

Inferring Inflation Expectations from Fixed-Event Forecasts*

Diego Winkelried
Universidad del Pacífico (Lima, Perú)

Often, expected inflation measured by surveys is available only as fixed-event forecasts. Even though these surveys do contain information of a complete term structure of expectations, direct inferences about them are troublesome. Records of fixed-event forecasts through time are associated with time-varying forecast horizons, and there is no straightforward way to interpolate such figures. This paper proposes an adaptation of the measurement model of Kozicki and Tinsley (2012) to suit the intricacies of fixed-event data. Using the Latin American Consensus Forecasts, the model is estimated to study the behavior of inflation expectations in four inflation targeters (Chile, Colombia, Mexico, and Peru). For these countries, the results suggest that the announcement of credible inflation targets has been instrumental in anchoring long-run expectations.

JEL Codes: C32, E37, E52.

1. Introduction

Expectations about future inflationary outcomes play a key role in macroeconomic analysis. For instance, the determination of aggregate prices in modern macroeconomic models is often summarized

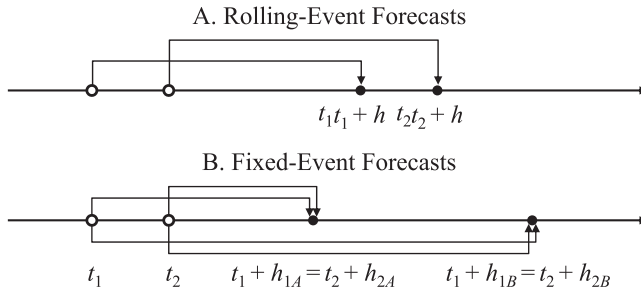
*I am indebted to Carlos Barrera, Carlos Casas, Eduardo Morón, Javier Torres, and Marco Vega for insightful discussion, and to participants at the XXXII Economists Meeting of the Central Reserve Bank of Peru and at the Viernes CIUP seminar of Universidad del Pacífico for their comments. I also gratefully acknowledge the suggestions of the associate editor and an anonymous referee for improving the content and clarity of this article. I would also like to thank Natalia Guerrero for her research assistance. I alone am responsible for any errors that may remain and for the views expressed in this article. Author contact: Faculty of Economics and Finance, Universidad del Pacífico, Av. Salaverry 2020, Lima 11, Perú. E-mail: winkelried.dm@up.edu.pe.

in expectations-augmented Phillips curves. Also, real interest rates, which provide a rough measure of the monetary policy stance once compared with suitable reference levels, can be computed by deducting inflationary expectations from nominal interest rates. Finally, many countries conduct their monetary policy within an explicit inflation-targeting regime, and the pillar of such a monetary arrangement is to maintain expectations anchored to a predetermined target. Being able to track and anticipate the behavior of expectations, thus, seems to be an important ingredient for the successful implementation of such a policy regime.

Economic surveys, where a number of participants are asked to produce forecasts of future inflation, provide a direct source of information on expected inflation.¹ As illustrated in figure 1, two types of data structures emerge.² The first one corresponds to “rolling-event” forecasts (REFs, henceforth), such as those recorded in the Livingston Survey for the United States or the Gallop Poll for the United Kingdom. Here, the survey collects h -period-ahead forecasts, so for every new release of the survey, the “event” to be forecast “rolls” forward. In other words, in an REF the horizon h is fixed, and the target date is always separated h periods from the forecast origin. The bulk of the literature that uses empirical measures of expectations from surveys, either to assess their rationality (e.g.,

¹An alternative approach is to deduce expected inflation from the difference between nominal and indexed bond yields. Such an approximation, nonetheless, is often biased and more volatile than survey measures, as it also captures factors that are not directly linked to inflation expectations, such as movements of risk or liquidity premia. Consistently with this critique, Ang, Bekaert, and Wei (2007) and Chernov and Mueller (2012) find that survey measures categorically outperform financial measures in forecasting inflation.

²We focus our analysis on *quantitative* (i.e., point) forecasts. A large literature, in contrast, favors using *qualitative* or *directional* forecasts (“would inflation be higher or lower?”), arguing that survey respondents are unable to produce trustworthy numerical predictions. Thus, the mapping from a qualitative perception to a precise quantitative figure is left to the econometrician, using either the so-called probability approach or the regression approach (see Pesaran and Weale 2006 for a comprehensive survey). This assumes, of course, that the econometrician is better suited to perform such mapping than the survey respondent. This may be true for consumer surveys, but the argument is weakened when the respondents are experts or informed professionals, as with the surveys used in this study.

Figure 1. Rolling-Event and Fixed-Event Forecasts

Davies and Lahiri 1999) or to use them in econometric models (e.g., Mehra and Herrington 2008; Mavroeidis, Plagborg-Møller, and Stock 2014), is based on this type of measure. This is so because expectations in theoretical macroeconomic models are formulated as REFs, and also because of the widespread popularity in applied work of sources such as the Livingston Survey.

The second data structure corresponds to “fixed-event” forecasts (FEFs, henceforth) which, in contrast to REFs, have received scant attention despite the fact that FEF data are available for a much larger number of countries. The “event” to be forecast is often an annual rate, for current and subsequent years. Specifically, following figure 1B, in period t_1 survey participants are asked to forecast inflation for period $t_1 + h_1$, an h_1 -period-ahead prediction; later on, in period $t_2 > t_1$ they are asked for a forecast of inflation *for the same date*, which now corresponds to an h_2 -period-ahead forecast, where $h_2 = h_1 - (t_2 - t_1) < h_1$. The forecast event is kept fixed throughout, while the forecasting horizon *shrinks* as the timeline approaches the event. The highly reputed company Consensus Economics Inc. conducts the Consensus Forecasts monthly poll among forecasters working in the private sector (in more than seventy countries), compiling their predictions of a range of economic variables and reporting them in a fixed-event format. Other widespread sources of FEFs of global and country-specific economic variables are the International Monetary Fund’s (IMF’s) World Economic Outlook, the World Bank’s Global Economic Prospects, the OECD’s Economic Outlook, and polls conducted by central banks.

Notwithstanding their availability, FEFs are an unexploited resource to describe the evolution of expected inflation.³ We conjecture that this is so because the time-varying nature of the forecast horizons makes comparisons over time troublesome. As suggested, it seems more natural to conceptualize expectations in macroeconomic models as REFs (for instance, to estimate an expectations-augmented Phillips curve or to compute real interest rates). And even though a survey that registers FEFs for horizons h_A and h_B (see figure 1) does contain information for expectations at any intermediate horizon $h_A \leq h \leq h_B$ (for instance, twelve-month-ahead expectations are implicitly contained in current- and next-year forecasts), there is no obvious way to interpolate such figures. The purpose of this paper is to develop an empirical model to explicitly infer expectations from data on actual inflation and FEFs.

The model is a version of the *shifting-endpoint* model advanced in Kozicki and Tinsley (1998, 2001, 2012), suitably modified to deal with the intricacies of the FEF data structure. We reckon that our modifications to Kozicki and Tinsley's framework widens its scope and should be useful in practice, especially to analyze expectations in countries for which only FEFs are available. Importantly, even though the model is fitted to irregularly sampled observations associated with time-varying forecast horizons, it is capable of producing a coherent and complete term structure of inflation expectations. Thus, the model provides an interpolation method that is internally consistent with inflation dynamics and survey information.

As stressed in Kozicki and Tinsley (1998, 2001), predictions derived from time-series models of inflation are likely to be poor proxies of expected inflation. Such predictions are based on historical data only and fail to accommodate structural changes in inflation dynamics that may be well reflected in survey measures.⁴ On the

³Since the seminal contribution of Nordhaus (1987), FEFs have become popular in testing forecasts' rationality or optimality (e.g., Clements 1997; Bakhshi, Kapetanios, and Yates 2005; Timmermann 2007). FEFs are in fact ideal for such purposes. Their evolution over time provides a stream of revisions whose correlation against various information sets can be directly assessed. A further advantage is that rationality can be tested without having the data of the target variable.

