Decaying Expectations: What Inflation Forecasts Tell Us about the Anchoring of Inflation Expectations*

Aaron Mehrotra and James Yetman
Bank for International Settlements

Well-anchored inflation expectations are considered to be a reflection of credible monetary policy. In the past, anchoring has typically been assessed using either long-run inflation surveys or breakeven inflation rates on financial assets with long maturities. Here we propose an alternative measure of inflation anchoring that makes full use of readily available, multiple-horizon, fixed-event forecasts. We show that a model where forecasts are assumed to diverge from a perceived long-run anchor towards actual inflation as the forecast horizon shortens fits the data well. It also provides simple estimates of the degree to which inflation expectations are anchored. We use our methodology to examine how inflation anchoring has evolved in forty-four economies.

JEL Codes: E31, E58.

1. Introduction

Well-anchored inflation expectations—where anchoring refers to both the level and the variability of anticipated future inflation—are important for the monetary transmission mechanism and can be

*The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank for International Settlements. Bat-el Berger, Steven Kong, Lillie Lam, and Jimmy Shek provided excellent research assistance. We are grateful to the co-editor Pierpaolo Benigno, three anonymous referees, Dietrich Domanski, Gregor Smith, Philip Turner, Frank Packer, Marco Lombardi, Pierre Siklos, and seminar participants at the Bank for International Settlements and Reserve Bank of New Zealand for helpful comments and suggestions. Author contact: Bank for International Settlements, Representative Office for Asia and the Pacific, 78th Floor, Two IFC, 8 Finance Street, Central, Hong Kong; e-mails: aaron.mehrotra@bis.org; james.yetman@bis.org.

considered to be a reflection of credible monetary policy.[1] If inflation expectations are not well anchored, forward-looking price- and wage-setting behavior may become a source of macroeconomic instability. In New Keynesian models, for example, well-anchored inflation expectations can significantly contribute to stabilizing actual inflation.

Assessing how well anchored inflation expectations are is complicated by the fact that actual inflation is subject to persistent shocks that drive inflation away from any anchor point. One solution to this problem is to assess anchoring by means of long-run inflation forecasts, beyond the horizon where shocks might have a measurable effect. Previous authors have done this mostly based on breakeven inflation rates drawn from asset prices with long maturities.[2] However, such rates are influenced by many factors other than inflation expectations.[3] Long-run inflation surveys can also be used. But, as we later discuss, only the mean forecast is typically available, so it is not clear how one can infer how tightly expectations are anchored using this measure.

In this paper, we propose a novel way to model the behavior of inflation expectations, and we use this to assess the degree of anchoring. We fit raw forecast data with a decay function, where inflation forecasts monotonically diverge from an estimated perceived long-run anchor in the direction of recent actual inflation as the forecast horizon shortens. We motivate this based on how inflation forecasts vary depending on the horizon for which they are made. Forecasts made sufficiently far in advance can appear to be anchored at a level that bears little relationship with actual inflation—both at the time that the forecast is made and for the period being forecast. For inflation-targeting economies, for example, the anchor level could correspond to the central bank's inflation target. However, this need not be the case. If the central bank lacks credibility, then long-horizon inflation forecasts may be anchored to some level that differs from the stated target, or may be heavily influenced by the actual inflation rate at the time the forecast is made. Regardless, as the

[1] This is especially true for inflation-targeting economies, where inflation forecasts can be viewed as an intermediate monetary policy target; see Svensson (1997).

[2] See, for example, European Central Bank (2012).

[3] See the discussion in Faust and Wright (2013).

forecasting horizon reduces, any role that a perceived anchor may have played in affecting inflation forecasts is likely to decrease as forecasters learn more about the realization of shocks to inflation for a given period. Our modeling strategy allows for all these possibilities. We also show that our empirical model is consistent with an autoregressive model of inflation.

Our decay function provides a parsimonious framework for fitting inflation forecasts that fully utilizes the multiple-horizon dimension of the available data. Furthermore, it allows us to generate an estimate of the perceived long-run anchor of inflation expectations.

We also evaluate potential changes in the estimated anchor over time by estimating rolling regressions over all available economy five-year-sample pairs. This yields time-varying estimates for the perceived anchor and its standard error.

Our data are from professional forecasters, collected and published by Consensus Economics. Our sample comprises median forecasts for each of forty-four economies. This includes a large number of emerging markets (in contrast with the advanced-economy focus prevalent in most previous related research) and is a considerably larger sample than similar studies use to evaluate the anchoring of inflation expectations. Indeed, the economies in our sample accounted for 88 percent of global GDP in 2014 at market exchange rates.

The research is related to the work by Kozicki and Tinsley (2012), who model the evolution of inflation expectations in the United States. They construct the term structure of inflation expectations using survey data. In their model, as in ours, as the forecast horizon increases, inflation expectations eventually converge to a fixed point. This endpoint may shift as the private sector adjusts their perceptions of the long-run target of the central bank.[4]

Our study is also related to other research examining the anchoring of inflation expectations. One strand of this research extracts measures of inflation expectations from high-frequency financial market data and investigates their link with macroeconomic variables, in particular economic news announcements. Gürkaynak et al.

[4]In earlier research, Kozicki and Tinsley (2001a, 2001b) show how evolving endpoints can be used to model the term structure of interest rates and the expectations hypothesis. Such shifts can be linked to agents' learning about the inflation target.

(2007) use daily data on nominal and inflation-indexed bonds and find evidence that long-run inflation expectations are better anchored in two inflation targeters (Canada and Chile) than in a non-inflation targeter (the United States).[5] Beechey, Johannsen, and Levin (2011) use daily data on inflation swaps and bond spreads to compare the anchoring of inflation expectations in the euro area with the United States. Galati, Poelhekke, and Zhou (2011) also use expectations measures derived from financial market data to investigate whether the international financial crisis affected inflation expectations. And the pass-through from short-term to long-term breakeven inflation rates is investigated using financial market data in Gefang, Koop, and Potter (2012), Jochmann, Koop, and Potter (2010), and Lemke and Strohsal (2013).

In another strand of research, anchoring is examined by comparing survey data on inflation expectations with movements in actual inflation. Levin, Natalucci, and Piger (2004) find that inflation expectations are less correlated with lagged inflation in economies with explicit inflation targets, while Clark and Nakata (2008) show that unexpected increases in inflation in recent years result in smaller increases in inflation expectations than twenty years ago.[6]

Our use of decay functions in modeling the behavior of inflation expectations is not completely new. Gregory and Yetman (2004) use a polynomial decay function and Blue Chip survey data to model the behavior of professional forecasters, focusing on the observation that forecasts made by different forecasters of the same outcome converge towards a consensus as the forecast horizon shortens. A related approach is the Bayesian learning model of Lahiri and Sheng (2008, 2010), who model the evolution of forecast disagreement across horizons. At long horizons, forecasters' prior beliefs are important. Then, as the forecast horizon shortens, newly arriving public information becomes more relevant. These studies incorporate all available forecast horizons into their empirical approach.[7]

[5] See also Gürkaynak, Levin, and Swanson (2010).

[6] Recently, Nason and Smith (2014) have used surveys of professional forecasters to measure the evolving trend (capturing long-run inflation expectations) and the cycle, in a trend-cycle model of inflation for the United States.

[7] See also Davies, Lahiri, and Sheng (2011) for a summary of the literature using three-dimensional panels that incorporate multiple forecasters, target dates, and forecast horizons. Here we only focus on the last two of these.

Decay functions have been discussed in relation to forecasts in other contexts as well. Faust and Wright (2013) find that a model with long-term and near-term expectations derived from surveys, together with a simple exponential decay path between them, does very well in terms of forecasting performance. Their exponential decay function is a special case of the Weibull distribution cumulative density function-based decay path that we consider here. However, we focus on capturing the behavior of expectations over the different forecast horizons, rather than attempting to provide *accurate* forecasts of future inflation.

We also note the methodological benefits of the proposed approach to analyze fixed-event forecast data. In empirical applications, it is admittedly easier to use fixed-horizon forecasts (for example, forecasts made each month for the following twelve-month period) than fixed-event forecasts (as in the case of the Consensus forecasts we examine, made for calendar years at varying horizons). Unfortunately, panels of fixed-horizon forecasts of comparable length are not available for many economies, limiting their applicability for large cross-country analyses.[8] Given this lack of data, a common approach has been to approximate fixed-horizon forecasts based on a weighted average of two fixed-event forecasts made for different periods (e.g., Dovern, Fritsche, and Slacalek 2012; Gerlach 2007; Kortelainen, Paloviita, and Viren 2011; Siklos 2013). But this approach has some limitations. First, it reduces the sample from which the observations are drawn, discarding valuable information. Second, the distribution of the weighted average of two forecasts, each made for different forecast horizons, is likely to have non-standard statistical properties.[9]

At the same time, we also acknowledge the limitations of the proposed approach. For instance, one needs multiple data points to estimate the perceived long-run anchor, potentially making it challenging to rapidly detect shifts in the degree of anchoring. This

[8]To our knowledge, comparable-length panels of fixed-horizon forecasts of inflation are only available for the United States (from the Federal Reserve Bank of Philadelphia) for up to one year ahead and for the euro area (from the European Central Bank) for both one and two years ahead.

[9]Limitations of this approach are acknowledged by some papers using this method. For example, Kortelainen, Paloviita, and Viren (2011) state that the “moving average process affects the properties of the data.”

stands in contrast to financial-market-based measures of inflation expectations that are typically available at a high frequency. And, regarding the interpretation of the coefficients, the standard error of the perceived anchor is related to the degree of anchoring, but is also likely to be influenced by sampling errors in the forecast data, as well as variation in the perceived long-run anchor over the sample and other possible sources of model misspecification. Finally, we note that the approach only uses forecast data at horizons of up to twenty-four months. If forecast behavior at longer horizons was different, our estimate for the perceived anchor could yield a biased measure of the anchor of longer-horizon expectations. This is in contrast to long-run inflation surveys, for example. Thus, our measure should be considered complementary to existing measures of anchoring.

This paper is structured as follows. The next section presents the methodology, outlines the decay function, and discusses its suitability for evaluating the behavior of inflation expectations. Section 3 applies the model to data from the United States, evaluates its robustness, and compares our estimated perceived anchor with other possible measures. Section 4 presents the results from applying our model to data from forty-three other economies. Section 5 concludes.

2. Methodology

We propose a parsimonious framework for fitting inflation forecasts that fully utilizes the multiple-horizon dimension of the data. The basic assumption behind our adopted functional form is that, if inflation expectations are well anchored at a particular level, inflation forecasts made sufficiently far in advance should be equal to the perceived anchor. Indeed, in an environment where inflation expectations are well anchored, there should exist some horizon beyond which long-run expectations are fixed and do not systematically respond to new data about economic conditions.[10] This suggests

[10]Long-run expectations could still change if, for example, a shock causes the degree of monetary policy credibility to change, or if the central bank announces a new level for the inflation target, consistent with the "moving endpoints" of Kozicki and Tinsley (2012).

that there are at least two dimensions to anchoring: both the level at which expectations are anchored in the long run and how tightly expectations are anchored at that level. Here, we focus primarily on the former.

As the forecast horizon shortens, even well-anchored inflation expectations will eventually start to deviate from their perceived long-run anchor towards the level of actual inflation. Forecasters gradually learn more about the realization of inflation shocks for a given period, for example. A slow adjustment could arise due to information about economic conditions being disseminated only slowly through the economy. This could result from costs of acquiring and processing new information, as in Devereux and Yetman (2003) and Mankiw and Reis (2002). Alternatively, the gradual adjustment could arise from the forecasters' prior beliefs, resulting in expectation stickiness, as postulated by Lahiri and Sheng (2008, 2010). Yet another explanation could be limited information-processing capacity (Sims 2003).

