Euro area inflation persistence in an estimated nonlinear DSGE model

Gianni Amisano
Università di Brescia

Oreste Tristani
European Central Bank

EABCN/CEPR, Swiss National Bank Workshop, Zurich 29-30 August 2006

Inflation persistence

- Existing literature focused on either statistical measures of persistence or linear Phillips curve type analyses (Angeloni *et al.*, 2005).

$$\widehat{\pi}_t = \frac{\iota}{1+\beta\iota}\widehat{\pi}_{t-1} + \frac{\beta}{1+\beta\iota}\mathsf{E}_t\widehat{\pi}_{t+1} + \frac{1-\zeta}{\zeta}\frac{1-\zeta\beta}{1+\beta\iota}\widehat{mc}_t$$

- Statistical measures: estimates of persistence depend on whether one allows for structural breaks (e.g. Levin and Piger, 2004). From Phillips curve: intrinsic persistence is probably on the high side (e.g. typically high estimates of "indexation parameter").

- Question: to what extent are these estimates affected by the linear set-up?

Why a nonlinear model (1)

- Because we can! Do linearised models miss features of the data?

- A nonlinear model can account for richer conditional dynamics than a linear model. Different patterns of persistence could be consistent with the same model. Loose intuition:

$$\widehat{\pi}_{t+1} = a\widehat{\pi}_t + b\widehat{\pi}_t^2 + \varepsilon_{t+1}$$

- Large deviations from the steady state happen. E.g. inflation 1980-2004: max 12.0 min 0.6, mean 4.2. If the linear model is an approximation, these deviations from the steady state may be large and induce biased estimates.

Why a nonlinear model (2)

- Nonlinear models provide sharper parameter estimates (An & Schorfeide, 2005). They also exploit more model restrictions, which mitigates the identification problems highlighted by Canova & Sala (2005): eg. variances.

- Within our framework we also look at the breaks/no breaks issue. Advantages:
  - multivariate framework;
  - structural interpretation of the breaks.

Results

- In the model without breaks, nonlinear terms are sizable: impulse responses are easily regime dependent. This model appears to be closer to the data.

- Differences in the impulse responses are statistically significant within our sample.

- Inflation "persistence" (in response to a change in the target) has declined over our sample.

Outline

- Sketch of the model.
- Estimation issues.
- Results.

Main ingredients of the model (1)

- Woodford (2003) set-up. Unit measure of households with utility separable in labour, with habit formation (Fuhrer, 2000). Serially correlated tax shock τ_t to allow for inflation-output trade-offs.

$$u\left(C_t, C_{t-1}, L_t\right) = \frac{\left(C_t - hC_{t-1}\right)^{1-\gamma}}{1-\gamma} - \int_0^1 \chi L_t\left(i\right)^{\phi} di$$

$$P_t C_t + B_t \leqslant (1-\tau_t)\left(\int_0^1 w_t\left(i\right) L_t\left(i\right) di + \int_0^1 \Xi_t\left(i\right) di\right) + W_t$$

Main ingredients of the model (2)

- Production function with serially correlated TFP shock A_t

$$Y_t\left(i\right) = A_t L\left(i\right)^{\alpha}$$

- Monopolistically competitive firms set prices subject to Calvo lottery: in each period, prices can be optimised with probability $1 - \zeta$, or set by an indexation rule (intrinsic inflation persistence) to

$$P_t = P_{t-1}\Pi_t^{\iota}\Pi^{1-\iota}$$

Main ingredients of the model (3)

Two possible rules

$$
\begin{aligned}
i_t &= (1-\rho_I)(\pi - \ln\beta) + \psi_\pi(\pi_t - \pi_t^*) + \psi_y(y_t - y_t^n) + \rho_I i_{t-1} + v_t^i \\
\pi_t^* &= (1-\rho_\pi)\pi + \rho_\pi \pi_{t-1}^* + v_t^{\pi^*}
\end{aligned}
$$

or

$$
\begin{aligned}
i_t &= (1-\rho_I)\Big((\pi_t^* - \ln\beta) + \psi_\pi(\pi_t - \pi_t^*) + \psi_y(y_t - y_t^n)\Big) + \rho_I i_{t-1} + v_t^i \\
\pi_t^* &= \pi_{t-1}^* + v_t^{\pi^*}
\end{aligned}
$$

