BANK FOR INTERNATIONAL SETTLEMENTS

BIS Working Papers

No 464

Decaying expectations: what inflation forecasts tell us about the anchoring of inflation expectations

by Aaron Mehrotra and James Yetman

Monetary and Economic Department

September 2014

JEL classification: E31, E58

Keywords: Inflation expectations, decay function, inflation targeting

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Decaying expectations: what inflation forecasts tell us about the anchoring of inflation expectations

Aaron Mehrotra and James Yetman[1]

Abstract

Well anchored inflation expectations are considered to be a reflection of credible monetary policy. In the past, anchoring has been assessed using either long-run inflation surveys or break-even inflation rates on financial assets with long maturities. But neither of these is ideal. 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 away from a 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. Based on our estimation results we argue that inflation expectations have become more tightly anchored over time in both inflation targeting economies and in those following other regimes. However, inflation targeting regimes have seen a greater change along three dimensions: the level of the anchor has fallen further; the tightness of anchoring has increased more; and the relationship between the anchor and actual inflation outcomes has weakened to a greater degree.

Keywords: Inflation expectations, decay function, inflation targeting

JEL classification: E31, E58

[1] Mehrotra: aaron.mehrotra@bis.org; Yetman: james.yetman@bis.org. Bank for International Settlements, Representative Office for Asia and the Pacific, 78th Floor, Two IFC, 8 Finance Street, Central, Hong Kong. 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 and Lillie Lam provided excellent research assistance. We are grateful to Dietrich Domanski, Gregor Smith, Philip Turner, Frank Packer, Marco Lombardi and seminar participants at the Bank for International Settlements for helpful comments and suggestions.

Contents

1. Introduction

Well-anchored inflation expectations – where anchoring refers to both the level and variability of anticipated future inflation – are important for the monetary transmission mechanism and are considered to be a reflection of credible monetary policy. If inflation expectations are not well anchored, forward-looking price and wage setting behaviour may become a source of macroeconomic instability. In many New Keynesian models, for example, well-anchored inflation expectations can significantly contribute to stabilising actual inflation.

Assessing how well anchored inflations 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 persistent shocks might have a measurable effect. Previous authors have done this based on either long-run inflation surveys or break-even inflation rates drawn from asset prices with long maturities.[2] But neither of these solutions is ideal. The former are generally published infrequently and for a limited number of economies, while the latter are influenced by many other factors than inflation expectations.[3]

In this paper, we propose a novel way to model the behaviour of inflation expectations, and 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 long-run anchor point in the direction of recent actual inflation as the forecast horizon shortens. Our approach is motivated by the assumption that inflation forecasts made sufficiently far in advance may 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 a constant 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 forecasting horizon reduces, any role that a long-run anchor may have played in affecting inflation forecasts is likely to decrease as forecasters learn more about the realisation of shocks to inflation for a given period. Our modelling strategy allows for all these possibilities.

Our decay function provides a parsimonious framework for fitting inflation forecasts that fully utilises the multiple-horizon dimension of the available data. Furthermore, it generates a small number of easy-to-interpret estimates that describe the anchoring of inflation expectations. These include both the level at which expectations are anchored in the long run and a measure of how tightly expectations are anchored to that estimated level. We interpret a precisely estimated anchor as indicating that long-run inflation expectations are tightly anchored. In contrast, weakly anchored expectations in our framework would imply that any estimated inflation anchor is poorly identified by, and therefore imprecisely

[2] See, for example, ECB (2012).

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

estimated in, the forecast data. This interpretation is clearly dependent on our model providing a good fit for the forecast data, which we also demonstrate.

The functional form that we use to model inflation expectations is based on the cumulative density function of the Weibull distribution. We show that this functional form cannot be rejected against a more general alternative functional form for most economies in our sample. Moreover, the estimates of the anchors of inflation expectations that we generate are broadly similar to break-even rates constructed from nominal and real bonds (where these are available), but are more stable over sub-samples. We further demonstrate the usefulness of the proposed methodology, contrasting the anchoring of inflation expectations in inflation targeting economies from those with other monetary policy regimes.

Our data are from professional forecasters, collected and published by Consensus Economics. Our sample comprises forecasts for 44 economies, of which 18 explicitly target inflation. This includes a large number of emerging market economies (in contrast with the advanced economy focus prevalent in previous 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% of global GDP in 2012 at market exchange rates.

Our estimates show that the tighter anchoring of inflation expectations in recent years is not limited to inflation targeting economies. However, inflation targeting is associated with greater changes in anchoring over time along three dimensions. First, the level of the anchor has declined by more over time for inflation targeters than for other economies. Second, the standard error of our estimate of the inflation anchor has decreased by more for inflation targeting economies than for others, suggesting that inflation targeters have seen a disproportionately large increase in the tightness with which long-run inflation expectations are anchored. And third, the estimated precision of the inflation anchor has become less affected by actual inflation volatility for the inflation targeters over time, but not for economies following other monetary regimes. This implies that the presence of a formal target may have helped to anchor inflation expectations beyond what can be explained by inflation outcomes, at least in inflation targeting economies.

The use of decay functions in modelling the behaviour of inflation expectations is not completely new. Gregory and Yetman (2004) use a polynomial decay function and Blue Chip survey data to model the behaviour 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.[4]

[4] See also Davies et al (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) conclude 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-based decay path that we consider here. However, we focus on capturing the behaviour of expectations over the different forecast horizons, rather than attempting to provide the most *accurate* forecasts of future inflation.

Our study is related to previous 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 (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 et al (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 et al (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 break-even inflation rates is investigated using financial market data in Gefang et al (2012), Jochmann et al (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 et al (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 20 years ago.[6]

Our paper also ties in with previous research examining the impact of inflation targeting on inflation expectations. Davis (2014) finds that the sensitivity of inflation expectations to oil price and inflation shocks has declined in inflation targeting economies, but not in others, in a sample of 36 economies. Johnson (2002, 2003) reports that the introduction of inflation targeting lowered the level of expected inflation. On the other hand, Cecchetti and Hakkio (2009) and Siklos (2013) do not find strong evidence that inflation targeting has affected the dispersion of inflation forecasts across forecasters. Our paper differs from these studies using survey data in that we focus on a long-run anchor *estimated* from inflation forecasts, rather than the forecasts themselves.

Finally, we note the methodological benefits of the proposed approach to analyse 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 12-month period) than fixed-event forecasts (as in the case of the Consensus forecasts we examine, made for calendar years at varying horizons).

5 See also Gürkaynak et al (2010).

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

Unfortunately, panels of fixed-horizon forecasts of comparable length are not available for many economies, limiting their applicability for large cross-country analyses.[7] 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 (eg Dovern et al., 2012; Gerlach, 2007; Kortelainen et al., 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.[8]

This paper is structured as follows. The next section presents the methodology, outlining the decay function and discussing its suitability for evaluating the behaviour of inflation expectations. Section 3 presents the results from applying our model to data from 44 economies, while Section 4 focuses on the differences between inflation targeters and those following other monetary regimes. Section 5 concludes the paper.

2. Methodology and data

2.1 Functional form

We propose a parsimonious framework for fitting inflation forecasts that fully utilises 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 their 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.[9] This suggests 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. If our framework fits the forecast data well, then we can estimate the latter econometrically as the standard error of the estimated anchor.

As the forecast horizon shortens, even well-anchored inflation expectations will eventually start to deviate from their long-run anchor towards the level of actual inflation. Forecasters gradually learn more about the realisation of shocks to inflation for a given period, for example. A slow adjustment could arise due to information about the economic conditions being disseminated only slowly through

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

[8] Limitations of this approach are acknowledged by some papers using this method. For example, Kortelainen et al (2011) state that the "moving average process affects the properties of the data."

[9] Long-run expectations could still change, for example, if a shock causes the degree of monetary policy credibility to change, or if the central bank announces a new level for the inflation target. We take the latter consideration into account in Section 4 by evaluating whether the adoption of inflation targeting brought about a change in the inflation rate at which inflation expectations are anchored in the long run.

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).

Forecasts for headline inflation at different horizons — Graph 1

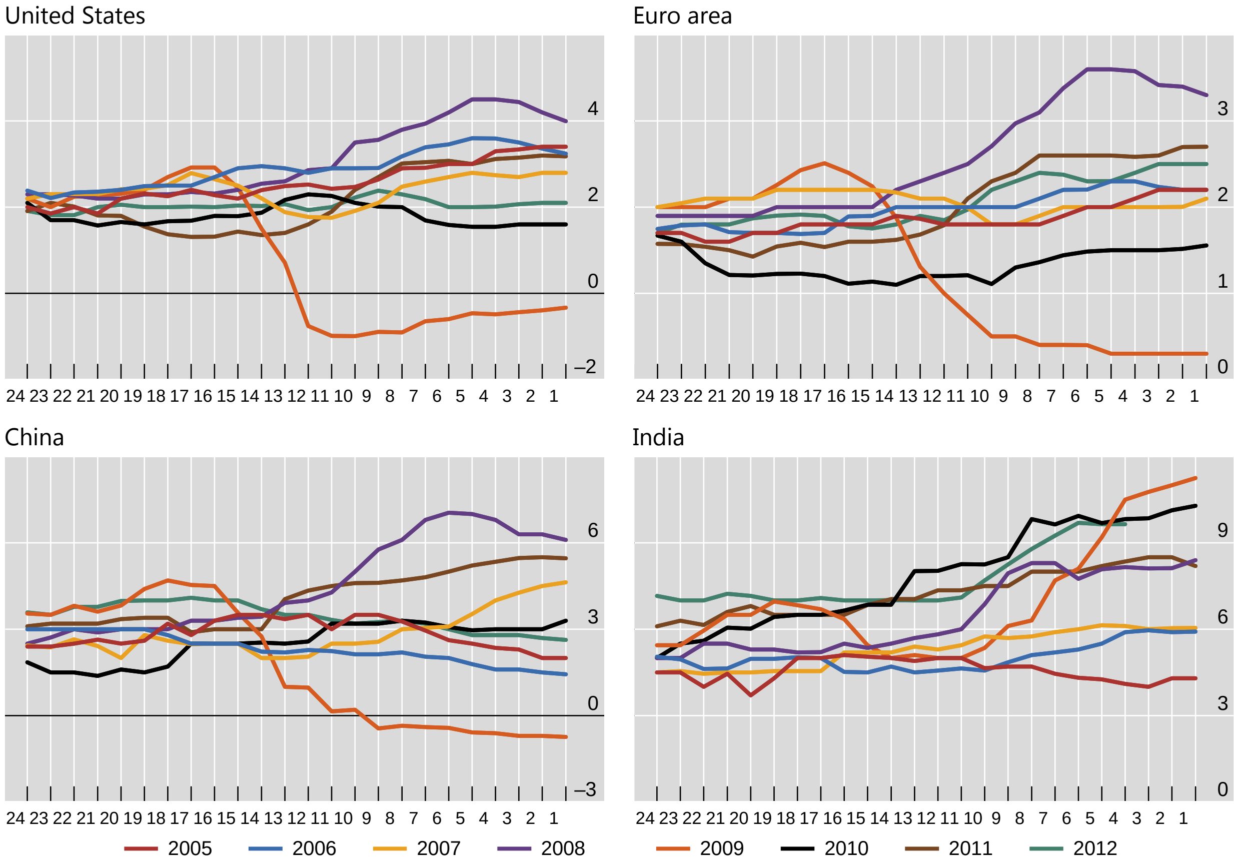

Note: The horizontal axis shows the forecast horizon, eg "24" indicates forecasts made 24 months before the completion of the calendar year being forecast. The vertical axis is the inflation rate in the CPI, measured in percent.

Source: Consensus Economics ©.

We apply our framework to forecast data from Consensus Economics. 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 24 monthly forecasts of the same outcome. Based on these fixed-event forecasts, Graph 1 illustrates the behaviour of median inflation forecasts across horizons from $h = 24$ (24 months before the completion of the year being forecast) to $h = 1$ (1 month before the completion of the year being forecast) for the period 2005–12 for the two largest advanced economies (United States and the euro area), as well as two large emerging economies (China and India), in our sample.[10] The figure confirms our prior regarding the behaviour 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, and eventually look like the typical distribution of inflation outcomes. The close resemblance between the 24-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 Graph 1 parallels the cross-country evidence provided by Isiklar and Lahiri (2007) regarding forecasts of GDP growth. They report that forecasts are very similar across the different economies at a 24 month horizon, and do not change very much during the first 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 24-month forecast horizons across the different years.

Given such observed behaviour of inflation forecasts, we 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:

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

In (1), h is measured in months before the end of the year that is being forecast. π^* is the level that long-run inflation expectations are anchored to and, conditional on the model fitting the data well, the standard error of the estimated π^* provides an estimate of the tightness with which long-run expectations are anchored. $\pi(t-h)$ is the level of inflation observed at the time when the forecast is made and $\varepsilon(t, t-h)$ is a residual term. To correct for the publication lag in inflation data, we use the 12-month growth rate in monthly CPI lagged by one month as the actual inflation rate. This also helps to address any potential endogeneity issues between expected inflation and inflation out-turns.

$\alpha(h)$ denotes a decay function. As already discussed, this has the property that, as the horizon shortens, there is greater weight on realised outcomes and less on the long-run anchor point. In particular, we assume 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:

$$\alpha(h) = 1 - \exp\left(-\left(\frac{h}{b}\right)^c\right). \tag{2}$$

This is based on the cumulative density function of the Weibull distribution.[11] Graph 2 illustrates the wide variety of possible decay paths that this functional form

[10] Similar graphs of the forecasts for the other 40 economies in the sample are shown as Graph A1 in the appendix.

[11] In Mehrotra and Yetman (2014), we briefly examine a more restricted version of the model (with $c = 0$) for a group of Asian economies. However, as we will see, this restriction can generally be rejected relative to the more general model examined here.

can generate for different values of b and c. With a small b parameter, for example, the function remains near 1 until the horizon gets close to zero. For high b, the opposite it 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 1 when c is higher, but only at forecast horizons above 4. Below that horizon, a higher c implies a more rapid decline in α.