We apply our framework to forecast data from Consensus Economics. These provide us with a large cross-country data set for a sufficiently long period of time, constructed using a consistent methodology. Indeed, while surveys of households and firms could potentially provide further insights into the anchoring of inflation expectations, the global nature of the data set provides consistency in measurement across a large number of economies. Inflation in the consumer price index (CPI) is used as the measure of headline inflation.[11]

Consensus Economics starts collecting forecasts for calendar-year inflation outcomes in January of the preceding year. They generally collect these forecasts each month until December of the year being forecast, for a total of twenty-four monthly forecasts of the same outcome. Based on these fixed-event forecasts, the left-hand panels of figure 1 illustrate the behavior of median inflation forecasts across horizons from $h = 24$ (twenty-four months before the completion of the year being forecast) to $h = 1$ (one month before the

[11] In addition, we provide estimates based on retail prices (RPIX) for the United Kingdom, since the official inflation target was specified in terms of the RPIX for part of the sample, and both forecasts and outcomes are available for this variable as well.

Figure 1. Forecasts of Headline Inflation (left-hand panels) and Gap between Forecasts and Recent Inflation at Different Horizons (right-hand panels)

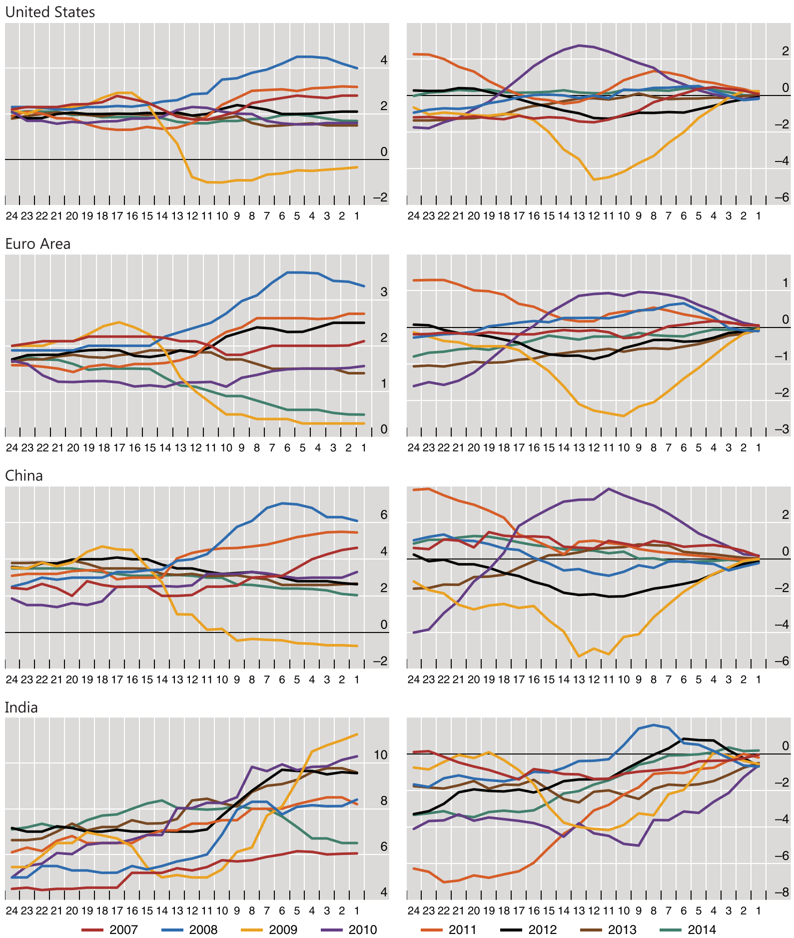

Sources: Consensus Economics; IMF, International Financial Statistics; Datastream; national data; BIS calculations.
Notes: Left-hand panels: each line shows forecasts of a given inflation outcome at different horizons. Right-hand panels: each line shows the gap between forecasts and the latest inflation outcome available at the time that the forecast was made. The horizontal axis shows the forecast horizon, e.g., "24" indicates forecasts made twenty-four months before the completion of the calendar year being forecast. The vertical axis is measured in percentage points. For India, data are based on March years (2014 = year ended March 2015); for all other countries, data are based on calendar years.

completion of the year being forecast) for the period 2007–14 for two large advanced economies (the United States and the euro area), as well as two large emerging economies (China and India). (Full-color versions of all figures in this text are available on the IJCB website.) The figure confirms our prior regarding the behavior of expectations. The forecasts for the different years do not vary much when the forecast horizon is long, but they start to deviate further from each other as the forecast horizon becomes shorter. The close resemblance between the twenty-four-month-ahead forecasts during a time period that includes the international financial crisis is particularly striking in the case of the United States, while we observe somewhat more divergence at longer horizons in the case of India.

The dynamic illustrated for inflation forecasts in the left-hand panels of figure 1 parallels the cross-country evidence provided by Isiklar and Lahiri (2007) regarding forecasts of GDP growth. They report that forecasts are generally very similar at a twenty-four-month horizon and do not change very much during the first few months as the forecast horizon starts to shorten. Comparable evidence is shown in Capistran and Lopez-Moctezuma (2014) for Mexico, where the forecasts by professional forecasters for both inflation and GDP growth are very similar at twenty-four-month forecast horizons across different years.

The right-hand panels of figure 1 illustrate another facet of forecast data that we wish to exploit. They display the difference between forecasts and the latest twelve-month inflation rate (at monthly frequency) available at the time that the forecasts were made. As the forecast horizon shortens, especially below twelve months, the gaps shrink and inflation forecasts look increasingly like the latest available inflation outcome. This is not surprising given that there is an overlap between the period being forecast and the period covered by the corresponding inflation outcome that increases as the horizon shortens. At horizons longer than twelve months there is a similar, albeit weaker, dynamic, although with notable exceptions (especially forecasts for 2009 made in late 2008 and early 2009). We impose a flexible functional form on the decay function that can accommodate this.

In addition to the empirical evidence in figure 1, we can motivate our empirical model of forecasts based on inflation following an autoregressive (AR) process. In its simplest form, consider the case where inflation at time t, π_t, equals the long-run perceived anchor π^{e*} and an error term ε_t, and the error term follows an AR(1) process, with $0 < \rho < 1$:

$$\pi_t = \pi^{e*} + \varepsilon_t, \tag{1}$$

$$\varepsilon_t = \rho\varepsilon_{t-1} + \eta_t. \tag{2}$$

The error term is intended to capture any shocks that take inflation away from its perceived long-run anchor point, while the autoregressive process allows for persistent shocks such that inflation adjusts only slowly. We can write

$$\begin{aligned} E_{t-j}(\pi_t) &= \pi^{e*} + \rho^j \varepsilon_{t-j} \\ &= \pi^{e*} + \rho^j(\pi_{t-j} - \pi^{e*}) \\ &= (1-\rho^j)\pi^{e*} + \rho^j \pi_{t-j}. \end{aligned} \tag{3}$$

The last expression implies that inflation forecasts are explained by a weighted average of the perceived anchor and the latest available inflation data at the time forecasts are made, where the weight on the anchor is increasing in the forecast horizon. The process shown in equations (1)–(3) can be extended to a more general (AR(p)) case where, under plausible assumptions, the weight on the anchor remains an increasing function of the forecast horizon (see the appendix).

We therefore model the expectations process for each economy as follows. The forecast of inflation for year t made at horizon h, denoted $f(t, t-h)$, is assumed to follow:

$$f(t, t-h) = \alpha(h)\pi^{e*} + [1 - \alpha(h)]\pi(t-h) + \varepsilon(t, t-h). \tag{4}$$

In (4), h is measured in months before the end of the year that is being forecast. π^{e*} is the level to which long-run inflation expectations are perceived to be anchored. $\pi(t-h)$ is the latest available level of annual inflation observed at the time when the forecast is made, and $\varepsilon(t, t-h)$ is a residual term.[12] Our measure of actual

[12]The declining weight on the perceived long-run anchor as the horizon shortens is comparable to the estimated dynamics arising from the Bayesian learning

inflation is the year-on-year growth rate of a twelve-month moving average of the CPI at monthly frequency, lagged by one month to allow for publication lag.[13] Using lagged inflation also helps to address any potential endogeneity issues between expected inflation and inflation outturns.

$\alpha(h)$ denotes a decay function. As already discussed, this has the property that, as the horizon shortens, there is greater weight on realized outcomes and less on the perceived long-run anchor point. In particular, we wish to assume a functional form that ensures that $\alpha(\infty) = 1$ and $\alpha(0) = 0$.

We are agnostic on the exact form that the decay function should take. Different economies are subject to different shock processes. Ideally, we would therefore like a flexible functional form that can embrace a wide range of possible paths as the horizon shortens. The candidate we consider is

$$\alpha(h) = 1 - \exp\left(-\frac{h}{b}^{c}\right). \tag{5}$$

This is based on the cumulative density function of the Weibull distribution.[14] Figure 2 illustrates the wide variety of possible decay paths that this functional form can generate for different values of b and c. With a small b parameter, for example, the function remains near one until the horizon gets close to zero. For high b, the opposite is true: $\alpha(h)$ may be close to zero for all short horizons. The c parameter potentially provides the decay function with some shape. For example, when $b = 4$, the function stays closer to one when c is higher, but only at forecast horizons above four. Below that horizon, a higher c implies a more rapid decline in α. For $c = 1$, (5) reduces to a simple exponential decay function.

model of Lahiri and Sheng (2008). They argue that forecasters give a lower weight to public information at longer horizons due to its lower perceived quality. As the forecast horizon shortens, the information becomes more accurate and certain, and so the weight placed by forecasters on public information (such as recent inflation) increases.

[13] At a horizon of zero months, this should coincide with the definition of annual inflation that forecasters are attempting to predict.

[14] In Mehrotra and Yetman (2014), we briefly examine a more restricted version of the model (with $c = 1$) for a group of Asian economies.

Figure 2. Weibull Decay Functions

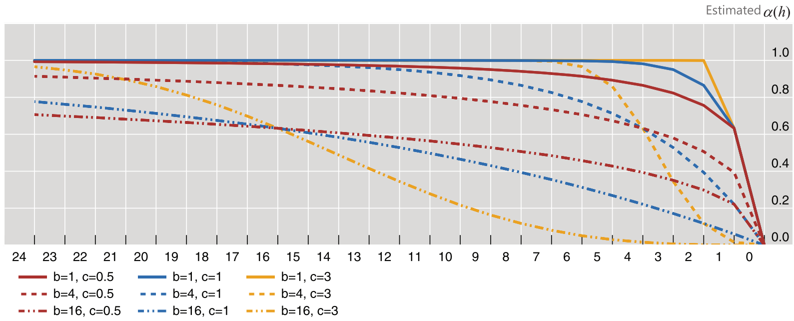

Source: Authors' calculations.
Note: The horizontal axis represents the forecast horizon h, which is the number of months before the end of the calendar year being forecast.

The variance of the residual in (4) is modeled using a flexible functional form that allows it to change with the forecast horizon as

$$V(\varepsilon(h,t)) = \exp(\delta_0 + \delta_1 h + \delta_0 h^2). \qquad (6)$$

This formulation allows the variance to vary across the forecasting horizon with minimal restrictions. Furthermore, the use of the exponential function in (6) ensures that the estimated variance is non-negative at all horizons for all possible values of the parameters.[15]

Forecasts made at different horizons for the same inflation outcome are likely to be highly correlated, especially if the horizons are close together. We explicitly model this, assuming that the correlation between residuals for forecasts of the same inflation rate, but made at two different horizons h and k, is given by

$$Corr(\varepsilon(t,t-h)),(\varepsilon(t,t-k)) = 1 - \Phi_1|h-k| - \Phi_2(h-k)^2. \qquad (7)$$

[15]Note that we are not imposing that the residual variance declines monotonically as the forecast horizon falls. If all forecasters agreed on the long-run anchor but disagreed on the interpretation of early incoming data that pertained to the inflation outcome, for example, this residual variance could increase as the forecast horizon shortened, at least at longer horizons. This is in contrast to forecast error variability, which is likely to decrease monotonically as the forecast horizon falls; see Isiklar and Lahiri (2007).

