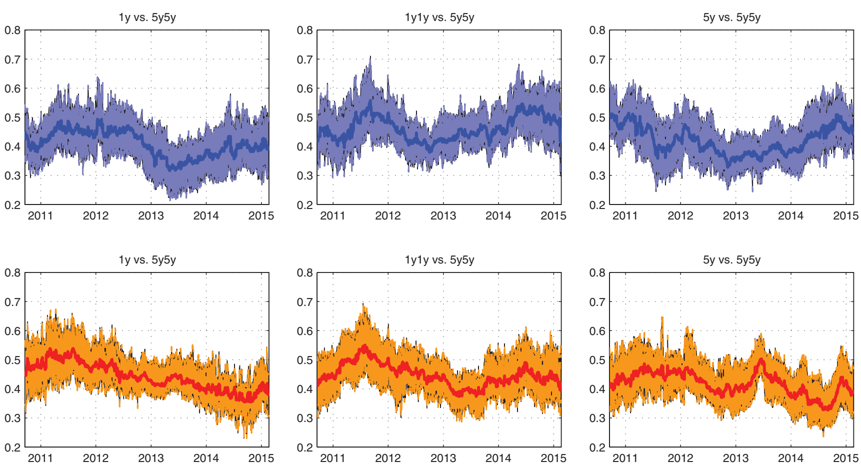

Notes: Pairs of short-term versus medium- to long-term inflation expectations are one year ahead versus seven years ahead, one year ahead versus ten years ahead, and two years ahead versus ten years ahead. Values are computed using 250 business days rolling windows; $\xi = 0.85, \tau = 0.75$. Bootstrapped confidence intervals are obtained with 1,000 replications. Sample: October 5, 2009 to February 19, 2015.

channeled to long-term views, igniting downward revisions in average expectations and upward revisions in uncertainty. These results point to a risk of downside de-anchoring of long-term inflation expectations from the "below but close to" 2 percent target.

5. Conclusions

In this paper we propose a new method to detect possible signs of de-anchoring of inflation expectations from the medium- to long-term objective of the monetary authority. Like the commonly used pass-through approach, our technique is totally market based and does not require any identification of the surprise component incorporated in inflation readings and other macroeconomic announcements. By looking at co-movements in the tails, we assess the sensitivity of long-term expectations to extreme shocks hitting short-term ones, both positive and negative. Although a departure of long-term expectations from the monetary policy target can occur also with stable short-term views, a high degree of co-movement with short-term expectations can be seen as a sufficient condition for de-anchoring.

Applying the new estimation technique of Taboga (2016) to daily quotes of inflation caps and floors for the euro area, we are able to recover the entire probability distributions assigned by market participants to future inflation at different horizons. Focusing on mean expectations and on the dispersion of market beliefs, we look at their evolution over time, making comparisons between the long- and short-end of each term structure. In addition, we calculate confidence bands around the mean of expected future inflation in order to assess the significance of the decline in short- and long-term expectations observed since 2012. Using inflation swaps and option implied-standard deviations, we also compute linear and tail correlations between short-and long-term expectations and dispersions around them; tail co-movements are measured by copula-based coefficients of tail dependence and by the TailCor indexes.

Computing confidence bands from the estimated distributions, we find that, at the end of 2014, the upper limit of such bands fell below 2 percent for maturities up to seven years, indicating that the negative gap between expectations and the 2 percent inflation rate was statistically significant at the 10 percent level at the end of our sample period. During 2014, average correlations between short- and long-term inflation expectations have increased sharply; moreover, downswings in short-term expectations, possibly reflecting bad macro news or worse-than-expected data readings, have been increasingly channeled to long-term views, igniting downward revisions in average expectations and upward revisions in uncertainty. Taken together, the evidence based on confidence bands and correlations leads one to conclude that some signs of downside de-anchoring of long-term inflation expectations from the 2 percent target are there and should not be overlooked.

While market-based measures of long-term inflation expectations have been unusually low in the euro area since 2013, during 2014 they also declined in the United States and in the United Kingdom.[15] Even though broader international trends may in part be responsible for this common pattern, the level of anchoring of inflation expectations could be different in the three economies. Further avenues of research could include extending the estimation of option-implied distributions to U.S. and U.K. data and assessing tail co-movements within countries.