Weibull decay functions

Graph 2

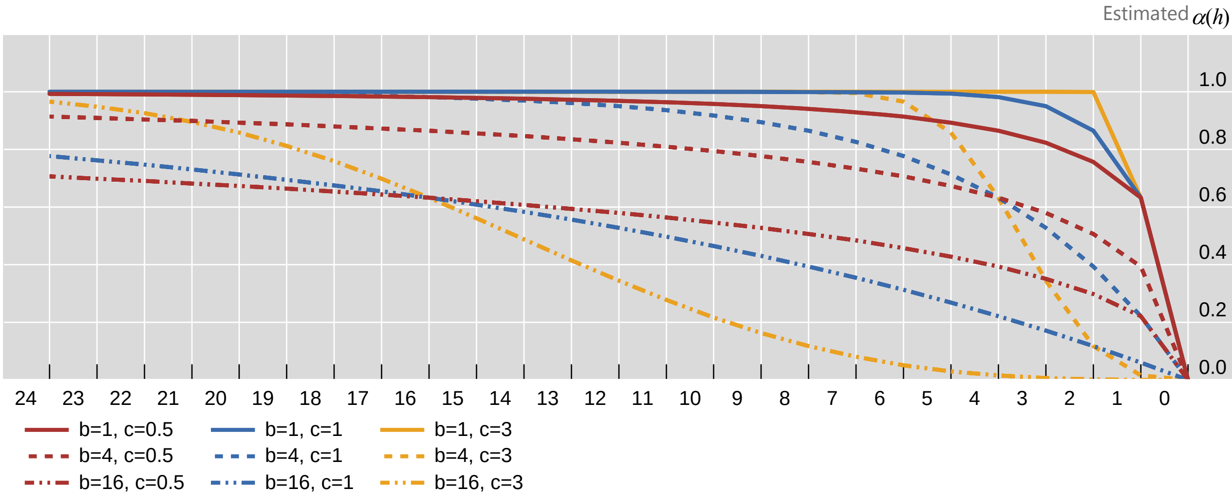

Note: Horizontal axis represents the forecast horizon h, which is the number of months before the end of the calendar year being forecast.

Source: Authors' calculations

We will also test this functional form against both more general and more restrictive alternatives. For $c = 1$, (2) reduces to a simple exponential decay function. The more general alternative we examine is $\alpha(h) = \alpha_h$, where there is a separate coefficient for each horizon and no restriction is imposed on the path of the α's across horizons.

The variance of the residual in (1) is modelled using a flexible functional form that allows it to change with the forecast horizon as:

$$V\big(\varepsilon(h,t)\big) = \exp(\delta_0 + \delta_1 h + \delta_2 h^2). \tag{3}$$

This formulation allows the variance to vary across the forecasting horizon with minimal restrictions.[12]

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:

[12] The use of the exponential function in (3) ensures that the fitted value of the variance is non-negative at all horizons for all possible values of the parameters. 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).

$$Corr\big(\varepsilon(t,t-h),\varepsilon(t,t-k)\big)=\phi_0-\phi_1|h-k|. \tag{4}$$

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

$$Cov\big(\varepsilon(t,t-h),\varepsilon(t,t-k)\big)=\left[\sqrt{V\big(\varepsilon(t,t-h)\big)V\big(\varepsilon(t,t-k)\big)}\right][\phi_0-\phi_1|h-k|]. \tag{5}$$

We note that in the existing empirical literature, instead of fully utilising the multiple-horizon nature of the data, it is common to use an approximation in order to convert fixed-event to fixed-horizon forecasts. As an example, in Dovern et al. (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 $k\in\{1,2,\dots,12\}$ and $k+12$ months. Fixed 12-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:

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

This approach implies that the 12-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 9/12 and 3/12, respectively. A similar approach is adopted in a number of other studies, for example Gerlach (2007), Kortelainen et al. (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 datapoints. 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 behaviour of forecasts at long horizons differs significantly from those at short horizons, as will be clear from our later empirical results.

2.2 The data

We apply our model to Consensus Economics forecasts of inflation, as these provide us with a large cross-country dataset 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 dataset implies that consistency in measurement across a large number of economies is essential.[13]

[13] Mankiw et al (2004) document that forecast disagreement differs substantially between consumers and expert forecasters. However, the evolution of disagreement does show similar time-series dynamics across the different groups of forecasters.

Inflation in the consumer price index (CPI) is used as our measure of headline inflation for all economies.[14] We include all economies for which forecaster-level forecasts of inflation and monthly series on inflation outcomes are available no later than 2008, yielding a sample of 44 economies.[15] Our sample is considerably larger than those considered in previous studies that have analysed long-run anchoring of inflation expectations. At market exchange rates and prices prevailing in 2012, the 44 economies included in our study accounted for 88% of world GDP.[16]

Data are available commencing with forecasts for the 1990 calendar year for some of the industrialised economies in the sample, while for emerging markets the sample periods generally start later.[17] Table 1 displays data availability and indicates when explicit inflation targeting was adopted where applicable.[18] Monthly median forecasts are constructed for the $h = 24$ to $h = 1$ months prior to the end of the year being forecast.

The model is estimated by maximum likelihood, economy-by-economy. We consider a wide range of possible starting values for each economy, and maximise the likelihood function using the hill-climbing method of Broyden, Fletcher, Goldfarb and Shanno (see Shanno, 1985, for details), until the estimates converge.

In some of our reported results, including all available forecasts of inflation from Consensus Economics for the relevant sample period means that we have an unbalanced panel. 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 part-way 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.

14 While it may have been preferable to use wholesale prices (WPI) for India due to its traditional significance as a headline inflation measure, Consensus Economics only collect forecasts for this indicator from 2002 onwards.

15 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.

16 Based on data from IMF WEO. 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.

17 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.

18 Slovakia and Spain also briefly followed inflation targeting but their monetary policy regimes changed again when they joined the euro area.

Data

Table 1

	Data available from	Year when inflation targeting adopted		Data available from	Year when inflation targeting adopted
Argentina	1993		Lithuania	1998	
Brazil	1990	1999	Malaysia	1990	
Bulgaria	1995		Mexico	1990	2001
Canada	1990	1991	Netherlands	1990	
Chile	1993	1999	Norway	1990	2001
China	1994		Peru	1993	2002
Chinese Taipei	1990		Philippines	1994	2002
Colombia	1993	1999	Poland	1990	1998
Croatia	1998		Romania	1995	2005
Czech Republic	1995	1997	Russia	1995	
Estonia	1998		Singapore	1990	
Euro area	1998		Slovakia	1995	
France	1990		Slovenia	1995	
Germany	1990		Spain	1990	
Hong Kong SAR	1990		Sweden	1990	1993
Hungary	1990	2001	Switzerland	1990	
India	1994		Thailand	1990	2000
Indonesia	1990	2005	Turkey	1995	2006
Italy	1990		Ukraine	1995	
Japan	1990		United Kingdom	1990	1992
Korea	1990	2001	United States	1990	
Latvia	1998		Venezuela	1993	

Sources: The years when inflation targeting was adopted are from Jahan (2012) which builds on Hammond (2011) and Roger (2010).

3. Results

We illustrate the estimated decay functions in Graph 3 for two large advanced economies (the euro area and the United States) and two large emerging market economies (China and India). These estimates are based on the sample period 2005–12. It is clear from the decay functions displayed in the left-hand panel of Graph 3 that inflation expectations start to move closer to actual inflation only gradually. Among the four economies, at the longest horizon of 24 months, the weights on the long-run anchor are similar across the four economies. Greater divergence is obtained after the forecast horizon has reduced to roughly 16 months. For example, the weight on the long-run anchor for India decreases rapidly and, at all horizons of 16 months and below, inflation expectations are driven more by actual inflation in this economy than for the other economies in the sample. On the other hand, the decay functions for the euro area and the United States are very similar in shape and level across the different forecast horizons.[19]

Decay function and estimated long-run anchor

Graph 3

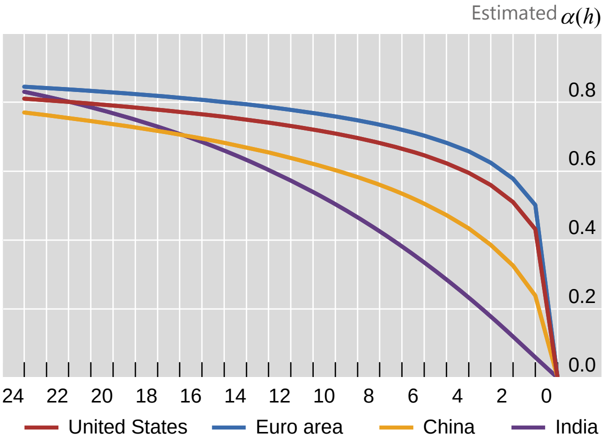

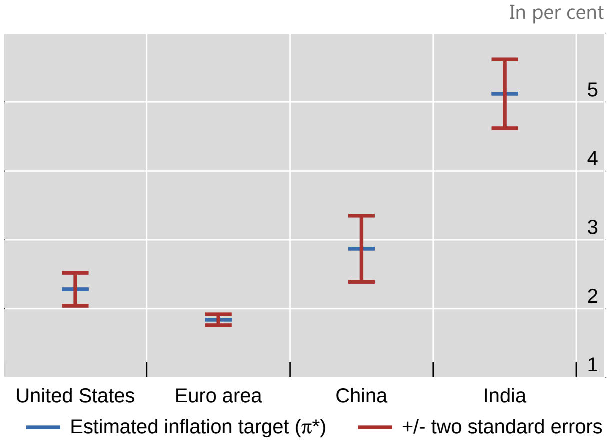

[1] Horizontal axis represents the forecast horizon *h*, which is the number of months before the end of the calendar year being forecast.

Source: Authors' calculations.

There are also differences between the economies in terms of the estimated level and standard error of the long-run anchor, π^*. For the economies shown in Graph 3, the estimated level of the anchor over this sample period is lowest for the euro area, at 1.84, which is consistent with the European Central Bank's definition of price stability as inflation "below, but close to, 2 percent". For the euro area, the standard error of the estimated π^*, capturing how tightly long-run expectations are anchored, is also low. For the United States, the estimated π^* is 2.28, while the one for China is 2.87. Finally, long-run inflation expectations are anchored at a level of 5.12 for India. China and India share a relatively large standard error (around 0.25).

The key parameter estimates for the 2005–12 period for all economies in our sample are given in Table 2. There is a wide variety of estimated parameters across economies, and one extreme outlier. Venezuela's estimated long-run inflation anchor is both the highest (32%) and the least precisely estimated (standard error of 6%). The implication is that expectations are poorly anchored for this economy, in contrast to most others. This observation is confirmed when we evaluate the decay functions for all economies in the sample, for the same time period. These are shown in Graph A2 in the appendix. For both Venezuela and Argentina the estimated decay path implies little decay suggesting that the model does not capture their inflation expectations dynamics well. Uniquely in our sample, in the

[19] We note that the declining weight on the long-run target is comparable to the estimated dynamics arising from the Bayesian learning 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.

case of Argentina, the estimate of c is zero across a wide range of possible starting values, all of which yield the same value of the likelihood function.[20]

Estimation results, 2005-12

Table 2

	b	s.e.(b)	c	s.e.(c)	π^*	s.e.(π^*)
Argentina	0.33	0.13	0.00	0.00	9.79	0.38
Brazil	0.17	0.61	0.23	0.17	4.67	0.09
Bulgaria	16.23	0.76	1.44	0.14	3.55	0.08
Canada	3.16	1.62	0.54	0.15	2.01	0.05
Chile	1.14	1.84	0.41	0.22	3.01	0.04
China	11.61	3.71	0.53	0.10	2.87	0.24
Chinese Taipei	5.99	1.51	0.57	0.10	1.68	0.09
Colombia	19.10	4.22	0.67	0.31	4.03	0.22
Croatia	16.22	1.81	0.68	0.12	2.61	0.03
Czech Republic	10.72	1.45	3.30	1.36	2.72	0.03
Estonia	13.95	1.70	0.76	0.14	2.88	0.09
Euro area	3.21	1.42	0.31	0.07	1.84	0.04
France	3.96	2.50	0.22	0.08	1.74	0.05
Germany	4.59	2.97	0.29	0.09	1.76	0.11
Hong Kong SAR	12.70	4.37	0.45	0.09	2.92	0.20
Hungary	10.03	1.77	1.31	0.28	3.64	0.12
India	13.98	1.81	1.06	0.25	5.12	0.25
Indonesia	7.70	1.49	0.71	0.10	6.03	0.05
Italy	2.67	1.36	0.36	0.08	1.87	0.04
Japan	1.01	2.94	0.14	0.14	0.15	0.15
Korea	0.43	1.02	0.20	0.12	2.90	0.08
Latvia	22.97	3.13	1.17	0.10	2.28	0.37
Lithuania	19.74	1.64	1.13	0.26	2.65	0.11
Malaysia	3.13	1.99	0.32	0.12	2.37	0.13
Mexico	12.96	1.51	2.67	1.13	3.69	0.07
Netherlands	0.73	1.06	0.25	0.13	1.73	0.12
Norway	3.74	1.86	0.43	0.12	2.06	0.10
Peru	1.52	2.62	0.23	0.19	2.42	0.13
Philippines	35.43	33.64	0.24	0.15	4.73	0.29
Poland	8.58	0.92	1.52	0.26	2.51	0.04
Romania	38.88	3.00	0.95	0.05	2.33	0.39
Russia	32.32	4.64	1.10	0.10	5.55	0.59
Singapore	13.39	2.30	0.69	0.12	1.79	0.11
Slovakia	5.55	2.33	0.40	0.12	2.56	0.12
Slovenia	47.57	10.28	0.51	0.12	2.20	0.07
Spain	7.60	4.08	0.24	0.07	2.55	0.11
Sweden	1.29	0.71	0.37	0.07	2.00	0.08
Switzerland	4.37	1.74	0.38	0.08	1.10	0.06
Thailand	1.15	1.13	0.26	0.10	2.63	0.14
Turkey	24.66	0.75	2.47	0.28	4.39	0.36
Ukraine	19.87	0.79	1.85	0.24	7.70	0.15
United Kingdom	2.39	1.53	0.33	0.11	2.16	0.10
United States	5.38	2.37	0.34	0.07	2.28	0.12
Venezuela	1.42	0.46	103.27	54.31	31.85	5.90

Note: s.e. indicates standard error.

Source: Authors' calculations.

[20] With $c = 0$, the "decay function" (2) no longer decays but instead reduces to the constant $(1 - \exp(-1))$.

