What drives long-run inflation expectations?*

(Work in progress; please do not quote or circulate without permission)

Carlos Carvalho	Stefano Eusepi	Emanuel Moench	Bruce Preston
PUC-Rio	NY Fed	Deutsche Bundesbank	Melbourne University

February, 2015

Abstract

According to both central bankers and economic theory, anchored inflation expectations are key to successful monetary policymaking. Yet, we know very little about the drivers of inflation expectations – especially about the long run. We explore a simple model of expectations formation and inflation determination in which the sensitivity of long-run inflation expectations to short-run inflation surprises is state-dependent. Price-setting agents act as econometricians trying to learn about average long-run inflation. They set sticky prices according to their views about future inflation, which hence feed into actual inflation. As in Marcet and Nicolini (2003), agents' estimates of long-run inflation move slowly in normal times, as they keep adding observations to the sample they consider. However, after being surprised repeatedly/largely enough, agents discard old observations and put more weight on recent developments. As a result, long-run inflation expectations become more sensitive to short-run news. We estimate the model using **only short-run** inflation forecasts from surveys, and find that it produces long-run forecasts that track survey measures quite closely. The estimated model has several uses:

1) It can tell a story of how inflation expectations got unhinged in the 1970s; it can also be used to construct a counterfactual history of inflation under anchored long-run expectations.

2) At any given point in time, it can be used to compute the probability of inflation or deflation scares.

3) If embedded into an environment with explicit monetary policy, it can also be used to study the role of policy in shaping the expectations formation mechanism.

JEL classification codes:

Keywords: inflation expectations, learning, survey data, long-run expectations, unanchoring

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, or the Deutsche Bundesbank. Emails: cvianac@econ.puc-rio.br, stefano.eusepi@ny.frb.org, emanuel.moench@gmail.com, bruce.preston@monash.edu. The views

1 Introduction

According to both central bankers and economic theory, anchored inflation expectations are key to successful monetary policymaking. Yet, we know very little about the drivers of inflation expectations – especially about the long run.

In this paper we explore a simple model of expectations formation and inflation determination in which the sensitivity of long-run inflation expectations to short-run inflation surprises is state-dependent. Price-setting agents act as econometricians trying to learn about average long-run inflation. They set sticky prices according to their views about future inflation, which hence feed into actual inflation.

The mechanism underlying the state-dependence of the intensity of inflation expections updating is very simple. As in Marcet and Nicolini (2003), agents' estimates of long-run inflation move slowly in normal times, as they keep adding observations to the sample they consider. However, after being surprised repeatedly/largely enough, agents discard old observations and put more weight on recent developments – that is, they switch to constant-gain learning. As a result, long-run inflation expectations become more sensitive to short-run news. This allows the model to produce episodes of unanchored expectations.

In the model we entertain, inflation expectations feed into actual inflation because of forward-looking Calvo (1983) price setting. We show that a reduced-form representation of the model boils down to a simple model of inflation with a time-varying drift, in the spirit of Cogley and Sbordone (2006) or the shifting-endpoints model of Kozicki and Tinsley (2013). However, there is a fundamental difference betweem our model and models with an exogenous inflation drift. In those models, innovations to the inflation drift are usually independent from short-run innovations to inflation. In our model, the reduced-form representation shows clearly that both short-run inflation movements and the average long-run level of inflation are driven by the same forcing process. This restriction could be expected to handicap the model in terms of fitting data on inflation and inflation expectations, relative to models with an exogenous inflation drift. It turns out that this is not the case. Our model with a tight restriction on the forcing processes fits the data essentially as well as specifications with an exogenous drift.

More importantly, when we estimate our model using only data on short-run inflation forecasts (up to two quarters ahead), it produces long-run inflation forecasts that fit survey data on long-run inflation expectations extremely well. This can be seen in Figure 1.