Appendix 1. LAD Method

Extraction of Option-Implied Distributions

In what follows, a quick description of the estimation method to derive option-implied probability distribution functions elaborated in Taboga (2016) is given. Let $\mathcal{I}$ be the stochastic value of the average annual inflation rate over a given time horizon. We assume that $\mathcal{I}$ has a discrete probability distribution with finite support, $R_{\mathcal{I}} = \{i_1, \dots, i_n\}$. In the absence of arbitrage, there exists an n-dimensional vector of positive state prices $\pi = (\pi_1, \dots, \pi_n)$ such that the price $\Pi(f)$ of any derivative contract on $\mathcal{I}$ having payoff $f = f(\mathcal{I})$ can be written as

$$\Pi(f) = \sum_{j=1}^{n} \pi_j f(i_j).$$

[15] Ciccarelli and Garcia (2015) investigate possible spillovers of inflation expectations across countries, finding substantial spillovers from euro-area long-term expectations onto international ones—in particular, U.S. ones—since August 2014.

In particular, the price of a zero-coupon cap with strike k and maturity $T = M$ years is given by

$$\sum_{j=1}^{n} \pi_j((1+i_j)^M - (1+k)^M)^+,$$

whereas the price of a zero-coupon floor with strike k and maturity $T = M$ years equals

$$\sum_{j=1}^{n} \pi_j((1+k)^M - (1+i_j)^M)^+.$$

Suppose we have the market quotes of N_C caps with strikes $\{k_1^C, \ldots, k_{N_C}^C\}$ and N_P floors with strikes $\{k_1^P, \ldots, k_{N_P}^P\}$. Let C be the $N_C \times 1$ vector of cap quotes and P be the $N_P \times 1$ vector of put quotes. Let F_C and F_P be $N_C \times n$ and $N_P \times n$ matrices of payoffs, defined as

$$F_{C,ij} = (s_j - K_i^C)^+ \ \text{ and } F_{P,ij} = (K_i^P - s_j)^+.$$

Having set

$$Y = \begin{bmatrix} C \\ P \end{bmatrix} \text{ and } X = \begin{bmatrix} F_C \\ F_P \end{bmatrix},$$

we can express the option prices as

$$Y = X\pi. \tag{1}$$

Since in practice market quotes encompass an error term,[16] the empirical version of equation (1) is

$$Y^0 = X\pi + \varepsilon, \tag{2}$$

where Y^0 is the vector of observed market prices and ε is a vector of pricing errors. Our goal is to estimate the vector of positive state prices π given that we observe Y^0 and we know the payoff

[16] Pricing errors can arise for various reasons, including the bounce between bid and ask quotes, price discreteness, and state prices due to illiquidity.

X; the risk-neutral probability distribution of $\mathcal{I}$ is then obtained by rescaling,

$$\rho = \frac{\pi}{\sum_{j=1}^{n} \pi_j}. \tag{3}$$

State prices are parameterized using a spline curve; Taboga (2016) shows that this is equivalent to imposing a set of linear equality restrictions. With no loss of generality, we assume that the support $R_{\mathcal{I}}$ of the distribution is equally spaced:

$$i_j = i_1 + (j-1)\delta, \quad \delta > 0 \text{ and } j = 1, \ldots, n.$$

Moreover, we assume that there exists a (piecewise cubic and twice continuously differentiable) spline function $\overline{\pi} : [s_1, s_n] \mapsto \mathbb{R}_+$ which interpolates the state prices:

$$\overline{\pi}(s_j) = \pi_j, \quad j = 1, \ldots, n.$$

The number of knot points of the spline function $\overline{\pi}$ is $N_T < n-4$; the first four elements of $R_{\mathcal{I}}$ cannot be knot points. De Boor's (1978) B-spline construction implies that the first derivative of $\overline{\pi}$ is piecewise quadratic, the second derivative is piecewise linear, the third is a stepwise constant function, and the fourth is a function that is zero everywhere except at knot points. The latter condition translates into a linear constraint on the state prices associated with knot points:

$$ND\pi = 0, \tag{4}$$

where N is a $(n-4-N_T) \times (n-4)$ selection matrix whose rows are vectors of the Euclidean basis of $\mathbb{R}^{n-4}$, $D = D_{n-3}D_{n-2}D_{n-1}D_n$, and D_k is the $(k-1) \times k$ first-difference matrix

$$D_k = \begin{bmatrix} -1 & 1 & 0 & 0 & \cdots & 0 \\ 0 & -1 & 1 & 0 & \cdots & 0 \\ \vdots & & & & & \vdots \\ \vdots & & & & & \vdots \\ 0 & 0 & \cdots & 0 & -1 & 1 \end{bmatrix}. \tag{5}$$