Solution method

- Second-order approximations (in logs) around the non-stochastic steady state (Klein, 2005).

$$\widehat{x}_{t+1} = \frac{1}{2}k_x + P\widehat{x}_t + \frac{1}{2}G\left(\widehat{x}_t \otimes \widehat{x}_t\right) + \varepsilon_{t+1}$$
$$\widehat{y}_t = \frac{1}{2}k_y + F\widehat{x}_t + \frac{1}{2}E\left(\widehat{x}_t \otimes \widehat{x}_t\right)$$

- Variables $M1 : y_t^o = [\pi_t, i_t, y_t]'$; $x_t = \left[\pi_{t-1}, y_{t-1}^{nat}, y_{t-1}, i_{t-1}, a_t, \pi_t^*, \tau_t, v_t^i\right]'$
$M2 : y_t^o = [\Delta\pi_t, i_t - \pi_t, y_t]'$;
$x_t = \left[\pi_{t-1}^* - \pi_{t-1}, y_{t-1}^{nat}, y_{t-1}, i_{t-1} - \pi_{t-1}, a_t, v_t^\pi, \tau_t, v_t^i\right]'$

Definition of inflation persistence

- Broad definition: how slow is the return of inflation to its long run value once it has deviated from it?

- Our implementation: how large and persistent is the response of inflation after a given impulse shock starting from different inflation levels?

- Idea: high inflation (could) become entrenched in expectations.

$$\begin{aligned}\pi_t &= (1-\iota)\,\pi + \iota\pi_{t-1} + f\left(\cdot, \iota\pi_t, \pi_{t+1}, \ldots\right) \\ \widehat{\pi}_t &= \frac{1}{2}k_\pi + F_\pi \widehat{x}_t + \frac{1}{2}E_\pi\left(\widehat{x}_t \otimes \widehat{x}_t\right)\end{aligned}$$

Estimation issues

- Dynamic system

$$\begin{aligned}\text{(measurement equation) } \mathbf{y}_t^o &= G(\mathbf{x}_t, \mathbf{v}_t, \boldsymbol{\theta}) \\ \text{(state equation)} \mathbf{x}_t &= H(\mathbf{x}_{t-1}, \mathbf{w}_t, \boldsymbol{\theta})\end{aligned}$$

- Filtering problem: not feasible analitycally. Filtering by simulation. Simplest filter: Particle Filter (PF). Arulampalam *et al*. (2002) , IEEE; Doucet *et al*. (2001); Fernandez-Villaverde and Rubio-Ramirez (2005); An and Schorfheide (2006).

Particle filter

- At each point in time, we start from a large number N of draws of $\mathbf{x}_t$ (a so called *swarm of particles*): $\mathbf{x}_t^{(i)}, i = 1, 2, ..., N$.

1. (a) Projection: draw one-step-ahead forecasted particles $\widetilde{\mathbf{x}}_{t+1}^{(j)}$ (i.e. simulate the state equation $p(\mathbf{x}_{t+1}|\mathbf{x}_t^{(j)}, \boldsymbol{\theta})$). In this way, the projection distribution is approximated by the sample

$$p(\mathbf{x}_{t+1}|\underline{\mathbf{y}}_t^o, \boldsymbol{\theta}) \approx \frac{1}{N} \sum_{j=1}^{N} p(\mathbf{x}_{t+1}|\mathbf{x}_t^{(j)})$$

(b) Update: now we need to take into account the information coming from the observation of $\mathbf{y}^o_{t+1}$ to construct $p(\mathbf{x}_{t+1}|\underline{\mathbf{y}}^o_{t+1}, \boldsymbol{\theta})$. Given $\widetilde{\mathbf{x}}^{(j)}_{t+1}$, we construct the probability of $\mathbf{y}^o_{t+1}$ through weights $w^{(j)}_{t+1} \propto p(\mathbf{y}^o_{t+1}|\mathbf{x}^{(j)}_{t+1}, \boldsymbol{\theta})$. The updated density $p(\mathbf{x}_{t+1}|\underline{\mathbf{y}}^o_{t+1}, \boldsymbol{\theta})$ is approximated by the importance sampler

$$\left(\widetilde{\mathbf{x}}^{(j)}_{t+1}, w^{(j)}_{t+1}\right), j = 1, 2, ..., N.$$

Sample mean of weights is consistent estimate of $p(\mathbf{y}^o_{t+1}|\underline{\mathbf{y}}^o_{t}, \boldsymbol{\theta})$.