It shows:

- model-implied 5-10 years ahead inflation forecasts (black line is the median forecast, green shades are the 30-70 percentiles)
- red circles are 10 years ahead survey forecast from the Michigan Survey
- green diamond are 10 years average cpi inflation forecasts from Decision-Makers Poll of portfolio managers

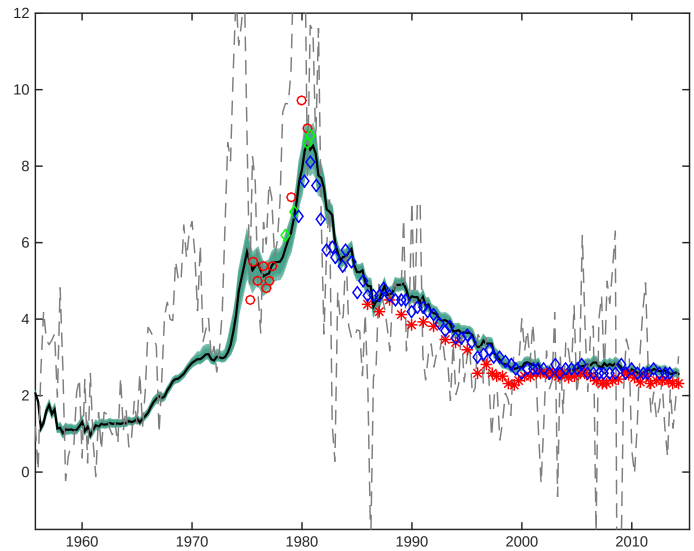

Figure 1: Long-run inflation forecasts produced by the estimated model.
Obs: model-implied 5-10 years ahead inflation forecasts (black line is the median forecast, green shades are the 30-70 percentiles); red circles are 10 years ahead survey forecast from the Michigan Survey; green diamond are 10 years average cpi inflation forecasts from Decision-Makers Poll of portfolio managers; blue diamond are 5-10 years ahead GDP deflator forecast from Blue Chip Economics Indicators; red stars are 5-10 years CPI inflation forecast from Blue Chip Financial Forecasts. actual CPI inflation (dashed grey line)

- blue diamond are 5-10 years ahead GDP deflator forecast from Blue Chip Economics Indicators[1]
- red stars are 5-10 years CPI inflation forecast from Blue Chip Financial Forecasts
- actual CPI inflation (dashed grey line)

Once again, these measures of long-run inflation expectations were **not** used in the estimation. The results are robust to a variety of changes to the specification (number of lags in latent processes, series used in estimation, addition of energy and commodity price series as "observable shocks" etc.)

The estimated model has several uses that we are currently exploring:

1) It can tell a detailed story of how inflation expectations got unhinged in the 1970s; it can also be used to construct a counterfactual history of inflation under anchored long-run expectations.

2) At any given point in time, it can be used to compute the probability of inflation or deflation scares.

3) If embedded into an environment with explicit monetary policy, it can also be used to study the role of policy in shaping the expectations formation mechanism.

1.1 Literature review

[To be added]

[1] We add 0.5% to adjust for the difference in sample mean relative to CPI inflation.

2 A simple model of expectations formation and inflation determination

In this section we present a simple reduced-form model of expectations formation and inflation determination. Because of nominal price rigidities, agents pricing decisions reflect their expectations of future inflation. This creates a link from the latter to actual inflation. In subsection 2.1 we derive this link explicitly in a model with Calvo (1983) pricing. That model also illustrates how our assumptions on agents' views about inflation amount to them being quite "close" to rational expectations – departing from them only because of the need to estimate the long-run average rate of inflation.

Price-setting agents perceive the law of motion of inflation to be

$$\text{PLM:} \qquad \pi_t = \bar{\pi}_t + \omega\varphi_t + e_t, \tag{1}$$

where π_t is inflation, e_t is a zero mean i.i.d. shock, φ_t is zero mean a stationary "short-run component", and $\bar{\pi}_t$ is the average level of inflation that they expect to prevail in the long run. That is what agents need to learn about (i.e., estimate). Other than that, they know the parameter ω and the law of motion for φ_t.

For simplicity, we adopt the usual "anticipated utility assumption": Although agents realize that they will be updating their estimate of $\bar{\pi}_t$ as data come in, the expect $\bar{\pi}_t$ to remain constant. Hence, they form expectations of inflation in the future as

$$\mathbb{E}_t\left[\pi_{t+T}\right] = \bar{\pi}_t + \omega\mathbb{E}_t\left[\varphi_{t+T}\right].$$

These agents set prices that will remain in place for at least a few periods. So, whenever an agent is reseting her price at a given time t, she embeds her expectations of inflation going forward into her price-setting decision. This makes actual inflation depend on price-setters' expectations of future inflation. Suppose that this dependence arises only through the perceived long-run inflation average, so that the actual law of motion of inflation is:

$$\text{ALM:} \qquad \pi_t = T_\pi\bar{\pi}_t + \omega\varphi_t + \tilde{u}_t, \tag{2}$$

where $\tilde{u}_t$ is a zero mean i.i.d. shock, φ_t and $\bar{\pi}_t$ have been defined before, and $T_\pi < 1$ is the parameter that controls the feedback from inflation expectations to actual inflation – it depends on the details of firms' optimal price-setting policies.

