Forecasting Inflation Risks in Latin America:

A Technical Note

Rodrigo Mariscal
Andrew Powell

Inter-American Development Bank

Department of Research and Chief Economist

TECHNICAL NOTES

No. IDB-TN-403

June 2012

Forecasting Inflation Risks in Latin America:

A Technical Note

Rodrigo Mariscal
Andrew Powell

Inter-American Development Bank

2012

Abstract[1]

There are many sources of inflation forecasts for Latin America. The International Monetary Fund, Latin Focus, the Economist Intelligence Unit and other consulting companies all offer inflation forecasts. However, these sources do not provide any probability measures regarding the risk of inflation. In some cases, Central Banks offer forecast and probability analyses but typically their models are not fully transparent. This technical note attempts to develop a relatively homogeneous set of methodologies and employs them to estimate inflation forecasts, probability distributions for those forecasts and hence probability measures of high inflation. The methodologies are based on both parametric and non-parametric estimation. Results are given for five countries in the region that have inflation targeting regimes.

JEL classifications: C53, E37
Keywords: Inflation forecast, Inflation risk, Latin America

[1] The opinions expressed herein are those of the authors and not necessarily those of the IDB, its Board of Executive Directors or the countries they represent. Rodrigo Mariscal is a consultant at the Inter-American Development Bank, Research Department, 1300 New York Ave NW, Washington, DC, 20577. Email: rmariscal@iadb.org. Andrew Powell is Principal Advisor at the Inter-American Development Bank, Research Department, 1300 New York Ave NW, Washington, DC, 20577. Email: andrewp@iadb.org.

1 Introduction

There are many sources of inflation forecasts for Latin America. The IMF, Latin Focus, the Economist Intelligence Unit (EIU) and other consulting companies all offer inflation forecasts. In general, however, while these sources provide forecasts, they do not make available probability measures regarding the risk of higher inflation. In some cases, Central Banks offer forecast and probability analyses, but Central Bank models are not fully transparent and different Central Banks have different methodologies. In this technical note, we attempt to develop a relatively homogeneous set of methodologies and employ them to estimate inflation forecasts, probability distributions for those forecasts and hence probability measures of inflation. The methodologies are based on both parametric and non-parametric estimation, and we employ them on five inflation targeting countries in the region.

We estimate two types of uni-variable models to describe the inflation dynamics of Latin American countries (LAC). We use Auto-Regressive Moving Average models choosing appropriate lag structures (ARMA(p,q) models), and we also estimate versions of these models with a fractional integration parameter to capture higher persistence effects (namely ARFIMA(p,d,q)), both with a constant dummy and with seasonal dummies to correct for seasonality and extreme values that the models cannot capture.

Given the persistence that we find in the data and certain results discussed below, we take the ARFIMA(p,d,q) models to forecast inflation processes. The application of the ARFIMA models that we use are based on Doornik and Ooms (2004).[2] To cope with the model selection uncertainties, we estimate several ARMAs and ARFIMAs and select the best of those using defined criteria (see below).

We use the annualized monthly inflation for all LAC. To define the sample period we test for structural breaks in the series using the Bai-Perron Test.[3] As we find that the ARFIMA models outperform the ARMA models in terms of our selection criteria, we use the ARFIMA to forecast inflation.

In estimating the probability distribution for future inflation, the problem of what distributional assumption to employ is confronted head on. Typically in econometric forecasting, forecast estimates are assumed to be normally distributed. We consider this typical case, but instead we estimate the probability distribution of the forecasts using bootstrap methods. Then, to calculate the risk probabilities we use a kernel density estimation of the bootstrap results. We find that the risk

[2] Our estimations are computed with the ARFIMA package 1.04 for Ox developed by Doornik and Ooms (2006).
[3] Bai and Perron (1998, 2003).

probabilities give more information to analyze inflation forecasts as they capture and distinct cases of high/low mean and high/low variance. They can also be used to make country comparisons.

In this note, we focus on presenting the results largely in an intuitive and graphical format. A technical appendix gives further details regarding the methodologies and more detailed results.

Literature Review

Several institutions forecast and analyze inflation in LAC, including several Central Banks, and they publish their results in inflation reports and other publications. However, Central Banks in the region employ different models and do not always make the models they use fully transparent. Moreover, while Central Banks and some other institutions publish inflation forecasts, they do not all publish measures of probability regarding the likelihood of higher inflation levels.

The literature associated with the ARMA models and the Box-Jenkins methodology is standard in textbooks.[4] To correctly estimate an ARMA model model, the series must be stationary. An issue with macroeconomic variables is to establish the order of integration. To deal with this kind of problem, the ARFIMA type of models are very useful. These models are more general as they try to establish the degree of integration rather than assume exante whether a series is stationary or is, say, I(1). The fractional integration model was first introduced by Granger and Joyeux (1980) and Hosking (1981) as a class of model that can capture the properties of long-memory series.[5] Diebold and Rudebusch (1989) show useful applications to model macroeconomic series. A good introduction to the ARFIMA model and an application to model and forecast inflation can be found in Doornik and Ooms (2004) and Bos, Franses, and Ooms (2001).

To calculate forecast intervals one usually relies on the normality of the forecast errors. This approach could be justified when one has large samples and a well-known normal distribution, but even under such circumstances the forecast error may not be normal distributed.[6] Besides, it may be that the distribution of inflation is not symmetric around its mean. First introduced by Efron (1979), the bootstrap technique is used to estimate unknown distributions in small samples or when a sampling distribution is analytically intractable. Veall (1987) and McCullough (1994) do an application of this technique to estimate forecast intervals. For a more recent and complete review of bootstrap methods see Horowitz (2001), Horowitz (2003) and Häirdle, Horowitz, and Kreiss (2003).

[4] See, for example, Davidson and MacKinnon (1993).
[5] See also Beran (1994) and the literature review from Baillie (1996), Henry and Zaffaroni (2002) and Robinson (2003).
[6] In general, the forecast is a product of *asymptotic normal variables* instead of a linear combination of *normal variables;* see McCullough (1994).

2 Forecasting Models

To forecast inflation we use two kinds of univariable models: the ARFIMA(p,d,q), which stands for autoregressive fractional integrated moving average model, and the ARMA(p,q), which stands for autoregressive (AR) and moving average (MA) model. Here "p" is the maximum order of the AR component, "q" is the maximum order of the MA component and "d" is the order of integration of the series, denoted $I(d)$. Our interest in taking these two models is to compare which one better describes the inflation processes in LAC so we could use the superior model for forecasting. We estimate that for most of the countries inflation data is best described by the ARFIMA model. In this section we present the facts behind the estimations.

We take the longest series of the monthly consumer price index (CPI) that we could find for Brazil, Chile, Colombia, Mexico and Peru.[7] We calculate inflation as $\pi_t = 1200 \log(P_t/P_{t-1})$, where π_t and P_t are the rate of inflation and the monthly CPI in period t. To define the sample for each country we run a structural break test for the mean of inflation. This test, known as Bai-Perron(Bai and Perron, 1998, 2003), takes a number of partitions given exogenously and uses a special algorithm to minimize the squared residuals and estimate the means, the break dates and the confidence intervals for the break dates. The test gives the break dates endogenously.