(c) This sample can be resampled using the weights $w^{(j)}_{t+1}$ to obtain a new set of particles with equal weight $\left(\mathbf{x}^{(j)}_{t+1}, 1\right), j = 1, 2, ..., N$.

Problems with the PF

- PF is based on a blind proposal. If candidate distribution too spread out, large number of draws given negligible weights $\Rightarrow$ poor numerical accuracy properties. Figure (A): first particle will be killed either by reweighting or by resampling.

- Sensitivity to outliers. Figure (B): particle 3 will get a unit weight. All the others killed by reweighting or by resampling

- No or low measurement error: $p(\mathbf{y}^o_{t+1}|\mathbf{x}^{(j)}_{t+1}, \boldsymbol{\theta})$ to compute weights, but if no measurement error this becomes degenerate.

figure A: PF at work, N=3

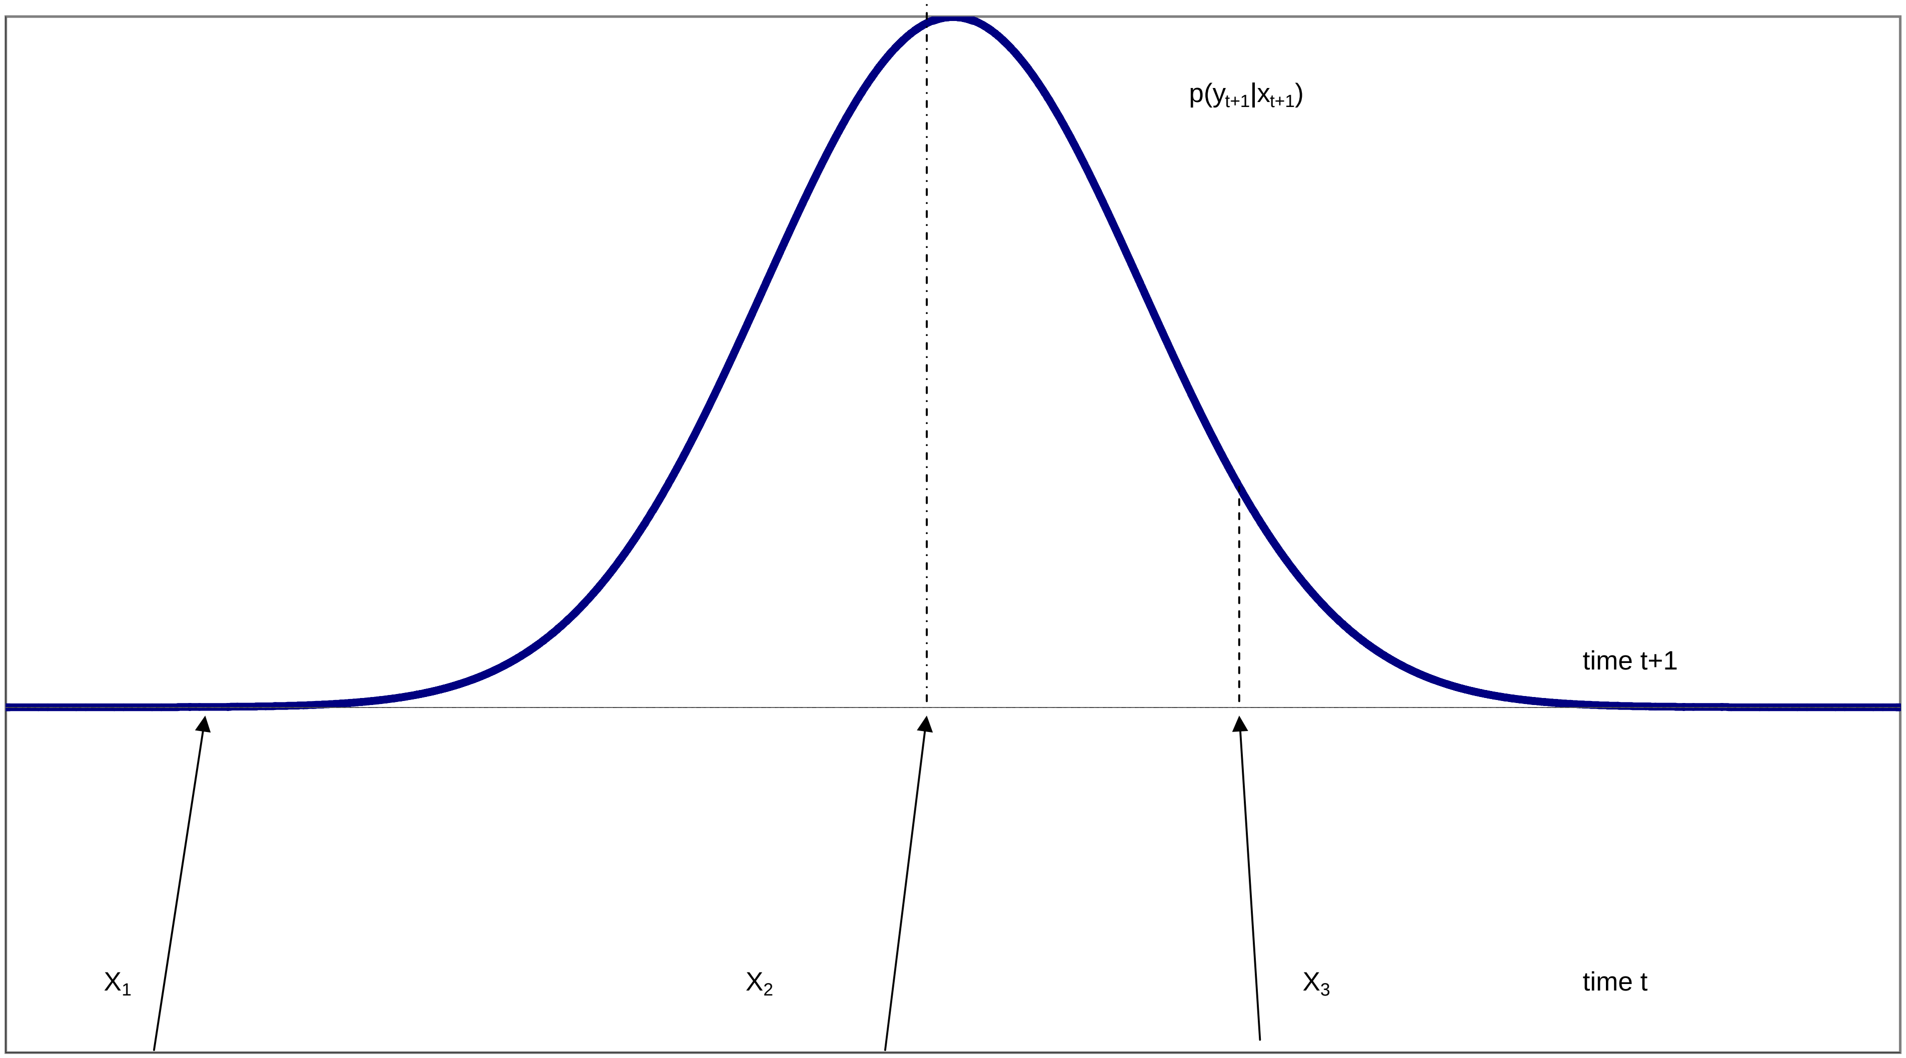

Figure B: PF at work, N=3

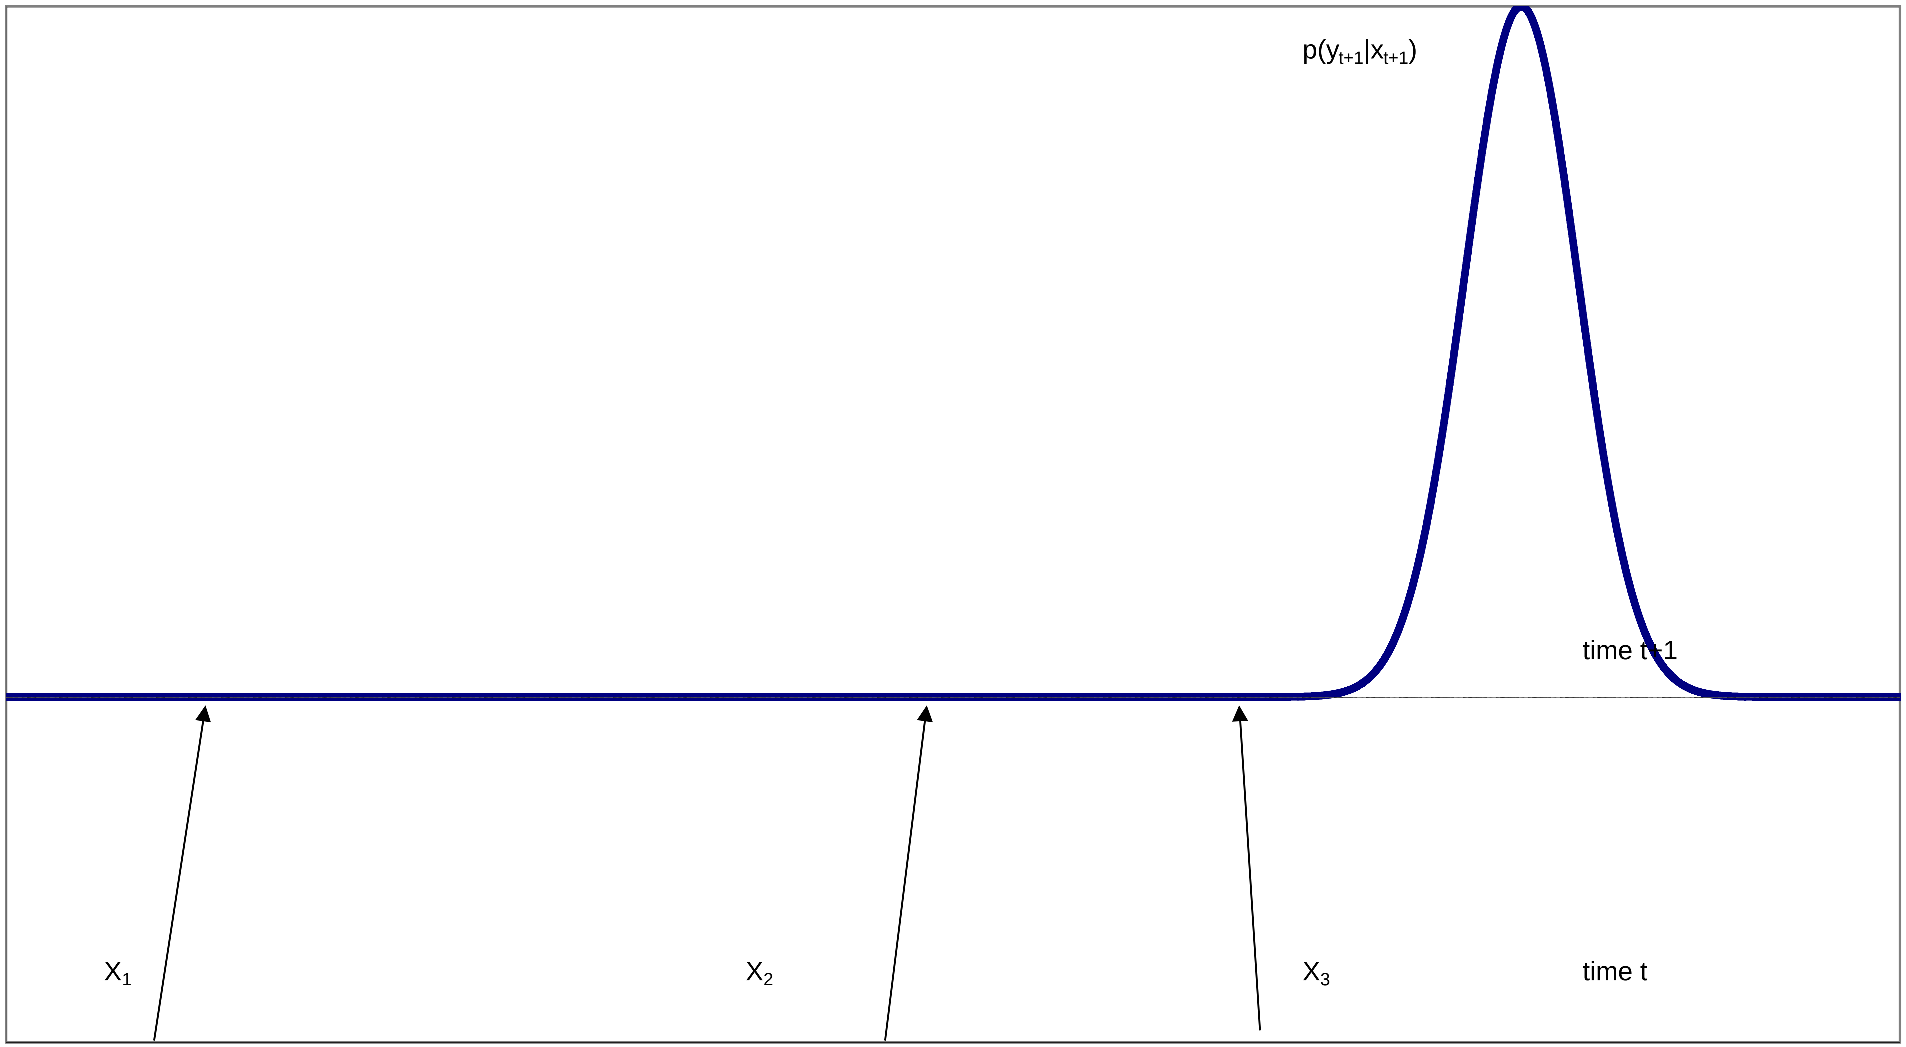

Our preferred alternative