Notice that agents' perceived law of motion and the actual law of motion for inflation differ only in that the time-varying intercept in the actual law of motion for inflation is $T_\pi\bar{\pi}_t$ rather than $\bar{\pi}_t$. Hence, if T_π is close to unity these differences will usually be small.

It remains to specify the process by which agents update their estimates of average long-run inflation. We borrow from Marcet and Nicolini (2003), and assume the following learning algorithm:

$$\bar{\pi}_t = \bar{\pi}_{t-1} + k_t^{-1} \times f_{t-1},$$

where

$$f_t = \pi_t - \mathbb{E}_{t-1}\pi_t.$$

As usual in models with learning of this kind, we avoid simultaneity issues by assuming that the $\bar{\pi}_t$ estimate (which affects current decisions) depends on the previous period's forecast error.

The learning gain $k_t > 1$ is determined by

$$k_t = \begin{cases} \bar{g}, \text{ if } \frac{\left|\mathbb{E}_{t-1}\pi_t - \mathbf{E}_{t-1}^M \pi_t\right|}{\sqrt{MSE^M}} > v \\ \\ k_{t-1} + 1, \text{ otherwise} \end{cases},$$

where $\mathbb{E}_{t-1}\pi_t$ is the agent's one-period-ahead forecast, $\mathbf{E}_{t-1}^M \pi_t$ is the model-consistent one-period-ahead expectation, and MSE^M is the mean squared error associated with one-period-ahead model-consistent expectations – i.e. $MSE^M = \mathbf{E}^M \left[\pi_t - \mathbf{E}_{t-1}^M \pi_t\right]^2$.

The state-dependent gain k_t captures agents' atempts to protect against "structural change". A constant gain ($\bar{g} > 1$) produces better forecasts when the economic environment changes, but it does not converge in a stationary environment. A decreasing gain (OLS-type) estimator, on the other hand, converges in stationary environments. As in Marcet and Nicolini (2003), the proposed learning algorithm uses OLS in periods or relative stability and switches to constant-gain when instability is "detected" – i.e., when $\left|\mathbb{E}_{t-1}\pi_t - \mathbf{E}_{t-1}^M \pi_t\right| > v \; MSE^M$. A stable environment here is an economy where inflation does not vary too much. So that the distance between agents' forecast and the model-consistent forecast tend to be small.

To obtain more intuition on the properties of this learning algorithm, write:

$$\begin{aligned} \left|\mathbb{E}_{t-1}\pi_t - \mathbf{E}_{t-1}^M \pi_t\right| &= \left(1 - \tilde{T}_\pi\right) \bar{\pi}_t \\ &= \left(1 - \tilde{T}_\pi\right) \left(\bar{\pi}_{t-1} + k_t^{-1} \times f_{t-1}\right) \\ &= \left(1 - \tilde{T}_\pi\right) \left[\bar{\pi}_0 + \sum_{\tau=0}^{t} k_\tau^{-1} f_{\tau-1}\right]; \text{ given } \bar{\pi}_0, \; f_{-1}. \end{aligned}$$

Hence, the distance $\left|\mathbb{E}_{t-1}\pi_t - \mathbf{E}_{t-1}^M \pi_t\right|$ tends to be large when forecast errors happen to be of the "same sign" for a few periods. This pushes $\bar{\pi}_t$ away from its long-run value (of zero or otherwise), and drives a larger wedge between agents' model of the economy and the truth.

One may wonder how agents can switch to a constant gain algorithm based on a criterion that involves model-consistent expectations – which they are assumed not to know. This make this assumption in the "as if" tradition in Economics. Agents indeed do not know $\mathbf{E}_{t-1}^M \pi_t$, but they conclude that there must be something unusual happening when their forecasts "go astray" in the sense that $\left|\mathbb{E}_{t-1}\pi_t - \mathbf{E}^M_{t-1}\pi_t\right|$ becomes too large. This assumption captures agents' efforts to deter model instability by using statistical tools to detect time-variation in their model's intercept. They key here is that this term depends on past forecast errors. In turn, the size of forecast errors depends on the gain being used by agents. This opens the door for multiple "learning equilibria", as discussed below.