The gradual adjustment of inflation expectations this embodies is in line with the observed empirical autocorrelation of inflation that decays only slowly (see Fuhrer and Moore 1995). In practical terms, the off-diagonal elements of the variance-covariance matrix take the form

$$Cov(\varepsilon(t,t-h)),(\varepsilon(t,t-k)) =$$
$$\left[\sqrt{V(\varepsilon(t,t-h))V(\varepsilon(t,t-k))}\right][1-\Phi_1|h-k|-\Phi_2(h-k)^2]. \quad (8)$$

We do not model the persistence of the errors through time, as some of the observations are missing for some economies. This leads to some loss of efficiency in terms of estimation. But, as we generally have twenty-four forecast horizons for each year, we find that our estimates are not lacking in precision. At the same time, we acknowledge that the assumptions made regarding the covariance structure could have an impact on the estimated π^{e*} (see section 3).

We capture potential time variation in the estimated anchor by estimating rolling regressions over all available five-year subsamples, yielding time-varying estimates for π^{e*}.

The model is estimated by maximum likelihood, economy by economy. We consider a wide range of possible starting values for each economy, and maximize the likelihood function using the hill-climbing method of Broyden, Fletcher, Goldfarb, and Shanno (see Press et al. 1988 for details), until the estimates converge. We then choose the estimates with the highest corresponding log-likelihood function value.

It is of interest to compare our approach with that by Kozicki and Tinsley (2012), who construct the term structure of inflation expectations using survey data. In their model, as in ours, inflation expectations converge to a fixed point if the forecast horizon is long enough. However, their model is one of inflation, and expectations are assumed to be consistent with the model, whereas our framework only considers the behavior of expectations. Relatedly, there are at least three important differences between our method and the one that they propose.

One difference concerns the type of data used and its implications for the analysis. Kozicki and Tinsley consider two to three different horizons: eight months, fourteen months, and, in a latter part of the sample, ten years. Arguably, the small number of widely varying forecast horizons make the modeling of the underlying inflation process highly relevant in Kozicki and Tinsley (2012), as a means of linking the forecasts. In contrast, we consider twenty-four horizons, all one month apart.

A second difference relates to the structure of the variance-covariance matrix. Kozicki and Tinsley (2012) assume no correlation between the degree to which the forecasts fit the model at different horizons—their model is assumed to fit equally well at all horizons. We assume a more general form of the variance-covariance matrix, where the diagonal elements vary across horizons. Moreover, the covariance elements are allowed to be non-zero, and they are a function of the gap between the forecast horizons. Again, the nature of the underlying data is relevant: in our case, the assumption of a diagonal variance-covariance matrix would be problematic, as there are a large number of horizons that substantially overlap.

A third difference concerns the estimation method for the fixed point. Kozicki and Tinsley (2012) use a Kalman filter that requires a long sample and assume that both the autoregressive coefficients and the variance-covariance matrix are constant over this sample. We use shorter rolling samples, assuming that the perceived long-run anchor remains constant for these shorter samples, but allow for variation in all estimated parameters from one rolling sample to the next.

Taken together, while there are similarities in the underlying approaches, we believe that our approach is sufficiently different from that in Kozicki and Tinsley (2012) to make a valuable contribution, and is more appropriate to the data that we focus on.

We noted in the introduction that it is common in the existing literature to use an approximation in order to convert fixed-event forecasts to fixed-horizon forecasts, instead of fully utilizing the multiple-horizon nature of the data. For example, in Dovern, Fritsche, and Slacalek (2012), the approximation is a weighted average of two fixed-event forecasts. Let $\hat{\pi}_{t+k|t}$ denote the k-month-ahead forecast for inflation based on information available at time t. The survey includes a pair of forecasts $\{\hat{\pi}_{t+k|t}, \hat{\pi}_{t+12+k|t}\}$ for each month, with horizons $k \in \{1, 2, \ldots, 12\}$ and $k + 12$ months. Fixed twelve-month horizon forecasts are then approximated as averages of the forecasts for the current and next calendar years, weighted by their shares in the forecast period:

$$\hat{\pi}_{t+12|t} = \frac{k}{12}\hat{\pi}_{t+k|t} + \frac{12-k}{12}\hat{\pi}_{t+12+k|t}.$$

This approach implies that the twelve-month-ahead forecast for inflation made in October 2012 is approximated by the sum of $\hat{\pi}_{2012M12|2012M10}$ and $\hat{\pi}_{2013M12|2012M10}$, with weights 3/12 and 9/12, respectively. A similar approach is adopted in a number of other studies, including Gerlach (2007), Kortelainen, Paloviita, and Viren (2011), and Siklos (2013).

Such an approximation inevitably involves a reduction in the sample size, typically by a factor of 0.5, because each transformed observation is based on two raw forecast data points. While Dovern and Fritsche (2008) provide evidence that the approach captures the *cross-sectional* dispersion of predictions well, the distribution of the weighted average of two forecasts, each made for different forecast horizons, is likely to have non-standard statistical properties. The behavior of forecasts at long horizons differs significantly from those at short horizons, as will be clear from our later empirical results.

Finally, while the approach we take here is arguably intuitive and generates a small number of interpretable estimates, there are also limitations. For example, multiple data points are required to estimate π^{e*}. Too short a rolling sample will result in an imprecisely estimated perceived anchor, but too long a rolling sample may mask possible break points in the inflation expectations process. The need for multiple data points also implies that the timeliness of the approach is much lower than inflation anchors derived from financial market data, for example.

There are also limitations related to the interpretation of the coefficients. Importantly, the standard error of the perceived long-run anchor is likely to be influenced by the degree of anchoring: weakly anchored expectations would tend to imply that our estimated anchor is poorly identified by, and therefore imprecisely estimated in, the forecast data. However, it is also influenced by how well the model fits the data and the stability of the perceived anchor across the sample. Another drawback of the approach is that it only uses forecast data with horizons of up to twenty-four months to estimate the long-run anchor. If forecast behavior differs significantly at

Figure 3. Decay Function and Perceived Inflation Anchor: United States

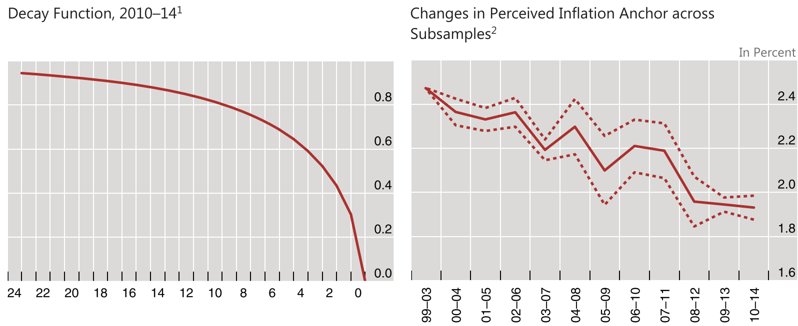

Source: Authors' calculations.
[1]The horizontal axis represents the forecast horizon h, which is the number of months before the end of the calendar year being forecast.
[2]The solid lines represent the rolling five-year sample estimates for the perceived long-run anchor, while the dotted lines are the 95 percent confidence intervals.

longer horizons than twenty-four months, the estimate for π^{e*} could be a biased measure of a perceived long-run anchor. This contrasts with longer-horizon surveys, such as the five-year forecast from Consensus Economics, which potentially provide a "cleaner" measure of long-run expectations.

3. Application to U.S. Data and Comparison with Other Measures of Anchoring

In this section, we first apply our measure to data from the United States, evaluate its robustness, and compare it with other measures of anchoring.

Figure 3 shows the benchmark estimates. The left-hand panel displays the estimated decay function for the 2010–14 sample period, and the right-hand panel shows changes in the estimated level and 95 percent confidence band of the perceived long-run anchor (π^{e*}) over time. It is clear from the shape of the decay function that inflation expectations start to move closer to actual inflation only gradually (left-hand panel). The b and c parameters associated with the shape of the decay function are 4.74 and 0.65, respectively. The estimated

π^{e*} has seen a gradual decline over time (right-hand panel). During the 2010–14 sample, the estimated π^{e*} is 1.9 percent and its standard error is 0.03.

We next demonstrate how well the model fits the data. To do so, for each rolling sample, we calculate the following ratios:

$$R_1 = 1 - \frac{\sum_{t,h}(f(t,t,-h) - \pi(t-h))^2}{\sum_{t,h} f(t,t-h)^2},$$

$$R_2 = 1 - \frac{\sum_{t,h}(f(t,t,-h) - \pi^{e*})^2}{\sum_{t,h} f(t,t-h)^2},$$

$$R_3 = 1 - \frac{\sum_{t,h}(f(t,t,-h) - \alpha(h)\pi^{e*} - [1-\alpha(h)]\pi(t-h))^2}{\sum_{t,h} f(t,t-h)^2}.$$

The first is a measure of the share of forecast variability that is explained by the most recently available level of inflation at the time the forecast is made. The second is the share of variability that is explained by the estimated perceived inflation anchor. And the third is the share of variability that is explained by our model, which is a generalization of the other two. As in the estimation of the model, actual inflation is lagged by one period in each of these calculations to account for publication lag in the inflation data.

While there is no mechanical link between our estimates and minimizing average unexplained variability (we use maximum likelihood, and the variance structure allows the importance of each observation to vary in the estimation), a well-behaved model should generally explain more forecast variability than the alternatives presented above. That is, R_3 should be larger than either R_1 or R_2. Indeed, this is what we find for the United States, for all sample periods (see table 3, in the appendix).

Are there benefits in assuming the flexible functional form yielded by the Weibull, compared with a simple exponential decay function? The left-hand panel of figure 8 (in the appendix) shows the estimated π^{e*} and its standard error for different subsamples, for both an exponential function and the Weibull. While in most periods the estimates are similar, in one subsample both the estimated perceived anchor and its standard error increase abruptly when the exponential form is used, even as there is no sharp change in the underlying forecast data. In that subsample, there is little decay (right-hand panel). Thus, the Weibull yields somewhat more plausible (and more precise) estimates. Moreover, as shown later in section 4, for some economies, allowing a flexible functional form results in decay paths that are quite different from an exponential shape.

Given that we can motivate our model with an autoregressive model of inflation, it is of interest to compare the estimated π^{e*} with the simple intercept of an autoregressive model of inflation. Here, we consider two different horizons, $h = 24$ and $h = 12$. We also make two different assumptions regarding the residuals. In the first case, errors are assumed to be iid, such that $\rho_j = 0$ in equation (10) (in the appendix). In the second specification, an AR(1) error structure is assumed. The left-hand panel of figure 9 (in the appendix) shows our π^{e*} anchor as the solid line. The dotted lines are alternatives based on autoregressive specifications, with forecasts at either twenty-four-month or twelve-month horizons. While many estimated anchors are similar across the different specifications, some rolling samples with an AR(1) term result in unstable estimates. Similar results are obtained for the second moments. The standard errors are nearly always smaller using the Weibull distribution than the other specifications (right-hand panel). This is not surprising, given that our preferred specification makes use of forecasts at all horizons, and so embodies more information.

Is there support in the data for the flexibility assumed for the error variance structure? In particular, equation (6) allows the residual variance to vary non-monotonically with the horizon. Considering all possible rolling samples for the United States, we indeed find evidence of non-monotonicity, with the variance generally largest at around the twelve-month horizon (figure 10, in the appendix). The highest variances occur in subsamples that include the observations for 2008 (i.e., for rolling samples from 2004–08 to 2008–12).

We also note that the estimates are relatively robust to alternative assumptions regarding the error covariance structure. Using the estimates for 2010–14 for the United States, we supplement the equation for the correlation of the residuals (equation (7)) in two ways. First, we add the term $-\Phi_3(h + k)$. This explicitly accounts for the possibility that the correlation of the residuals is either higher or lower at longer forecast horizons, all else equal. Second, we also add the term $-\Phi_4(h + k)^2$, allowing for a non-linear impact of the length of the forecast horizon on the correlation of the residuals.

The estimate of the anchor is robust across these specifications: the estimated π^{e*} changes by less than 0.01. Its standard error is also little changed by the inclusion of $-\Phi_3(h+k)$. However, adding the term $-\Phi_4(h+k)^2$ results in a decline in the standard error from 0.03 to 0.01.