This test allows us to identify the periods when the inflation settle around a more recent state and would let the models capture the dynamics of inflation in a more recent and stable period. Interestingly for all LAC, this recent period is the one with the lowest level of inflation on average, and for all it started in the early or late 1990s. The results of the test are presented in Table 1 (page 8). Our selected sample starts at the first observation after the last break date.

We estimate 4,096 ($= 2^6 \times 2^6$) ARFIMAs and 4,095 ARMAs per country. This amount account for all the possible combinations of the $AR(p)$ and $MA(q)$ parameters from zero and up to six. Unfortunately the process is time consuming, and our CPU memory was insufficient to take the combinations up to more than six.[8]

Our criteria for selecting the better model are, in order of importance, the following: the model must have statistically significant coefficients, it must pass all the standard test (normality, ARCH, Portmaneau), it must have the lowest information criteria (SC, HQ, AIC) and finally, it must have the lowest Root Mean Square Error (RMSE) in the pseudo out-of-sample forecast (described

[7] Our sources are for Brazil: Broad National Consumer Price Index (IPCA); for Chile: IPC General Base Promedio 2009; for Colombia: Indice de precios al consumidor (IPC) ; for Mexico: IPC Por objeto del gasto Nacional; and for Peru: Indice de Precios al Consumidor (Indice Base 2009=100).

[8] To give an example, for an ARFIMA, the number of equations to be estimated for all the possible combinations of parameters from zero to 12 are 16,777,216.

below). In general, we find that the ARFIMA models outperform the ARMA models or perform at least do as well as they do.

In Figure 1 (page 10), we plot the Schwarz (1978) and Rissanen (1978) criteria (SC) against our classification of models according to one of our selection standards. We find that using this criterion, the best ARFIMAs outperform the best ARMAs for Chile. For other countries like Brazil and Mexico, ARFIMAs do as well as ARMAs do, while for Colombia and Peru, the ARMAs do better than ARFIMAs.

We then measure the RMSE of the best ARFIMAs against the best ARMAs. We calculate the RMSE using a *dynamic pseudo out-of-sample forecast*[9] for four horizons and a rolling sample. First we estimate the selected model with one fourth of the whole sample, which we refer to as the initial sample (s_0). Then we forecast out of the sample for $h = \{3, 7, 14, 21\}$ steps ahead. Next we add to the initial sample one forward observation and forecast again out of the sample for $h = \{3, 7, 14, 21\}$ steps ahead. We stop until we cover the whole feasible sample, that is from $s = s_0, \ldots, S$.[10] We calculate the RMSE as follows:

$$\mathrm{RMSE}_h = \left[\frac{1}{S_h} \sum_{s=s_0}^{S_h} \left(\pi_{s+h|s} - \hat{f}_{s+h|s} \right)^2 \right]^{1/2} \tag{1}$$

We find that in general the ARFIMA models do better than the ARMAs and the AR(1)s. For Chile and Colombia, the ARFIMA does better in every horizon and even against an AR(1). Brazil's ARFIMA also does better than the ARMA but not much better than the AR(1), and actually for the 21 horizon it does worse. The ARFIMAs for Mexico and Peru do not do better in every case against the ARMA, but they do better for long horizons, or at least not worse than any ARMA or AR(1). The results of this exercise are presented in Figure 2 on page 11.

Finally, we use the likelihood ratio (LR) to test that none of the ARFIMAs is encompassed by its nested ARMA model;[11] we find that for all the ARFIMA models, this is not the case (see Table 2 on page 9).

All the previous analysis accounts for model uncertainties. Now to account for the uncertainties in the coefficients and in the disturbances, we use the observed sample errors to construct an "artificial" sample of the errors and the dependent variable. We re-estimate both the coefficients and forecast with this new dependent variable. This task is done several times taking the sample

[9] We mean this: $\hat{f}_{T+h+1|T} = \hat{\beta}\hat{f}_{T+h|T}$, for a given $\hat{\beta}$ estimated with a sample T.

[10] As we are forecasting out-of-sample our feasible sample will vary with the forecasting horizons and will be smaller than the whole sample.

[11] Since the $I(0)$ and $I(1)$ are particular cases in the ARFIMA model, one can perform LR tests for the null $d = 0$ or $d = 1$ against the alternative $d \in (0, 1)$.

errors randomly without replacement so that we end up with a whole range of coefficients and forecasts. This method is known as bootstrapping. Finally we use a kernel density estimation to produce the probability density of the forecasts. For a more technical explanation see the Section *Bootstrapping Forecast* in the Appendix (page 16). In Figure 3 (page 12) we plot the out-of-sample inflation forecasts for the five countries using the ARFIMA models with the bootstrap confidence interval.[12]

There are several aspects to note about these forecasts. First, the forecast estimate that the annual inflation for the next periods will be bigger than the last observed but lower than those in other periods. Second, there are countries that have larger confidence intervals than others, like Brazil and Chile. This fact reflects that the variance of inflation in the sample is larger for those countries. And finally, for some countries (like Brazil, Colombia and Mexico) the forecast density is less symmetric than other countries' density, which means that the probabilities of values above the mean are different from the values below the mean (see Figure 4, page 13).

3 Probability Measurement

We are basically interested in end-of-year inflation, so we present the probability measures for December of 2011 (one year ahead) and 2012 (two years ahead). As discussed above, our preferred model is that stemming from the ARFIMA methodology. The probabilities for the ARFIMA models are displayed for the five countries in Figure 5 (page 14) below. We use a Gaussian kernel density to estimate the probability of high inflation for the annual inflation rate.[13]

There are three things to note. First, the further to the right a curve is, the higher the forecast inflation level, and it also follows that the country may have a greater probability of experiencing high inflation. Second, the steeper the curve, the less the predicted variance, and we will be less likely to be "surprised" with high values of inflation. Finally, when the two-year probability curve is below the one-year probability curve, it means either that the inflation forecast is rising in the first year but settling in the second year, or that the forecast is converging to the unconditional mean more rapidly; this convergence will vary from model to model and will depend on the persistence of the models.

The estimations in Figure 5 (page 14) also allow us to make comparisons across countries. For example, comparing Brazil and Peru, it can be seen that Peruvian inflation *stochastically domi-*

[12] We transform the monthly forecast to annual forecast; the procedure for doing so is discussed in the Appendix (page 17).
[13] For more detail on the calculation, see the Section *Transformation and Kernel Density Estimation* in the Appendix (page 17).

nates[14] Brazil's (and all other countries' inflation) for the next two years. In other words, the entire distribution of Peruvian inflation is to the left of other countries' inflation distributions considering a two year horizon.

An other interesting case is Colombia's inflation versus Chile's inflation at the close of 2011. Both mean forecasts are roughly the same, but Colombia's probability function appears as less risky, in the sense that is less likely to have inflation rates greater than 5 percent. In this case we say that Colombia's inflation *second-order stochastically dominates* Chile's inflation.

The models also predict that inflation in Colombia and Brazil will be higher than the previous year, with more or less the same variance. Mexico's inflation will be higher than the previous year as well, but its variance will increase. On the other hand, Chile's and Peru's inflation will be less, though Chile's variance will increase.

4 Conclusions

We show a way to calculate inflation probabilities using kernel densities and forecasts from ARFIMA models. In our forecast we try to account for different uncertainties associated with the forecasting exercise, which include model selection, parameter stability and error sample distribution.

Note that these forecasts are unconditional and for that reason the conclusions should be interpreted as a short-run analysis. A long-run analysis should take into account other variables and institutional arrangements and policy actions that are influencing the path of inflation.

