Bounded rationality and heterogeneous expectations in macroeconomics

Massaro, D.

Chapter 4

Behavioral Heterogeneity in U.S. Inflation Data

4.1 Introduction

In recent years, pricing behavior has been described in the context of models that incorporate both nominal rigidities and optimizing agents with rational expectations.[1] One of the most popular versions of New Keynesian pricing models is derived from Calvo (1983) and it implies a forward-looking inflation equation (a "New Keynesian Phillips curve", NKPC henceforth) of the form

$$\pi_t = \delta E_t \pi_{t+1} + \gamma mc_t \ , \tag{4.1}$$

which relates inflation, π_t, to next period's expected inflation rate and to real marginal costs, mc_t.[2] An important implication of this model is that there is no intrinsic inertia in inflation, in the sense that there is no structural dependence of inflation on its own lagged values. As a result, this specification has often been criticized on the grounds that it can not account for the important empirical role played by lagged dependent variables in inflation regressions (see e.g., Rudd and Whelan (2005a,b) for a recent discussion). This critique resulted in various proposals

[1] See Clarida, Galí, and Gertler (1999) for a survey, and Woodford (2003) for a detailed treatment.

[2] Roberts (1995) shows that Eq. (4.1) can be derived from a number of different models of price rigidity.

for so-called "hybrid" variants of the NKPC, which take the form

$$\pi_t = \theta E_t \pi_{t+1} + (1-\theta)\pi_{t-1} + \gamma mc_t \ . \tag{4.2}$$

Hybrid models have been theoretically motivated in several ways. Fuhrer and Moore (1995) assume an alternative contracting specification in which workers bargain over relative real wages; Christiano, Eichenbaum, and Evans (2005) use a variant of the Calvo model in which firms that are unable to reoptimize their price instead index it to past inflation; Galí and Gertler (1999) assume the existence of a group of backward looking price setters.[3]

Starting from different theoretical formulations, a significant strand of research focused on assessing the empirical relevance of the forward looking component in Eq. (4.2), generating mixed results. Here we briefly summarize some of the evidence obtained in previous studies.

Galí and Gertler (1999) focus on estimates of θ obtained from directly fitting Eq. (4.2) using GMM. Under this procedure $E_t\pi_{t+1}$ is replaced with π_{t+1} and the model is estimated using instruments for π_{t+1}. Employing this technique, Galí and Gertler (1999) estimate large values of θ and conclude that rational forward-looking behavior plays an important role in determining U.S. inflation.

Sbordone (2005) follows Christiano, Eichenbaum, and Evans (2005) by assuming that firms that are not allowed to reset prices through the Calvo random drawing are nonetheless allowed to index their current prices to past inflation, and they do so by some fraction $\rho \in [0, 1]$. She then estimates the closed form solution of the model

$$\pi_t = \rho\pi_{t-1} + \zeta \sum_{j=0}^{\infty} \delta^j E_t mc_{t+j} \ ,$$

where inflation is a function of lagged inflation and expected future real marginal costs. Using a two-step distance estimator that exploits an auxiliary autoregressive representation of the data as in Campbell and Shiller (1987) to estimate the present value form of the inflation dynam-

[3]In the specification of Galí and Gertler (1999) the weights of lagged and expected future inflation are not constrained to sum to unity, unless the time discount factor $\delta = 1$.

ics model, Sbordone (2005) finds that the forward-looking component is quantitatively more relevant than the backward-looking component, confirming thus the results of Galí and Gertler (1999).

Lindé (2005) estimates a New-Keynesian sticky price model using a full information maximum likelihood approach and suggests a hybrid version of the NKPC where forward-looking behavior is significant but about equally or less important than backward-looking behavior.

Fuhrer (1997) considers the model developed in Fuhrer and Moore (1995), which extends the staggered contracting framework of Taylor (1980) in a way that imparts persistence to the rate of inflation. Using a maximum likelihood estimation procedure he concludes that forward looking behavior plays essentially no role in observed inflation dynamics.

Rudd and Whelan (2006) estimate closed form solutions of model (4.2) given by

$$\begin{aligned} \Delta\pi_t &= \lambda_1 \sum_{j=0}^{\infty} \lambda_2 E_t mc_{t+j} && \text{for } \theta \leq 1/2 \\ \pi_t &= \mu_1 \sum_{j=0}^{\infty} E_t mc_{t+j} + \mu_2 \pi_{t-1} && \text{for } \theta > 1/2 \ , \end{aligned}$$

where $\lambda_1, \lambda_2, \mu_1$, and μ_2 are functions of θ. Using both VAR-based methods and GMM estimation, they find no significant evidence of rational forward-looking behavior in U.S. data.

One possible explanation for the mixed evidence on the empirical relevance of rational forward-looking behavior stemming from previous tests of sticky price models, may be rooted, as put forward by Rudd and Whelan (2006), in the reliance of these models on a strict form of rational expectations (RE henceforth). Rudd and Whelan (2006) conclude that:

"...further research in this area is probably best aimed toward developing models that deviate from the standard rational expectations framework in favor of alternative descriptions of how agents process information and develop forecasts".

In this paper we take this criticism seriously and propose a model of inflation dynamics characterized by heterogeneous boundedly rational agents and evolutionary selection of forecasting strategies. Standard New Keynesian models of price setting are based on the assumptions that:

(i) prices are sticky; (ii) agents make optimal decisions given their beliefs about future inflation; (iii) individual expectations are formulated in a rational (i.e., model-consistent) way.

Empirical studies suggest that a significant degree of price stickiness is present in the U.S. economy, providing thus a rationale for firms trying to make predictions about future inflation when setting current prices. In our analysis we keep the assumption of sticky prices and optimizing behavior (given individual beliefs), so that expected future inflation has an important influence on current inflation, but we depart from standard models by assuming heterogeneous subjective expectations.

Heterogeneity in individual expectations has been abundantly documented in the literature. Fore example, Frankel and Froot (1987, 1990), Allen and Taylor (1990) and Ito (1990) find that financial experts use different forecasting strategies to predict exchange rates. More recently, Carroll (2003), Mankiw, Reis, and Wolfers (2003), Branch (2004) and Pfajfar and Santoro (2010) recently provided empirical evidence in support of heterogeneous expectations using survey data on inflation expectations, while Hommes, Sonnemans, Tuinstra, and van de Velden (2005), Adam (2007), Pfajfar and Zakelj (2010), Assenza, Heemeijer, Hommes, and Massaro (2011), and Hommes (2011) find evidence for heterogeneity in learning to forecast laboratory experiments with human subjects.

We construct a framework with monopolistic competition, staggered price setting and heterogeneous firms. Our stylized model includes two types of price setters. The first type are *fundamentalists*, or forward-looking, who believe in a present-value relationship between inflation and real marginal costs. The evolution of inflation depends on the discounted sum of expected future values of real marginal costs, and fundamentalists use a VAR methodology to forecast the evolution of the driving variable as in Campbell and Shiller (1987). The second type are *naive*, or backward-looking agents who use the simplest backward-looking rule of thumb, naive expectations (i.e., their forecast coincides with the last available observation), to forecast future inflation.