- Conditional particle filter (CPF) of Ionides (2003). Instead of drawing $\mathbf{x}_{t+1}^{(j)}$from $p(\mathbf{x}_{t+1}|\mathbf{x}_t^{(j)}, \boldsymbol{\theta})$ and udpate (reweight) using $p(\mathbf{y}_{t+1}^o|\mathbf{x}_{t+1}^{(i)}, \boldsymbol{\theta})$:
  - sample $\mathbf{x}_{t+1}^{(j)}$ from $p(\mathbf{x}_{t+1}|\mathbf{x}_t^{(j)}, \mathbf{y}_{t+1}^o, \boldsymbol{\theta})$ (i.e. condition explicitly on $\mathbf{y}_{t+1}^o$)
    – problem: $p(\mathbf{x}_{t+1}|\mathbf{x}_t^{(j)}, \mathbf{y}_{t+1}^o, \boldsymbol{\theta})$ and $p(\mathbf{y}_{t+1}^o|\mathbf{x}_t^{(i)}, \boldsymbol{\theta})$ are not known in closed form $\Rightarrow$ use linearisation of measurement equation around $E(\mathbf{x}_t|\underline{\mathbf{y}}_t^o, \boldsymbol{\theta})$.
  - udpdate using weights $p(\mathbf{y}_{t+1}^o|\mathbf{x}_t^{(i)}, \boldsymbol{\theta})$.
- This would work even in the absence of measurement error. Intuition: Figure (C)