⁴This is indeed the case in our empirical application below. Survey measures in all the Latin American countries of our sample correctly anticipated the trend decline in inflation during the 1990s. This is remarkable, since purely backward-looking or adaptive expectations would tend to systematically overpredict inflation in such a context.

other hand, forecasts from surveys are often responsive to the latest inflation observations for short-term horizons and relatively unresponsive for long-term horizons. This pattern is difficult to reconcile with model-based multi-period forecasts. If the model fitted to the inflation process exhibits strong mean reversion to a fixed level, forecasts for all horizons would be too insensitive to recent inflation, whereas if the inflation model is persistent (say, because it contains a unit root), then forecasts would be excessively sensitive to inflation news.

The shifting-endpoint model is a compromise whose multi-period forecasts reflect the main dynamic features of the inflation process and also account for the behavior of available survey data. The endpoint is an unobservable variable that measures the perceptions of long-run inflation held by economic agents at a given point in time, and that directly affects actual inflation and its forecasts for arbitrary horizons. Survey records are taken as error-ridden versions of such forecasts, and so the inferred inflation expectations would exploit all available information from inflation and surveys.

The remainder of the document is organized as follows. Section 2 presents the shifting-endpoint model, discusses its main properties, and describes some variants. The model can be represented in state-space form, and the moving nature of the forecast horizons in FEFs implies that this representation is time varying. The Kalman filter can be used for estimation and inference, even when survey expectations are sampled less frequently than inflation. Section 3 presents empirical results for four Latin American countries, using data from the Latin American Consensus Forecasts survey. The evolution of the predicted endpoints reveals the key role that expectations have played in these countries, first to reduce inflation to single-digit levels and then to keep it stable, as well as the importance of the central bank's credibility in anchoring long-run expectations. This section also illustrates the potential usage of the proposed shifting-endpoint model to infer about inflation expectations in real time, an empirical exercise that is likely to be valuable for the routine conduct of monetary policy in an inflation-targeting framework. Finally, section 4 presents concluding remarks and questions for future research.

2. Econometric Framework

Next, we describe the workings of the shifting-endpoint model for inflation and discuss its statistical treatment. We also present model variants that can be evaluated within the same analytical framework. In what follows, $x_{t|s}$ denotes the expectation of a random variable x_t , formed by economic agents conditional on the information up to and including period s .

2.1 Shifting-Endpoint Model for Inflation

The limiting conditional forecast of inflation is given by

$$\mu_t = \lim_{h \rightarrow \infty} \pi_{h|t}, \quad (1)$$

where t denotes the time subscript of the information set on which expectations are conditioned. Thus, given information up to and including period t , μ_t measures the perceived level at which inflation would eventually stabilize, i.e., the *endpoint*.

We assume that long-run expectations are formed in a weakly rational manner, in the sense of Nordhaus (1987): changes in perception are unpredictable or, more formally, μ_t is a martingale with respect to its own past. In other words, if agents can anticipate future changes to their long-run perceptions, then such changes should be immediately incorporated in their current perceptions, as in the law of iterated expectations. This behavior can be satisfactorily modeled by assuming that the endpoint follows a random walk,

$$\mu_t = \mu_{t-1} + \nu_t, \quad (2)$$

where ν_t is an innovation satisfying $\nu_{t|s} = 0$ for $s < t$. The endpoint μ_t is treated as an unobservable variable, and the main purpose of the analysis is to infer about its state using inflation and survey data.

The notion of a varying endpoint for expected inflation can be easily accommodated in a parametric forecasting model. Let π_t denote inflation at period t and suppose that the expectations in period t are formed using information up to and including period $t - 1$. Define also $\Phi(L) = \phi_1 + \phi_2 L + \dots + \phi_{p-1} L^{p-2} + \phi_p L^{p-1}$ as a polynomial in the lag operator ($L^k \pi_t = \pi_{t-k}$), and assume that

the roots of $1 - \Phi(z)z = 0$ lie outside the unit circle (we make some allowance for a unit root in $1 - \Phi(z)z$ later). Inflation dynamics are captured by

$$\pi_t = \Phi(L)\pi_{t-1} + (1 - \Phi(1))\mu_{t-1} + \epsilon_t, \quad (3)$$

where ϵ_t is an inflation shock, assumed to be uncorrelated with ν_t at all lags and leads. Under the aforementioned assumptions, $\pi_t - \mu_{t-1}$, the deviation of inflation from its latest perceived long-run level, is a zero-mean stationary process. By construction, given the information of period $t - 1$, inflation is expected to converge to μ_{t-1} as the forecasting horizon increases. Thus, as is further discussed in section 2.5, the dynamic specification (3) allows us to disentangle the effects of the shifting endpoint from short-run fluctuations in inflation.

2.2 Incorporating Survey Data

It is convenient to write the forecasting model in companion form. Let

$$z_t = \begin{bmatrix} \pi_t \\ \pi_{t-1} \\ \vdots \\ \pi_{t-p+1} \end{bmatrix} \quad \text{and} \quad C = \begin{bmatrix} \phi_1 & \phi_2 & \cdots & \phi_{p-1} & \phi_p \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{bmatrix}, \quad (4)$$

where z_t is a p -vector with current and past inflation data and C is the companion matrix. Let also I_p be the $p \times p$ identity matrix, i_p be a p -vector of ones, and e_1 be a unit p -vector with a one in the first element and remaining elements zero. Then, (3) corresponds to the first row of the companion form

$$z_t = Cz_{t-1} + (I_p - C)i_p\mu_{t-1} + e_1\epsilon_t. \quad (5)$$

Multi-step forecasts of inflation based on this model are $\pi_{t+h|t-1} = e_1'z_{t+h|t-1}$, where

$$z_{t+h|t-1} = C^{h+1}z_{t-n} + (I_p - C^{h+1})i_p\mu_{t-1}. \quad (6)$$

Survey data registered in period t represent participants' forecasts conditional on inflation up to and including period $t - 1$. Moreover, consider that the survey collects information for m different forecasts F_{it} , each associated with a different horizon that may be time varying h_{it} ($i = 1, 2, \dots, m$). Taking the inflation forecast from the shifting-endpoint model as an approximation of the survey expectation, we may write

$$F_{it} = \pi_{t+h_{it}|t-1} + \varepsilon_{it}, \quad (7)$$

where ε_{it} is an approximation error that captures discrepancies between the implicit (and unknown) forecasting model of survey participants and the shifting-endpoint model, as well as possible measurement errors in survey data. It is important to note that the errors in (7) do not reflect differences between actual inflation and forecasts. This is because both the model-based forecasts and survey records are conditioned on $t - 1$, which implies that $\text{cov}(\varepsilon_t, \varepsilon_{it}) = \text{cov}(\nu_t, \varepsilon_{it}) = 0$ for $i = 1, 2, \dots, m$. Thus, unlike forecast errors, there is no reason to expect them to be serially correlated; hence, $\text{cov}(\varepsilon_{it}, \varepsilon_{js}) = 0$ for $t \neq s$ and any i or j .

On the other hand, revisions of forecasts of different target dates made at the same time are likely to be highly correlated (Clements 1997; Bakhshi, Kapetanios, and Yates 2005). News at period t that leads to a revision in the forecast at horizon h_{it} is also likely to produce a revision in the forecast at horizon h_{jt} ($i \neq j$). Following (7), even though an important part of these co-movements would be accounted for by changes in the endpoint, we expect the approximation errors to be contemporaneously correlated $\sigma_{ij} = \text{cov}(\varepsilon_{it}, \varepsilon_{jt}) \neq 0$. Note that whether σ_{ij} is zero or not is a straightforward testable hypothesis that can be formally addressed.

2.3 State-Space Form

The statistical treatment of the model is based on its state-space representation (see Harvey 1989, ch. 5). The Kalman filter yields linear least-squares predictions of the unobserved endpoint μ_t based on current and past observations, along with their corresponding

mean square errors.⁵ Moreover, given the parameters of the model (ϕ_1, \dots, ϕ_p , the variances of ν_t and ϵ_t , and the covariances among $\varepsilon_{1t}, \dots, \varepsilon_{mt}$), a Gaussian likelihood function can be evaluated from the one-step-ahead prediction errors produced by the Kalman filter. This function can be maximized numerically, thereby providing (quasi) maximum-likelihood estimates of the unknown parameters.