How different is our proposed perceived anchor from other measures of anchoring for the United States? One alternative measure is breakeven inflation rates—an indicator of inflation expectations based on financial asset prices. As breakeven rates are available at a higher frequency than forecasts, they could provide a more timely indication of shifts in anchoring than our measure. But previous literature has also raised concerns about financial-market-based measures of expectations. One issue is their volatility.[16] A second concern is that the breakeven rates between real and nominal bonds are affected by various factors: expected inflation, inflation risk premiums, liquidity premiums, and technical market factors (Hördahl 2009). Especially at times of crisis, liquidity risk premiums could play a dominant role, complicating the extraction of inflation expectations from data on nominal and real bonds (see also Bauer and Rudebusch 2015; Liu et al. 2015).[17]

In figure 4, we compare our estimated π^{e*} for all rolling five-year subsamples against the average breakeven inflation rates over the periods where they are both available. The daily variation in the breakeven rates is taken into account by plotting +/− two standard deviations around the mean; similarly, uncertainty around the estimated perceived anchor is illustrated by 95 percent confidence intervals.

Figure 4 illustrates the excessive volatility of breakeven rates. The high variation is particularly prevalent in their wide confidence bands, as shown in the left-hand panel.[18] For every subsample, breakeven rates are more volatile than our estimate of π^{e*}. However, both measures are less volatile than actual inflation outcomes,

[16]Faust and Wright (2013) contend that inflation expectations extracted from inflation-indexed financial market instruments are too volatile to represent either rational long-run inflation expectations or the target of a central bank.

[17]Shen (2006) provides earlier evidence about the impact of liquidity risk premiums on breakeven inflation rates.

[18]Note also the widening of the confidence bands during the recent crisis period.

Figure 4. Comparing Breakevens with Other Measures: United States

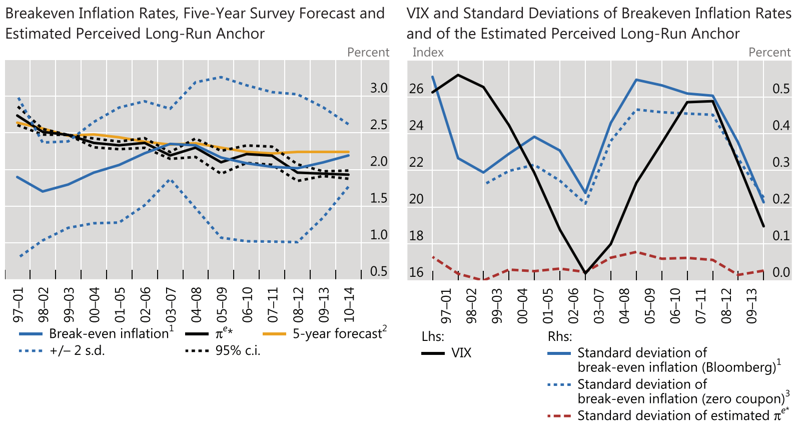

Sources: Gürkaynak, Sack, and Wright (2008) data updated from https://www.federalreserve.gov/econresdata/researchdata/feds200805.xls; Bloomberg; Consensus Economics; national data; authors' calculations.

[1]Calculated using data from Bloomberg. Breakeven rates are derived using generic Treasury yield and generic government inflation-indexed bonds at the ten-year maturity for 1997–98, generic breakeven rates at the ten-year maturity thereafter.

[2]Average of forecasts made in October during each five-year period.

[3]Calculated using updated data set based on Gürkaynak, Sack, and Wright (2008).

as would be expected of indicators of perceived long-run inflation anchors (see table 4, in the appendix).

The volatility of breakeven rates is positively correlated with a measure of financial market risk aversion, the VIX index, illustrating the potential relevance of financial market factors (figure 4, right-hand panel). The correlation coefficient between the mean value for the VIX index and the standard deviation of breakeven rates, across the various subsamples, is 0.4. In contrast, the correlation between the mean value for the VIX index and the standard deviation of the estimated π^{e*} is very close to zero. These results are robust to using a measure of breakevens based on nominal and zero-coupon yield curves, instead of the generic Bloomberg measure (figure 4, right-hand panel and table 4, which appears in the appendix). The wide confidence bands for breakeven rates are also obtained if we use five-year, five-year forward breakeven inflation rates, although the sample is shorter in this case (see the right-hand panel of figure 11, in the appendix).[19]

Another alternative measure of anchoring is given by long-run surveys. The level of our estimated π^{e*} is similar to the five-year-ahead forecasts from Consensus Economics (figure 4, left-hand panel).[20] The largest difference between the two measures, across the different subsamples, amounts to 0.31 percentage points. However, the five-year-ahead forecast falls within the 95 percent confidence bands of the estimated π^{e*} in only one-third of the subsamples.

Our estimate of the perceived inflation anchor offers several advantages over long-run forecasts. First, the π^{e*} estimates contain information that is not included in long-run surveys: whether an anchor can be identified in the data and how precisely the anchor is estimated. In contrast, there is no measure of uncertainty available for the five-year forecasts from Consensus Economics, as only the mean forecasts are published. Second, for some other economies, long-run expectations from surveys have remained constant for many consecutive years, suggesting that their information value may be limited at times.[21] In addition, in periods of high uncertainty, forecasters may be reluctant to offer long-run forecasts.[22] However, it is possible that the five-year-ahead expectations may yield a better point forecast of the long-run inflation anchor in the presence of persistent inflation shock processes, as π^{e*} is estimated using forecast data only up to twenty-four months.

In sum, the analysis performed for the United States suggests that our measure can be considered complementary to existing measures of anchoring.

[19] The left-hand panel of figure 11 (in the appendix) contains the equivalent graph based on U.K. data, and is qualitatively similar.

[20] The five-year-ahead expectations are those made in October for the calendar year five years ahead. For some economies included in section 4, the long-run forecasts are made in September. The years shown in the graph correspond to the year in which the long-run forecast is made.

[21] For example, in the case of Canada, five-year-ahead Consensus Forecasts for CPI inflation remained unchanged at 2.0 percent for seven successive years.

[22] For example, in the Consensus Economics database, there is no six- to ten-year-ahead inflation forecast for Japan made in early 2001. The April 2001 report states that “few of our panellists were prepared to predict economic trends beyond 2005 at this time.”

4. Multiple-Country Evidence

Next, we use our measure to evaluate the anchoring of inflation expectations in a large panel of economies. We include all economies for which forecasts of inflation for individual forecasters, and corresponding monthly series on year-on-year inflation outcomes, are available no later than 2008. This yields a sample of forty-four economies (including the United States).[23] Our sample is considerably larger than those considered in previous studies that have analyzed long-run anchoring of inflation expectations. At market exchange rates and prices prevailing in 2014, the forty-four economies included in our study accounted for 88 percent of world GDP.[24]

Data are available commencing with forecasts for the 1990 calendar year for some of the industrialized economies in the sample, while for emerging markets the sample periods generally start later.[25] Table 1 displays data availability.

In some of our reported results, including all available forecasts of inflation from Consensus Economics for the relevant sample period means that we have missing observations in our time-series data. For example, for most Latin American economies, forecasts are only available for even months at the beginning of the sample, before switching to odd months, and finally all months partway through 2002. For some Eastern European economies, forecasts are only available for even months before 2007. We take explicit account of this in our estimation by setting the contribution to the likelihood function to zero for missing observations.

We illustrate the estimated decay functions in figure 5 for two large advanced economies (the euro area and Japan) and two large

[23]We require forecasts at the individual forecaster level in order to construct the median forecasts that we use in our study. Where possible, we then backdate these series using average forecasts published by Consensus Economics for earlier years. Australia and New Zealand are excluded from our panel because of a lack of monthly data on inflation outcomes.

[24]Based on data from the International Monetary Fund's World Economic Outlook. For this calculation, the euro area is regarded as one economy and the individual euro-area member states included in our sample are omitted to avoid double-counting.

[25]For a few advanced economies there are also three forecasts, made in late 1989, for the 1989 calendar year. We do not include these in our sample.

Table 1. Data

Country	Data Available From
Argentina	1993
Brazil	1990
Bulgaria	1995
Canada	1990
Chile	1993
China	1994
Chinese Taipei	1990
Colombia	1993
Croatia	1998
Czech Republic	1995
Estonia	1998
Euro Area	1998
France	1990
Germany	1990
Hong Kong SAR	1990
Hungary	1990
India[1]	1994
Indonesia	1990
Italy	1990
Japan	1990
Korea	1990
Latvia	1998
Lithuania	1998
Malaysia	1990
Mexico	1990
Netherlands	1990
Norway	1990
Peru	1993
Philippines	1994
Poland	1990
Romania	1995
Russia	1995
Singapore	1990
Slovakia	1995
Slovenia	1995
Spain	1990
Sweden	1990
Switzerland	1990
Thailand	1990
Turkey	1995
Ukraine	1995
United Kingdom	1990
United States	1990
Venezuela	1993

[1]Indian data (forecasts and outcomes) are based on March years (all other economies are based on calendar years).

Figure 5. Decay Function and Estimated Perceived Long-Run Anchor, 2010–14

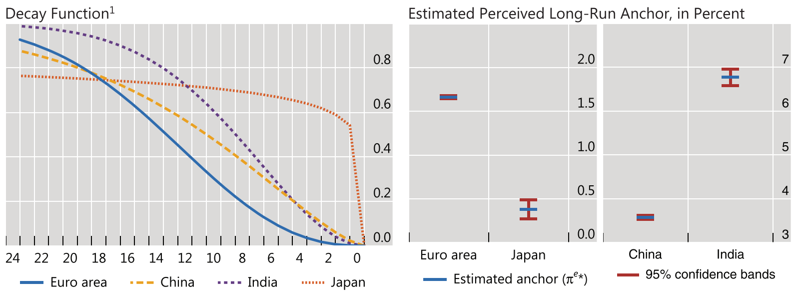

Source: Authors' calculations.
[1]The horizontal axis represents the forecast horizon h, which is the number of months before the end of the calendar year being forecast.

emerging market economies (China and India). These estimates are based on the most recent five-year sample period, 2010–14.

At the longest horizon of twenty-four months, the weight on the perceived long-run anchor is the lowest for Japan but exceeds 0.75 for all four economies. The decay functions for China and India are similar in shape across the different forecast horizons. Japan stands out from the other economies in that the weight on the perceived long-run anchor stays broadly constant over the longer forecast horizons, before declining rapidly at the very shortest horizons.

There are also differences between the economies in terms of the estimated level and standard error of the perceived long-run anchor, π^{e*}. For the economies shown in figure 5, the estimated level of the anchor over 2010–14 is lowest for Japan, at 0.38 percent. However, the standard error is higher in Japan than in the euro area. The level of the perceived anchor for the euro area, at 1.7 percent, is arguably consistent with the European Central Bank's definition of price stability as inflation "below, but close to, 2 percent." The estimated π^{e*} for China is 3.6 percent. Finally, inflation expectations are anchored at a level of 6.8 percent for India, with a relatively large standard error (of around 0.10 percent).

Rolling regression estimates starting with the 1999–2003 subsample are shown in figure 6. The perceived long-run anchor has been relatively stable over time in the euro area, but there is somewhat

Figure 6. Changes in Estimated Perceived Inflation Anchor in Four Major Economies (in percent)

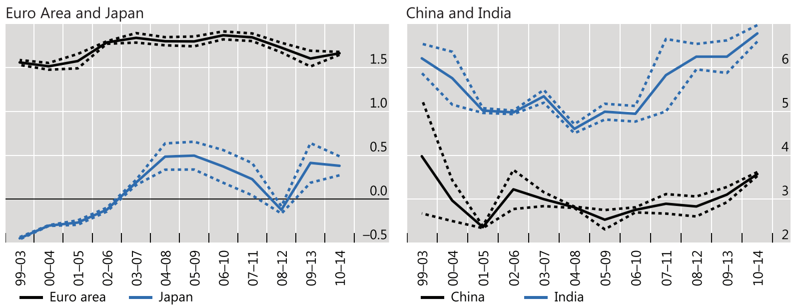

Source: Authors' calculations.
Note: The solid lines represent the rolling five-year sample estimates for the perceived long-run anchor, while the dotted lines are the 95 percent confidence intervals.

more variation in Japan, where the long-run anchor has fluctuated around 0 percent. The standard errors for Japan have also varied over time. Echoing the result from the previous graph, the π^{e*} estimates for Japan display larger standard errors than those for the euro area for samples starting in the mid-2000s.[26]

Key parameter estimates for the 2010–14 period for all economies are given in table 2. A perceived long-run anchor cannot be identified for Argentina and Venezuela during this time period, implying that expectations were poorly anchored.[27] For the other economies, the mean estimate for π^{e*} is 3.0 percent and the median is 2.6 percent;

[26]Kamada, Nakajima, and Nishiguchi (2015) show that Japanese households' long-run (five-year) inflation expectations are insensitive to actual inflation, in contrast to shorter-run expectations. While one could therefore interpret the fluctuations in our estimated perceived anchor as reflecting the fact that the longest forecast horizon used in our methodology is only twenty-four months, the five-year-ahead Consensus Forecasts for Japan also shift over time. For example, the five-year-ahead forecast for inflation made in October 2003 (0.5 percent) is 1 percentage point lower than the forecast made in October 2006 (1.5 percent). However, the levels of the five-year forecasts for Japan are higher than the point estimates of our perceived long-run anchor for every five-year rolling sample that we estimate, and also systematically higher than inflation outcomes.