All the empirical studies on forward- versus backward-looking behavior in inflation dynamics cited above, take the distribution of heterogeneous firms as fixed and exogenous. However, recent empirical analysis suggest that this assumption is overly restrictive. For example, Zhang, Osborn, and Kim (2008), Kim and Kim (2008) and Hall, Han, and Boldea (2011) find evidence for multiple structural breaks in the relative weights of forward- and backward-looking firms, while Carroll (2003) and Mankiw, Reis, and Wolfers (2003) show that the distribution of heterogeneity evolves over time in response to economic volatility. Furthermore, Frankel and Froot (1991), Bloomfield and Hales (2002), Branch (2004), Assenza, Heemeijer, Hommes, and Massaro (2011) and Hommes (2011), among others, provide evidence that the proportions of heterogeneous forecasters evolve over time as a reaction to past forecast errors using survey data as well as experimental data. On the basis of this empirical evidence, we endogenize the evolution of the distribution of heterogeneous firms by assuming that agents can switch between different forecasting regimes, depending on recent prediction performance of their forecasting rules as in Brock and Hommes (1997).

Obviously we are not the first to introduce a dynamic predictor selection mechanism in macroeconomic models. Recent theoretical papers analyzing inflation dynamics under endogenous selection of expectation rules include, among others, Brock and de Fontnouvelle (2000), Tuinstra and Wagener (2007), Anufriev, Assenza, Hommes, and Massaro (2008), Brazier, Harrison, King, and Yates (2008), Branch and McGough (2010), De Grauwe (2010), Branch and Evans (2010).

The main novelty of our paper consists in the estimation of a NKPC with heterogeneous expectations and endogenous switching between different beliefs using U.S. macroeconomic data. To our knowledge, there are only a few empirical applications that attempt to estimate heterogenous agents models with fully-fledged switching mechanism. Those attempts consider the S&P500 market index (Boswijk, Hommes, and Manzan (2007)), commodity markets (Reitz and Westerhoff (2005, 2007)), the Asian equity markets (De Jong, Verschoor, and Zwinkels (2009)), the DAX30 index options (Frijns, Lehnert, and Zwinkels (2010)), and the U.S. housing market (Kouwenberg and Zwinkels (2010)).

Moreover, our paper contributes to the debate about the empirical relevance of forward-looking behavior in inflation dynamics. In fact our model, although with a different behavioral interpretation, is similar to the hybrid models estimated by Sbordone (2005) and Rudd and Whelan (2006) among others. Our measure of fundamental expectation is constructed in the same way as Sbordone (2005) and Rudd and Whelan (2006) obtain their estimation of the discounted sum of expected future values of real marginal costs in the closed-form solution of the model, i.e., using the Campbell and Shiller methodology, while the expectations of naive firms account for lagged value of inflation in the hybrid specification of the NKPC. The main difference stems from the time-varying weights assigned to fundamentalists and naive price setters, evolving over time according to past forecasting performances.

The results of our analysis provide empirical evidence for behavioral heterogeneity in U.S. inflation dynamics. Moreover, the data support the hypothesis of an endogenous mechanism relating predictors choice to their forecasting performance. In fact, our results suggest that the degree of heterogeneity varies considerably over time, and that the economy can be dominated temporarily by either forward-looking or backward-looking behavior.

These findings have important implications for monetary policy. Standard policy recommendations based on determinacy under RE, may not be a robust criterion for policy advices in the presence of heterogeneous expectations. In fact, recent papers have shown that multiple equilibria, periodic orbits and complex dynamics can arise in presence of dynamic predictor selection, even if the model under RE has a unique stationary solution (see Anufriev, Assenza, Hommes, and Massaro (2008), Branch and McGough (2010), and De Grauwe (2010) among others).

The paper is structured as follows. Section 4.2 derives a NKPC with heterogenous expectations and endogenous switching dynamics. Section 4.3 presents the estimation results and describes the fit of the model. Section 4.4 discusses the robustness of the empirical results to alternative forecasting models for the driving variable in the NKPC and to alternative measures of real marginal costs. Section 4.5 concludes.

4.2 The model

This section derives a NKPC with heterogeneous, potentially nonrational expectations and endogenous switching between forecasting strategies.

4.2.1 The NKPC with heterogeneous expectations

We consider a model with a continuum of monopolistic firms, indexed by i, which produce differentiated goods. The demand curve for the product of firm i takes the form:

$$Y_{i,t} = (P_{i,t}/P_t)^{-\eta} Y_t,$$

where η is the Dixit-Stiglitz elasticity of substitution among differentiated goods, Y_t is the aggregator function defined as $Y_t = [\int_0^1 Y_{i,t}^{(\eta-1)/\eta} di]^{\eta/(\eta-1)}$, and P_t is the aggregate price level defined as $P_t = [\int_0^1 P_{i,t}^{1-\eta} di]^{1/(1-\eta)}$. Nominal price rigidity is modeled by allowing, in every period, only a fraction $(1-\omega)$ of the firms to set a new price along the lines of Calvo (1983). Firms that reset prices maximize expected discounted profits, which are given by

$$\max_{P_{i,t}} E_t^i \sum_{j=0}^{\infty} \omega^j Q_{t,t+j} \left(\frac{P_{i,t}}{P_{t+j}} - MC_{t+j} \right) \left(\frac{P_{i,t}}{P_{t+j}} \right)^{-\eta} Y_{t+j},$$

where E_t^i denotes the subjective expectation of firm i, $Q_{t,t+j}$ is the stochastic discount factor and MC_t are real marginal costs of production. Linearizing the first order conditions of this problem around a zero inflation steady state delivers

$$p_{i,t} = (1-\omega\delta) E_t^i \sum_{j=0}^{\infty} (\omega\delta)^j mc_{t+j} + \omega\delta E_t^i \sum_{j=0}^{\infty} (\omega\delta)^j \pi_{t+j+1}, \tag{4.3}$$

where δ is the time discount factor, $\pi_t = p_t - p_{t-1}$, and lower case letter denote log-deviations from steady state.

Optimal pricing decisions involve subjective forecasts of marginal costs and inflation, hence we will have that firms with different expectations will set different prices. The average price set by optimizing firms is given by

$$p_t^* = \sum_i n_{i,t} p_{i,t},$$

where $n_{i,t}$ denotes the fraction of firms characterized by the subjective expectation operator E_t^i at time t. Given the assumed staggered price setting mechanism, the aggregate price level evolves as a convex combination of the lagged price level p_{t-1} and the average optimal reset price p_t^* as follows

$$p_t = \omega p_{t-1} + (1-\omega) p_t^*. \tag{4.4}$$

Under the assumption of a representative firm with rational expectations, Eqs. (4.3) and (4.4) can be used to derive the standard NKPC

$$\pi_t = \delta E_t \pi_{t+1} + \gamma mc_t \ , \tag{4.5}$$

where $\gamma \equiv (1-\omega)(1-\delta\omega)\omega^{-1}$. Deriving an equation for inflation similar to Eq. (4.5) is not entirely obvious when expectations are heterogeneous. Optimal pricing decisions depend on forecasted marginal costs and inflation. Inflation is determined by other agent's pricing decisions and their forecasts. As a result, optimal price setting would require forecasting the marginal costs and inflation forecasts of others, see e.g., Woodford (2001a). Denote the average forecast of individuals as $\overline{E}_t = \sum_i n_{i,t} E_t^i$ and suppose the following holds:

Condition 1

$\overline{E}_t \overline{E}_{t+1} x_{t+j} = \overline{E}_t x_{t+j}$ for $j \geq 0$ and $x \in \{mc, \pi\}$.

Condition 1, also known as the "tower property" in probability theory, requires that the average forecast at time t of the future average forecast of a certain variable x at time $t+1$ is equal to the average forecast of variable x at time t, that is, agents, on average, do not expect predictable changes in average future expectations of a certain macroeconomic variable x. Stated differently, Condition 1 corresponds to the law of iterated expectations at the aggregate level.