Goodness-of-fit results

Table 3

	Pseudo R^2	p-values for different specifications		
		Exponential vs Weibull	Exponential vs flexible	Weibull vs flexible
Argentina	0.25	0.000	0.001	0.042
Brazil	0.12	0.003	0.024	0.136
Bulgaria	0.80	0.000	0.004	0.078
Canada	0.21	0.113	0.350	0.431
Chile	0.13	0.093	0.053	0.076
China	0.49	0.027	0.001	0.002
Chinese Taipei	0.29	0.003	0.030	0.162
Colombia	0.63	0.284	0.167	0.168
Croatia	0.52	0.032	0.112	0.219
Czech Republic	0.46	0.000	0.005	0.283
Estonia	0.49	0.186	0.144	0.162
Euro area	0.32	0.000	0.000	0.282
France	0.26	0.000	0.000	0.265
Germany	0.26	0.000	0.000	0.214
Hong Kong SAR	0.44	0.000	0.000	0.014
Hungary	0.52	0.201	1.000	1.000
India	0.82	0.793	0.082	0.063
Indonesia	0.49	0.026	0.683	0.889
Italy	0.25	0.000	0.000	0.005
Japan	0.18	0.004	0.000	0.000
Korea	0.13	0.005	0.000	0.000
Latvia	0.51	0.108	0.994	0.999
Lithuania	0.54	0.634	0.021	0.016
Malaysia	0.10	0.000	0.177	0.853
Mexico	0.01	0.300	0.000	0.000
Netherlands	0.10	0.000	0.001	0.069
Norway	0.21	0.009	0.004	0.017
Peru	0.10	0.013	0.281	0.570
Philippines	0.25	0.008	0.112	0.330
Poland	0.38	0.018	0.960	0.999
Romania	0.71	0.426	0.538	0.517
Russia	0.61	0.331	0.231	0.225
Singapore	0.44	0.061	0.238	0.351
Slovakia	0.30	0.000	0.000	0.121
Slovenia	0.61	0.003	0.116	0.450
Spain	0.39	0.000	0.000	0.092
Sweden	0.14	0.000	0.000	0.001
Switzerland	0.31	0.000	0.000	0.005
Thailand	0.16	0.000	0.001	0.024
Turkey	0.23	0.000	0.000	0.389
Ukraine	0.85	0.000	0.000	0.120
United Kingdom	0.22	0.000	0.000	0.003
United States	0.27	0.000	0.000	0.000
Venezuela	0.82	0.000	0.000	0.011

Source: Authors' calculations.

Our interpretation of the estimated standard error of π^* as a measure of the degree to which inflation expectations are anchored depends on our framework providing a good description of the forecast data. We next demonstrate this in two ways, in Table 3. First, we compute a pseudo R^2 following McKelvey and Zavoina (1975) as:

$$Pseudo\ R^2 = \frac{\sum_{t,h}(f(t,t-h) - \alpha(h)\pi^* - [1-\alpha(h)]\pi(t-h))^2}{\sum_{t,h}(f(t,t-h) - \bar{f})^2},$$

where $\bar{f} = (1/TH)\sum_{t,h} f(t,t-h)$. We also estimate a more restrictive version of the model, where $c = 1$ so that (2) reduces to a simple exponential decay function (labelled "exponential"), and a more general alternative where no restrictions are placed on the value of $\alpha(h)$ from one horizon to the next ("flexible"). We then compute likelihood ratio tests between each of these functional forms.

The average pseudo R^2 for our sample of economies is 0.37. At the same time, there is considerable variation, from a high of 0.85 (Ukraine) to a low of 0.01 (Mexico). With the exception of Mexico, this measure of goodness-of-fit exceeds 0.10 for all economies in the sample, indicating that our simple model explains a considerable share of the variability in our panels of inflation forecasts.

When we compare the Weibull-based decay function against the exponential or the more general form, the evidence suggests that the Weibull distribution fits the data well. Based on likelihood ratio tests, the exponential model can be rejected against the Weibull for 73% of economies at the 5% level (61% of economies at the 1% level). The exponential model can be rejected against the more flexible model that allows for an independent α_h at each horizon for 61% (55%) of economies at the same significance levels. In contrast, at equivalent levels, the Weibull is only rejected against the more general model in 34% (21%) of cases.

One concern is that our estimates of the inflation anchors may be compromised since they are obtained by considering only data at forecast horizons of up to 24 months. If forecasters significantly change their views over longer horizons than that, our estimated inflation anchor may be a poor proxy for long-run inflation expectations. We test for this directly by comparing our estimates for π* against long-term Consensus forecasts made over the same period. The latter are forecasts for average inflation for 6-10 years ahead but, in contrast to the forecast data that we used above, are made only twice a year (April and October for most economies; March and September for economies in emerging Europe). We include those 37 economies for which long-term forecasts are available no later than October 2004. We then compute the average for all available long-term forecasts made during 2004-12, economy by economy. We compare these against the estimated π* in Graph 4, together with a 45° line to aid comparison.[21]

It is clear from Graph 4 that the average long-run forecasts are highly correlated with the estimated π*. If we were to include a 95% confidence band of the average long-run forecast in the graph, this would enclose our estimated π* in 29 of the 37 economies. The exceptions are mainly those economies at either extreme. For those economies with the lowest estimated inflation anchors, long-run forecasts indicate that inflation is expected to be higher 6-10 years ahead than data for the coming 24 months imply; for those with the highest estimated anchors, the reverse is true. Aside from these extremes, our methodology, using forecast horizons up to 24 months, produces estimates for the inflation expectations anchor that are generally comparable to those obtained from averaging long-run inflation forecasts.

[21] Note that forecasts of inflation outcomes for 2005–12, made at the shorter horizons and used to estimate the π*'s in the graph, were made between January 2004 and December 2012. Hence the results reported in Graph 4 are based on forecasts made over the same periods.

π* vs. long term forecasts Graph 4

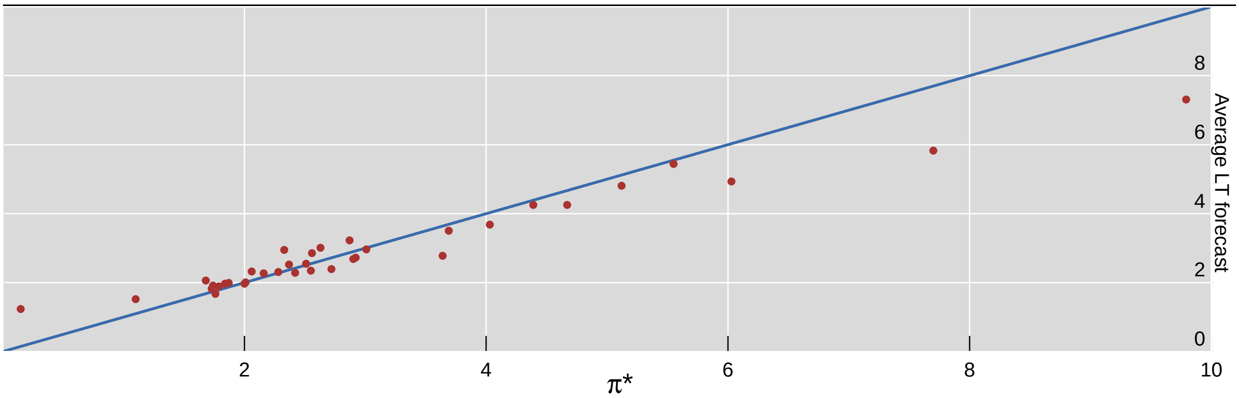

Note: The graph shows the estimates of π* and long-term Consensus forecasts. The long-term forecasts are for average inflation 6–10 years ahead and are made twice a year. The graph includes the results for the 37 economies for which long-term forecasts are available no later than October 2004 with the exception of Venezuela, an extreme outlier.

Source: Consensus Economics ©; Authors' calculations.

Given our model, long-term inflation expectations should be very stable, at least in economies with relative macroeconomic stability and/or well-established policy frameworks, as expectations converge to their long-run anchor as the horizon increases. In contrast, measures of inflation expectations extracted from inflation-indexed financial market instruments tend to be more volatile, and are arguably too volatile to represent either rational long-run inflation expectations or the target of a central bank, as Faust and Wright (2013) contend. Hördahl (2009) argues that there are four components that constitute the break-even rates between real and nominal bonds: expected inflation, inflation risk premia, liquidity premia and technical market factors. Especially at times of crisis, liquidity risk premia could play a major role, complicating the extraction of inflation expectations from data on nominal and real bonds. Even the United States' fixed income market, which is the deepest, was affected by such factors at the time of the international financial crisis.

We compare our estimates of the inflation anchor with break-even inflation rates for the four economies where both real and nominal bonds have existed for the longest and where the markets for these instruments are considered to be the most liquid: Canada, Sweden, the United Kingdom and the United States. The break-even inflation rate is obtained as the difference between nominal and real bond rates, using daily data. We then compare the estimated π* for several five-year subsamples against the average break-even inflation rates over the same periods. Graph 5 shows that our measure is more stable over time than the alternative based on break-even inflation rates. This is especially so in the case of Canada, where real bonds are purchased primarily by pension funds and life insurance companies, in part due to favourable tax treatment (Côté et al, 1996). This pushes down yields on real bonds so that break-even inflation rates in Canada tend to be higher than our estimated π*.[22]

[22] For Japan, data are only available for a five-year period 2008-12 (and hence are not shown in Graph 5), obtained as the difference between 10-year nominal and index-linked yields. However, for some years, the break-even inflation rates for Japan are negative, which is arguably not credible as a measure of long-term inflation expectations.

Break-even inflation rates and estimated long-run anchor

In per cent

Graph 5

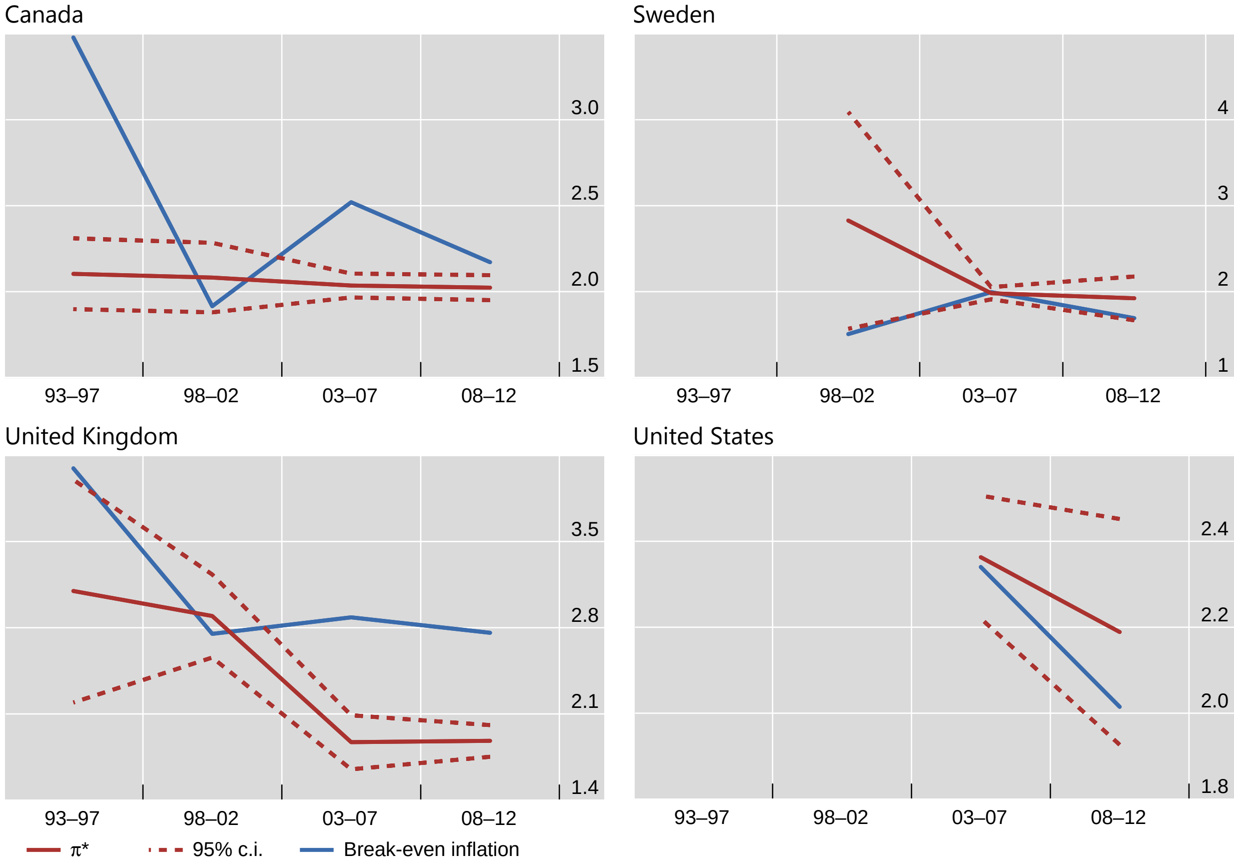

Note: For Canada, break-even inflation is obtained as the long-term nominal benchmark government bond yield less inflation-linked government bond yield; for Sweden, nominal bond yield (breakeven comparator from Bloomberg) less inflation-linked bond yield (all maturities); for the United Kingdom, generic break-even rate at the 10-year maturity; for the United States, yield on 10-year US Treasuries less 10-year index-linked government bond yield.

Source: Bloomberg; national data; authors' calculations.

4. Application to the effects of inflation targeting

To further illustrate the usefulness of our empirical methodology, we use it to investigate the effects of inflation targeting on the anchoring of inflation expectations. We first compare our estimates of the anchor in inflation targeting economies with their official targets. To do this, we estimate our model using a sample starting in the year an economy adopted a policy of inflation targeting and compare our estimated inflation anchor against the average announced target since adoption.

Table 4 suggests that our framework yields estimated anchors that are similar to the average of the midpoint of announced inflation targets of these economies. In particular, the estimated anchor is in most cases within one percentage point from the announced target, consistent with the typical range of announced inflation targets. Larger discrepancies than this are observed only in the cases of Colombia, the Czech Republic and Romania.[23] These results can be interpreted as providing evidence in support of the hypothesis that the inflation targets of these economies have been generally credible.

Estimated long-run anchor and average official inflation targets

Table 4

	Estimated π* since adoption	Average target since adoption		Estimated π* since adoption	Average target since adoption
Brazil	5.48 (0.14)	4.75	Norway	2.02 (0.07)	2.50
Canada	2.20 (0.11)	2.08	Peru	3.11 (0.15)	2.21
Chile	3.61 (0.05)	3.00	Philippines	4.97 (0.24)	4.20
Colombia	10.66 (0.16)	4.96	Poland	2.94 (0.18)	3.42
Czech Republic	4.49 (0.18)	3.17	Romania	2.33 (0.39)	3.98
Hungary	3.40 (0.22)	3.62	Sweden	2.35 (0.11)	2.00
Indonesia	6.03 (0.05)	5.18	Thailand	2.67 (0.11)	1.75
Korea	2.95 (0.06)	3.00	Turkey	6.15 (0.19)	5.25
Mexico	3.68 (0.23)	3.38	United Kingdom	2.94 (0.11)	2.27

Note: Standard errors in parentheses.