The law of motion of the unobservable endpoint (2) constitutes the scalar transition (or state) equation. Upon stacking the shifting-endpoint model (3), the conditional forecasts (6), and their relationship with various data points from surveys (7), we obtain the measurement equations

$$y_t = A_t z_{t-1} + B_t \mu_{t-1} + w_t, \quad (8)$$

where y_t is an $(m+1)$ -vector that contains current date information on inflation and surveys; A_t and B_t are, respectively, an $(m+1) \times p$ matrix and $(m+1)$ -vector of coefficients; and w_t is an $(m+1)$ -vector of measurement errors. The entries of A_t and B_t depend on h_{it} and thus may be time varying, capturing the shrinking nature of the forecasting horizon for FEFs. This is the most important difference with respect to the model analyzed by Kozicki and Tinsley (2001, 2012).

More explicitly, let $a(h) = e_1' C^{h+1}$ (a p -row vector) and $b(h) = 1 - a(h)i_p$ (a scalar). Then,

$$\begin{bmatrix} \pi_t \\ F_{1t} \\ F_{2t} \\ \vdots \\ F_{mt} \end{bmatrix} = \begin{bmatrix} a(0) \\ a(h_{1t}) \\ a(h_{2t}) \\ \vdots \\ a(h_{mt}) \end{bmatrix} z_{t-1} + \begin{bmatrix} b(0) \\ b(h_{1t}) \\ b(h_{2t}) \\ \vdots \\ b(h_{mt}) \end{bmatrix} \mu_{t-1} + \begin{bmatrix} \epsilon_t \\ \varepsilon_{1t} \\ \varepsilon_{2t} \\ \vdots \\ \varepsilon_{mt} \end{bmatrix}, \quad (9)$$

where each vector and system matrix in (8) is implicitly defined.

It is worth mentioning that although measurement equations (9) are designed to accommodate m different survey responses at different horizons, they are flexible enough to adapt to many other

⁵If desired, a refined prediction conditioned on all the available information in the sample can be obtained by means of a smoothing algorithm that uses the output of the Kalman filter.

structures. One case would be if the surveys contain both REFs and FEFs: for REFs the horizons are fixed $h_{it} = h_i$, whereas for FEFs the forecasting horizons will vary with t in a deterministic fashion. For instance, in surveys such as the Consensus Forecasts, the data from month M_t refer to forecasts by the end of year i , with $i = 1$ being the current year, so the forecasting horizons (in months) evolve deterministically as $h_{it} = 12i - (M_t - 1)$.

Surveys may also report *averages* of forecasts over R different horizons. In this case, the corresponding measurement equation would read $F_t = (1/R) \sum_{r=1}^R [a(h_{rt})z_{t-1} + b(h_{rt})\mu_{t-1}] + \varepsilon_t$. We did attempt to incorporate information of this kind in our empirical application below, but with marginal effects on the results. Thus, we ignore such forecasts and stick to the model (9) as it stands.

On the other hand, we have implicitly interpreted F_{1t}, \dots, F_{mt} as aggregate responses for different forecasting horizons. This is just for expositional convenience, since (9) can handle responses made by different agents, by simply allowing subindex i to denote a forecaster/horizon pair. For instance, for two forecasters and two horizons, F_{1t}, F_{2t} can represent the results for the first participant, and F_{3t}, F_{4t} for the second, with $h_{1t} = h_{3t}$ and $h_{2t} = h_{4t}$ indicating that there are only two target horizons. In this case, more structure to the measurement error covariances, in the spirit of Davies and Lahiri (1999), may be appropriate—for instance, by letting $\varepsilon_{it} = R_i + H_t + \text{error}_{it}$, where R_i and H_t are, respectively, respondent and horizon effects.⁶

2.4 Missing Data

Generally, survey data are available less frequently than inflation data (which are assumed to be available for all periods in the sample), so for certain periods some of the observations in the y_t vector will be missing. The treatment of missing observations is perhaps one of the clearest advantages of using the Kalman filter to process

⁶A further extension of the model allows expectations to be formed at different moments. The measurement equation for an expectation formed with information up to period $t - n$ ($n \geq 1$) is $F_t = a(h_t + n - 1)z_{t-n} + b(h_t + n - 1)\mu_{t-n} + \varepsilon_t$. Thus (8) becomes $y_t = A_t Z_t + B_t \bar{\mu}_t + w_t$, where Z_t and $\bar{\mu}_t$ stack, respectively, the lags of z_t and μ_t associated with each information set. The number of columns of the now sparse matrices A_t and B_t should vary accordingly.

time-series models, as the filter requires only minor modifications to deal with such a problem (see Harvey 1989, p. 144).

Consider the case when only $\bar{m}_t < m$ forecasts in y_t are available at period t , and let D_t be the $(\bar{m}_t + 1) \times (m + 1)$ selection matrix that collects non-missing elements; so, for instance, $y_t^* = D_t y_t$ is an $(\bar{m}_t + 1)$ -vector that contains observable data. Note that D_t is formed by $\bar{m}_t + 1$ rows of the identity matrix of order $m + 1$. Then, upon pre-multiplying the measurement equations (8) by D_t ,

$$y_t^* = A_t^* z_{t-1} + B_t^* \mu_{t-1} + w_t^*, \quad (10)$$

where A_t^* , B_t^* , and w_t^* contain the rows of A_t , B_t , and w_t , respectively, that correspond to the available observations in y_t^* . The state variable μ_t is not affected by the transformation, so the Kalman filter can be applied normally to the observable system (10) in period t . The missing observations contribute neither to the state predictions nor to the likelihood function.

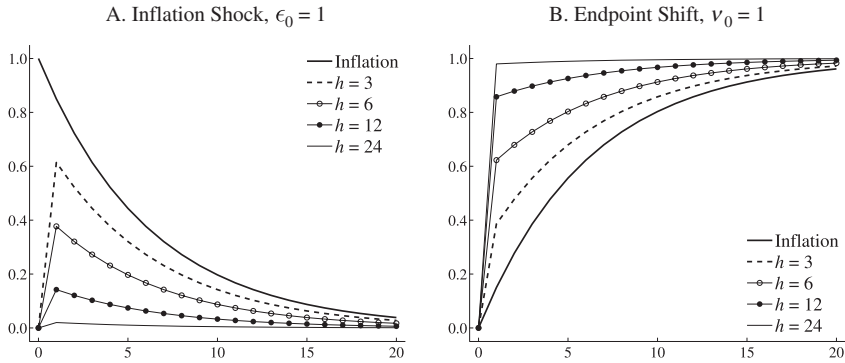
2.5 Signal Extraction and the Term Structure of Expectations

From (6) and (7), we can readily verify that

$$\pi_{t+h|t-1} = a(h)z_{t-1} + b(h)\mu_{t-1}, \quad (11)$$

so that expected inflation is estimated as a linear combination of the latest realizations of inflation z_{t-1} and the perceived long-run level μ_{t-1} . Note $\lim_{h \rightarrow \infty} a(h) = 0$ and $\lim_{h \rightarrow \infty} b(h) = 1$. Thus, as h increases, expectations converge from a short-term forecast dominated by recent history to the endpoint (1). By the same token, equation (7) implies that (roughly) $F_{it} \simeq \pi_{t-1} + \varepsilon_{it}$ for small h_{it} , and $F_{it} \simeq \mu_{t-1} + \varepsilon_{it}$ for large h_{it} . Therefore, in the process of extracting the signal μ_t , more importance is given to long-term survey expectations (when available), even though inflation data and shorter-term expectations contain also information about this hidden state. Kozicki and Tinsley (2012) provide a more detailed discussion on the information content of the observable variables. See also the discussion in section 3.3 below.