The LAD estimator of state prices is based on equation (2) and on the set of linear restrictions (4). An estimator minimizing absolute pricing errors is preferred to a least-squares estimator because of its computational convenience and robustness to outliers, which are known to contaminate data on option prices. The LAD estimator $\hat{\pi}_{LAD}$ of the state prices is the solution of the minimization problem

$$\hat{\pi}_{LAD} = \arg\min_{\pi} \sum_{i=1}^{N_C+N_P} w_i |Y_i^0 - X_i \pi| \qquad (6)$$
$$\text{s.t.} \quad ND\pi = 0, \pi \geq 0,$$

where Y_i^0 and X_i are the rows of Y^0 and X, respectively, and w_i are weights assigned to pricing errors. In our estimates we set $w_i = 1/\sqrt{Y_i^0}$: this choice applies a dampening factor to deeply out-of-the-money options, which tend to have larger pricing errors in percentage terms.

The minimization problem can be written as a linear programming (LP) problem:

$$\min_z d^T z \qquad (7)$$
$$\text{s.t.} \quad Az = b, z \geq 0,$$

where

$$d = \begin{bmatrix} w \\ w \\ 0 \end{bmatrix} \quad \text{and} \quad z = \begin{bmatrix} \varepsilon^+ \\ \varepsilon^- \\ \pi \end{bmatrix}$$

are $(2N_C + 2N_P + n) \times 1$ vectors, w is the $(N_C + N_P) \times 1$ vector of weights, ε^+ and ε^- are the positive and negative parts of the $(N_C + N_P) \times 1$ vector of pricing errors, and

$$A = \begin{bmatrix} I & -I & X \\ 0 & 0 & ND \end{bmatrix} \quad \text{and} \quad b = \begin{bmatrix} Y^0 \\ 0 \end{bmatrix}.$$

The solution of the LP problem can be found by standard and computationally inexpensive LP algorithms. The LAD estimator $\hat{\pi}_{LAD}$ is then given by the last n components of the LP solution $\hat{z}$.

Once we have computed the LAD estimator $\hat{\pi}_{LAD}$, we can get a new estimator $\hat{\pi}_U$ fulfilling a unimodality condition. Since the risk-neutral distributions are obtained by rescaling the state prices, the unimodality of π implies the unimodality of the risk-neutral distribution ρ. Let

$$\varphi(\pi) = \arg\max_i \pi_i \quad \text{and} \quad g(\pi) = (g_1(\pi), \ldots, g_{n-1}(\pi))$$
$$\text{s.t. } g_i(\pi) = \begin{cases} 1 & \text{if } i < \varphi(\pi) \\ -1 & \text{if } i \geq \varphi(\pi) \end{cases}.$$

The set of vectors that satisfy unimodality is $U = \{\pi \in \mathbb{R}^n_+ : (D_n\pi) \circ g(\pi) \geq 0\}$, where D_n is the first-difference matrix defined in (5) and $\circ$ denotes the Hadamard or entrywise product. The unimodal LAD estimator $\hat{\pi}_U$ is then the solution of the minimization problem

$$\min_\pi \sum_{i=1}^{N_C+N_P} w_i |Y_i^0 - X_i\pi| \tag{8}$$
$$\text{s.t.} \quad ND\pi = 0, \pi \geq 0 \text{ and } \pi \in U.$$

A way to solve the problem (8) using a Bayesian version of the LAD estimator is detailed in Taboga (2016).

Inflation Swap Rates in Terms of State Prices

Let s_M be the fixed leg of a zero-coupon inflation swap with maturity M years. Let $\mathcal{I}_M$ be the (stochastic) annual rate of inflation over the next M years. Taking expectations under the risk-neutral measure Q, the following condition must hold:

$$\mathbb{E}_0^Q[D_M((1+s_M)^M - (1+\mathcal{I}_M)^M)] = 0, \tag{9}$$

where D_M is the discount factor for the time interval $[0, M]$. Rewriting equation (9) in terms of state prices and taking into account that $\mathbb{E}_0^Q[D_M] = \sum_j \pi_j$, we get

$$\left(\sum_{j=1}^n \pi_j\right)(1+s_M)^M - \sum_{j=1}^n (1+i_j)^M \pi_j = 0.$$

Since the risk-neutral distribution d is given by $d_j = \pi_j / \sum_k \pi_k$, there follows that

$$s_M = \left(\sum_{j=1}^{n} d_j (1 + i_j)^M \right)^{1/M} - 1. \qquad (10)$$

For $M = 1$, this equivalence boils down to

$$s_1 = \sum_{j=1}^{n} d_j i_j \, ;$$

the inflation swap rate equals the mean of the option-implied distribution d. For $M > 1$, equation (10) states that the inflation swap rate is a non-linear function of the probability distribution extracted from inflation options having the same maturity.