Sources: National sources; authors' calculations.

Inflation targeting also coincided with a fall in the level at which inflation expectations are anchored in 12 out of the 14 inflation targeting economies in our sample for which sufficient data are available (Table 5). In this exercise, we estimate the model for the five years immediately preceding and following the adoption of inflation targeting, respectively. This change in the level of the anchor between periods is in most cases statistically significant. Moreover, the standard error of the anchor decreased in 11 out of the 14 economies following the adoption of inflation targeting, consistent with inflation expectations becoming more tightly anchored.

It is also of interest to compare changes in the degree of anchoring over time in inflation targeting economies with those following other policy regimes. Table 6 shows median results for the estimation samples 1999–2005 and 2006–12, respectively. By 2006, all inflation targeting economies in our sample had adopted these frameworks, providing a relatively clean experiment.

Regarding the level of the estimated anchor, inflation expectations actually appear to be anchored at a lower level in the non-inflation targeters in both sub-samples. However, the results reported in Table 6 also indicate that the level of the inflation anchor fell further in inflation targeting economies between the two sub-samples than for other economies. In addition, the median standard error of the estimated inflation anchor was higher for inflation targeters than others in the first sub-sample, but lower in the second. Similar results are obtained if we focus on mean estimates rather than medians. Both of these results suggest that inflation

[23] In the case of Colombia, the high inflation observed at an early part of the sample accounts for the discrepancy. If we repeat the estimation with a start date of 2005 (instead of 2000), our estimate for π* is 4.03 (instead of 10.66).

targeting economies have seen greater changes in the behaviour of inflation expectations than other economies.[24]

Estimation results before and after inflation targeting

Table 5

	Prior to IT		Since IT adoption			Prior to IT		Since IT adoption	
	π*	s.e.	π*	s.e.		π*	s.e.	π*	s.e.
Brazil	37.09	23.22	5.61	0.12	Norway	3.11	0.11	2.16	0.07
Chile	5.37	0.24	3.63	0.03	Peru	5.46	0.23	2.41	0.06
Colombia	16.31	0.65	9.76	0.15	Philippines	7.12	0.34	6.61	1.21
Hungary	2.50	1.23	3.57	0.29	Poland	9.72	0.13	2.38	0.40
Indonesia	7.63	0.42	6.11	0.08	Romania	0.92	1.17	3.65	0.22
Korea	4.46	0.04	2.96	0.02	Thailand	5.05	0.12	2.50	0.15
Mexico	11.04	0.88	2.68	0.51	Turkey	14.48	3.43	5.42	0.29

Note: Prior to IT refers to the five years immediately preceding the adoption of an inflation targeting framework; Since IT adoption refers to the five years immediately after the adoption of inflation targeting. The year when inflation targeting was adopted is excluded from both samples. Economies with less data than five years before IT adoption are excluded. s.e. indicates standard error.

Sources: National sources; authors' calculations.

To give a sense of the full distribution of forecast anchors at different time periods, Graph 6 plots the estimates of the inflation anchors and their standard errors for these two sub-samples for inflation targeters and non-inflation targeters separately. The horizontal axis shows the estimates of the inflation anchor, and the vertical axis its standard error. For inflation targeters (top row), there has been a considerable tightening of the distribution of inflation anchors and their standard errors between the two sub-samples, as both the level and standard error of the estimated inflation anchor have declined over time. In contrast, the decline along these same dimensions is much more limited in the case of economies with other monetary policy regimes (bottom row).

Estimation results, 1999–2005 and 2006–12

Table 6

	Inflation targeters		Non-inflation targeters	
	Coefficient	Standard error	Coefficient	Standard error
π* (1999-2005)	3.60	0.23	2.51	0.18
π* (2006-2012)	2.95	0.11	2.35	0.13

Note: Medians across economies. Excluding Venezuela.

[24] Our framework can also be used to analyse differences *within* the group of inflation targeters. As an example, when applied to the group of "old" inflation targeters (that adopted the regime prior to 2000) and more recent inflation targeters (that adopted the regime post-2000), the results suggest statistically significant differences between the two groups. From a simple average across economies, for a sample of 2005–12, the estimated π* is lower for the "old" targeters (2.73 vs 3.60) and more tightly estimated (standard error of 0.06 vs 0.18). More detailed results are available upon request.

Another way previous research has evaluated the anchoring of inflation expectations is by comparing the possible disconnect between actual inflation and the long-run anchor, both in economies that are targeting inflation and in economies with other monetary regimes (eg Levin et al, 2004). We use our estimates to evaluate how the anchor is affected by volatility in actual inflation in the two groups of economies, and how this has changed over time. If long-term inflation expectations are tightly anchored, actual inflation volatility should have little impact on the standard error of the estimated anchor. In Graph 7, we compare the correlation between inflation volatility and the tightness of the estimate of π* in the inflation targeting economies (top panels) and the non-inflation targeters (bottom panels). Inflation volatility is measured by the standard deviation of monthly y-o-y CPI inflation rates. The two sample periods we consider are 1999–2005 and 2006–12.

Relationship between estimated level and standard error of the inflation anchor — Graph 6

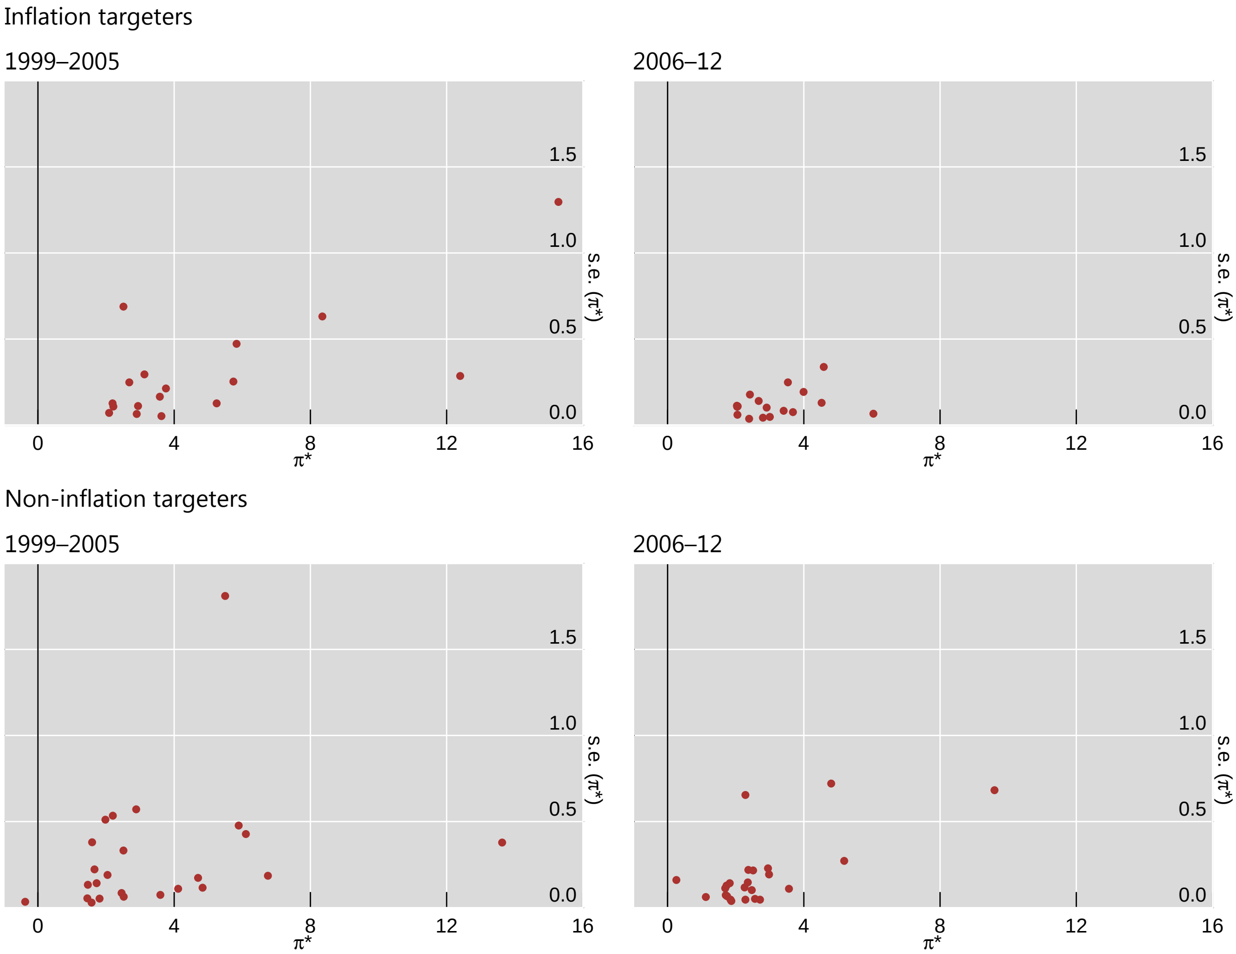

Note: The horizontal axis is the estimated inflation anchor, and the vertical axis the standard error. Excluding Turkey, Ukraine and Venezuela, each of which had estimated inflation anchors exceeding 60% in the first sub-sample.

Source: Authors' calculations.

These results suggest that our measure of how tightly long-run expectations are anchored has become less correlated with actual inflation volatility over time in the inflation targeting economies (top panels) but, if anything, have become more strongly associated with inflation volatility in the group of other economies (bottom panels). The correlation between actual inflation volatility and the standard error of the estimated anchor is very weak in the inflation targeting economies in 2006–12 compared to the earlier sample. The results are robust to excluding, in addition to Venezuela (an extreme outlier), all economies beyond +/-2 standard deviations from the sample mean (not shown). The results are similar to those reported by Levin et al. (2004) who find that there is a stronger correlation between private-sector inflation forecasts and actual inflation in economies without an inflation target than in economies that have explicit inflation objectives.[25]

Relationship between inflation volatility and the standard error of the estimated π* Graph 7

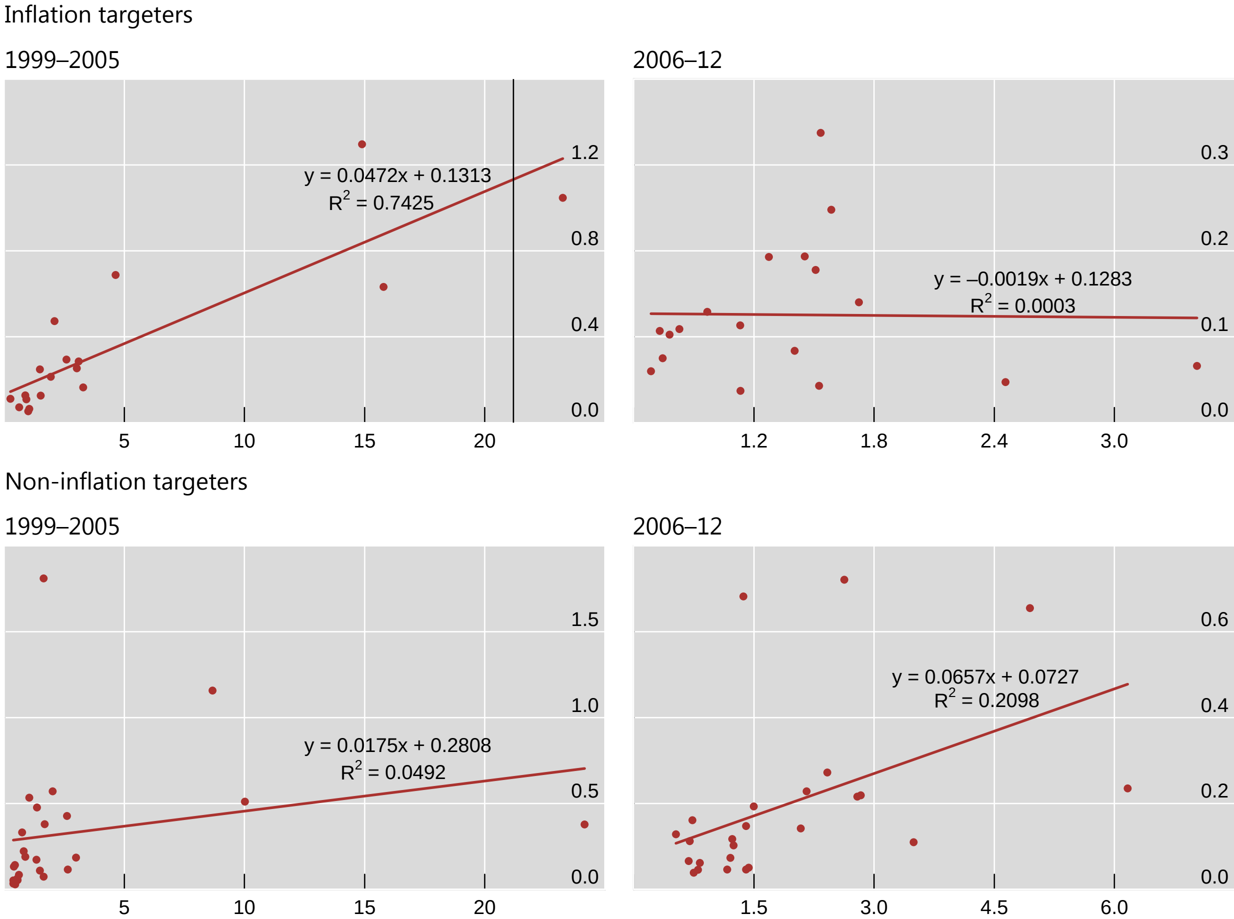

Note: The horizontal axis is inflation volatility, measured as the standard deviation of the monthly year-on-year inflation rate during 1999–2005 and 2006–12. The vertical axis the standard error of the estimated inflation anchor. Venezuela is excluded.

Source: Authors' calculations.

[25] Levin et al (2004) focus on a smaller country sample comprised of five inflation targeters and seven non-inflation targeters, all of which are advanced economies. Moreover, they compare actual inflation forecasts and a moving average of past inflation, while we focus on an *estimated* long-run anchor and the *volatility* of actual inflation.

5. Conclusion

In this paper, we have proposed a novel methodology to model the behaviour of inflation expectations. Forecast data are modelled using a decay function such that inflation forecasts monotonically diverge from a long-run anchor towards actual inflation as the forecast horizon shortens. Our methodology provides a parsimonious framework for fitting inflation forecasts that fully utilises the multiple-horizon dimension of the data. Further, we find empirical support for the model in a large panel of forecast data for 44 economies.