These properties of the model are illustrated in figure 2. For $\Phi(L) = \Phi(1) = 0.85$, the figure shows the response of inflation and expectations formed with information up to period $t - 1$, to the two

Figure 2. Inflation and Expectations Responses to Shocks

Note: The horizontal axis shows the number of periods (months) after the shock occurs at period 0. The vertical axis shows the response of inflation and expectations in percent terms.

shocks of the system. An inflation shock $\epsilon_0 = 1$, panel A, produces a persistent deviation in the inflation rate from the endpoint. Given the transitory nature of the shock, for short-run horizons (say, $h \leq 6$) expected inflation reflects the temporarily higher current inflation; in contrast, expectations for longer horizons (say, $h \geq 12$) are virtually unaffected, as agents correctly anticipate that the influence of the shock would vanish almost completely by the end of the forecast horizon. On the other hand, expectations are sensitive to shifts in the endpoint $\nu_0 = 1$, panel B. Such a shock causes a sluggish response in inflation, which displays a smooth transition towards its new long-run level, that in turn affects short-run expectations; on the other hand, the shift passes through immediately and almost completely to long-run expectations, as agents can anticipate that the change in the endpoint will be completely transferred to inflation by the end of the forecast horizon.

The state-space form includes error-ridden versions of (11) for time-varying horizons where survey data is available. Thus, even though survey data may be limited to infrequently sampled observations and to selected horizons, once the model is estimated, it can be used to construct a complete term structure of expected inflation. Put differently, in practice we only observe some points of the responses shown in figure 2, but the model allows us to interpolate the entire profile of responses in a coherent, data-consistent way.

2.6 Model Variants

In our empirical exploration, we also analyze three variants of the model within the same state-space framework, in order to highlight the most important features of the shifting-endpoint model that uses both inflation and survey data. We are particularly interested in assessing the model's ability to simultaneously explain inflation and expectations data under competing specifications.

The first variant assumes a constant endpoint (CE), $\mu_t = \mu$. This can be achieved by setting $\text{var}(\nu_t) = 0$ and treating the initial condition $\mu_0 = \mu$ as an additional parameter to be estimated. Provided that $\Phi(1) < 1$, inflation is treated as a stationary, mean-reverting process.

In contrast, the second variant imposes a unit root (UR) to the autoregressive polynomial in (3), $\Phi(1) = 1$. In this case $C^h i_p = i_p$ for any $h > 0$, and the companion form (5) reduces to $z_t = C z_{t-1} + e_1 \epsilon_t$. Since it turns out that $B_t = 0$ in all measurement equations, μ_t is not identified and cannot be computed with the Kalman filter. However, μ_t can be reformulated to be the limiting forecast which continues to exist: it can be verified that $\lim_{h \rightarrow \infty} C^h = \bar{C}$, where $C\bar{C} = \bar{C}$, and so $\mu_t = e_1' \bar{C} z_t$. The endpoint is a moving average of the most recent inflation observations, and Kozicki and Tinsley (1998) show that, following the definition in (1), it corresponds to the permanent component of the Beveridge-Nelson decomposition of the inflation equation. The Kalman filter can still be used to evaluate the likelihood function in this case, under $\text{var}(\nu_t) = 0$.

The last variant is a model of inflation that ignores the information from surveys. The result is a univariate generalization of the local level (LL) model described in Harvey (1989, ch. 2), where inflation is decomposed as the sum of a random walk (the endpoint) and a zero-mean stationary component. This variant can be easily treated by considering a single measurement equation, the first row in (9), or by making the variances of the approximation errors arbitrarily large, $\text{var}(\varepsilon_{it}) = \kappa \rightarrow \infty$ for all i .

3. Application to Latin American Countries

After a long history of high inflation, during the 1990s many Latin American countries adopted a series of reforms that would eventually

bring inflation down to single-digit levels (see, *inter alia*, Mishkin and Savastano 2001). Even though experiences may differ in the detail, many commonalities across countries can be identified (see Corbo and Schmidt-Hebbel 2001; Quispe-Agnoli 2001). First, to facilitate the reduction of inflation, and to isolate monetary policy from political pressures, these countries granted independence to their central banks at an early stage of the stabilization effort. Then, economic authorities would adopt a rudimentary form of inflation targeting, typically by simply announcing numerical targets (or official forecasts), as a first attempt to anchor market expectations. The process of disinflation would be gradual as the central banks improved their credibility and built a reputation as inflation targeters. Once stable levels of inflation were reached, the central banks would adopt a fully fledged inflation-targeting regime, characterized by the announcement of long-run targets and the abandonment of any other nominal anchor (typically, currency depreciation or money growth).

We estimate the shifting-endpoint model of section 2 using data from four successful Latin American inflation targeters: Chile, Colombia, Mexico (all of which adopted the regime in 1999), and Peru (which adopted the regime in 2002).⁷ Following Vega and Winkelried (2005), the central banks began announcing numerical targets long before the definite adoption date: 1991 in Chile, 1995 in Colombia and Mexico, and 1994 in Peru. Thus, by analyzing the evolution of the estimated endpoints, our empirical application aims to assess the role of expectations in the initial disinflation and the subsequent periods of price stability.

3.1 Data

Monthly inflation corresponds to the percent variation of the officially targeted consumer price index. These data and the numerical targets come from each central bank's website.

Survey data, on the other hand, are extracted from the well-known Latin American Consensus Forecasts reports. These reports

⁷Brazil is also an important Latin American inflation targeter. Its inflationary experience is different from those of the listed countries, enough to deserve special treatment. Hence, for the sake of brevity, the results for the Brazilian case (which are available upon request) are not reported.

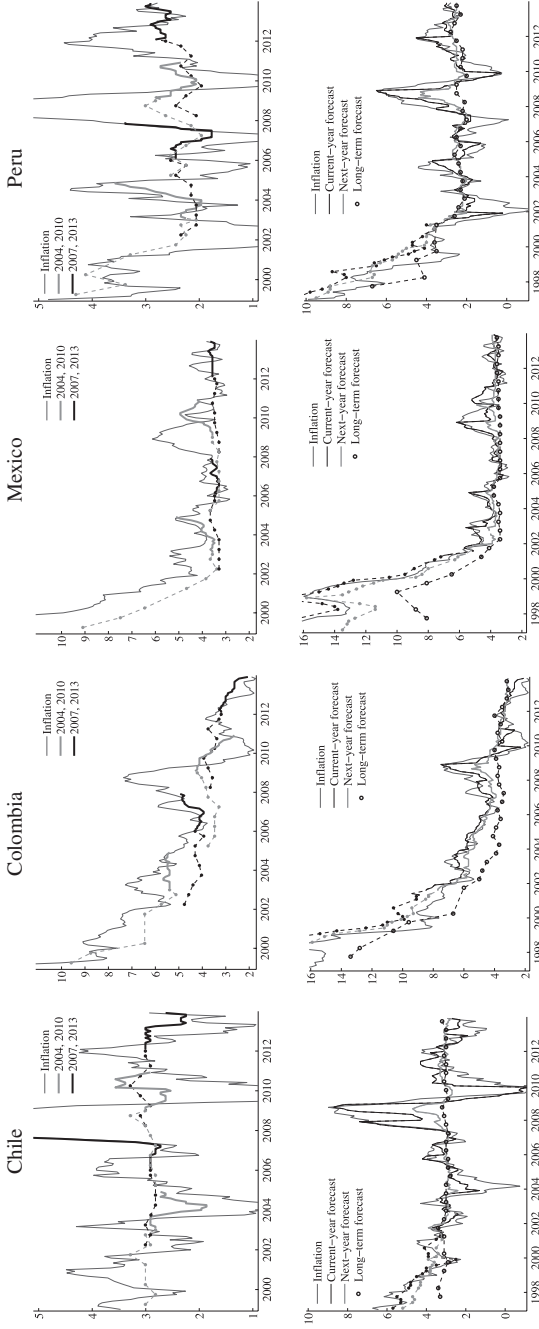
are available bi-monthly (alternate, even months) between April 1993 and April 2001, and monthly thereafter (since May 2001). Each report surveys a number of prominent financial and economic analysts and publishes their individual forecasts as well as simple descriptive statistics. The focus of our analysis is on the mean (i.e., the “consensus”) forecast F_{it} . The surveys are conducted by the middle of the month, after the inflation figure for the previous month is released. Thus, in agreement with the shifting-endpoint model assumptions, in period t expectations are formed with information up to period $t - 1$.