Appendix 2. Copula Functions

DEFINITION 2 (Copula Function). *A copula is an n-dimensional distribution function $C : [0,1]^n \rightarrow [0,1]$ of a random vector $(U_1, \ldots, U_n)$, where the marginal law of U_i is the uniform distribution on $[0,1]$ for all $i \in \{1, \ldots, n\}$.*

Copula functions are very popular in the study of multivariate distribution functions thanks to their role in imposing a dependence structure on predetermined marginal distributions. Their importance derives from Sklar's theorem, which proves that any multivariate distribution function can be characterized by a copula and that copula functions, together with univariate marginal distribution functions, can be used to construct multivariate distribution functions.

THEOREM 1 (Sklar's Theorem). *Let H be an n-dimensional distribution function with marginals $F_1, \ldots, F_n$. Then an n-copula C exists such that, for each $x \in R^n$,*

$$H(x_1, \ldots, x_n) = C(F_1(x_1), \ldots, F_n(x_n)).$$

If the marginals $F_1, \ldots, F_n$ are all continuous, then C is unique; otherwise C is univocally determined on $(RanF_1 \times RanF_2 \times RanF_n)$

(where $RanF_i$ denotes the rank of F_i). Conversely, if C is an n-copula and $F_1, \ldots, F_n$ are distribution functions, then the function H defined above is an n-dimensional distribution function with marginals $F_1, \ldots, F_n$.

The proof of this theorem can be found, e.g., in Nelsen (2006).

The main feature of Sklar's theorem is that for continuous multivariate distribution functions, the univariate marginals and the multivariate dependence structure can be separated and the dependence structure can be represented by a copula.

Let F be a univariate distribution function. Let us recall that the generalized inverse of F is defined as $F^{-1}(t) = \inf\{x \in \mathbf{R} | F(x) \geq t\}$ for each t in $[0,1]$, with the usual convention that $\inf(\emptyset) = -\infty$.

An important corollary of Sklar's theorem, which is fundamental in the study of copulas and their applications, is the following:

Corollary 1. *Let H be an n-dimensional distribution function with continuous marginals $F_1, \ldots, F_n$ and copula C. Then for each $u \in [0,1]^n$,*

$$C(u_1, \ldots, u_n) = H(F_1^{-1}(u_1), \ldots, F_n^{-1}(u_n)).$$

In the following we recall the Student's t-copula that we use in the paper.

Definition 3 (Student's t-Copula). *The Student's t-copula can be written as*

$$C_{\rho,\nu}(u,v) = \int_{-\infty}^{t_\nu^{-1}(u)} \int_{-\infty}^{t_\nu^{-1}(v)} \frac{1}{2\pi(1-\rho^2)^{1/2}} \times \left\{1 + \frac{x^2 - 2\rho xy + y^2}{\nu(1-\rho^2)}\right\}^{-(\nu+2)/2} dsdt,$$

where ρ and ν are the parameters of the copula, and t_ν^{-1} is the inverse of the standard univariate Student's t-distribution with ν degrees of freedom, expectation 0, and variance $\frac{\nu}{\nu-2}$.

Student's t-copula allows for joint fat tails. Increasing the value of ν decreases the tendency to exhibit extreme co-movements. The

Student's t-dependence structure supports joint extreme movements regardless of the marginal behavior of the individual variables.

Appendix 3. TailCor Measures

Let X_{jt} be the j-th element of the random vector $\mathbf{X}_t$. Denote by Q_j^τ its τ-th quantile for $0 < \tau < 1$, and let $IQR_j^\tau = Q_j^\tau - Q_j^{1-\tau}$ be the τ-th interquantile range. Let Y_{jt} be the standardized version of X_{jt}:

$$Y_{jt} = \frac{X_{jt} - Q_j^{0.50}}{IQR_j^\tau}.$$

By standard trigonometric arguments, the projection of (Y_{jt}, Y_{kt}) onto the 45-degree line is

$$Z_t^{(jk)} = \frac{1}{\sqrt{2}}(Y_{jt} + Y_{kt}),$$

and the tail interquantile range is

$$IQR^{(jk)\xi} = Q^{(jk)\xi} - Q^{(jk)1-\xi},$$

where $Q^{(jk)\xi}$ is the ξ-th quantile of Z_t^{jk}. The larger ξ is, the further we explore the tails.