We demonstrate the usefulness of the proposed methodology by comparing the anchoring of inflation expectations in inflation targeting with that in non-inflation targeting economies. Our estimates show that inflation expectations have become more tightly anchored over time in both inflation targeting economies and in those following other regimes. However, inflation targeting regimes have seen a larger change along three dimensions: the level of the anchor has fallen further; the tightness of anchoring has increased more; and the relationship between the anchor and actual inflation outcomes has weakened to a greater degree. The last finding suggests that inflation targets serve to anchor expectations beyond what can be explained by inflation outcomes. We are also able to reconcile our results related to the level of the anchor with inflation forecasts at much longer forecast horizons of 6-10 years and estimates of long-run expected inflation extracted from high frequency financial data.

There are at least two lines of possible future work that could build on this approach. First, we have focused on the long-run anchor of inflation expectations, but our framework also provides evidence of the anchoring of expectations at shorter horizons that we have not explored here, characterised by our estimates of b and c. Second, our use of median forecasts for each economy in this paper ignores forecaster-level data. A natural extension for future research would be to model the behaviour of inflation expectations at a forecaster level using our functional form. This could be used to identify differences in expectations behaviour both across individual forecasters and for the same individual forecaster across time, and provide further insights into the anchoring of inflation expectations.

References

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, 104–129.

Capistran, C and G Lopez-Moctezuma (2014): "Forecast revisions of Mexican inflation and GDP growth", *International Journal of Forecasting* 30(2), 177–191.

Cecchetti, S G and C Hakkio (2009): "Inflation targeting and private sector forecasts", *NBER Working Papers* no 15424.

Clark, T E and T Nakata (2008): "Has the behaviour of inflation and long-term inflation expectations changed?", *Economic Review*, Federal Reserve Bank of Kansas City, Quarter One.

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, Lahiri K and X Sheng (2011): "Analyzing three-dimensional panel data of forecasts", In M P Clements and D F Hendry (eds), *Oxford Handbook on Economic Forecasting*, 473–495. Oxford: Oxford University Press.

Davis, J (2014): "Inflation targeting and the anchoring of inflation expectations: cross-country evidence from Consensus Forecasts", *Federal Reserve Bank of Dallas Globalization and Monetary Policy Institute Working Paper* no 174.

Devereux, M and J Yetman (2003): "Predetermined prices and the persistent effects of money on output", *Journal of Money, Credit and Banking* 35(5), 729–741

Dovern, J and U Fritsche (2008): "Estimating fundamental cross-section dispersion from fixed event forecasts", *Discussion Paper* no 1, University of Hamburg.

Dovern, J, Fritsche U and J Slacalek (2012): "Disagreement among forecasters in G7 countries", *The Review of Economics and Statistics* 94(4), 1081–1096.

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", *Handbook of Economic Forecasting 2A*, 2-56, G Elliott and A Timmermann (eds), Elsevier, Amsterdam

Fuhrer, J and G Moore (1995): "Inflation Persistence", *Quarterly Journal of Economics* 110(1), 127–160.

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 & Data Analysis* 56(11), 3120–3133.

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), 461–473.

Gürkaynak, R, Levin A T, Marder A N and E T Swanson (2007): "Inflation targeting and the anchoring of inflation expectations in the western hemisphere", *Federal Reserve Bank of San Francisco Economic Review*, 25–47.

Gürkaynak, R S, Levin A 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): 1208–1242.

Hammond, G (2011): "State of the Art of Inflation Targeting", Centre for Central Banking Studies Handbook no 29, London: Bank of England.

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), 167–187.

Jahan, S (2012): "Inflation Targeting: Holding the Line", *Finance & Development*, March, International Monetary Fund.

Jochmann, M, G Koop and S M Potter (2010): "Modeling the dynamics of inflation compensation", *Journal of Empirical Finance* 17(1), 157–167.

Johnson, D R (2002): "The Effect of Inflation Targeting on the Behavior of Expected Inflation: Evidence from an Eleven-Country Panel", *Journal of Monetary Economics* 49(8), 1493–519.

Johnson, D R (2003): "The Effect of Inflation Targets on the Level of Expected Inflation in Five Countries", *Review of Economics and Statistics* 55(4), 1076–81.

Kortelainen, M, Paloviita, M and M Viren (2011): "Observed inflation forecasts and the new Keynesian macro model", *Economics Letters* 112(1), 88–90.

Lahiri, K and X Sheng (2008): "Evolution of forecast disagreement in a Bayesian learning model", *Journal of Econometrics* 144(2), 325–340.

Lahiri, K and X Sheng (2010): "Learning and heterogeneity in GDP and inflation forecasts", *International Journal of Forecasting* 26(2), 265–292.

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 Papers* no 383.

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.

Mankiw, N G, R Reis and J Wolfers (2004): "Disagreement about inflation expectations", In M Gertler and K Rogoff (eds): *NBER Macroeconomics Annual* 2003, 209–48.

McKelvey, R D and W Zavoina (1975): "A statistical model for the analysis of ordinal level dependent variables", *Journal of Mathematical Sociology* 4(1), 103–120.

Mehrotra, A and J Yetman (2014): "How anchored are inflation expectations in Asia? Evidence from surveys of professional forecasters", *BIS Papers* no 77, 181-191

Nason, J M and G W Smith (2014): "Measuring the slowly evolving trend in US inflation with professional forecasts", *CAMA Working Paper* 7/2014, January.

Roger, S (2010): "Inflation Targeting Turns 20", *Finance & Development*, March, International Monetary Fund.

Shanno, D F (1985): "On Broyden-Fletcher-Goldfarb-Shanno Method", *Journal of Optimization Theory and Applications* 46(1), 87–94.

Siklos, P L (2013): "Sources of Disagreement in Inflation Forecasts: An International Empirical Investigation", *Journal of International Economics* 90(1), 218–231.

Sims, C A (2003): "Implications of rational inattention", *Journal of Monetary Economics* 50(3), 665–690.

Forecasts of headline inflation at different horizons

Graph A1

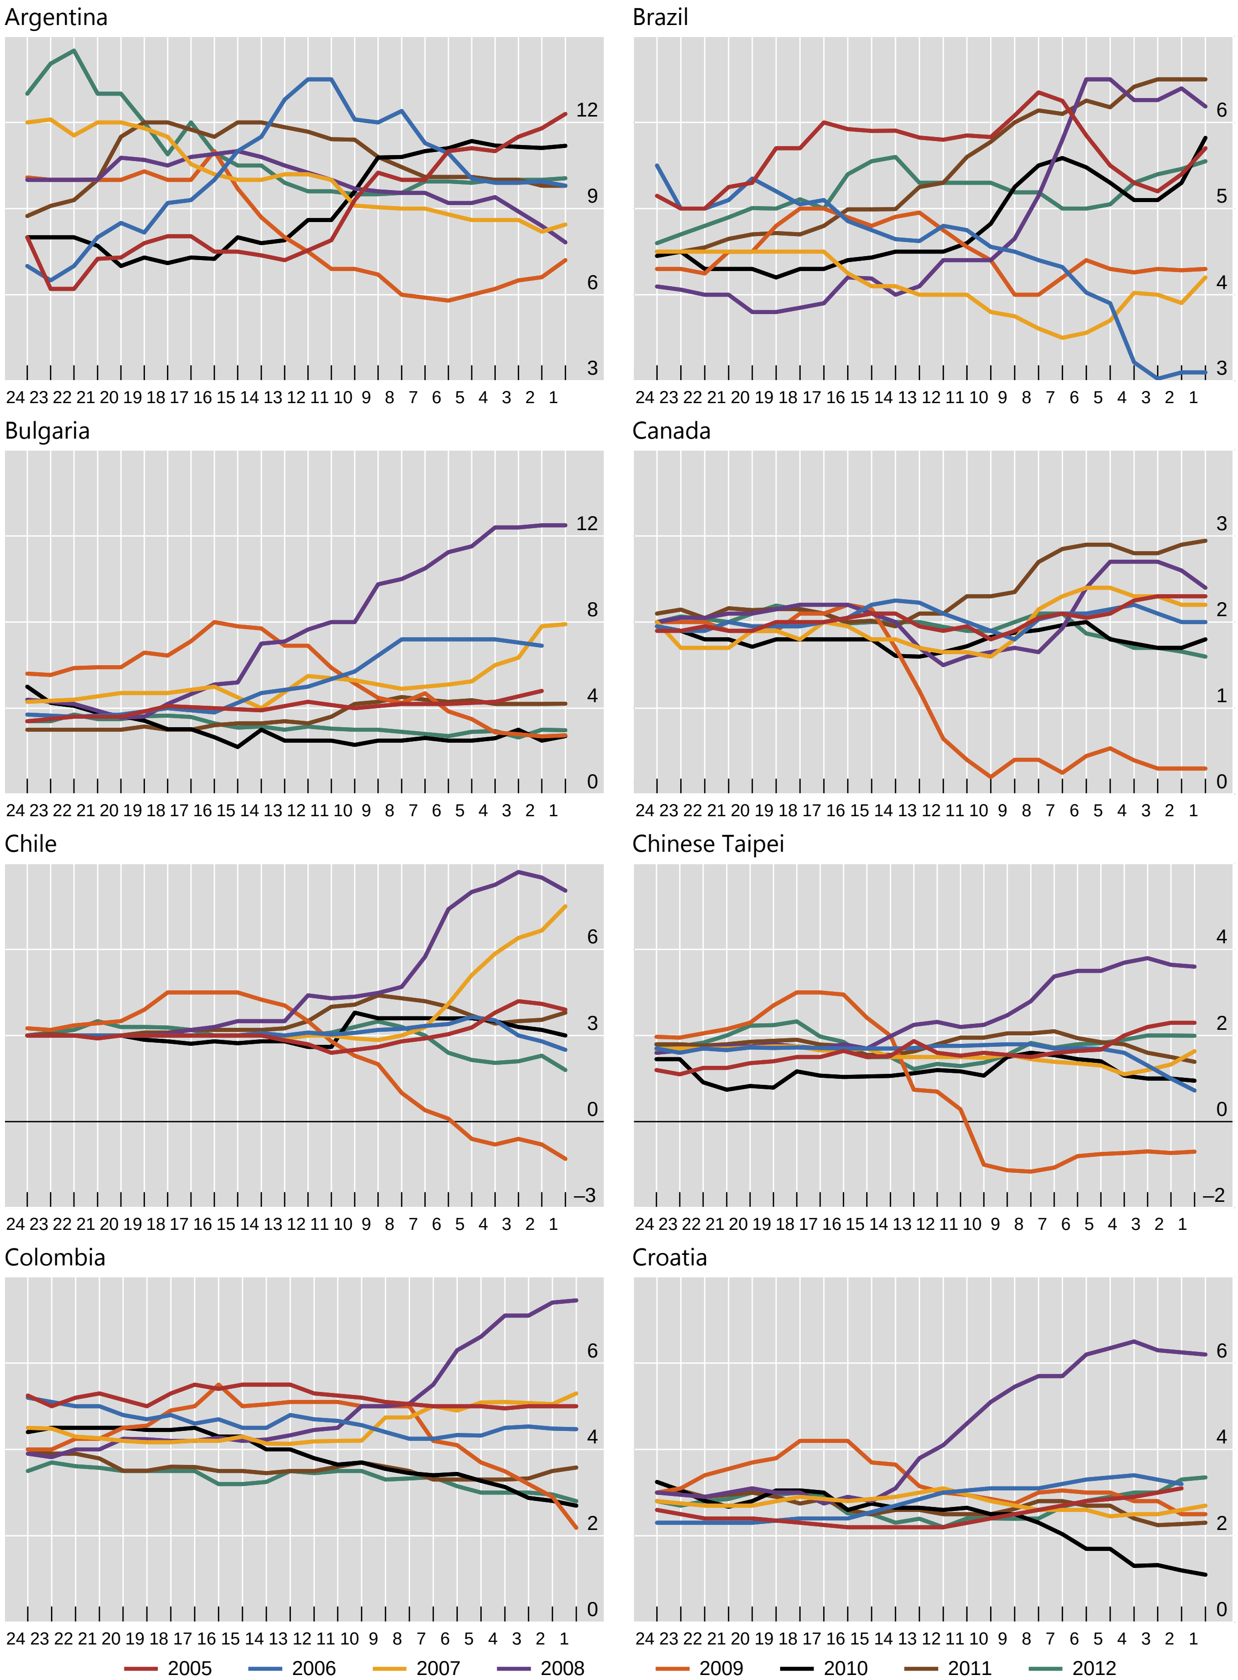

Horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.