Appendix 4.A shows that whenever Condition 1 is satisfied, it is possible to aggregate individual pricing rules (4.3) and obtain a Phillips curve of the form

$$\pi_t = \delta \overline{E}_t \pi_{t+1} + \gamma mc_t \ . \tag{4.6}$$

In what follows we will assume that Condition 1 is satisfied and we will use a NKPC equation as in (4.6) for our empirical investigation. Condition 1 is similar to the assumptions made by Branch and McGough (2009) and Adam and Padula (2011) in order to derive a NKPC of the form of (4.6) in a context of subjective heterogeneous expectations.

4.2.2 Evolutionary selection of expectations

We assume that agents can form expectations by choosing from I different forecasting rules, and we denote by $E_t^i \pi_{t+1}$ the forecast of inflation by rule i. The fraction of individuals using the forecasting rule i at time t is denoted by $n_{i,t}$. Fractions are updated in every period according to an *evolutionary fitness* measure. At the beginning of every period t agents compare the realized relative performances of the different strategies and the fractions $n_{i,t}$ evolve according to a discrete choice model with multinomial logit probabilities (see Manski and McFadden (1981) for details), that is

$$n_{i,t} = \frac{\exp(\beta U_{i,t-1})}{\sum_{i=1}^{I} \exp(\beta U_{i,t-1})}, \tag{4.7}$$

where $U_{i,t-1}$ is the realized fitness metric of predictor i at time $t-1$, and the parameter $\beta \geq 0$ refers to the *intensity of choice* and it reflects the sensitivity of the mass of agents to selecting the optimal prediction strategy. Brock and Hommes (1997) proposed this model for endogenous selection of expectation rules. The key feature of Eq. (4.7) is that strategies with higher fitness in the recent past attract more followers. The case $\beta = 0$ corresponds to the situation in which differences in fitness can not be observed, so agents do not switch between strategies and all fractions are constant and equal to $1/I$. The case $\beta = \infty$ corresponds to the "neoclassical" limit in which the fitness can be observed perfectly and in every period all agents choose the best predictor.

A strong motivation for switching among forecasting rules can be found in empirical works on individual expectations. Frankel and Froot (1991) find that professional market participants in the foreign exchange markets expect recent price changes to continue in the short term, while they expect mean reversion to fundamental value in the long term. Moreover, Frankel and Froot (1991) report survey evidence showing that professional forecasting services in the foreign exchange markets rely both on technical analysis and fundamental models, but with changing weights over time, and the weights appear to depend strongly on recent forecasting performances. Branch (2004) finds evidence for dynamic switching between alternative forecasting strategies that depends on the relative mean squared errors of the predictors using survey data on inflation expectations. In addition, Bloomfield and Hales (2002), Assenza, Heemeijer, Hommes, and Massaro (2011) and Hommes (2011) document experimental evidence that participants switch between forecasting regimes conditional on recent forecasting performances.

4.2.3 A simple two-type example

We assume that agents can choose between two forecasting rules to predict inflation, namely fundamentalist and naive. The first rule, fundamentalist, is based on a present-value description of the inflation process. When all agents have rational expectations, repeated application of equation Eq. (4.6) gives

$$\pi_t = \gamma \sum_{k=0}^{\infty} \delta^k E_t mc_{t+k} \ . \tag{4.8}$$

We refer to (4.8) as the *fundamental inflation*. Fundamentalists use the expression (4.8) to forecast future inflation. In particular, leading (4.8) one-period ahead we get

$$\pi_{t+1} = \gamma \sum_{k=1}^{\infty} \delta^{k-1} E^f_{t+1} mc_{t+k} \ , \tag{4.9}$$

where E^f_{t+1} denotes fundamentalists forecast. Applying the expectation operator E^f_t on both sides we get

$$E^f_t \pi_{t+1} = \gamma \sum_{k=1}^{\infty} \delta^{k-1} E^f_t mc_{t+k} \ . \tag{4.10}$$

From a behavioral point of view, fundamentalists can be considered as agents who believe in rational expectations and use the closed form solution of the model to forecast the inflation path. If Eq. (4.6) were the true data generating process and if all agents in the economy were of the same type, then the inflation path implied by fundamental expectations would coincide with the inflation path under rational model-consistent expectations.[4]

In order to characterize the fundamental forecast (4.10) we use the VAR methodology of Campbell and Shiller (1987). Assuming that the forcing variable mc_t is the first variable in the multivariate VAR

$$Z_t = AZ_{t-1} + \epsilon_t, \tag{4.11}$$

we can rewrite the sum of discounted future expectations of marginal costs (4.10) as

$$E^f_t \pi_{t+1} = \gamma \sum_{k=1}^{\infty} \delta^{k-1} E^f_t mc_{t+k} = \gamma e_1' (I - \delta A)^{-1} A Z_t,$$

where e_1' is a suitably defined unit vector.[5]

The second rule, which we call naive, takes advantage of inflation persistence and uses a simple backward-looking forecasting strategy:

$$E^n_t \pi_{t+1} = \pi_{t-1} \ . \tag{4.12}$$

Notice however that, although being an extremely simple rule, the naive forecasting strategy is optimal when the stochastic process is a random walk; hence for a near unit process, as in the case of inflation, naive expectations are almost optimal.

[4] In fact, substituting the fundamental forecast in Eq. (4.6), we get $\pi_t = \delta\gamma \sum_{k=1}^{\infty} \delta^{k-1} E^f_t mc_{t+k} + \gamma mc_t = \gamma \sum_{k=0}^{\infty} \delta^k E^f_t mc_{t+k}$, which corresponds to the inflation path implied by Eq. (4.8) when the discounted sums of current and future expected marginal costs are estimated in the same way.

[5] Technically, because the discounted sum of real marginal costs starts at $k = 1$, we measure it using $(I - \delta A)^{-1} A Z_t$ instead of $(I - \delta A)^{-1} Z_t$.

The specific choice of the set of forecasting rules, namely fundamental and naive, will enable us to compare the outcome of our analysis with the results of previous empirical works based on the hybrid Phillips curve specification. In fact, our measure of fundamental expectation is similar to the measure used by Sbordone (2005) and Rudd and Whelan (2006), among others, to estimate the discounted sum of expected marginal costs when fitting the closed-form solution of the hybrid model (4.2) to U.S. data, while the backward-looking component introduced in different ways in RE models is accounted for by the expectations of rule of thumb firms. The main difference between traditional hybrid specifications of the NKPC and our model is the fact that the weights assigned to forward-looking and backward-looking component are endogenously varying over time. We in fact assume that agents can switch between the two predictors based on recent forecasting performance. Defining the absolute forecast error as

$$FE_t^i = |E_{t-1}^i \pi_t - \pi_t| \ ,$$

with $i = f, n$, we can then define the evolutionary fitness measure as

$$U_{i,t} = -\frac{FE_t^i}{\sum_{i=1}^{I} FE_t^i} \ . \tag{4.13}$$

The evolution of the weights of different heuristics is then given by Eq. (4.7).
Denoting the fraction of fundamentalists as $n_{f,t}$ we can summarize the full model as

$$\pi_t = \delta(n_{f,t} E_t^f \pi_{t+1} + (1 - n_{f,t}) E_t^n \pi_{t+1}) + \gamma mc_t + u_t \ , \tag{4.14}$$

where

$$\begin{aligned}
E_t^f \pi_{t+1} &= \gamma e_1' (I - \delta A)^{-1} A Z_t \\
E_t^n \pi_{t+1} &= \pi_{t-1} \\
n_{f,t} &= \frac{1}{1 + \exp\left(\beta \left(\frac{FE_{t-1}^f - FE_{t-1}^n}{FE_{t-1}^f + FE_{t-1}^n}\right)\right)} \\
FE_{t-1}^i &= |E_{t-2}^i \pi_{t-1} - \pi_{t-1}|, \qquad \text{with} \ \ i = f, n \ .
\end{aligned}$$

4.3 Estimation results

This section describes data and methodology used to estimate the nonlinear switching model derived in the previous section.