To have a probability of higher inflation does not mean that the national authorities or the Central Banks are failing in their functions; indeed, these probability distributions reflect shocks that by their nature are unanticipated and we do not therefore incorporate the effects of new policy actions that might be taken if shocks were to materialize. Rather, we see this analysis as a homogeneous way to consider the ex ante risks of higher inflation in the region at the current time.

[14] We say that a probability function $F(\pi)$ first-order stochastic dominates (is unambiguously better than) another probability function $G(\pi)$, if $F(\pi) \leq G(\pi)$ for all π. For more detail on this see Hadar and Russell (1969), Rothschild and Stiglitz (1970).

Table 1. Break Dates from Structural Change Test

		Brazil			
		(January 1948 - March 2011)			
Coefficients	23.84	63.23	224.49	18.41	6.10
Std.Err.	2.89	9.75	58.42	2.86	0.88
Break dates	Jul-1975	Jan-1985	Jun-1994	Jan-1997	
		Chile			
		(January 1947 - March 2011)			
Coefficients	25.33	86.67	17.36	4.16	
Std.Err.	2.17	13.08	1.47	0.44	
Break dates	Jun-1971	Jan-1981	Oct-1993		
		Colombia			
		(July 1954 - March 2011)			
Coefficients	9.71	21.28	5.82		
Std.Err.	1.44	1.32	1.22		
Break dates	Jan-1973	Jun-1998			
		Mexico			
		(January 1969 - March 2011)			
Coefficients	9.89	20.71	65.46	18.40	5.07
Std.Err.	1.95	2.43	7.39	3.85	0.55
Break dates	Sep-1975	Dec-1981	Mar-1988	Feb-1999	
		Peru			
		(February 1950 - March 2011)			
Coefficients	9.35	119.94	15.86	2.32	
Std.Err.	0.86	21.84	3.84	0.51	
Break dates	Aug-1982	Aug-1991	Oct-2000		

Notes:

Bai and Perron (1998, 2003) test for multiple structural breaks.

We use a constant as regressor, correct for serial correlation, a trimming of 0.15 and allow up to 5 breaks.

Table 2. Likelihood Ratio Test

Country	$\mathcal{L}_{ur}$	$\mathcal{L}_r$	$\chi^2(1)$	Prob.
Brazil	-427.79	-438.87	22.16	0.000
Chile	-584.31	-606.95	45.29	0.000
Colombia	-358.61	-367.01	16.78	0.000
Mexico	-315.23	-339.66	48.86	0.000
Peru	-303.93	-322.79	37.71	0.000

Notes:

$\mathcal{L}_{ur}$ is the likelihood of the unrestricted model: $I(d)$.

$\mathcal{L}_r$ is the likelihood of the restricted model: $I(0)$.

$\chi^2(1)$ is the Chi-square statistic with one degree of freedom.

Table 3. Resume Table

	Forecast[1]		$Pr(\pi) \geq 5\%$		$Pr(\pi) \geq 10\%$	
Country	Dic 2011	Dic 2012	Dic 2011	Dic 2012	Dic 2011	Dic 2012
Brazil	7.0	6.6	0.988	0.895	0.026	0.005
Chile	4.2	3.5	0.230	0.122	0.000	0.000
Colombia	4.0	5.3	0.066	0.632	0.000	0.000
Mexico	3.6	3.9	0.076	0.140	0.000	0.000
Peru	3.1	2.8	0.004	0.002	0.000	0.000

Notes:

Annual inflation forecasts and probabilities.

[1] Mean of the bootstrap distribution.