The typical issue of the Latin American Consensus Forecasts provides forecasts for current-year and next-year inflation. Denoting the year of the survey as year 1, these forecasts correspond to the measurements F_{1t} and F_{2t} . We refer to these variables jointly as “short-term” forecasts. In addition, the April and October issues include also “long-term” (up to ten years ahead) forecasts. Forecasts for years 3 to 6 ($F_{3t}, F_{4t}, F_{5t}, F_{6t}$) are explicitly published, whereas forecasts for longer horizons (from year 7 to year 11) are reported as an average. In practice, such averages show little variation with respect to F_{6t} and are thus discarded from our analysis with no consequential effects on our results. Thus, depending on the month, the measurement equations have a minimum of $m = 2$ horizons and a maximum of $m = 6$. As mentioned earlier, the forecasting horizons evolve deterministically as $h_{it} = 12i - (M_t - 1)$, where M_t denotes the month of the survey.

Our sample spans from February 1997 to December 2013. Thus, for each country, the estimations are based on 203 observations on inflation, 178 observations on “short-term” forecasts, and 34 observations on “long-term” forecasts.

Figure 3 displays some of the data. The first row plots the evolution of the FEFs as they approach the inflation outturn (to avoid clutter, only the events 2004, 2007, 2010, and 2013 are displayed). For a given event, each point corresponds to a different survey. The second row presents the data as they enter the regression model: the short-term forecasts F_{1t} and F_{2t} , and the longest-term forecast F_{6t} , all of them reported at the same time t . Each month is associated with a different forecast horizon, so the strong swings in F_{1t} and F_{2t} are due to a change in the event to be forecast (say, $h_{1t} = 1$ in December and then $h_{1t} = 12$ in January).

Figure 3. Inflation, Selected Fixed-Event Forecasts, and Moving-Horizon Forecasts



Note: The dots in the figure represent forecasts that are irregularly sampled (long-term forecasts in all of the sample period, and short-term forecasts until April 2001). To ease visualization, these dots are connected with linearly interpolated values represented by discontinuous lines (the interpolations are not used in the estimation). Within each row, the scale of the vertical axis is the same for Chile/Peru and Colombia/Mexico.

A preliminary analysis of these surveys provides prima facie evidence of the adequacy of the shifting-endpoint model to explain the behavior of measured expectations. An important implication of the model is that the deviation of inflation from the endpoint is expected to be a mean-reverting process around zero. Also, short-term forecasts would be influenced by recent news in inflation, whereas long-term forecasts should be determined by the endpoint. Consider the regression equation

$$F_{it} - F_{6t} = \alpha_i + \beta_i(\pi_{t-1} - F_{6t}) + \text{error}_t, \quad (12)$$

where π_{t-1} denotes the latest inflation figure (the forecast origin), and F_{6t} is the the forecast corresponding to the longest horizon available, which serves as a rough proxy of the moving endpoint. The shifting-endpoint model is to be regarded as a reasonable description of how expectations are formed if both α_i and β_i approach zero as i increases. This is exactly the pattern that emerges in table 1. The estimates decrease with i in all cases and often lose statistical significance for $i \geq 3$ onwards (note that the sample size for these estimations is constrained by the thirty-four observations available for long-term forecasts, so the results should be taken as indicative rather than conclusive).

3.2 Estimation Results

Next, we present the estimation results of the shifting-endpoint model (SE), and its variants (CE, UR, and LL). We use a diffuse prior (i.e., setting the initial state variance to a very large number) to initialize the Kalman filter. In all cases, the lag length of the autoregressive model for inflation is set to $p = 13$, which is the value that minimizes the Schwarz information criterion. Also, the null hypothesis that the approximation errors in (7) are not contemporaneously correlated was contrasted with a likelihood-ratio test, and categorically rejected in all instances. Hence, these covariances are estimated unrestrictedly.

For each model, table 2 presents the sum of the estimated autoregressive coefficients $\Phi(1)$ as a measure of persistence of inflation around the endpoint, which is restricted to $\Phi(1) = 1$ in the UR model; the estimated standard deviation of the inflation shock, $\text{std}(\epsilon_t)$; the estimated standard deviation of the endpoint shock,

Table 1. Mean Reversion of Forecasts

	$F_1 - F_6$	$F_2 - F_6$	$F_3 - F_6$	$F_4 - F_6$	$F_5 - F_6$
Chile	α β	0.181 (0.070)* 0.116 (0.035)*	0.083 (0.047)* 0.026 (0.027)	0.055 (0.039) 0.007 (0.021)	0.013 (0.023) 0.004 (0.016)
Colombia	α β	0.451 (0.194)* 0.413 (0.110)*	0.276 (0.180) 0.262 (0.128)*	0.166 (0.163) 0.159 (0.112)	0.036 (0.096) 0.066 (0.048)
Mexico	α β	0.101 (0.113) 0.732 (0.058)*	0.028 (0.054) 0.204 (0.020)*	0.019 (0.037) 0.096 (0.016)*	0.016 (0.034) 0.037 (0.017)*
Peru	α β	0.422 (0.156)* 0.544 (0.111)*	0.201 (0.099)* 0.059 (0.062)	0.077 (0.053) 0.045 (0.032)	0.015 (0.046) 0.015 (0.032)

Notes: The table shows least-squares estimates of equation (12), using the thirty-four available observations for long-term forecasts. HAC standard errors are shown in parentheses. * denotes coefficients different from zero at a 5 percent significance level.

Table 2. Estimation Results

	Model Specification				std(ε_t)	RMSE Inflation	RMSE Survey Forecasts			
	$\Phi(1)$		std(ε_t)	std($\Delta\mu_t$)			1	2	3	≥ 4
	CE	UR	LL	SE						
Chile	CE	0.931 (0.022)	0.379	0.000	0.473	0.38	0.88	0.93	0.36	0.24
	UR	1.000 (0.000)	0.387	0.000	1.417	0.39	1.07	1.37	1.62	1.64
	LL	0.544 (0.043)	0.311	0.286	1.038	0.46	0.64	1.06	1.17	1.24
	SE	0.677 (0.040)	0.405	0.175	0.844	0.43	0.57	0.78	0.56	0.54
Colombia	CE	0.979 (0.011)	0.309	0.000	1.322	0.33	0.99	0.98	1.50	1.62
	UR	1.000 (0.000)	0.333	0.000	1.570	0.34	1.07	1.03	1.44	1.69
	LL	0.480 (0.054)	0.242	0.257	1.122	0.43	0.41	0.83	1.11	1.45
	SE	0.713 (0.051)	0.324	0.193	0.925	0.36	0.53	0.62	0.77	1.06
Mexico	CE	0.981 (0.007)	0.278	0.000	0.798	0.28	0.60	0.72	0.60	0.80
	UR	1.000 (0.000)	0.292	0.000	1.750	0.29	0.82	1.29	1.80	2.21
	LL	0.564 (0.051)	0.219	0.241	1.263	0.34	0.58	1.01	1.68	2.08
	SE	0.790 (0.031)	0.291	0.161	0.987	0.30	0.51	0.66	0.76	1.01
Peru	CE	0.949 (0.015)	0.334	0.000	0.745	0.34	0.94	0.91	0.76	0.59
	UR	1.000 (0.000)	0.346	0.000	1.037	0.35	0.97	0.96	1.26	1.25
	LL	0.511 (0.041)	0.268	0.234	0.819	0.41	0.50	0.82	1.05	1.11
	SE	0.693 (0.038)	0.369	0.143	0.606	0.39	0.50	0.60	0.78	0.74