2.1 A model with explicit price-setting decisions

This section develops model where the link from expectations to actual inflation assumed above arises explicitly in the context of Calvo (1983) pricing, as in Preston (2005). For expositional simplicity we rely on a model without indexation (which we add later, for estimation).

The model starts the New Keynesian Phillips Curve under arbitrary expectations, as in Preston (2005):

$$\pi_t = \kappa \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \hat{s}_T + \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \beta (1-\alpha) \pi_{T+1} + u_t, \tag{3}$$

where, as before, $\mathbb{E}_t$ denotes agents' expectations, $\kappa = \frac{(1-\alpha\beta)(1-\alpha)}{\alpha}\zeta$, β is the discount factor, α is the infrequency of price changes and ζ is the Ball and Romer (1990) index of real rigidities. The NKPC is subject to shocks u_t, and marginal cost s_t evolves according to

$$\lambda\pi_t + s_t = \varphi_t, \tag{4}$$

where $\lambda > 0$ is a parameter, and φ_t is exogenous and evolves according to

$$\varphi_t = \rho_1\varphi_{t-1} + \rho_2\varphi_{t-2} + \epsilon^{\varphi}_t. \tag{5}$$

This setup can be mapped into a popular specification in the literature. In the simple new Keynesian model we can interpret (4) as deviations from a targeting rule under discretion with a constant target of zero inflation. The looser interpretation is that we allow some feed-back from inflation to the marginal costs, as a reduced-form way to capture some policy responses. This specification (with $\lambda > 0$) is useful to control the inflation to drift (discussed below).

What do agents know? To keep the framework as simple as possible, and in line with the previous section, we assume they know everything except for the long-run mean of inflation. Agents estimate a long-run inflation mean $\bar{\pi}_t$ in every period. They form expectations of marginal costs and inflation using the correct law of motion for the marginal cost

$$\lambda\pi_t + s_t = \varphi_t \quad \rightarrow \quad s_t = \varphi_t - \lambda\pi_t,$$

and the following equilibrium *perceived* law of motion (PLM) for inflation:

$$\pi_t = \bar{\pi}_t + \omega_{\pi 1}\varphi_{t-1} + \omega_{\pi 2}\varphi_{t-1} + e_t,$$

where ω denote the rational expectations coefficients (to be solved for by matching coefficients).

It is useful to express equations in terms of deviations from a constant mean π^C. So every variable in the model is expressed in terms of deviation from this known constant. This implies, for example, that $\bar{\pi}_t$ denotes perceived deviations from π^C. This save notation and it is equivalent to the case where agents actually learn the mean, at least if we assume (as we do here) that the actual long-term mean of inflation remains constant at π^C. So we can **alternatively** then define

$$\pi_t = \ln(\frac{P_t}{P_{t-1}}) - \pi^C.$$

Notice finally that to keep the model simple, under this alternative formulation we do assume indexation to π^C, so that the Phillips curve coefficients do not depend on the level of inflation.

2.1.1 Rational expectations

Under rational expectations, the solution to this model is well known. The model simplifies to

$$\pi_t = \kappa s_t + \beta \mathbf{E}_t \pi_{t+1} + u_t,$$

where $\mathbf{E}_t$ denotes rational expectations. Substituting the marginal cost using (4) gives

$$\pi_t = \kappa \left(\varphi_t - \lambda \pi_t\right) + \beta \mathbf{E}_t \pi_{t+1} + u_t,$$

or, rearranging,

$$\pi_t = \tilde{\kappa} \varphi_t + \tilde{\beta} \mathbf{E}_t \pi_{t+1} + \tilde{u}_t,$$

where $\tilde{\kappa} = \frac{\kappa}{1+\lambda\kappa}$ and $\tilde{\beta} = \frac{\beta}{1+\lambda\kappa}$. So this is the same as a standard NKPC model with exogenous marginal cost, but with modified parameters. Iterating forward we get