Figure 1. Model Selection Based on Information Criteria

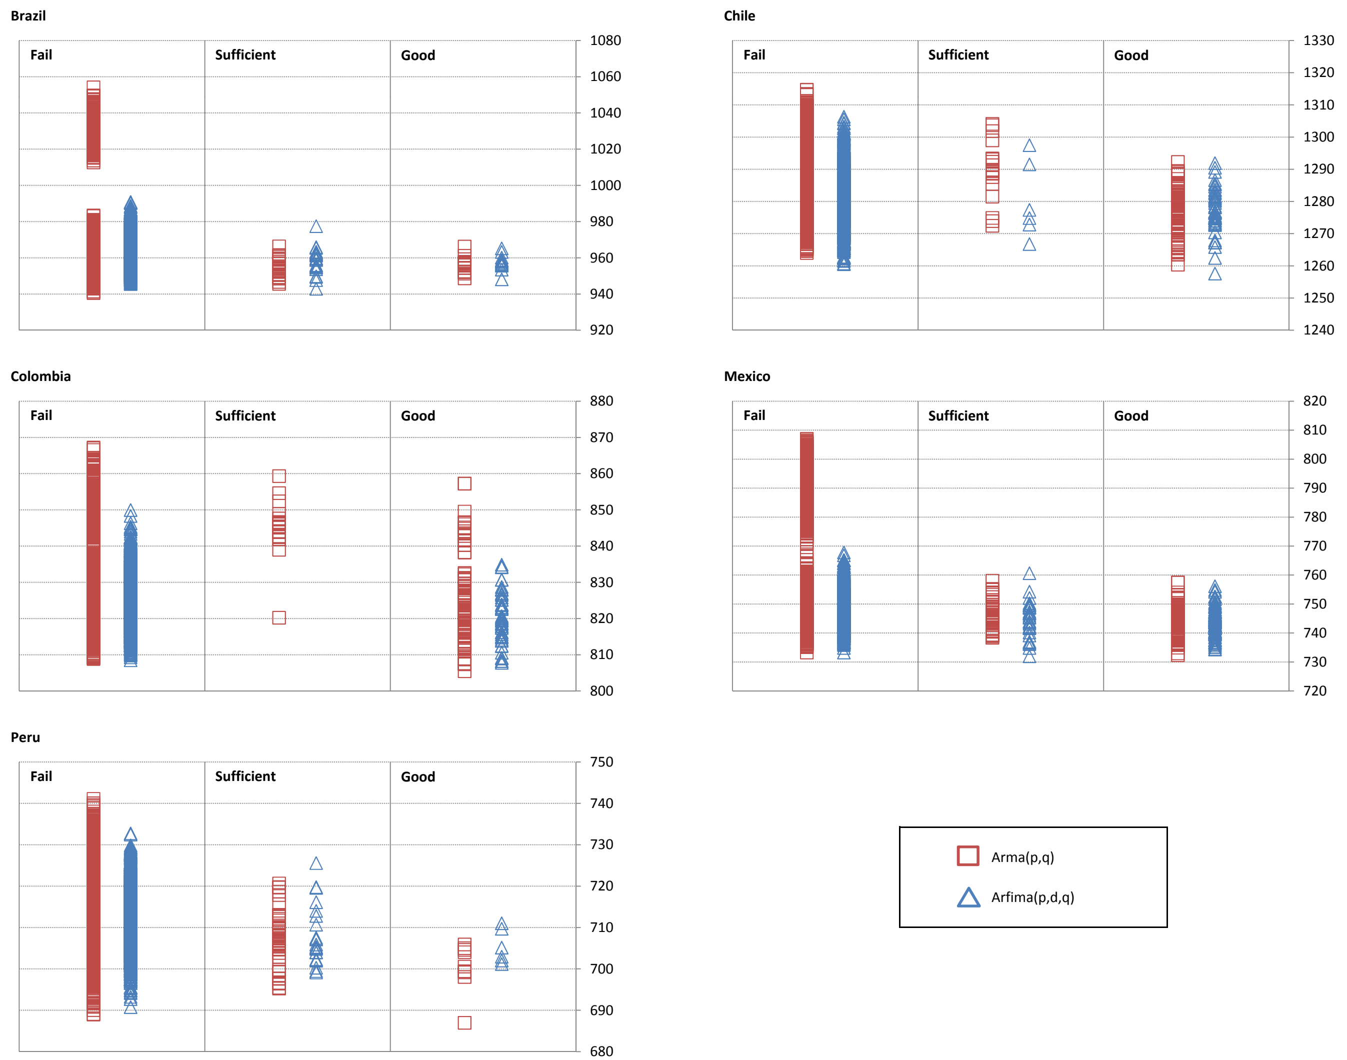

Note: Schwarz Criteria (SC) for the ARMA and the ARFIMA models for all the possible combinations of the $AR(p)$ and $MA(q)$ parameters from zero and up to six. A *Fail* model is one that doesn't have statistically significant coefficients or doesn't pass all the standard tests (normality, ARCH, Portmaneau) at 10% of significance. A *Sufficient* model is not a *Fail* model: have statistically significant coefficients and at least passes all the standard tests at 90% of significance. A *Good* model is a *Sufficient* model that at least passes all the standard tests at 95% of significance.

Figure 2. Model Selection Based on RMSE

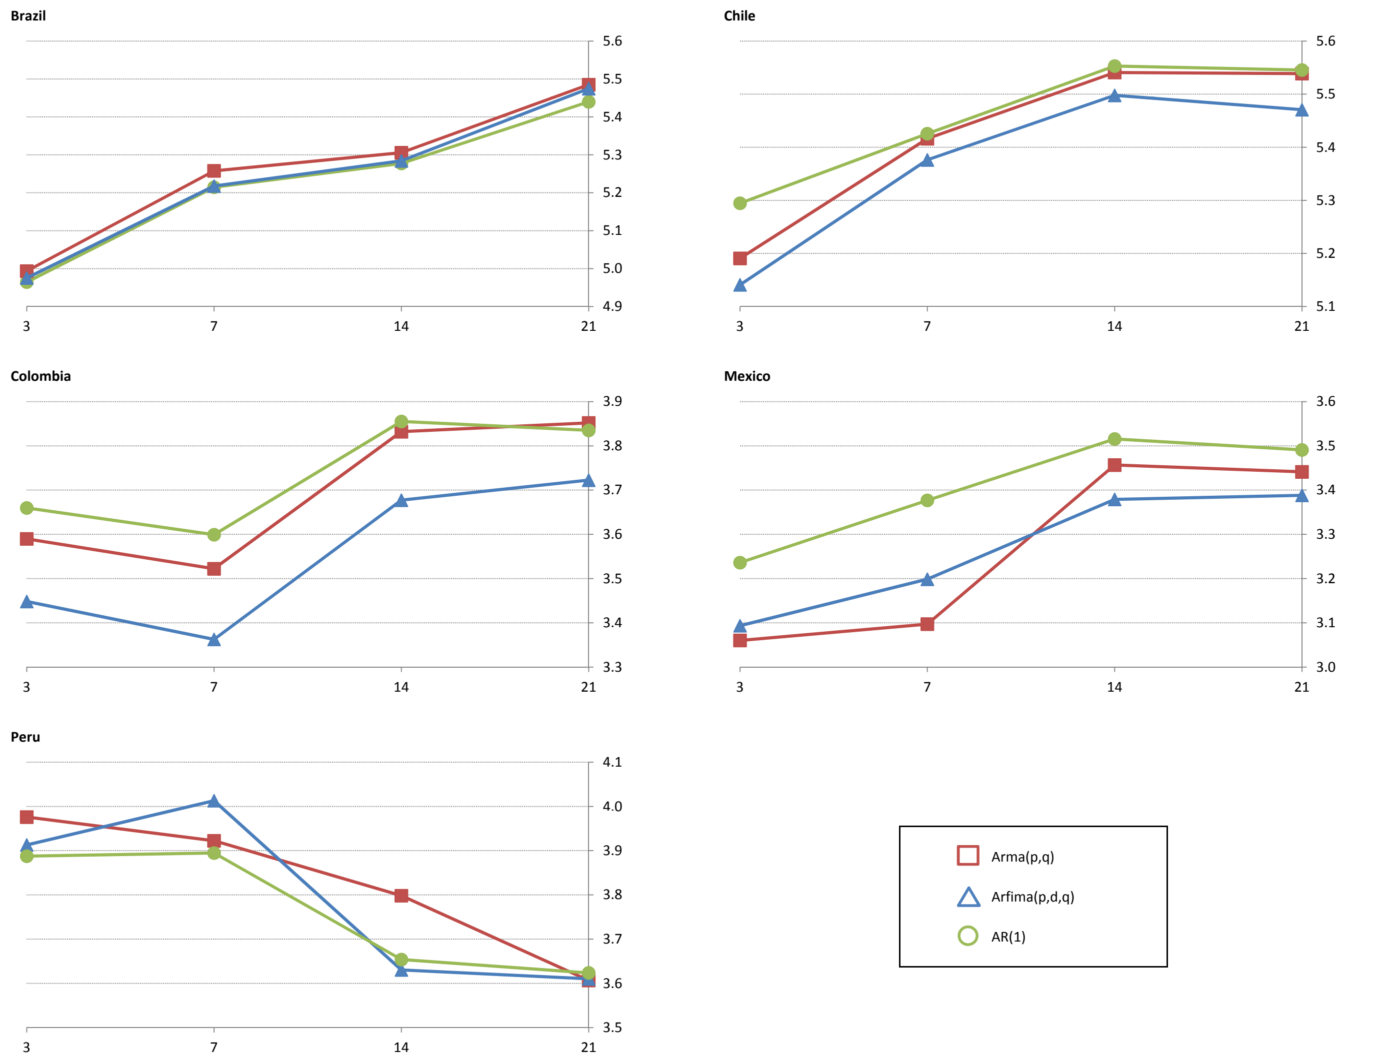

Note: RMSE of the 3, 7, 14 and 21 h step-ahead forecast starting with one fourth of the sample and rolling the forecast one step forward to the end of the sample.