TailCor is then defined as follows (Ricci and Veredas 2013):

DEFINITION 4 (TailCor). *Under technical assumptions, TailCor between X_{jt} and X_{kt} is*

$$TailCor^{(jk)\xi} := s_g(\xi, \tau) IQR^{(jk)\xi},$$

where $s_g(\xi, \tau)$ is a normalization such that under Gaussianity and linear uncorrelation $TailCor^{(jk)\xi} = 1$, the reference value.

A table with values of $s_g(\xi, \tau)$ for a grid of reasonable variables for τ and ξ can be found in Ricci and Veredas (2013), appendix T.

When interest lies in the tail of one side of the distribution, downside TailCor and upside TailCor can be used:

DEFINITION 5 (Downside TailCor). *Downside TailCor is defined as*

$$TailCor^{(jk)\xi-} := s_g(\xi, \tau) IQR^{(jk)\xi-},$$

where $IQR^{(jk)\xi-} = Q^{(jk)0.50} - Q^{(jk)1-\xi}$.

Definition 6 (Upside TailCor). *Upside TailCor is defined as*

$$TailCor^{(jk)\xi+} := s_g(\xi,\tau) IQR^{(jk)\xi+},$$

where $IQR^{(jk)\xi+} = Q^{(jk)\xi} - Q^{(jk)0.50}$.

The estimation procedure consists of four simple steps that can be followed under technical assumptions:

(i) Standardize X_{jt} and X_{kt}.
(ii) Estimate the *IQR* of the projection: $I\hat{Q}R^{(jk)\xi}_{\hat{Z},T}$.
(iii) Find the normalization $s_g(\xi,\tau)$ from the table.
(iv) Compute $Tai\hat{l}Cor^{(jk)\xi}_{\hat{Z},T} = s_g(\xi,\tau) I\hat{Q}R^{(jk)\xi}_{\hat{Z},T}$.

References

Antunes, A. 2015. "Co-movement of Revisions in Short- and Long-Term Inflation Expectations." *Economic Studies* (Banco de Portugal) 1 (1): 1–19.

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

Bauer, M., and J. Christensen. 2014. "Financial Market Outlook for Inflation." FRBSF Economic Letter No. 2014-14 (May 12).

Christoffersen, P., V. Errunza, K. Jacobs, and H. Langlois. 2012. "Is the Potential for International Diversification Disappearing? A Dynamic Copula Approach." *Review of Financial Studies* 25 (12): 3711–51.

Ciccarelli, M., and J. A. Garcia. 2015. "International Spillovers in Inflation Expectations." ECB Working Paper No. 1857.

de Boor, C. R. 1978. *A Practical Guide to Splines.* Springer-Verlag.

Ehrmann, M. 2014. "Targeting Inflation from Below — How Do Inflation Expectations Behave?" Working Paper No. 52, Bank of Canada.

Embrechts, P., F. Lindskog, and A. McNeil. 2003. "Modelling Dependence with Copulas and Applications to Risk Management." In *Handbook of Heavy Tailed Distributions in Finance*, ed. S. Rachev. Elsevier.

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

Gürkaynak, R., B. Sack, and E. Swanson. 2005. "The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models." *American Economic Review* 95 (1): 425–36.

Kocherlakota, N. 2013. "Optimal Outlooks." Speech at the Conference on Extracting and Understanding the Risk Neutral Probability Density from Option Prices, NYU Stern School of Business, New York City, September 20.

Lyziak, T., and M. Paloviita. 2016. "Anchoring of Inflation Expectations in the Euro Area: Recent Evidence Based on Survey Data." ECB Working Paper No. 1945.

Nelsen, R. 2006. *An Introduction to Copulas*. Second Edition. New York: Springer.

Pericoli, M. 2012. "Expected Inflation and Inflation Risk Premium in the Euro Area and in the United States." Working Paper No. 842, Bank of Italy.

Ricci, L., and D. Veredas. 2013. "TailCor." Mimeo.

Scharnagl, M., and J. Stapf. 2015. "Inflation, Deflation, and Uncertainty: What Drives Euro-Area Option-Implied Inflation Expectations, and Are They Still Anchored in the Sovereign Debt Crisis?" *Economic Modelling* 48: 248–69.

Smith, T. 2012. "Option-Implied Probability Distributions for Future Inflation." *Quarterly Bulletin* (Bank of England) 52 (3).

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: 33–48.

Taboga, M. 2016. "Option-Implied Probability Distributions: How Reliable? How Jagged?" *International Review of Economics and Finance* 45 (September): 453–69.