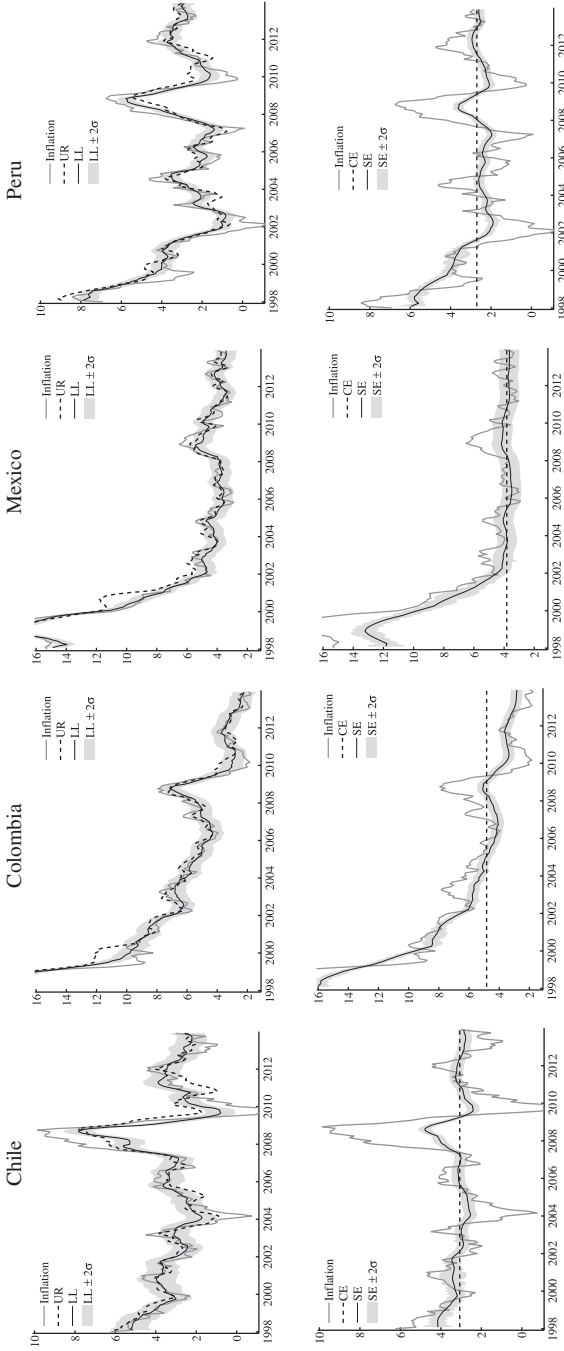
Notes: The table shows maximum-likelihood estimates, using data from February 1997 to December 2013. All models use $p = 13$. CE: constant-endpoint model; UR: unit-root model; LL: local-level model; SE: shifting-endpoint model. $\Phi(1)$: sum of the autoregressive coefficients, $\Phi(1) = \sum_{i=1}^p \phi_i$, (robust standard errors in parentheses) restricted to $\Phi(1) = 1$ in the UR model; $\text{std}(\varepsilon_t)$: standard deviation of the inflation shock; $\text{std}(\Delta\mu_t)$: standard deviation of the endpoint shock, restricted to $\text{std}(\Delta\mu_t) = 0$ in the CE and UR models; $\text{std}(\varepsilon_t)$: average standard deviation of the approximation errors of the $m = 6$ survey measures; “RMSE” is the root mean square error of the one-step-ahead predictions produced by the Kalman filter for inflation and survey forecasts for years 1 (current), 2 (next), 3, and ≥ 4 (subsequent) (the column labeled “ ≥ 4 ” shows the average RMSE for the remaining three long-term forecasts).

$\text{std}(\Delta\mu_t)$, which is restricted to $\text{std}(\Delta\mu_t) = 0$ in the CE and UR models; and, to save space, the *average* standard deviation of the approximation errors of the $m = 6$ survey measures, $\text{std}(\varepsilon_t)$. In addition, as measures of fit, the table presents the root mean square error (RMSE) of the one-step-ahead predictions produced by the Kalman filter for the observed data, i.e., inflation and survey forecasts. In the case of the LL model, we proceed in two steps: first, the inflation parameters were estimated using only the first measurement equation; second, the remaining output in the table was obtained *conditional on the first-stage estimates*. Figure 4 shows the predicted endpoints, evaluated at the maximum-likelihood estimates.

There are various similarities in the estimations across countries. In all cases, the UR model tracks actual inflation quite closely, but it performs rather poorly when it comes to survey data. The RMSEs for inflation are amongst the lowest across models (ranging from 0.29 in Mexico to 0.39 in Chile), whereas the RMSEs for the survey forecasts are the highest (ranging, for long-term forecasts, from 1.64 in Chile to 2.21 in Mexico).

This finding is related to the mean-reverting properties of expectations discussed in table 1 which, by construction, the UR model completely overlooks. In this respect, the LL model constitutes an improvement upon the UR model. Figure 4 reveals that the predicted endpoint in the LL model can be regarded as a de-noised and smoother version of the moving-average endpoint of the UR model; inflation expectations revert quickly from the observed data to such an estimated trend (the sum of autoregressive coefficients is $\Phi(1) \simeq 0.5$ in all cases), which improves the model's performance to account for the variability of short-term forecasts. Compared with the UR model, the RMSE of the LL model for the current-year inflation forecast (i.e., year 1) reduces sizably from 1.07 to 0.64 in the case of Chile, from 1.07 to 0.41 in Colombia, from 0.82 to 0.58 in Mexico, and from 0.97 to 0.50 in Peru. Long-term forecasts, nonetheless, are also poorly predicted in the LL model, as survey data appear to be much less sensitive to inflation news than what is implied in the LL model. Thus, for long-term forecasts, the reductions in RMSE compared with the UR model are quite modest: from 1.64 to 1.24 in the case of Chile, from 1.69 to 1.45 in Colombia, from 2.21 to 2.08 in Mexico, and from 1.25 to 1.11 in Peru.

Figure 4. Predicted Endpoints by Model Variant



Note: CE: constant-endpoint model; UR: unit-root model; LL: local-level model; SE: shifting-endpoint model. For the CE, the figure displays the estimated value of μ . For the UR, the moving-average endpoint, $\mu_t = e_1' C' z_t$ (see section 2.6). Only point estimates are reported for these models in order to avoid overloading the graphs. For the LL and SE models, the figure shows the smoothed predictions along with a 95 percent confidence interval. The scale of the vertical axis is the same for Chile/Peru and Colombia/Mexico.

On the other hand, in the CE model the sum of autoregressive coefficients $\Phi(1)$ is significantly higher than in the LL and SE models (where this quantity is estimated unrestrictedly). In the latter specifications, some of the persistence in observed inflation is attributed to the dynamics of a time-varying mean, and some to short-run deviations from this mean. Since inflation in the CE model is assumed to revert to an imposed constant level, the process is unsurprisingly estimated as highly persistent. Thus, short-term forecasts (especially for the current year) are similar to those of the UR model: the RMSEs for the current-year forecast range from 0.60 to 0.99 in the CE model, as compared with 0.82 to 1.07 in the UR model. However, with the exception of Colombia, the performance of the CE in fitting survey data improves rapidly and dramatically for long-term forecasts. A unit increase in i (the year index) implies an increase of $h_{i+1,t} - h_{it}$ = twelve months in the associated forecast horizon, so despite the large values of $\Phi(1)$, current inflation becomes less influential even for small values of i . As a result, the CE estimates turn out to be closer to the sample average of long-term forecasts than to the sample average of observed inflation (see figure 4), and the CE model does remarkably well in explaining long-term forecasts, three years ahead and beyond, in Chile (with an RMSE of 0.36 for the forecast three years ahead and of 0.24 for the “ ≥ 4 ” forecast), Mexico (0.60 and 0.80), and Peru (0.76 and 0.59). For these countries, inflation and expectations have been fluctuating (albeit, persistently) within a relatively narrow range for most of the sample period. In contrast, inflation trends downwards during all of the sample period in the Colombian case, making the CE model unsuitable.

The SE model is a compromise between the LL and the CE models. Survey information produces a smoother endpoint than the LL model, as shown in figure 4, which is manifested in a reduction of the standard deviation of the endpoint shocks from $\text{std}(\Delta\mu_t) \simeq 0.25$ to $\text{std}(\Delta\mu_t) \simeq 0.17$ and, more importantly, to significant increases in the noise-to-signal ratios $q = \text{var}(\epsilon_t)/\text{var}(\Delta\mu_t)$ from $q \simeq 1$ to $q \in [2.8, 6.7]$. Given the actual inflation persistence, a smoother endpoint is traded with an increase in $\Phi(1)$, from $\Phi(1) \simeq 0.5$ to $\Phi(1) \simeq 0.7$. With this, the SE model accounts for short-term forecasts’ variability as much as the LL model, while clearly outperforming the LL model to explain long-term forecasts. When compared with the CE model, the SE model performs considerably better when

it comes to short-term forecasts (years 1 and 2 in the table) but, except in the Colombian case, it is outperformed when predicting long-term forecasts. However, the improvement in the short-term fit of the SE model seems to more than compensate for the moderate deterioration in fitting longer-term expectations. As we shall discuss below, the improvement brought by the SE model is on the forecasting horizons that can be regarded as relevant for monetary policy analysis.