Figure 3. Annual Inflation Forecast

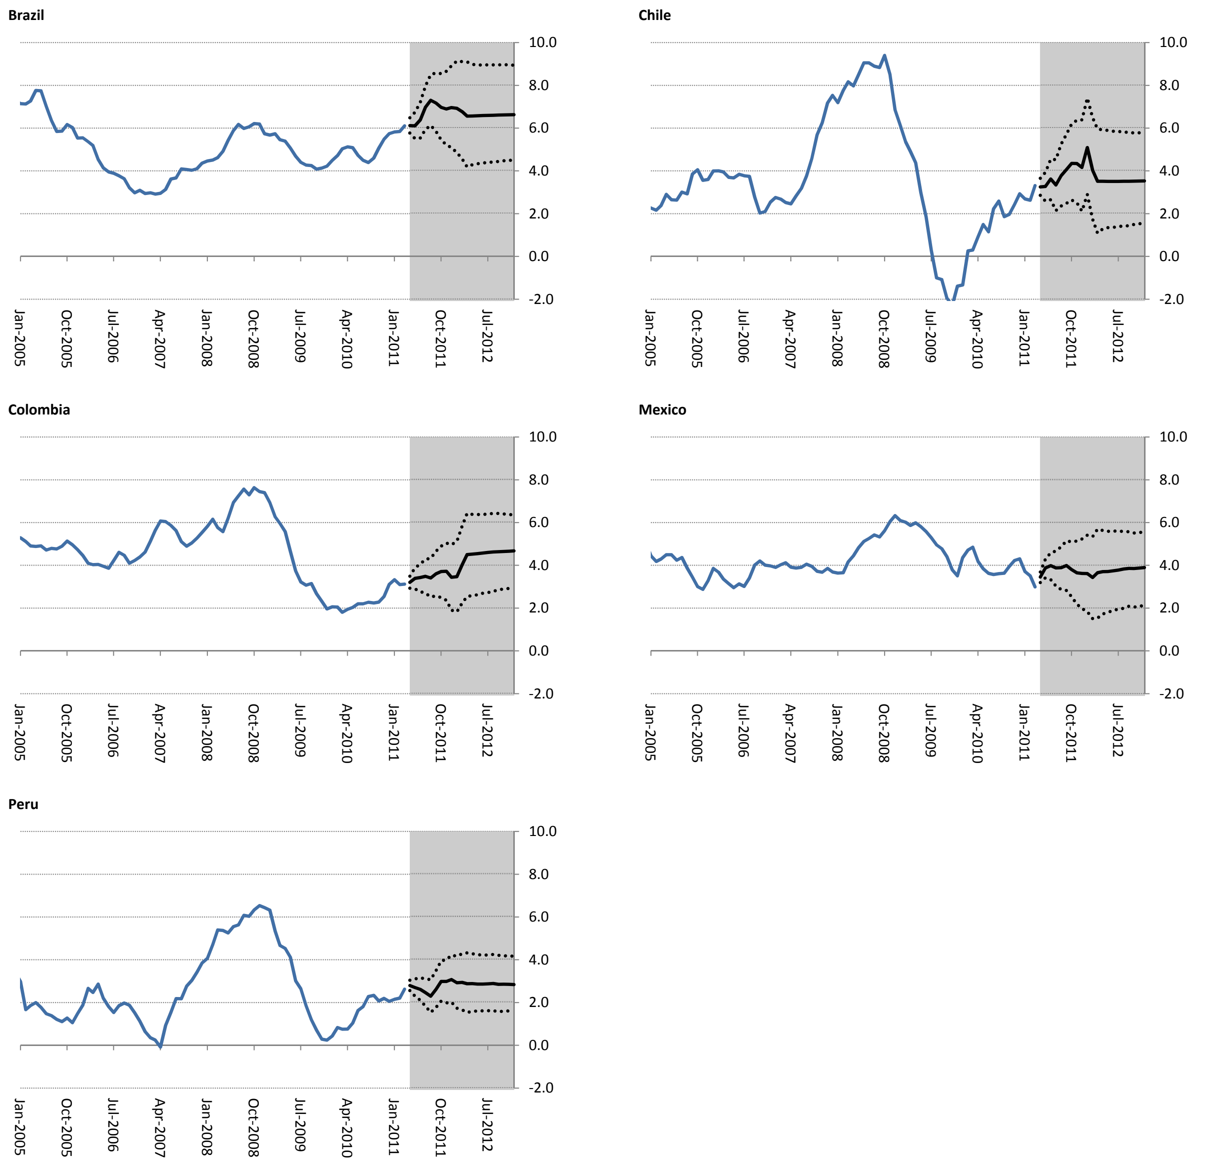

Note: Out-of-sample forecast with the ARFIMA models. The shaded area marks the beginning of the forecast (April 2011) and the dotted lines report the bootstrap 90% forecast interval.