Forecasts of headline inflation at different horizons (continued) Graph A1

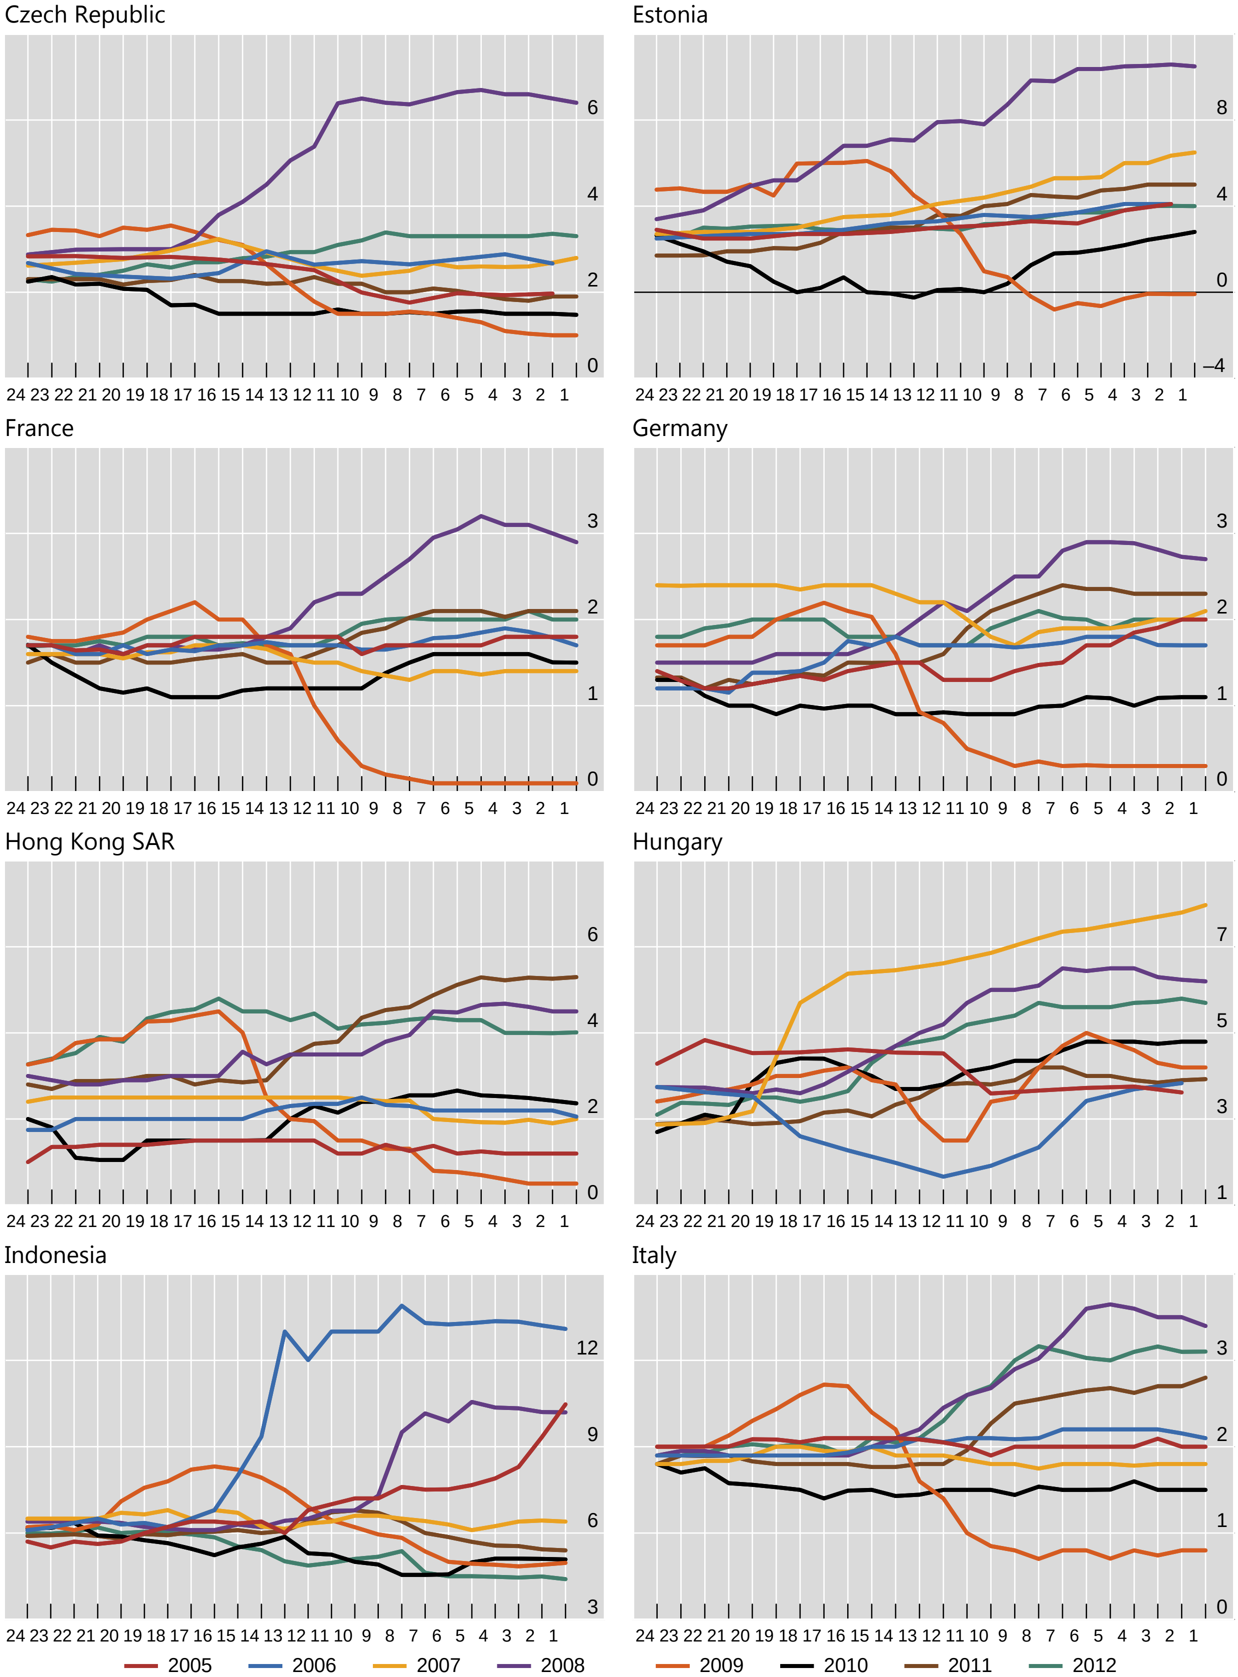

Horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.

Forecasts of headline inflation at different horizons (continued) Graph A1

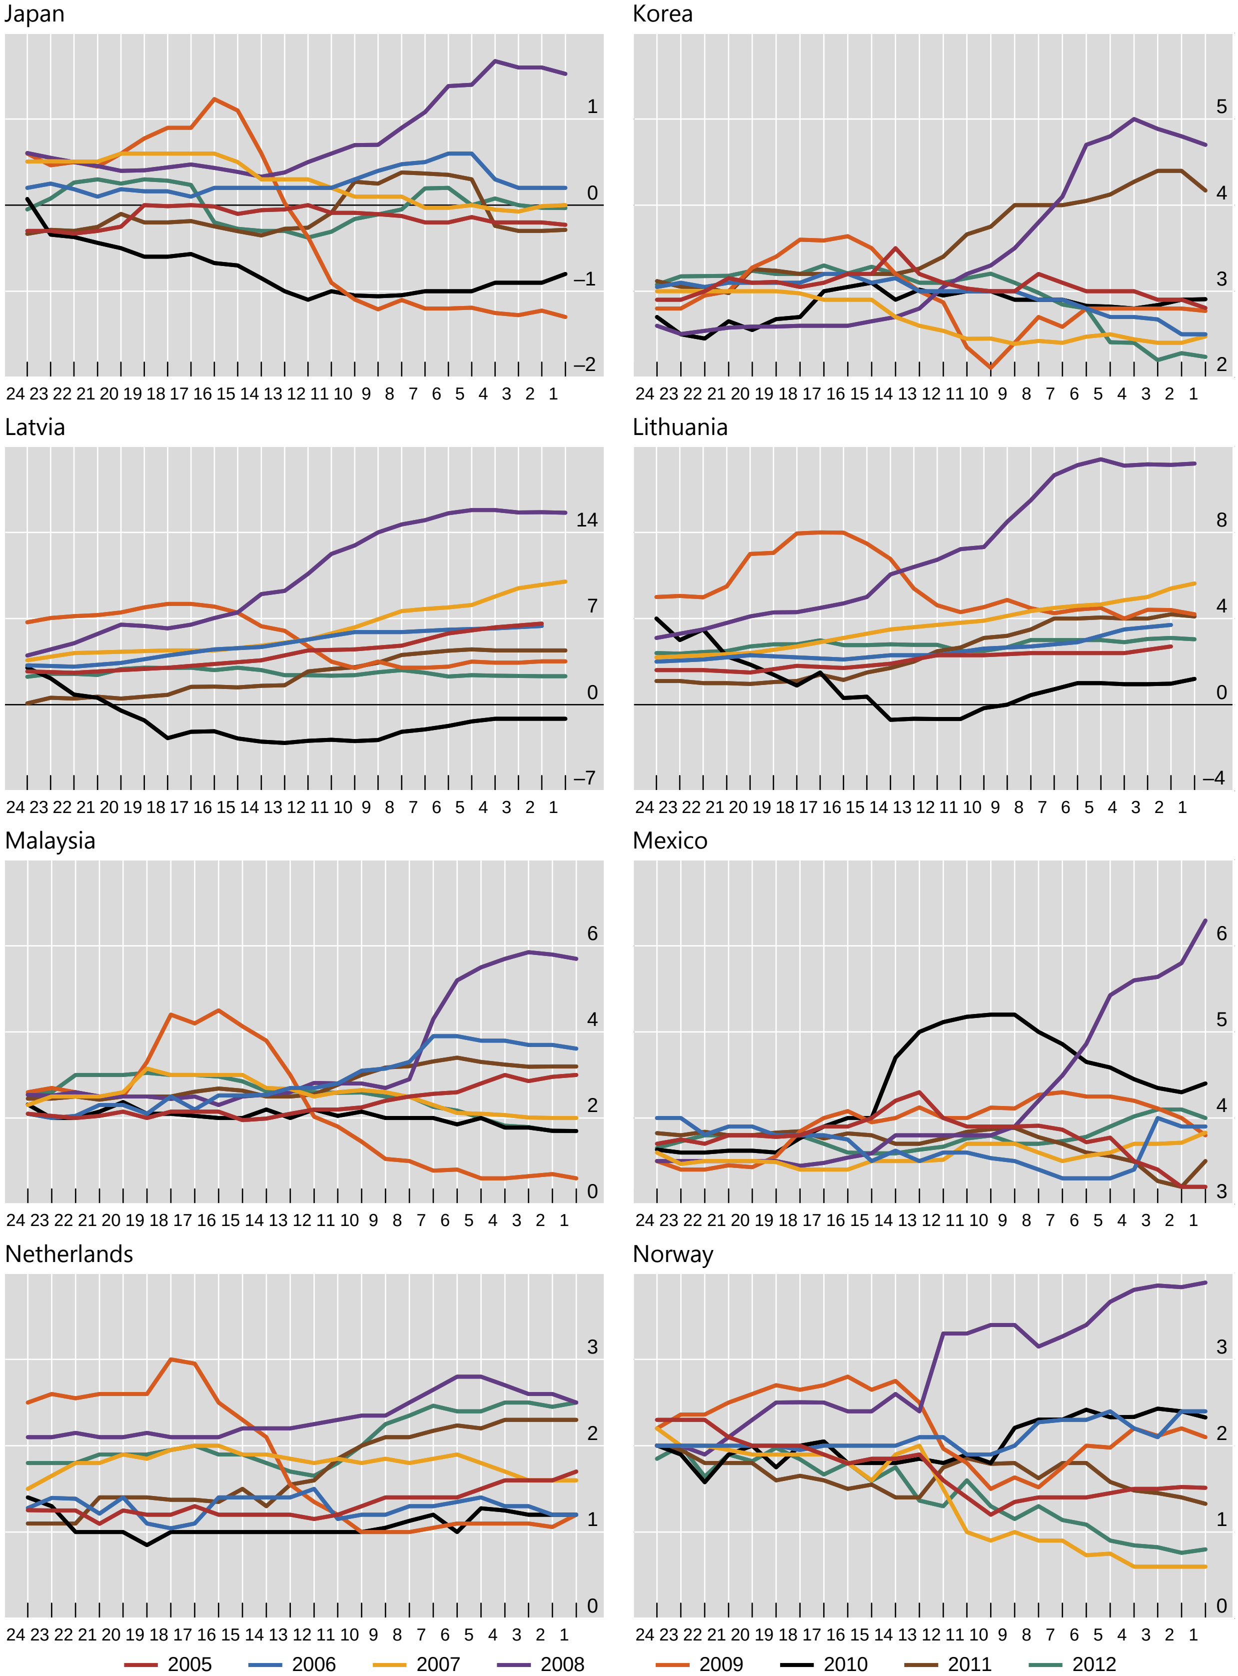

Horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.