3.3 *Signal Extraction*

It has been argued that the SE model is able to disentangle short-run inflation shocks from long-run expectation shocks. Furthermore, it was argued that the SE would interpolate a complete profile of expectations from sparsely sampled survey information. Next, we present details on the computations involved in the prediction of the endpoint in the SE model.

From the Kalman filter, the endpoint is determined following the recursion $\mu_t = \mu_{t-1} + K_t v_t$, where K_t is the Kalman gain and v_t is the vector of prediction errors of the observed data. Thus, K_t measures how these errors are weighted to arrive at an updated inference about the unobserved state μ_t . To get a better understanding of how the endpoint is finally determined in our application, table 3 presents the implied Kalman gains in the SE model for selected dates, and for the case of Peru (an identical pattern is found for the remaining countries).

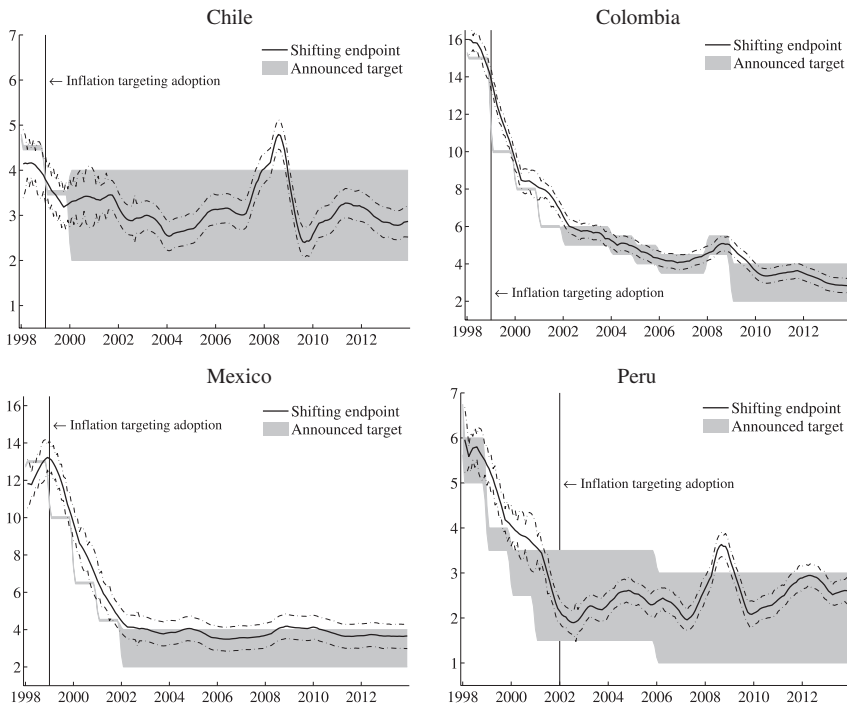
By construction, the gain is exactly zero for missing data (all survey data in March 2001 and long-term forecasts in February 2001, September 2006, and March 2013), which follows from the adjustments described in section 2.4. On the other hand, upon comparing the results from February to April 2001, September to October 2006, or March to April 2013, it can be seen that whenever long-term forecasts are available, the Kalman filter gives more weight to them, and less to actual inflation, to predict the endpoint. This is exactly what one would want: long-horizon expectations should provide more information about the endpoint and, thus, should receive more weight in the filtering process. When only short-term forecasts are available, the gains for F_1 and F_2 are of comparable magnitude and always higher than that of inflation. When long-term forecasts

Table 3. Kalman Gains in the Peruvian SE Model

	Inflation	F_1	F_2	F_3	F_4	F_5	F_6
February 2001	0.1036	0.1767	0.1661	0	0	0	0
March 2001	0.1187	0	0	0	0	0	0
April 2001	0.0654	0.1410	0.0720	0.0167	0.0364	0.0669	0.1364
September 2006	0.0640	0.0985	0.0947	0	0	0	0
October 2006	0.0528	0.0992	0.0550	0.0160	0.0220	0.0568	0.1034
March 2013	0.0621	0.1047	0.0993	0	0	0	0
April 2013	0.0492	0.1062	0.0542	0.0126	0.0274	0.0504	0.1028

Notes: The Kalman gain measures how prediction errors of the observed data are weighted in order to get an updated prediction of μ_t . The table shows these weights for selected dates. The gain is set to zero for missing values. Survey data are not available for March 2001, whereas long-term forecasts are only available in April and October.

Figure 5. Predicted Endpoints and the Announcement of Inflation Targets



Note: The scale of the vertical axis is the same for Chile/Peru and Colombia/Mexico.

become available, the gain of F_2 halves while the gain of F_6 , the furthest forecast, increases to a level close to the gain of F_1 . Interestingly, since the prediction errors are typically larger for F_6 than for F_1 (see the RMSE results in table 2), the absolute magnitude of the correction would also be larger for F_6 than for F_1 . Thus, when F_6 is observed, the Kalman filter updates the endpoint prediction towards it.

3.4 Inflation Targeting and Disinflation

Figure 5 shows the evolution of the predicted endpoints along with the inflation targets for each country *shifted forward* one year. For

instance, Chile's inflation target was 3.5 percent for 2000, and this value appears during 1999 in the figure.⁸ The purpose is to illustrate the joint evolution of the endpoint and the *announcements* of the inflation targets, and thus we are assuming that the central banks have announced their targets one year in advance. This was the actual practice during disinflation; however, all central banks have progressively increased the horizon for the targets as inflation has reached long-run levels.

During disinflation, roughly until 2002, the announcements have clearly served as a benchmark for private economic agents' long-run forecasts. Shocks to the endpoints bear an almost one-to-one correspondence to changes in the announced targets.

From 2002 onwards, long-run expectations have lain within the target ranges in almost all instances. The remarkable exception is year 2008, especially for Chile and Peru. During 2007 and most of 2008, these economies were hit by a sequence of large food price and oil shocks that deviated inflation from its target. The shocks were persistent enough that they affected long-run perceptions. However, the endpoint returned rapidly to the inflation targets by late 2008 as the result of both aggressive increases in the monetary policy rates and the start of the global financial crisis. It has remained within the target range ever since.

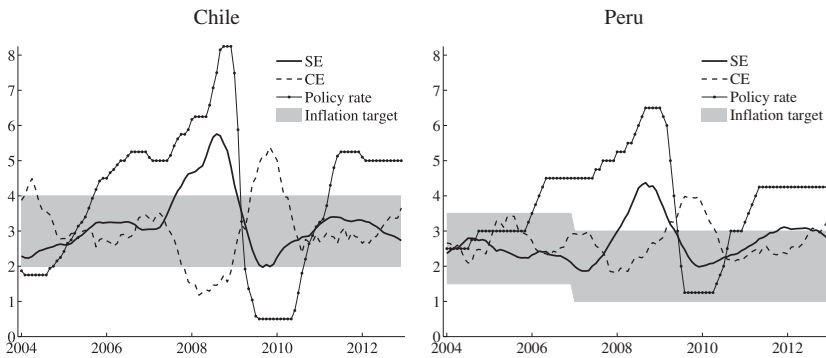
Given these dynamics, our results support the conclusion that the countries in our sample are good examples of successful inflation-targeting experiences (see Corbo and Schmidt-Hebbel 2001). Their central banks have managed to establish a credible regime of stable inflation with anchored expectations.

3.5 The Shifting-Endpoint Model as an Early-Warning Device

The SE model seems useful to better understand how inflation expectations evolve over time by exploiting available survey information. In this section we explore the real-time properties of this model and how it summarizes relevant information for the day-to-day conduct of monetary policy. To this end, we focus on the cases of

⁸Chile announced point targets until 2000; thereafter, the inflation targets have been published as ranges (actually, central values surrounded by a symmetric tolerance level). The same occurred in Colombia and Mexico until 2002. Peru has announced target ranges during all of the sample period.