4.3.1 Data description

We use quarterly U.S. data on the inflation rate, the output gap, unit labor costs, the labor share of income, hours of work and consumption-output ratio, from 1960:Q1 to 2010:Q4. Inflation is measured as log-difference of CPI. Output gap is measured as quadratically detrended log-real GDP. We use unit labor costs, labor share of income, detrended hours of work and detrended consumption-output ratio time series for nonfarm business sector in the construction of the VAR model (4.11). A more detailed description of data sources and variables definition is given in Appendix 4.B.

4.3.2 The fit of the model

In this section we discuss the empirical implementation of model (4.14). In the "baseline" specification (the one used in the results reported below) we use the output gap as driving variable, and we use a four-lag, two-variables VAR in output gap and labor share of income to compute fundamental expectations.[6] The selection procedure of the VAR specification is extensively documented in Appendix 4.C. Standard unit root tests reveal that the labor share of income is an I(1) process, therefore the VAR model includes the rate of change of the labor share of income.[7] Denoting by Y_t the vector of dependent variables, $Y_t = [y_t,\ \Delta lsi_t]'$, the vector Z_t is defined as $Z_t = [Y_t,\ Y_{t-1},\ Y_{t-2},\ Y_{t-3}]'$. Although being parsimonious, our VAR specification captures about 94% of output gap volatility (see Table 4.1). The parameters of the matrix A are estimated by OLS, and the discount factor δ is fixed to the standard value

[6] Section 4.4 discusses the sensitivity of the results to the use of labor share as driving variable and to alternative specifications of the forecasting VAR.

[7] Observations in period 1959:4 were taken as initial conditions to build the series in first differences.

0.99. Model (4.14) is then estimated using non-linear least squares (NLS). Table 4.1 presents the results and diagnostic checks are reported in Appendix 4.C.

Table 4.1: NLS estimates of model (4.14)

Parameter	β	γ
Estimate	3.975^{***}	0.005^{**}
Std. error	1.100	0.002
R^2 from Inflation Equation	0.767	
R^2 from Output Gap VAR Equation	0.945	

Notes: Standard errors are computed using White's heteroskedasticity-consistent covariance matrix estimator (HC-CME). *, **, *** denote significance at the 10%, 5%, and 1% level.

All coefficients have the correct sign and are significant at least at the 5% level. The positive sign and the significance of the intensity of choice parameter β implies that agents switch towards the better performing forecasting rule, based on its past performance.[8] The positive sign and the significance of parameter γ is a rather interesting result. It has been quite difficult to obtain parameter estimates with the correct sign and of a plausible magnitude when output gap is used as driving variable for the inflation process. Fuhrer and Moore (1995) and Galí and Gertler (1999), for example, find a negative and insignificant estimate of γ when real marginal costs are approximated by detrended output. The results in Table 4.1 show that taking into account non-rational heterogeneous expectations seems to help establishing a plausible link between output and inflation dynamics via the NKPC. Interestingly, Adam and Padula (2011) reach the same conclusion by estimating a NKPC using data from the Survey of Professional Forecasters as proxy for expected inflation.

The series of inflation predicted by (4.14) for the estimated values of β and γ is plotted in Fig. 4.1, as dashed line, versus the actual series (solid line).[9]

Overall the predicted inflation path tracks the behavior of actual inflation quite well (the R^2 from inflation equation (4.14) is about 0.77, see Table 4.1).

Our results are, in some respects, similar to findings obtained in previous empirical works.

[8]The order of magnitude of β is more difficult to interpret as it is conditional on the functional form of the performance measure U.

[9]The graphs are in deviation from the mean.

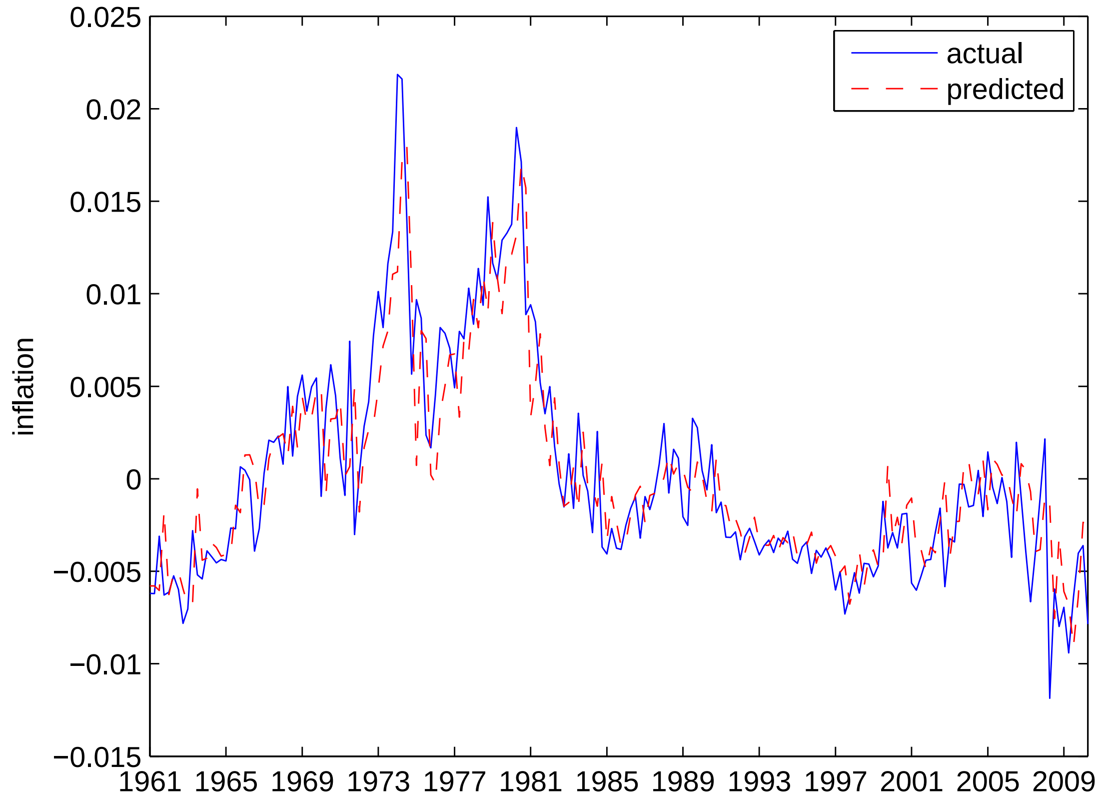

Figure 4.1: Actual vs. predicted inflation

In particular, Galí and Gertler (1999) and Sbordone (2005) find that models derived from the assumption of heterogeneous price setting behavior are capable of fitting the level of inflation quite well. However, Rudd and Whelan (2005a) and Rudd and Whelan (2006) show that this good fit reflects the substantial role that these models still allow for lagged inflation, and that forward-looking components play no discernable empirical role in determining inflation.

Our NKPC specification allows for time-varying weights assigned to fundamentalists and naive price setters. Having estimated model (4.14), we are now ready to assess the relative importance over time of forward-looking versus backward-looking components in inflation dynamics. Table 4.2 displays descriptive statistics of the weight of the forward-looking component n_f.