Figure D: CPF at work, N=3

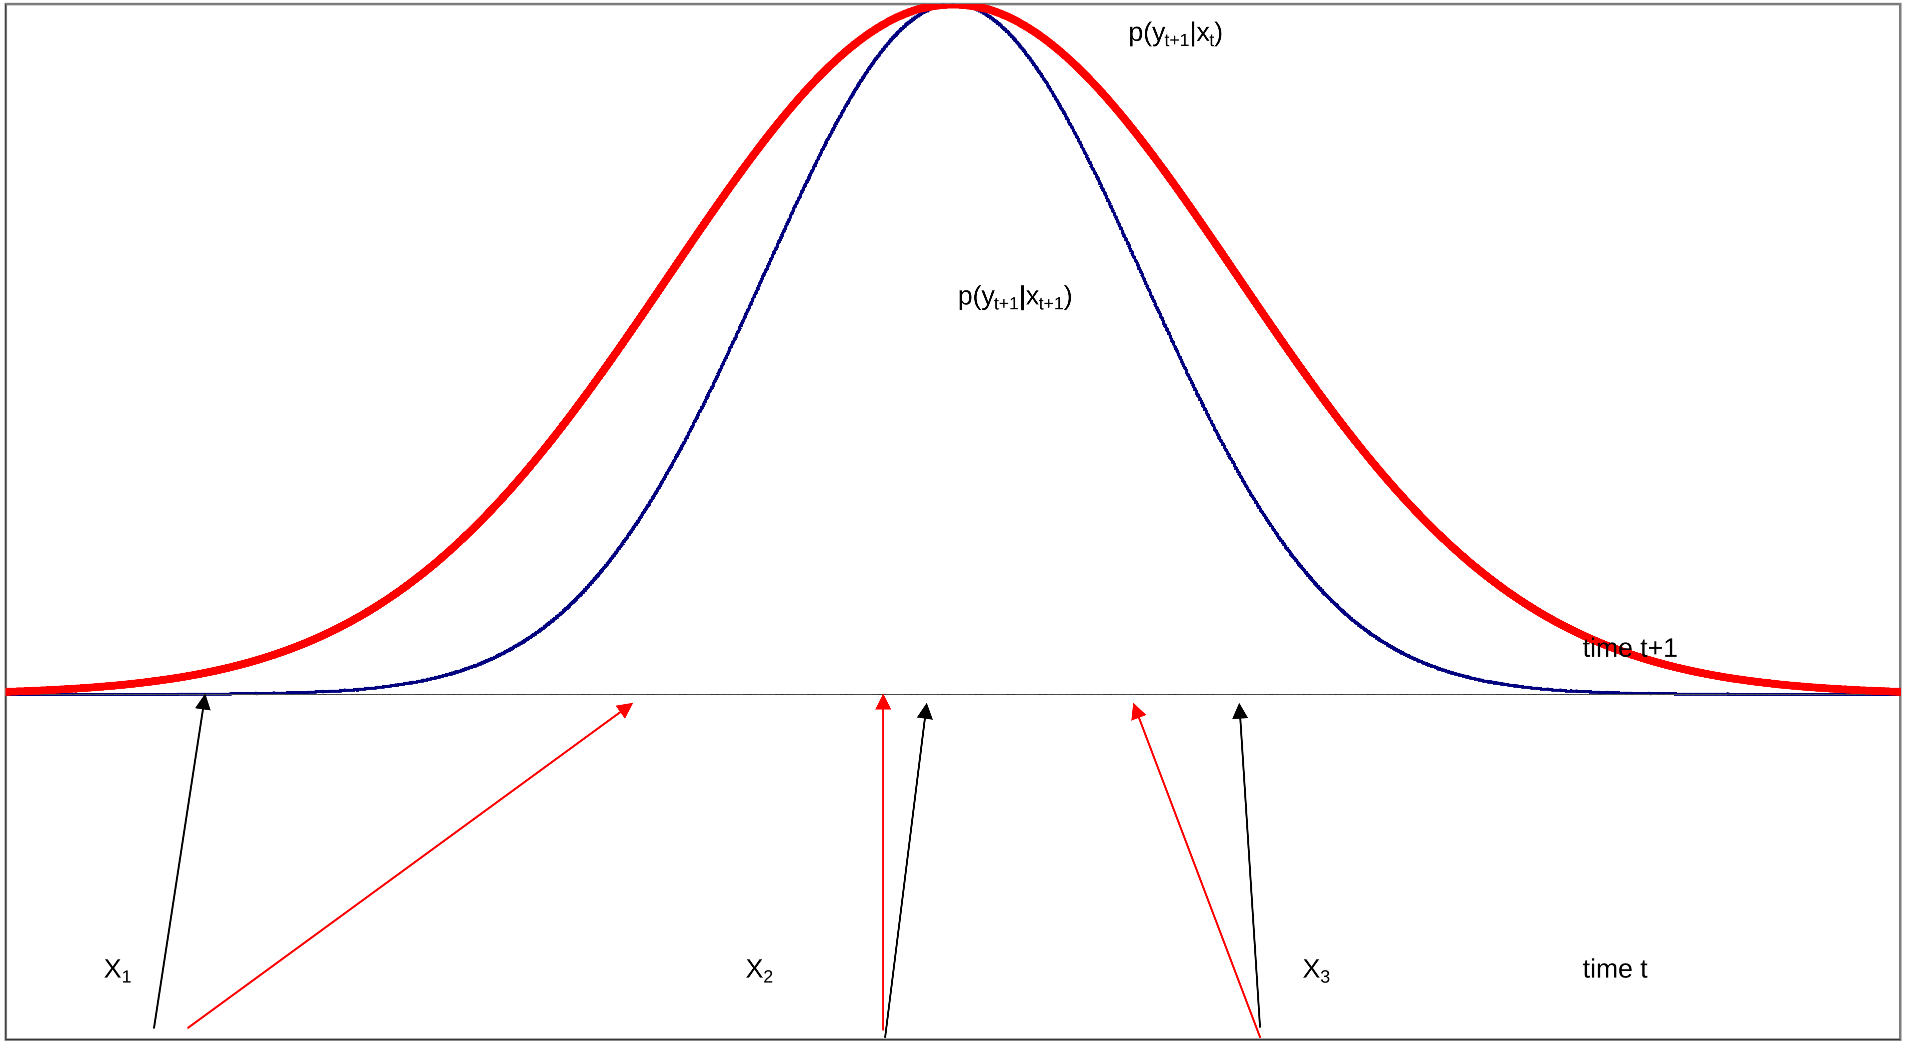

Inference

- Inference on the parameters of the model: we use MH.

- Special problems for priors of standard deviation. Use *prior predictive* approach:

  - draw parameter values from prior and compute model solution and moments of the stationary distribution of observables;

  - use pre sample information (1970s) to assess plausibility of prior and tune prior hyperparameters accordingly.

Data

- Observables: y, i, π (or y, $i - \pi$, $\Delta\pi$) taken from the Area Wide Model database (see Fagan, Henry and Mestre, 2005).

- Raw data, only output is detrended.

- The estimation period runs from 1980Q1 to 2004Q4.

Results: first vs. second order models

- Estimates of the second order model tend to be more precise, particularly for M2. Lower posterior standard deviations of most parameter estimates in the quadratic case. The distribution is much more concentrated especially for γ and, in M2, for h and ψ_π.

- Nonlinear solutions do not produce striking changes in mean parameter estimates (with exceptions: nonlinear ι higher in M1; lower γ, higher ϕ in M2). More important differences across models: e.g. ϕ much higher in M2 than in M1.

Nonlinearities

- In M2, changes in estimated inflation target are permanent. $\Rightarrow$ The SS moves and observed deviations of inflation from steady state are mostly negligible $\Rightarrow$ No nonlinear effects visible in sample.

- In M1. Estimated steady state of the inflation objective is around 2.7 percent $\Rightarrow$ at the beginning and at the end of our sample period, relatively larve deviations from the SS occur $\Rightarrow$ Nonlinear effects visible in sample.