Figure 4. Probability Distribution of Annual Inflation

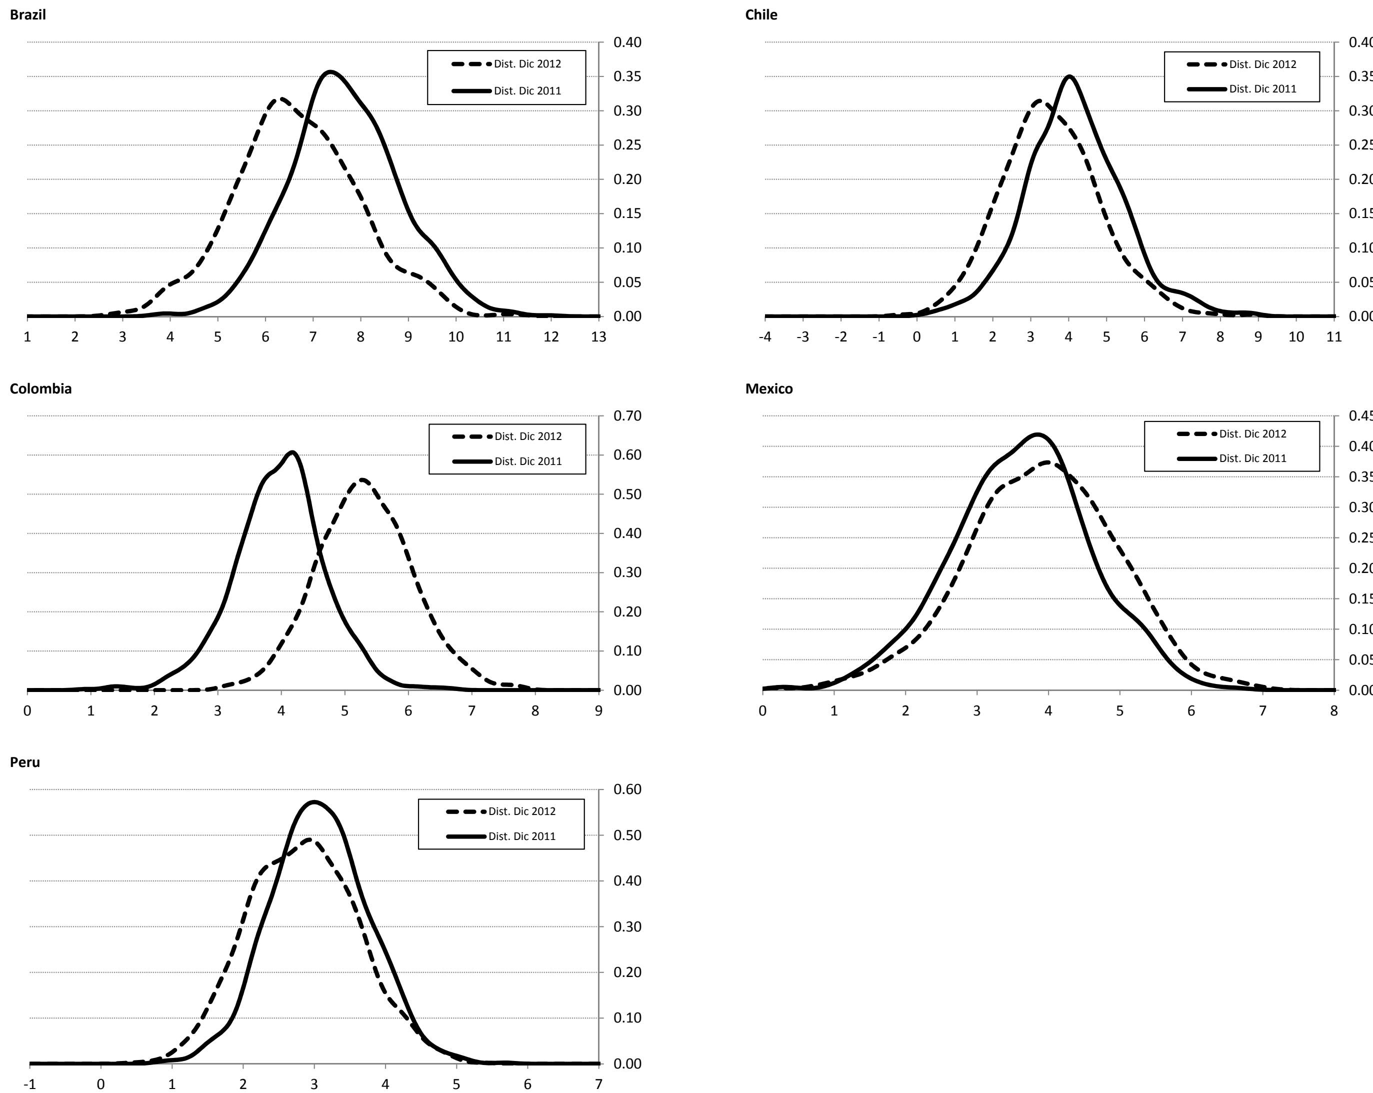

Note: Kernel density of the out-of-sample forecast with the ARFIMA models.

Figure 5. Probability of Inflation

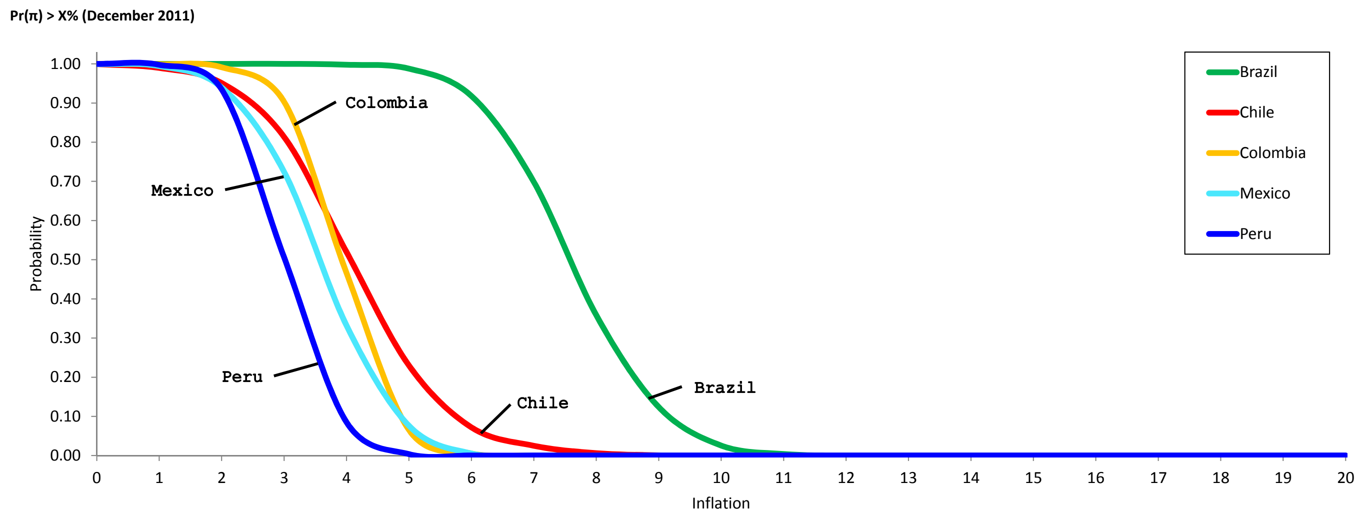

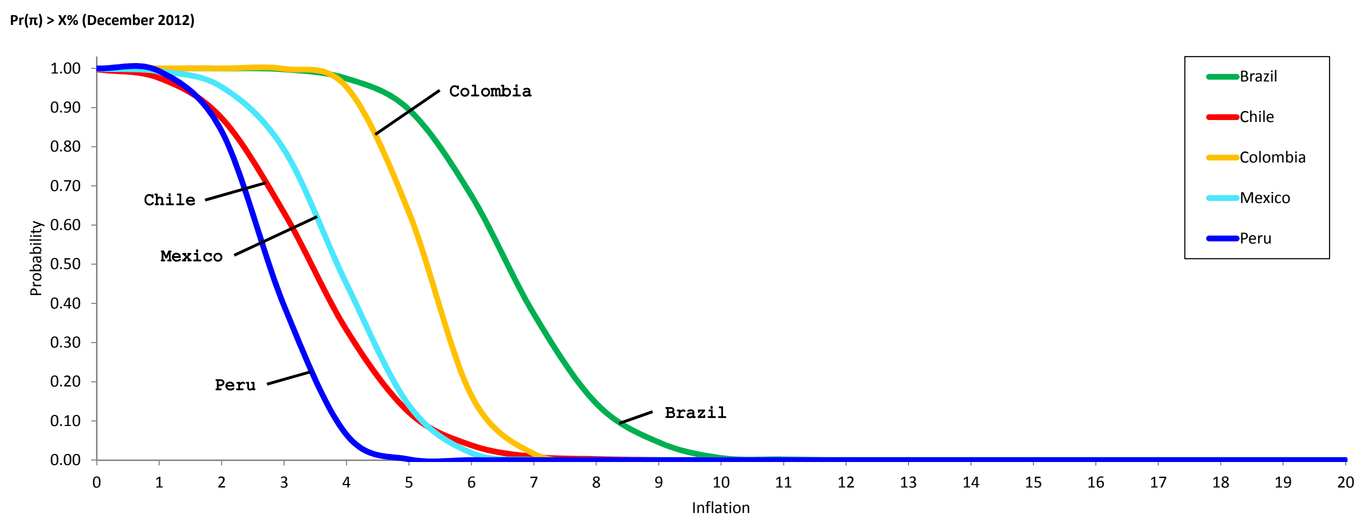

Note: Probability that annual inflation exceed a given value. The calculations are taken from the bootstrap forecast density of the ARFIMA models.