Table 4.2: Descriptive statistics of weight n_f

Mean	0.338
Median	0.190
Maximum	0.976
Minimum	0.019
Std. Dev.	0.322
Skewness	0.774
Kurtosis	2.078
Auto-corr. Q(-1)	0.383

On average, the majority of agents use the simple backward-looking rule (with mean fraction $1 - 0.34 = 0.66$). However, the spread between the minimum and the maximum indicates that the market can be dominated by either forward-looking or backward-looking agents. Moreover the autocorrelation of the series n_f, about 0.38, indicates that agents do not change quickly their strategy, suggesting a certain degree of inertia in the updating process.

Fig. 4.2 shows the time series of the fraction of fundamentalists, i.e., the forward-looking component in our NKPC specification, the time series of the distance of actual inflation from the fundamental solution, and a scatter plot of the fraction of fundamentalists against the relative forecast error of the naive rule.

It is clear that the fraction of fundamentalists varies considerably over time with periods in which it is close to 0.5 and other periods in which it is close to either one of the extremes 0 or 1. For example, immediately after the oil crisis of 1973, the proportion of fundamentalists drops almost to 0. Soon after the difference between inflation and fundamental value reaches its peak in 1974:Q3/1974:Q4, the estimated weight of the forward-looking component shoots back up to about 0.8. During the second oil crisis inflation was far above the fundamental, causing more and more agents to adopt a simple backward-looking rule to forecast inflation. Fundamentalists dominated the economy in the late-80s, while from 1992 until 2003, inflation stayed continuously well below the fundamental, causing the weight of fundamentalists to fall. From 2004 until the early stages of the recent global financial crisis, the proportion of fundamentalists stayed, on average, around 0.5, reaching peaks of about 0.8. In the aftermath of the crisis we observe that n_f declines with the fundamentalist rule losing its forecast accuracy when actual inflation falls below its fundamental value in the last quarter of 2008, and then it increases again to around 0.7 in the last two quarters of 2010.

The bottom panel of Fig. 4.2 presents a scatter plot of the relative forecast error of the naive rule, $(FE^n - FE^f)/(FE^n + FE^f)$, versus the fraction of fundamentalist agents, n_f. Due to the positive estimated value of β this line slopes upwards, such that a more accurate fundamentalist forecast results in a higher weight n_f. The S-shape is induced by the logit function in Eq. (4.7).

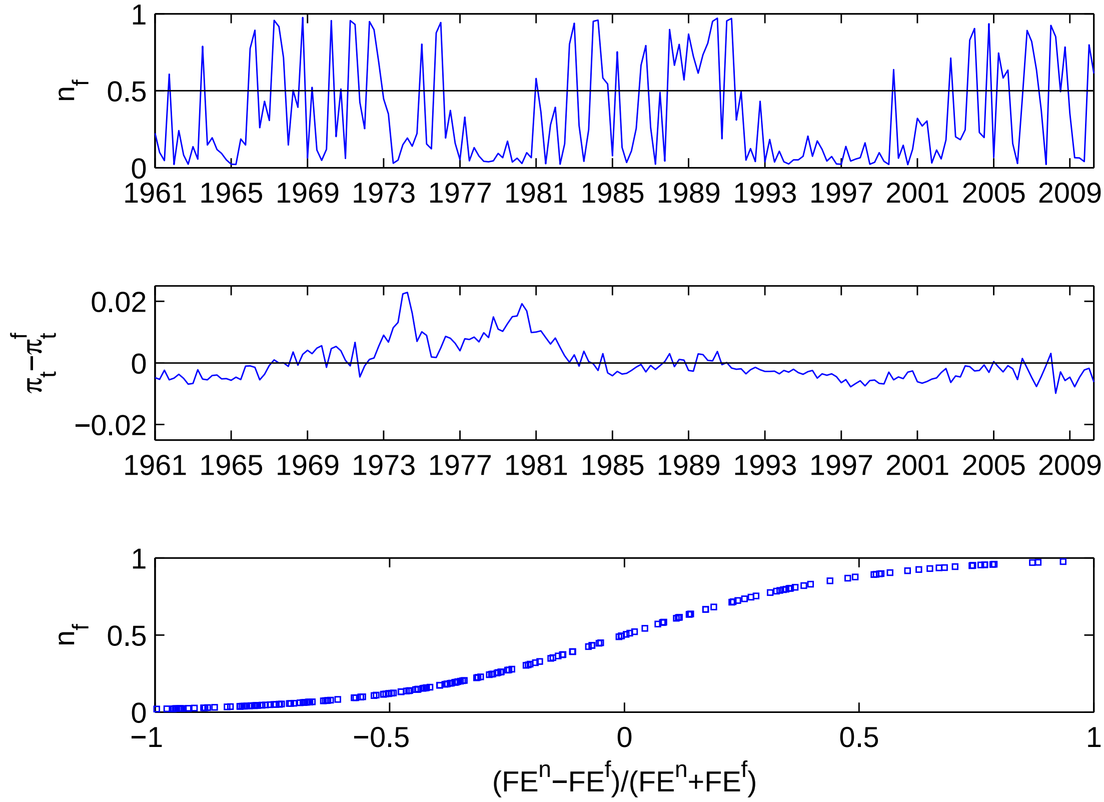

Figure 4.2: **Top panel:** Time series of the fraction of fundamentalists $n_{f,t}$. **Middle panel:** Distance between actual inflation and fundamental solution. **Bottom panel:** Scatter plot of the weight $n_{f,t}$ versus the relative forecast error of the naive rule.

The analysis conducted in this section shows that the evolutionary switching model fits the data quite nicely. The positive sign and the significance of the intensity of choice parameter, β, implies that the endogenous mechanism that relates predictors choice to their past performance is supported by the data. We also find that the ability of the discounted sum of expected future output gap values to predict the empirical inflation process varies considerably over time. In fact the spread between the minimum and the maximum value of n_f, i.e., the fraction of fundamentalists, shows that the economy can be dominated by either forward-looking or backward-looking behavior. Moreover, even though the market is, on average, dominated by agents using a simple heuristic to predict inflation, fundamentalists, or forward-looking components, still have a significant impact on inflation dynamics.

4.3.3 Specification tests and out-of-sample forecasting

In order to assess the validity of our baseline model, which will be denoted by H_1 for the purposes of this section, we test it against four alternative specifications: a model with heterogeneous agents and exogenous estimated fixed weights ($n_{f,t} \equiv \widehat{n}_f$), which is similar to the model estimated by Rudd and Whelan (2006) and Sbordone (2005), and it is denoted by H_2; a static model with heterogeneous agents in which we let $\beta = 0$ ($n_{f,t} \equiv 0.5$), which corresponds to the model of Fuhrer and Moore (1995), denoted by H_3; a model with homogeneous fundamentalists agents ($n_{f,t} \equiv 1$), which corresponds to the RE closed form solution of the standard NKPC without backward looking component, denoted by H_4; a model with homogeneous naive agents ($n_{f,t} \equiv 0$), which recalls the old backward-looking Phillips curve and it is denoted by H_5. Given that, with the exception of model H_3 which obtains by setting $\beta = 0$, the competing models are nonnested, we will use nonnested hypothesis testing procedures. In particular, we construct a test for the adequacy of our nonlinear specification with endogenous switching in explaining inflation dynamics (null hypothesis) against the alternative specifications mentioned above. Nonnested hypotheses tests are appropriate when rival hypotheses are advanced for the explanation of the same economic phenomenon. We will follow the procedure described in Davidson and MacKinnon (2009) and compute a heteroskedasticity-robust P test of H_1, based on the following Gauss-Newton regression:

$$\widehat{\mathbf{u}} = \widehat{\mathbf{X}}\mathbf{b} + \widehat{\mathbf{Z}}\mathbf{c} + \text{residuals} \;, \tag{4.15}$$

where $\widehat{\mathbf{u}}$ are residuals from the NLS estimates of model H_1, $\widehat{\mathbf{X}} \equiv \mathbf{X}(\widehat{\theta})$ is the matrix of derivatives of the nonlinear regression function corresponding to the dynamics implied by model H_1, evaluated at the NLS estimates of the model, $\widehat{\theta} = [\widehat{\beta}, \widehat{\gamma}]$, and $\widehat{\mathbf{Z}}$ is a matrix collecting the differences between the fitted values from H_2, H_3, H_4, H_5, and the fitted values from H_1. The P test is based on the joint significance of the regressors in $\widehat{\mathbf{Z}}$, i.e., $\mathbf{c} = 0$. We report the results of the test in Table 4.3, and refer the reader to Davidson and MacKinnon (2009), p. 284, for details on the construction of the heteroskedasticity-robust test. The results of the nonnested