Figure 6. Eighteen-Month-Ahead Inflation Expectations and Monetary Policy Interest Rates



Chile and Peru, with emphasis on the period from 2007 to 2009. Even though these two cases are generally regarded as successful inflation-targeting experiences, our previous findings indicate that expectations were temporarily unanchored and moved away from their target ranges by late 2007/early 2008 (cf. figure 5). As mentioned, before the international financial crisis in late 2008, both economies experienced, on one hand, a series of persistent food and fuel price shocks and, on the other hand, a sequence of positive demand shocks (a byproduct of extremely favorable terms of trade) that eventually passed through inflation and its expectations (see Cespedes, Chang, and Velasco 2014 for a general review).

Based on the (plausible) assumption that the relevant monetary policy forecasting horizon lies between one and two years, figure 6 presents the evolution of the eighteen-month-ahead inflation expectations implied by the SE and CE models, computed as in equation (11).⁹ From the estimation results in table 2, these are the two models that fare best in accounting for the behavior of long-run expectations. Besides, on a priori grounds, the CE model seems to

⁹Here, we seek to mimic the real-time inference that could be made by a researcher whose information set ends in period t . Thus, unlike the predictions presented in figures 4 or 5 that correspond to *smoothed* predictions, i.e., made with all information in our sample, the interpolated expectations in figure 6 are based on *filtered* or *updated* predictions on the endpoint, i.e., made with information up to period t .

be an adequate alternative to the SE model for the cases and samples selected in this analysis, where inflation has been fluctuating around the target ranges.

Yet, it is quite apparent that the SE predictions lead the CE predictions with a considerable lag. Recall that whereas the endpoint is able to respond to expectation news in the SE model, it is held constant in the CE model. Hence, before the financial crisis, medium-term inflation expectations seemed to have increased substantially, a fact recorded by the surveys and thereby reflected in the dynamics of the endpoint predicted by the SE model and its expectations at various horizons. On the contrary, the expectations of the CE model reflect the transition of current inflation to a given steady-state level, which is determined by very long-run forecasts, and so for the relevant horizons for monetary policy analysis the CE model expectations resemble actual inflation too much. Similar dynamics are recorded after late 2008, where the endpoint and expectations rapidly return to the target ranges in the SE model, and more sluggishly in the CE model. This differentiated behavior can also be observed in table 2, where the RMSE of the CE model is substantially higher than that of the SE model for short-term forecasts (year 1: respectively, 0.88 versus 0.57 in the case of Chile, and 0.94 versus 0.50 for Peru) and medium-term forecasts (year 2: respectively, 0.98 versus 0.78 for Chile, and 0.91 versus 0.60 for Peru).

Figure 6 also shows the path of the monetary policy interest rates. Being inflation targeters, it is reasonable to suppose that the movements of the policy rate reflect the perceptions about future inflation of the central bank, using all available information. In both countries, the strong correlation (or at least the directional co-movements) between the actual policy rates and the expectations inferred from the SE model is quite noticeable. There are many ways to rationalize such co-movements. For instance, private expectations may be responding to a high (low) inflation forecast made by the central bank; such a forecast may be driving an increase (decrease) in the interest rate and, as long as its effects are recorded in the surveys, may also be signaling an increase (decrease) in the inferred endpoint. On the other hand, the central bank may be directly responding to private expectations, especially when it suspects that they may have been affected by long-lived shocks. In either way, this exercise illustrates the extent to which the SE model is capable of summarizing,

in a satisfactory way, real-time information regarded as relevant for monetary policy decisions. This information, of course, is captured by the FEFs of inflation.

4. Closing Remarks

Fixed-event forecasts provide a widespread, yet unexplored, source of inflation expectations in many countries. The main difficulty is that the very structure of the FEFs, especially the fact that they correspond to moving forecast horizons, hinders their direct applicability in empirical work. To overcome this hindrance, and to infer about the term structure of inflation expectations from FEFs, we have proposed an extended version of the shifting-endpoint model of Kozicki and Tinsley (2012). Even though the resulting model is time varying, it can be easily handled with the Kalman filter, even for irregularly sampled survey expectations. By fitting the shifting-endpoint model to Latin American data, we conclude that it is able to jointly account for the behavior of actual inflation and survey records, often outperforming competing specifications. Our empirical exploration also suggests that survey FEFs provide a valuable source of information on expected inflation, complementary to that contained in historical records of inflation. Finally, the extended shifting-endpoint model may be a useful tool to track the behavior of inflation and its expectations, as reflected in the FEF records, in real time.

Given the availability of FEFs, exploring alternative methods to readily and effectively use such data in econometric models is likely to have important practical implications. Moreover, it would be of interest to compare the performance of FEFs versus REFs, or to assess the informational content of one type of forecast relative to the other, for countries where both are available. We hope our analysis proves to be a meaningful contribution to this promising research agenda.

References

- Ang, A., G. Bekaert, and M. Wei. 2007. "Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better?" *Journal of Monetary Economics* 54 (4): 1163–1212.

- Bakhshi, H., G. Kapetanios, and T. Yates. 2005. "Rational Expectations and Fixed-Event Forecasts: An Application to UK Inflation." *Empirical Economics* 30 (3): 539–53.
- Cespedes, L. F., R. Chang, and A. Velasco. 2014. "Is Inflation Targeting Still on Target? The Recent Experience of Latin America." *International Finance* 17 (2): 185–208.
- Chernov, M., and P. Mueller. 2012. "The Term Structure of Inflation Expectations." *Journal of Financial Economics* 106 (2): 367–94.
- Clements, M. P. 1997. "Evaluating the Rationality of Fixed-Event Forecasts." *Journal of Forecasting* 16 (4): 225–39.
- Corbo, V., and K. Schmidt-Hebbel. 2001. "Inflation Targeting in Latin America." Working Paper No. 105, Central Bank of Chile.
- Davies, A., and K. Lahiri. 1999. "Re-examining the Rational Expectations Hypothesis Using Panel Data on Multi-period Forecasts." In *Analysis of Panels and Limited Dependent Variable Models*, ed. C. Hsiao, K. Lahiri, L. F. Lee, and M. H. Pesaran, 226–54 (chapter 9). Cambridge University Press.
- Harvey, A. C. 1989. *Forecasting, Structural Time Series and the Kalman Filter*. Cambridge University Press.
- Kozicki, S., and P. A. Tinsley. 1998. "Moving Endpoints and the Internal Consistency of Agents' Ex Ante Forecasts." *Computational Economics* 11 (1–2): 21–40.
- . 2001. "Shifting Endpoints in the Term Structure of Interest Rates." *Journal of Monetary Economics* 47 (3): 613–52.
- . 2012. "Effective Use of Survey Information in Estimating the Evolution of Expected Inflation." *Journal of Money, Credit and Banking* 44 (1): 145–69.
- Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52 (1): 124–88.
- Mehra, Y. P., and C. Herrington. 2008. "On the Sources of Movements in Inflation Expectations: A Few Insights from a VAR Model." *Economic Quarterly* (Federal Reserve Bank of Richmond) 94 (2, Spring): 121–46.
- Mishkin, F. S., and M. A. Savastano. 2001. "Monetary Policy Strategies for Latin America." *Journal of Development Economics* 66 (2): 415–44.
- Nordhaus, W. D. 1987. "Forecasting Efficiency: Concepts and Applications." *Review of Economics and Statistics* 69 (4): 667–74.

- Pesaran, M. H., and M. Weale. 2006. "Survey Expectations." In *Handbook of Economic Forecasting*, Vol. 1, ed. G. Elliot, C. W. J. Granger, and A. Timmermann, 715–76 (chapter 14). Elsevier.
- Quispe-Agnoli, M. 2001. "Monetary Policy Alternatives for Latin America." *Economic Review* (Federal Reserve Bank of Atlanta) 86 (3): 43–53.
- Timmermann, A. 2007. "An Evaluation of the World Economic Outlook Forecasts." *IMF Staff Papers* 54 (1): 1–33.
- Vega, M., and D. Winkelried. 2005. "Inflation Targeting and Inflation Behavior: A Successful Story?" *International Journal of Central Banking* 1 (3): 153–75.