Forecasts of headline inflation at different horizons (continued) Graph A1

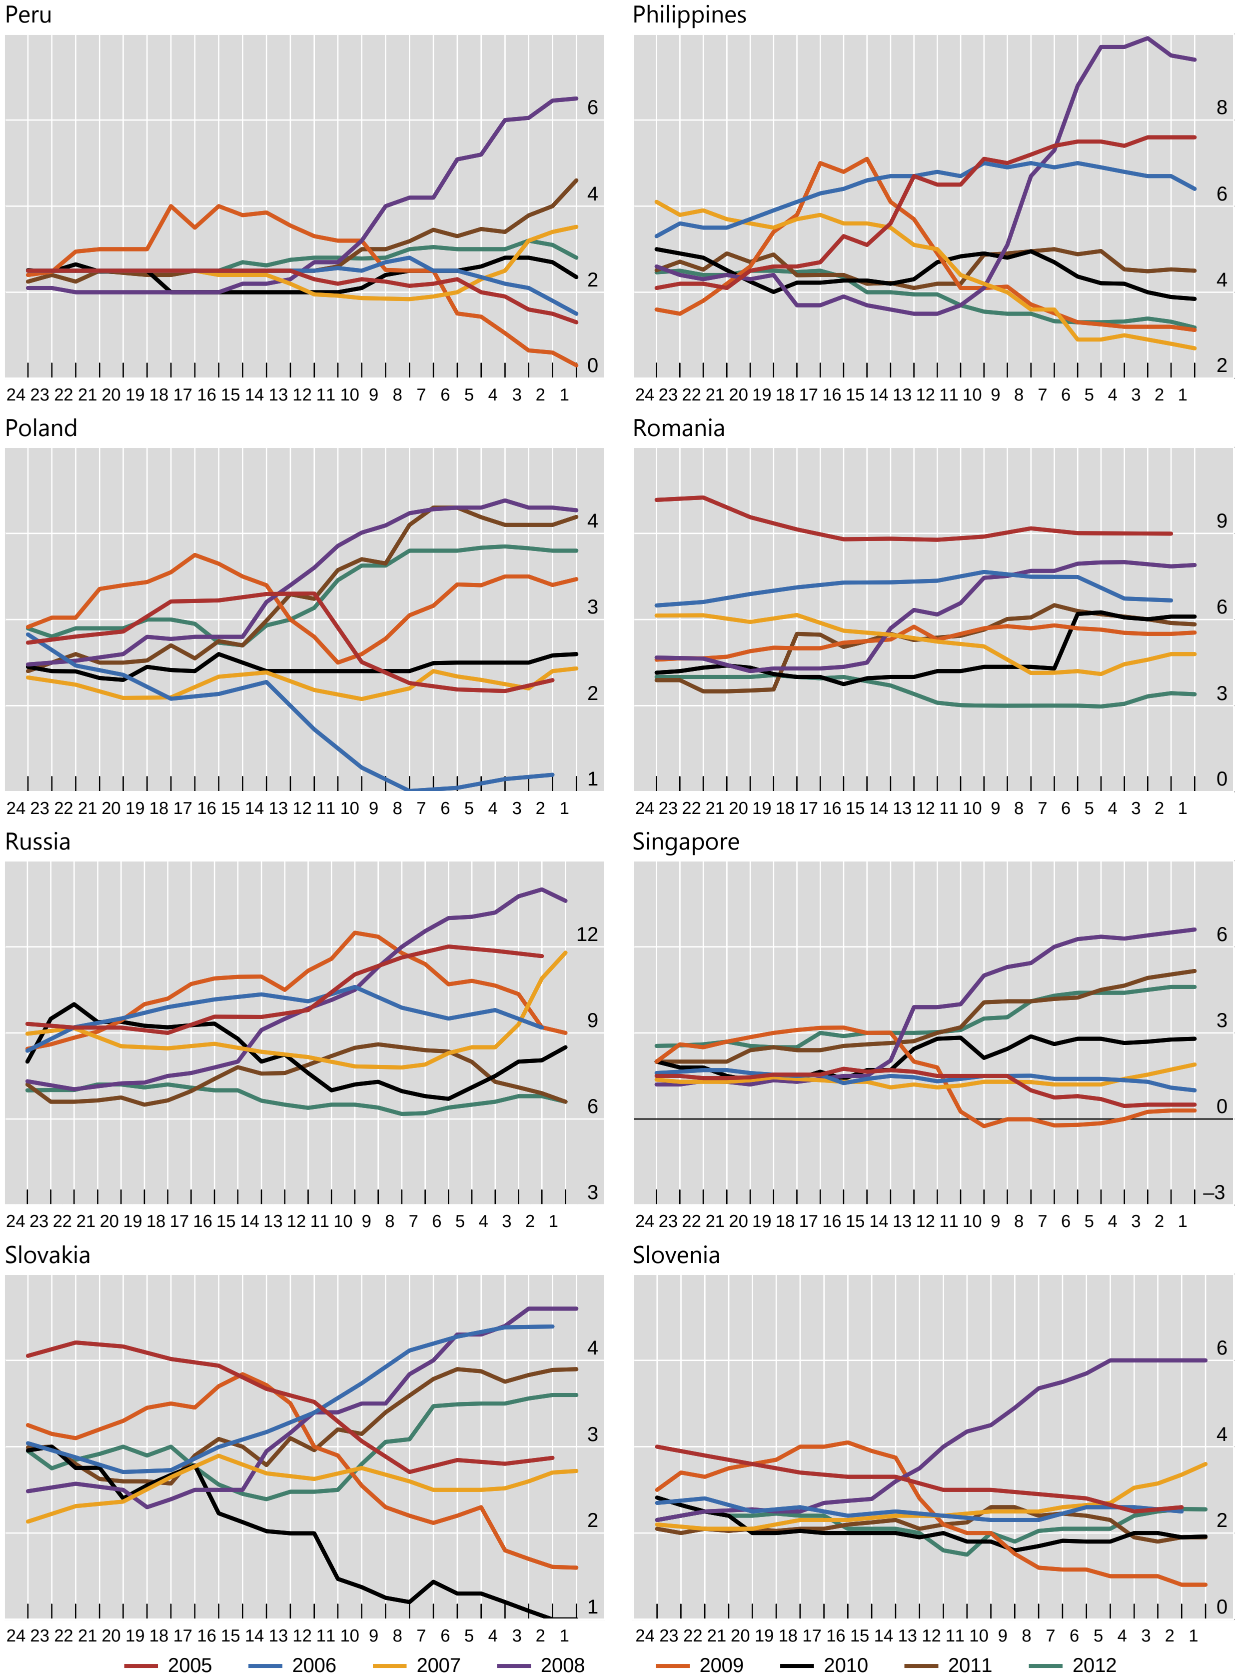

Horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.

Forecasts of headline inflation at different horizons (continued) Graph A1

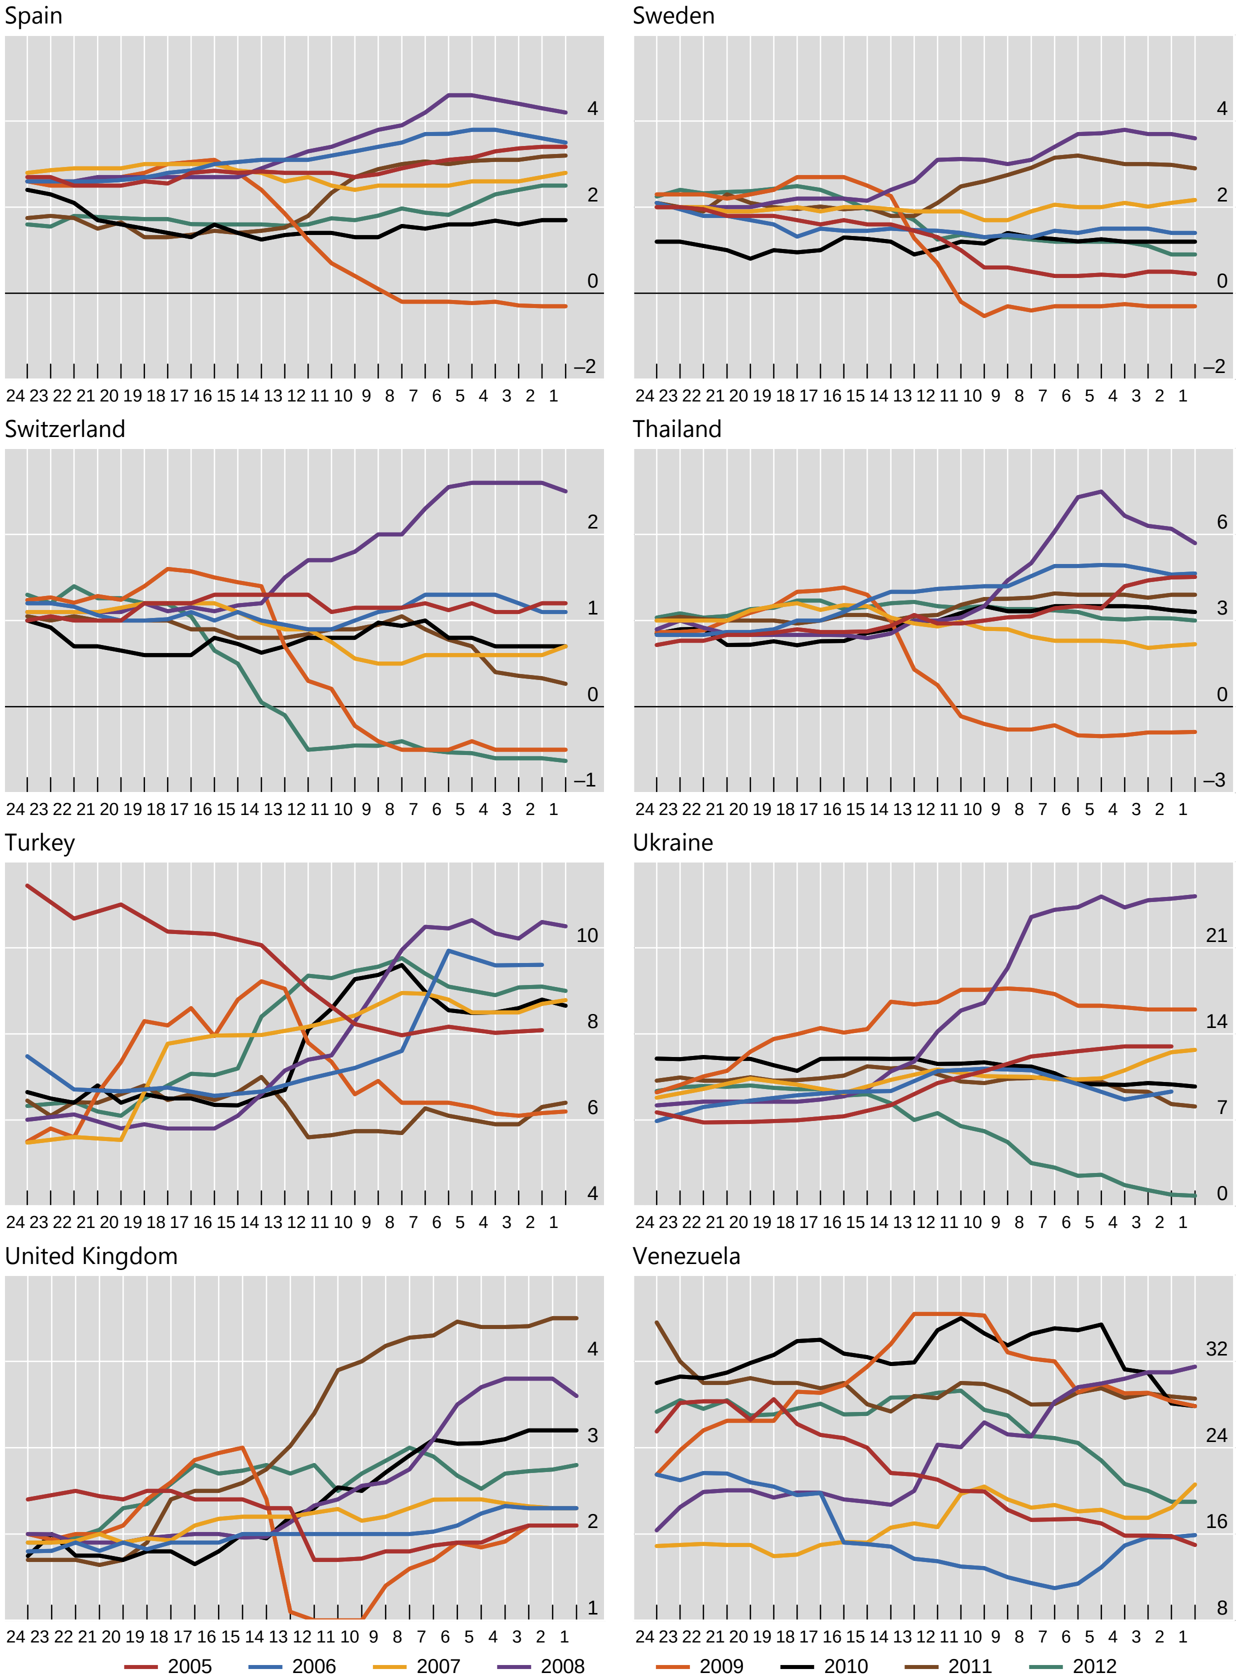

Horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.
Source: Consensus Economics ©.

Decay function and estimated long-run anchor

Graph A2

Decay function[1]

Estimated $\alpha(h)$

Estimated long-run anchor

In per cent

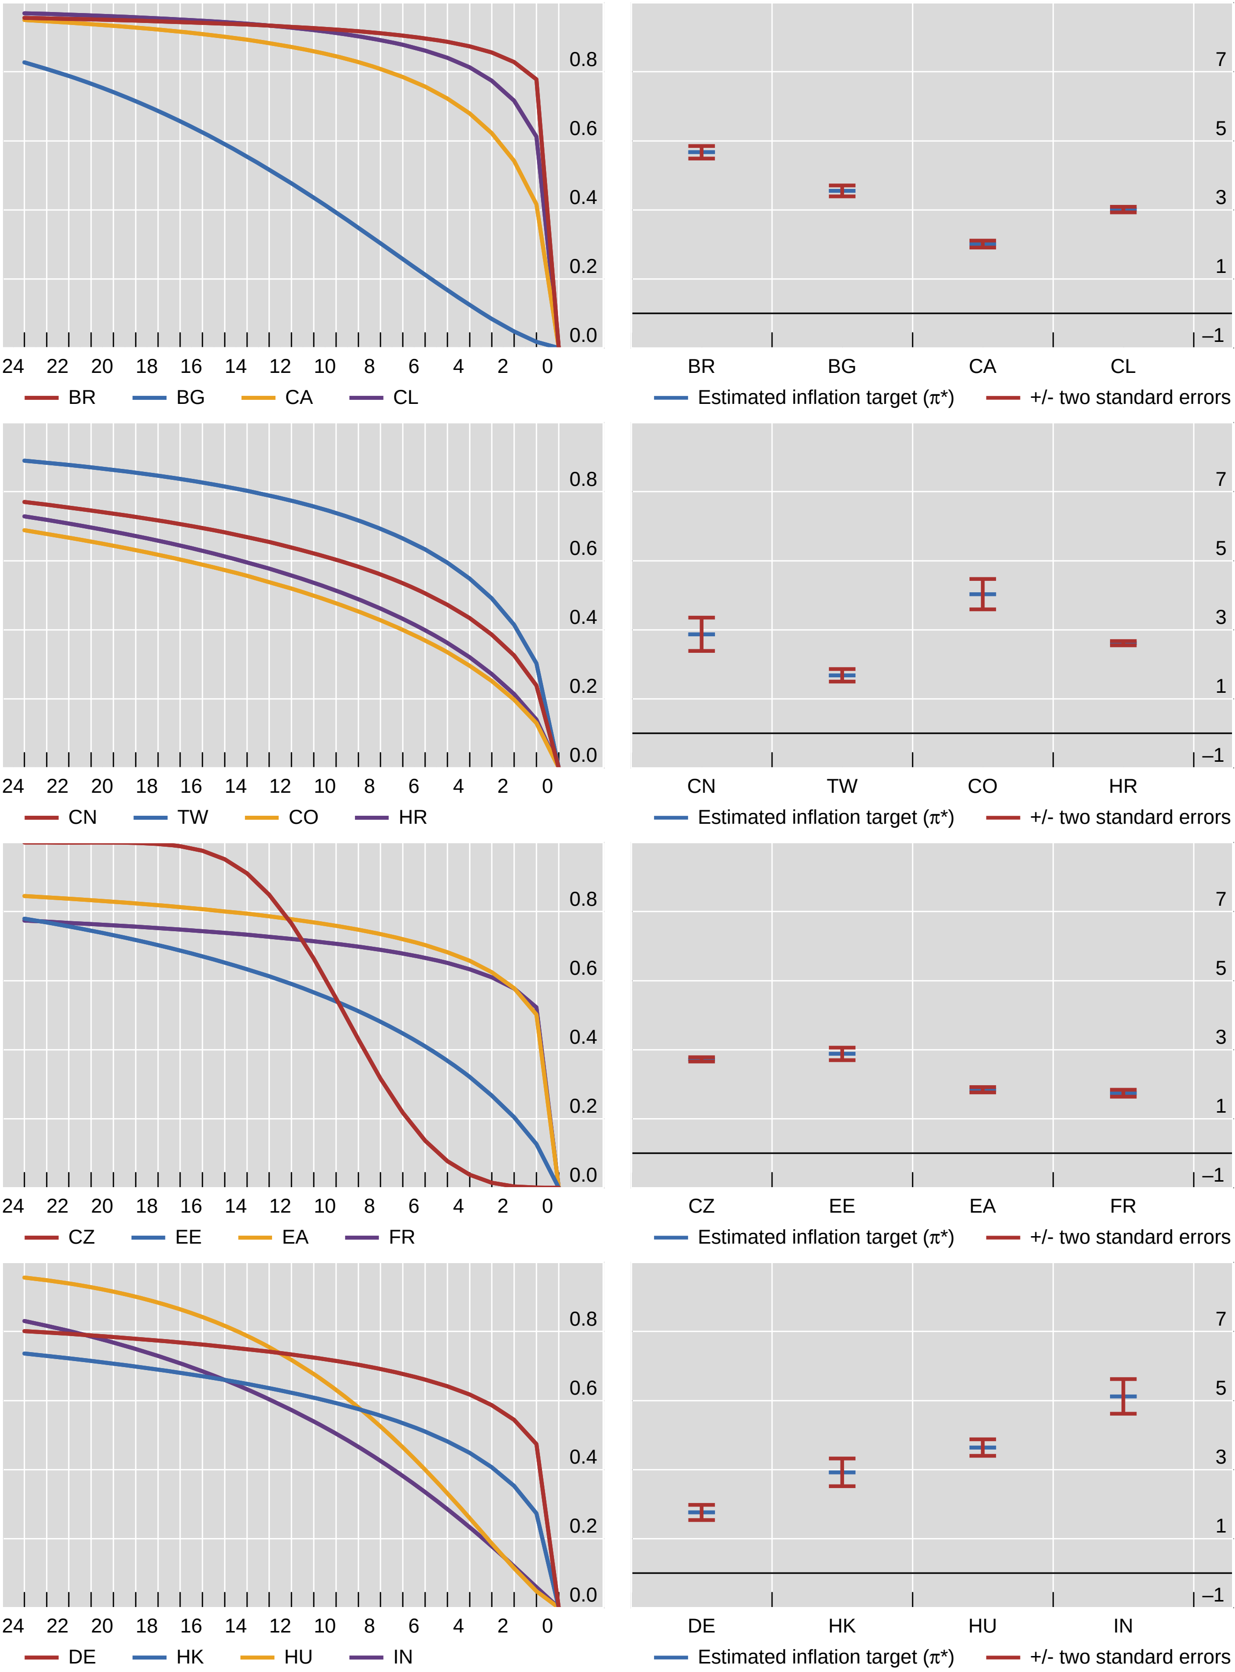

Decay function and estimated long-run anchor (continued)