Table 4.3: Nonnested hypotheses test

Null hypothesis: $\mathbf{c} = 0$			
test-statistic	5.359	Prob. $\chi^2(4)$	0.252

hypotheses test suggest that there is not a statistically significant evidence of departure from the null hypothesis (adequacy of nonlinear switching model H_1) in direction of alternative explanation of inflation dynamics (models H_2, H_3, H_4, and H_5). Stated differently, the test does not provide evidence in favor of alternative specifications of statistical models when tested against our baseline nonlinear switching model.[10]

As an additional test of the validity of our model with heterogeneous agents, we contrast its forecasting accuracy with the four alternative models. All models are initially estimated over the restricted sample 1960:Q1 - 2001:Q4, and evaluated over the out-of-sample period 2002:Q1 - 2010:Q4. Forecasts are created using an expanding window, meaning that each model is first estimated over the sample 1960:Q1 - 2001:Q4. Subsequently, inflation is forecasted up to one year ahead depending on the forecast horizon, which we vary from 1 to 4 quarters. The models

[10]For completeness, we also compared the switching model to the nested static model without switching ($\beta = 0$) using a likelihood ratio test. We rejected the null of a restricted static model at the 1% level on the basis of the test statistic $2\Delta LL = 91.80^{***}$, where ΔLL denotes the log-likelihood difference.

are then re-estimated with one extra observation and a new set of forecast is generated. We repeat this process to generate a number of 33 out-of-sample forecasts per horizon. The comparison of forecasting accuracy is assessed using the ratio of the average forecasting accuracy of our benchmark switching model over the average forecasting accuracy of the alternative models. A ratio less than one implies better performance for the benchmark model. Forecasting performance is measured using the mean absolute error and the mean squared error. Table 4.4 presents the forecast performance ratios. The results in Table 4.4 show that, overall, the fore-

Table 4.4: Comparison of out-of-sample forecast errors

Horizon	MAE				MSE			
q	$n_{f,t} \equiv \widehat{n}_f$	$n_{f,t} \equiv 0.5$	$n_{f,t} \equiv 1$	$n_{f,t} \equiv 0$	$n_{f,t} \equiv \widehat{n}_f$	$n_{f,t} \equiv 0.5$	$n_{f,t} \equiv 1$	$n_{f,t} \equiv 0$
1	0.963	0.973	0.679	0.914	0.943	0.975	0.545	0.868
2	1.075	0.968	0.805	0.971	1.168	1.000	0.712	0.983
3	1.096	0.975	0.887	0.957	1.223	0.978	0.805	0.945
4	1.230	0.918	0.873	1.017	1.639	0.942	0.806	1.245

Notes: Ratios of the forecast error of the model with heterogeneous agents and endogenous switching over the forecast error of the model with heterogeneous agents and constant estimated weights ($n_{f,t} \equiv \widehat{n}_f$), the model with constant equal weights ($n_{f,t} \equiv 0.5$), the model with a homogeneous fundamentalists agents ($n_{f,t} \equiv 1$), and the model with homogeneous naive agents ($n_{f,t} \equiv 0$). MAE is mean absolute error and MSE is mean squared error.

casts of the model with heterogeneous beliefs and evolutionary switching are more accurate than most of the existing models in the literature, with the exception of the model with constant estimated fractions, for which we report a ratio well above one at a horizon of four quarters.

4.4 Robustness analysis

The empirical analysis that we presented is conditional upon the assumptions that output gap is well forecast by our baseline VAR specification, and that the output gap itself is a good approximation to real marginal costs. In this section we address the issue of how sensitive our results are to alternative specifications of the VAR forecasting model, and to different measures of real marginal costs.

4.4.1 Robustness to the specification of the VAR model

In order to choose the baseline forecasting system, we started from a broad model that recalls the baseline specifications of previous empirical works (see, e.g., Woodford (2001b) and Rudd and Whelan (2005a)) and then restricted the number of variables to include in our VAR as documented in Appendix 4.C. However, one doesn't necessarily have to exclude from the information set other variables that may help forecasting the output gap beyond the contribution of the rate of change of the labor share of income. Therefore, to investigate how sensitive our results are to the specification of the fundamentalists' forecasting system, we augmented our baseline VAR model by including hours of work and consumption-output ratio.[11] These variables have been used in the VAR specifications considered by Rudd and Whelan (2005a) and Sbordone (2002). Table 4.5 reports results from alternative VAR forecasting models for the output gap.

Table 4.5: Estimation results using alternative VAR for output gap

	VAR specification			
	$\begin{bmatrix} y_t \\ \Delta lsi_t \end{bmatrix}$	$\begin{bmatrix} y_t \\ \Delta lsi_t \\ h_t \end{bmatrix}$	$\begin{bmatrix} y_t \\ \Delta lsi_t \\ c_t/y_t \end{bmatrix}$	$\begin{bmatrix} y_t \\ \Delta lsi_t \\ h_t \\ c_t/y_t \end{bmatrix}$
β	3.975***	4.037***	4.110***	4.265***
	(1.100)	(1.098)	(1.123)	(1.143)
γ	0.005**	0.005	0.005**	0.006***
	(0.002)	(0.003)	(0.002)	(0.002)
R^2 from Inflation Equation				
	0.767	0.765	0.769	0.769
R^2 from Output Gap VAR Equation				
	0.943	0.942	0.951	0.952

Notes: $y_t \equiv$ output gap, $\Delta lsi_t \equiv$ labor share growth, $h_t \equiv$ detrended hours of work, $c_t/y_t \equiv$ detrended consumption-output ratio. Optimal lag length in VAR specifications: $l_1 = l_2 = l_3 = l_4 = 4$. Standard errors are computed using White's heteroskedasticity-consistent covariance matrix estimator (HCCME). *, **, *** denote significance at the 10%, 5%, and 1% level.

As Table 4.5 shows, the estimates presented in section 4.3.2 are robust to alternative VAR

[11] Hours are quadratically detrended total hours of work in the non-farm business sector, while consumption-output ratio is linearly detrended.

specifications for the output gap. The alternative models provide a good description of the empirical output gap process and the point estimates of coefficients β and γ do not substantially change. The only relevant difference consists in the insignificance of the coefficient γ in the model of column 2, Table 4.5, although the t-statistic ($t = 1.52$) is well above one.