Appendix

Estimating and Forecasting an ARFIMA *Model*[15]

Let y_t be a random variable, then the $\text{ARFIMA}(p,d,q)$ model representation is:

$$\Phi(L)(1-L)^d(y_t-\mu_t)=\Theta(L)\varepsilon_t, \qquad t=1,\ldots,T. \tag{2}$$

Where $\Phi(L)=(1-\phi_1 L-\ldots-\phi_p L^p)$ is the AR polynomial, $\Theta(L)=(1+\theta_1 L+\ldots+\theta_q L^q)$ is the MA polynomial and L is the lag operator. The letters p and q just take integers values but d can take any real number. All roots of $\Phi(L)=0$ and $\Theta(L)=0$ lie outside the unit circle and don't have common roots between them, this condition state that the model is stationary. We assume that $\varepsilon_t \sim \text{NID}(0,\sigma_\varepsilon^2)$ and write μ_t for the mean of y_t and define $z_t=y_t-\mu_t$[16]. The fractional difference operator is defined by the following binomial expression:

$$(1-L)^d=\sum_{j=0}^{\infty}\binom{d}{j}(-L)^j=1-dL+\frac{d(d-1)}{2!}L^2-\frac{d(d-1)(d-2)}{3!}L^3+\ldots \tag{3}$$

When $-0.5<d<0.5$ the process is stationary and invertible. Then when $0<d<0.5$ the process has *long memory* and the autocovariance function has hyperbolic decay, ck^{2d-1} as $k\to\infty$, rather than an exponential decay. An other interest case is when $0.5\leq d<1$, in this case the process is non-stationary but still mean-reverting. When $d=1$ the process is called non-stationary, *infinite memory* or $I(1)$, and finally when $d=0$ the process is called stationary, *short memory* or $I(0)$. When the process is invertible, the MA and the AR representations of z_t are calculated as follows:

$$z_t=\sum_{j=0}^{\infty}a_j L^j\varepsilon_{t-j}=A(L)\varepsilon_t. \tag{4a}$$

$$z_t-\sum_{j=1}^{\infty}b_j L^j z_{t-j}=B(L)z_t=\varepsilon_t. \tag{4b}$$

Note that when $d<1$, then $a_j\to 0$ as $j\to\infty$, conventionally $a_0=1$ and $b_0=1$. We define $A(L)=\Phi(L)^{-1}(1-L)^{-d}\Theta(L)$ and $B(L)=\Theta(L)^{-1}\Phi(L)(1-L)^d$.

[15] For more detail of the procedure and the algorithms presented here see Doornik and Ooms (2006) and Doornik and Hendry (2009, Chapter 13).

[16] The treatment of the mean (μ_t) can be set in three different ways: fixing it at some predefined value, setting it as a deviation from the sample mean, or using no mean but include a constant in the regression.

To estimate the parameters, we use the Exact Maximum Likelihood (EML). An important thing to note is that the function of likelihood is maximized using Broyden?Fletcher?Goldfarb?Shanno (BFGS) with numerical derivatives and stationarity is imposed, which means that the procedure reject the parameter values ether when $\widehat{d} \leq -5$, when $\widehat{d} > 0.49999$ or when $|\widehat{\rho}_i| \geq 0.9999$, where $\widehat{\rho}_i$ are the roots of the AR polynomial[17].

The program uses the estimates from the EML to calculate the dynamic forecast (best linear prediction) of $z_{T+H|T}$ and $y_{T+H|T}$ using the following algorithms:

$$\widehat{z}_{T+H|T} = \sigma_\varepsilon^{-2}(\gamma_{T-1+H} \cdots \gamma_H)(\Sigma_T)^{-1}\mathbf{z} = \mathbf{q}'\mathbf{z}. \tag{5a}$$

$$\widehat{y}_{T+H|T} = \mathbf{q}'(\mathbf{y}-\mu) + \mu_{T+H|T} + \mathbf{x}'_{T+H|T}\widehat{\beta}. \tag{5b}$$

Where σ_ε^2 is the error variance, γ_i is the autocovariance $\mathrm{E}[(y_t-\mu)(y_{t-i}-\mu)]$, Σ_T is the autocovariance matrix, $\mathbf{z} = (z_1,\ldots,z_T)'$, $\mathbf{y} = (y_1,\ldots,y_T)'$, $\mu = (\mu_1,\ldots,\mu_T)'$, $\mathbf{x}$ is a matrix of exogenous variables, T is the actual sample size and H is the forecast horizon. This is a dynamic forecast which means that depending on the length of H, the right-hand variables may be predicted.

Bootstraping Forecasts

Let inflation be a random variable with an ARFIMA data generating process (DGP)[18]:

$$\pi_t = \sum_{j=1}^{T} b_j L^j \pi_{t-j} + \mathbf{x}_t'\beta + \varepsilon_t, \qquad t = 1,\ldots,T. \tag{6}$$

We assume that $\varepsilon_t \sim \mathrm{NID}(0,\sigma_\varepsilon^2)$ and define $B(L) = \Theta(L)^{-1}\Phi(L)(1-L)^d = \sum_{j=0}^{\infty} b_j L^j$. Estimating this equation by EML we get the $T \times 1$ residual vector, $\widehat{\varepsilon}_t = \pi_t - \sum_{j=1}^{T}\widehat{b}_j L^j \pi_{t-j} + \mathbf{x}_t'\widehat{\beta}$. Then we take a sample of size $\tau = 100$ from $\widehat{\varepsilon}_t$. This sample is randomly selected from a uniform distribution without replacement[19]. This way we get a pseudo-random sample of residuals $\widehat{\varepsilon}_t^*$ of size $\tau \times 1$. Because $\mathbf{x}_t'$ has seasonal variables, the sort order matters, so we just take a sub-matrix of size $\tau \times k$ with the same row order. Next, as with the bootstrap residuals, we take a sample of size

[17] We also try other methods of estimation like Modified Profile Likelihood (MFL) and Nonlinear Least Squares (NLS). We find that for most of the countries the "d" parameter is rather stable.

[18] Instead of defining a mean (μ_t) or letting the estimations take it as a deviation from the sample mean, we set it to zero and use a constant in the regression. In addition, we include in $\mathbf{x}_t$ seasonal dummies and impulse dummies that correct for some outliers.

[19] To assure a consistent bootstrap distribution with the underlying time series model, the residuals have to be homoskedastic and not autocorrelated. We select a subsample without replacement to get rid of the autocorrelation from the specification and avoid inconsistency due to violation in continuity. The consistency depend on the continuity of the mapping between the population distribution of the residuals and the distribution of the parameters. The parameters distribution have to be continuous with the errors. See Horowitz (2001), Section 2.1 and 2.2, and Brownstone and Valletta (2001), Section 1.2 and the literature cited. For a discussion of why is this type of bootstrap better than a block-bootstrap see Horowitz (2003), Section 4.

τ from $\widehat{\pi}_t$, randomly selected from a uniform distribution without replacement. With this "new" variables we construct the pseudo-random sample for the dependent variable $\widehat{\pi}_t^*$ as follow:

$$\widehat{\pi}_t^* = \sum_{j=1}^{\tau} \widehat{b}_j L^j \widehat{\pi}_{t-j} + \mathbf{x}^{*\prime}_t \widehat{\beta} + \widehat{\varepsilon}_t^*, \qquad t = 1, \ldots, \tau. \tag{7}$$

Finally we re-estimate the ARFIMA model and re-forecast (dynamically)[20] but now using the bootstrap dependent variable $\widehat{\pi}_t^*$. We repeat this process a large number of times ($J = 1,000$) and each time we collect the parameter values and the h-step ahead forecast. At the end of the procedure we get a bootstrap parameter values matrix $[\widehat{\Phi}^*(L), \widehat{d}^*, \widehat{\Theta}^*(L)]$ of $J \times k$ and a matrix of forecasts $(\widehat{\pi}_{t+h})$ of $J \times H$.

We construct the forecast confidence interval with the percentiles of the bootstrap forecast distribution F_h^*. For any h-step ahead, the forecast interval is $\mathrm{P}\{F_h^*(\alpha) \leq \widehat{\pi}_{t+h} \leq F_h^*(1-\alpha)\} = 1 - 2\alpha$. In this case for $J = 1000$ and $\alpha = 0.05$, the 90% interval is determined with the 50^{th} value and the 950^{th} value in the h^{th} column of the forecasts matrix.