$$\begin{aligned}
\pi_t &= \tilde{\kappa} \mathbf{E}_t \sum_{T=t}^{\infty} \tilde{\beta}^{T-t} \varphi_T + \tilde{u}_t \\
&= \tilde{\kappa} e_1' \left(I - \tilde{\beta}\Phi\right)^{-1} \boldsymbol{\varphi}_t + \tilde{u}_t \\
&= \left(\omega^*_{\pi 1} \rho_1 + \omega^*_{\pi 2}\right) \varphi_{t-1} + \omega^*_{\pi 1} \rho_2 \varphi_{t-2} + \omega^*_{\pi 1} \epsilon^{\varphi}_t + \tilde{u}_t,
\end{aligned}$$

where

$$\begin{aligned}
\omega^*_{\pi 1} &= \frac{\tilde{\kappa}}{1 - \tilde{\beta}\left(\rho_1 + \tilde{\beta}\rho_2\right)} \\
\omega^*_{\pi 2} &= \frac{\tilde{\kappa}\tilde{\beta}\rho_2}{1 - \tilde{\beta}\left(\rho_1 + \tilde{\beta}\rho_2\right)}.
\end{aligned}$$

2.1.2 Arbitrary expectations

First we substitute for the marginal cost in the Phillips curve to get

$$\pi_t = \kappa\left(\varphi_t - \lambda\pi_t\right) + \kappa\alpha\beta\mathbb{E}_t\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\left(\varphi_{T+1} - \lambda\pi_{T+1}\right)$$
$$+\mathbb{E}_t\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\beta\left(1-\alpha\right)\pi_{T+1} + u_t.$$

This can be simplified to

$$\begin{aligned}\pi_t &= \tilde{\kappa}\varphi_t + \tilde{\kappa}\alpha\beta\mathbb{E}_t\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\varphi_{T+1} \\ &\quad + \left[\tilde{\beta}\left(1-\alpha\right) - \tilde{\kappa}\lambda\alpha\beta\right]\mathbb{E}_t\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\pi_{T+1} + \tilde{u}_t.\end{aligned}$$

The simplest case that illustrates the logic of the solution assumes AR(1) marginal cost: $\rho_2 = 0$. Then, using the agents PLM for inflation:

$$\pi_t = \bar{\pi}_t + \tilde{\kappa}\left(1 - \tilde{\beta}\rho_1\right)^{-1}\varphi_t + e_t.$$

We can then obtain the forecasts (under the usual “anticipated utility” assumption for $\bar{\pi}_t$):

$$\mathbb{E}_t\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\varphi_{T+1} = \frac{\rho_1}{1-\rho_1\alpha\beta}\varphi_t$$
$$\mathbb{E}_t\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\pi_{T+1} = \frac{1}{1-\alpha\beta}\bar{\pi}_t + \frac{\tilde{\kappa}\rho_1}{\left(1-\tilde{\beta}\rho_1\right)\left(1-\rho_1\alpha\beta\right)}\varphi_t$$

and replace them into the NKPC to obtain

$$\pi_t = T_\pi\bar{\pi}_t + \tilde{\kappa}\varphi_t + \tilde{\kappa}\frac{\alpha\beta\rho_1}{1-\rho_1\alpha\beta}\varphi_t + \left[\frac{\tilde{\beta}\left(1-\alpha\right) - \lambda\tilde{\kappa}\alpha\beta}{\left(1-\tilde{\beta}\rho_1\right)\left(1-\rho_1\alpha\beta\right)}\right]\tilde{\kappa}\rho_1\varphi_t + \tilde{u}_t,$$

where

$$T_\pi = \frac{\tilde{\beta}\left(1-\alpha\right) - \lambda\tilde{\kappa}\alpha\beta}{1-\alpha\beta} < 1.$$

By definition of rational expectations, the actual law of motion for inflation is then

$$\pi_t = T_\pi\bar{\pi}_t + \tilde{\kappa}\left(1-\tilde{\beta}\rho_1\right)^{-1}\varphi_t + \tilde{u}_t.$$

As before, the ALM displays a time-varying intercept that differs from the PLM's. Notice that, if marginal costs respond strongly to inflation (higher λ), T_π decreases and so the feedback effect from learning is muted.