4.4.2 Robustness to alternative measures of marginal costs

Our benchmark model considers a traditional output gap measure, defined as the deviation of log real GDP from a quadratic trend, as the driving variable in the inflation process. Previous tests of sticky-price models under RE have reported that the NKPC provides a poor description of the actual inflation process when output gap is used as a proxy for real marginal costs (see, e.g., Fuhrer and Moore (1995) and Rudd and Whelan (2005a, 2006) among others). As an alternative to the standard approach, a number of researchers have suggested using the labor's share of income as driving variable in the NKPC. The motivation for this measure stems from the fact that the micro-foundations underpinning the NKPC imply that the correct driving variable for inflation is actually real marginal cost. Some theoretical restrictions are then required in order for real marginal costs to move with the output gap. Using average unit labor costs (nominal compensation divided by real output) as a proxy for nominal marginal cost results in the labor share of income (nominal compensation divided by nominal output) as a proxy for real marginal cost. Even though empirical implementations of this variant of the NKPC generated mixed evidence,[12] we estimate, as a second robustness exercise, the evolutionary switching model using the (log of) labor's share of income as driving variable. As noted in section 4.3, the labor share process presents a unit root when considered over the full sample 1960:Q1-2010:Q4. In order to avoid spurious correlations and facilitate comparison with earlier works, we restrict the estimation sample to 1960:Q1-2001:Q4.[13] The estimation results reported in Table 4.6 show that the estimated coefficients are significant and have the correct sign. Moreover, the point

[12]Galí and Gertler (1999), Woodford (2001b), Sbordone (2002) and others report that predicted inflation series based on labor share fit actual inflation well, while Rudd and Whelan (2005a,b, 2006) and others show that even the labor share version of the model provides a poor description of the inflation process.

[13]Standard unit root tests motivate this choice.

estimates are of the same order of magnitude as in the output gap VAR specification.

Table 4.6: Estimation results using alternative VAR for labor share of income

	VAR specification			
	$\begin{bmatrix} lsi_t \\ y_t \end{bmatrix}$	$\begin{bmatrix} lsi_t \\ y_t \\ h_t \end{bmatrix}$	$\begin{bmatrix} lsi_t \\ y_t \\ c_t/y_t \end{bmatrix}$	$\begin{bmatrix} lsi_t \\ y_t \\ h_t \\ c_t/y_t \end{bmatrix}$
β	3.755^{***}	3.689^{***}	4.037^{***}	3.930^{***}
	(1.033)	(1.069)	(1.140)	(1.138)
γ	0.009^{***}	0.007^{***}	0.010^{***}	0.008^{***}
	(0.002)	(0.001)	(0.002)	(0.002)
R^2 from Inflation Equation				
	0.808	0.811	0.811	0.811
R^2 from Labor Share VAR Equation				
	0.823	0.820	0.821	0.820

Notes: lsi_t $\equiv$ labor share of income, y_t $\equiv$ output gap, h_t $\equiv$ detrended hours of work, c_t/y_t $\equiv$ detrended consumption-output ratio. Optimal lag length in VAR specifications: $l_1 = l_2 = l_3 = l_4 = 2$. Standard errors are computed using White's heteroskedasticity-consistent covariance matrix estimator (HCCME). *, **, *** denote significance at the 10%, 5%, and 1% level.

Overall, the results presented in this section suggest that our analysis is robust to different VAR forecasting models for the driving variable in the inflation process. Moreover, using the labor share of income as an alternative measure of real marginal costs, does not significantly alter the main results.

4.5 Conclusions

Over the past decade it has become relatively well accepted that the purely forward-looking NKPC cannot account for the degree of inflation inertia observed in the data. In response, the profession has increasingly adopted hybrid models in which lagged inflation is allowed to have an explicit role in pricing behavior. This reformulation of the basic sticky-price model has recently provoked a heated debate as to the extent of forward- versus backward-looking behavior, with little consensus after years of investigation. Most of the empirical studies on the topic take the distribution of heterogeneous pricing behavior as constant and exogenously given. Recent works on structural stability in short-run inflation dynamics in the U.S. have provided statistical evidence of multiple structural breaks in the relative weights of forward- and backward-looking firms. Moreover, empirical studies based on survey data as well as experimental data, provided evidence that the proportions of heterogeneous forecasters evolve over time as a reaction to past forecast errors. In the light of this empirical evidence, we have proposed a model of monopolistic price setting with nominal rigidities and endogenous evolutionary switching between different forecasting strategies according to their relative past performances. Importantly, heterogeneous firms hold optimizing behavior given their subjective expectations on future inflation. In our stylized framework, fundamentalist firms believe in a present-value relationship between inflation and real marginal costs, as predicted by standard RE models, while naive firms use a simple rule of thumb to forecast future inflation. Although with a different behavioral interpretation, our measure of fundamental expectation mirrors the measure of forward-looking expectations in commonly estimated RE models, while the expectations of naive firms account for the lagged value of inflation in the hybrid specification of the NKPC. The difference with traditional tests of sticky-price models arises from the introduction of time-varying weights and endogenous switching dynamics.

We estimated our behavioral model of inflation dynamics on quarterly U.S. data from 1960:Q1 to 2010:Q4. Our estimation results show statistically significant behavioral heterogeneity and substantial time variation in the weights of forward- and backward-looking price setters. The data gave considerable support for the parameter restrictions implied by the theory. In particular, the intensity of choice was found to be positive, indicating that agents switch towards the better performing rule according to its past performance, and inflation was positively affected by real marginal costs. These results were found to be independent from whether detrended output or the labor share of income were used as a measure of real marginal costs. In addition, the heterogeneous agent model outperforms several well known benchmark models in an assessment of competing out-of-sample forecast.

Our findings have important monetary policy implications. Recent papers have shown that multiple equilibria, periodic orbits and complex dynamics can arise in New Keynesian models under dynamic predictor selection, even if the model under RE has a unique stationary solution. Given the statistical evidence found in our results for heterogeneous expectations and evolutionary switching, determinacy under RE may not be a robust recommendation and that monetary policy should be designed to account for potentially destabilizing heterogeneous expectations.

Appendix 4.A NKPC with heterogeneous expectations

Under a Calvo pricing mechanism, the pricing rule of firm i is given by equation (4.3):

$$p_{i,t} = (1-\omega\delta)E_{i,t}\sum_{j=0}^{\infty}(\omega\delta)^j mc_{t+j} + \omega\delta E_{i,t}\sum_{j=0}^{\infty}(\omega\delta)^j \pi_{t+j+1}. \tag{4.16}$$

Aggregating (4.16) over different firms and using that $\sum_i n_{i,t}p_{i,t} = p_t^* = \omega(1-\omega)^{-1}\pi_t$, we obtain a NKPC of the form

$$\pi_t = \overline{E}_t \sum_{j=0}^{\infty}(\omega\delta)^j \left[\frac{(1-\omega)(1-\omega\delta)}{\omega} mc_{t+j} + (1-\omega)\delta\pi_{t+j+1}\right], \tag{4.17}$$

where the operator $\overline{E}_t$, denoting the average expectation at time t, is defined as $\overline{E}_t = \sum_i n_{i,t}E_{i,t}$.[14] Leading (4.17) one-period ahead and applying the operator $\overline{E}_t$ on both sides yields

$$\overline{E}_t\pi_{t+1} = \overline{E}_t\overline{E}_{t+1}\sum_{j=1}^{\infty}(\omega\delta)^{j-1}\left[\frac{(1-\omega)(1-\omega\delta)}{\omega} mc_{t+j} + (1-\omega)\delta\pi_{t+j+1}\right]. \tag{4.18}$$

Assuming that Condition 1 holds and defining

$$X_{t+1} \equiv \sum_{j=1}^{\infty}(\omega\delta)^{j-1}\left[\frac{(1-\omega)(1-\omega\delta)}{\omega} mc_{t+j} + (1-\omega)\delta\pi_{t+j+1}\right],$$

we can rewrite

$$\overline{E}_t\pi_{t+1} = \overline{E}_t X_{t+1} \tag{4.19}$$

$$\pi_t = (1-\omega)\delta\overline{E}_t\pi_{t+1} + \frac{(1-\omega)(1-\omega\delta)}{\omega}mc_t + \omega\delta\overline{E}_t X_{t+1}. \tag{4.20}$$

Substituting (4.19) into (4.20) and defining $\gamma \equiv (1-\omega)(1-\delta\omega)\omega^{-1}$, we finally get

$$\pi_t = \delta\overline{E}_t\pi_{t+1} + \gamma mc_t .$$

[14] Preston (2005) and Massaro (2011) use a NKPC of the same form as (4.17), where inflation depends on expectations over an infinite horizon, for the analysis of monetary policies.