Graph A2

Decay function[1] Estimated $\alpha(h)$

Estimated long-run anchor In per cent

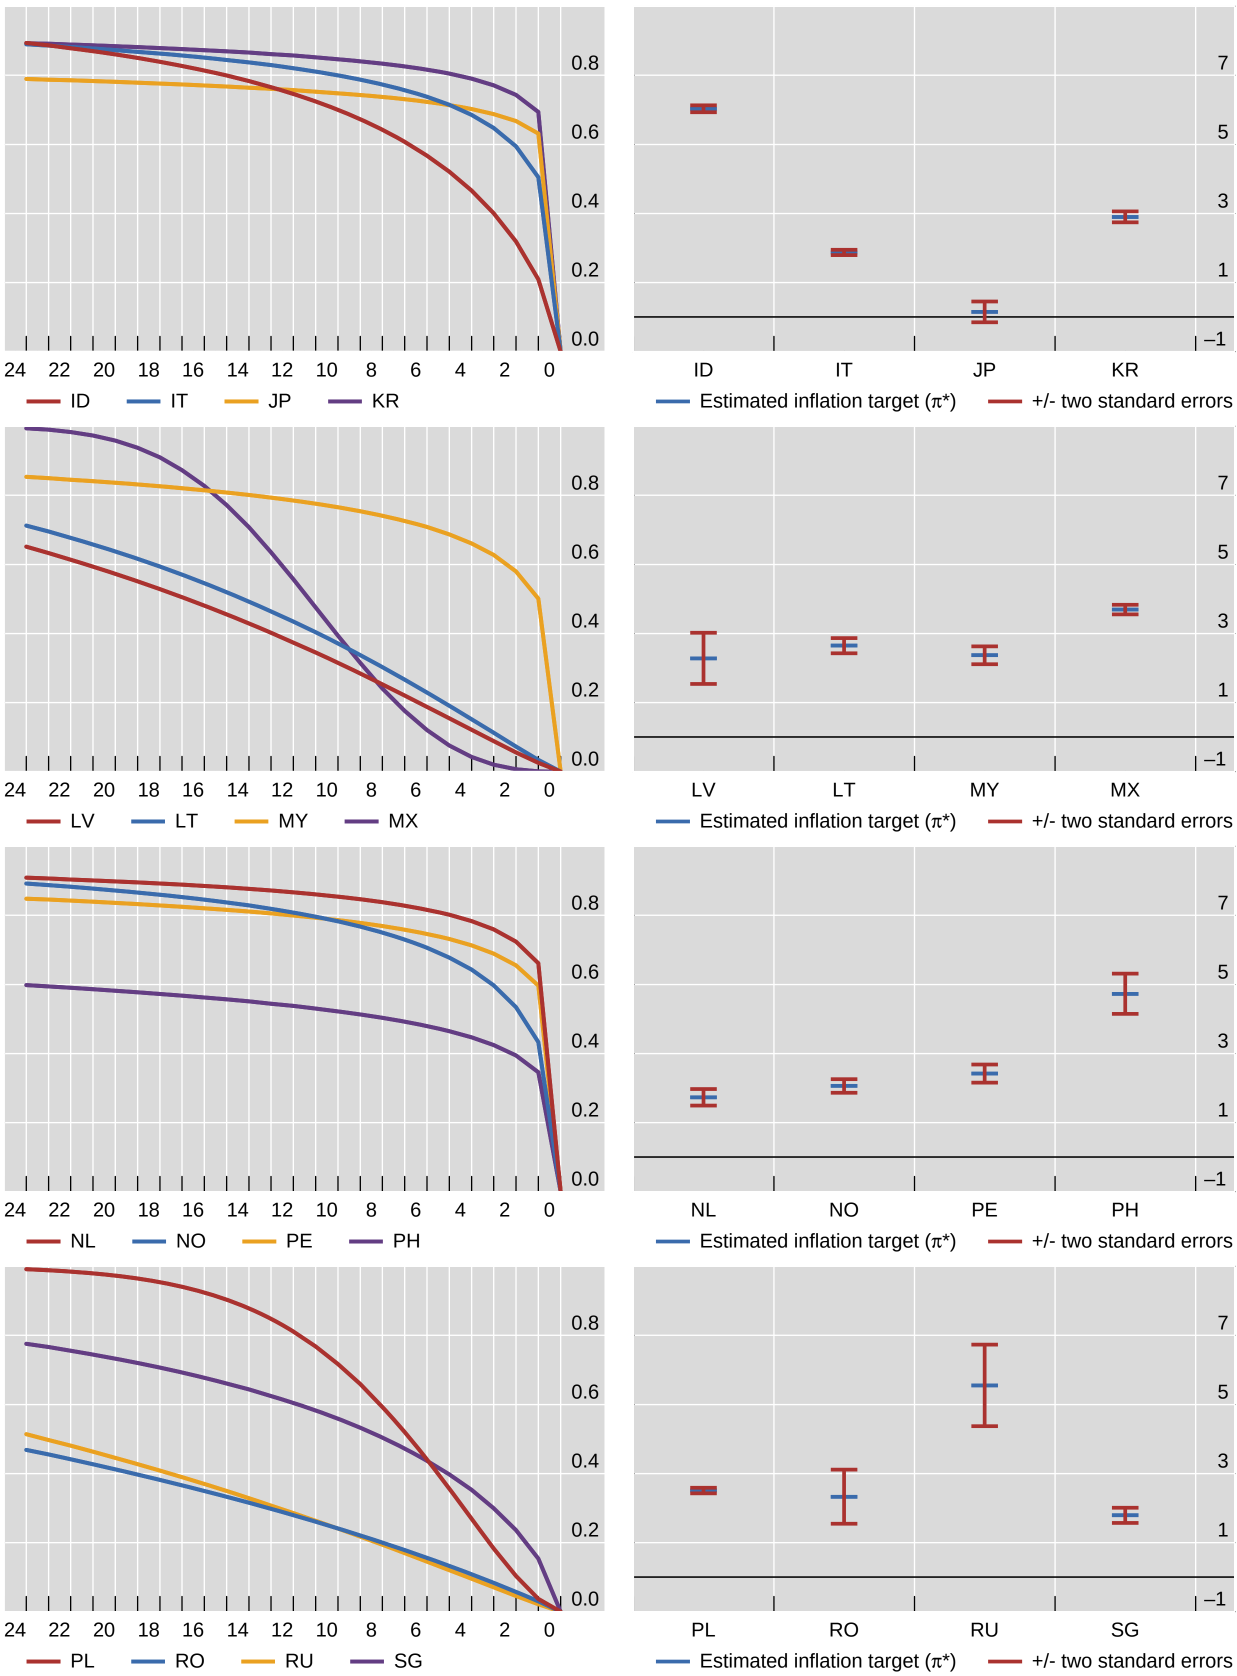

Decay function and estimated long-run anchor (continued) Graph A2

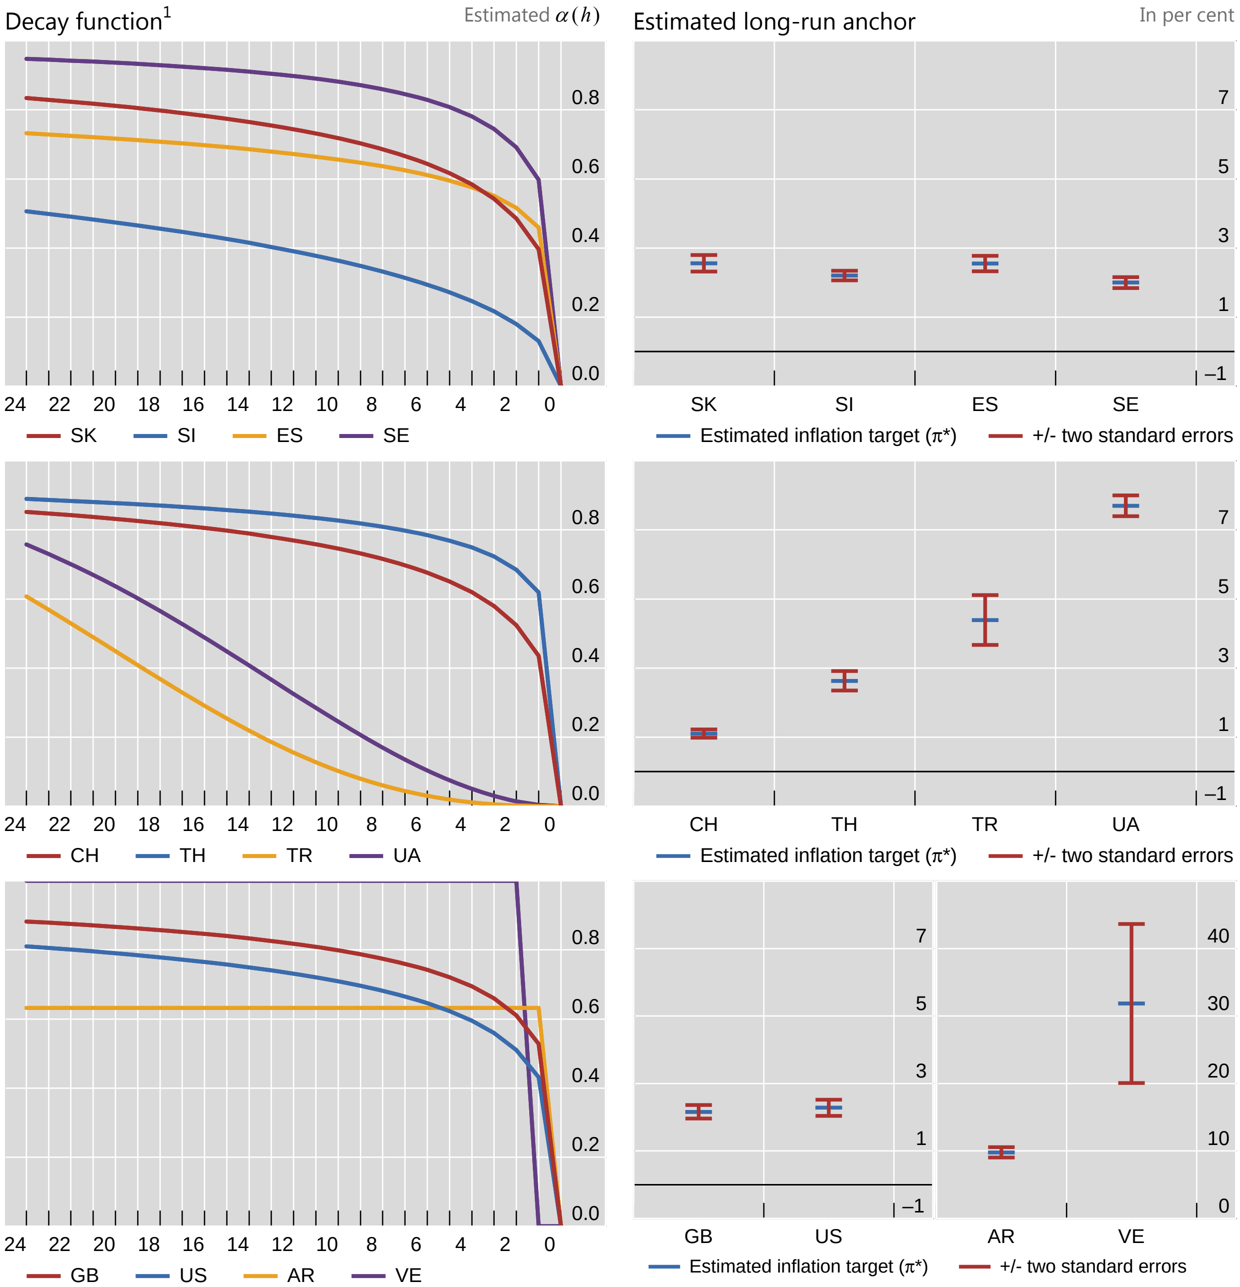

BR = Brazil; BG = Bulgaria; CA = Canada; CL = Chile; CN = China; TW = Chinese Taipei; CO = Colombia; HR = Croatia; CZ = Czech Republic; EE = Estonia; EA = Euro area; FR = France; DE = Germany; HK = Hong Kong SAR; HU = Hungary; IN = India; ID = Indonesia; IT = Italy; JP = Japan; KR = Korea; LV = Latvia; 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; AR = Argentina; VE = Venezuela

[1] Horizontal axis represents the forecast horizon, defined as the number of months before the end of the calendar year being forecast.

Source: Authors' calculations.