[27]The estimates of b and c are such that the weight on the perceived anchor is close to zero for all forecast horizons up to twenty-four months.

Table 2. Estimation Results, 2010–14

Country	b	c	π^{e*}	S.E. (π^{e*})
Argentina	Not Identified			
Brazil	0.67	17.50	5.06	0.04
Bulgaria	13.06	1.14	3.29	0.03
Canada	6.20	0.54	2.14	0.03
Chile	0.05	0.24	2.98	0.00
China	14.22	1.41	3.57	0.02
Chinese Taipei	5.28	2.54	1.75	0.00
Colombia	11.34	0.48	3.57	0.03
Croatia	10.46	0.60	3.03	0.05
Czech Republic	3.25	0.23	1.41	0.08
Estonia	11.66	0.79	2.31	0.07
Euro Area	16.06	2.39	1.66	0.01
France	11.91	3.46	1.61	0.01
Germany	12.88	2.82	1.79	0.00
Hong Kong SAR	3.73	0.34	3.68	0.08
Hungary	18.83	5.34	3.06	0.07
India	10.89	1.87	6.77	0.10
Indonesia	19.39	0.08	6.05	0.02
Italy	31.92	0.61	1.78	0.02
Japan	3.60	0.19	0.38	0.06
Korea	12.96	1.59	3.09	0.01
Latvia	10.53	0.44	1.68	0.14
Lithuania	8.35	0.87	2.19	0.05
Malaysia	7.76	0.27	2.69	0.13
Mexico	13.81	9.86	3.68	0.02
Netherlands	41.96	0.24	1.69	0.10
Norway	8.81	2.17	1.79	0.02
Peru	1.15	0.66	2.67	0.04
Philippines	0.86	5.81	4.01	0.03
Poland	13.95	3.02	2.60	0.04
Romania	19.33	1.19	3.05	0.08
Russia	0.87	6.94	6.88	0.15
Singapore	17.83	1.22	2.48	0.05
Slovakia	11.17	2.38	2.99	0.01
Slovenia	13.40	3.98	2.14	0.02
Spain	16.06	4.37	1.65	0.02
Sweden	14.50	2.61	1.88	0.04
Switzerland	16.41	2.10	1.08	0.09
Thailand	12.54	1.30	3.24	0.03
Turkey	6.34	0.75	6.77	0.01
Ukraine	15.71	1.30	8.68	0.13
United Kingdom	3.44	0.47	1.89	0.03
United States	4.74	0.65	1.93	0.03
Venezuela	Not Identified			

Source: Authors' calculations.
Note: S.E. indicates standard error.

the mean and median standard errors for π^{e*} are 0.05 percent and 0.03 percent, respectively. The decay functions for all economies in the sample, for the same time period, are shown in figure 12, in the appendix.

How good is the fit of the estimated model in the large panel of economies? Using identical measures as in section 3 for the United States, we find that the model (ratio R_3) explains more forecast variability than the alternatives R_1 (variability explained by the recent level of inflation) and R_2 (variability explained by the estimated perceived anchor). This is the case for most of the 762 country rolling-sample pairs that we consider. Excluding the few cases where no inflation anchor can be identified in the data, R_3 is larger than R_1 in 96 percent of cases and larger than R_2 in 97 percent of cases. Put another way, the median R_1 is 0.96 versus 0.87 for each of the other two.[28]

We can go further and calculate these ratios for a limited number of horizons to highlight the difference in the role of the perceived anchor versus lagged inflation at different horizons. At the longest four horizons, the median R_3 (0.98) is only slightly larger than R_2 (0.96) but much larger than R_1 (0.78), indicating that the anchor is playing the main role in fitting the data. In contrast, at the shortest four horizons, the median R_1 and R_3 are the same to two decimal places (0.98), versus a much lower R_2 (0.79), illustrating the preponderance of recent inflation over the perceived anchor in explaining near-term inflation forecasts.

The full set of estimates of the perceived anchor, and 95 percent confidence bands, for all five-year rolling samples, are displayed in figure 13, in the appendix. Where no anchor can be identified (because $\alpha(h)$ is close to zero at all horizons up to twenty-four months), this is indicated by an "x" close to the horizontal axis. These cases, constituting around 4 percent of all rolling samples, are those in which a model with any weight on an anchor does not fit the data well. This typically occurs when an economy experiences very high levels of inflation within the five-year rolling sample—in most cases with peak inflation exceeding 100 percent—and includes

[28]The panel of economies considered here includes the United States. The detailed results for goodness of fit for each country rolling-sample pair are not reported here to save space, but are available on request from the authors.

rolling samples for Argentina, Brazil, Mexico, Russia, Turkey, and Venezuela. The inability to identify an anchor in these cases is not surprising, given that inflation forecasts are unlikely to be anchored when inflation outcomes are high and volatile.

More generally, the estimated perceived anchors display a lot of variation over the rolling samples. Some economies have seen dramatic declines in the estimated anchor, from double digits to less than 5 percent (e.g., Bulgaria, Columbia, Hungary, Mexico, and Poland). In most cases, inflation forecasts have been stably anchored at levels below 5 percent in recent rolling samples. But there are a few countries where inflation expectations remain anchored to levels above 5 percent (Brazil, India, Indonesia, Russia, Turkey, and Ukraine) or an anchor cannot be identified in the latest rolling samples (Argentina and Venezuela). Finally, in some economies the perceived anchor has fallen in the most recent rolling samples (including the Czech Republic and Spain), suggesting that recent low-inflation outcomes in these economies are becoming embedded in inflation expectations.

As a comparison, figure 13 also displays the average long-run (five-year) inflation forecasts for the same rolling samples, where available from Consensus.[29] In some cases, five-year inflation forecasts look a lot like the estimated perceived anchors, including in China, Germany, India, and Mexico. However, that is not always the case. In particular, the longer-term inflation forecasts appear to imply some reversion to the mean: in cases where the inflation rate is relatively high by global standards (e.g., Indonesia and Ukraine), the longer-term inflation forecasts tend to be systematically lower than the estimated perceived anchors. In contrast, where the inflation rate is relatively low (e.g., the euro area, Japan, and Switzerland since 2000), the reverse is true. And the confidence band around the estimated anchor provides information that cannot be deduced from the long-term inflation forecast. For example, the dramatic widening of the confidence band for Malaysia in the aftermath of the Asian crisis indicates less-stable inflation expectations. Similarly, the confidence band for Japan is now wider than one decade ago.

[29]For the United Kingdom the forecasts are of the RPIX inflation rate for 1999–2003 and the CPI inflation rate thereafter.

Figure 7. Breakeven Inflation Rates and Estimated Perceived Long-Run Anchor (in percent)

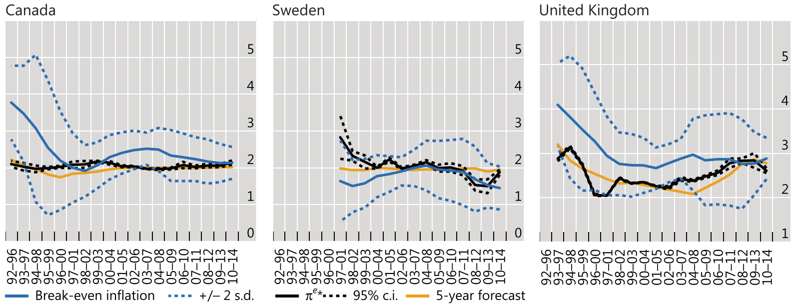

Sources: Bloomberg; Consensus Economics; national data; authors' calculations.
Notes: For Canada, breakeven inflation is obtained as the long-term nominal benchmark government bond yield less inflation-linked government bond yield; for Sweden, nominal bond yield (breakeven comparator from Bloomberg) less inflation-linked bond yield (all maturities); for the United Kingdom, generic breakeven rates at the ten-year maturity. For five-year forecasts, average of forecasts made in September/October during each five-year period. The five-year forecast for the United Kingdom refers to the RPIX for 1999–2003 and to the consumer price index from 2004 onwards.

In figure 7, we also include the breakeven rates for three economies: Canada, Sweden, and the United Kingdom.[30] Similarly to the United States, in these countries both real and nominal bonds have existed for a relatively long period, and markets for these instruments are considered to be rather liquid.[31]

Echoing the findings in section 3, breakeven inflation rates for the three economies are quite volatile. This is illustrated both by variation in the average breakeven rate over the sample and the

[30]Note that there is a small definitional difference between the displayed measures of inflation in the case of the United Kingdom: breakeven rates purported compensate for retail price (RPI) inflation, while the perceived long-run anchor is estimated based on forecasts and outcomes of RPIX inflation (excluding the effect of changes in interest rates).

[31]In a recent paper, De Pooter et al. (2014) construct financial-market-based measures of inflation expectations for Brazil, Chile, and Mexico from nominal and inflation-linked bonds in order to examine the anchoring of long-run inflation expectations, and they report improved anchoring in all three countries. Another country with available data on nominal and inflation-linked bonds is Japan, but only starting in 2008.

wide confidence bands. The confidence bands tend to widen during the international financial crisis. Moreover, there is some evidence of bias in the average levels of breakeven inflation rates relative to our estimated perceived anchor. This is especially the case for the United Kingdom and Canada, where real bonds are purchased primarily by pension funds and life insurance companies, in part due to favorable tax treatment (Côté et al. 1996). This pushes down yields on real bonds so that breakeven inflation rates in Canada tend to be higher than our estimated π^{e*}.[32]

5. Conclusion

In this paper, we have proposed a novel methodology to assess the behavior of inflation expectations. Forecast data are modeled using a decay function such that inflation forecasts monotonically diverge from a perceived long-run anchor towards actual inflation as the forecast horizon shortens. Our methodology provides a parsimonious framework for fitting inflation forecasts that fully utilizes the multiple-horizon dimension of the data. It also generates a small number of estimates that can be used to draw inference about the degree of anchoring of inflation expectations. We find empirical support for the model in a large panel of forecast data for forty-four economies.

There are at least two lines of possible future work that could build on our modeling approach. First, we have focused on the long-run anchor of inflation expectations, but we could also use our framework to examine the behavior of expectations at shorter horizons that we have not explored here, characterized by our estimates of b and c. Second, our use of median forecasts for each economy in this paper may hide some important differences across forecasters. A natural extension for future research would be to model the behavior of inflation expectations at a forecaster level using our functional form. This could be used to identify differences in expectations behavior both across individual forecasters and for the same forecaster across

[32] See also Christensen, Dion, and Reid (2004), who find that the breakeven rates in Canada have been higher and more variable than survey measures of expectations during some years, due to movements in risk premiums and other factors not related to inflation expectations.

time, and provide further insights into the anchoring of inflation expectations.[33]

Appendix. An Autoregressive Model of Inflation Expectations

The AR(p) case can be written as

$$\pi_t = \pi^{e*} + \varepsilon_t, \tag{9}$$

$$\varepsilon_t = \sum_{j=1}^{p} \rho_j \varepsilon_{t-j} + \eta_t. \tag{10}$$

In this case, expectations can be written as

$$\begin{aligned} E_{t-j}(\pi_t) &= \pi^{e*} + f_{j0}\varepsilon_{t-j} + f_{j1}\varepsilon_{t-j-1} + \cdots + f_{jp-1}\varepsilon_{t-j-p+1} \\ &= \pi^{e*} + f_{j0}(\pi_{t-j} - \pi^{e*}) + f_{j1}\varepsilon_{t-j-1} + f_{j2}\varepsilon_{t-j-2} \\ &\quad + \cdots + f_{jp-1}\varepsilon_{t-j-p+1} \\ &= (1 - f_{j0})\pi^{e*} + f_{j0}\pi_{t-j} + f_{j1}\varepsilon_{t-j-1} + f_{j2}\varepsilon_{t-j-2} \\ &\quad + \cdots + f_{jp-1}\varepsilon_{t-j-p+1} \\ &= (1 - f_{j0})\pi^{e*} + f_{j0}\pi_{t-j} + e_{t,t-j}, \end{aligned}$$

where the f's are functions of the ρ's and $e_{t,t-j} = \sum_{k=1}^{p-1} f_{jk}\varepsilon_{t-j-k}$.