Appendix 4.B Data sources

Below we describe the data sources and the data definitions used in the paper.

Inflation is constructed using the quarterly Price Indexes for GDP from the March 2011 release of the NIPA Table 1.1.4, 1960:Q1 - 2010:Q4, which can be downloaded at http://www.bea.gov/national/nipaweb/SelectTable.asp.

Output gap is constructed using the quarterly real GDP from the March 2011 release of the NIPA Table 1.1.3, 1960:Q1 - 2010:Q4, which can be downloaded at http://www.bea.gov/national/nipaweb/SelectTable.asp.
To construct our measure of the output gap we take logs and quadratically detrend.

Unit labor costs are constructed using the Bureau of Labor Statistics quarterly Unit Labor Costs series PRS85006113, 1960:Q1 - 2010:Q4, for the nonfarm business sector. The series can be downloaded at http://data.bls.gov/, under the heading Major Sector Productivity and Costs Index.

Labor share of income is constructed using the Bureau of Labor Statistics quarterly Labor Share Income series PRS85006173, 1960:Q1 - 2010:Q4, for the nonfarm business sector. The series can be downloaded at http://data.bls.gov/, under the heading Major Sector Productivity and Costs Index.

Hours of work are constructed using the Bureau of Labor Statistics quarterly Hours series PRS85006033, 1960:Q1 - 2010:Q4, for the nonfarm business sector. The series can be downloaded at http://data.bls.gov/, under the heading Major Sector Productivity and Costs Index. To construct our measure of the hours of work we take logs and quadratically detrend.

Consumption-output ratio is constructed using the quarterly real GDP from the March 2011 release of the NIPA Table 1.1.3, 1960:Q1 - 2010:Q4, which can be downloaded at http://www.bea.gov/national/nipaweb/SelectTable.asp.
To construct our measure of the consumption-output ratio we take logs and linearly detrend.

Appendix 4.C Econometric procedure

Here we provide details on the econometric procedure adopted to estimate model (4.14).

Baseline VAR specification

The first step concerns the choice of the baseline VAR specification to estimate the matrix A, needed to construct the forecasts of fundamentalists,

$$E_t^f \pi_{t+1} = \gamma e_1'(I - \delta A)^{-1} A Z_t.$$

We started with a very broad VAR model in the output gap (y_t), unit labor costs (ulc_t), the labor share of income (lsi_t), and the inflation rate (π_t).[15] As shown in Table 4.7, standard unit root tests show that unit labor costs and labor share of income are I(1) processes.[16] Therefore

Table 4.7: Augmented Dickey-Fuller unit root tests on ulc_t and lsi_t

H_0:	ADF test statistic	p-value
ulc_t has a unit root	0.0860	0.9970
lsi_t has a unit root	−2.6689	0.2508

we estimated VAR models which include the rate of change of unit labor costs (Δulc_t) and of labor share (Δlsi_t). The number of lags was chosen optimally on the basis of the comparison of standard information criteria, namely the sequential modified LR test statistic (LR), the Akaike information criterion (AIC), the Schwarz information criterion (SIC) and the Hannan-Quinn information criterion (HQ). We then performed pairwise Granger causality tests and proceeded iteratively, eliminating insignificant regressors, highest p-value first. We found evidence that

[15] This specification extends the baseline specifications of previous empirical works, e.g., Woodford (2001b) and Rudd and Whelan (2005a), by adding lagged inflation in the output gap equation. However, we later eliminate it from the VAR specification together with the rate of change of unit labor costs since we found evidence that none of them Granger cause the output gap.

[16] The presence of a unit root in the labor share time series was not detected in previous empirical works such as Woodford (2001b) and Rudd and Whelan (2005a). This is due to the fact that our dataset incorporates observations until 2010:Q4. Unit root tests performed on the same sample considered by Woodford (2001b) and Rudd and Whelan (2005a) confirms the results found by these authors, i.e., the presence of a unit root in lsi_t is rejected.

neither inflation nor the rate of change of unit labor costs Granger cause the output gap, therefore we excluded the variables Δulc_t and π_t from the VAR and we chose a four-lag bivariate VAR in the output gap and labor share of income growth as our baseline specification.[17] The Portmanteau test reports no autocorrelation in the residuals up to the 20th lag (p-value $Q(20) = 0.796$) and, although being parsimonious, the baseline VAR captures about 95% of output gap volatility ($R^2 = 0.945$).

Denoting by Y_t the vector of dependent variables, $Y_t = [y_t, \Delta lsi_t]'$, the vector Z_t is defined as $Z_t = [Y_t, Y_{t-1}, Y_{t-2}, Y_{t-3}]'$. The matrix A denotes then the matrix of OLS estimates of the baseline VAR, obtained by regressing Z_t on Z_{t-1}.

NLS estimation

Here we describe the estimation of model (4.14) by Non-linear Least Squares (NLS). For notational convenience we rewrite model (4.14) as

$$\pi = \mathbf{x}(\theta) + \mathbf{u}, \tag{4.21}$$

where $\theta = [\beta, \gamma]$ is the vector of parameters to be estimated, and the scalar function $\mathbf{x}(\theta)$ is a nonlinear regression function corresponding to the dynamics implied by model (4.14). We estimate model (4.21) using NLS, see Table 4.1 in section 4.3.2, and perform diagnostic checks on the residuals. Table 4.8 reports the result of the White test for heteroskedasticity. Given

Table 4.8: Heteroskedasticity test

H_0: homoskedasticity			
F-statistic	8.058	Prob. $F(3, 194)$	0.000
Obs*R^2	21.94	Prob. $\chi^2(3)$	0.000

the presence of heteroskedasticity, we perform the heteroskedasticity-robust test for serial autocorrelation proposed by Davidson and MacKinnon (2009). The procedure is based on the

[17] The lag order of 4 was selected by 3 out 4 criteria, namely the LR, the AIC, and the HQ.

following Gauss-Newton regression:

$$\widehat{\mathbf{u}} = \widehat{\mathbf{X}}\mathbf{b} + \widehat{\mathbf{Z}}\mathbf{c} + \text{residuals}, \tag{4.22}$$

where $\widehat{\mathbf{u}} = \pi - \mathbf{x}(\widehat{\theta})$ denotes the NLS residuals from (4.21), $\mathbf{X} \equiv \mathbf{X}(\widehat{\theta})$ denotes the vector of derivatives of $\mathbf{x}(\theta)$ with respect to θ, computed in $\widehat{\theta}$, and $\widehat{\mathbf{Z}}$ is a matrix collecting lagged values of the NLS residuals $\widehat{\mathbf{u}}$. In order to test for $\mathbf{c} = 0$, we use the test statistic suggested by Davidson and MacKinnon (2009), distributed as $\chi^2(p)$ under the null hypothesis, where p denotes the order of serial correlation being tested.[18] The results reported in Table 4.9 show the absence of serial correlation in the residuals up to the 20th lag.

Table 4.9: Serial correlation test

H_0: no serial correlation ($\mathbf{c} = 0$)			
test-statistic ($p = 20$)	21.51	Prob. $\chi^2(p)$	0.368

[18]See Davidson and MacKinnon (2009), p. 284, for details about the construction of the test.