- Impulse responses computed as deviation from the path without shocks .

Two starting point scenarios: A) state variables at their values at beginning of the sample; B) state variables at their values at the end of the sample.

Impulse responses

- In M1, starting from a high-inflation (wrt low-inflation) environment:
  - A positive inflation target shock has much more pronounced effects and the return to the mean dynamics is much slower.
  - A policy shock has less impact on inflation.
  - A technology shock has smaller, but more persistent effects on inflation.
  - A cost-push shock has a smaller impact.
- In M2, no differences are observed.

Our last word?

- Results are still preliminary: up to 20,000 particles; up to 30,0000 MCMC iterations. Robustness checks:
  - Model: nominal contracts à la Sheedy (2005); slight variants of the policy rule.
  - Simulations as in KKSS: $\widehat{x}_t = P\widehat{x}_{t-1} + \frac{1}{2}k_x + G\left(x^1_{t-1} \otimes x^1_{t-1}\right) + \sigma\varepsilon_t$.
  - Evaluation: autocovariance functions.
- Parameter estimates are quite stable, but small changes can have large effects on the reduced form coefficients.

Conclusion

- In nonlinear models, impulse responses are regime dependent. Differences visible within our sample.
- A shock to inflation (induced by a temporary change in the target) has more pronounced effects starting from a high inflation environment.
- Future directions: explore asset price implications of DSGE models.