2.1.3 Allowing for indexation

If we allow for indexation to past inflation, we need only make small changes to the model above.

$$\begin{gathered}\pi_t = \gamma_p \pi_{t-1} + \iota_p \left(1-\gamma_p\right) \bar{\pi}_t + \\ + \kappa \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \hat{s}_T + \\ \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \beta (1-\alpha) \left(\pi_{T+1} - \gamma_p \pi_T - \iota_p \left(1-\gamma_p\right) \bar{\pi}_{T+1}\right) + u_t\end{gathered}$$

where again

$$\kappa = \frac{(1-\alpha\beta)(1-\alpha)}{\alpha}\zeta$$

and the marginal cost is now evolving according to

$$\lambda \left(\pi_t - \gamma_p \pi_{t-1}\right) + s_t = \varphi_t.$$

Recall that inflation and the perceived long-term mean are defined in terms of deviation from the constant π^C. I this case $\left(1-\gamma_p\right)(1-\iota_p)$ fraction of indexation comes form π^C and $0 \leq \iota_p \leq 1$. In detail, the indexation term in *levels* can be expressed as:

$$\Pi_{t-1}^{\gamma_p} \bar{\Pi}_{t-1}^{\iota_p (1-\gamma_p)} \left(\Pi^C\right)^{(1-\iota_p)1-\gamma_p}.$$

Under rational expectations we can write

$$\tilde{\pi}_t = \kappa s_t + \beta \mathbf{E}_t \tilde{\pi}_{t+1} + u_t$$

where

$$\tilde{\pi}_t = \pi_t - \gamma_p \pi_{t-1}.$$

Note that ι_p does not matter matter because $\bar{\pi}_t = 0$ (or $\bar{\Pi}_t = \Pi^C$). Indexation implies a $\left(1-\gamma_p\right)$ weight on π^C.

As above, we can write

$$\tilde{\pi}_t = \kappa \left(\varphi_t - \lambda \tilde{\pi}_t\right) + \beta \mathbf{E}_t \tilde{\pi}_{t+1} + u_t$$

so that

$$\tilde{\pi}_t = \tilde{\kappa} \varphi_t + \tilde{\beta} \mathbf{E}_t \tilde{\pi}_{t+1} + \tilde{u}_t.$$

The solution is then

$$\pi_t = \gamma_p \pi_{t-1} + \tilde{\kappa} \mathbf{E}_t \sum_{T=t}^{\infty} \tilde{\beta}^{T-t} \varphi_T + \tilde{u}_t.$$

Under general expectations the model becomes

$$\pi_t = \gamma_p \pi_{t-1} + \iota_p \left(1-\gamma_p\right) \bar{\pi}_t +$$
$$+\kappa \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \hat{s}_T +$$
$$+\mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \beta (1-\alpha) \left(\pi_{T+1} - \gamma_p \pi_T - \iota_p \left(1-\gamma_p\right) \bar{\pi}_{T+1}\right) + u_t.$$

This can be rearranged as

$$\left[1 + \tilde{\beta}(1-\alpha)\gamma_p - \tilde{\kappa}\alpha\beta\lambda\gamma_p\right] \pi_t = \gamma_p \pi_{t-1} + \frac{\iota_p \left(1-\gamma_p\right)}{(1+\lambda\kappa)} \bar{\pi}_t +$$
$$+\tilde{\kappa} \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \varphi_T +$$
$$+\left[\tilde{\beta}(1-\alpha) - \tilde{\kappa}\alpha\beta\lambda\right] \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \left(1-\gamma_p \alpha\beta\right) \pi_{T+1} +$$
$$-\beta(1-\alpha) \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \iota_p \left(1-\gamma_p\right) \bar{\pi}_{T+1} + u_t.$$

As before we have to figure out the expectations. Agents' PLM is now given by

$$\pi_t = \left(1-\gamma_p\right) \bar{\pi}_t + \gamma_p \pi_{t-1} + \omega_{\pi,1}\varphi_t + \omega_{\pi,2}\varphi_{t-1} + \tilde{u}_{t.}$$

In order to understand the logic of the solution, let's focus on π_t and $\bar{\pi}_t$, and let's set $\varphi_t = 0$ (this is without loss of generality because φ_t is exogenous). The discounted sum for inflation forecasts is then

$$\begin{aligned} \mathbb{E}_t \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \pi_{T+1} &= \mathbb{E}_t \left[\sum_{T=t}^{\infty} (\alpha\beta)^{T-t} + \gamma_p \sum_{T=t}^{\infty} \left(\alpha\beta\gamma_p\right)^{T-t}\right] \bar{\pi}_t \\ &= \left(\frac{1}{1-\alpha\beta} - \frac{\gamma_p}{1-\alpha\beta\gamma_p}\right) \bar{\pi}_t + \gamma_p \left(1-\alpha\beta\gamma_p\right)^{-1} \pi_t. \end{aligned}$$