Under plausible assumptions about the parameters—for example, that $0 < \rho < 1$ and $\rho_{j+1} < \rho_j$—we have $f_{jk} > f_{j+1k}$ and $f_{jk} > f_{jk+1}$. This implies that the weight on the perceived anchor is increasing in the forecast horizon and the variance of the error term is decreasing monotonically in the forecast horizon. Regarding the latter condition, the estimated model allows for a monotonically decreasing error variance, but no such restriction is imposed in the analysis.

[33] See Yetman (2017) and Hattori and Yetman (2017) for recent evidence along these lines, for forecasts of the Canadian and U.S. economies, and the Japanese economy, respectively.

Figure 8. Weibull Function and Exponential Decay Function: United States

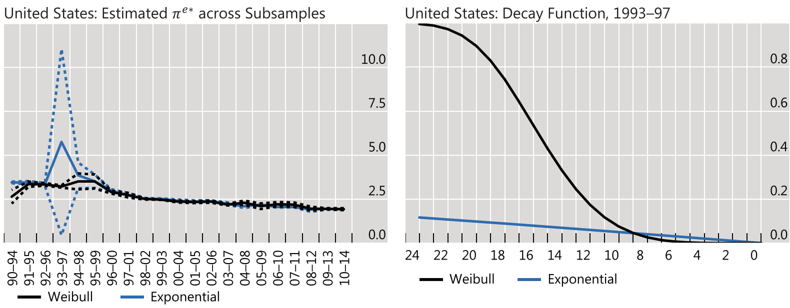

Source: Authors' calculations.
Note: The dotted lines are the 95 percent confidence intervals.

Figure 9. Weibull Function and an Autoregressive Model for Inflation: United States

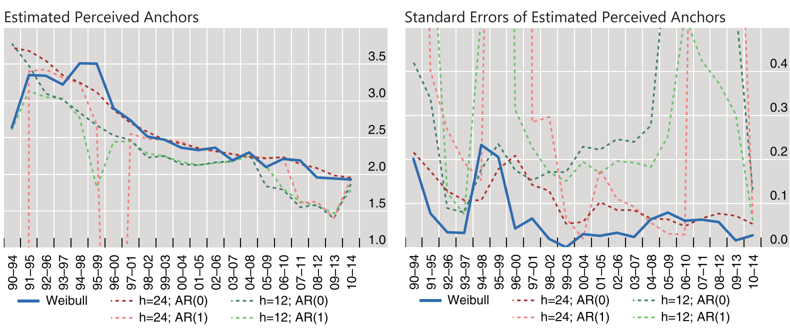

Source: Authors' calculations.
Note: See the main text and the appendix for a definition of the autoregressive model of inflation.

Figure 10. Variance of Residual

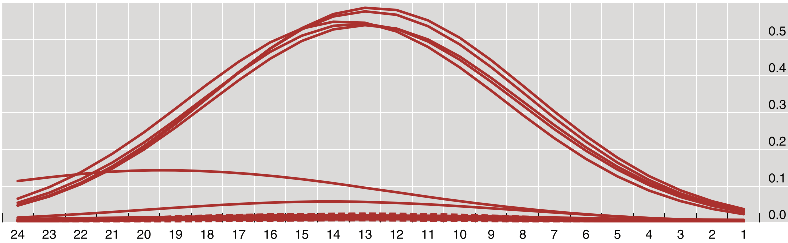

Note: The graph shows the variance of the estimated residual for all the five-year subsamples for the United States, at different forecast horizons.

Figure 11. Five-Year, Five-Year Forward Breakeven Inflation Rates and Estimated Perceived Long-Run Anchor (in percent)

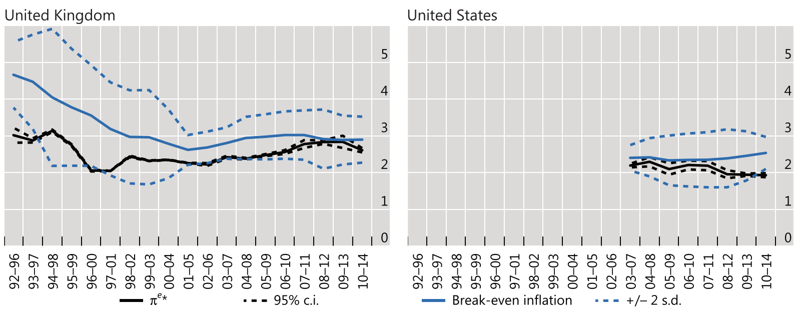

Sources: Bank of England; Federal Reserve Bank of St Louis; national data; authors' calculations.
Note: United Kingdom: five-year, five-year breakeven inflation rate implied from the gilt market as calculated by the Bank of England; United States: five-year, five-year forward inflation expectation rate from the Federal Reserve Bank of St. Louis.

Figure 12. Forecasts of Headline Inflation at Different Horizons[1]

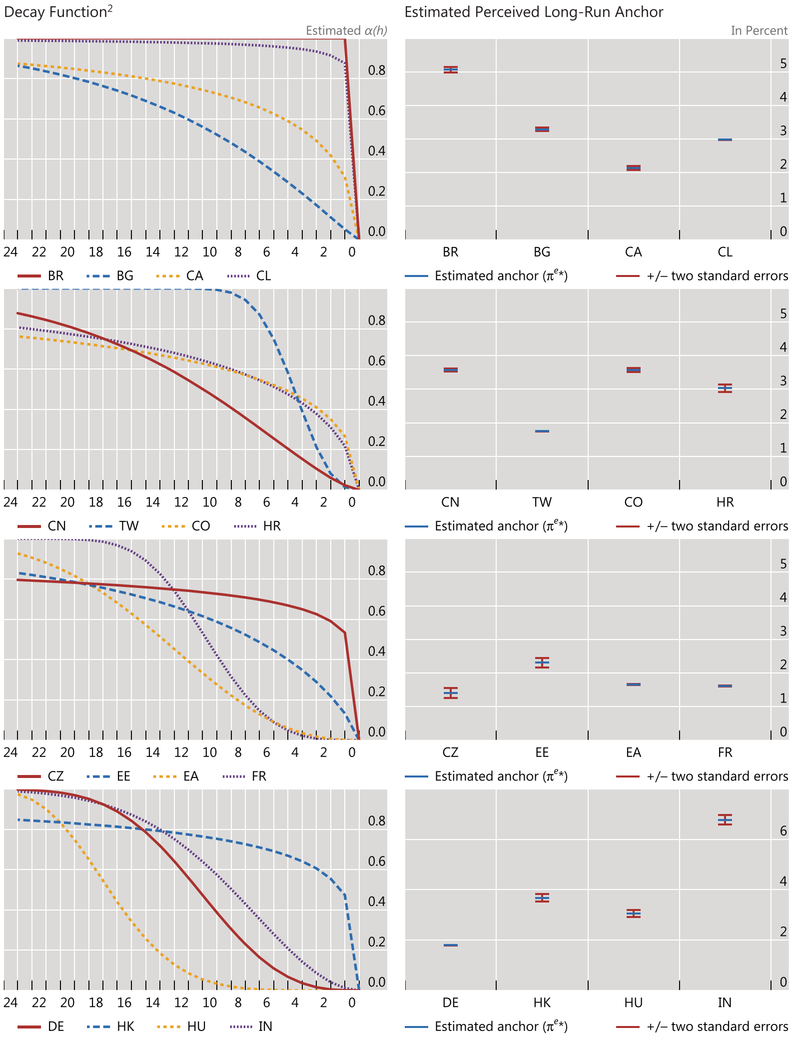

(continued)

Figure 12. (Continued)

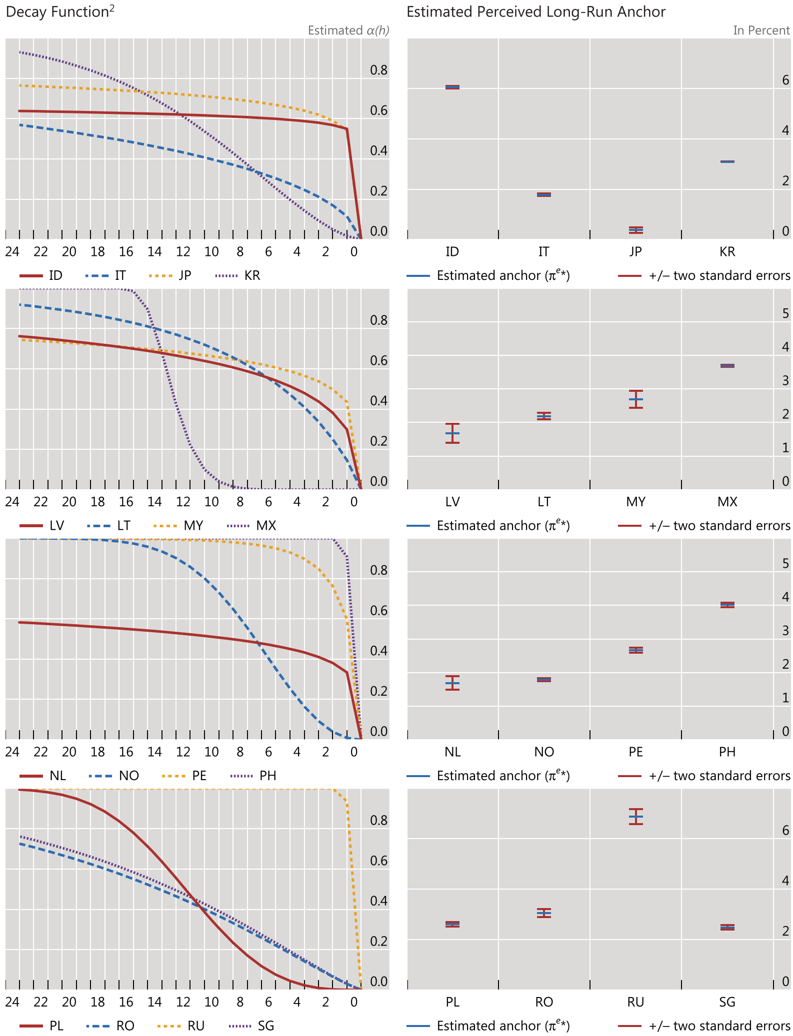

(continued)

Figure 12. (Continued)