Transformation and Kernel Density Estimation

The kernel density is a non-parametric estimation of the density function of a random variable. In this section we describe how we perform this estimation for the annual inflation forecast and how we get form the monthly annualized forecast to the annual inflation forecast.

We define the following identities:

$$\pi_t^1 \equiv 1200 \log(P_t/P_{t-1}) \qquad \text{(Monthly Annualized Inflation)}$$
$$\pi_t^{12} \equiv 100 \log(P_t/P_{t-12}) \qquad \text{(Annual Inflation)}$$

[20] The details of the forecast are in the section *Estimating and Forecasting an* ARFIMA *Model*, page 15 of the Appendix.

The annual inflation can expressed as $\pi_t^{12} \equiv \frac{1}{12}\left(\pi_t^1 + \pi_{t-1}^1 + \pi_{t-2}^1 + \ldots + \pi_{t-11}^1\right)$. Using this expression we calculate the expected value and the variance of the forecast.

$$\begin{aligned}
\mathrm{E}_t\left[\pi_{t+h}^{12}\right] &= \mathrm{E}_t\left[\frac{1}{12}\left(\pi_{t+h}^1 + \pi_{t+h-1}^1 + \ldots + \pi_{t+h-11}^1\right)\right], \\
&= \frac{1}{12}\left(\mathrm{E}_t\left[\pi_{t+h}^1\right] + \mathrm{E}_t\left[\pi_{t+h-1}^1\right] + \ldots + \mathrm{E}_t\left[\pi_{t+h-11}^1\right]\right). \\
\mathrm{Var}_t\left[\pi_{t+h}^{12}\right] &= \mathrm{Var}_t\left[\frac{1}{12}\left(\pi_{t+h}^1 + \pi_{t+h-1}^1 + \ldots + \pi_{t+h-11}^1\right)\right], \\
&= \frac{1}{12^2}\sum_{i=0}^{11}\mathrm{Var}_t\left[\pi_{t+h-i}^1\right] + \frac{2}{12^2}\sum_{j=i+1}^{11}\sum_{i=0}^{10}\mathrm{Cov}_t\left[\pi_{t+h-i}^1, \pi_{t+h-j}^1\right].
\end{aligned}$$

The kernel estimator has the following form:

$$\widehat{f}_w(y) = (Tw)^{-1}\sum_{t=1}^{T} K\left(\frac{y - \widehat{y}_t^s}{w}\right). \tag{10}$$

Where w is the bandwidth, y is the random variable and $K(\cdot)$ is the kernel function. In this case we use the Gaussian Kernel. We use the bandwidth value recommended in Silverman (1986, p. 48)[21].

The results of the kernel estimation delivers discrete values of the density function estimation $\widehat{f}(\pi_t)$ for a range of values in the function's domain say, $\pi_{t_0}, \pi_{t_1}, \ldots, \pi_{t_N}$, with certain increments in this range that we define as $\delta = \|\pi_{t_i} - \pi_{t_j}\|, i \neq j$. We obtain the cumulative distribution function estimate $\widehat{F}(\pi_t)$ with the following formula:

$$\widehat{F}(\pi_t) = \sum_{\{j:\pi_{t_j} \geq \pi_t\}} \delta\,\widehat{f}(\pi_{t_j}). \tag{11}$$

Note that we add all the density values for inflations grater than π_t as we like to calculate the probability that inflation exceed some π_t value. Because we are dealing with discrete values, the density estimation is going to have some errors and will not sum to one exactly. The size of the errors will depend on the size of the δ increments (0.1 in our case) and on the symmetry of the curve. The errors will be positive when the slope of the density curve is positive and negative when the slope of the density curve is negative.

[21] $\overline{w} = 0.9T^{-1/5}\min\{\widehat{\sigma}_y^2, \mathrm{iqr}/1.34\}$, where iqr is the interquartile range.

References

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

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

BAILLIE, R. T. (1996): "Long memory processes and fractional integration in econometrics," *Journal of Econometrics*, 73(1), 5–59.

BERAN, J. (1994): *Statistics for long-memory processes*. Chapman & Hall.

BOS, C. S., P. H. FRANSES, AND M. OOMS (2001): "Inflation, forecast intervals and long memory regression models," *Tinbergen Institute Discussion Paper*.

BROWNSTONE, D., AND R. VALLETTA (2001): "The Bootstrap and Multiple Imputations: Harnessing Increased Computing Power for Improved Statistical Tests," *Journal of Economic Perspectives*, 15(4), 129–241.

DAVIDSON, R., AND J. G. MACKINNON (1993): *Estimation and Inference in Econometrics*. Oxford University Press.

DIEBOLD, F. X., AND G. D. RUDEBUSCH (1989): "Long memory and persistence in aggregate output," *Journal of Monetary Economics*, 24(2), 189–209.

DOORNIK, J. A., AND D. F. HENDRY (2009): *Econometric Modelling - PcGive 13: Volume III*. Timberlake Consultants.

DOORNIK, J. A., AND M. OOMS (2004): "Inference and Forecasting for ARFIMA Models With an Application to US and UK Inflation," *Studies in Nonlinear Dynamics & Econometrics*, 8(2).

——— (2006): "A package for estimating forecasting and simulating Arfima Models: Arfima package 1.04 for Ox," *Discussion paper*.

EFRON, B. (1979): "Bootstrap Methods: Another Look at the Jackknife," *The Annals of Statistics*, 7(1), 1–26.

GRANGER, C. W. J., AND R. JOYEUX (1980): "An Introduction to Long-Memory Time Series Models and Fractional Differencing," *Journal of Time Series Analysis*, 1(1), 15–29.

HADAR, J., AND W. R. RUSSELL (1969): "Rules for Ordering Uncertain Prospects," *American Economic Review*, 59(1), 25–34.

HÄIRDLE, W., J. HOROWITZ, AND J.-P. KREISS (2003): "Bootstrap Methods for Time Series," *International Statistical Review*, 71(2), 435–459.

HENRY, M., AND P. ZAFFARONI (2002): "The long range dependence paradigm for macroeconomics and finance," Discussion Papers 0102-19, Columbia University, Department of Economics.

HOROWITZ, J. L. (2001): "Chapter 52 The Bootstrap," in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. E. Leamer, vol. 5, pp. 3159–3228. Elsevier.

——— (2003): "The Bootsrap in Econometrics," *Statistical Science*, 18(2), 211–218.

HOSKING, J. R. M. (1981): "Fractional differencing," *Biometrika*, 68(1), 165–176.

MCCULLOUGH, B. D. (1994): "Bootstrapping forecast intervals: An application to AR(p) models," *Journal of Forecasting*, 13(1), 51–66.

RISSANEN, J. (1978): "Modeling by shortest data description," *Automatica*, 14(5), 465–471.

ROBINSON, P. M. (ed.) (2003): *Time Series with Long Memory*. Oxford University Press.

ROTHSCHILD, M., AND J. E. STIGLITZ (1970): "Increasing risk: I. A definition," *Journal of Economic Theory*, 2(3), 225–243.

SCHWARZ, G. (1978): "Estimating the Dimension of a Model," *The Annals of Statistics*, 6(2), 461–464.

SILVERMAN, B. W. (1986): *Density Estimation for Statistics and Data Analysis*. Chapman & Hall.

VEALL, M. R. (1987): "Boostrapping the Probability Distribution of Peak Electricity Demand," *International Economic Review*, 28(1), 203–212.