Substituting in the Phillips curve we get

$$\pi_t = \left(1-\gamma_p\right) \tilde{T}_\pi \bar{\pi}_t + \gamma_p \pi_{t-1} + \left(\omega^*_{\pi 1}\rho_1 + \omega^*_{\pi 2}\right) \varphi_{t-1} + \omega^*_{\pi 1}\rho_2 \varphi_{t-2} + \omega^*_{\pi 1}\epsilon^{\varphi}_t + \tilde{u}_t,$$

where

$$\begin{aligned}\omega^*_{\pi 1} &= \frac{\tilde{\kappa}}{1-\tilde{\beta}\left(\rho_1+\tilde{\beta}\rho_2\right)} \\ \omega^*_{\pi 2} &= \frac{\tilde{\kappa}\tilde{\beta}\rho_2}{1-\tilde{\beta}\left(\rho_1+\tilde{\beta}\rho_2\right)},\end{aligned}$$

and

$$\tilde{T}_\pi = \iota_p\left(\frac{1}{1+\lambda\kappa}-\frac{\tilde{\beta}\left(1-\alpha\right)}{1-\alpha\beta}\right)+\left(\tilde{\beta}\left(1-\alpha\right)-\tilde{\kappa}\alpha\beta\lambda\right)\frac{1-\gamma_p\alpha\beta}{1-\gamma_p}\left(\frac{1}{1-\alpha\beta}-\frac{\gamma_p}{1-\alpha\beta\gamma_p}\right),$$

which differs from agents' PLM

$$\pi_t = \left(1-\gamma_p\right)\bar{\pi}_t+\gamma_p\pi_{t-1}+\omega_{\pi,1}\varphi_{t-1}+\omega_{\pi,2}\varphi_{t-2}+e_t$$

only by $\tilde{T}_\pi$.[2]

2.2 Relationship with models with an exogenous "inflation drift"

Having shown how the reduced-form model presented in the beginning of this section can be obtained from a model with explicit price-setting, we now relate it to models with an exogenous "inflation drift".

Revert to the model without indexation:

$$\pi_t = T_\pi\bar{\pi}_t+\tilde{\kappa}\mathbf{E}_t\sum_{T=t}^{\infty}\tilde{\beta}^{T-t}\left(\varphi_T+u_T\right) \tag{6}$$

where agents PLM is

$$\pi_t = \bar{\pi}_t+\omega\varphi_{t-1}+e_t,$$

and where ω is consistent with the correct evaluation of the PDV in (6). The actual inflation process can then be written as

$$\pi_t = T_\pi\bar{\pi}_t+\mathbb{Z}_t,$$

where

$$\mathbb{Z}_t = \tilde{\kappa}\mathbf{E}_t\sum_{T=t}^{\infty}\tilde{\beta}^{T-t}\left(\varphi_T+u_T\right).$$

[2] The coefficients $\omega_{\pi,1}$ and $\omega_{\pi,2}$ from the agents' PLM can be solved for by matching coefficients with those in the ALM for inflation.

The agents' updating algrithm for updating mean inflation can be written as

$$\begin{aligned}\bar{\pi}_{t+1} &= \bar{\pi}_t + k_t^{-1}\left(\pi_t - \mathbf{E}_{t-1}\pi_t\right) \\ &= \bar{\pi}_t + k_t^{-1}\left(T_\pi \bar{\pi}_t + \mathbb{Z}_t - \bar{\pi}_t - \omega\varphi_{t-1}\right) \\ &= \bar{\pi}_t + k_t^{-1}\left[\left(T_\pi - 1\right)\bar{\pi}_t + \mathbb{Z}_t - \mathbf{E}_{t-1}\mathbb{Z}_t\right] \\ &= \left[1 + k_t^{-1}\left(T_\pi - 1\right)\right]\bar{\pi}_t + k_t^{-1}\left(\mathbb{Z}_t - \mathbf{E}_{t-1}\mathbb{Z}_t\right).\end{aligned}$$