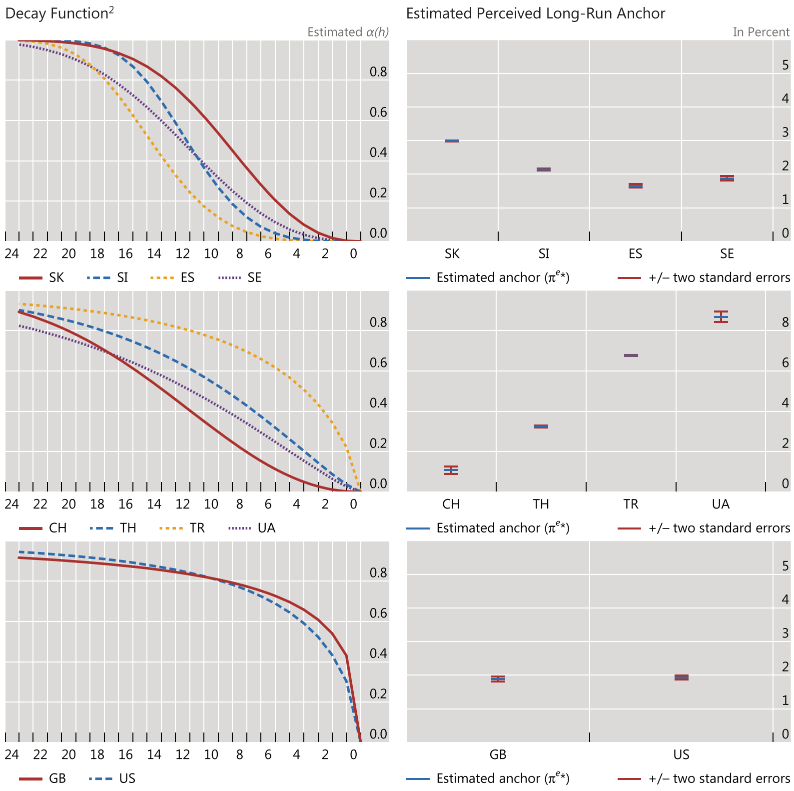

Sources: Authors' calculations.
Note: BR = Brazil; BG = Bulgaria; CA = Canada; CL = Chile; CN = China; TW = Chinese Taipei; CO = Colombia; HR = Croatia; CZ = Czech Republic; EE = Estonia; EA = Euro area; FR = France; DE = Germany; HK = Hong Kong SAR; HU = Hungary; IN = India; ID = Indonesia; IT = Italy; JP = Japan; KR = Korea; LV = Latvia; LT = Lithuania; MY = Malaysia; MX = Mexico; NL = Netherlands; NO = Norway; PE = Peru; PH = Philippines; PL = Poland; RO = Romania; RU = Russia; SG = Singapore; SK = Slovak Republic; SI = Slovenia; ES = Spain; SE = Sweden; CH = Switzerland; TH = Thailand; TR = Turkey; UA = Ukraine; GB = United Kingdom; US = United States.
[1]Based on an estimation period of 2010–14.
[2]The horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.

Figure 13. Inflation: Estimated Perceived Anchors and Five-Year Forecasts (in percent)

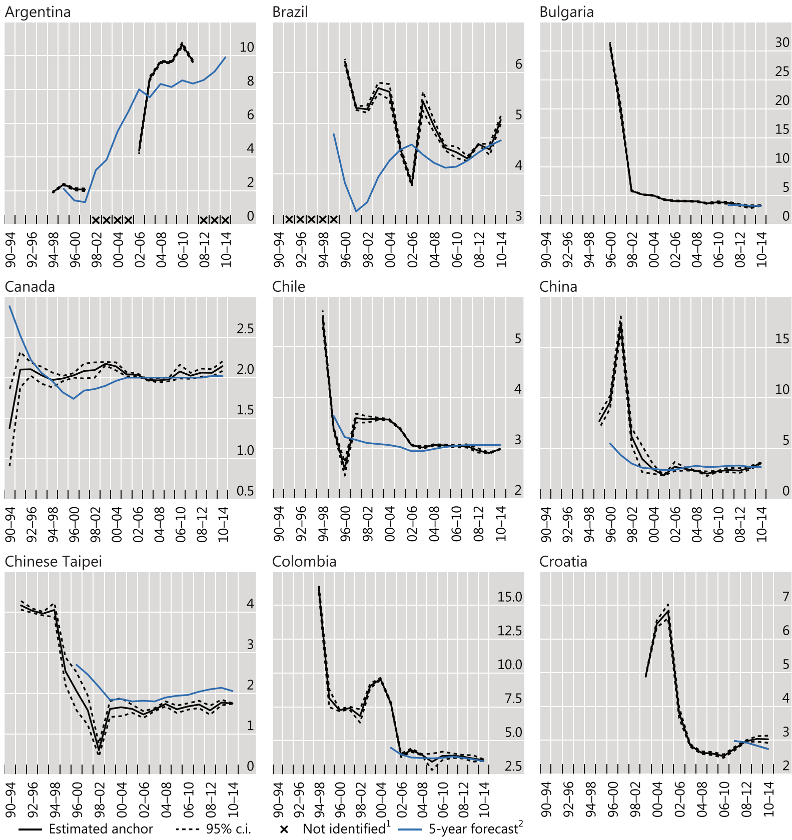

(continued)

Figure 13. (Continued)

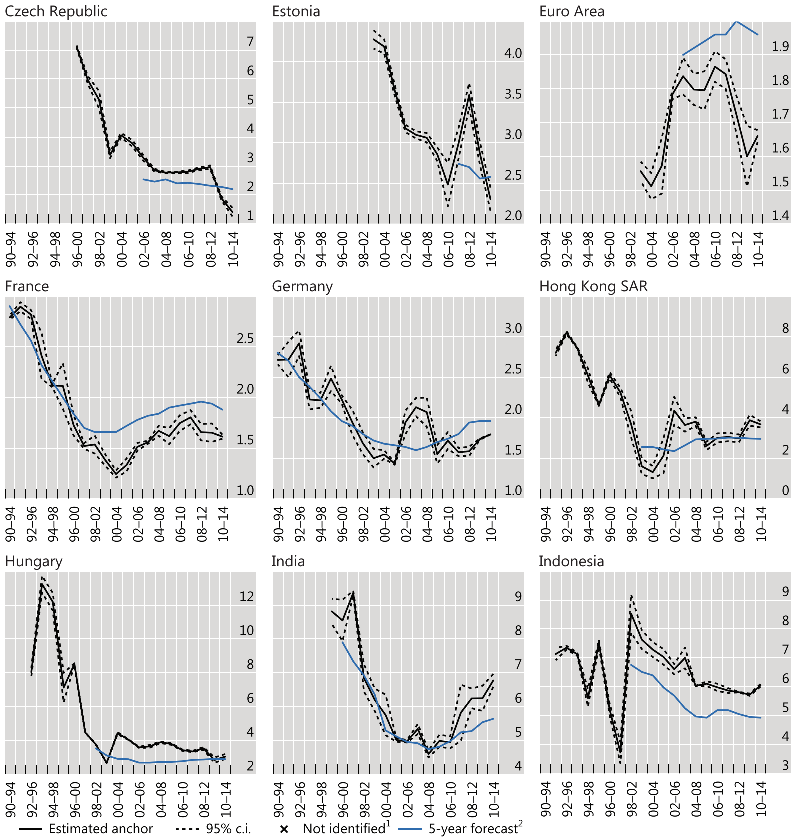

(continued)

Figure 13. (Continued)

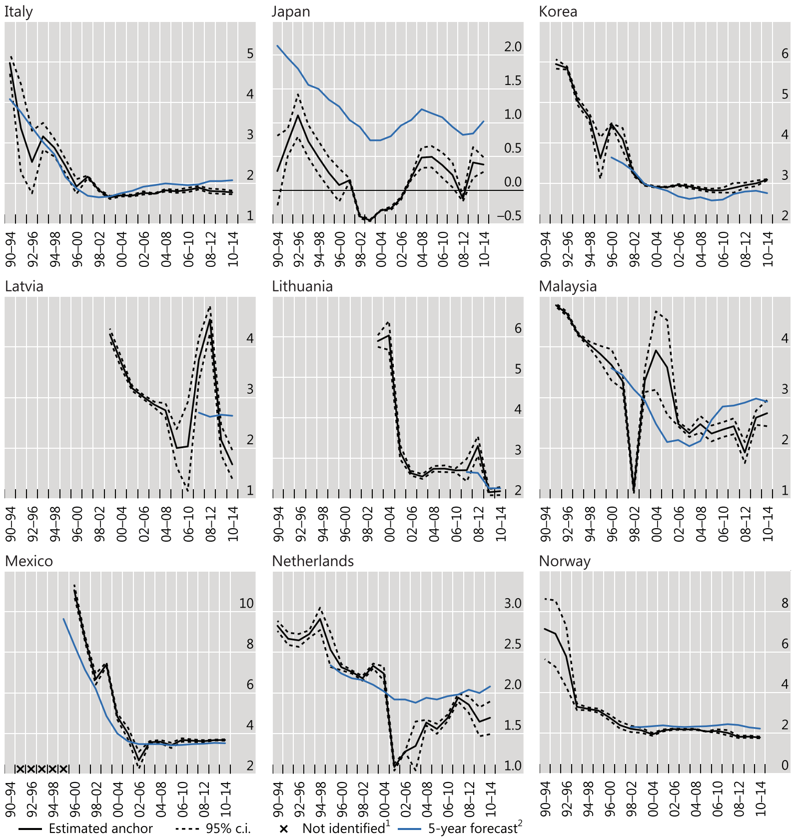

(continued)

Figure 13. (Continued)

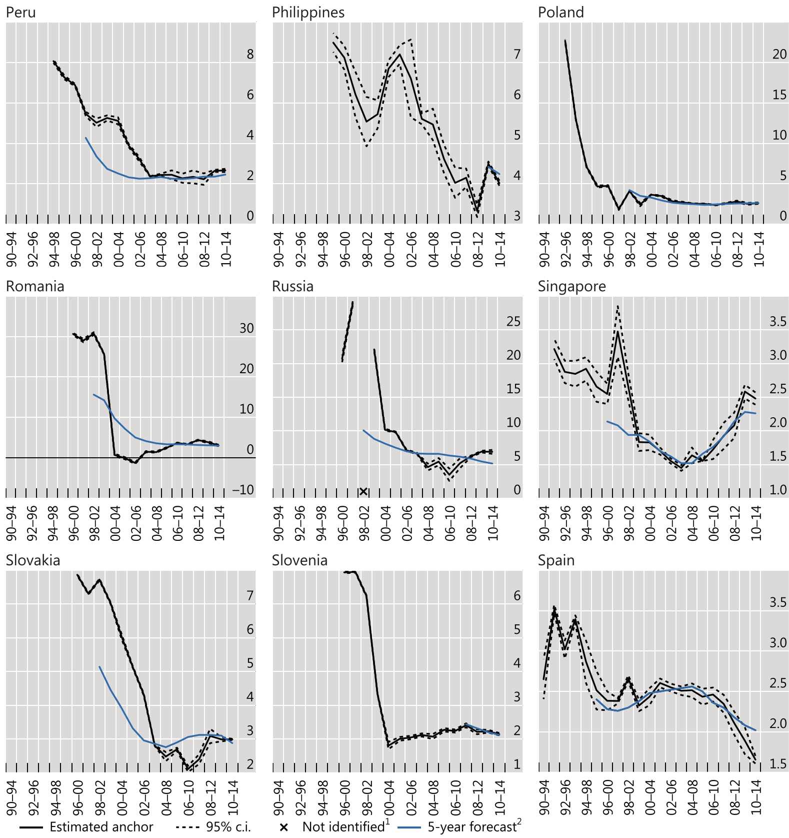

(continued)

Figure 13. (Continued)

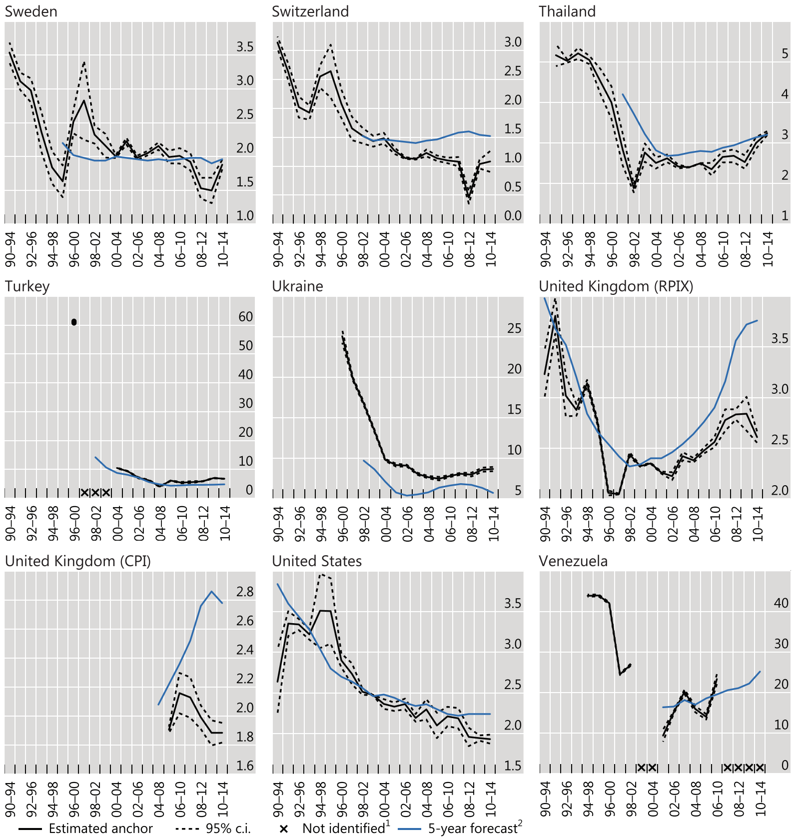

Sources: Consensus Economics; national data; authors' calculations.

[1]Not identified indicates that the model estimates imply little or no weight on an anchor; instead, a model where forecasts only reflect recent inflation fits the data best.

[2]Average for each five-year window. For five-year forecasts, average of forecasts made in September/October during each five-year period. For the United Kingdom, two estimates of the perceived anchor are generated based on CPI and RPIX inflation. The five-year forecast for the United Kingdom refers to the RPIX until the period 1999–2003 and to the consumer price index from 2004 onwards.

Table 3. Measures of Goodness of Fit, United States: Ratios

Sample	R_1	R_2	R_3
1990–94	0.965	0.918	0.966
1991–95	0.964	0.989	0.996
1992–96	0.988	0.984	0.996
1993–97	0.987	0.960	0.994
1994–98	0.980	0.887	0.991
1995–99	0.978	0.863	0.993
1996–2000	0.977	0.957	0.992
1997–2001	0.949	0.939	0.973
1998–2002	0.942	0.967	0.986
1999–2003	0.941	0.960	0.978
2000–04	0.949	0.961	0.980
2001–05	0.949	0.954	0.982
2002–06	0.946	0.964	0.984
2003–07	0.954	0.918	0.985
2004–08	0.801	0.800	0.858
2005–09	0.704	0.777	0.843
2006–10	0.636	0.757	0.815
2007–11	0.632	0.747	0.806
2008–12	0.478	0.768	0.793
2009–13	0.767	0.960	0.973
2010–14	0.864	0.957	0.970

Note: See text for details.

Table 4. Standard Deviations of Inflation and Inflation Expectations Measures, United States

Sample	Inflation Outcomes	Breakevens	π^{e*}
1997–2001	0.657	0.555	0.065
1998–2002	0.699	0.333	0.019
1999–2003	0.656	0.294	0.001
2000–04	0.595	0.345	0.030
2001–05	0.616	0.392	0.026
2002–06	0.695	0.354	0.033
2003–07	0.580	0.239	0.024
2004–08	0.631	0.429	0.063
2005–09	1.175	0.547	0.078
2006–10	1.410	0.531	0.060
2007–11	1.310	0.509	0.062
2008–12	1.330	0.503	0.057
2009–13	1.046	0.376	0.016
2010–14	0.767	0.213	0.027

Notes: The standard deviation of inflation outcomes for each subsample is based on monthly year-on-year inflation rates. The standard deviation of inflation breakevens for each subsample is based on daily data on breakeven rates from Bloomberg. The standard deviation of π^{e*} is based on the estimates described in the paper.

Table 5. Measures of Breakeven Rates, United States

Sample	Bloomberg	Zero-Coupon Yields	5y-5y Forward
1997–2001	1.898 (0.789, 3.008)		
1998–2002	1.702 (1.035, 2.369)		
1999–2003	1.795 (1.207, 2.384)	2.088 (1.566, 2.609)	
2000–04	1.959 (1.268, 2.650)	2.214 (1.617, 2.811)	
2001–05	2.062 (1.278, 2.845)	2.295 (1.664, 2.926)	
2002–06	2.223 (1.515, 2.931)	2.402 (1.858, 2.946)	
2003–07	2.351 (1.873, 2.829)	2.469 (2.050, 2.888)	2.403 (2.055, 2.751)
2004–08	2.332 (1.474, 3.191)	2.450 (1.689, 3.211)	2.419 (1.898, 2.940)
2005–09	2.166 (1.072, 3.260)	2.301 (1.369, 3.232)	2.334 (1.658, 3.010)
2006–10	2.087 (1.024, 3.150)	2.239 (1.322, 3.156)	2.345 (1.627, 3.063)
2007–11	2.037 (1.018, 3.056)	2.210 (1.301, 3.119)	2.353 (1.603, 3.104)
2008–12	2.020 (1.013, 3.027)	2.192 (1.288, 3.095)	2.392 (1.601, 3.182)
2009–13	2.097 (1.346, 2.849)	2.239 (1.574, 2.904)	2.457 (1.789, 3.126)
2010–14	2.192 (1.766, 2.619)	2.290 (1.839, 2.742)	2.536 (2.101, 2.971)

Note: The table shows the means over the subsamples, using daily data, with the mean +/− two standard deviations in parentheses.

References

Bauer, M. D., and G. D. Rudebusch. 2015. "Optimal Policy and Market-based Expectations." FRBSF Economic Letter No. 2015-12 (April 13).

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.

Capistran, C., and G. Lopez-Moctezuma. 2014. "Forecast Revisions of Mexican Inflation and GDP Growth." *International Journal of Forecasting* 30 (2): 177–91.

Christensen, I., F. Dion, and C. Reid. 2004. "Real Return Bonds, Inflation Expectations, and the Break-even Inflation Rate." Working Paper No. 2004-43, Bank of Canada.

Clark, T. E., and T. Nakata. 2008. "Has the Behavior of Inflation and Long-Term Inflation Expectations Changed?" *Economic Review* (Federal Reserve Bank of Kansas City) (First Quarter).

Côté, A., J. Jacob, J. Nelmes, and M. Whittingham. 1996. "Inflation Expectations and Real Return Bonds." *Bank of Canada Review* (Summer): 41–53.

Davies, A., K. Lahiri, and X. Sheng. 2011. "Analyzing Three-dimensional Panel Data of Forecasts." In *Oxford Handbook on Economic Forecasting*, ed. M. P. Clements and D. F. Hendry, 473–95. Oxford: Oxford University Press.

De Pooter, M., P. Robitaille, I. Walker, and M. Zdinak. 2014. "Are Long-term Inflation Expectations Well Anchored in Brazil, Chile, and Mexico?" *International Journal of Central Banking* 10 (2): 337–400.

Devereux, M., and J. Yetman. 2003. "Predetermined Prices and the Persistent Effects of Money on Output." *Journal of Money, Credit and Banking* 35 (5): 729–41.

Dovern, J., and U. Fritsche. 2008. "Estimating Fundamental Cross-section Dispersion from Fixed Event Forecasts." Discussion Paper No. 1, University of Hamburg.

Dovern, J., U. Fritsche, and J. Slacalek. 2012. "Disagreement Among Forecasters in G7 Countries." *Review of Economics and Statistics* 94 (4): 1081–96.

European Central Bank. 2012. "Assessing the Anchoring of Longer-term Inflation Expectations." *Monthly Bulletin* (July): 65–78.

Faust, J., and J. Wright. 2013. "Forecasting Inflation." In *Handbook of Economic Forecasting*, Vol. 2A, ed. G. Elliott and A. Timmermann, 2–56 (chapter 1). Amsterdam: Elsevier.

Fuhrer, J., and G. Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110 (1): 127–60.

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

Gefang, D., G. Koop, and S. M. Potter. 2012. "The Dynamics of UK and US Inflation Expectations." *Computational Statistics and Data Analysis* 56 (11): 3120–33.

Gerlach, S. 2007. "Interest Rate Setting by the ECB, 1999-2006: Words and Deeds." *International Journal of Central Banking* 3 (3): 1–45.

Gregory, A. W., and J. Yetman. 2004. "The Evolution of Consensus in Macroeconomic Forecasting." *International Journal of Forecasting* 20 (3): 461–73.

Gürkaynak, R., A. T. Levin, A. N. Marder, and E. T. Swanson. 2007. "Inflation Targeting and the Anchoring of Inflation Expectations in the Western Hemisphere." *FRBSF Economic Review* (Federal Reserve Bank of San Francisco), 25–47.

Gürkaynak, R. S., A. 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.

Gürkaynak, R. S., B. Sack, and J. H. Wright. 2008. "The TIPS Yield Curve and Inflation Compensation." Finance and Discussion Series Paper No. 2008-05, Board of Governors of the Federal Reserve System.

Hattori, M., and J. Yetman. 2017. "The Evolution of Inflation Expectations in Japan." BIS Working Paper No. 647.

Hördahl, P. 2009. "Disentangling the Drivers of Recent Shifts in Break-even Inflation Rates." *BIS Quarterly Review* (March): 10–11.

Isiklar, G., and K. Lahiri. 2007. "How Far Ahead Can We Forecast? Evidence from Cross-country Surveys." *International Journal of Forecasting* 23 (2): 167–87.

Jochmann, M., G. Koop, and S. M. Potter. 2010. "Modeling the Dynamics of Inflation Compensation." *Journal of Empirical Finance* 17 (1): 157–67.

Kamada, K., J. Nakajima, and S. Nishiguchi. 2015. "Are Household Inflation Expectations Anchored in Japan?" Working Paper No. 15-E-8, Bank of Japan.

Kortelainen, M., M. Paloviita, and M. Viren. 2011. "Observed Inflation Forecasts and the New Keynesian Macro Model." *Economics Letters* 112 (1): 88–90.

Kozicki, S., and P. A. Tinsley. 2001a. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics* 47 (3): 613–52.

———. 2001b. "Term Structure Views of Monetary Policy under Alternative Models of Agent Expectations." *Journal of Economic Dynamics and Control* 25 (1–2): 149–64.

———. 2012. "Effective Use of Survey Information in Estimating the Evolution of Expected Inflation." *Journal of Money, Credit and Banking* 44 (1): 145–69.

Lahiri, K., and X. Sheng. 2008. "Evolution of Forecast Disagreement in a Bayesian Learning Model." *Journal of Econometrics* 144 (2): 325–40.

———. 2010. "Learning and Heterogeneity in GDP and Inflation Forecasts." *International Journal of Forecasting* 26 (2): 265–92.

Lemke, W., and T. Strohsal. 2013. "What Can Break-even Inflation Rates Tell Us about the Anchoring of Inflation Expectations in the Euro Area?" Mimeo.

Levin, A. T., F. B. Natalucci, and J. M. Piger. 2004. "Explicit Inflation Objectives and Macroeconomic Outcomes." ECB Working Paper No. 383.

Liu, Z., E. Vangelista, I. Kaminska, and J. Relleen. 2015. "The Informational Content of Market-based Measures of Inflation Expectations Derived from Government Bonds and Inflation Swaps in the United Kingdom." Working Paper No. 551, Bank of England.

Mankiw, N. G., and R. Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

Mehrotra, A., and J. Yetman. 2014. "How Anchored Are Inflation Expectations in Asia? Evidence from Surveys of Professional Forecasters." *BIS Papers* 77 (March): 181–91.

Nason, J. M., and G. W. Smith. 2014. “Measuring the Slowly Evolving Trend in US Inflation with Professional Forecasts.” CAMA Working Paper No. 7/2014 (January).

Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. 1988. *Numerical Recipes in C.* First Edition. Cambridge University Press.

Shen, P. 2006. “Liquidity Risk Premia and Breakeven Inflation Rates.” *Economic Review* (Federal Reserve Bank of Kansas City) (Second Quarter): 29–54.

Siklos, P. L. 2013. “Sources of Disagreement in Inflation Forecasts: An International Empirical Investigation.” *Journal of International Economics* 90 (1): 218–31.

Sims, C. A. 2003. “Implications of Rational Inattention.” *Journal of Monetary Economics* 50 (3): 665–90.

Svensson, L. E. O. 1997. “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets.” *European Economic Review* 41 (6): 1111–46.

Yetman, J. 2017. “The Evolution of Inflation Expectations in Canada and the US.” *Canadian Journal of Economics* 50 (3): 711–37.