Hence, the model can be represented as

$$\begin{aligned}\pi_t &= T_\pi \cdot \bar{\pi}_t + \mathbb{Z}_t \\ \bar{\pi}_{t+1} &= \rho_{\bar{\pi},t}\bar{\pi}_t + k_t^{-1}\left(\mathbb{Z}_t - \mathbf{E}_{t-1}\mathbb{Z}_t\right)\end{aligned}$$

where

$$\rho_{\bar{\pi},t} = \left[1 - k_t^{-1}\left(1 - T_\pi\right)\right]$$

and where, under our assumptions, $\mathbb{Z}_t$ is an exogenous process. In the simple AR(1) case above (ρ_2) this process is

$$\begin{aligned}\mathbb{Z}_t &= \tilde{\kappa}\left(1 - \tilde{\beta}\rho_1\right)^{-1}\varphi_t + \tilde{u}_t \\ &= \tilde{\varphi}_t + \tilde{u}_t.\end{aligned}$$

Hence, our model can be seen to correspond to a simple model of inflation with an exogenous drift. However, differently from other specifications with a time-varying inflation intercept, in which innovations to the drift and short-run inflation shocks are independent, in our model the same $\mathbb{Z}_t$ process drives the actual law of motion for inflation and determines the innovations that lead agents to update their assessment of average long-run inflation.

3 Estimation and results

We estimate the model with indexation. The quarterly data used in the estimation are CPI inflation, Livingston survey 6-months ahead CPI forecasts (available 1955-2014Q2), and 1- and 2-quarter ahead CPI forecasts from SPF available since 1983Q2. The full sample is 1955-2014Q2, and the model involves missing observations and unobserved processes. It is estimated by casting the system in state-space form, appending the equations for the inflation forecasts used as observables, and employing full-information Bayesian estimation methods (details to be added).

Note that we only discipline the estimation of the model with short-run inflation forecasts. Hence, a comparison of the long-run inflation forecasts produced by the model with long-run survey forecasts is a legitimate "test" of the model as a description of long-run inflation expectations.

Results of the estimated model are presented in Figure 1 in the Introduction. It shows long-run inflation forecasts produced by the model, and various measures of long-run inflation expectations that were not used in the estimation.

Figure 2 shows the path of the learning gain implied by the estimated model. The solid black line shows the point-by-point median across the distribution of paths implied by the estimated model. The dashed red line is the mode, and the bands depict the 30-70 percentiles.

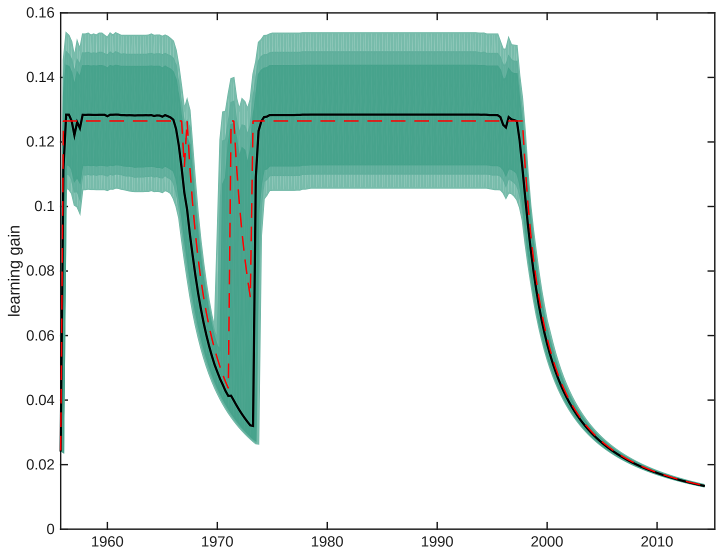

Figure 2: Learning gain in the estimated model.
Obs: solid black line shows the point-by-point median across the distribution of paths implied by the estimated model; dashed red line is the mode, and the bands depict the 30-70 percentiles.

The estimated model tells a story in which inflation expectations got unacnhored in the 1970s, after a series of short-run inflation surprises. Agents kept using a constant gain learning algorithm for quite a while – until the late 1990s. At that point, they reverted back to a decreasing gain procedure, and as a result the sensitivity of long-run inflation expectations decreased. This explains why long-run inflation expectations have been remarkably stable since the 2000s, despite large swings in inflation.

4 Conclusion

[To be added]

References

[1] [To be added]

5 Appendix

[To be added]